JUNIPER NETWORKS | Engineering Simplicity

Juniper Mist WAN Assurance Configuration Guide

Published
2026-02-06

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

Juniper Mist WAN Assurance Configuration Guide

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

About This Guide

Use this guide to learn about key features of Juniper Mist WAN Assurance and to understand the configuration choices that are available through the Juniper Mist™ cloud portal.

1

CHAPTER

Get Started

IN THIS CHAPTER

Overview of Juniper Mist WAN Assurance

SUMMARY

Get familiar with the purpose, features, and benefits of Juniper Mist™ WAN Assurance.

IN THIS SECTION

The WAN Edge is the demarcation point for your enterprise network to reach the outside world. This boundary is a crucial security and troubleshooting hotspot. The WAN Edge can be a simple border between your enterprise network and the outside world. However, the WAN Edge can also be a Juniper® SD-WAN driven by Mist AI™ device. Options include Juniper® SRX Series Firewalls, Juniper® Networks Session Smart™ Routers, or a cloud solution like Juniper Secure Edge.

Figure 1: Juniper Mist WAN Assurance

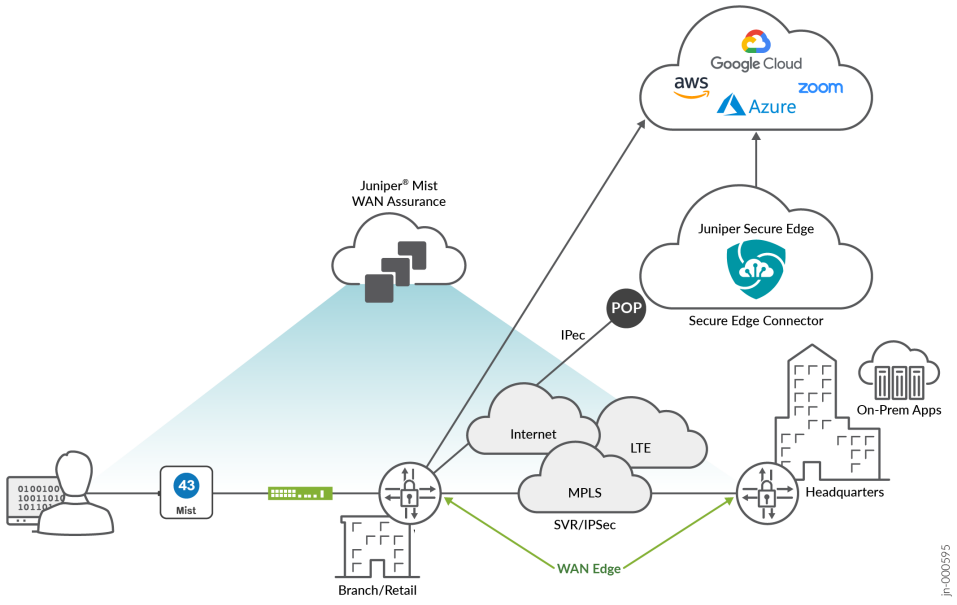

The WAN Edge transforms with Juniper's AI-driven SD-WAN solution and acts as your centralized policy enforcement point (PEP). Combined with the Juniper Mist WAN Assurance cloud service, the

Juniper SD-WAN solves many of the legacy SD-WAN solutions' security, monitoring, and troubleshooting challenges. Bring deployment, monitoring, and troubleshooting across your network by integrating Juniper Mist Wired Assurance, Juniper Mist Wireless Assurance, and now Juniper Mist WAN Assurance. Juniper Mist WAN Assurance securely connects branch offices with Juniper® Session Smart™ Routers or Juniper® SRX Series Firewalls across the SD-WAN.

Figure 2: Juniper AI-Driven SD-WAN Solution

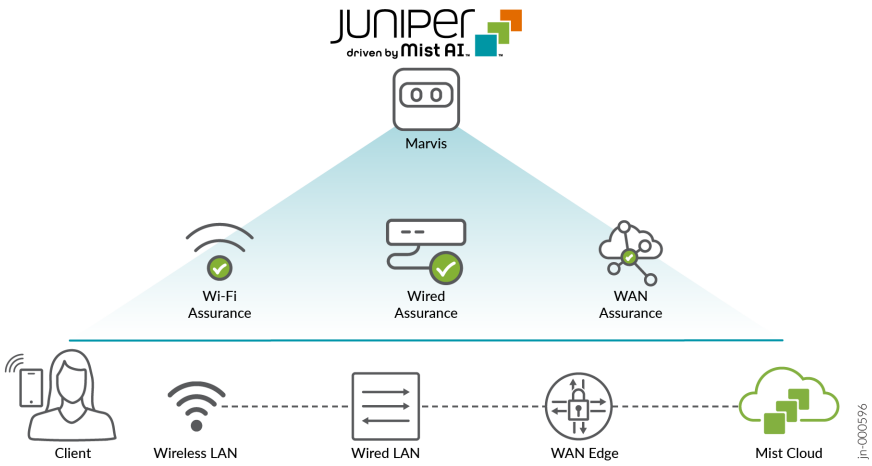

Intersite Connectivity (SD-WAN)

Your WAN Edge transforms when integrated with Juniper® SD-WAN driven by Mist AI™. Your WAN Edge device becomes fast, secure, and application-aware with Juniper Mist WAN Assurance. Software-defined WAN traffic from remote sites travels through an abstracted overlay across less expensive broadband service providers. This connectivity design replaces expensive legacy MPLS solutions. Edge devices deliver stateful failovers across various connection types, including MPLS, broadband, satellite, and LTE. This real-time switch for critical applications is imperceptible to the user. Juniper Mist WAN Assurance also brings visibility to your WAN Edge with targeted insights for health, tunnel activity, connectivity, and active sessions. By creating that software-defined space, you can influence traffic at the application level with greater control for access and security.

Figure 3: SD_WAN Intersite Connectivity

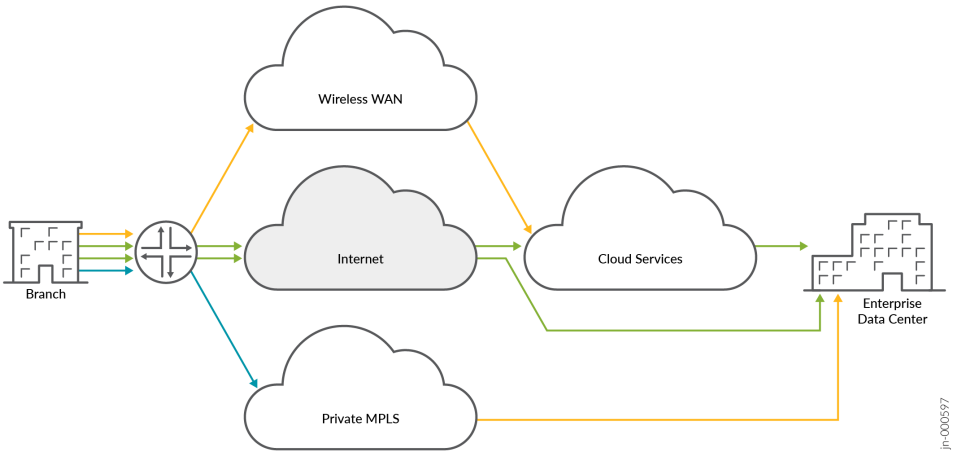

Juniper Mist WAN Assurance Service Level Expectations (SLEs)

Mist's Predictive Analytics and Correlation Engine (PACE) provides data science and machine learning to understand the end-user experience. The WAN SLE metrics are: WAN Edge Health, WAN Link Health, and Application Health. Juniper Mist WAN Assurance identifies the root cause of WAN issues impacting user experiences. Service-level expectations enable simpler operations, better visibility into end-user experiences, and simplify monitoring and troubleshooting your network.

Figure 4: Juniper Mist WAN Assurance Service Level Expectations (SLEs)

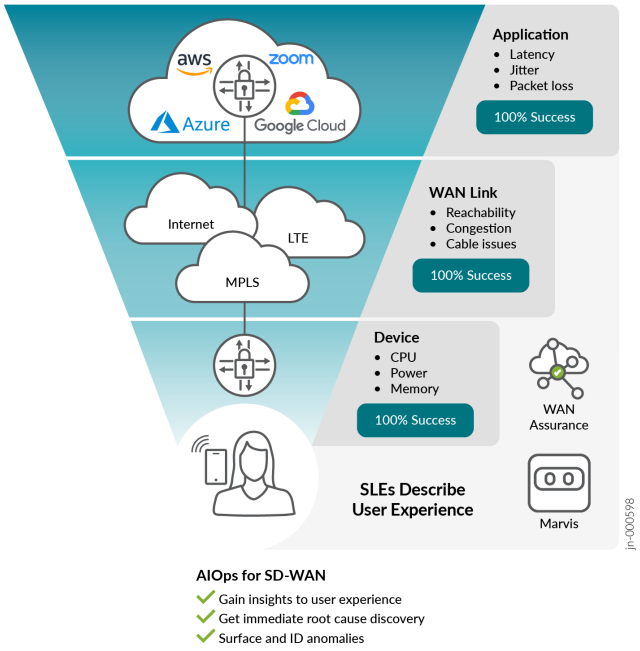

Mist Management Model

Juniper's AI-driven SD-WAN solution is a single management platform for branch Wireless, Wired, and SD-WAN. Juniper SD-WAN zero-touch provisioning (ZTP), life cycle, and configuration are done through a single Mist dashboard.

Figure 5: Mist Management Model

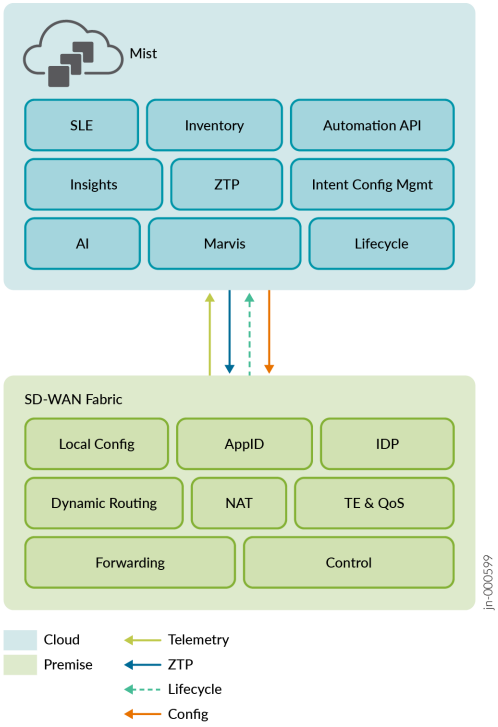

Watch the following video for an overview of the Juniper Mist WAN Assurance feature.

Video: WAN Assurance - How's user experience? Really?

Juniper Mist WAN Assurance Configuration Hierarchy

IN THIS SECTION

Introduction to Configuration Hierarchy

Juniper Mist WAN Assurance Configuration

For network administrators, it's essential to understand that each piece of the puzzle builds your network's policies, security, and connectivity in Juniper Mist WAN Assurance cloud service. A full SD-WAN deployment requires each part to complete intersite connectivity. Mist automatically translates your traffic intent to configurations for WAN Edge devices using Mist's intent-based networking (IBN) model. Each part works together to build complex interface assignments, security, routing policies, and—depending on the platform—destination zones. Therefore, understanding the Mist intent model is crucial as we dive into the configuration hierarchy for Juniper Mist WAN Assurance.

Intent-Based Routing

IBN solves several problems. For example, consider the need for secure communication between two networks. An intent model states that secure communication requires a secure tunnel between Network A and Network B. In this scenario, a network administrator identifies which traffic uses the tunnel and describes other desired general properties. But an operator wouldn't specify or even know how to build a tunnel. To implement a tunnel, you must know how many devices to secure, how to make BGP advertisements, and which features and parameters to turn on. By contrast, an IBN system automatically generates an entire configuration of all devices based on the service description. It then provides ongoing assurance checks between the intended and operational state of the network, using closed-loop validation to continuously verify the configuration's correctness. IBN is a declarative network operation model. It contrasts with traditional imperative networking, which requires network engineers to specify the sequence of actions needed on individual network elements and creates significant potential for error.

Intent-based model key characteristics:

- Do not require as much explicit direction as traditional network models require.
- Build policies based on which network goes to which application.
- Configure Juniper Mist WAN Assurance **Networks** and **Applications** organization-wide.
- Push relevant configurations only.

- Configure only the **Applications** that a device uses. If a device doesn't use an **Application**, the intent-based network doesn't configure it on that device.

Let's look at the example of configuring DHCP on a LAN and assume that the interface is already configured and assigned to a zone.

Required steps in the Junos CLI:

- Navigate to the Junos system services level and enable the DHCP-local-server for your interface.
- Navigate to the Junos system address assignment and create an address pool specifying the target network, the range of addresses for the pool, the default gateway, and any other DHCP attributes.
- Navigate to your security zone and enable host-inbound traffic for the DHCP system service to allow the SRX Series to process DHCP requests from clients.

The above steps require multiple configuration lines spread across three configuration hierarchies at a minimum.

The same workflow is significantly streamlined in Mist:

- First, navigate to your LAN configuration and open it for editing.
- Next, enable the DHCP Server radio button to unlock the configuration and populate the required fields (IP Start, IP End, and gateway).
- Save the LAN configuration and then save the device configuration.

Configuration Hierarchy Elements

Organization-Wide Configuration Elements

The top of the Mist configuration is called your Mist organization. These elements impact your entire software-defined wide area network (SD-WAN) deployment. The different components at this configuration level become building blocks for sources and destinations across your deployment. Once identified, traffic requests associate a sender and the desired destination appropriately. The elements help build different Juniper Mist WAN Assurance deployment components depending on your platform. Identifying the source and destination will build IPsec tunnels across the WAN and associated security zones on the Juniper® SRX Series Firewall. These components on the Juniper® Networks Session Smart™ Router become the corresponding source and destination to help build the Secure Vector Routing (SVR) metadata exchange. The two platforms approach the challenge of SD-WAN uniquely, which makes it important to know your Juniper Mist WAN Assurance platform.

Networks

The Juniper Mist WAN Assurance **Network** is the “who” in the Mist intent-driven paradigm. **Networks** are sources of the request in your network. Networks enable you to define groups of “Users.” Once you create this element in your Mist design, the network is defined for use across the entire organization.

Characteristics of Networks on the Juniper® Networks Session Smart™ Router:

- Mist **Networks** create Tenants in the background for SVR.
- The Session Smart Router identifies Tenants at the logical interface (Network Interface).
- LAN and WAN interface configurations identify your Tenant (request source).

Characteristics of Networks on the Juniper® SRX Series Firewall:

- **Networks** create Address books used as the source for **Security Policies** and Advanced Policy Based Routing (APBR) Policies.
- Configurations are applied to the device if an **Application Policy** is configured.
- For the LAN, the zone's name is derived from the name of the specified network.
- For the WAN, the zone’s name is based on the name of the WAN.

Route Advertisement (Advertise via Overlay)

WAN Assurance is about the abstraction of the transport network into the SD-WAN. You can advertise networks through SD-WAN for control and reachability with route advertisement. Then established networks in your LAN segments can be advertised across the overlay. Setting up these networks generates the source addresses for service policies. Network Address Translation (NAT) for the source and destination can route traffic to your users if needed.

The purpose of SD-WAN is intersite connectivity. Therefore, networks can be advertised through overlay to enable reachability between your SD-WAN devices. With this setting, your network will share the address across the WAN so other devices know how to reach it. You can advertise to other spokes or to hub LAN neighbors. For more information about this feature, see "Network Settings" on page 94.

Access to Mist Cloud

Mist is a full-stack solution. Only some of your devices are WAN Edge or SD-WAN routers. Specific devices will want access to the Mist Cloud to leverage other solutions like Wireless and Wired Assurance on Wireless APs and switches. **Access to Mist Cloud** will automatically generate specific firewall/policy rules enabling the devices to phone home to Mist without needing an explicit application policy. However, you don’t want this level of access on all devices behind the WAN Edge in an SD-WAN

deployment because doing so can pose a policy challenge for routers. Generally, select **Access to Mist Cloud** for APs or switches, so that you can monitor and troubleshoot these devices from the Mist portal.

Enabling access to the Mist cloud ensures that anything sitting behind the WAN Edge can reach the Mist cloud without needing to express policies for connectivity manually. Ports and protocols for this setting include the following:

- TCP/443
- DNS/53
- SSH/2200
- NTP/123
- Syslog/6514
- ICMP

Users

Don't let the label fool you. **Users** does not represent a single user on your network. **Users** are subsets of subnets or indirectly connected subnets. Since **Networks** are "who," think of **Users** as a subdivision of that network identity. There are often universal rules to treat networks the same. For example, for 99% of your traffic, you want sessions to do the same thing. But what about when you're blocking access to a corporate network from a guest network, and only one IP needs printer access? In this situation, add a **User**. For those familiar with the Session Smart Routing platform, compare a User to a Tenant. You also might create **Users** to define indirect prefixes on the network.

- **Users** can define granular permissions. For example, your LAN segment may need Internet access, but you must restrict it to a particular network device. So here, you'd create an access policy around that desktop.
- You sometimes need to reach indirectly connected prefixes behind a router on the LAN segment. For example, picture a router behind a device that connects multiple devices to an outside application. In this case, you can add users to a network that you've specifically configured as a "network not directly attached." For more information, see "Network Settings" on page 94.

Applications

Applications comprise the "what" in the Mist intent-driven model paradigm. **Applications** are what your network delivers. **Applications** represent traffic destinations and are named for what a client would access, like a "database" or the "Internet." Once you create this element in your Mist design, the **Application** is defined for use across the entire organization.

Characteristics of **Applications** on the Juniper® Networks Session Smart™ Router

- Mist applications create services in the background for SVR.
- **Applications** can be ports, protocols, prefixes, custom domains, or app names from the built-in AppID library.

Ports, protocols, and prefixes are where all the policy revolves.

- **Custom Apps** are a set of ports, protocols, or prefixes.
- **Apps** map to the Internet app-id.
- **URL Categories** are force-point URLs.

Characteristics of Applications on the Juniper® SRX Series Firewall

- **Applications** determine the destination used in a security policy.
  - A prefix of 0.0.0.0/0 with protocol "any", is resolved to *any* within the Juniper Mist WAN Assurance policy. No address book or application is necessary.
- **Custom Apps** on the WAN Edge use the SRX Series on-box engine "type" and are a combination of an address book and applications.
- **Apps** map to the SRX Series Layer 7 AppID engine.
- **URL Categories** are force-point URLs.

Traffic Steering

Traffic Steering is the "how" in the Mist intent-driven model paradigm. **Traffic Steering** is how you define the different paths that traffic can take to reach its destination. If traffic to an application has multiple paths, you can restrict the paths to a subset of paths and configure an order of preference. You can also load and balance numerous streams across the available paths.

Characteristics of **Traffic Steering** on the Juniper® Networks Session Smart™ Router:

- The Juniper® Session Smart Router™ is session-based and uses continuous path monitoring techniques for both underlay and overlay paths to find the best available path for any application.
- There are three steering strategies for the SSR Series:
  - Ordered: This is the default—go in order of the list. Active paths at the top take priority. If a path goes down, then move to the next active path in the list. This creates an ordered list.
  - Weighted: Allows you to set your desired order based on weight. For example, two weighted paths, both set to 5, result in ECMP sessions across the two paths. On the other hand, two weighted paths, with one set to 5 and the other set to 10, result in ordered steering, with sessions taking the lower-weight path first.

    - ECMP: Fully load balance traffic with an equal-cost multipath algorithm. Sessions will be split evenly across all available paths.

- Unlike the SRX Series Firewalls, traffic Steering is not required on an application policy for the SSR if there is already a route for the traffic in its RIB. There are situations where configuring traffic steering on an application policy will result in undesirable behavior. See "Internet Backhaul Through an SSR Hub" on page 167 for more information.

- The SSR supports traffic steering policies that can steer traffic in two ways:

    - Towards the overlay with various options for steering this traffic over different WAN paths using Secure Vector Routing (SVR). For traffic steering in the overlay, Mist WAN Assurance relies on BGP to route traffic between SSR devices. You can leverage this behavior to exchange and propagate routes between your existing network(s) and your SSR device(s).

    - Locally routed out one or more specific interfaces, which is common for local breakout (underlay) traffic. For customers who do not want to perform dynamic routing with the SSR, or customers without existing dynamic routing solutions, see "Internet Backhaul Through an SSR Hub" on page 167 for details.

- For application policies that have a block action, do not enter any traffic steering

- SSR employs a deny-by-default behavior. You do not need to create blocking policies unless a particular network object already has access to a broader application and you wish to limit a specific range within that address space.

Characteristics of **Traffic Steering** on the Juniper® SRX Series Firewall:

- The Juniper® SRX Series Firewall is zone-based, and the destination zone is determined by the paths configured within a **Traffic Steering** policy.

- **Traffic Steering** configures forwarding-type routing instances and the relevant routing policy to import routes. For your SRX Series, this routing instance is used in APBR.

- There are several steering strategies for the SRX Series:

    - Ordered: Default, go in order of the list. The top takes priority, then failover to the next. Creates an ordered list.

    - Weighted: Allows you to set your desired order based on weight. For example, two weighted paths, both set to 5 results in ECMP across the two paths. On the other hand, two weighted paths with one set to 5 and the other set to 10 results in ordered steering with traffic taking lower weight path first.

    - ECMP: Fully load balance traffic with an equal-cost multipath algorithm. Traffic will be split evenly across all available paths.

Application Policy

The “who,” “what,” and “how” come together with **Application Policy**. The Mist intent-driven model simplifies manually generating routes and security policies through Junos OS on the SRX Series with thousands of lines of code. It also simplifies deploying a Session Smart Router for those transitioning from a Conductor-based Session Smart deployment to WAN Assurance. You no longer need explicit permissions and interface assignments to get up and running. WAN Assurance is zero trust. This feature is both implied and part of the intent-driven model. You must explicitly grant permission to allow a **Network** to access to an **Application**. Otherwise, it will not route.

Order only matters when egressing your local network on the Juniper® Networks Session Smart™ Router. The Session Smart Router uses the most specific matches. As a result, traffic steering isn’t necessary for local traffic. Also, using a block in your **Traffic Steering** does not work with SVR, as it undermines the proprietary process. If you don't want a device, subnet, or network to access an **Application**, don't create traffic steering for that device.

Characteristics of **Application Policy** on the Juniper® SRX Series Firewall:

The steering path determines the destination zone in the SRX series. Please ensure that policies have **Traffic Steering** assigned because the order of policies matters when working with the SRX Series. As a traditional zone-based firewall, it uses a list of rules that generate filters and policies. Most specific rules should be at the top of the Application Policy list on the SRX Series.

Scaling Your Network: Automation in Mist

WAN Edge Templates

Once basic configuration elements of SD-WAN are in place, Mist enables you to deploy new WAN Edge devices through WAN Edge templates. All that previous configuration can be templated with WAN Edge templates. These templates work for a standalone Edge device to a full SD-WAN deployment with hundreds of sites. The automation process removes errors and simplifies deploying multiple spoke sites and headends.

Templates reduce or eliminate common configuration tasks and remove human error when configuring multiple devices. WAN Edge templates:

- Enforce standards across a deployment.
- Ensure that all your network devices point to the same DNS (8.8.8.8).
- Provide predictable behavior because they use the same Network Time Protocol (NTP) for synchronization and logging. (This also affects specific certificates.)
- Simplify troubleshooting and management.

Figure 6: WAN Edge Template

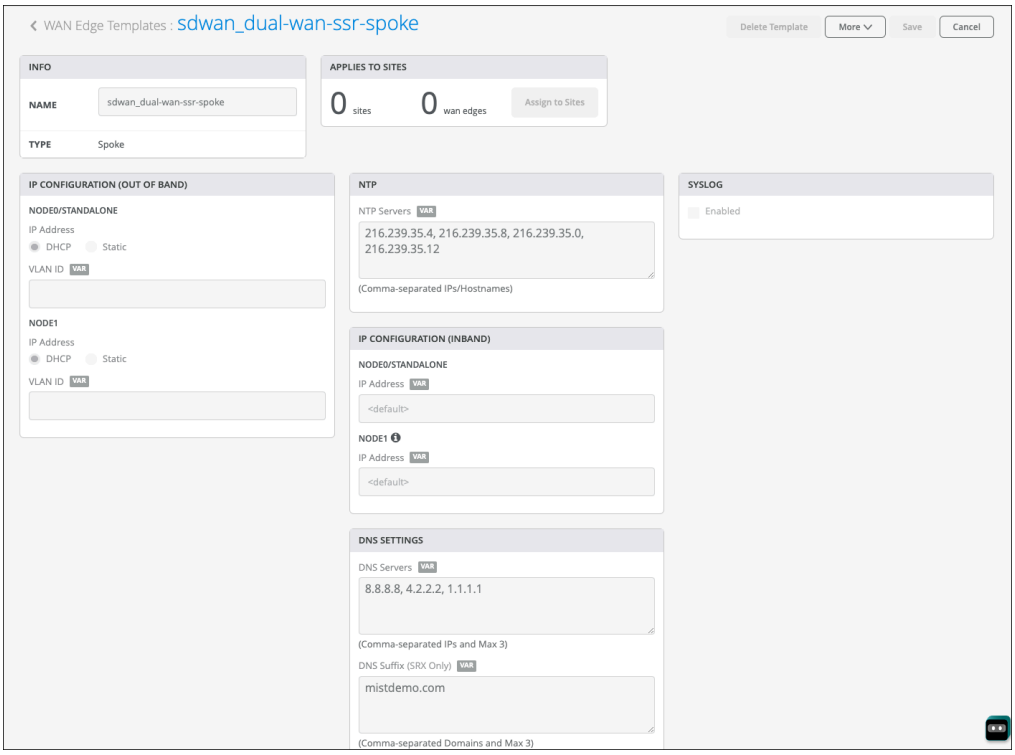

However, WAN Edge templates do more than automate tasks. You might use a template to standardize a configuration that *can be* applied consistently across sites even if you don't *actually* deploy all features at all sites. For example, you might not need a guest network at every site, but by including the configuration in the template, you're reserving that interface. If future plans call for a guest network, the interface is ready to be used.

These templates also allow for:

- Bulk orders of hardware for ports and site groups through specific models.
- Specific use cases and traffic flows.
- Different corporate LAN networks.
- Guest networks.

WAN Edge templates automatically configure repetitive information like an IP, gateway, or VLAN. In addition, WAN Edge templates can include traffic steering, access policies, routing preferences, and any additional configuration you'd like to standardize. Remember that you'll need a prefix, NAT, or other local information for WAN and LAN connectivity.

Hub Profiles

Hub profiles work with WAN Edge templates. Hubs are not at the Edge and are universally unique throughout your network. Hubs affect how Mist builds the overlay network. Each branch and remote office build the SD-WAN communication to the hub. Topology is determined by overlay endpoints that make up a single overlay. Every hub WAN interface creates an overlay endpoint for spokes. Spoke WAN interfaces map the appropriate hub WAN interfaces, defining the topology. This is the abstraction of the transport network. Because the two platforms for WAN Assurance solve the abstraction differently, you need to understand their nuances when building that overlay network.

Juniper® SRX Series Firewall

The SRX Series overlay SD-WAN combines a virtual router for route separation and IPsec tunnels for secure transit traffic. WAN configurations determine the topology and build the overlay network. One thing to note is that you can implement only one overlay per organization. However, you can have many paths within this overlay across multiple types of transport and securely isolate and forward traffic. For SRX Series devices, the overlay combines a security zone, virtual router, and IPsec tunnels.

Juniper® Networks Session Smart™ Router

The Session Smart overlay SD-WAN is your neighborhood, which involves proprietary communication through BFD on port 1280 for liveness and jitter, latency, and loss between Session Smart peers. When you configure a WAN interface on a hub profile, it creates an overlay hub endpoint. On the Session Smart Router, the endpoint is the receiving end of the SVR.

A few things happen when you map a spoke WAN Interface to the overlay hub endpoint. The spoke will establish peer connectivity and identify the neighborhoods and vectors for SVR, which is the Session Smart abstraction of the transport network.

A final note on the overlay: The SRX Series and Session Smart Routers cannot exist in a single overlay. These devices can be paired through BGP at the hub, but their solutions to create intersite connectivity are unique and cannot operate together in the same overlay. If you have migration plans, identify which routes need advertisement and advertise at the hub.

Keep the following hub profile considerations in mind:

- Hub profiles must be built first, so spoke templates know where to connect.
  - Hubs must have static IPs for overlay endpoints.
  - The overlay endpoint configuration is exposed in the WAN Edge spoke template.
- There is no limit to the number of hubs you can incorporate within these guidelines:
  - One hub per datacenter
  - Two hubs for redundancy (HA clusters)

Spokes choose the primary hub through traffic steering and an **Application Policy**. Zero-touch-provisioning (ZTP) requires DHCP (for physical implementation) unless ZTP is done and then migrated to the destination network. You can pre-stage the devices manually, too.

Figure 7: Hub Profile

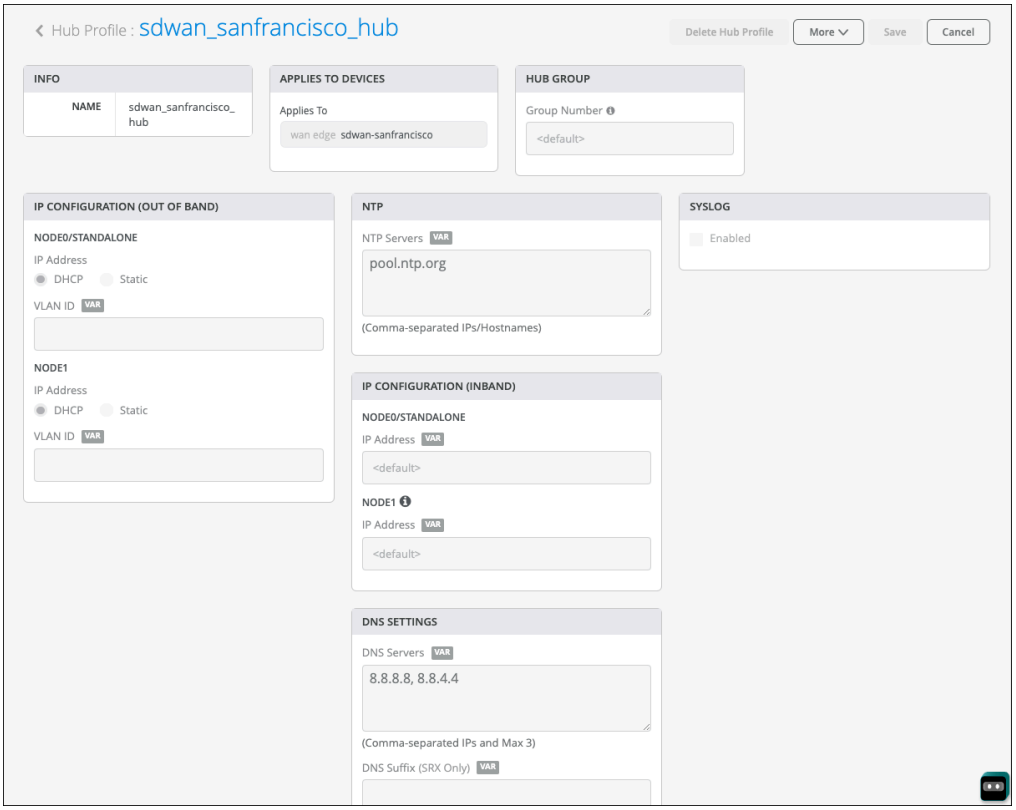

Site Variables

Site variables are configured on a per-site basis. When planning a network holistically, you can create standard templates for specific WAN Edges and WAN Edge clusters. Ideally, you have only one WAN Edge device per site (or a single logical WAN Edge if the device is clustered). Since variables can differ per site, administrators use them in templates or the WAN Edge configuration page. The transformation happens when the configuration is rendered and pushed to the device.

Keep the following site variable considerations in mind:

- The syntax for variables matches Jinja2 and is contained within double curly brackets, like this: {{variableName}}
- The UI enforces the leading and trailing curly brackets as part of the name.
- Site variable limitations:

- No spaces in the variable.
- No special characters (except underscore) within the variable field.
- Variables can be used only in one field and cannot specify an entire prefix.

For example, 10.88.88.88/24 would need at least two variables, one for the IP address (10.88.88.88) and another for the prefix length (24).

Figure 8: Site Variables

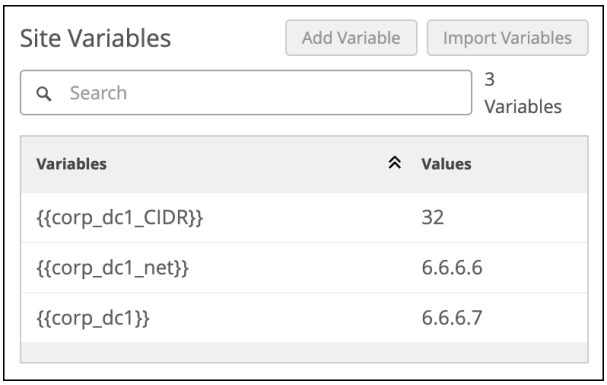

The best way to use the real power of templates is with site variables. Many configuration items are required to deploy the hardware. It makes sense to combine the WAN Edge templates and site variables. Consider the following situation where you can define entire IP subnets of the first three octets, leaving minimal configuration at each device:

Create standard templates and place variables in standard interfaces like your WAN in either of these ways:

- With a WAN1PFX variable, let's say {{192.168.170}}, and in the WAN field on the Configuration page, it would be {{WAN1PFX}}.1 for the local IP and {{WAN1PFX}}.2 for the gateway.
- You could define a {{WAN1IP}} and {{WAN1_GW}} pair of variables; however, there are places where the subnet may be reused but not the specific IP.

Figure 9: Site Variables in WAN Configuration

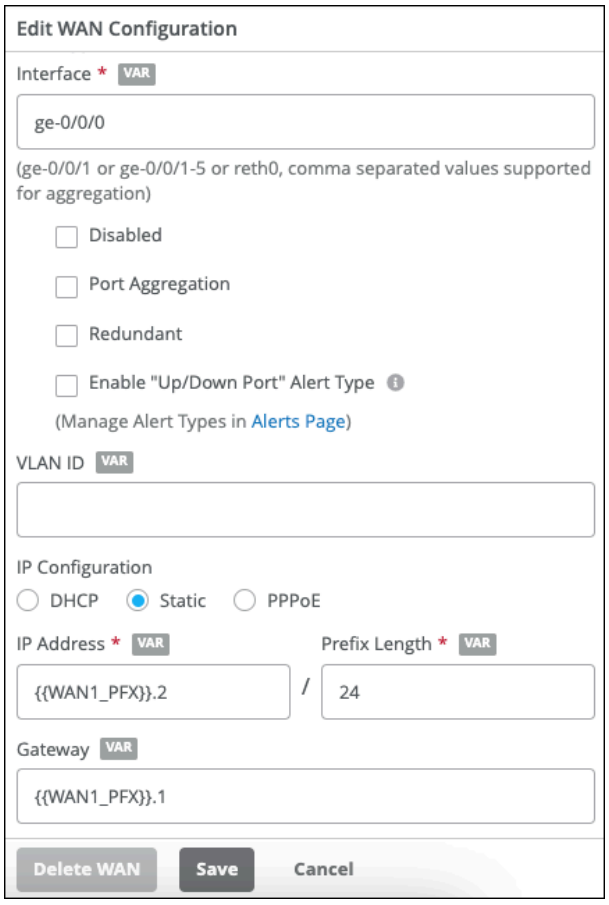

Another robust use case is the magic octet, where the third octet becomes a variable, and that variable might also apply to multiple fields. For example, a {{SITEID}} variable might be used for both the third octet and a VLAN tag. In that case, the network prefix might be 192.168. {{SITEID}}.1/24 with the {{SITE_ID}} VLAN ID. Remember that although WAN Edge templates apply only to the WAN Edge, site variables also apply to switches and APs. The purpose of the automation is to simplify deployments and increase reusability.

Introduction to Applying Templates

Remember that a site is a collection of all your assets in a single location. It is implied that there will only be a single WAN Edge. A key feature of Mist management through the Juniper Mist AI is your ability to use configuration templates to group WAN Edges and make bulk updates. Templates provide uniformity and convenience, while the hierarchy provides scale and granularity.

Import and Export Templates

No solution covers all circumstances. You can have multiple templates. To save time, clone a template. Then customize the cloned copy by modify it as needed.

Figure 10: Export or Clone Template

Modifying Template Sample

```
{
"type": "standalone",
"ip_configs": {
  "LAN": {
    "type": "static"
    "ip":
"{{LAN1_PFX}}.1" ,
      "netmask": "/24*
     }
  },
  "dhcpd_config" : {},
  "dns_servers" : [
    "8.8.8.8"
   ],
....
}
```

Overriding the Template

Templates apply to sites, which apply to devices. Templates are used to standardize configurations, but exceptions always exist. Rather than create a slightly different template for one site, you override the template configuration on the device.

Figure 11: Override Template Settings

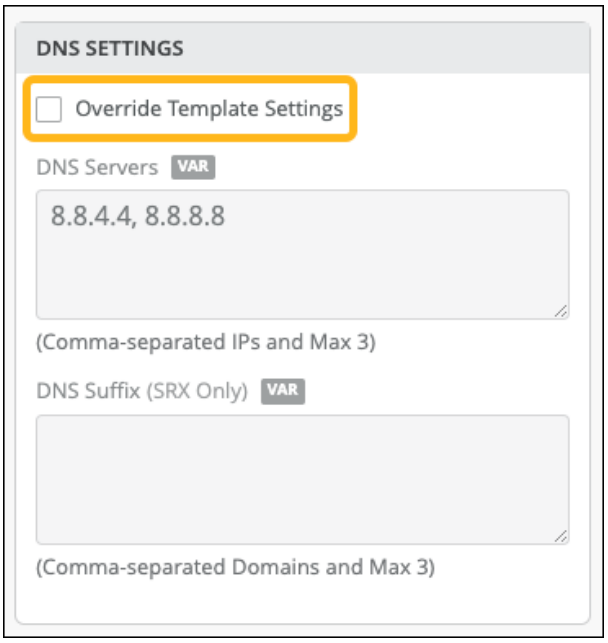

If you need to override the template, you can enable the Override Template Settings option for the required configuration blocks on a per-device basis. Figure 12 on page 20 shows how you can override the DNS and the **Application Policy** but none of the other settings, such as WANs, LANs, or NTP servers.

Figure 12: Application Policy

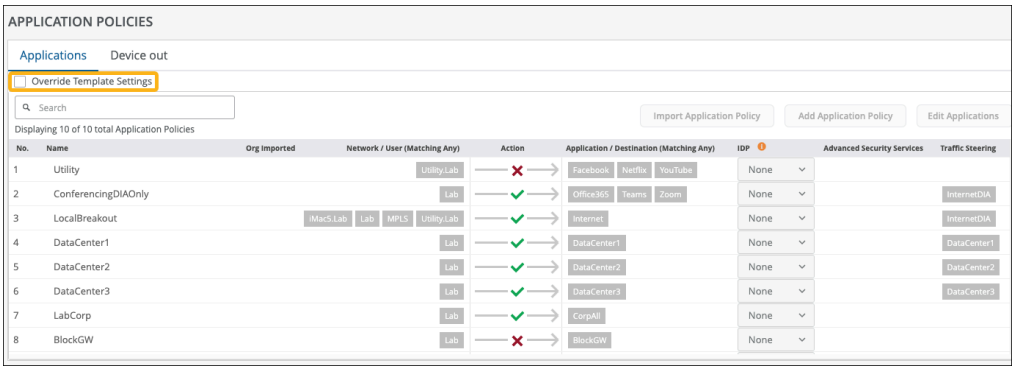

The screenshot illustrates an all-or-nothing action. When you override the template settings, this configuration longer inherits any application policies from the WAN Edge template.

You need to have one of the following role assigned to you in order to override the configuration:

- Super User

- Network Admin (All Sites Access)
- Network Admin (Site Group or Specific Sites Access)

Organizational-Level Application Policy

Figure 13: Organizational-Level Application Policy

Figure 13 on page 21 shows **Application Policy** configuration option at the organization level.

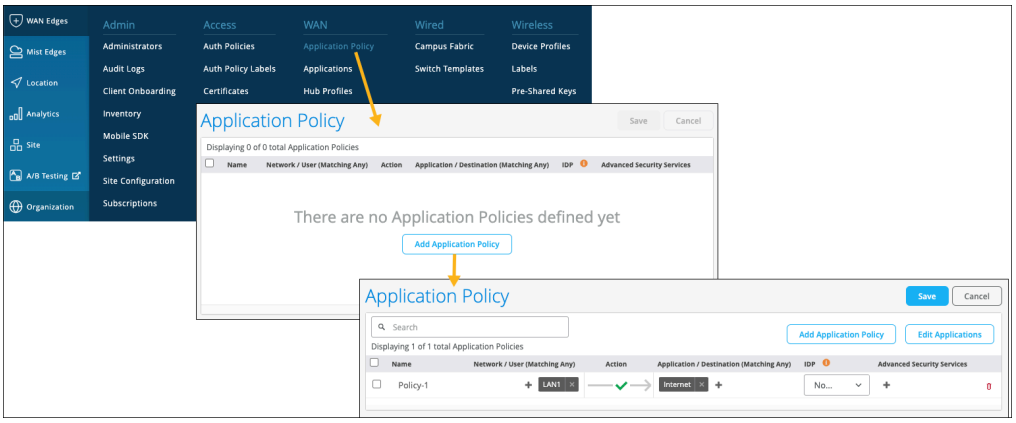

Although templates save time on deploying multiple devices, you might have various templates to account for different device models or slightly different configurations. You can create the same **Application Policy** on each template, but consider using an organization-level **Application Policy** as a shortcut. With an organizational-level **Application Policy**, you can create importable application rules into WAN Edge templates and hub profiles for large-scale network topologies.

Let's explore some best practices and restrictions for using an organizational-level **Application Policy**. Give each organization-level **Application Policy** a globally unique name, or you'll get errors when saving the configuration. The imported policy has all the fields dimmed because it is not meant to be modified. There's no organizational-level **traffic steering**, which makes sense, because traffic steering applies to local connections and intent.

Consider applying an organizational-level **Application Policy** to a LAN block or one subnet. If you create a LAN "supernet" of a 10/8, the policy will allow anything sourced from 10/any to reach the Internet, meaning it would work for all your sites. This is why planning is crucial. Design your network to streamline troubleshooting with similar traffic patterns regardless of deployment. For example, some sites have LTE, and traffic must egress there on that site instead of others. In addition, some sites are standalone, and others are SD-WAN. A universal policy could apply to both by telling traffic steering on standalone sites to go out the WAN to the underlay, whereas the sites are going to the overlay for the SD-WAN spokes.

In summary, the use case for an organization-level policy is to describe network-wide traffic patterns regardless of the site; as a policy, you define what is and is not allowed. Then, when applied to the site or template (which applies to places), you add the steering portion giving you the final piece of the puzzle.

WAN Design Considerations

Figure 14 on page 22 shows the workflow for WAN Edge provisioning.

Figure 14: WAN Edge Provisioning Workflow

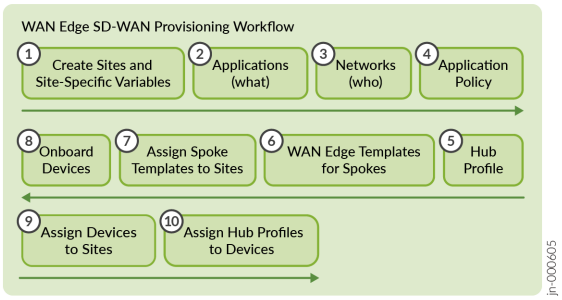

Reviewing the blocks that make up the completed project is essential to deploy an SD-WAN, as follows:

1. Think about “who” (**Network**) makes up the source of requests in your organization.
2. Consider “what” destinations (**Applications**) users access.
3. Where do those elements go in your organization? Consider site types.
4. Finally, consider “how” (**Traffic Steering**) those users get and gain access to their traffic destinations.
5. Now you can use the power of Mist AI, templates, and variables for scale.

SD-WAN Provisioning

The order of operations matters. When preparing to implement SD-WAN provisioning, complete tasks in this order:

1. First, plan your network with templates, thinking about the deployment holistically.

2. Hub profiles must come before WAN Edge spoke templates.

3. Design with your **Applications** (destinations of traffic) first, and then **Networks** (who).

You can analyze apps and become more granular later.

1. Make sure you know your networks (sources of traffic).

Networks inform policy and traffic steering.

- Apply the appropriate **Application Policy** to both ends (spokes and hubs).

Strive for end-to-end reachability when establishing overlay endpoints. Keep in mind that you can't connect an isolated MPLS endpoint to an Internet endpoint.

RELATED DOCUMENTATION

Juniper Mist WAN Assurance Platform Considerations | 23

Overview of Juniper Mist WAN Assurance | 2

Juniper Mist WAN Assurance Platform Considerations

SUMMARY

Understand the various WAN Edge platform options to ensure that Juniper Mist™ WAN Assurance meets your organization's specific needs, current network architecture, and future goals.

IN THIS SECTION

Introduction

Juniper Mist WAN Assurance is a comprehensive SD-WAN solution to streamline your operations and optimize your end user experience. It does so using tunnel-free overlay technology, integrated security, cloud management with zero-touch provisioning, and AI-driven operations. Choosing the right WAN Edge platform is crucial for leveraging the full benefits of Juniper Mist WAN Assurance. This topic aims to assist organizations in making informed decisions aligned with their specific needs, current network architecture, and future goals.

Platform Selection Criteria

For WAN Edge devices in Mist WAN Assurance, Juniper Networks® Session Smart™ Routers and Juniper Networks® SRX Series Firewalls can be used as platform in a deployment.

When selecting a WAN Edge platform, it's essential to consider how each option fits with your organization's requirements:

- **Session Smart Routers** are recommended for deployments aiming to maximize network efficiency and performance, particularly in dynamic environments.
- **SRX Series Firewalls** are suited for organizations seeking to maintain strong security postures while gradually transitioning to SD-WAN.

Ultimately, your platform decision within Juniper Mist WAN Assurance should align with your organization's specific needs, current network architecture, and future goals. Under Mist, both platforms provide simplified management and Zero Touch Provisioning (ZTP) for easy setup, ideal for branches with limited IT support, ensuring a reliable and scalable network. The platform portfolio supports a variety of WAN types (MPLS, broadband, 3G/4G LTE) and speeds (1/10/25/40/100 Gbps). Whether you prioritize the cutting-edge routing capabilities, application performance and efficiency of the Session Smart Routers or the traditional security and transitional benefits of the SRX Series, both paths lead to a more intelligent, responsive, and user experience centric network.

Session Smart™ Router in WAN Assurance

When selecting a platform to use as WAN Edge in Juniper Mist, generally the Session Smart Routers are recommended for most SD-WAN deployments. The Session Smart Routers excels in several areas which are key to the successful implementation and operation of SD-WAN:

- Network efficiency—By forming overlays using secure vector routing (SVR) technology, an open standard for tunnel-free forwarding between router peers, Session Smart Routers ensures your

bandwidth is used as efficiently as possible. It maximizes performance for users and applications without unnecessary overhead.

- Rapid Failover—During instances of network failure or poor network quality, the Session Smart Routers gives the fastest possible failover, minimizing downtime and maintaining connectivity for critical applications and services.
- Rich Telemetry and Insights—The session-based architecture of the Session Smart Routers not only facilitates efficient data forwarding but also provides a wealth of telemetry data. This offers deep insights into the performance of users, applications, and the network itself. It serves as the foundation for Marvis and AI models in the Mist cloud, enabling continuous improvement and optimization of your network.
- Integrated Security—A router with security at its foundation, the Session Smart Routers provides Layer 3-Layer 4 security policy by default, and optional Layer 5-Layer 7 where advanced security at the branch is needed. Integrations with Security Service Edge (SSE) providers and the Session Smart Routers are simplified using Mist driven orchestration for simplicity. This offers a seamless transition to a modern cloud based Secure Access Service Edge (SASE) architecture.

For more SSR information, including datasheets, quick start guides, and hardware guides, go to Juniper Mist Supported Hardware—SSR Series Devices.

SRX Series in Juniper Mist WAN Assurance

While the Session Smart Routers Series is recommended for most deployments, the Juniper Networks SRX Series presents an ideal alternative for environments where traditional security mechanisms and transitional benefits are prioritized.

Key features include:

- Enhancing Existing Deployments—If your current infrastructure already utilizes SRX Series Firewalls, integrating them with Mist WAN Assurance can amplify their capabilities through the addition of cloud and AI technologies. This not only extends the life and utility of your current investments but does so in a way that aligns with modern, intelligent network management practices.
- Evolutionary Transition to SD-WAN—For organizations looking to transition to SD-WAN at a measured pace, starting with SRX Series as your WAN Edge platform within Mist WAN Assurance can offer a familiar yet powerful stepping stone. This allows your team to gradually adapt to SD-WAN technologies and principles while still gaining immediate benefits from AI and cloud management features.

For more SRX information, including datasheets, quick start guides, and hardware guides, go to Juniper Mist Supported Hardware—SRX Series Firewalls.

Request Help with a New Deployment

SUMMARY

If you need help with a new deployment of Juniper Mist™ devices, follow these steps to submit your request.

The Juniper Mist Support team provides onboarding assistance to help customers with new deployments. You'll get assistance with initial setup, configuration, and basic troubleshooting.

NOTE:

- These services do not include network design.
- Submit your request at least 48 hours in advance of your preferred appointment time.
- Available only for wireless, switching, and SD-WAN deployments.

To request help with a new deployment:

1. Click the question icon near the top-right corner of the Juniper Mist portal, and then click **Support Tickets**.

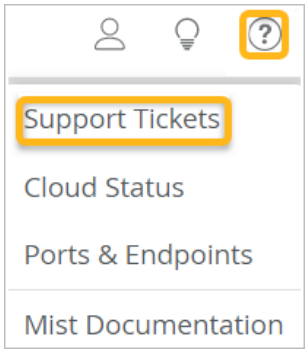

2. Click **Create a Ticket**.
3. For **Technology**, select the technologies that are applicable to your network and are included in the onboarding support program.

NOTE: Onboarding help is available only for wireless, switching, and SD-WAN deployments.

4. For **Ticket Type**, select **Onboarding Help for New Deployment**.

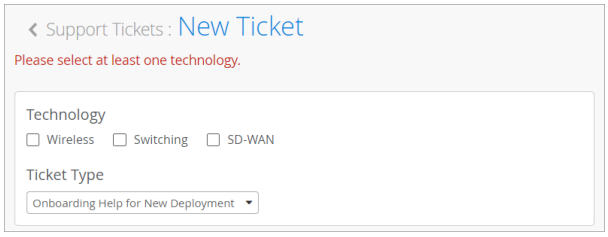

5. Enter a short **Ticket Summary** and a detailed **Description**.
6. Under **Checklist**, answer all the questions.

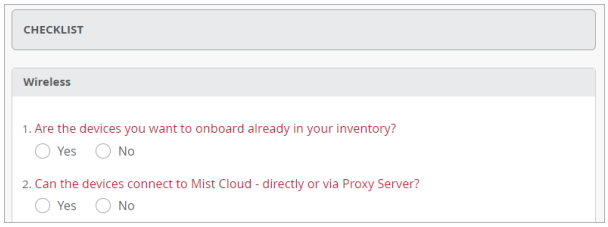

> **NOTE**: The checklist includes each technology that you selected at the top. Complete all questions in all sections that appear.
>
> As you complete the checklist, you can use suggested hyperlinks for self-help. If you find your answers in these links, you can cancel submitting this ticket.

7. Under **Schedule**, select the date, time, and time zone for your onboarding help session.

> **NOTE**: If you need to change your appointment after you submit your ticket, you can go to your list of open tickets, select this ticket, and reschedule.

8. Click **Submit** at the top right corner of the page.

   If this button is unavailable, ensure that you've entered all required information. Required fields have red labels.

The support team will contact you to conduct the help session.

Prepare to Configure Your WAN (Major Prerequisites)

SUMMARY

Complete these essential tasks to set up your organization and sites and onboard your WAN Edge devices, before you configure your WAN.

Table 1: Deployment Tasks and Links

Category	Task	More Information
Preliminaries	Gather information about your network, applications, user needs, devices, IP addresses, VLAN IDs, and so on.	

Table 1: Deployment Tasks and Links *(Continued)*

Category	Task	More Information
Organization Setup	Before you can configure your WAN or onboard your devices, you need to complete these tasks in the Juniper Mist™ portal: • Create your organization. • Set up at least one site. • Activate your subscriptions, including Juniper Mist WAN Assurance. • Configure your firewall to allow outbound Juniper Mist traffic. Recommended, but not required before you configure your WAN: • Add user accounts for other personnel who are working with you to deploy Juniper Mist. You can even enable limited access for the personnel who are installing devices. • Set up various security options as needed. For example, manage certificates, disable Juniper Mist support access, or enable Single Sign-On.	• Juniper Mist Quick Start • Firewall Configuration: Juniper Mist IP Addresses and Ports • Security Options

Table 1: Deployment Tasks and Links *(Continued)*

Category	Task	More Information
Additional Site Setup for WAN Edge Devices	• Root Password—For each site, enter a root password to use for your WAN Edge devices. If you don't define a root password in your site configuration, then the moment you activate a device to be managed by Mist Cloud, it will receive a randomly assigned root password. • Additional Licenses for SRX—For any sites where you're deploying SRX Series Firewalls, you'll also need an AppSecure and APP Track license. And you'll need to enable the APP Track license option in the WAN Edge Application Visibility section of your site configuration.	• Site Configuration Settings in the Juniper Mist Management Guide • Application Security User Guide for Security Devices
Device Installation	Install and claim your devices into your organization.	• SSR Datasheets, Quick Starts, and Hardware Guides • SRX Datasheets, Quick Starts, and Hardware Guides • Manage Your Inventory in the Juniper Mist Portal (in the Juniper Mist Management Guide)

2

CHAPTER

WAN Assurance Design

IN THIS CHAPTER

WAN Assurance Design Quick Start

SUMMARY

Use the information in this topic to understand what use case most closely matches your deployment type and navigate to the related content that is provided to learn how to get started.

IN THIS SECTION

With Juniper Mist™ WAN Assurance, there are a variety of design options that can be used to provide optimal routing for your specific deployment type. This topic describes some of the most commonly used design topologies for WAN Assurance. Read the information below to see which option best describes your deployment needs. The "What Do You Want to Do?" on page 36 table in this topic points you to additional resources that you can use to get started configuring your specific deployment pattern.

Hub and Spoke

This topology type includes a datacenter or other large site (hub) and branches (spokes). If your Networks, Users, and Applications need different accesses/access policies, this is the topology type you should follow. This requires a Spoke WAN Edge template to be configured where you can configure your hub and spoke topology.

Let's say you have several bank branches (spokes) in your deployment that need to reach out to the datacenter (hub). Each of the teller's desks need access to applications, such as the point of sale system, which is located at the data center, as well as applications that reside on the internet. Or maybe the security cameras from each branch location of your deployment need access to the surveillance system. If this sounds like your deployment, see "Configure a WAN Edge Template" on page 87 and "Configure Path Selection from Hub-to-Spoke with Traffic Steering" on page 139.

Note that Juniper Mist WAN Assurance provides flexibility with regard to the location of your hub. You can have hubs located physically in a datacenter, virtually in cloud environments, or in collocation facilities.

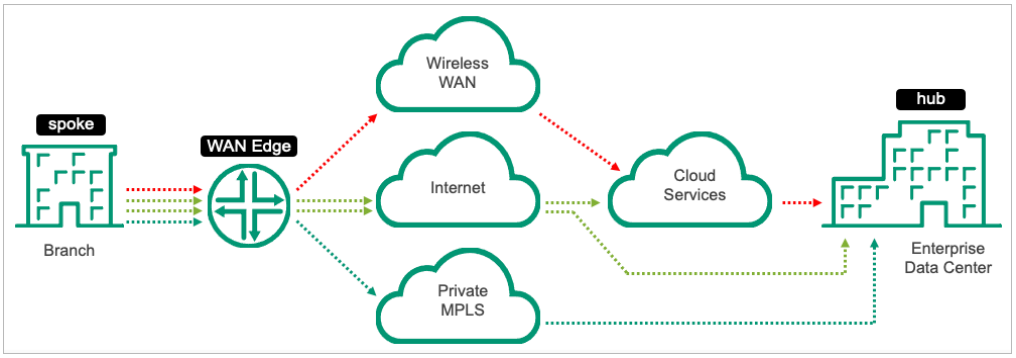
spoke
Branch
WAN Edge
Wireless WAN
Internet
Private MPLS
Cloud Services
hub
Enterprise Data Center

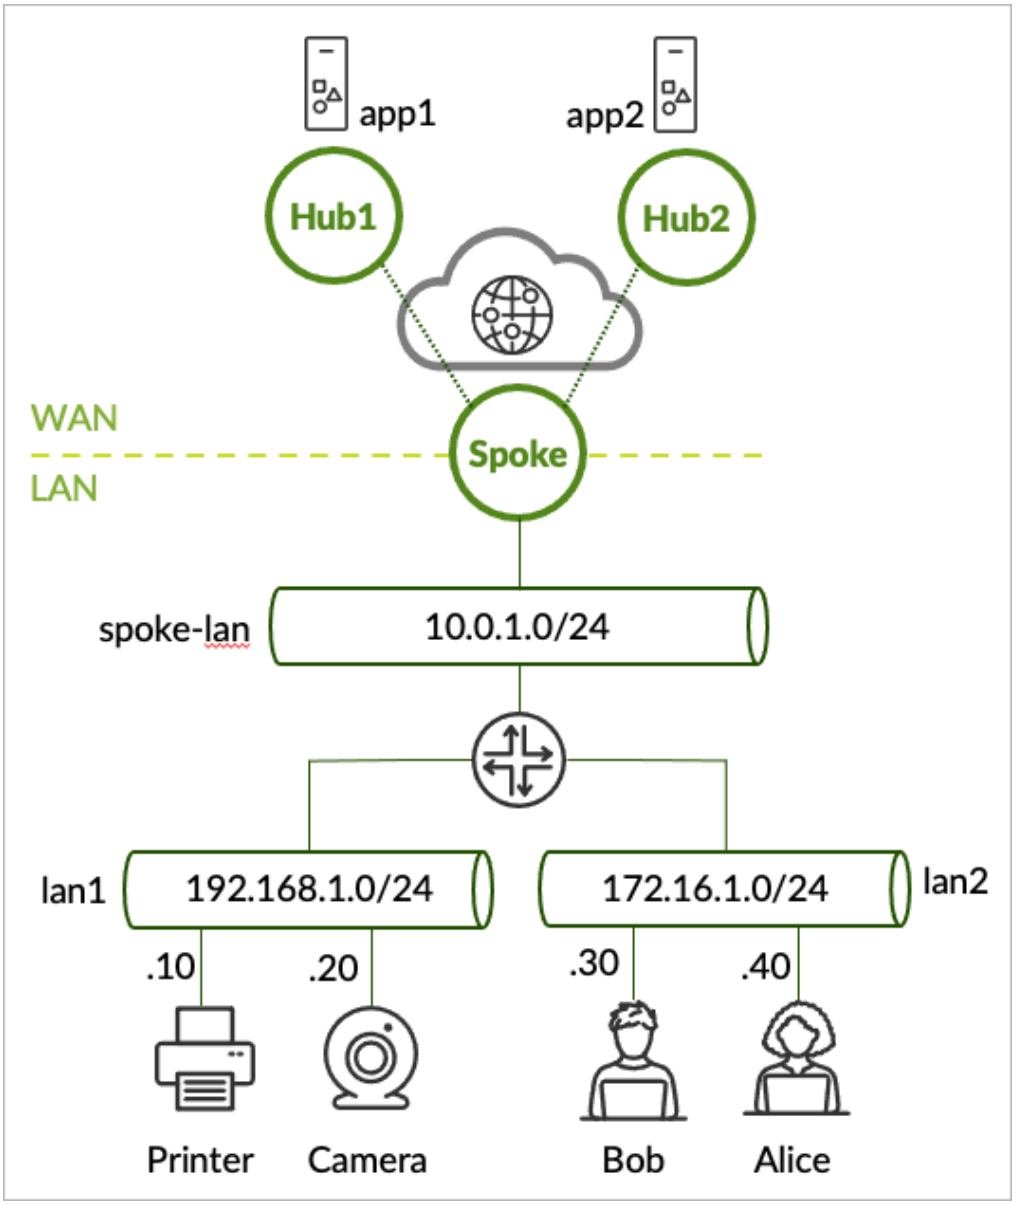
app1
app2
Hub1
Hub2
WAN
LAN
Spoke
spoke-lan
10.0.1.0/24
lan1
192.168.1.0/24
172.16.1.0/24
lan2
.10
.20
.30
.40
Printer
Camera
Bob
Alice

Mesh

A hub and spoke topology tends to cover most use cases. However, for customers who have concerns about latency between sites, a mesh WAN topology can provide interconnectivity across multiple sites and devices with minimal latency. This is ideal for large deployments with various locations to achieve path optimized connectivity. A mesh topology provides lower latency paths by providing more direct connections between sites.

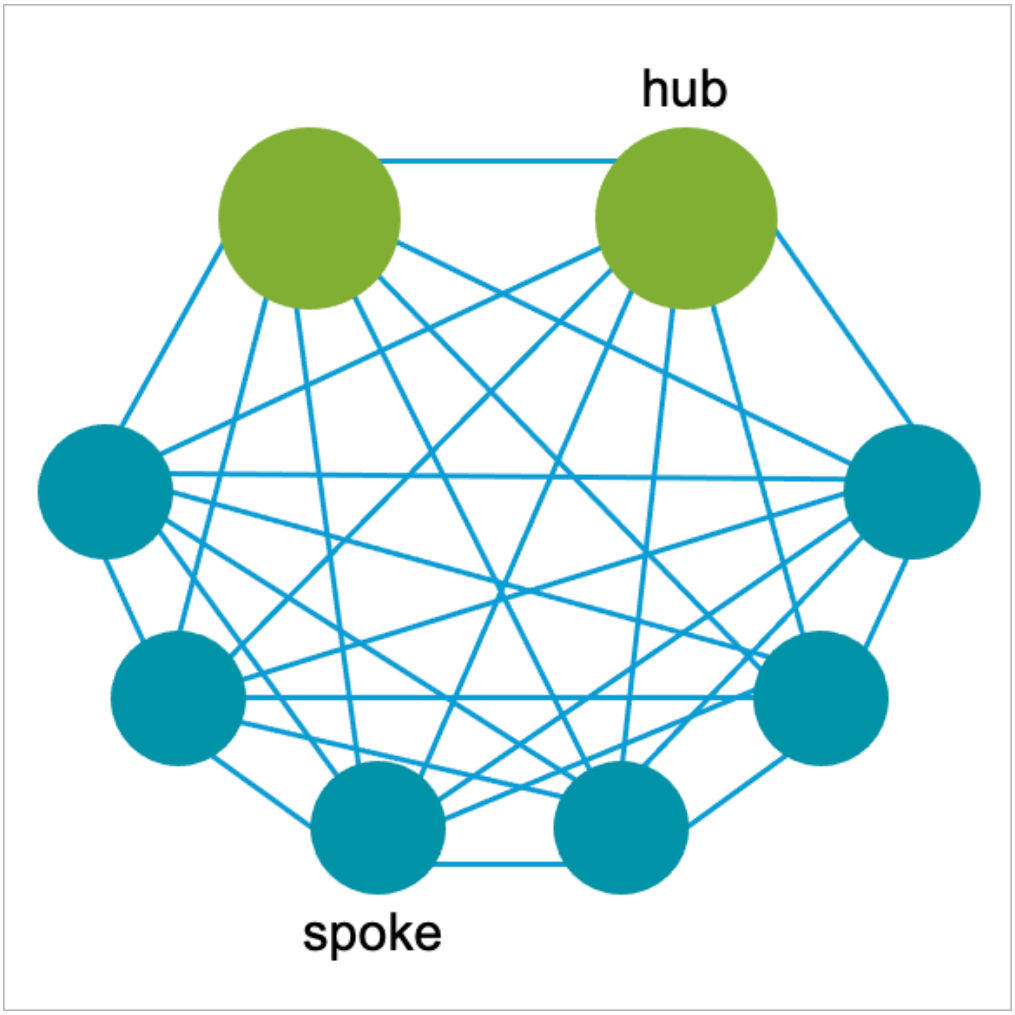

NOTE: As mesh topologies scale out, they can become more expensive from a device perspective. A hub and spoke topology is recommended if you do not have excessive traffic traversing between sites.

You create "hub profiles" on page 82 for WAN Edge devices at hub sites. You create "WAN Edge templates" on page 87 for WAN Edge devices at spoke sites. Hub WAN interfaces create overlay

endpoints for spokes. Spoke WAN interfaces map the appropriate Hub WAN interfaces, defining the topology. Hub profiles drive the addition and removal of paths on your overlay.

NOTE: Devices that are part of a hub and spoke overlay cannot be included in a mesh topology, and vice versa.

Security Service Edge in Cloud with Standalone Sites

In this topology type, you have standalone sites that need to reach the Security Service Edge (SSE) in the cloud. While your security edge resides in the cloud, you have standalone WAN Edge devices at each of your sites, such as at each of the individual coffee shops in a coffee shop chain. In this scenario, your standalone WAN Edge devices steer applications needing advanced security from the cloud to the SSE in the cloud with the ability to break out other applications directly to the internet.

You must configure a standalone WAN Edge Template to accommodate your standalone WAN Edge devices at your sites. See "Configure a WAN Edge Template" on page 87.

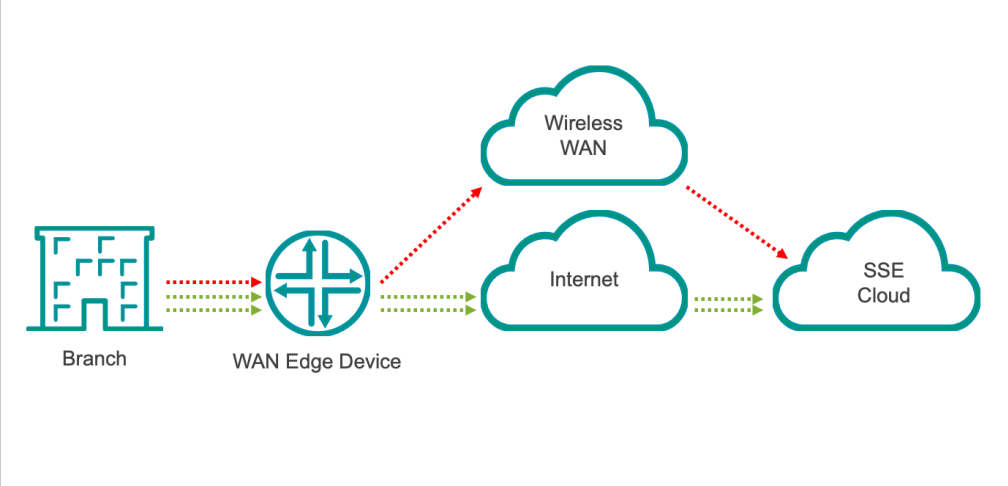

Standalone WAN Edge Devices

In this topology design, there is no hub or cloud that your WAN Edge devices send traffic to, but rather, you use standalone WAN Edge devices at your individual sites and use the on-box security components that come standard on those devices. This deployment type requires you to configure a "Standalone" WAN Edge template. See "Configure a WAN Edge Template" on page 87.

What Do You Want to Do?

Table 2: Top Tasks

Design/Topology Type	Use these resources:
Hub and Spoke Deployment *This topology type includes a datacenter or other large site (hub) and branches (spokes). Follow this topology type if your Networks and Users need different accesses. This requires you to configure a Spoke WAN Edge template.*	• "Configure a WAN Edge Template" on page 87 • "Configure Path Selection from Hub-to-Spoke with Traffic Steering" on page 139
Mesh Topology *For customers who have concerns about latency between sites, a mesh WAN topology can provide interconnectivity across multiple sites and devices with minimal latency. This is ideal for large deployments with various locations to achieve path optimized connectivity.*	• "Hub Profiles" on page 82 • "Configure a WAN Edge Template" on page 87
Security Service Edge in Cloud with Standalone Sites *In this topology type, you have standalone sites that need to reach the Security Service Edge (SSE) in the cloud. Your standalone WAN Edge devices steer applications needing advanced security from the cloud to the SSE in the cloud.*	• "Configure a WAN Edge Template" on page 87
Standalone WAN Edge Devices *In this topology design, there is no hub or cloud that your WAN Edge devices send traffic to, but rather, you use standalone WAN Edge devices at your individual sites and use the on-box security components that come standard on those devices.*	• "Configure a WAN Edge Template" on page 87

RELATED DOCUMENTATION

SRX JVD

WAN Assurance Advanced Design Overview

SUMMARY

Get started designing Full Stack and high availability deployments for your WAN Edge devices.

IN THIS SECTION

WAN Assurance Design Overview

The WAN Assurance design chapter of this guide takes you through the advanced configurations for your hub and spoke SD-WAN deployments. There are unique guides for high availability and Full Stack implementations. The Full Stack design combines AI-driven wired, wireless, and WAN solutions. The high availability design allows for redundancy between your WAN Edge devices. Both designs require you to have completed a basic WAN Assurance configuration. For an overview of the basic WAN Assurance designs, see "WAN Assurance Design Quick Start" on page 32.

Full Stack Design for WAN Edge Devices

WAN Edge devices deployed in Juniper Mist™ WAN Assurance for Full Stack deployments follow the "Full Stack Design for WAN Edge Devices" on page 40. This design requires an SD-WAN deployment administrator to add WAN Edges, switches, and APs for an end-to-end Full Stack solution that combines Juniper Mist WAN, wired, and wireless.

High Availability Design for WAN Edges

WAN Edge devices deployed in Juniper Mist™ WAN Assurance can be configured for high availability deployments to ensure redundancy between interfaces and devices. You must first configure the hub and spoke topology referenced in this guide (see "Configure a WAN Edge Template" on page 87 and "Configure Path Selection from Hub-to-Spoke with Traffic Steering" on page 139).

After you've configured the hub and spoke topology, follow the procedure in the "High Availability Design for WAN Edge Devices" on page 60 topic to configure interface and device redundancy for your WAN Edges.

Routing Configuration Design

SUMMARY

Learn about different WAN Assurance routing designs to see what configuration best suits your deployment needs.

IN THIS SECTION

This topic provides an overview of the different routing designs in Juniper Mist™ WAN Assurance. Each section guides you to the documentation that you must follow in order to complete the configuration for each. Different types of routing must be configured depending on your use case.

Basic Overlay

The basic overlay design is where there is no dynamic routing that you need to integrate with. In this type of configuration, you are simply sending traffic to the overlay, or breaking out traffic locally to your WAN links. With this, dynamic routing to the overlay is automatically taken care of for you by Mist when the "Advertise to the Overlay" checkbox is selected in the Network configuration. See "Configure a Hub Profile" on page 82.

Dynamic Routing

With the dynamic routing design, you already have dynamic routing set up with your datacenter, or with a WAN provider that you are integrated with. This use case is for those who are integrating with an existing network, such as if you have a network with configured Mist-managed devices, and are now looking to configure BGP or OSPF on top of it.

- "BGP" on page 178
- "OSPF" on page 172
- Appendix: Building an Extended Topology with Hub Overlay and BGP Peering

Dynamic Routing for Full Stack Campus Fabric

With this design, you already have dynamic routing set up with your datacenter, or with your WAN provider integration. This use case is for those who are integrating with an existing network, and are now looking to run dynamic routing such as BGP or OSPF into the campus fabric as part of your Full Stack campus fabric deployment.

- "BGP" on page 178
- "OSPF" on page 172
- Campus Fabric WAN Router Integration—Use Case and Reference Architecture

RELATED DOCUMENTATION

SSR JVD

SRX JVD

Full Stack Design for WAN Edge Devices

SUMMARY

Use this Full Stack guide to set up your Juniper WAN Edge devices in concert with Juniper Mist Access Points (APs) deployed in Wireless Assurance, and Juniper EX Series Switches deployed in Wired Assurance.

IN THIS SECTION

Overview

IN THIS SECTION

The Juniper Mist Full Stack enables you to expand your network capabilities by integrating Mist Access Points (APs), EX Switches, and WAN Edge devices. This design brings all your network devices into a cohesive onboarding, monitoring, and troubleshooting dashboard.

The Full Stack design begins with WAN Edge devices deployed in Mist WAN Assurance. After completing the Juniper Validated Design (JVD) topology (see SSR JVD and SRX JVD), you should already have a WAN Edge device deployed in a hub and spoke network. The WAN Edge device serves as the

foundation for building out your entire network with the Full Stack. The Full Stack is specifically designed for a branch that utilizes Juniper equipment.

For successful implementation of the Full Stack, you'll need at least one Juniper EX Switch to onboard into the Mist cloud. If you plan to do advanced testing with virtual circuits, two EX Switches is ideal. Additionally, you can incorporate a Mist AP into the setup to enhance the wireless capabilities of the network. Onboarding those into your LAN network for Mist management gives administrators the ability to monitor and manage their WAN Edges, switches, and APs all from the Mist dashboard.

The Figure 1 on page 41 shows the topology of the Juniper Mist Full Stack with WAN Assurance at its core. The devices in green are meant to show the different types of devices that make up a Full Stack. You can have a number of APs, Switches, and WAN Edges in a single Full Stack deployment.

Figure 15: Juniper® Mist Validated Design - Mist Full Stack (WAN Assurance with Wireless and Wired Assurance)

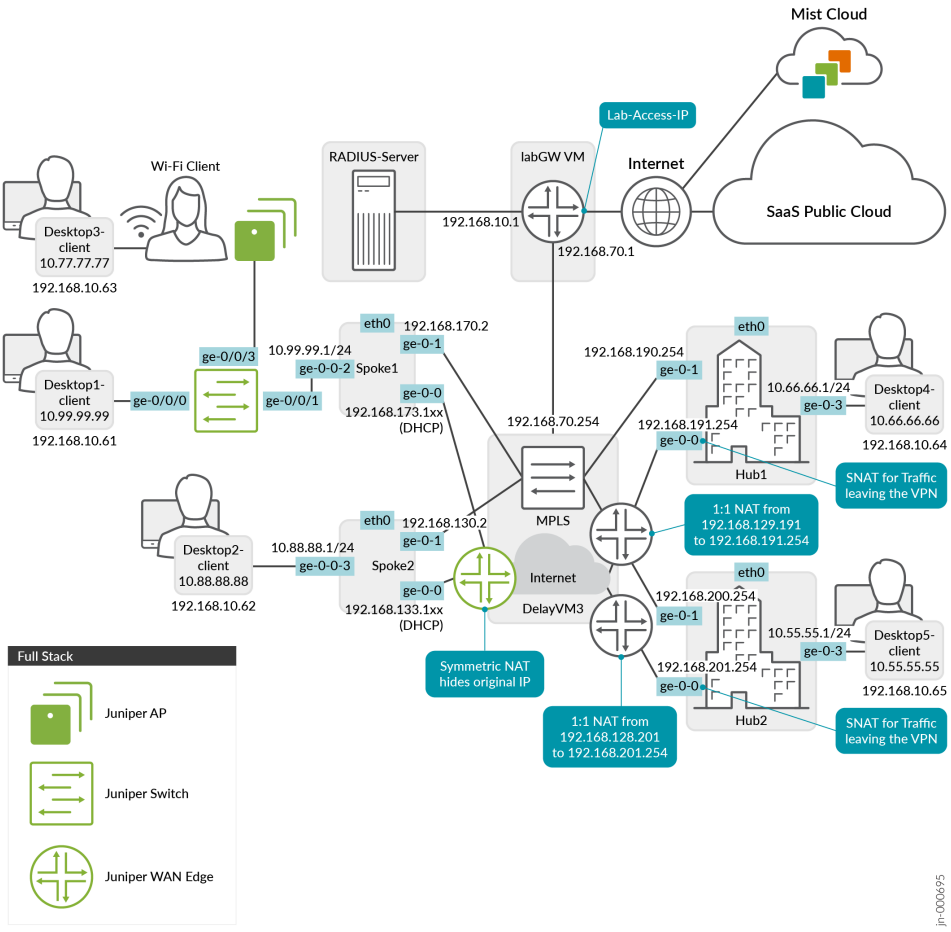

Requirements and Considerations

To get started, you'll need to alter some of the interfaces found in the Juniper Mist WAN Configuration Guide topology. We'll show you how to do this using a Spokes configuration template found in the Mist WAN Configuration Guide. See "WAN Assurance Configuration Overview" on page 75.

NOTE: Devices are expected to get out to the Mist cloud using the on-box Internet Service Provider (ISP) link. The device will attempt to retrieve a DHCP address from that ISP. Your devices will connect to Mist on these endpoints by default out of the box. However, if the device is behind a firewall, it may not be able to reach the Mist cloud. See *Juniper Mist Firewall Ports and IP Addresses for Firewall Configuration* for the firewall ports to open based on your device.

The device uses a hostname to connect on and needs to resolve that hostname using outbound DNS access to 8.8.8.8. Along with the previously mentioned DNS, the SSR will attempt to use 1.1.1.1 and the SRX will attempt to use 8.8.4.4. These are used as the DNS servers over the ISP link mentioned above to resolve the endpoints. Once the device connects to the Mist cloud, your configured DNS is used by the device.

If your SRX Series Firewall is showing as disconnected when it is online and reachable locally, you can troubleshoot the issue using "Troubleshoot Disconnected SRX Series Firewalls" on page 411.

ATTENTION: Follow the sections and steps below as your step-by-step procedure for completing a Full Stack setup.

Onboard a WAN Edge Device

This procedure assumes that you have already onboarded your WAN Edge device to the Mist cloud. If you need to onboard a WAN Edge device, follow the steps outlined in Cloud Ready SSR Devices Quick Start Guide or Cloud Ready SRX Series Firewalls Quick Start Guide, then return to this procedure.

Create a Site

This procedure assumes that you have already created a site, which you will later assign your WAN Edge template to in order to complete the Full Stack design. If you need to create a new site, follow the steps outlined in *Configure a Site*, then return to this procedure.

Create a Hub Profile

If the WAN Edge device in your Full Stack design is part of a hub and spoke topology, you must "configure a hub profile" on page 82 for any WAN Edge devices at hub sites. Hub profiles create an overlay and assign a path for each WAN link on the overlay.

You can also "Create a Hub Profile by Cloning" on page 84 an existing one to save time.

> **NOTE**: You'll create hub profiles for WAN Edge devices at hub sites. You'll create WAN Edge templates for WAN Edge devices at spoke sites. Hub WAN interfaces create overlay endpoints for spokes. Spoke WAN interfaces map the appropriate Hub WAN interfaces, defining the topology. Hub profiles drive the addition and removal of paths on your overlay.

However, if your topology does not use an overlay, skip to the "Create a New Spokes Configuration Template" on page 43 section below.

Create a New Spokes Configuration Template

The WAN Edge device at the spoke site in your Full Stack design is configured via the WAN Edge template. You can create a new spoke WAN Edge template, or clone an existing spoke template and then make the necessary changes.

The most efficient way to configure a WAN Edge template is to create a spoke WAN Edge Template from device-model. This automatically sets the configuration for you, and you can adjust the configuration in any way necessary.

1. In the Juniper Mist™ portal, click **Organization** > **WAN** > **WAN Edge Templates**.
   A list of existing templates, if any, appears.
   To learn how to create a spoke template, see "Configure a WAN Edge Template" on page 87.

   Alternatively, device-specific templates automatically assign many aspects of the configuration, including WAN and LAN interfaces for connectivity as well as traffic steering profiles and applications policies to set rules for the traffic on your device.

   > **NOTE**: The device-model template you select does not have to match the model of the device hardware itself. You are not confined to the configuration that device-specific templates automatically configure for you. You can customize the templates in any way that you need.

To configure a device-specific template, navigate to **Organization** > **WAN** > **WAN Edge Templates**, then in the New Template pop-up window, enter a name, select a template type, and finally, select **Create from Device Model**.

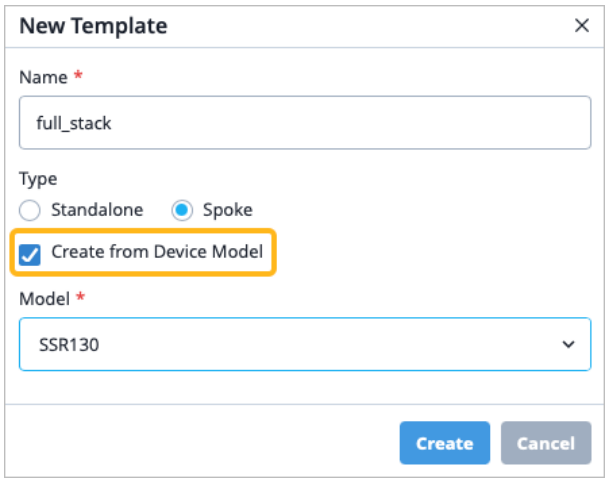

2. ***Optional:*** You can also clone an existing template. Click **More** and select **Clone**.

Figure 16: Selecting Clone Option for Template

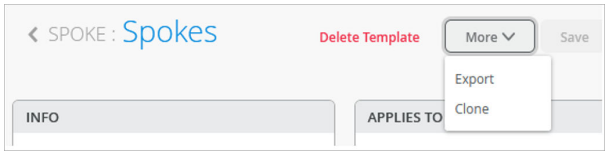

a. Enter a name, such as **full-stack**, and click **Clone**.

Figure 17: Saving Cloned Template

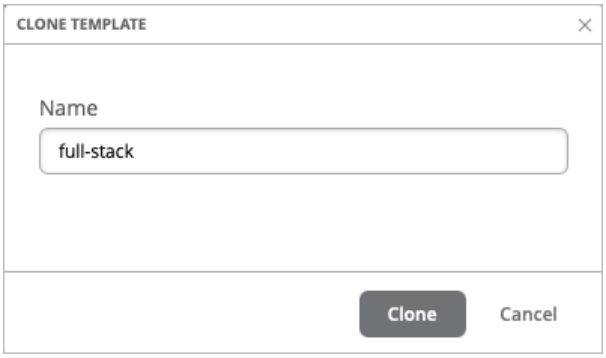

TIP: Refresh your browser after cloning. This ensures objects displayed are truly refreshed.

Edit the Template as Needed

When you create a WAN Edge template from a device model, aspects of the configuration are automatically set for you, but you can adjust any of that configuration as needed. This section provides you with some examples of configuration that are commonly edited or added to the template as part of the Full Stack configuration.

1. On the WAN interface configuration section, edit or add any WAN interfaces as needed.

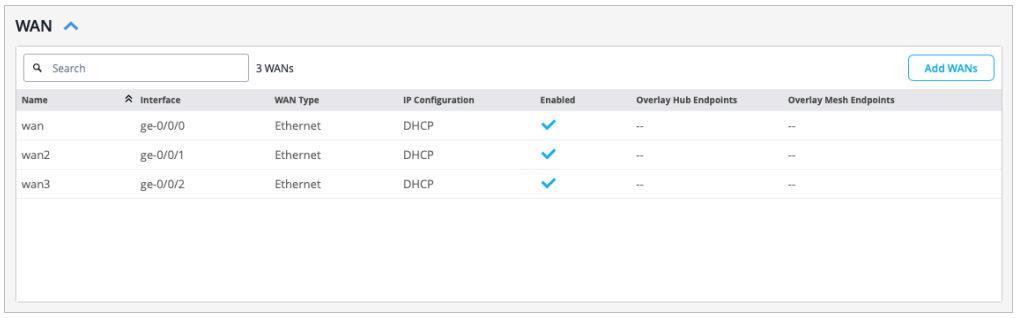

NOTE: The ge-0/0/0 WAN interface that was configured is what you plug into the ISP so that the ISP can then provide you with an address via DHCP. This DHCP address allows the device to use that provided WAN link to phone home to the Mist cloud to get the configuration from Mist using the ports and endpoints listed in *Juniper Mist Firewall Ports and IP Addresses for Firewall Configuration*.

Juniper Mist offers interface flexibility, meaning that you can have the same WAN interface configured on multiple devices. For example, you may want to configure the same interface on two different devices so that if one device fails, the other still has a physical link to your ISP. You can configure this in the **Interface** field of the WAN Configuration side-panel simply by entering comma separated values. This allows you to group these interfaces together to achieve redundancy.

2. On the LAN interface configuration section, note the LAN interface that got configured.

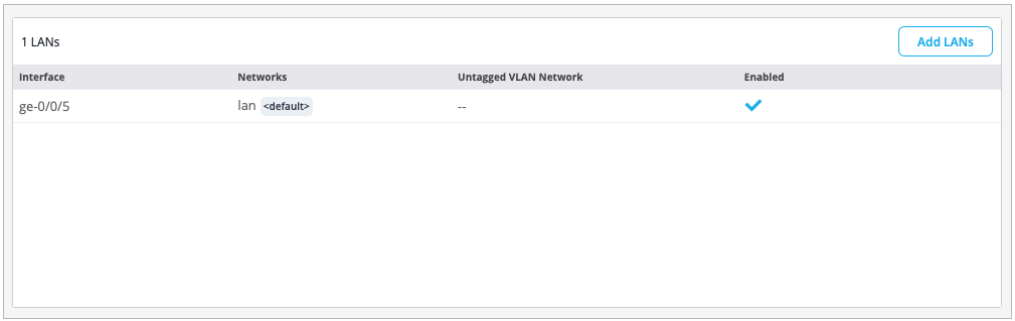

Select the LAN interface to open the Edit LAN Configuration side-panel.

Notice that the Network field is already configured with the LAN network. This is the network that automatically gets created as part of the device-model template and ensures that all of the devices in the Full Stack are able to access the Mist cloud in the same way using the same policies. There is a DHCP server running on this LAN interface that hands out addresses in the 192.168.1.0/24 address space.

You can navigate to the Network configuration from the left menu > **Organization** > **Networks**, then select the appropriate network.

NOTE: The Access to MIST Cloud checkbox is selected by default as part of the Network configuration. This setting allows other endpoints on the network to access the same services and policies that are built into the WAN Edge device so that the devices can automatically connect to the Mist cloud.

If you navigated to the Network configuration, you can navigate back to the WAN Edge template, from the left menu, navigate to **Organization** > **WAN Edge Templates**, then select the appropriate template.

Figure 18: Modify LAN Interface Configuration

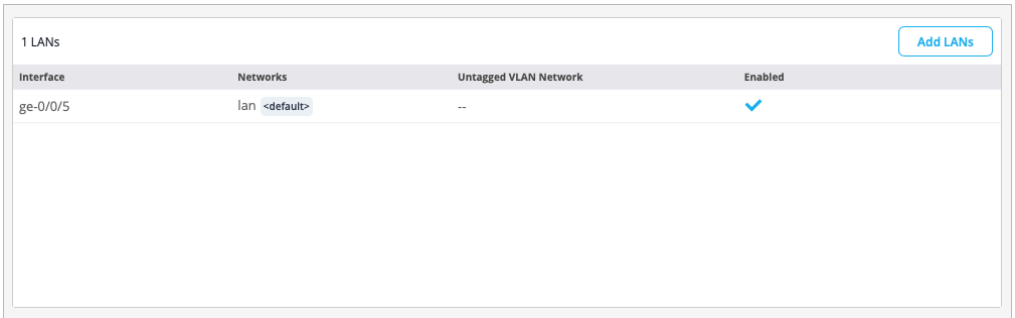

3. Customize the LAN interface as needed, or add new LAN interfaces. Step 4 below is a sample of how you may configure a new LAN interface with a guest network.

4. You may find that you need to add a guest network to a LAN interface to allow guest users onto your corporate Wi-Fi. You first create the guest network, then add a new LAN interface that you configure that guest network on.

   a. Navigate to the left menu to **Organization** > **WAN** > **Networks**.

   b. In the top-right corner of the page, click **Add Networks**, then fill in the fields. You can use the samples in Table 1 below to guide you.

Table 3: Sample Guest Network Configuration Fields

Field	How to Configure
Name	Guest
Subnet IP Address	Enter the subnet IP address you want this network to use (Example: 172.16.1.0).
Prefix Length	Enter the prefix length for the IP Subnet (Example: 24)
VLAN ID	100
Access to MIST Cloud checkbox	Deselect the checkbox, as the guest network traditionally does not have devices that need to connect to the Mist cloud.

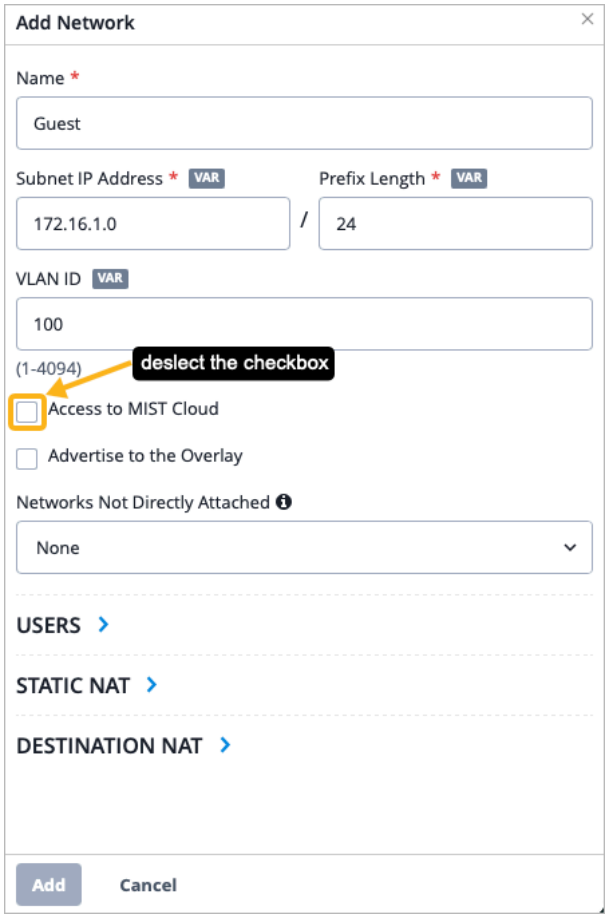

a. Click **Add** from the bottom of the Add Network side-panel.

b. Navigate back to the WAN Edge template.

c. From the LANs section, click **Add LANs**, then fill in the fields. You can use the samples in Table 2 on page 48 below to guide you.

Table 4: Sample LAN Configuration Fields for Guest Network

Field	How to Configure
Interface	ge-0/0/6
Network	Guest (select your newly created network from the drop-down list).

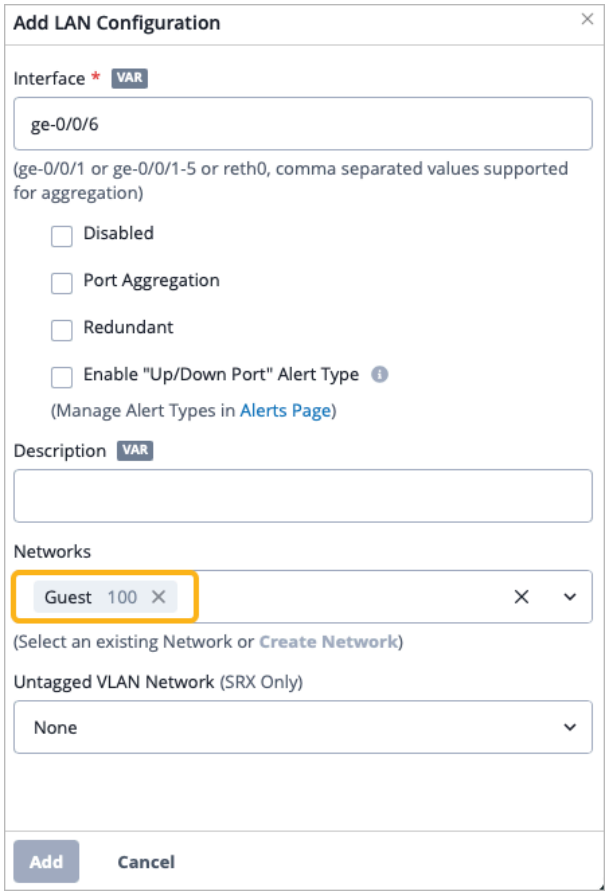

You must also create an Application Policy to allow guest traffic. Continue with the procedure to ensure completeness of your Full Stack design, then create an application policy as described in step 7 below.

a. Click **Add** from the bottom of the **Add LAN Configuration** side-panel.

b. Navigate to the IP Config section and click **Add IP Config** and add the guest network and the corresponding IP address.

5. Configure DHCP for the guest network. Navigate to the **DHCP Config** section of the WAN Edge template.

a. Click **Add DHCP Config**.

b. Select the guest network from the **Network** dropdown. Configure the fields. You can use the samples in Table 3 on page 50 below to guide you.

Table 5: Sample DHCP Configuration Fields for Guest Network

Field	How to Configure
Name	Guest
DHCP	Server
IP Start	Refer to the IP Start in this DHCP configuration (Example: 172.16.1.2).
IP End	Refer to the IP End in this DHCP configuration (Example: 172.16.1.254).
Gateway	Refer to the gateway address for this DHCP configuration (Example: 172.16.1.1).
DNS Servers	Refer to the DNS Server that your network will use to translate IP addresses (Example: 8.8.8.8,1.1.1.1).

You should now see your new LAN configuration in the list.

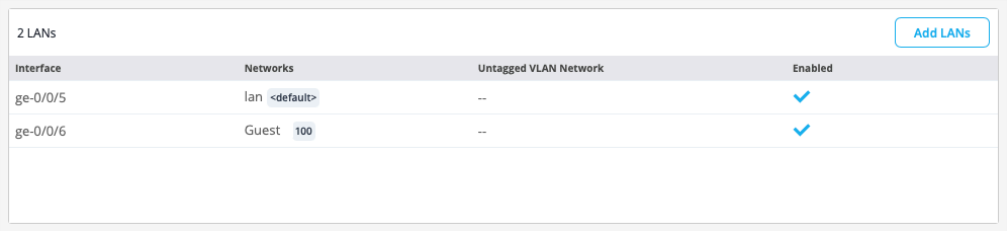

The guest network must also be configured on any switches and APs that are part of your Full Stack design.

6. Navigate to the **Traffic Steering** section of the template. Notice a traffic steering rule was configured for you already as part of the device-specific template, and is configured to have traffic use the WAN path.

   Edit the existing traffic steering or add new traffic steering rules as needed.

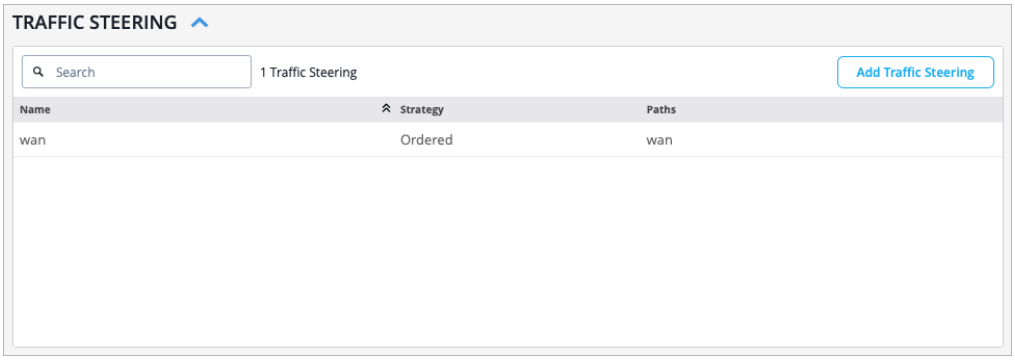

Recall that there are multiple WAN interfaces configured as part of the device-specific template. You may find yourself using those other WAN interfaces when you configure a second WAN path in Traffic Steering. Or, you may find that you need to configure another traffic steering policy to send a certain type of traffic out of the other WAN link.

To learn how to create a new traffic steering policy, see "Traffic Steering Rules" on page 134 and "Configure Path Selection from Hub-to-Spoke with Traffic Steering" on page 139.

7. Navigate to the **Application Policies** section of the template. If you used a device-specific template, notice that an application policy was configured for you. You can edit this policy as needed, or you can create a new one.

Figure 19: Application Policy Configuration

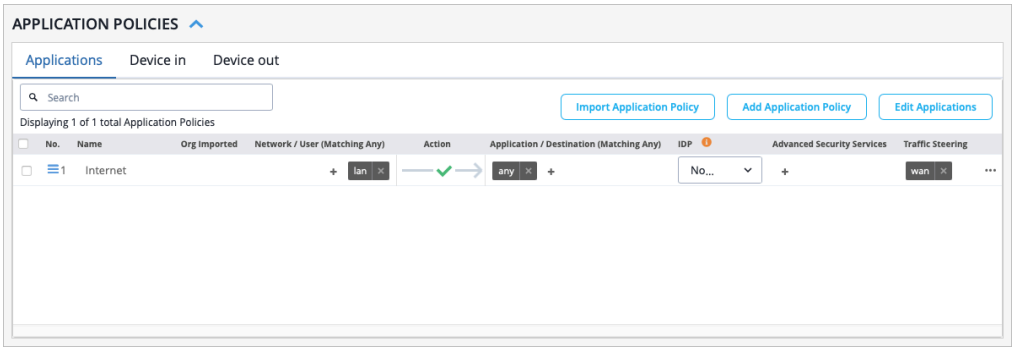

You must now create an application policy to allow guest traffic as per the Guest network configured in step 4 above. Examples are in the sample table below.

a. From the Application Policies section, click **Add Application Policy**.

b. Fill in the various fields. You can use the samples in Table 4 on page 51 below to guide you.

Table 6: Sample Application Policy for Guest Network Traffic

Field	How to Configure
Name	guest-internet
Network	Guest
Action	Allow
Application/Destination	Any

Table 6: Sample Application Policy for Guest Network Traffic *(Continued)*

Field	How to Configure
Traffic Steering	wan

You should now see your new application policy that allows guest traffic.

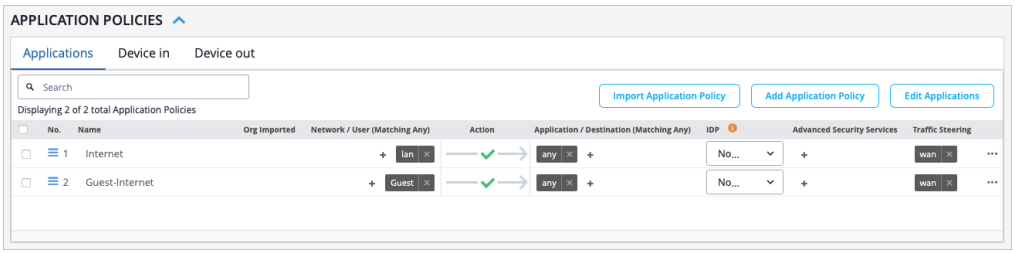

For more information on how to create a new application policy, see "Application Policies" on page 144.

8. Click **Save** in the top right corner of the template to save your changes.

Assign the New Template to a Site

1. At the top of the WAN Edge Templates page, click **Assign to Sites** under the Assign to Sites section.
2. Follow the procedure in "Assign Sites" on page 91.
3. Review the Site column on the WAN Edge Templates page as shown in Figure 6 on page 52.

Figure 20: Details of WAN Edge Template

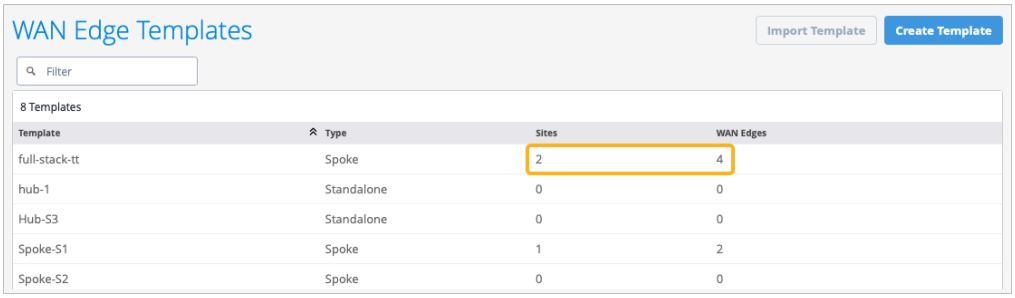

Assign the WAN Edge Devices to a Site

As part of completing the Full Stack design, you must now assign the WAN Edge devices to the same site you applied your template to in the previous section. This ensures that the WAN Edge devices in your Full Stack get the necessary configuration.

1. From the left menu, navigate to **Organization** > **Site** > **Inventory**. You should see your devices listed there. Notice that they do not have a site assigned to them.
2. Select the devices that you want to configure as part of your Full Stack.
3. Click the **More** button in the top right corner of the page, then select **Assign to Site**.

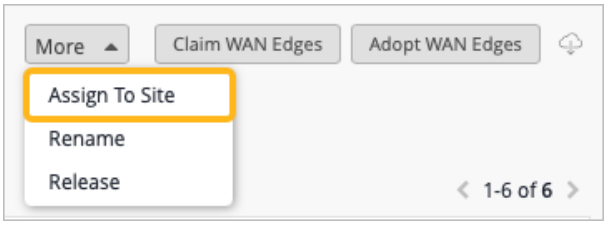

4. Select a site from the drop-down in the **Assign WAN Edges** pop-up. You should select the same site that your WAN Edge template is assigned to.
5. Click **Assign to Site** at the bottom of the pop-up.
   You should now see that the devices have been assigned to the site.

Add Your Switch to the Full Stack Design

Now it is time to onboard your switch and add it to your Full Stack design. For details on how to onboard your switch, refer to:

- The product documentation for your switch in the Juniper TechLibrary.
- To get a new cloud-ready EX switch up and running in the Juniper Mist cloud portal, follow the onboarding steps and links in Cloud-Ready EX and QFX Switches with Mist and *Onboard Switches to Mist Cloud*.

As part of completing the Full Stack design, you must now assign the switch to the same site as the other devices in the Full Stack, and must select the ports which the other devices will connect to your switch on. You must also configure the same guest network you configured on your WAN Edge.

1. Assign the switch template to the site.
2. From the left menu, navigate to **Organization** > **Admin** > **Inventory**.
3. In the Inventory page, ensure the inventory view is set to **org (Entire Org)** so that you see all your devices.

Figure 21: EX Series Switch in Inventory

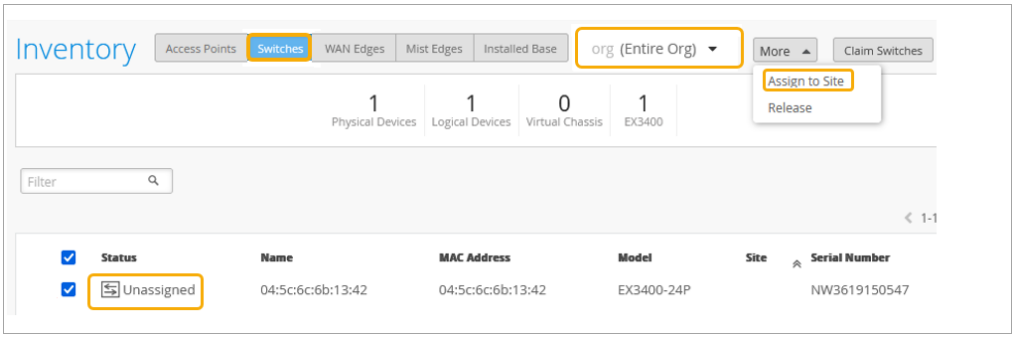

4. Select your new switch, then click **More** > **Assign to Site**.
5. In the Assign Switches pop-up:
   - Select the appropriate site.
   - Disable the **Manage configuration with Mist** option. You can enable this option at a later stage if required.

Figure 22: Select the Site for the Switch

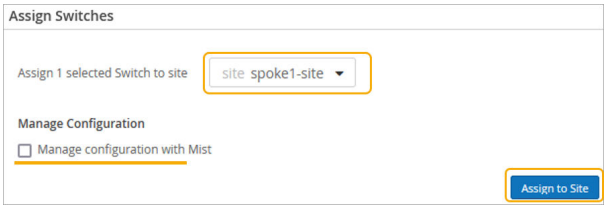

6. Click **Assign to Site**.
7. Confirm the change in the Assign Switches pop-up once you assign the device to the site.

   You can see the site name under **New Site**.

Figure 23: Assigned Switch to Site Details

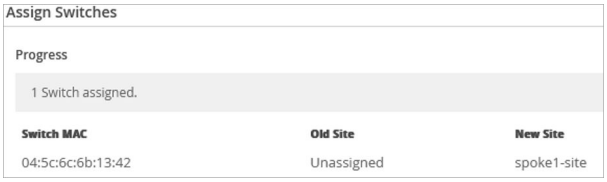

8. From the left menu, navigate to **Switches** and select the same site that you used in the previous steps. The page displays the list of switches assigned to the site.
9. Click the switch to open the switch configuration page.

Figure 24: Select Assigned Switch for Modification

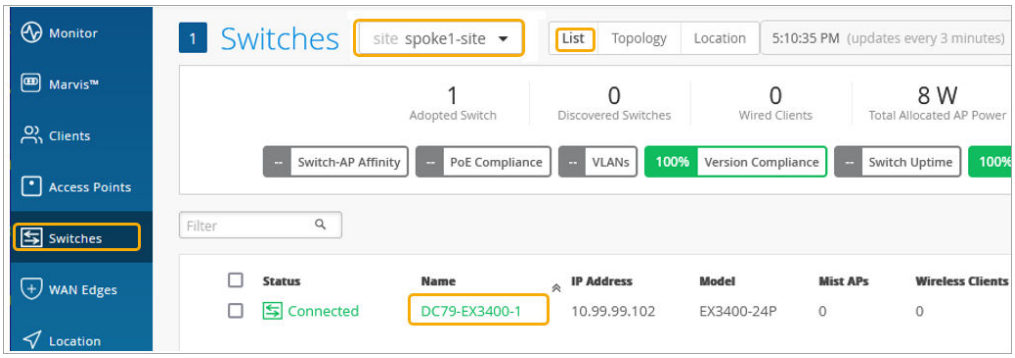

10. Verify the device name, then scroll down to the **Switch Configuration** section and check **Enable Configuration Management**.

Figure 25: Configuration of Assigned Switch

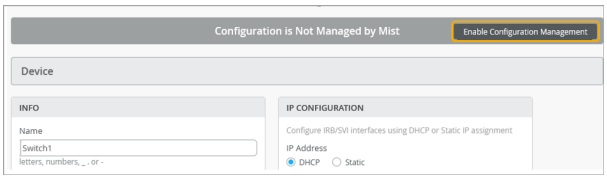

You must now select the ports the other devices in the Full Stack will use to connect to your switch.

11. Navigate to the **Port** section, then click **Add Port Configuration**.

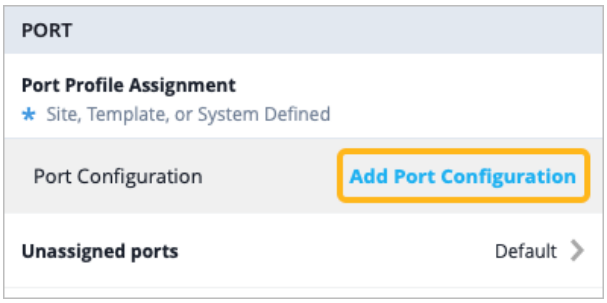

12. In the New Port Configuration page, configure the following options to indicate the port the other devices will connect to the switch on:

- Set the **Port IDs** as **ge-0/0/0**.
- Select the existing **Configuration Profile** as **Uplink**, then select the checkbox to save the changes.

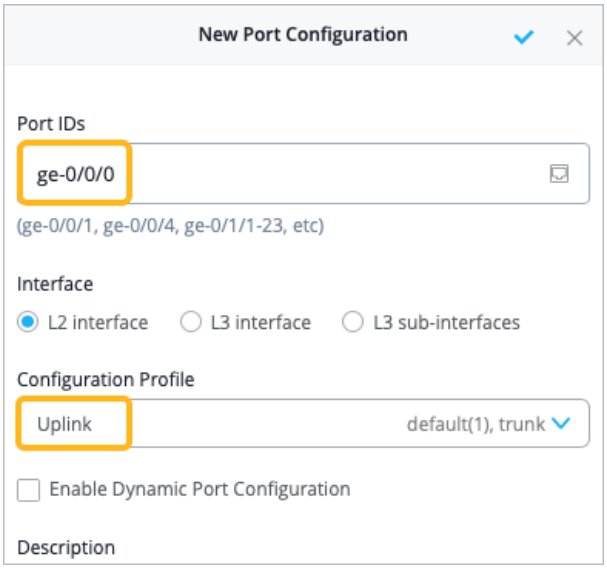

13. Add another port configuration, this time for your AP. Click **Add Port Configuration**, then:
    - Set the **Port IDs** as **ge-0/0/1**.
    - Select the existing **Configuration Profile** as **AP**. Select the checkmark to save the changes.

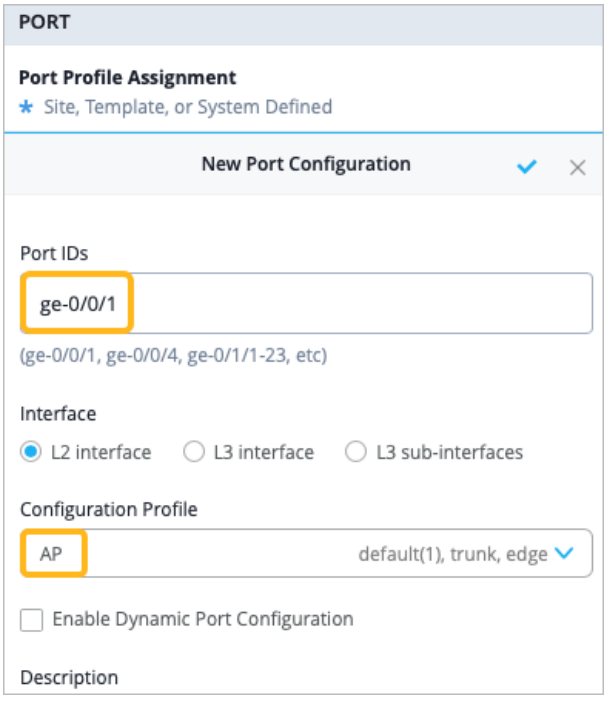

14. Navigate to the Networks section and click **Add Network** to add the same guest network you configured on your WAN Edge. You can use the samples in Table 5 on page 57 below to guide you.

Table 7: Sample Guest Network Configuration in Switch Templates

Field	How to Configure
Name	Guest
VLAN ID	100
IPv4 Subnet	Enter the subnet IP address you want this network to use (Example: 172.16.1.0).

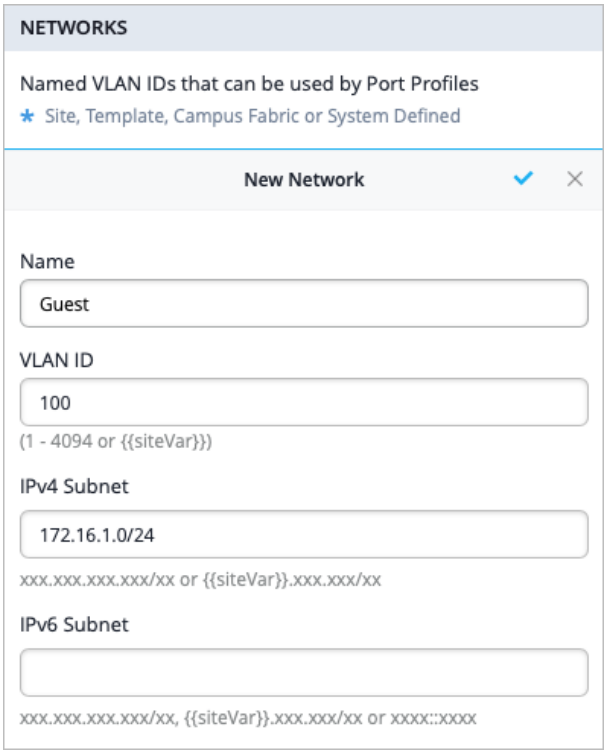

15. Save your changes.

You've now added a Juniper switch to your Mist Full Stack deployment.

Add Your APs to the Full Stack Design

Now it is time to onboard your AP and add it to your Full Stack design. Refer to:

- The product documentation in the Juniper TechLibrary for details on how to onboard your AP.
- For the steps to get a new AP up and running in the Mist cloud, see the Juniper Mist Access Points Quick Start Guide.

To complete the Full Stack design, you must now assign the switch to the same site as the other devices in the Full Stack. You must also configure your AP with the same guest network you configured on your WAN Edge and switch.

1. Add your AP to the same site that the WAN Edge device and switch in your Full Stack were assigned to. See *Assign APs to Sites*. You can also use *Automatically Assign Devices to Sites* to automatically take care of tedious steps such as this for you.
2. Configure the guest network. From the left menu, navigate to **Organization** > **Wireless** > **WLAN Templates**, then select the appropriate template or create a new one.
3. In the WLAN Template, navigate to the **WLANs** section, then click **Add WLAN**.

   NOTE: You must create an untagged LAN network to match that of your WAN Edge.

4. Add the same guest network you created in the previous steps. You can use the samples in Table 6 on page 58 below to guide you. For more information on how to configure WLAN templates, see *Configure a WLAN Template*, *Adding a WLAN*, and *WLAN Options*.

   Table 8: Sample Guest Network Configuration in WLAN Templates

Field	How to Configure
SSID	Guest
Security Type	Choose a security type that best fits your deployment.
VLAN	Tagged
VLAN ID	100
Guest Portal	If guest users on your network will need to sign in to get internet access, make the appropriate selection. For example, you can *Add a Custom Guest Portal to a WLAN*, *Use an External Portal for Guest Access*, or *Use an Identity Provider for Guest Access*.

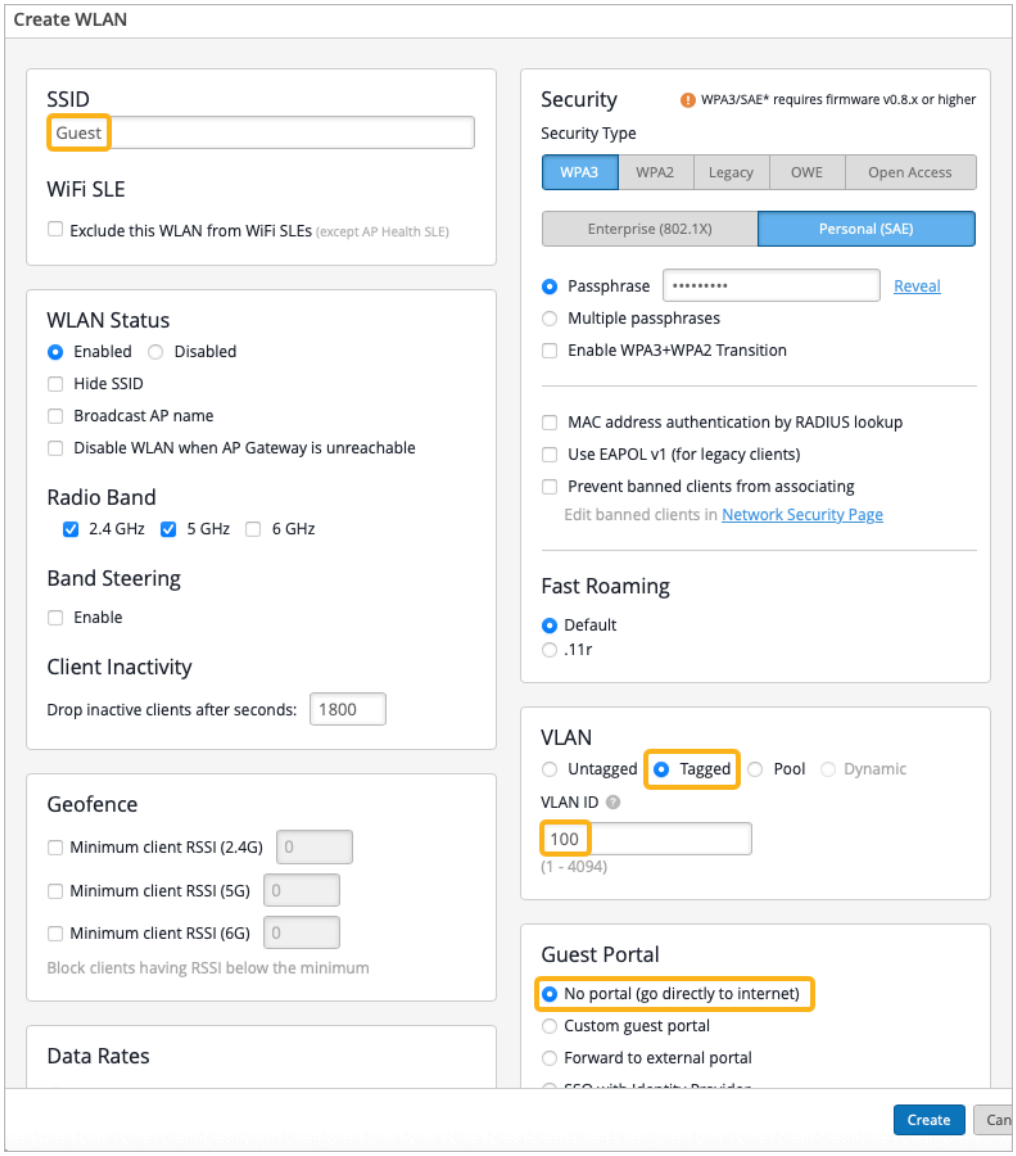

5. Click **Create** at the bottom of the Create WLAN window.
6. **Save** the WLAN template.
   You have now completed the Full Stack design in which your WAN Edges, Switches, and APs are interconnected to provide a cohesive dashboard.

SEE ALSO

WAN Assurance Configuration Overview | 75

Use Site Variables to Streamline Configuration | 77

Define Applications and Networks | 80

Application Settings | 98

High Availability Design for WAN Edge Devices

SUMMARY

Use this design guide to configure high availability for your WAN Edge devices in Juniper Mist WAN Assurance.

IN THIS SECTION

The High Availability (HA) design for WAN Edge devices is for administrators who want to ensure that interfaces or whole devices can take over for one another in the event of a failure in their WAN Edge

deployments. This is for administrators who want to deploy HA WAN Edge devices at the Edge, but not for Whitebox setups.

In this documentation, you'll find step-by-step guidance for setting up an HA hub and spoke deployment using Juniper® Mist WAN Assurance. This builds upon the hub and spoke topology referenced in this guide. After you've configured that topology, use the steps below to set up WAN Edge devices for HA.

If you need to set up your hub and spoke topology, see "Configure a WAN Edge Template" on page 87 and "Configure Path Selection from Hub-to-Spoke with Traffic Steering" on page 139, then return to this procedure.

NOTE: The device hardware used in an HA pair must be identical. For example, an HA pair with two SSR120 Routers is compatible for HA, whereas an HA pair with one SSR120 and one SSR130 is incompatible.

Overview

You will deploy a highly available Hub and Spoke as shown in Figure 26 on page 61. In the figure below, we see the Juniper Mist WAN Assurance High Availability topology. Nodes of an HA pair can be configured to run as active/standy or active/active. You can also set up your hubs for HA if needed.

Figure 26: Juniper Mist WAN Assurance High Availability Topology

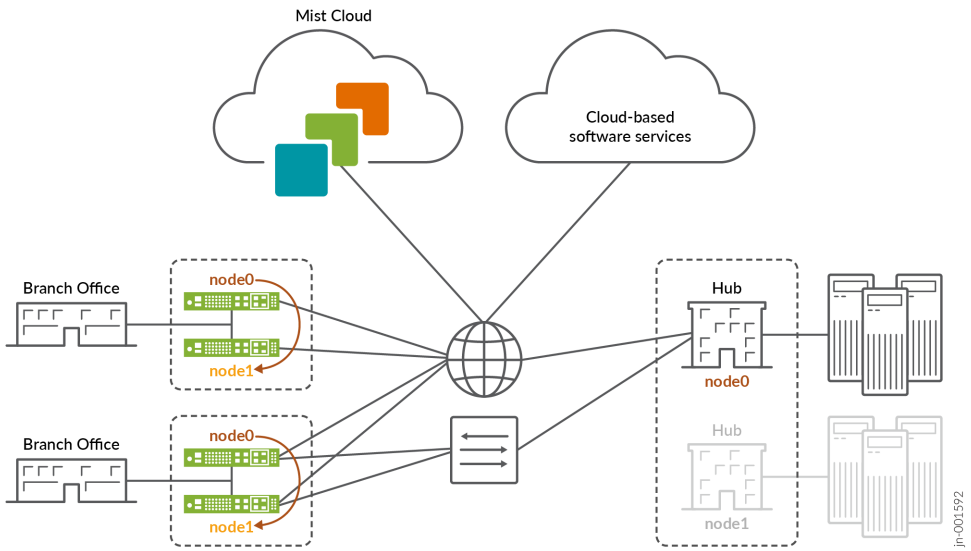

Interfaces

The Interfaces use the following pattern for each node:

Node0: ge-0/0/x

Node1: ge-1/0/x

For SRX devices, follow the naming pattern found in SRX Cluster Slot Numbering and Logical Interface Naming.

For SSR devices, follow the naming pattern found in SSR Device Default Port Layout.

WAN Interfaces for HA hubs require static IP addresses. Spokes reach out across the overlay to these WAN interface endpoints.

WAN Interfaces for HA

Each path and Node in an HA network require their own designated WAN interface. This ensures active/active usage, meaning that these interfaces stay active and engaged, no matter what. WAN interfaces on spoke devices can contain either a static IP address or be linked to a DHCP-lease, giving you flexibility in how you manage them.

LAN Interfaces for HA

You'll need to define the LAN interfaces for both HA hubs and spokes as redundant interfaces, and then specify the interfaces together as **ge-0/0/x, ge-1/0/x**.

Redundant Interfaces are only Active/Passive, meaning the active WAN Edge device will ARP for the IP address configured on the interface.

NOTE: The redundant interfaces must be in the same Layer 2 domain and need a single static IP address for Session Smart Routers (SSRs). These interfaces will have a shared MAC address. Based on the device, the system decides who will be node0 and who will be node1. SRX Series Firewalls do not need the same layer 2 domain for redundant interfaces.

- The lowest MAC address will be selected for node0.
- For redundant interfaces, you can define which **node** is the primary, but we recommend leaving the default to node0 for consistency.

Prerequisites for High Availability

This procedure assumes that you have already:

- Onboarded your WAN Edge device to the Mist cloud. If you need to onboard a WAN Edge device, follow the steps in Cloud Ready SSR Devices Quick Start Guide or Cloud Ready SRX Series Firewalls Quick Start Guide.
- Created a site, which you will assign your WAN Edge template to later in this HA workflow. If you need to create a new site, see the Juniper Mist Management Guide.
- Configured **Networks**, **Applications**, **Variables**, **Hub Profiles** and **WAN Edge Templates**.
  - If any of these steps are new to you, follow the applicable links in "WAN Assurance Configuration Overview" on page 75 before returning to this procedure.

> **NOTE**: Both devices in an HA pair must be on the same firmware version.

ATTENTION: Follow the sections and steps below as your step-by-step procedure for completing an HA setup.

Connect HA Synchronization and Fabric Links

It's important to be aware of the two specific Ethernet interfaces that handle HA synchronization and fabric data exchange on the supported devices.

The HA synchronization link ensures that the two devices are chronologically synchronized and can swap appropriately in the event of an interface or device failure. The synchronization interface serves as the back-or-midplane of a chassis-based router.

The fabric interface is a forwarding interface between two nodes in a router and is used for forwarding data when the ingress interface and egress interface for a given session are active on different nodes.

The synchronization and fabric interfaces are usually the two last ports of the system. You must wire them back-to-back with direct patch cables, as these are physical connections between the two nodes that are collocated in the same datacenter. See Figure 2 on page 64 below.

NOTE: To have a functional HA cluster, you must connect dedicated ports for the ha sync and fabric interfaces. To understand which ports to use for these, see SRX Cluster Slot Numbering and Logical Interface Naming and SSR Device Default Port Layout.

Figure 27: Redundant Nodes in a High Availability Cluster

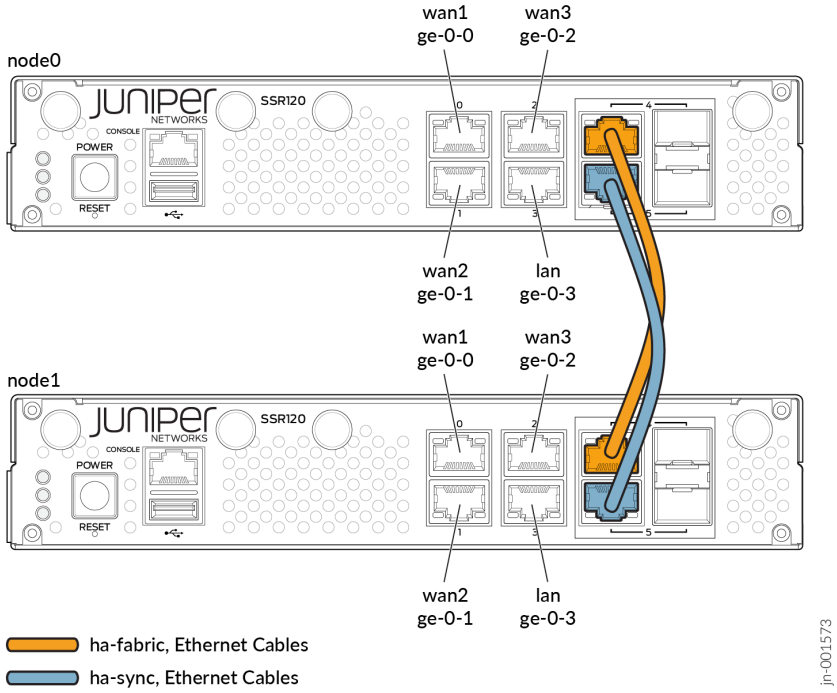

Configure Nodes for Redundancy

In an HA design, the nodes of an HA pair must be redundant so that they can successfully take over for one another in the event of a failover. This means that the interfaces must be configured identically on each of the nodes in the pair.

To do this, you must either edit an existing WAN Edge template and update it with the appropriate node interface configuration, or you can "Create a new template" on page 87 with the necessary node interface configuration. This procedure demonstrates how to configure a template for HA.

NOTE: The steps below show how to configure nodes for redundancy in a WAN Edge template, but the same steps can be applied to your hub profiles if you have any hubs in your high availability design.

1. Navigate to the WAN Edge template.
2. Scroll down to the WAN section and click on the WAN interface that you want to configure for redundancy.
3. In the **Interface** field, enter the names of the redundant nodes. You can enter the names with a comma separating them (Example: ge-0/0/7, ge-1/0/7).
4. Select the **Redundant** checkbox. If you are configuring HA for SSRs, skip ahead to step 5.

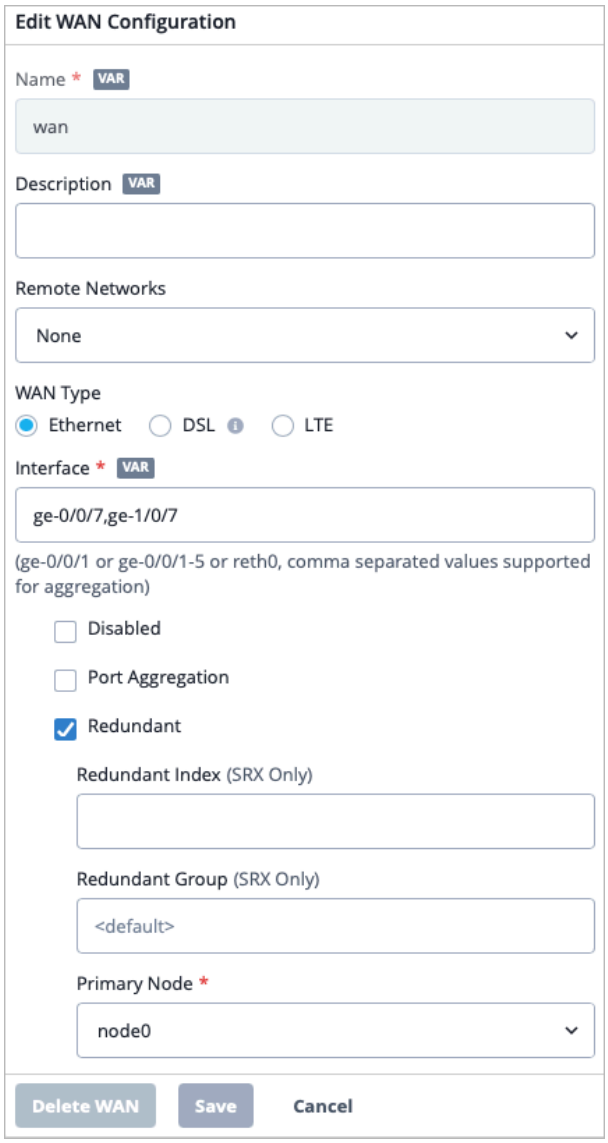

a. Enter a value in the **Redundant Index (SRX Only)** field.

b. Enter the same value in the **Redundant Group (SRX Only)** field. For more information, see Chassis Cluster Redundancy Groups and Chassis Cluster Redundancy Group Failover.

> **NOTE**: You can have an interface failover as part of another redundancy group. Specify the redundancy group number in the Redundant Group field on that interface so that the interface fails over with the specified group.

5. From the **Primary Node** drop-down, select the node that this interface belongs to. The node you select here will be the "active" node in the HA setup.

> **NOTE**: The steps in this procedure focus on the active/passive HA design, but if you want interfaces to run active/active, you must configure one interface with a Primary Node of Node0, and the other interface with a Primary Node of Node1. Then, you must configure "traffic steering" on page 69 rules that send traffic out both interfaces.

6. Click **Save** on the WAN Configuration Window.
7. Continue editing and adding WAN and LAN interfaces for redundancy as needed in the template.
8. Click **Save** in the top right corner of the page to save the template.

Customize the Fabric Interface if needed

There are certain circumstances where you may need to customize the fabric interface, for example, if your fabric port needs to be set to higher capacity. If you do not need to customize the fabric link, skip ahead to the "Assign the Template to a Site" on page 67 section.

The fabric link can only be customized from the API, and you must be logged into the Mist portal in order to use the REST API Explorer.

1. Login to the Mist portal.
2. Click the **?** button in the top right corner, then click **API Documentation**.
3. From the Site section of the table of contents, click the HA Cluster link. Use the documentation to learn how to customize the fabric link of your HA cluster in the API. The documentation demonstrates using the SRX, but is applicable to any type of WAN Edge device.
4. Navigate to the API URL for your global region to customize the fabric interface according to the instructions. See *API Endpoints and Global Regions* .
5. If you need to learn more about how to use the API, see *Use the Django Web Interface to Make API Changes* , *Additional RESTful API Documentation*, and the Juniper Mist API Reference.

Assign the Template to a Site

Now you must assign the template to the same site that you will assign the two devices in your HA pair to.

Follow the steps in "Assign Sites" on page 91, then return to this procedure and navigate to the "Assign HA Devices to a Site" on page 67 section below.

Assign the HA Devices to a Site and Create an HA Cluster

WAN Edge devices can be configured to operate as an HA cluster. An HA cluster is where a pair of devices can be connected together and configured to operate as a single device to provide high availability. You may want to cluster an existing WAN Edge device with a brand new WAN Edge device for the purpose of high availability. With Mist, you can cluster WAN Edge devices automatically.

1. Navigate to **Organization** > **Site** > **Inventory**. You should see your recently onboarded devices listed there. Notice that they do not have a site assigned to them.
2. Select the two devices that you want to configure as an HA pair.
3. Click the **More** button in the top right corner of the page, then select **Assign To Site**. You will only be able to assign a site to two devices at the same time when those devices are the same hardware model.

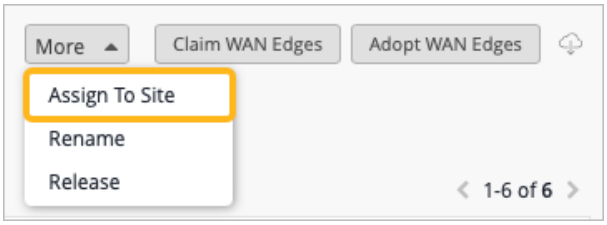

4. Select a site from the drop-down in the **Assign WAN Edges** pop-up. You should select the same site that your WAN Edge template for HA is assigned to.
5. Select the **Create Cluster** checkbox. Once that checkbox is selected, the **Manage configuration with Mist** option is automatically selected.

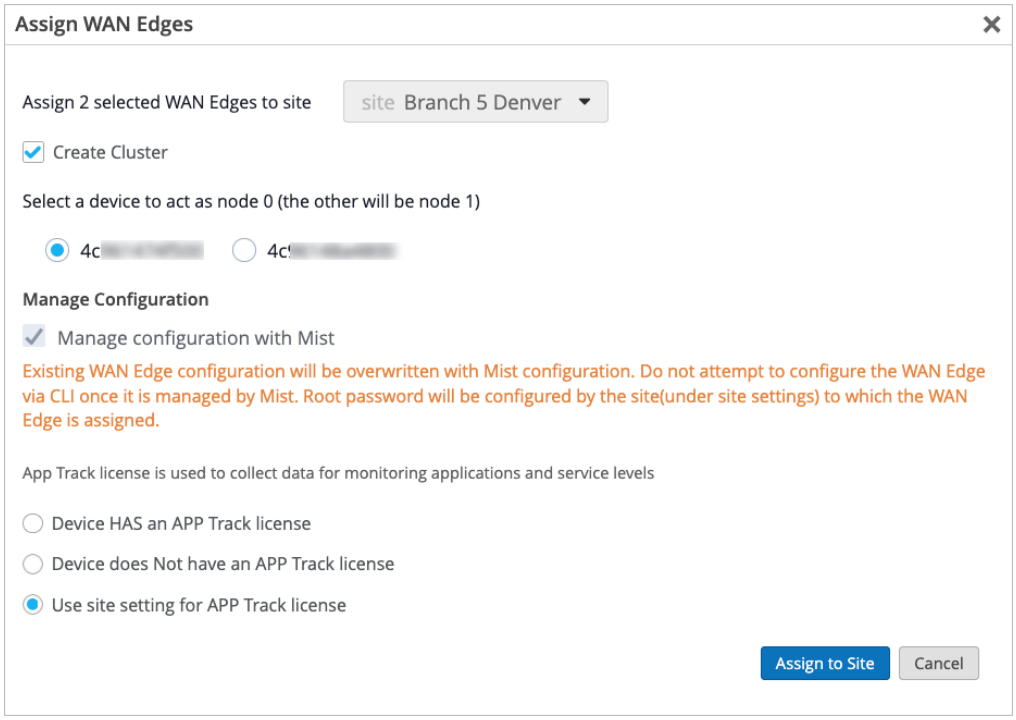

6. Select the MAC address for the device that you want to act as node0 in the cluster. The other will act as node1.
7. Click **Assign to Site** at the bottom of the pop-up.
   You should now see that the two devices have been assigned to the site.
   At this point, Mist reboots the devices from standalone mode to cluster mode.

   Mist runs the correct commands on each box individually to build them into a cluster. This process can take between three and fifteen minutes to complete depending on the platform. After this, the device is connected to the Mist cloud.
8. Verify that your devices have been clustered. A double graphic to the left of the device in the Inventory list indicates that the devices have successfully been clustered. You can also expand the row to see the MAC addresses for the two devices.

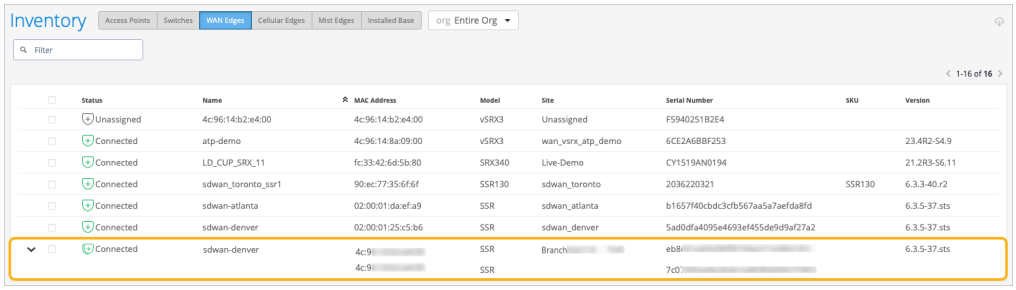

Configure Traffic Steering Rules

Traffic steering rules direct the flow of data traffic from one location or device to another. These rules help control how data packets are routed within a network, ensuring efficient and optimized data delivery.

To learn more about Traffic Steering, see "Traffic Steering Rules" on page 134 and "Configure Path Selection from Hub-to-Spoke with Traffic Steering" on page 139.

Navigate to the Traffic Steering section of the template and edit the existing traffic steering rules, or add new rules as needed. Refer to the links provided at the beginning of this section.

Table 1 on page 69 below provides a sample configuration of one traffic steering rule. One rule can contain multiple paths, allowing you to specify which path you want as the primary path for traffic to take when reaching an application, and a secondary path that traffic can take if the first path goes down. Once you create an "application policy" on page 70 and apply the traffic steering rule to it, traffic will follow that rule when attempting to access the application.

Table 9: Sample Traffic Steering Rule for HA

Name	Overlay
Strategy	ECMP
Paths	h1-wan0, h1-wan1

In Figure 3 on page 69 below, there are multiple traffic steering rules configured. The primary path traffic will take is the first path listed, and the second path listed is the secondary path that traffic can failover to if needed (from left to right).

Figure 28: Traffic Steering Rule Configuration

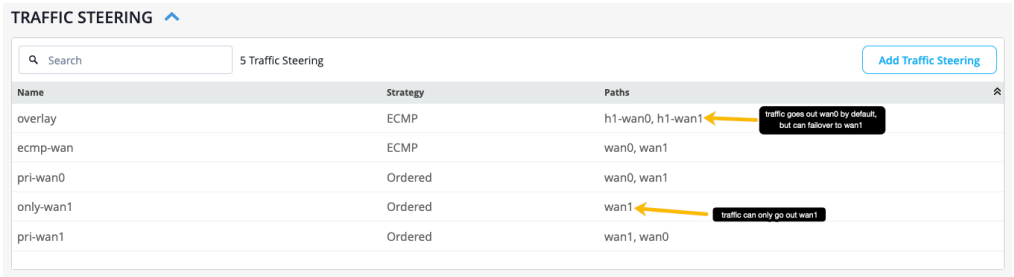

REMEMBER: If you want interfaces to run active/active, you must configure one interface with a Primary Node of Node0, and the other interface with a Primary Node of Node1. Then, you must configure two traffic steering rules, one rule to send traffic out one interface, and another rule to send traffic out the other interface.

Configure Application Policies

You will now modify the application policies in your WAN Edge template to define which networks and users can access which applications, and which traffic steering policies are used. For more information on how to create a new application policy, see "Application Policies" on page 144.

Navigate to the Application Policies section of the template. You can edit any existing policies as needed, or you can create a new one.

Table 2 on page 70 provides a sample configuration of one application policy.

Table 10: Sample Application Policy for HA Traffic

Field	How to Configure
Name	guest-internet
Network	spoke-guest
Action	Allow
Application/Destination	any
Traffic Steering	only-wan1

In Figure 4 on page 71, the application policies are configured with the traffic steering rules from Figure 3 on page 69. In the "public-dns" application policy, traffic will use wan0 as the primary path and wan1 as the backup path as defined by the "pri-wan0" traffic steering rule. The "guest-internet" application policy states that traffic can only take the wan1 path.

Figure 29: Application Policy Configuration

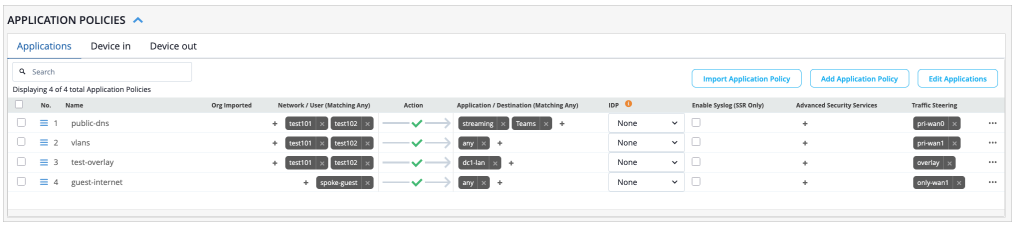

Configure a High Availability Cluster (Video Walkthrough)

This video walks you through how to create a highly available cluster. It captures the same steps that were presented to you earlier in this document, this time in video format.

ATTENTION: The steps in this video apply to *any* type of WAN Edge device being clustered for high availability.

Video: How to Configure SRX Clustering with Mist

SEE ALSO

WAN Assurance Configuration Overview | 75

Use Site Variables to Streamline Configuration | 77

Network Settings | 94

Application Settings | 98

Application Policies | 144

Configure a Hub Profile | 82

Configure a WAN Edge Template | 87

Assign Sites | 91

Assign Sites | 91

Replace a WAN Edge Node in a High Availability Cluster

SUMMARY

You can replace a WAN Edge device in a high availability cluster setup with a few simple steps.

You can replace a WAN Edge device within a high availability cluster setup with just a few simple steps.

Before you replace a WAN Edge node from the cluster, you must:

- Remove the cluster fabric cables from the node being replaced and connect it to the new replacement node.
- Make sure that the replacement WAN Edge is both the same model as the device being replaced.
- Make sure you upgrade the firmware version to match that of the existing device in the cluster.

NOTE: Replacing a node in a high availability setup causes minimal impact on network services. Therefore, we recommend that you plan to do this task during a maintenance window.

To replace a WAN Edge:

1. In the Juniper Mist portal, on the left menu, go to **Organization** > **Admin** >**Inventory** and select the **WAN Edges** tab.
   The page displays a list of WAN Edge devices. You can set the view as **org (Entire Org)** or **Site** in the Inventory page.
2. Click the high availability pair that you want to replace a node on to open the details in a new page.
3. Select **Utilities** > **Replace WAN Edge**.

Figure 30: Replace WAN Edge

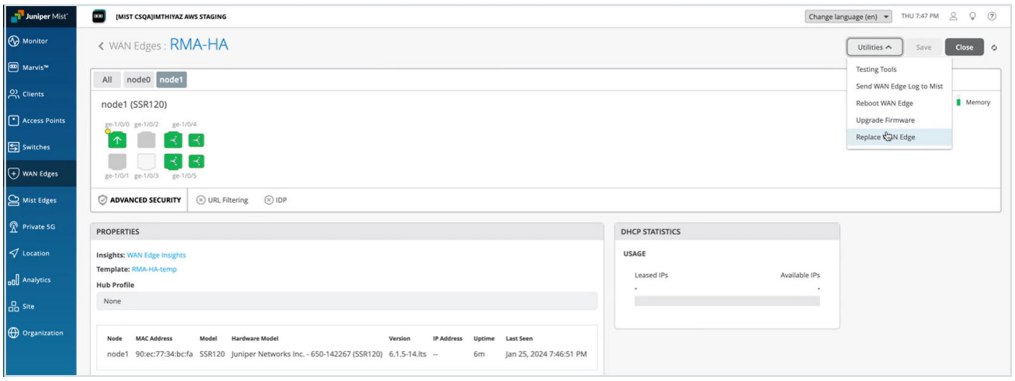

4. On the Replace WAN Edge window, select the old WAN Edge node that you want to replace and select the new replacement device's MAC address from the **MAC Address of available WAN Edge** drop-down list.

Figure 31: Replace WAN Edge with Another Device

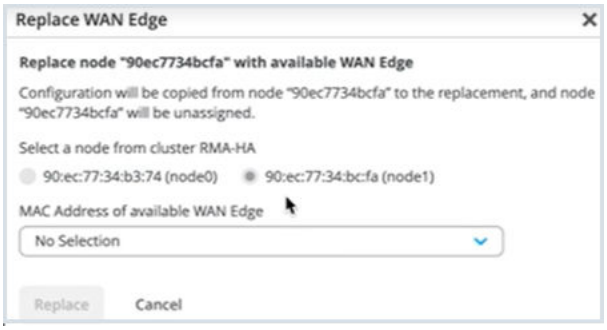

After you click **Replace**, allow about 15 minutes for the replacement procedure to complete.

Refresh your browser and check under WAN Edges to find out if your high availability setup is updated and available as a part of the inventory.

3

CHAPTER

WAN Edge Configuration

IN THIS CHAPTER

WAN Assurance Configuration Overview

SUMMARY

After you've completed the prerequisite tasks and designed your WAN, use this information to guide the configuration process.

IN THIS SECTION

Configuration Workflow

The configuration process consists of a set of inter-related tasks.

Figure 32: WAN Configuration Workflow

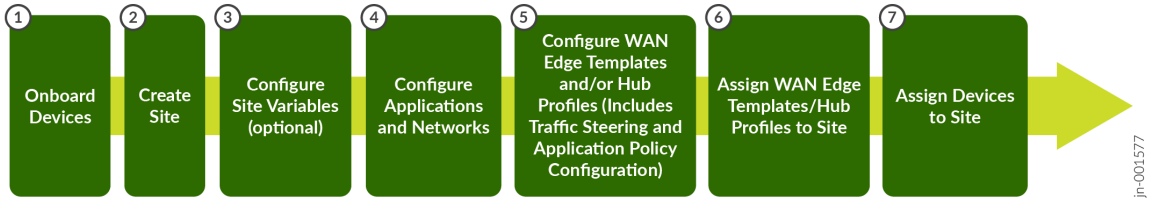

Table 11: WAN Configuration Workflow

Task	Description	More Information
Create variables to use in your configurations.	If you haven't already done so, add your sites to your organization. Then, to streamline the setup process, define site variables for each site. You use these variables later in the templates for WAN Edge devices and the hub profile.	"Define Site Variables" on page 77
Define your applications and network.		"Define Applications and Networks" on page 80

Table 11: WAN Configuration Workflow *(Continued)*

Task	Description	More Information
Create application policies.	Application policies determine which networks or users can access which applications, and according to which traffic steering policy.	"Application Policies" on page 144
Create WAN Edge templates.	Create WAN Edge templates to streamline and standardize device configuration across your sites.	"Configure WAN Edge Templates" on page 87
Create hub profiles.	Create hub profiles for standalone or clustered devices to automate overlay path creation.	"Configure Hub Profiles" on page 82
Configure any standalone devices.	If you are managing a device without a hub profile or WAN Edge template, going to **WAN Edge** > **WAN Edge**, select the device, and enter the settings.	Configuration Reference chapter of this guide
Assign devices, templates, and hub profiles to sites.	Onboard your devices by assigning them to a site. Complete the onboarding by attaching hub profiles and spoke templates to the respective hub sites and spoke sites. This final step brings the topology together.	"Assign Sites" on page 91

Optional Configuration Tasks

You can perform the following tasks on your devices for providing additional security measures:

- Set Up Secure Edge Connectors—Perform traffic inspection by Secure Edge for the WAN Edge devices managed by Juniper Mist Cloud portal. To get started, see "Juniper Mist Secure Edge Connector Overview" on page 222.
- Configure IDP-Based Threat Detection—Monitor the events occurring on your network and proactively stop attacks and prevent future attacks. See "Intrusion Detection and Prevention" on page 158.

Upgrade software on your device to take advantage of new enhancements.

- Upgrade Software—Upgrade the software on your device through the Juniper Mist portal in a few simple steps. See "Upgrade WAN Edge Device" on page 214.

Use Site Variables to Streamline Configuration

SUMMARY

Follow these steps to add site variables that you can then use to streamline the configuration process for your WAN deployment.

Site variables help simplify your configuration process and provide flexibility for large-scale deployments.

As you set up templates, profiles, and other configurations, you can insert a variable, such as {{SUBNET_1}} to represent a value such as a subnet address. You can use variables in the API and in the Juniper Mist portal. On various configuration screens, look for the **VAR** label to identify fields that accept variables.

You store the definitions for your variables in your site configuration. When you associate a template or profile with a site, Juniper Mist plugs in the site-specific values and pushes them to the devices.

NOTE: With the API, you can easily see the organization-wide list of variables by using **GET /api/v1/orgs/:org_id/vars/search?var=***.

To define site variables:

1. From the left menu, click **Organization** > **Admin** > **Site Configuration**.
2. Scroll down to the **Site Variables** section.
3. Click **Add Variable**.
4. In the pop-up window, enter a name surrounded by double curly brackets, specify the value, and then click **Save**.

   For help with the variable format, see the samples in the table below these steps.

5. Continue until you've added all the variables that you need for this site.

   This table shows sample variables for five sites. Three of the sites are the locations for spokes (spoke1-site, spoke2-site, and spoke3-site.) Two are the locations for hubs (hub1-site and hub2-site). As you explore other topics in this chapter, you'll see additional examples involving these same sites and variables. You'll see how the variables are used to configure networks, applications, hubs, and hub-to-spoke traffic steering.

   Table 12: Variables for Five-Site Configuration Example

Site Name	Variable	Value
spoke1-site	{{SPOKE_LAN1_PFX}}	10.99.99
	{{SPOKE_LAN1_VLAN}}	1099
	{{WAN0_PFX}}	192.168.173
	{{WAN1_PFX}}	192.168.170
spoke2-site	{{SPOKE_LAN1_PFX}}	10.88.88
	{{SPOKE_LAN1_VLAN}}	1088
	{{WAN0_PFX}}	192.168.133
	{{WAN1_PFX}}	192.168.130
spoke3-site	{{SPOKE_LAN1_PFX}}	10.77.77
	{{SPOKE_LAN1_VLAN}}	1077
	{{WAN0_PFX}}	192.168.153
	{{WAN1_PFX}}	192.168.150
hub1-site	{{HUB1_LAN1_PFX}}	10.66.66
	{{HUB1_LAN1_VLAN}}	1066
	{{WAN0_PFX}}	192.168.191

Table 12: Variables for Five-Site Configuration Example *(Continued)*

Site Name	Variable	Value
	{{WAN1_PFX}}	192.168.190
	{{WAN0_PUBIP}}	192.168.129.191
	{{WAN1_PUBIP}}	192.168.190.254
hub2-site	{{HUB2_LAN1_PFX}}	10.55.55
	{{HUB2_LAN1_VLAN}}	1055
	{{WAN0_PFX}}	192.168.201
	{{WAN1_PFX}}	192.168.200
	{{WAN0_PUBIP}}	192.168.129.201
	{{WAN1_PUBIP}}	192.168.200.254

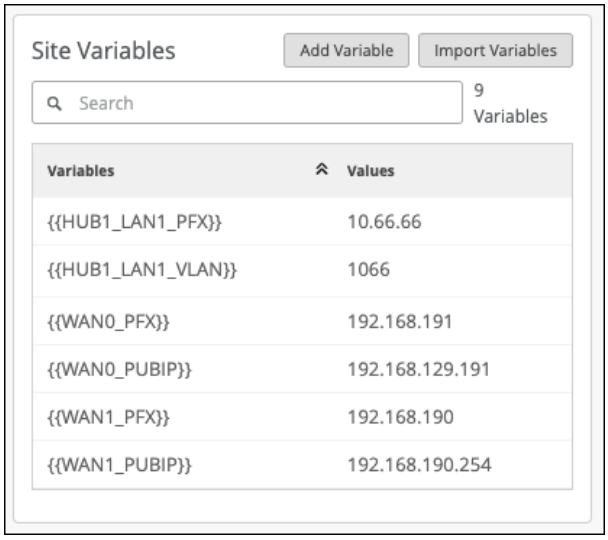

6. Click **Save** at the top-right corner of the Site Configuration page.

You can use variables in various fields when configuring your WAN. When you make site assignments for the profiles and templates that use these variables, then the specified values will be sent down to the devices.

Define Applications and Networks

SUMMARY

Before you configure templates, hub profiles, or standalone devices, define your networks and applications.

IN THIS SECTION

Define Applications

You can think of *applications* in Juniper Mist as *what destinations* your network traffic goes to. Examples include network services, SaaS, private subnets, and cloud workloads.

Defining applications is a prerequisite for creating application policies.

You can define applications by entering your own customized settings, selecting brand-name applications, or assigning URL categories.

To configure applications:

1. From the left menu, click **Organization** > **WAN**> **Applications**.
2. Click **Add Applications** in the top-right corner of the **Applications** page.
3. Enter a name and description.
   - **Name**—Enter a unique name for the application. You can use up to 32 characters for naming the application including alphanumerics, underscores, and dashes.
   - **Description**—Enter a description of the application and context.
4. Select the application type, and enter the relevant settings.

   You can specify applications in various ways. For help with these options, see "Application Settings" on page 98.
5. Click **Add** at the bottom of the settings panel.

Define Networks

SUMMARY

Follow these steps to define the networks for your WAN Edge devices.

You can think of *networks* in Juniper Mist as segments of your user and device population. Networks are attached to interfaces and are used as sources in application policies. A network also is a gateway for traffic.

Creating networks is a prerequisite for other WAN configuration tasks, so you'll do these tasks early in the configuration workflow.

To configure networks:

1. From the left menu, click **Organization** > **WAN** > **Networks**.
   A list of existing networks appears.
2. Click **Add Networks** in the top-right corner of the Networks page.
3. Enter the settings in the **Add Network** side panel.
   For help with the options, see "Network Settings" on page 94.
4. At the bottom of the side panel, click **Add**.
   The screen displays the newly created networks.

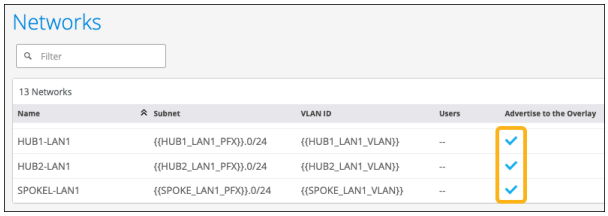

Configure a Hub Profile

SUMMARY

Follow these steps to set up a hub profile to configure each hub device.

IN THIS SECTION

Each hub device in a Juniper Mist™ cloud topology must have its own profile. Hub profiles are a convenient way to create an overlay and assign a path for each WAN link on that overlay in Juniper WAN Assurance.

NOTE: You'll create hub profiles for WAN Edge devices at hub sites. You'll create WAN Edge templates for WAN Edge devices at spoke sites. Hub WAN interfaces create overlay endpoints for spokes. Spoke WAN interfaces map the appropriate Hub WAN interfaces, defining the topology. Hub profiles drive the addition and removal of paths on your overlay.

When you create a hub profile for a WAN Edge device, the Mist cloud generates and installs the SSL certificates automatically. It also sets up WAN uplink probes for failover detection.

As a time saver, you can create a hub profile by cloning an existing one.

Before You Begin

- "Prepare to Configure Your WAN (Major Prerequisites)" on page 28
- "Define Applications and Networks" on page 80

Create a Hub Profile (Without Cloning)

Use this procedure to get started setting up your hub profile. From this high-level overview, you can jump to more detailed information later in this topic.

To get started with your hub profile:

1. From the left menu, select **Organization** > **WAN** > **Hub Profiles**.
   A list of existing profiles, if any, appears.
2. Click **Create Profile** in the top-right corner of the page.

   > **NOTE**: You can also create a hub profile by importing a JavaScript Object Notation (JSON file) using the **Import Profile** option.

3. Enter a descriptive name to identify this profile, and then click **Create**.
   The name must be 2-32 characters long. It can contain only letters, numbers, and underscores.
4. Work your way through the remaining sections of the hub profile, saving your work as you go.
   Tips:

   - When a configuration panel appears at the side of the screen, enter the settings, and then click the **Add** or **Save** button at the bottom of the panel.
   - If you click a button in a panel and a new sub-section opens, enter the settings, and then click the check mark in the title bar of the sub-section.

     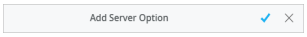

   - In fields with a VAR label, you can enter variables to represent values such as IP addresses, VLAN IDs, and more. As you type, the field displays any matching variables as defined in site configurations across the organization.

     

     For more information, see "Use Site Variables to Streamline Configuration" on page 77.
   - After completing each major section of a template, click **Save** at the top-right corner of the template page. When the page reloads, showing the list of templates, simply click the template to resume your work.
   - For help with the various sections of the hub profile, see the Configuration Reference chapter of this guide.
5. Click **Save** at the top-right corner of the Hub Profile page.

Create a Hub Profile by Cloning

As a time saver, you can create a hub profile by cloning an existing one. Then modify the cloned profile so that it has the required settings.

To create a hub profile by cloning:

1. In the Juniper Mist cloud portal, click **Organization** > **WAN** > **Hub Profiles**.
2. Click the hub profile that you want to clone.
3. In the top-right corner of the profile page, click **More** and then click **Clone**.

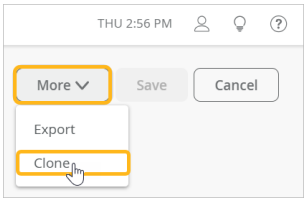

4. Enter a name for the new hub profile, and then click **Clone**.

> **TIP**: If you see any errors while naming the profile, refresh your browser.

5. Edit the cloned settings as needed for this hub profile.

Create a Hub-to-Hub Overlay

Use this feature to create a peer overlay path between two hub devices. You can then set up traffic steering rules for the use of this path. You might use a hub-to-hub path as a preferred route for data center traffic originating from sites. Or you might use it as a failover path for a hub-to-spoke connection.

When you create a hub-to-hub overlay, you add corresponding hub endpoints to the hub profiles. For example:

- On Hub A, add an endpoint for one of Hub B's WAN interfaces.
- On Hub B, add an endpoint for one of Hub A's WAN interfaces.

> **NOTE**: The selected interfaces can be two matching WAN interfaces that exist in both hub profiles or two different WAN interfaces that are unique to each device.

Before You Begin: Configure the hub profiles for both hubs (Hub A and Hub B).

To create a hub-to-hub overlay:

1. Go to the hub profile for one hub (Hub A) and add an endpoint for the other hub (Hub B):
   a. Click the WAN interface that you want to use for this hub-to-hub overlay.
   b. At the bottom of the **Edit WAN Configuration** window, click **Add Hub-to-Hub Endpoints**.

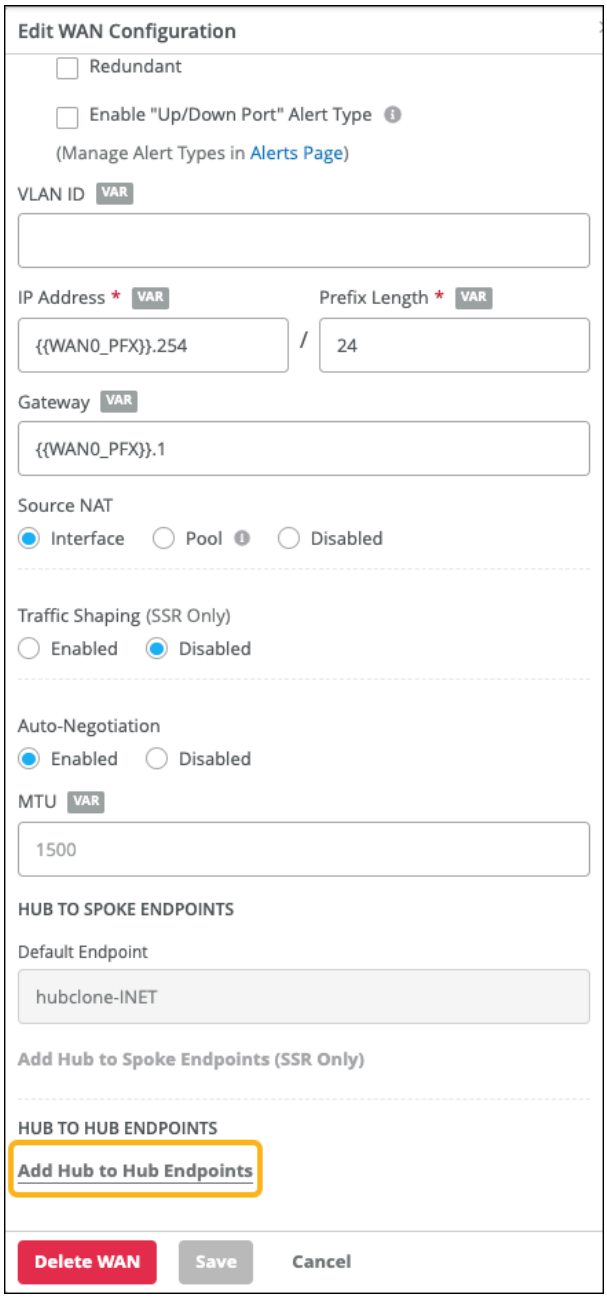

   c. From the drop-down menu, select a WAN interface on Hub B that you want to use for this hub-to-hub overlay.

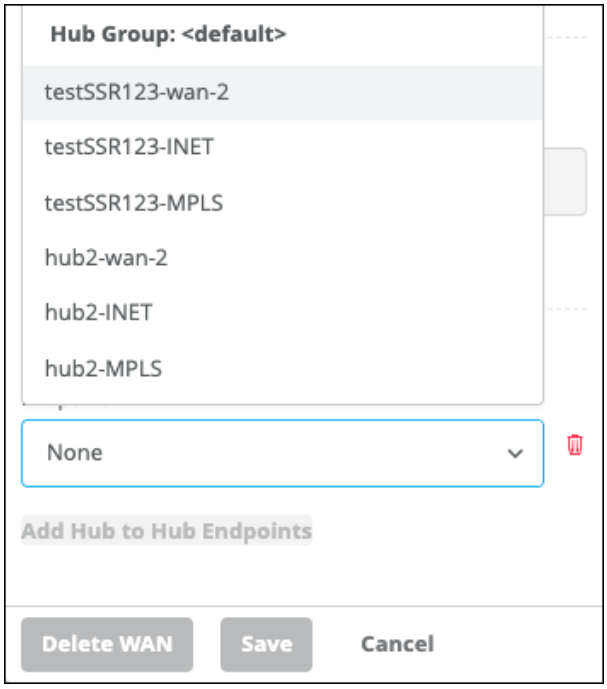

If needed, you can click **Add Hub-to-Hub Endpoints** again to add more endpoints.

d. Click **Save** at the bottom of the panel to save the WAN configuration.

In the WAN section of the hub profile, you'll see the new endpoints in the Hub to Hub Endpoints column.

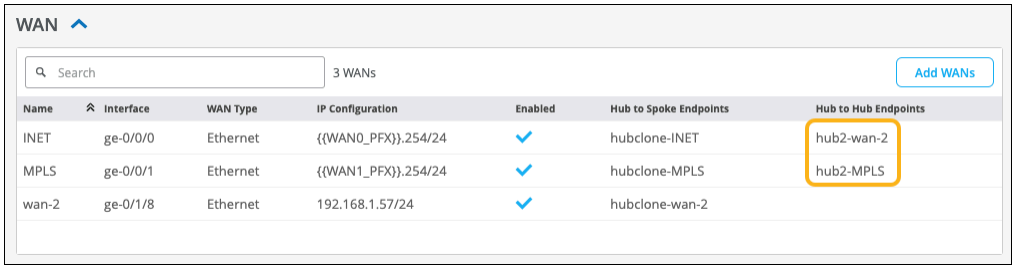

e. Click **Save** at the top-right corner of the hub profile page.

2. Go to the hub profile for other hub (Hub B) and add an endpoint for the first hub (Hub A):

a. Click the WAN interface that you want to use for this hub-to-hub overlay.

b. At the bottom of the **Edit WAN Configuration** window, click **Add Hub-to-Hub Endpoints**.

c. From the drop-down menu, select a WAN interface on Hub A that you want to use for this hub-to-hub overlay.

If needed, you can click **Add Hub-to-Hub Endpoints** again to add more endpoints.

d. Click **Save** at the bottom of the panel to save the WAN configuration.

In the WAN section of the hub profile, you'll see the new endpoint in the Hub to Hub Endpoints column.

e. Click **Save** at the top-right corner of the hub profile page.

You now have an overlay between the two hubs (Hub A and Hub B).

Configure a WAN Edge Template

SUMMARY

Streamline your configuration process by creating templates for your WAN Edge devices. For even quicker configuration, use pre-defined, device-specific templates.

IN THIS SECTION

By using WAN Edge templates, you can define common spoke characteristics such as WAN interfaces, traffic-steering rules, and access policies. When you assign a template to a site, Juniper Mist sends the configuration to the devices. When changes are needed, simply change the template, and all devices receive the new configurations. With this process you can easily configure and manage a large number of devices while ensuring consistency across your network infrastructure.

NOTE: For hub-and-spoke configurations, you use WAN Edge templates to configure WAN Edge devices at spoke sites. You use hub profiles to configure WAN Edge devices at hub sites.

Before You Begin

- "Prepare to Configure Your WAN (Major Prerequisites)" on page 28
- "Define Applications and Networks" on page 80

Create a WAN Edge Template By Defining Your Own Settings

Create a template to define common spoke characteristics that you want to deploy across your organization. Associate templates with sites, and then onboard devices to sites. When onboarded, each device automatically gets the configuration from the associated template.

NOTE: Follow this procedure if you want to define your own settings. If you'd prefer to use pre-defined, model-specific templates, see "Create a Template with Pre-Defined Device-Specific Settings" on page 89 later in this topic.

To configure a WAN Edge template:

1. From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**.
2. Click **Create Template** in the top-right corner of the page.

   NOTE: You can also create a WAN Edge template by importing a JavaScript Object Notation (JSON) file using the **Import Profile** option.

3. In the New Template window, enter a name for the template (up to 64 characters), click **Spoke** for the type, and then click **Create**.

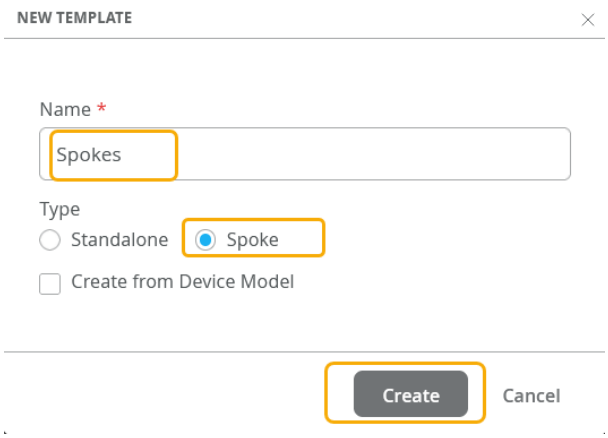

4. At the top of the page, select the sites to be configured with this template, and enter other general settings.
   - You can apply the same template to multiple sites. If a site already had a template assigned to it, the newly assigned template will now apply instead.
   - For help with other sections at the top of the configuration page, see "General WAN Edge Settings" on page 110.
5. Work your way through the remaining sections of the template, saving your work as you go.

   Tips:

- When a configuration panel appears at the side of the screen, enter the settings, and then click the **Add** or **Save** button at the bottom of the panel.
- If you click a button in a panel and a new sub-section opens, enter the settings, and then click the check mark in the title bar of the sub-section.

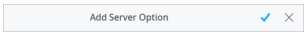

- In fields with a VAR label, you can enter variables to represent values such as IP addresses, VLAN IDs, and more. As you type, the field displays any matching variables as defined in site configurations across the organization.

  For more information, see "Use Site Variables to Streamline Configuration" on page 77.

- After completing each major section of a template, click **Save** at the top-right corner of the template page. When the page reloads, showing the list of templates, simply click the template to resume your work.
- For help with the various sections of the template, see the Configuration Reference chapter of this guide.

Create a Template with Pre-Defined Device-Specific Settings

Create a template with predefined, device-specific settings to apply standard configurations for Juniper devices.

These templates are unique for each device model. After you name your template and select the device model, no additional input is needed.

The WAN Edge device specific templates provide basic network configuration in a single step. Templates also allow for re-usable and consistent configuration across your WAN Edge devices. The template provides device-specific, preconfigured WAN interfaces, LAN interfaces, a traffic steering policy, and an application policy. All you have to do is name the template and select the device type.

For example, SSR120 WAN Edge template generates several values, including Ethernet interfaces for **LAN** and **WAN** with relevant **DHCP** and **IP** values:

- wan ge-0/0/0
- wan2 ge-0/0/1

- wan3 ge-0/0/2
- lan ge-0/0/3

To select a device-specific WAN Edge template:

1. From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**.
2. Click **Create Template** in the top-right corner to open a new template page.
3. Enter a descriptive name to identify template.
4. Select the **Create from Device Model** check box.
5. Select your device model.

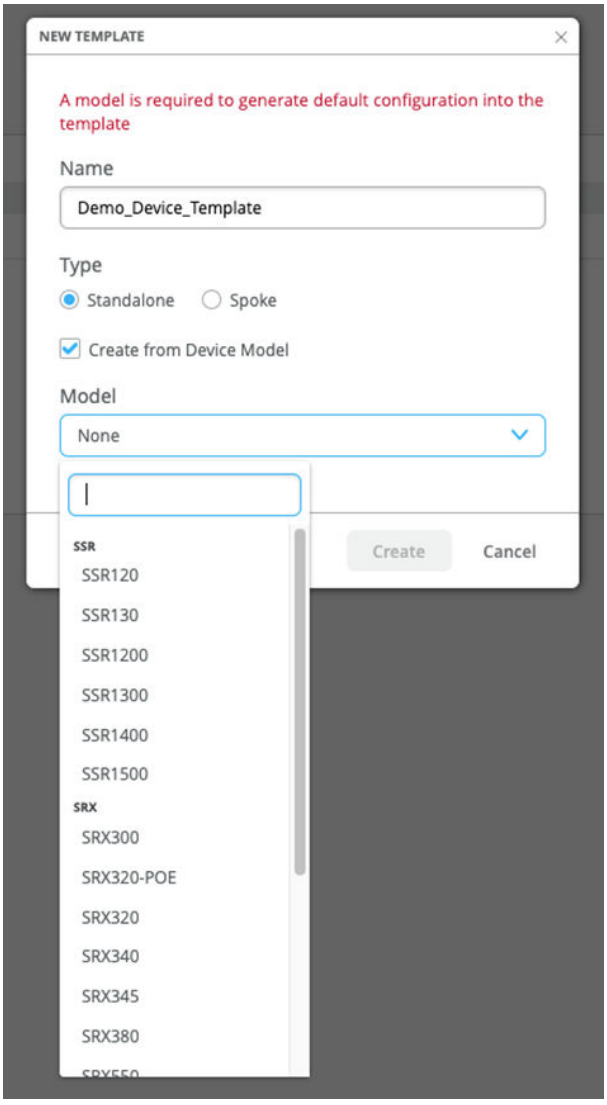

6. Click **Create** at the bottom of the configuration panel.
   The template is now listed on the WAN Edge Templates page. You now have a working WAN Edge template that you can apply to many sites and devices across your organization.
7. On the WAN Edge Templates page, select the check box for the template.

8. Click **Assign to Site** at the top of the template page.
9. Select the site where you want to apply this template.
10. Click **Apply**.

Onboard the device and assign it to the site.

Assign Sites

SUMMARY

Complete the configuration process by assigning templates, hub profiles, and devices to sites.

IN THIS SECTION

In a Juniper Mist™ network, you must onboard your WAN Edge devices by assigning them to sites. Assign your hub profiles, WAN Edge templates, and standalone devices to sites. This final step brings the topology together.

Before You Begin

- Create your WAN Edge template, hub profile, or standalone device configuration. For help, see "WAN Assurance Configuration Overview" on page 75.
- Install your device and claim or adopt it into your Juniper Mist organization. For more onboarding help, see SSR Series Devices and Cloud-Ready SRX Firewalls .

Assign Sites

1. Navigate to your WAN Edge template, hub profile, or standalone WAN Edge device.
2. Make the site assignment:

- In a WAN Edge template—Click **Assign to Sites**. Then click the plus button (+), select the site(s) to add, then click Apply. Save the template.

  > **NOTE**: A site can't have two WAN Edge templates. If a site already has a template assigned to it, the newly assigned template will replace the existing one.

- In a hub profile—Under **Applies to Devices**, click the drop-down list, and then click a site. Save the profile.
- For a standalone device—From the left menu, select **Organization** > **Admin** > **Inventory**. Select the check box next to the device, click the **More** menu, and then click **Assign to Site**. In the pop-up window, select a site, and then click **Assign to Site** at the bottom of the window.

Device Configuration and Automatic Rollback

After a device is adopted or claimed, it automatically receives configurations from the Juniper Mist portal.

If the configuration is valid, the device waits five minutes and then checks its connection to the Mist cloud. If the connection is good, the configuration proceeds.

If there are issues:

- Invalid Configuration—If a new configuration is invalid, the device rejects it and returns to the last valid configuration.
- Poor Connectivity—If the device loses connectivity after the five-minute wait period, it rolls back to the previous configuration. This auto-rollback feature ensures that the device quickly returns to a working configuration.

4

CHAPTER

Configuration Reference

IN THIS CHAPTER

Network Settings

SUMMARY

Explore the options that you can configure when creating a network.

IN THIS SECTION

Finding the Network Page

From the left menu, click **Organization** > **WAN** > **Networks**.

NOTE: For step-by-step instructions, see "WAN Assurance Configuration Overview" on page 75.

Configuration Options

Table 13: Network Options

Fields	Description
Name	Enter a unique name for the network. The name can contain alphanumeric characters, underscores, and hyphens, and must be less than 32 characters long.
Subnet IP Address	Enter the network IP address. You can use absolute values (example: 192.0.2.0) or variables. Example: {{SPOKE_LAN1_PFX}}.0. For more information, see "Use Site Variables to Streamline Configuration" on page 77.
Prefix Length	Enter the length of the address prefix, from 0 through 32. You can also use variables for the prefix length. Example: {{PFX1}}

Table 13: Network Options *(Continued)*

Fields	Description
VLAN ID	Enter the VLAN ID that is associated with the network. If your device is using an untagged interface, you should use 1 as the VLAN ID instead of the variable.
Access to Mist Cloud	Check this option to allow other devices within the defined network to connect to the Mist cloud for management through the Session Smart Router. We recommend this option if you intend to have Juniper switches or Juniper APs on this network. WAN Edge devices have built-in policies for connecting to Mist cloud services using WAN interfaces. By enabling this option, you permit the network to use these built-in policies. Use case example: Management LAN networks for Mist-managed devices.

Table 13: Network Options *(Continued)*

Fields	Description
Advertise to the Overlay	Select this option to advertise the network to the hub devices through the overlay tunnels. By using this feature, you gain the ability to advertise over the SD-WAN, enabling enhanced control and reachability. You can enable or disable these additional options: • **Advertise to Other Spokes**—This option leverages BGP community strings to allow the route to be sent to other spokes via the hub as a BGP route reflector. The WAN Edge device will originate a static route out to the attached networks. You might not need this option in these situations: • If you already have an aggregate route from the hub, you likely do not need to make this route visible to all spokes. • If you want the network to advertise the prefix only to hubs (not to other spokes), disable this option. • **Advertise to Hub LAN BGP Neighbor**—This option leverages BGP community strings to allow this route to be advertised into BGP at the edge of the SSR fabric (hub LAN, spoke LAN, or spoke SSE BGP). The WAN Edge device will originate a static route out to the attached networks. • **Advertise to Hub LAN OSPF Neighbor (SRX Only)**—This option leverages OSPF to advertise the network prefix to any LAN OSPF neighbor at the hub. • **Override Prefix to Advertise**—Enable this option when the prefix to advertise to the hubs is different from the original network. For example, you might use this option when enabling NAT. If you select this option, also enter the IP Address and Prefix Length.

Table 13: Network Options *(Continued)*

Fields	Description
Summarize	Select from the following route summarization options: • **Hub Overlay Summarization**—Enable the network to summarize the network prefix advertised to the overlay. For example: Juniper Mist portal can summarize 192.168.1.0/24 to 192.168.0.0/16. This feature limits the number of BGP updates received by a hub from each spoke and sent by the hub back to all the other spokes. • **Hub LAN BGP Summarization**—Enable the network to summarize the network prefix advertised to the LAN BGP neighbor. For example: Juniper Mist portal can summarize 192.168.1.0/24 to 192.168.0.0/16. • **Hub LAN OSPF Summarization (SRX Only)**—Enable the network to summarize the network prefix advertised to the LAN OSPF neighbor. For example: Juniper Mist portal can summarize 192.168.1.0/24 to 192.168.0.0/16. • **Route Summarization**—Summarize local routes toward overlay. You can specify the IP addresses and prefix length of the summarized routes. Session Smart Routers support summarization when the network is attached to the spoke only.
Networks not directly attached *(SSR Only)*	Select the networks that are not directly connected networks that arrive on this network assigned to a LAN. For example, enable this option for indirectly connected prefixes behind a router on a LAN segment. You'd also use this option in cases where there are geo-redundant hubs that could potentially provide access to the same core network space.

Table 13: Network Options *(Continued)*

Fields	Description
Users	Additional networks or users. Example: remote networks or users connected to the main network. Click **Add User** and then enter the **Name** and **IP Prefixes**.
Static NAT	Perform a one-to-one static mapping of the original private host source address to a public source address. Click the **Add Static NAT** option and enter the **Name** , **Internal IP**, **External IP** and select the option to apply to outgoing traffic on the **Underlay** or **Overlay**. For SRX, also enter the **WAN Name**.
Destination NAT	Translate the destination IP address of a packet. Click the **Add Destination NAT** option and enter the **Name** , **Internal IP**, **Internal Port**, **External IP**, **External Port** and select the option to apply to outgoing traffic on the **Underlay** or **Overlay**. **SRX Only:** also enter the **WAN Name**.

Application Settings

SUMMARY

Define the applications that you want to use in your application policies, which allow or deny access to network destinations.

IN THIS SECTION

Enable Application Visibility (SRX Only) | 104

Custom Applications

When you select the Custom Apps as the application type, you'll define traffic destinations by IP address range, domain name, and protocol. This approach is useful for internal services, specific IP ranges, and unique protocols used in your network.

Table 14: Custom Settings

Field	Description
IP Address	Enter the network IP address, including prefix (if any) of the application. You can enter multiple destination IP addresses or domain names separated by a comma.
Domain Names	Enter the domain name of the application, such as juniper.example.com. The domain name is used in cloud breakout profiles to generate the fully qualified domain name (FQDN). The cloud security providers use the FQDN to identify the IPsec tunnels.
Protocol, Port Number, and Port Ranges	Select a protocol and specify the port ranges (start and end ports) that the application is using. If you select Custom (SRX Only) as the protocol, also enter a protocol number from 1-254. **NOTE**: If you need to add more protocols for this application, click the + button at the top-right corner of the protocol section.

Consider using variables to represent values. If you see the VAR label next to a field heading, you can enter variables in that field to match destinations without having to enter specific values.

It's a good practice to create an application with a less specific address. For example, you might have a data center with a 10.0.0.0/8 prefix. A more specific address can be contained within this prefix for more specific path selection.

The following table provides examples.

Table 15: Custom Application Examples

Custom Application	IP Address	Description
ANY	0.0.0.0/0	Use this address to match all or any IPv4 address destinations. The IP address 0.0.0.0 also serves as a placeholder address.
SPOKE-LAN1	10.0.0.0/8	A match criterion for all IP addresses inside the corporate VPN.
HUB1-LAN1	10.66.66.0/24	A match criterion for all IP addresses attached at the LAN-interface of the Hub1 device.
HUB2-LAN1	10.55.55.0/24	A match criterion for all IP addresses attached at the LAN interface of the Hub2 device.

URL Categories

Applications represent the endpoints users are trying to reach—these can be IPs, domain names, or URLs. URL patterns help define these destinations more flexibly. The Juniper Mist cloud provides a list of URL categories based on types (example: shopping, sports) and grouped by severity (all, standard, strict). You can use the URL categories to define an application. Additional sub-categories offer even more granular filtering for application creation. You can select a single or multiple URL categories for an application.

For example: You can create an application and name it "social media" and select URL categories as Social Networking or Instant Messaging. Then you can create policies to block or restrict access during work hours.

When adding an application, you can use the URL Categories type to define destinations by categories such as entertainment, shopping, and sports.

NOTE: This option requires an IDP/URL Filtering license. It is packaged with some devices and can be purchased in a security bundle.

For example, create an application called Social Media and select URL categories Social Networking and Instant Messaging. Later, on the Application Policies page, create policies to block or restrict access to these URLs during work hours.

Figure 33: URL Category Examples

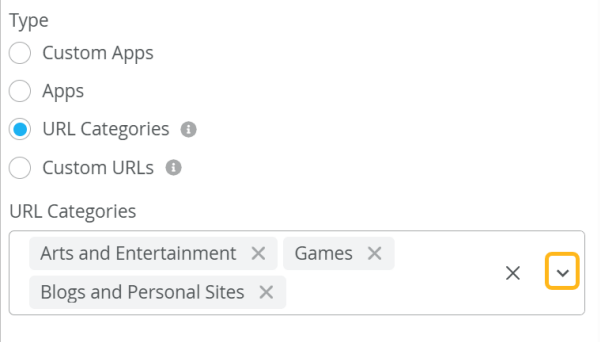

NOTE: For Session Smart Routers, we recommend configuring Applications with URL categories on spoke devices only, and not on hub devices.

Custom URLs

When defining applications, you can enter custom URLs for services not covered by predefined applications.

You can enter:

- Exact domain names. `example.com`
- Wildcard domains using an asterisk. Example: `*.example.com`

  With this approach, you can group related services under one application. For example, `*.google.com` includes Gmail, Drive, Meet, and other Google sites.
- Use a comma separator to specify multiple URLs.
- You can specify up to 15 URL patterns for an application.

- Only the * wildcard is supported.
- You can view the supported patterns by hovering the mouse over the tooltip icon. Note that you can use the *https://abc.com* pattern only for SRX Series devices.

Figure 34: Custom URLs

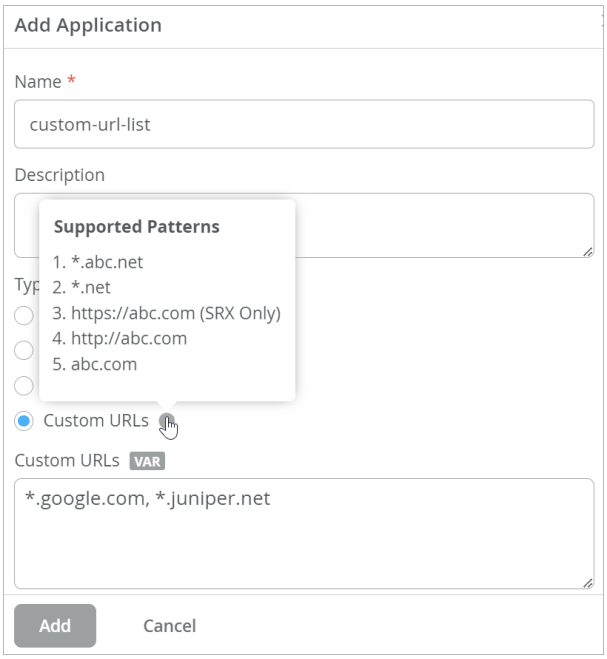

Advanced Settings

Under **Advanced Settings**, specify the **Traffic Type**. Keep **Default** for general traffic, select a preset traffic type, or select **Custom**.

- If you select a preset traffic type, you'll see the values for settings such as failover (SSR only), latency, jitter, and loss.

> **NOTE**: For Apps and URL Categories, you can only select a specific traffic type after you select the **Override Settings** check box.

- If you select **Custom** as the traffic type, also select a **Traffic Class**, and then adjust the preset values as described in the following table.

Table 16: Advanced Settings

Fields	Description
Failover Policy (SSR Only)	*Applies only to SSRs* • Revertible—Traffic automatically switches back to the primary link when the primary link recovers. • Non-Revertible—Requires manual intervention to revert to the primary link. When traffic switches to the secondary link due to primary link failure, it does not automatically revert back to the primary link. • None—Disable session failover. If the primary link on your device fails to meet the Service Level Agreement (SLA), existing sessions remain on the primary link, while new sessions will be redirected to the secondary link. When the primary link recovers and meets the SLA, existing sessions on the secondary link will continue, and any new sessions will start on the primary link. This behavior remains consistent even if the entire link goes down.
Traffic Class	These options provide granular control over traffic prioritization. Specify the priority for this application. • Best Effort—No special treatment, suitable for noncritical data. • Medium—Prioritized over Best Effort, used for non-latency-sensitive applications. • High—Critical applications with low latency requirements. • Low—Background or non-urgent traffic
DSCP Class (SSR Only)	*Applies only to traffic through SSRs* The Differentiated Services Code Point (DSCP) value tags packets for specialized handling across the network. When you select a traffic class (Best Effort, High, Medium, or Low), the applicable default DSCP Class value is displayed as a help text. You can choose to override it to fine-tune this application to your specifications. Range: 0-63
Maximum Latency	By setting maximum latency in milliseconds, you can ensure that delay-sensitive applications like voice and video meet your performance requirements. Based on this threshold, SD-WAN avoids links with excessive delay. Range: 0-4294967295 (milliseconds)
Maximum Jitter	You can constrain jitter, or the variation in latency, by setting a maximum value in milliseconds. Based on this threshold, SD-WAN selects stable links to maintain predictable performance. Range: 0-4294967295 (milliseconds)

Table 16: Advanced Settings *(Continued)*

Fields	Description
Maximum Loss	For further fine-tuning, you can specify the maximum acceptable percentage of packet loss to maintain application reliability. Based on this threshold, SD-WAN avoids links with high packet loss rates. Range: 0-100

Enable Application Visibility (SRX Only)

SUMMARY

Follow these steps to enable application-aware security services.

IN THIS SECTION

The Juniper Networks Application Security (AppSecure) feature is a suite of application-aware security services for the Juniper Networks® SRX Series Firewalls. AppSecure enables you to see the applications on your network and learn how they work. It enables you to observe their behavioral characteristics and assess their relative risk, which allows the Juniper Mist™ cloud to track and report applications passing through the device.

Before You Begin

Consult this list to ensure that you have the licenses and application signatures necessary to enable application visibility.

- You need a valid AppSecure license on your SRX Series Firewall to use the feature. Use the `show system license` command to check if your device has the license. For details about license requirements and installation, see Juniper Licensing User Guide.
- We recommend using the latest version of application signatures. To install the latest version of application signatures, run the following commands on your device:

1. Download the application signature package version on your device. The command downloads the latest version of the package.

```
user@host> request services application-identifications download
Please use command "request services application-identification download status"
 to check status
```

```
user@host> request services application-identifications download status
Application package 3410 is downloaded successfully.
```

2. Install the application signature package version on your device.

```
user@host> request services application-identification install
Please use command "request services application-identification install status" to check
status
and use command "request services application-identification proto-bundle-status" to check
protocol bundle status
```

3. Verify the application signature package version installed on your device.

```
user@host> show services application-identification version
Application package version: 3410
```

For more details, see Predefined Application Signatures for Application Identification.

You can see the application signature version in the Juniper Mist cloud portal of your device under the **SECURITY SERVICES** panel.

Figure 35: Check Application Security (AppSecure) Version

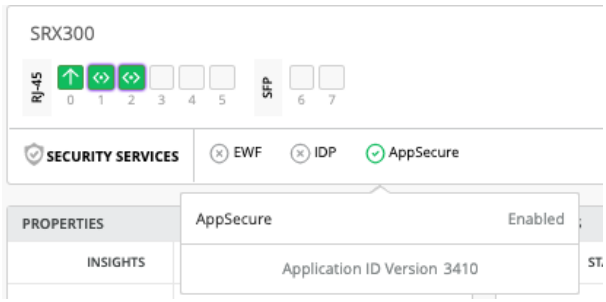

Enable Application Visibility During Device Adoption

If you're onboarding new devices, you can enable application visibility as part of the normal device adoption workflow. This option is available in the site assignment settings.

To enable application visibility while assigning a device to a site:

1. From the left menu, click **Organization** > **Admin**> **Inventory**.
2. Click the **WAN Edges** button at the top of the Inventory page.
3. Click the **Adopt WAN Edges** button the top-right corner of the page.

Figure 36: WAN Edge Adoption Commands

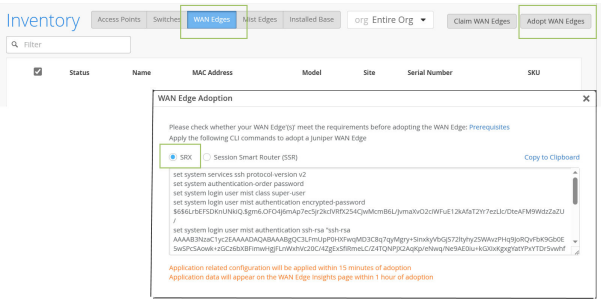

Juniper Mist generates a code snippet in the **WAN Edge Adoption** window.

4. Ensure that you've selected the SRX option, and click **Copy to Clipboard**.
5. Close the pop-up window.
6. Go to the CLI for your SRX Series Firewall, enter configuration mode, paste the code, and commit the configuration.

   This code creates the following settings on your SRX Series Firewall:

   - Enable SSH.
   - Create a Juniper Mist cloud user.
   - Create a device ID and credentials.
   - Set up the outbound SSH client and associated timers.

   After you commit the configuration on your SRX Series Firewall, the device appears on the Inventory page in the Juniper Mist portal.
7. On the Inventory page, select the check box for the newly added SRX Series Firewall, then click **More** at the top of the page, and then click **Assign to Site**.
8. In the Assign Gateways pop-up window, enter these settings:

   a. Select the site.

b. To manage the configuration in Juniper Mist, select the **Manage configuration** check box.

c. Select the appropriate option to describe your AppTrack license.

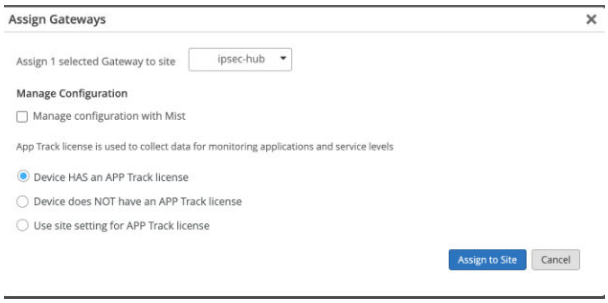

- **Device has an App Track license**—Application visibility is already enabled on the device.
- **Device does NOT have an App Track license**—The device does not have application security license.
- **Use site setting for App Track license**—Enable application visibility under site setting options.

d. Click **Assign to Site**.

9. **For a device-based license or no license**—Complete these additional steps if you selected **Device has an App Track license** or **Device does NOT have an App Track license** in the gateways assignment window.

a. On the Inventory page, click your newly assigned SRX.

b. On the device details page, scroll down to the **Application Visibility** section.

c. Select the same license option that you selected in the site assignment window: **Device has an App Track license** or **Device does NOT have an App Track license**.

d. Click **Save** at the top right corner of the device details page.

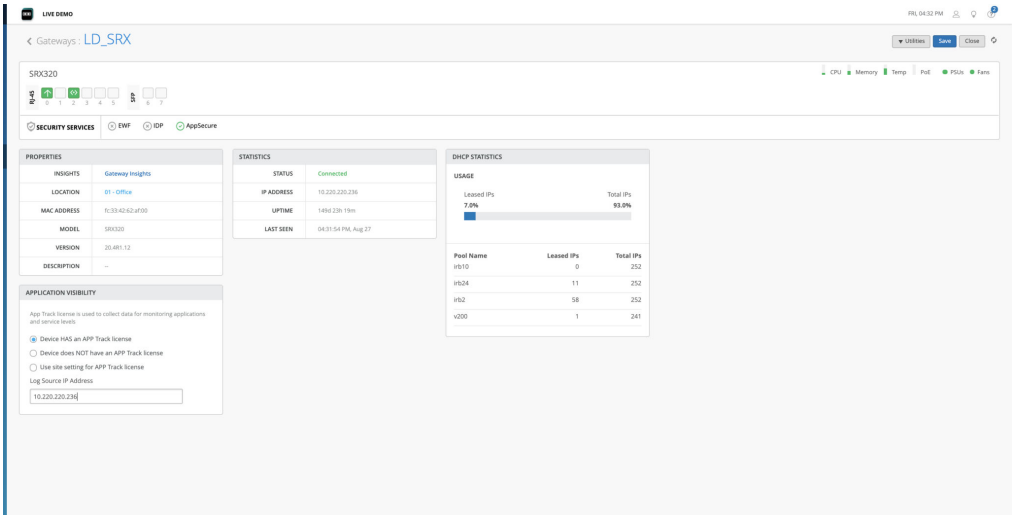

10. **For a site-based license**—Complete these additional steps if you selected **Use site setting for App Track license** in the gateway assignment window.

    a. From the left menu, select **Organization** > **Admin** > **Site Configuration**.

    b. Select the site that you assigned to the newly added SRX device.

    c. Scroll down to the **WAN Edge Advanced Security** section.

    d. To enable application visibility, select the check box for **My SRX devices have an App Track license**.

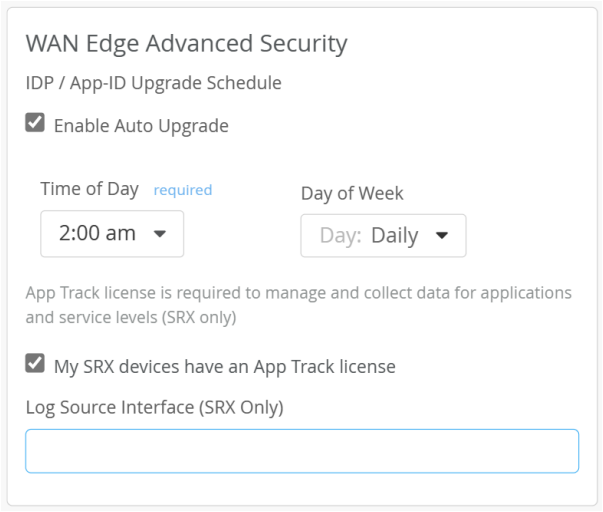

    e. In the **Log Source Interface** box, enter the IP address of the interface to use as the source address for log messages.

    This interface needs connectivity to the cloud or Internet. It acts as the source address for log messages for the application session records.

    f. Save the site configuration.

You can verify API messages of **/sites/site-id/setting** to see the following options, depending on whether you selected or unselected **My SRX devices have an App Track License**:

- • The **“gateway_mgmt“: {“app_usage“: True}** message indicates that the check box is selected.
  - The **“gateway_mgmt“: {“app_usage“: False}** message indicates that the check box is not selected.

Example:

```
GET /api/v1/sites/232527fe-4126-40bb-8c78-2c8d1dfed043/setting
HTTP 200 OK
Allow: OPTIONS, GET, PUT
Content-Type: application/json
Vary: Accept
```

```
{
   "switch_mgmt": {
      "root_password": "mist123"
   },
   <<< API OUTPUT TRIMMED >>>
   "zone": {
      "autozones_enabled": false,
      "autozones_rssi": -70
   },
   "gateway_mgmt": {
      "app_usage": true,
      "security_log_source_interface": "ge-0/0/0"
   },
   "id": "86f13595-9599-48a7-8c26-ad98a702b9e5",
   "for_site": true,
   "site_id": "232527fe-4126-40bb-8c78-2c8d1dfed043",
   "org_id": "001f3ef8-d69d-4780-b9c3-7a1f3cb123f0",
   "created_time": 1599493540,
   "modified_time": 1600069580
```

NOTE: The `gateway_mgmt` section appears only if you used the site settings option when enabling application visibility.

Enable Application Visibility After Initial Onboarding

Use this procedure if you want to enable application visibility on devices that you previously adopted into your organization and assigned to a site.

To enable application visibility on an SRX Series Firewall that you already assigned to a site:

1. From the left menu, select **Organization** > **Admin** > **Site Configuration**.
2. Select the site.
3. Scroll down to the **WAN Edge Advanced Security** section.

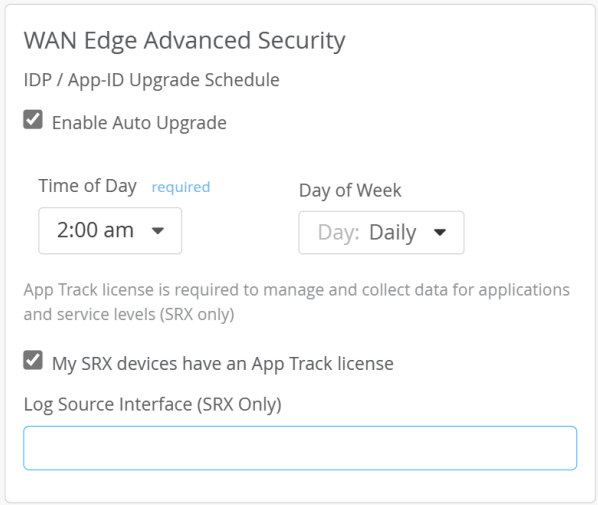

4. To enable application visibility, select the box for **My SRX devices have an App Track license**.
5. For **Log Source Interface**, enter the IP address of an interface on SRX Series Firewall that has connectivity to the cloud or Internet.

   This interface acts as the source address for log messages for the application session records.
6. Click **Save**.

To view the applications details, click **Monitor** > **Service Levels**. Click the **Insights** tab, and then scroll down to the **Applications** section to get details about applications usage.

General WAN Edge Settings

SUMMARY

Use this information to enter the settings at the top of a WAN Edge template, hub profile, or device configuration.

IN THIS SECTION

Navigating to a Template, Profile, or Device Configuration

To get started with these settings, you need to go to the appropriate configuration page:

- For a WAN Edge template—From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create one.
- For a hub profile—From the left menu, select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create one.
- For an individually managed WAN Edge device—From the left menu, select **WAN Edges** > **WAN Edges**. Click a device.

General Settings Overview

These settings appear near the top of the screen when you're viewing a WAN Edge template, a hub profile, or a device configuration.

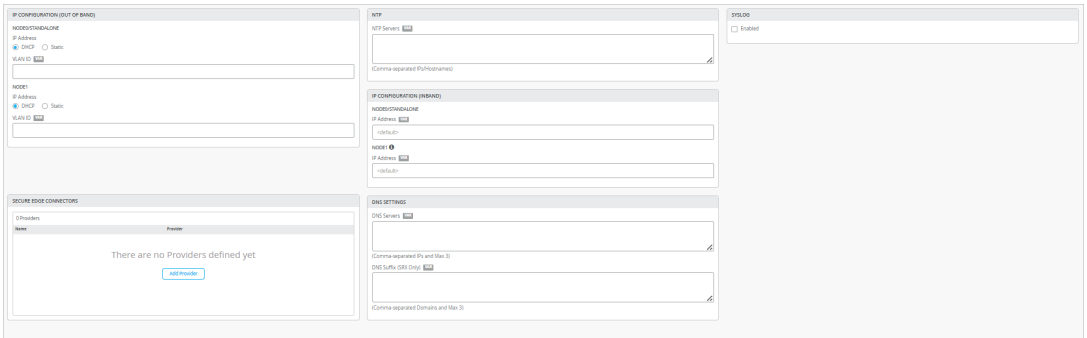

Support for IPv6

You can use IPv6 addresses to configure the network overlay for the WAN Edge device, as well as other settings. These include the WAN interface, the LAN interface, Application Policies that use IPv6 applications, as well as NTP and DNS services.

On the WAN Edge Insights page, statistics are available for selected IPv6 events and properties.

Note that IPv6 support varies between SSR and SRX devices, and not every IP address field available in the Mist dashboard necessarily supports IPv6 addresses. The Mist dashboard itself will provide the best indication of whether or not IPv6 is supported in a given circumstance. In addition to the Mist dashboard, you can use the Mist API to enable IPv6 support and configure the overlay.

For the WAN connection to the Mist cloud, only the IPv4 address is supported.

Field Descriptions

Table 17: Settings in the General Section

Fields	Description
Info	The name of the template, hub profile, or devices
Applies to Sites *(for templates) or* Applies to Device, Hub Group*(for hub profiles)*	See: • "Configure a WAN Edge Template" on page 87 • "Configure a Hub Profile" on page 82
IP Configuration (Out Of Band)	The out of band management details to manage the device. • **IP Address**—Specify the type of IP address assignment: **DHCP** or **Static**. If you select static, enter the IP address, subnet mask, and default gateway. • **VLAN ID**—Enter a value between 1 and 4094. For a high availability cluster, enter the same IP Address type and VLAN ID for both nodes, but different IP addresses.
NTP	Enter the IP address or hostname of the Network Time Protocol (NTP) server. NTP allows network devices to synchronize their clocks with a central source clock on the Internet.
DNS Settings	Enter a comma-separated list of IP addresses of Domain Name System (DNS) servers. Network devices use the DNS-name servers to resolve hostnames to IP addresses.
Syslog	Optionally, enable system logging. For help, see "Syslog Options" on page 113.

Table 17: Settings in the General Section *(Continued)*

Fields	Description
Secure Edge Connectors	Secure Edge performs traffic inspection for the WAN Edge devices managed by Juniper Mist Cloud portal. This section of the page displays all the defined providers. To add more, click **Add Provider**. Complete all required fields, and then click **Add**. For help, see "Juniper Mist Secure Edge Connector Overview" on page 222. **NOTE**: For the Secure Edge Connector configuration to take effect, you'll also need to enable an application policy with traffic steering. For help, see "Traffic Steering Rules" on page 134.

Syslog Options

SUMMARY

Use this information to send system log messages to local files, external hosts, users, or the console. Also configure content types, archive settings, and more. (Available logging options vary, depending on which settings you're configuring, such as syslog for a switch, a WAN Edge device, or other.)

Table 18: Files

Field	Description
Add File	Use the **Add File** link to specify a named file in the local file system. Syslog messages that match the specified criteria will be sent to this file. Cycle through this process as needed to specifying multiple files for different matching criteria.

Table 18: Files *(Continued)*

Field	Description
File Name	Enter a name for the file.
Match	To filter messages, enter a text string. Only log messages that contain this string will be saved to the specified file.
Explicit Priority	Select this option to include priority information in the file.
Structured Data	Select this option to use structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.
Archive	Set the maximum number of files to store, and the maximum file size in megabytes (m) or bytes.
Contents	To specify the types of log information to capture, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be saved to the specified file. For more information about the content options, see Table 21 on page 116 later in this topic.

Table 19: Hosts

Field	Description
Add Host	Use the **Add Host** link to specify an external server that is configured as a system log message host. Syslog messages that match the specified criteria will be sent to this server. Cycle through this process as needed to specifying multiple hosts for different matching criteria.

Table 19: Hosts *(Continued)*

Field	Description
Host	Enter the host.
Port	Enter the port number, from 1 through 65535.
Match	To filter messages, enter a text string. Only log messages that contain this string will be saved to the specified file.
Explicit Priority	Select this option to include priority information in the file.
Structured Data	Select this option to use structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.
Archive	Set the maximum number of files to store, and the maximum file size in megabytes (m) or bytes.
Contents	To specify the types of log information to capture, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be saved to the specified file. For more information about the content options, see Table 21 on page 116 later in this topic.

Table 20: Users

Field	Description
Add User (link)	On the Users tab, you're sending log messages to the terminal session of one or more users. To add a user, click this link, and the enter the user information. Cycle through this process as needed to specify multiple users for different matching criteria.

Table 20: Users *(Continued)*

Field	Description
User	Enter one or more usernames, separated by spaces, or enter an asterisk (*) to include all users.
Match	To filter messages, enter a text string. Only log messages that contain this string will be sent to the specified users.
Contents	To specify the log information to capture for this user, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be sent to the specified users. For more information about the content options, see Table 21 on page 116 later in this topic.

Table 21: Console Options and Content Types

Field	Description
Add Content (link)	Use the **Add Content** link to specify the types of log messages to capture. You'll cycle through this process multiple times to add different content types. For example, let's say you want to capture (1) critical authorization events, (2) warnings for authorization events, (3) critical change logs, (4) change log errors, and (5) ftp errors. For this example, you'd create five content types. **NOTE**: You can add content types in a few different places. For example, if you're working on the Console tab, you're specifying the content to send to the console. If you're working on the User tab, you're specifying the content to capture for that user.
Facility	Select a log event that you want to capture in this content type.

Table 21: Console Options and Content Types *(Continued)*

Field	Description
Level	For the specified log event, select the severity level to capture in this content type.

Table 22: Archive Tab

Field	Description
Files	Specify the maximum number of files (1-1000) to store in the system log. After this maximum is reached, log messages are archived.
Size	Specify the maximum size (65536 - 1073741824) of the system log. Select either **m** (megabytes) or **bytes**. After this maximum is reached, log messages are archived.

Table 23: General Tab

Field	Description
Time Format	Use this option to include the milliseconds, the year, or both in the syslog timestamps. • To add an option, click the plus (+) button, then click the option in the list. • To remove an option, click the X.
Routing Instance	By default, system logging traffic is sent from the management interface on your device and its associated routing instance. However, the Junos syslog client is completely VRF aware. If a server is reachable through a virtual routing and forwarding (VRF) instance, the syslog client can send log messages to the server. Use this option to specify a routing instance to use.
Network	Use this option to specify a network instance to use.

WAN Interfaces

SUMMARY

Use this information to enter the settings in the WAN section of a WAN Edge template, hub profile, or device configuration.

IN THIS SECTION

Navigating to a WAN Configuration

You'll find WANs on the configuration page for your WAN Edge templates (for spokes), hub profiles, and individually managed devices.

- For a WAN Edge template—From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create a new one. Scroll down to the **WAN** section.
- For a hub profile—From the left menu, select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create a new one. Scroll down to the **WAN** section.
- For an individually managed WAN Edge device—From the left menu, select **WAN Edges** > **WAN Edges**. Click a device. Scroll down to the **WAN** section.

WAN Configuration Overview

A WAN interface is the interface out to the Internet.

By default, Juniper Mist sets up the WNA Edge device for management connectivity to the Mist cloud by enabling DHCP on the default WAN interface.

WAN Settings

SUMMARY

On the configuration page, the WAN section lists the WANs that you've defined. To add a WAN, click **Add WANs**, enter the settings in the side panel, and then click **Add** at the bottom of the panel.

Table 24: Settings for the WAN Configuration Panel

Field	Description
Name	Enter a name to identify this interface. You can enter up to 32 letters, numbers, underscores, and dashes. The name must start and end with a letter or a number.
Description	Enter a description for this interface.
WAN Type	Select the type of WAN link. If you select LTE, also enter the **LTE APN** and **LTE Authentication** (described later in this table). **NOTE**: The DSL option is for SRX devices only.

Table 24: Settings for the WAN Configuration Panel *(Continued)*

Field	Description
Interface	Enter the name of the interface (such as ge-0/0/1, ge-0/0/1-5, or reth0). Or, for aggregation, enter a comma-separated list of interfaces. Then, select additional interface options as needed: • **Disabled**— Administratively disable the LAN port. If you disable a physical interface (for example, ge-0/0/1), all associated sub-interfaces or VLAN interfaces (for example, ge-0/0/1.200 and ge-0/0/1.100) will go down. If you disable an aggregated Ethernet interface (for example, ae1) on a WAN Edge device, only the ae interface will go down; the underlying physical interfaces in the ae bundle will remain up. • **Port Aggregation**—Group Ethernet interfaces to form a single link layer interface. You can use this for Link Aggregation Control Protocol (LACP) configuration. With LACP, devices exchange heartbeats, providing faster failover. • **Disable LACP**—Disable LACP interface. LACP is enabled by default when you enable port aggregation. You can disable it when needed, for example if you are onboarding a new device or disconnecting a device for RMA. • **Enable Force Up**—Choose this option prior to onboarding a device that is connected to the LAN port via Link Aggregation Control Protocol (LACP). For example, when onboarding a new switch to the Mist cloud, the switch is not already provisioned for LACP. • **Redundant**— *(available for all WAN Types except DSL)* Select this option to enable redundancy. Also select the **Primary Node**. The drop-down list includes the nodes specified in the IP Config section of the configuration.

Table 24: Settings for the WAN Configuration Panel *(Continued)*

Field	Description
	For SRX only, also enter the **Redundant Index** and **Redundant Group**. • **Enable "Up/Down Port" Alert Type**—Select this option to designate this interface as a critical WAN Edge port for alerts. After you save this configuration, you also need to enable the Critical WAN Edge Port Up and Critical WAN Edge Port Down alerts on the Alerts Configuration page. Then you'll receive alerts whenever this port goes down or comes back up. For more information, see the Alerts chapter of the Juniper Mist AI-Native Operations Guide. • Enable Scheduled Speed Tests—Select this option to allow Marvis to run self-driving speed tests on the selected WAN interfaces automatically during low activity times. If you enable this feature, be aware of these factors: • This feature requires a Marvis for WAN subscription for each device that you want to run Marvis self-driving speed tests on. • It's selected by default when you select Ethernet as the WAN type, and is recommended for these interfaces. • It's *not recommended* for devices with built-in LTE. The reason is that this feature uses bandwidth and can be costly if you pay for bandwidth usage through a service provider. If you have LTE interfaces connected to your Ethernet interfaces, as on Cradlepoint devices, do not enable this option. • If you enable this feature on this interface, also enable it at the organization level. To do so, go to **Organization** > **Settings** > **WAN Speed Test Scheduler**.

Table 24: Settings for the WAN Configuration Panel *(Continued)*

Field	Description
LTE APN *(if WAN Type is LTE)*	*(Optional for SRX Series Firewalls and mandatory for Session Smart Routers)* If you selected LTE as the WAN type, also complete this field. On SRX Series Firewalls, the LTE Mini-Physical Interface Module (Mini-PIM) provides wireless WAN support on the SRX300 Series and SRX550 High Memory Services Gateways. The Mini-PIM contains an integrated modem and operates over 3G and 4G networks. The Mini-PIM can be installed in any of the Mini-PIM slots on the devices. For help with installation, see https://www.juniper.net/documentation/us/en/hardware/lte-mpim-install/topics/task/lte-mpim-hardware-intalling.html. You need to set up an LTE interface on your WAN Edge device and insert the Subscriber Identity Module (SIM) in the LTE card. In the LTE APN field, enter the access point name (APN) of the gateway router. The name can contain alphanumeric characters and special characters. (Mandatory when enabling LTE with Session Smart Routers).
LTE Authentication *(if WAN Type is LTE)*	If you selected LTE as the WAN type, select an authentication type for the APN configuration: • PAP—Select this option to use Password Authentication Protocol (PAP) as the authentication method. Provide User name and Password. • CHAP—Select this option to use Challenge Handshake Authentication Protocol (CHAP) authentication as the authentication method. Provide User name and Password. • None (Default)—Select this option if you do not want to use any authentication method.

Table 24: Settings for the WAN Configuration Panel *(Continued)*

Field	Description
DSL Type *(if WAN Type is DSL)*	Currently, there is only one option, which is pre-selected for you: VDSL
VLAN ID	*(Not applicable if WAN Type is LTE)* Enter the VLAN ID for this interface.

IP Configuration	*(Not applicable if WAN Type is LTE)* Select IP configuration options.
	• DHCP—Select this option to have your WAN configuration use Dynamic Host Configuration Protocol (DHCP).
	• Static—Select this option to assign your WAN a static, unchanging IP address. Enter the IP Address, Prefix Length, and Gateway address.
	• PPPoE—Select this option to have your WAN to use Point-to-Point Protocol over Ethernet (PPPoE). Choose from the following authentication options:
	  • PAP—Select this option to use Password Authentication Protocol (PAP) as the authentication method. Provide User name and Password.
	  • CHAP—Select this option to use Challenge Handshake Authentication Protocol (CHAP) authentication as the authentication method. Provide User name and Password.
	  • None (Default)—Select this option if you do not want to use any authentication method.

Table 24: Settings for the WAN Configuration Panel *(Continued)*

Field	Description
Source NAT	Use Network Address Translation (NAT) along with advertising the public IP address unless the WAN address is a publicly routable address. Select Source NAT options: • Interface—NAT using source interface. • Pool—NAT using defined IP address pool. • Disabled—Disable source NAT
Traffic Shaping	Select **Enabled** or **Disabled**. Enable traffic shaping for better bandwidth control and traffic prioritization. Traffic shaping is driven by forwarding class parameters derived from application configuration. If you enable traffic shaping, you need to also configure a transmit cap so that the actual traffic rate aligns with the configured shaping rate, regardless of the physical interface bandwidth. If you enable traffic shaping on Session Smart Routers, you can also specify the traffic shaping percentage for an interface. In other words, you can define the maximum bandwidth (as a percentage) that should be allocated to high-, medium-, and low-priority traffic, as well as best-effort traffic. By default, traffic shaping is disabled.
Auto Negotiation	Select **Enabled** or **Disabled**.
MTU	Enter an MTU value between 256 -9192. Default is 1500.

Table 24: Settings for the WAN Configuration Panel *(Continued)*

Field	Description
Endpoints	Add endpoints to create overlays to use in your traffic steering rules. For example, create different overlays for normal traffic, critical traffic, and other use cases. • Hub to Hub Endpoints—Click **Add Hub to Hub Endpoints**. Select a WAN interface on the other hub. For example, if you're configuring the hub profile for Hub A, and you want a hub-to-hub overlay to Hub B, select a WAN interface on Hub B. • Hub to Spoke Endpoints (SSR Only)—Click **Add Hub to Spoke Endpoints**. Enter the custom endpoint for this overlay. • Overlay Mesh Endpoints (SSR Only) —Click **Add Overlay Mesh Endpoints**. Select the Mesh Name, Path, and BFD Profile.

LAN Interfaces

SUMMARY

Use this information to enter the settings in the LAN section of a WAN Edge template, hub profile, or device configuration.

IN THIS SECTION

Navigating to a LAN Configuration

You'll find LANs on the configuration page for your WAN Edge templates (for spokes), hub profiles, and individually managed devices.

- For a WAN Edge template—From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create one. Scroll down to the **LAN** section.
- For a hub profile—From the left menu, select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create a new one. Scroll down to the **LAN** section.
- For an individually managed WAN Edge device—From the left menu, select **WAN Edges** > **WAN Edges**. Click a device. Scroll down to the **LAN** section.

LAN Settings Overview

Associate your LAN with the appropriate port on the device, and add network services as needed.

The LAN section of the configuration includes these sub-sections:

- "IP Config" on page 127
- "DHCP Config" on page 127
- "Custom VR" on page 130
- "LAN Configurations" on page 131

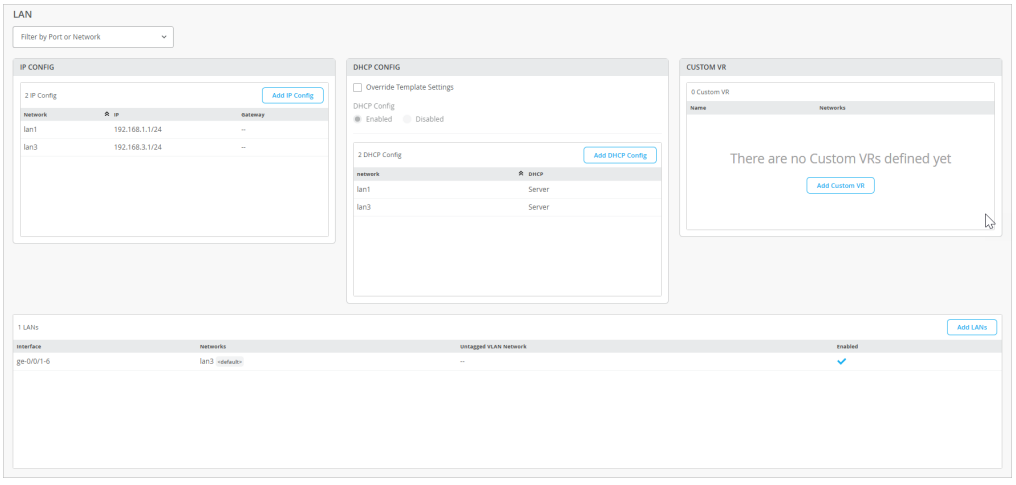

IP Config

SUMMARY

In the IP Config section, you see the networks, IP addresses, and redirect gateways (applicable to Session Smart Routers only) for the LANs that you've defined. To add a LAN, click **Add IP Config**, enter the settings in the side panel, and then click **Add** at the bottom of the panel.

Table 25: Settings in the IP Config Panel

Field	Description
Network	Select an available network from the drop-down menu. This list includes all networks that you previously added on the Networks page. For help, see "Network Settings" on page 94.
IP Address	Enter the IPv4 address for this interface.
Prefix Length	Enter the prefix length for the interface.
Redirect Gateway *(SSR Only)*	*(For Session Smart Routers only)* Enter the IP address of redirect gateway .

DHCP Config

SUMMARY

In the DHCP Config section, you can enable or disable DHCP and see the DHCP configs that you've defined. To add a DHCP config, click **Add DHCP Config**, enter the settings in the side panel, and then click **Add** at the bottom of the panel.

Table 26: Settings in the DHCP Config Panel

Field	Description
Network	Select an available network from the drop-down menu. This list includes all networks that you previously added on the Networks page. For help, see "Network Settings" on page 94.
DHCP Type	Select DHCP Server or DHCP Relay. Then complete the asterisked fields.
DHCP Relay Settings	
Servers	Enter a comma-separated list of IP addresses for the relay servers.
DHCP Server Settings	
IP Start, IP End	Use these two fields to specify the range of IP addresses to use.
Gateway	Enter the IP address of the network gateway.
Maximum Lease Time	Specify the maximum duration of a DHCP lease, from 3600 seconds (1 hour) to 604800 seconds (1 week).
DNS Servers	Enter IP address of the Domain Name System (DNS) server.

Table 26: Settings in the DHCP Config Panel *(Continued)*

Field	Description
Server Options	The Server Options table displays the server options that you've already defined. Click an existing option to edit, or click **Add Option**. The settings include: • **Code**—Select a code from the list. The list displays the option numbers with brief descriptions. The description indicates the type of value you will need to enter in the Value field. • **Type**—This becomes a read-only field based on what you already selected in the Code field (this field indicates the requirements for the selected code). For example, if you selected Option 15 (domain-name), the type is FQDN, which requires a fully qualified domain name in the Value field. • **Value**—Enter an appropriate value. For example, if you selected Option 67 (boot-file-name) for the code, the type is string, and the value must be a string that specifies the boot file name. After completing the fields, click the check mark in the Add Server Option title bar.

Table 26: Settings in the DHCP Config Panel *(Continued)*

Field	Description
Static Reservations	This section lists the addresses in the DHCP range that are reserved for devices that require static IP addresses. Click an existing reservation to edit, or click **Add Reservation**. The settings include: • **Name**—Enter a descriptive name to identify this reservation. The name can include up to 64 characters. It can contain only letters, numbers, underscores, and dashes. • **MAC Address**—Enter the MAC address of the device that will use this static IP address. • **IP Address**—Enter the IP address to be reserved. This address must be within the range that you identified in the IP Start and IP End fields (near the top of the DHCP Config panel). After completing the fields, click the check mark in the Add Static Reservation title bar.

Custom VR

SUMMARY

If you've set up Virtual Routing and Forwarding (VRF) instances for your WAN, you also can configure VRF Route Leaking (propagation) to share route information across these instances. It provides route isolation and segmentation from other routes. A common implementation of this feature is to isolate guest, PCI, or IoT network. Put each network in a separate VRF instance.

Table 27: Settings for the Custom VR Panel

Field	Description
Name	Enter a descriptive name to identify this virtual router.
Network	Select an available network from the drop-down menu. This list includes all networks that you previously added on the Networks page. For help, see "Network Settings" on page 94.
Extra Routes	This section lists the extra routes that you've defined for this VR. Click an existing route to edit, or click **Add Extra Routes**. The settings include: • **Prefix**—Enter the prefix IP address in the following format: *xxx.xxx.xxx.xxx/xx*. • **Gateway**—Enter the IP address of the gateway.

LAN Configurations

SUMMARY

The LAN section lists the LANs that you've defined. To add a LAN, click **Add LANs**, enter the settings in the side panel, and then click **Add** at the bottom of the panel.

Table 28: Settings for the LAN Configuration Panel

Field	Description
Interface	Enter the interface, multiple interfaces separated by commas, or a range of interfaces. The number of interfaces depends on the device. For example, if the device only supports four interfaces, you can add up to four in the LACP bundle. Examples: • Single—ge-0/0/3 • Comma-separated list—ge-0/0/10,ge-0/0/11,ge-0/0/12,ge-0/0/13 • Range—ge-0/0/20-23 Then select additional interface options as needed: • **Disabled**— Administratively disable the LAN port. If you disable a physical interface (for example, ge-0/0/1), all associated sub-interfaces or VLAN interfaces (for example, ge-0/0/1.200 and ge-0/0/1.100) will go down. If you disable an aggregated Ethernet interface (for example, ae1) on an SSR device, only the ae interface will go down; the underlying physical interfaces in the ae bundle will remain up. • **Port Aggregation**—Group Ethernet interfaces to form a single link layer interface. You can use this for Link Aggregation Control Protocol (LACP) configuration. With LACP, devices exchange heartbeats, providing faster failover. • **Disable LACP**—Disable LACP interface. LACP is enabled by default when you enable port aggregation. You can disable it when needed, for example if you are onboarding a new device or disconnecting a device for RMA.

Table 28: Settings for the LAN Configuration Panel *(Continued)*

Field	Description
	• **Enable Force Up**—Choose this option prior to onboarding a device that is connected to the LAN port via Link Aggregation Control Protocol (LACP). For example, when onboarding a new switch to the Mist cloud, the switch is not already provisioned for LACP. Setting **Enable Force Up** temporarily disables LACP and forces the first specified Ethernet interface to the *up* state, which in turn allows the switch to connect to the Mist cloud using zero-touch provisioning (ZTP). Then it will retrieve the configuration files needed to complete the onboarding. In this way, it gets configured, including LACP provisioning. • **Redundant**— Enable redundancy. Then enter the **Redundant Index** and **Redundant Group** (SRX only) and the Primary Node. The drop-down list includes the nodes specified in the IP Config section of the configuration. • **Enable "Up/Down Port" Alert Type**—Select this option to designate this interface as a critical WAN Edge port for alerts. After you save this configuration, you also need to enable the Critical WAN Edge Port Up and Critical WAN Edge Port Down alerts on the Alerts Configuration page. Then you'll receive alerts whenever this port goes down or comes back up. For more information, see the Alerts chapter of the Juniper Mist AI-Native Operations Guide.
Description	Enter a description to identify this interface.
Network	Select a network from the list. When you do, the remaining configuration is filled in automatically. The Network list includes all networks that you previously added on the Networks page. For help, see "Network Settings" on page 94.

Table 28: Settings for the LAN Configuration Panel *(Continued)*

Field	Description
Untagged VLAN Network *(SRX Only)*	If this interface is for an SRX, select the network to use for untagged VLAN traffic.

Traffic Steering Rules

SUMMARY

Use this information to enter the settings in the Traffic Steering section of a WAN Edge template, hub profile, or device configuration.

IN THIS SECTION

Navigating to a Traffic Steering Rule

You'll find traffic steering policies on the configuration page for your WAN Edge templates (for spokes), hub profiles, and individually managed devices.

- For a WAN Edge template—From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create a new one. Scroll down to the **Traffic Steering** section.
- For a hub profile—From the left menu, select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create a new one. Scroll down to the **Traffic Steering** section.
- For an individually managed WAN Edge device—From the left menu, select **WAN Edges** > **WAN Edges**. Click a device. Scroll down to the **Traffic Steering** section.

Traffic Steering Overview

You can think of *traffic steering* in Juniper Mist as *how* the WAN Edge device uses your LAN and WAN interfaces to connect your users to your applications. Define traffic steering rules to define the various paths that application traffic can traverse to reach its intended destination.

When an application has multiple available paths, you can further restrict the traffic to a subset of those paths and establish a preference order. Additionally, traffic steering enables the loading and balancing of multiple streams across the available paths for optimal performance and redundancy.

Traffic steering is where you define the different paths that application traffic can take to traverse the network. The paths that you configure within traffic steering determine the destination zone. For any traffic steering rule, you must define the paths for traffic to traverse and strategies for utilizing those paths.

NOTE: If you're planning to create application policies for an SRX Series Firewall, you'll also need to create traffic steering policies. Every application policy needs to be associated with a traffic steering rule.

Order of Traffic Steering Rules

Be aware of these factors when defining rules.

- SRX, being a zone-based firewall, determines the destination zone based on the paths configured within the traffic steering policy. This integration ensures that security policies are enforced in conjunction with the defined traffic paths.
- SSR employs a proprietary traffic steering solution known as Secure Vector Routing. This innovative approach determines the next hop and the optimal vector to the destination, leveraging advanced algorithms and network intelligence to make informed routing decisions.

Settings for Traffic Steering

Table 29: Settings for Traffic Steering

Field	Description
Name	Enter a descriptive name to identify this traffic steering rule. The name can contain up to 32 letters, numbers, underscores, and dashes. It must start and end with a letter or number.
Strategy	• Ordered—Starts with a specified path and fails over to backup path(s) when needed. • Weighted—Selects the paths with lowest cost; if paths have equal costs, trraffic is load-balanced. • Equal-cost multipath—Load balances traffic equally across multiple paths.
Paths	Displays all the paths that you've already defined. Click a path or click **Add Paths**. Then, enter the path settings, and click the check mark in the Add Paths title bar. The settings include: • Type and Name—Select a type from the drop-down list. Based on the selected type, select the network, name, or provider. • Cost*(applicable to the Weighted strategy only)*—Enter a number or variable to indicate the relative cost of this path. The path with the lowest cost will be prefered over those with higher costs.

Configure an In-Band IP Address and Traffic Steering (SSR Only)

SUMMARY

Take this approach if you host your management services in data centers that are accessible only through the overlay.

When no source NAT rules are applied, sessions for management services such as DNS, NTP, and Mist traffic from the device will use the in-band IP address when exiting through the hub LAN.

To provide granular control over routing of management traffic that originates from the WAN Edge device, you can enable traffic steering. Choose from various options (ordered, ECMP, or weighted) across both overlay and underlay.

You can also use traffic steering to define backup paths, such as a cellular WAN link, which ensures uninterrupted management connectivity.

NOTE: This feature applies only to SSR Series Routers. For SRX Series Firewalls, the IP is auto generated.

Before You Begin: Create your WAN Edge template, hub profile, or standalone device configuration. For help, see "WAN Assurance Configuration Overview" on page 75.

To configure an in-band IP address and traffic steering:

1. Navigate to your WAN Edge template, hub profile, or standalone WAN Edge device.
2. In the IP Configuration (In-Band) section, enter the IP address.

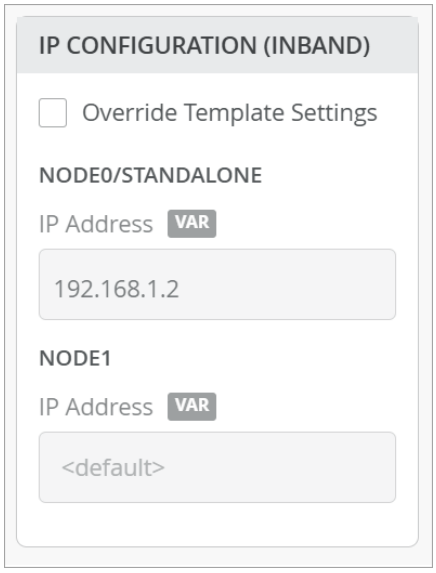

NOTE: In the case of a high availability cluster, enter two unique IP addresses (one for each node).

3. In the Traffic Steering section, click **Add Traffic Steering**, and then enter the settings.

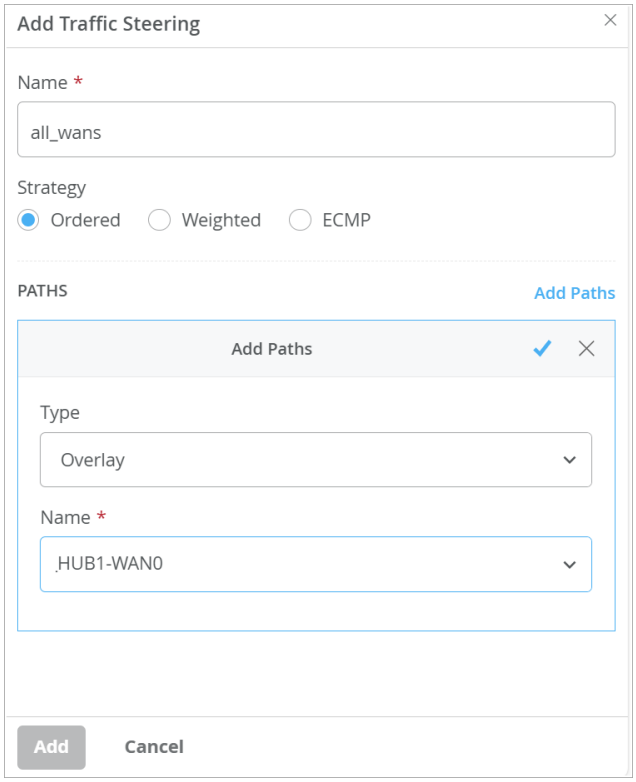

- **Name**—Enter a name for the traffic steering rule.
- **Strategy**—Select the strategy (Ordered or ECMP or Weighted).

- **Paths**—Click **Add Paths**, select **Overlay** as the type, and select the previously created hub.

> **NOTE**: An Overlay Hub Endpoint must be defined in the WAN configuration prior to adding paths. This enables you to make a selection in the overlay Name field.

4. Click **Add** to save the traffic steering details.

> **NOTE**: Mist traffic cannot be sent to the overlay; do not select the type as **Overlay** for any Mist traffic when you create a traffic steering rule.

5. In the Application Policies section, complete these steps:

   a. Click the **Device Out** tab.

   You'll see built-in network policies provided by Juniper Mist that are configured by default.

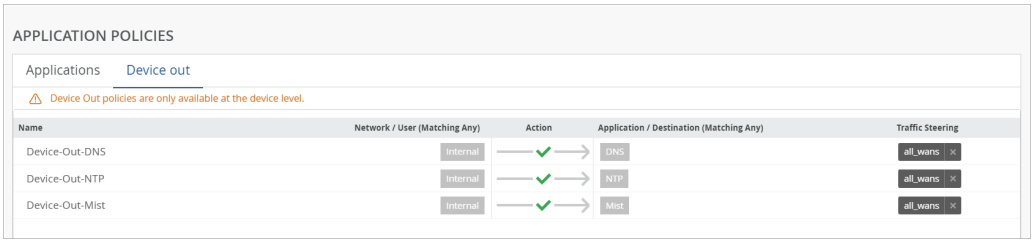

   b. In the **Traffic Steering** column, remove the **<default>** traffic sterring profile, then add the traffic steering profile you created in the previous step.

6. Click **Save** to save the changes in the template.

Configure Path Selection from Hub-to-Spoke with Traffic Steering

SUMMARY

Follow these steps to set up traffic steering by identifying hub-to-spoke paths.

Juniper Mist™ allows you to influence path selection for traffic going from hub-to-spoke. This comes in handy when you have multiple spokes that the hub is trying to reach and you want to have granular control over which spoke interface traffic will arrive on. This configuration gives you full path control in both directions.

Before You Begin

Configure your hub profiles.

To configure hub-to-spoke traffic steering:

1. From the left menu, select **Organization** > **WAN** > **Hub Profiles**.
2. Select the profile that you want to modify.
3. In the WAN section of the Hub Profile page, click the WAN interface that you want to define the logical spoke endpoints for.
4. In the **Edit WAN Configuration** panel, under **HUB TO SPOKE ENDPOINTS**, complete these steps:
   a. Notice that the **Default Endpoint** has automatically been defined for you. This represents the endpoints you are connecting to on the hub.
   b. To specify the spokes that you want traffic to arrive on, click **Add Hub to Spoke Endpoints (SSR Only)**.
   c. Specify the endpoints.

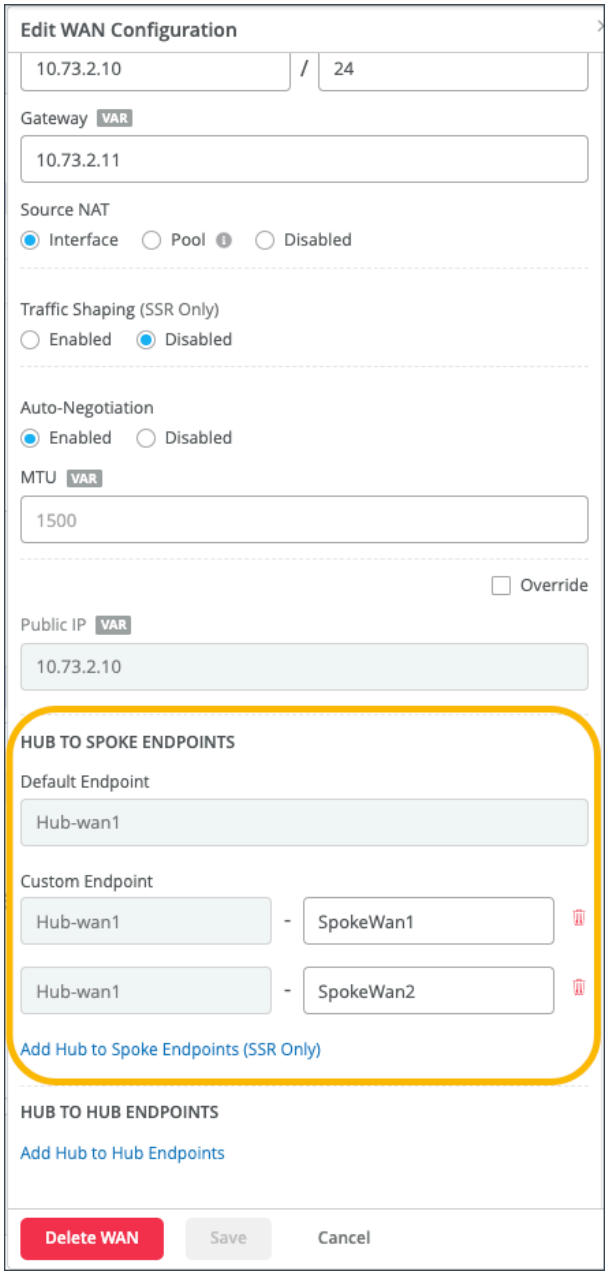

d. Click **Save** at the bottom of the window.

5. In the **TRAFFIC STEERING** section of the Hub Profile page, click **Add Traffic Steering**.
6. Enter a Name and Strategy, and then click **Add Paths**.
7. Keep the defaulted **Type** as **Overlay**.
8. For **Name**, select the name of the first hub-to-spoke endpoint you created.

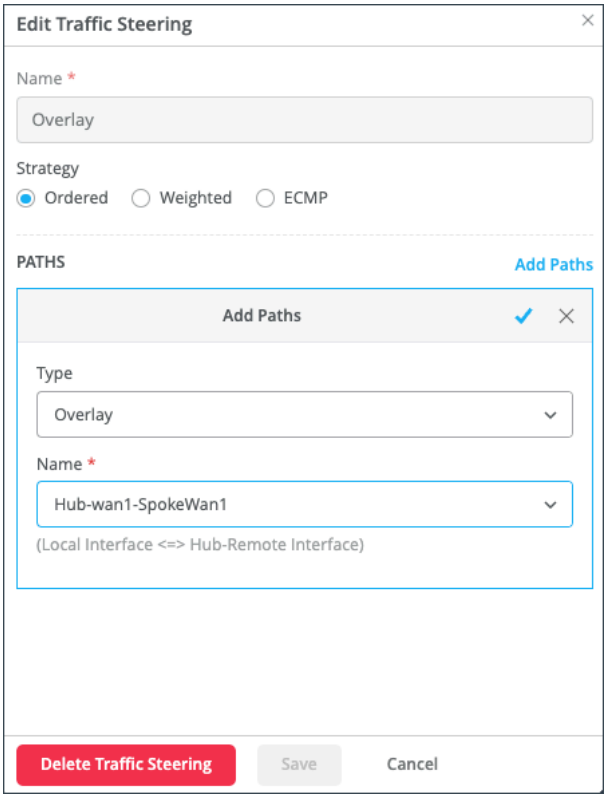

a. Click the blue check mark to save the path.

b. Select **Add Paths** again and keep the defaulted **Type** as **Overlay**.

c. In the **Name** field, select the name of the second hub-to-spoke endpoint you created.

d. Click **Save** at the bottom of the panel.

e. Click **Save** at the top-right corner of the hub profile.

You must now configure this on the spoke in order to accomplish hub-to-spoke traffic steering from end to end.

9. From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**.
10. Select your spoke template.
11. Click the WAN configuration that you want to modify.
12. In the settings panel, under **OVERLAY HUB ENDPOINTS**, click the endpoint for the selected WAN configuration (Example: Hub-wan1-SpokeWan1).

    You're selecting the overlay endpoint to make the connection between hub and spoke.
13. Click **Save**.

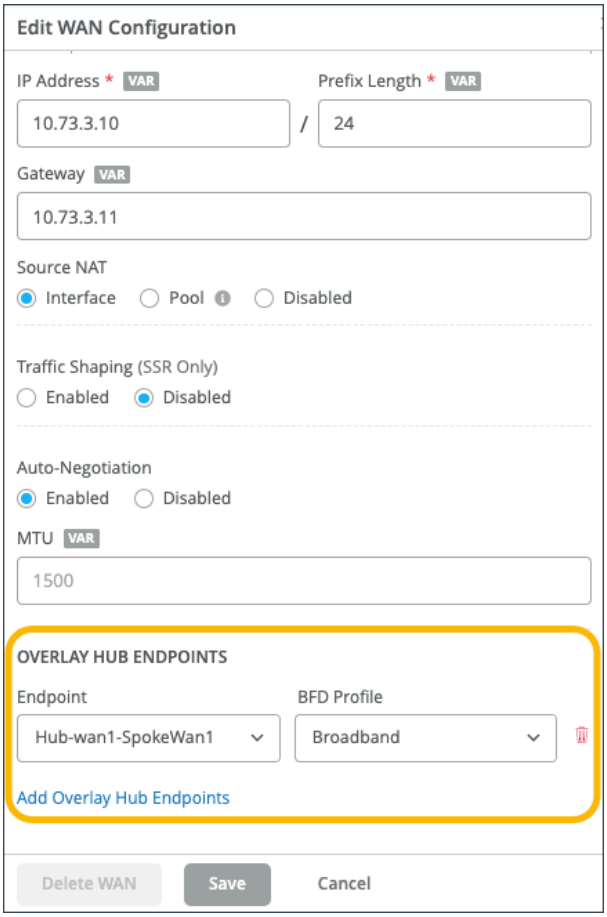

14. Select the second WAN configuration (Example: WAN2) and then set the corresponding endpoint (Example: Hub-wan1-SpokeWan2). Click **Save**.

    In the WAN section, in the Overlay Hub Endpoints column, you can see the endpoints you configured which indicate how you want your spokes to connect to the hub.

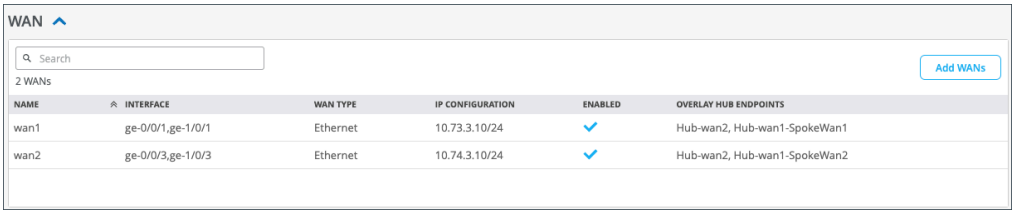

15. In the **TRAFFIC STEERING** section, click the overlay traffic steering policy that you want to configure.
16. In the **Edit Traffic Steering** panel, select **Add Paths**.
17. Keep **Overlay** as the **Type**.
18. Under **Name**, click the name of the first hub-to-spoke endpoint you created (Example: Hub-wan1-SpokeWan1).

    a. Click the blue check mark to save the path.

b. Click **Add Paths** again.

c. Keep **Overlay** as the **Type**.

d. In the **Name** field, select the name of the second hub-to-spoke endpoint you created (Example: Hub-wan1-SpokeWan2).

e. Continue to add paths as needed.

> **NOTE**: The order of path preference depends on what you configure for a Traffic Steering Strategy. You can configure a Strategy of Ordered, Weighted, or Equal-cost Multipath (ECMP). For help, see "Traffic Steering Rules" on page 134.

f. Click **Save** at the bottom of the panel.

g. Click **Save** at the top-right corner of the WAN Edge Template.

> **NOTE**: To accomplish the above behavior for spoke to hub traffic, configure Traffic Steering on the spoke.

Application Policies

SUMMARY

Use this information to add application policies to your organization, a WAN Edge template, a hub profile, or a device configuration.

IN THIS SECTION

Application Policies Overview

IN THIS SECTION

You can think of *application policies* in Juniper Mist as pulling together networks, applications, and traffic steering to determine *where* sessions are delivered. You're creating intention-driven policies by setting rules that govern access.

Use application policies to define which networks and users can access which applications, and which traffic-steering policies are used.

Understanding Policy Scope

You can define application policies at various levels to ensure that your policies meet the various needs of your organization and the many network users within it.

When you define an application policy at the organization-level, you can import and use the policy in multiple WAN Edge templates or in hub profiles. That is, you can follow the “define once, use multiple times” model.

For more particular use cases involving specific sites and devices, you can define an application policy in a WAN Edge template, hub profile, or device configuration. In this case, the scope of the policy is limited to the sites and devices that you specify. You cannot re-use this type of policy in other templates or profiles.

Design Policies to Support Traffic Flows

Before you create your application policies, consider your traffic-flow requirements. This example shows traffic flows for a corporate VPN.

Figure 37: Traffic Flow and Distribution

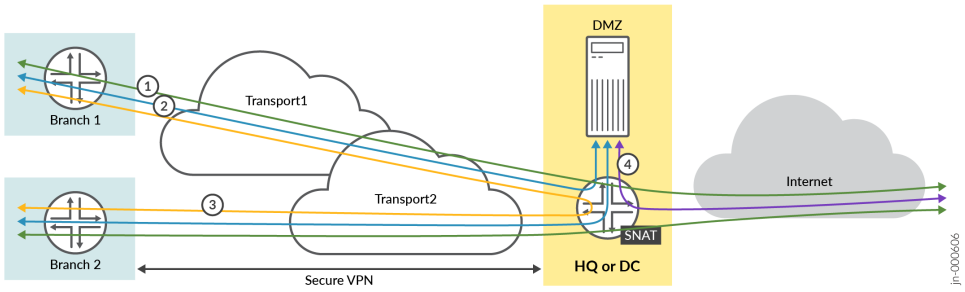

For this scenario, you need the following application rules:

- Policy 1—Allows traffic to flow between the spoke sites and the hub, as depicted by the blue lines.
- Policy 2—Allows traffic to flow between the spoke devices and the Internet via the hub device. In this case, the hub applies source NAT to the traffic and routes traffic to a WAN interface, as defined in the hub profile. This rule is general, so you should place it after the specific rules. Juniper Mist cloud evaluates and applies application policies in the order in which the policies are listed.
- Policy 3—Allows spoke-to-spoke traffic, as depicted by the yellow lines. This traffic goes through the corporate LAN via the overlay.

> **NOTE**: This may not be feasible in the real world except on expensive MPLS networks with managed IPs. Managed IPs send traffic directly from spoke-to-spoke. This type of traffic usually flows through a hub device.

- Policy 4—Allows traffic to flow between the demilitarized zone (DMZ) and the Internet, as depicted by the purple line.

The following example shows how policies appear in the Application Policies area of a configuration screen.

Figure 38: Sample Application Policies

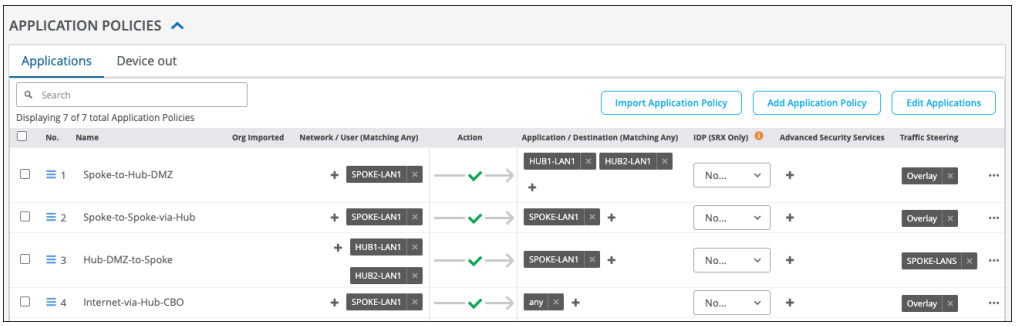

Table 30: Application Policies Configuration

S.No	Rule Name	Network	Action	Destination	Steering
1	Spoke-to-Hub-DMZ	SPOKE-LAN1	Pass	HUB1-LAN1 + HUB2-LAN1	Overlay
2	Spoke-to-Spoke-via-Hub	SPOKE-LAN1	Pass	SPOKE-LAN1	Overlay
3	Hub-DMZ-to-Spoke	HUB1-LAN1 + HUB2-LAN1	Pass	SPOKE-LAN1	SPOKE-LANS
4	Internet-via-Hub-CBO	SPOKE-LAN1	Pass	ANY	Overlay

NOTE: General Notes

- We recommend placing the most specific rules at the top.
- We recommend placing global rules towards the end of the policy rules list.
- Use the same name for the network on both sides for Session Smart Router for traffic to traverse between a hub and a spoke. The network name for the Session Smart Router must be identical to the security tenant used for traffic isolation. Because of this, the network name must match on both sides.
- Application policies that allow traffic from WAN to LAN must be configured with a network that has Destination NAT configured. This configuration enables the application policy to be pushed to the SRX Series firewall.

SSR Notes

- For SSR, you can list the policies in any order.
- For SSR, traffic steering is optional. When you use Session Smart Router, the system announces all routes on each LAN interface using the iBGP-based route distribution.
- Use the same name for the network on both sides for Session Smart Router for traffic to traverse between a hub and a spoke. The network name for the Session Smart Router must be identical to the security tenant used for traffic isolation. Because of this, the network name must match on both sides.
- Application policies that allow traffic from WAN to LAN must be configured with a network that has Destination NAT configured. This configuration enables the application policy to be pushed to the SRX Series firewall.

SRX Notes

- For SRX, you must associate a traffic steering policy to the application policy.
- Application policies that allow traffic from WAN to LAN must be configured with a network that has Destination NAT configured. This configuration enables the application policy to be pushed to the SRX Series firewall.

Navigating to Application Policies

From the left menu of the Juniper Mist portal, navigate as follows:

- Organization-level policy—Select **Organization** > **WAN** > **Application Policy** to create a policy at the organization level.
- Template-specific policy—Select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create a new one. Scroll down to the **Application Policies** section.
- Profile-specific policy—Select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create a new one. Scroll down to the **Application Policies** section.
- Device-specific policy—Select **WAN Edges** > **WAN Edges**. Click a device. Scroll down to the **Application Policies** section.

Using Same IP Addresses/Prefixes in Networks and Applications

In the application policies configuration, **Network/Users** belong to the source zone, and **Applications/Destination** belong to the destination zone.

You can use the same IP addresses and prefixes for both networks and applications when you define them for different purposes. That is, they act as a source in one policy and as a destination in another policy.

Consider the policies in Figure 39 on page 149.

Figure 39: Application Policies Details

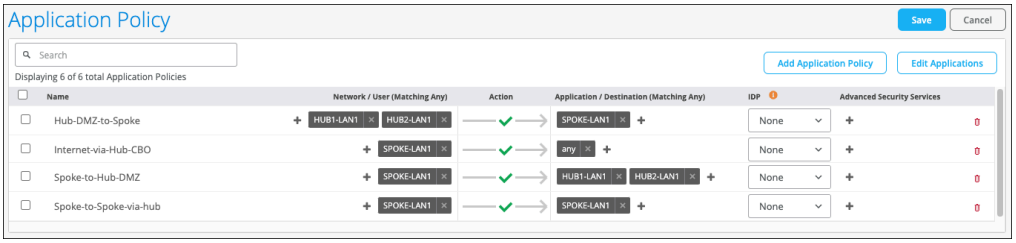

Here, you have a **Network/Users** SPOKE-LAN1 that has an IP address 192.168.200.0/24 for a spoke LAN interface. The screenshot shows that the following policies are using the same network in different ways:

- **Spoke-to-Spoke-via-Hub**—This policy allows inbound and outbound spoke-to-spoke traffic through a hub. Here, we defined *SPOKE-LAN1* as both a network and as an application.
- **Spoke-to-Hub-DMZ**—This policy allows spoke-to-hub traffic. Here, we defined *SPOKE-LAN1* as a network.
- **Hub-DMZ-to-Spoke**—This policy allows hub-to-spoke traffic. Here, we defined *SPOKE-LAN1* as an application.

Configure an Application Policy

Before You Begin: Before you define application policies, you must first create networks and applications. If you want to include traffic-steering or IDP profiles, add them as well. For help with these items, see:

- "Network Settings" on page 94
- "Application Settings" on page 98

- "Traffic Steering Rules" on page 134
- "Intrusion Detection and Prevention" on page 158

You can configure a policy for an application or a device.

To configure an application policy:

1. In the Application policy section of the configuration page, click the tab for the type of policy that you want
   - Applications—Configure a policy for any traffic, to or from the WAN Edge device.
   - Device In—(Available if you're on the device configuration page) For applications that require inbound ICMP traffic, configure a policy to allow ICMP traffic from a remote network to the WAN Edge device.
     a. Select the **Device In** tab, as shown below, and then click **Add Application Policy**.
     b. Under **Network / User** choose the remote network that is the source of the ICMP traffic, and then **ICMP** for the **Destination**.

     Figure 40: Policy to Allow Remote ICMP Traffic

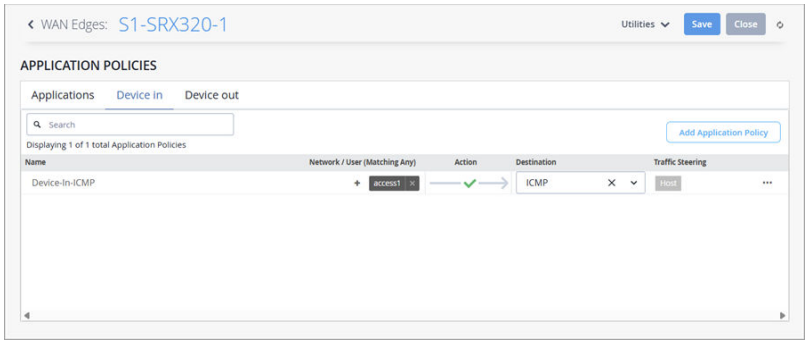

   - Device Out—Configure a policy for how traffic leaves the device. For example, you can set a traffic steering policy that states the device can only access the Mist cloud using a cellular network instead of your default link.

Figure 41: Policy to Traffic Out using a Cellular Network

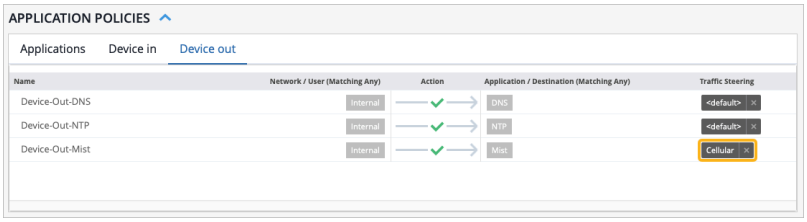

2. Click **Add Application Policy**.

   A policy appears in the policy list, with a name such as *Policy-1* or *Policy-2*.

3. Enter a name to replace the default name.
4. Work your way across the screen to select the networks and users, action (allow or block), applications, and other options.

   For more information, see "Settings for Application Policies" on page 151.

5. Click **Save** at the top-right corner of the configuration page.

Settings for Application Policies

Table 31: Settings for Application Policies

Field	Description
No.	*Not available in organization-level policies.* The position of the application policy. You can drag a policy up or down to change the number.
Name	Name of the policy. You can enter up to 32 letters, numbers, underscores, and dashes.
Network/User	Networks are the sources of the requests in your network. You can select one or more networks and users from the list.
Action	The action taken by this policy (allow or block traffic from source to destination). To change the action, click the icon, and then select from the list.

Table 31: Settings for Application Policies *(Continued)*

Field	Description
Application/Destination	The applications that this policy allows or blocks access to. You can select one or more applications from the list.

Table 31: Settings for Application Policies *(Continued)*

Field	Description

IDP	(Optional) Intrusion Detection and Prevention (IDP) profiles. If you have purchased the necessary IDP/URL filtering licenses and subscriptions, you can select one of the IDP profiles:
	• **None**—No IDP profile applied.
	• **Critical Only (SRX)**—The Critical-Only profile is suitable for critical-severity attacks. When the system detects a critical attack, this profile takes appropriate action. We recommend the Critical – Only SRX profile for SRX300 line of firewalls.
	• **Standard**—Standard profile is the default profile and represents the set of IDP signatures and rules recommended by Juniper Networks. The actions include:
	  • Close the client and server TCP connection.
	  • Drop current packet and all subsequent packets.
	• **Strict**—Strict profile contains a similar set of IDP signatures and rules as the standard profile does. However, when the system detects an attack, the strict profile actively blocks any malicious traffic or other attacks detected in the network.
	• **Alert**
	—Alert profile generates alerts only and does not take any additional action. Alert profiles are suitable only for low severity attacks. The IDP signature and rules are the same as in the standard profile.
	• **Recommended (SRX)**—Contains only the attack objects tagged as recommended by Juniper Networks. All rules have their Actions column set to take the recommended action for each attack object.

Table 31: Settings for Application Policies *(Continued)*

Field	Description
	• **Server Protection (SRX)**—Designed to protect servers. To be used on high memory devices with 2 GB or more of memory. • **Client Protection (SRX)**—Designed to protect clients. To be used on high memory devices with 2 GB or more of memory. • **Client and Server Protection (SRX)**—Designed to protect both clients and servers. To be used on high memory devices with 2 GB or more of memory. • **Custom**—Select one of your custom profiles, as listed in the drop-down menu. For help creating profiles, see "Configure Advanced Threat Prevention Features (SRX Only)" on page 197. The IDP profile applied in your application policy performs traffic inspection to detect and prevent intrusions on the allowed traffic.

Table 31: Settings for Application Policies *(Continued)*

Field	Description
Advanced Security Services	Enable secure AI-native Edge features in your application policy. Ensure that the required device-side licenses are in place. • Anti-virus—You can create an anti-virus configuration and associate it with an application policy. You can either select from a set of predefined configurations (Default, HTTP(S) Only, and No FTP), or create a custom anti-virus configuration. **NOTE**: When configuring antivirus profiles for Session Smart Routers, ensure your device is running version 6.3.5 or higher. • Secure Sockets Layer (SSL) Forward Proxy (SRX Only)—SSL forward proxy acts as an intermediary, performing SSL encryption and decryption between the client and the server. SSL forward proxy is a transparent proxy; that is, it performs SSL encryption and decryption between the client and the server, but neither the server nor the client can detect its presence. The following SSL Proxy profiles are available based on the cipher category: Weak, Medium, and Strong. • Security Intel (SRX Only)—Juniper's curated security intelligence feeds, SecIntel, provide dynamic and automatic updates to identify and block malicious domains, URLs, and IP addresses. In Juniper Mist, SecIntel profiles, enable you to block malicious and unwanted traffic such as Command and Control (C&C) communications, compromised IP addresses or IP subnets, and domains connected to malicious activities. • Anti-Malware (SRX Only)—Juniper Networks Anti-Malware is a security solution that uses cloud-sourced data to protect against advanced cybersecurity threats. This feature detects and blocks malware and unwanted files on the network before they reach an endpoint. In Juniper Mist, you

Table 31: Settings for Application Policies *(Continued)*

Field	Description
	can create anti-malware profiles for WAN Edge devices, detailing which files need cloud analysis and the steps to take when malware is detected. For help enabling these features, see "Anti-Virus Profiles for WAN Edge Devices" on page 193 and "Configure Advanced Threat Prevention Features (SRX Only)" on page 197.
Traffic Steering	*Not available in organization-level policies.* A traffic-steering profile defines the traffic path or paths for the policy. These profiles are required for deploying the policy to a WAN Edge spoke device or a hub device. The Traffic Steering field is not available for organization-level application policies. When you define an application policy directly inside a WAN Edge template or hub profile, you need to specify the order number and traffic-steering options. For more information, see "Traffic Steering Rules" on page 134.

Reordering and Deleting Application Policies

Reordering application policies allows you to move the policies around after they have been created.

Note the following about policy order for Session Smart Routers:

- New policies appear at the end of the policy list.
- Policy order is *not* important for Session Smart Routers. Session Smart Routers evaluate all policies simultaneously and apply the most specific matching policy to each session. The order of application policies in the policy list does not affect their evaluation or application.
- On SRX Series Firewalls, which function as Zone-Based Firewalls, the device evaluates polices in a top down manner. Ordering a more specific policy above a less specific policy will change which

policy matches. You need to position block policies at the top of the sequence to ensure that the matching traffic is blocked, rather than being evaluated by other rules.

You can organize your policies in any order according to your management needs without affecting their application.

Select a policy and use Up Arrow or Down Arrow to change the order. You can change the policy order anytime.

Figure 42: Changing Policy Order

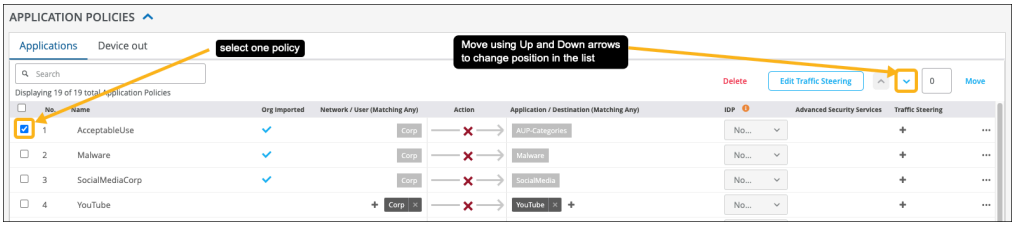

To delete an application policy, select the application policy you want to delete, and then click the **Delete** option that appears in red on the top right side of the pane.

Monitoring Breakout Paths (Beta)

You can monitor breakout paths with the Application Path Insights graph on the WAN Edge Insights dashboard.

NOTE: This feature is available to Beta participants only.

To improve your network monitoring experience with Session Smart Routers, Juniper Mist switches local breakout traffic from one path to another when the path doesn't meet the associated SLA requirements for the link latency, jitter, and loss parameters.

The SSR devices compare the SLA parameters (latency, jitter, and loss) for all the local breakout paths against the thresholds configured for these parameters for each application. Whenever a set threshold is breached (that is, a local breakout path fails to meet the associated SLA requirements), the traffic shifts to another path based on the traffic steering configuration. Any such shifts in traffic are displayed on the Application Path Insights graph on the WAN Edge Insights dashboard.

In the following example, you see a traffic shift from the ge-0/0/0 interface to the ge-0/0/5 interface due to a SLA threshold breach.

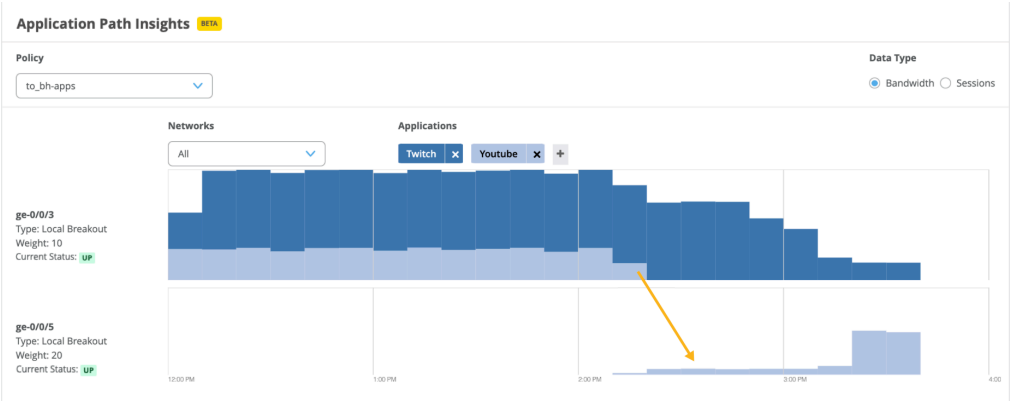

Intrusion Detection and Prevention

SUMMARY

Enhance network security by applying Intrusion Detection and Prevention (IDP) profiles.

IN THIS SECTION

Getting Started with IDP Profiles

Intrusion *detection* involves monitoring network events and identifying incidents and threats. Intrusion *prevention* follows up on the findings of the detection process by taking countermeasures. IDP profiles consist of many attack signatures, each with their own severity and recommended actions. When you apply a profile to an application policy, you enable that appropriate safeguards.

With the required licenses, IDP is available for all WAN Edge devices that are configured as spokes. IDP is also supported on SRX devices that are configured as hubs (SRX only).

NOTE:

- IDP is a calculation-heavy feature. You will likely see performance degradation on entry-level WAN Edge devices if you enable IDP in your policies.
- IDP requires licenses. For help, see IDP Basic Configuration (in the Junos OS Intrusion Detection and Prevention User Guide).

WARNING: When you activate the IDP feature for the first time on a spoke-device, we recommend you to plan it in a maintenance window. The start of the IDP engine and inclusion into the path from LAN to WAN (that is, service-chaining) might take a few minutes and might also interrupt ongoing communications.

Add an IDP Profile to an Application Policy

To enable intrusion detection and prevention, apply an IDP profile to an application policy.

Before You Begin: Configure your networks, templates, profiles, devices, and application policies. For help, see "WAN Assurance Configuration Overview" on page 75.

To add an IDP profile to an application policy:

1. Navigate to your WAN Edge template, hub profile, or standalone WAN Edge device.
2. In the **Applications Policies** section, click **Add Application Policy** (for a new policy), or scroll to an existing policy.

   NOTE: For help with a new application policy, see "Application Policies" on page 144.

3. In the **IDP** column, select an IDP profile.

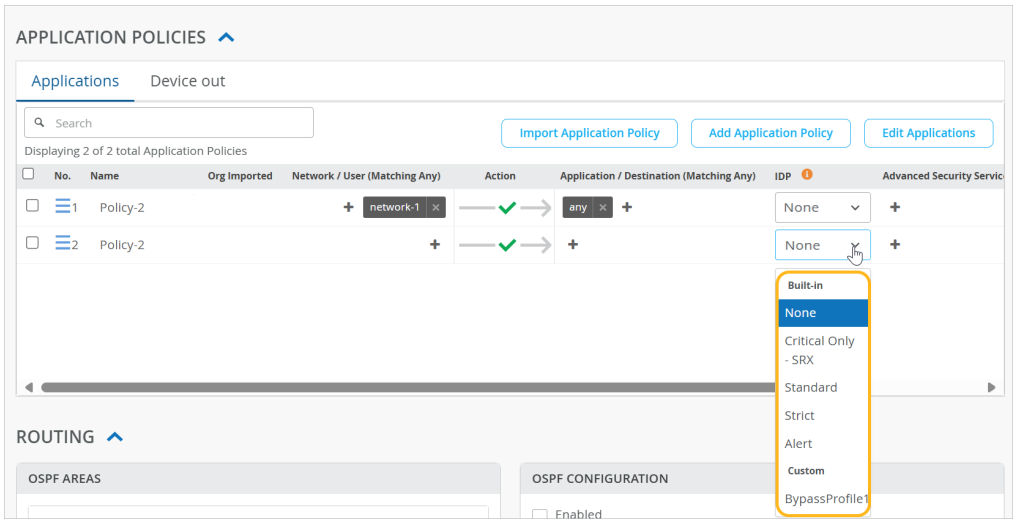

Options include:

Table 32: IDP Options

IDP Option	Description
None	No IDP profile applied.
Standard	Standard profile is the default profile and represents the set of IDP signatures and rules recommended by Juniper Networks. The actions include: • Close the client and server TCP connection. • Drop current packet and all subsequent packets.
Strict	Strict profile contains a similar set of IDP signatures and rules as the standard profile does. However, when the system detects an attack, the strict profile actively blocks any malicious traffic or other attacks detected in the network.
Alert	Alert profile generates alerts only and does not take any additional action. Alert profiles are suitable only for low severity attacks. The IDP signature and rules are the same as in the standard profile.
Critical Only - SRX	The Critical-Only profile is suitable for critical-severity attacks. When the system detects a critical attack, this profile takes appropriate action. We recommend the Critical – Only SRX profile for SRX300 line of firewalls.
Recommended - SRX	Contains only the attack objects tagged as recommended by Juniper Networks. All rules have their Actions column set to take the recommended action for each attack object.
Server-Protection - SRX	Designed to protect servers. To be used on high memory devices with 2 GB or more of memory.
Client-Protection - SRX	Designed to protect clients. To be used on high memory devices with 2 GB or more of memory.
Client-And-Server-Protection - SRX	Designed to protect both clients and servers. To be used on high memory devices with 2 GB or more of memory.
Custom	Select one of your custom profiles, as listed in the drop-down menu. For help creating profiles, see Configure Advanced Threat Prevention Features (SRX Only).

> **NOTE**: If you've created bypass profiles, they also appear in the Custom Profiles section of the IDP drop-down list. For help, see "Customize IDP with Bypass Profiles" on page 164(later in this topic).

4. Click **Save** at the top-right corner of the template.

The selected IDP profile is applied.

Running a Simulator to Test Your IDP Settings

You can test the effects of your IDP-based policies by launching sample attacks. You can use tools such as Nikto in Kali Linux, which has a variety of options available for security-penetration testing.

Before You Begin: Set up a virtual machine (VM) desktop (desktop1) in a sandbox or lab environment, and install a simple security scanner for web servers, such as Nikto. Nikto is an open-source webserver and web application scanner. For example, you can run Nikto against an unhardened Apache Tomcat webserver (or its equivalent) that is local to your lab. In this test, you can send plain or unencrypted HTTP requests for IDP inspection.

The following sample shows a process where you install the tool, check the presence of the HTTP server, and then launch the attacks.

```
virsh console desktop1
apt-get update
apt-get install -y nikto
# Using your individual Lab-Access-IP we test if the labinternal
# Apache Tomcat Server of the Apache guacamole container is avail
wget http://172.16.77.155:8080
-2022-09-16 15:47:32- http://172.16.77.155:8080/
Connecting to 172.16.77.155:8080… connected.
HTTP request sent, awaiting response… 200
Length: unspecified [text/html]
Saving to: ‘index.html’

index.html [ <=> ] 10.92K -.-KB/s in 0s

2022-09-16 15:47:32 (85.3 MB/s) - ‘index.html’ saved [11184]

# Now start our security scanner for the first time
nikto -h http://172.16.77.155:8080
- Nikto v2.1.5
---------------------------
+ Target IP: 172.16.77.155
```

```
+ Target Hostname: 172.16.77.155
+ Target Port: 8080
+ Start Time: 2022-09-16 15:48:22 (GMT0)
---------------------------
+ Server: No banner retrieved
+ The anti-clickjacking X-Frame-Options header is not present.
+ No CGI Directories found (use ‘-C all’ to force check all possible dirs)
+ Server leaks inodes via ETags, header found with file /favicon.ico, fields: 0xW/21630
0x1556961512000
+ OSVDB-39272: favicon.ico file identifies this server as: Apache Tomcat
+ Allowed HTTP Methods: GET, HEAD, POST, PUT, DELETE, OPTIONS
+ OSVDB-397: HTTP method (‘Allow’ Header): ‘PUT’ method could allow clients to save files on the
web server.
+ OSVDB-5646: HTTP method (‘Allow’ Header): ‘DELETE’ may allow clients to remove files on the
web server.
+ /examples/servlets/index.html: Apache Tomcat default JSP pages present.
+ Cookie JSESSIONID created without the httponly flag
+ OSVDB-3720: /examples/jsp/snp/snoop.jsp: Displays information about page retrievals, including
other users.
+ OSVDB-3233: /manager/manager-howto.html: Tomcat documentation found.
+ /manager/html: Default Tomcat Manager interface found
+ 6544 items checked: 1 error(s) and 10 item(s) reported on remote host
+ End Time: 2022-09-16 15:50:03 (GMT0) (101 seconds)
---------------------------
+ 1 host(s) tested
```

Run the security scanner. You'll notice that the scanner takes longer to run because it detects more errors and less events.

```
nikto -h http://172.16.77.155:8080
- Nikto v2.1.5
---------------------------
+ Target IP: 172.16.77.155
+ Target Hostname: 172.16.77.155
+ Target Port: 8080
+ Start Time: 2022-09-16 16:01:51 (GMT0)
---------------------------
+ Server: No banner retrieved
+ The anti-clickjacking X-Frame-Options header is not present.
+ No CGI Directories found (use ‘-C all’ to force check all possible dirs)
+ Server leaks inodes via ETags, header found with file /favicon.ico, fields: 0xW/21630
0x1556961512000
```

```
+ OSVDB-39272: favicon.ico file identifies this server as: Apache Tomcat
+ Allowed HTTP Methods: GET, HEAD, POST, PUT, DELETE, OPTIONS
+ OSVDB-397: HTTP method (‘Allow’ Header): ‘PUT’ method could allow clients to save files on the
web server.
+ OSVDB-5646: HTTP method (‘Allow’ Header): ‘DELETE’ may allow clients to remove files on the
web server.
+ /examples/servlets/index.html: Apache Tomcat default JSP pages present.
+ 6544 items checked: 5657 error(s) and 6 item(s) reported on remote host
+ End Time: 2022-09-16 16:05:27 (GMT0) (216 seconds)
---------------------------
+ 1 host(s) tested
```

Viewing Event Information

You can view the generated events by navigating to **Site** > **WAN** > **Security Events**.

On the Security Events page, you'll see all generated events, as shown in this example.

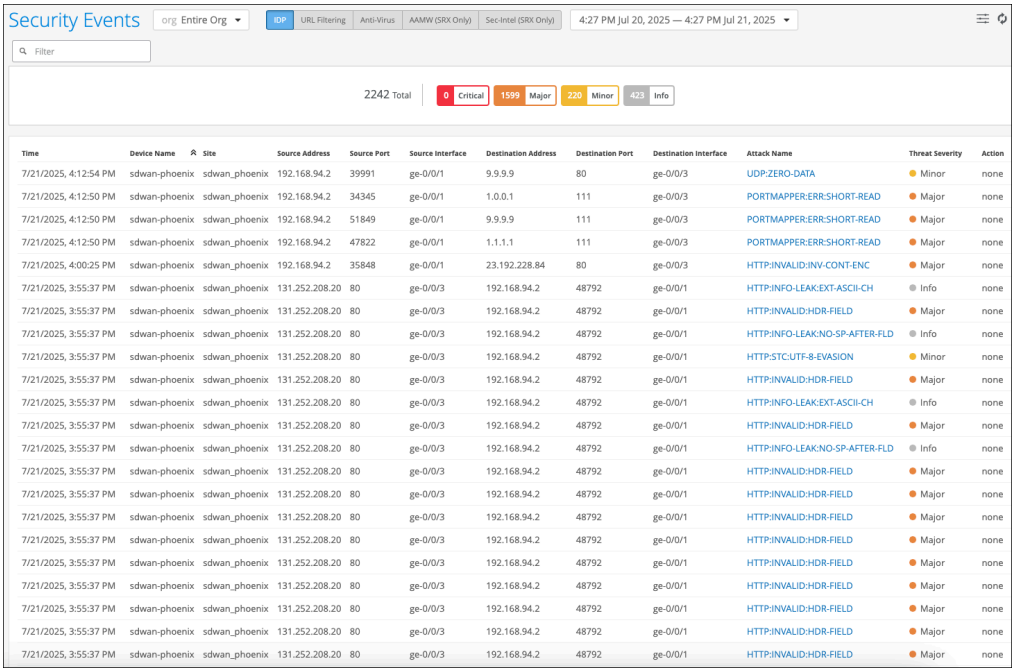

The Security Events page includes these helpful features:

- To filter the table—Above the table, click a button to filter the list by severity level, such as Critical or Minor.
- To view attack details—Click the hyperlink in the Attack Name column to see information about the attack and the default actions.

You can view more information about an event and the IDP actions by clicking the hyperlink.

In the previous example, you used passive logging for the events by using IDP profile type Alerts. Next, use IDP profile type Strict to stop or mitigate the events. When you use the Strict profile, the IDP engine closes TCP connections against the detected attacks.

You can follow the same process as shown in the sample. However, this time you change the spoke device template and change the IDP profile from **Alert** to **Strict**, as shown in this example.

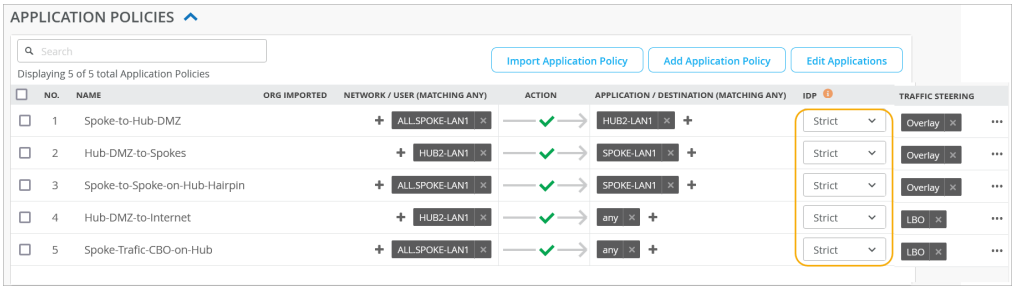

This example shows that for some events, the action is to close the session to mitigate the threats.

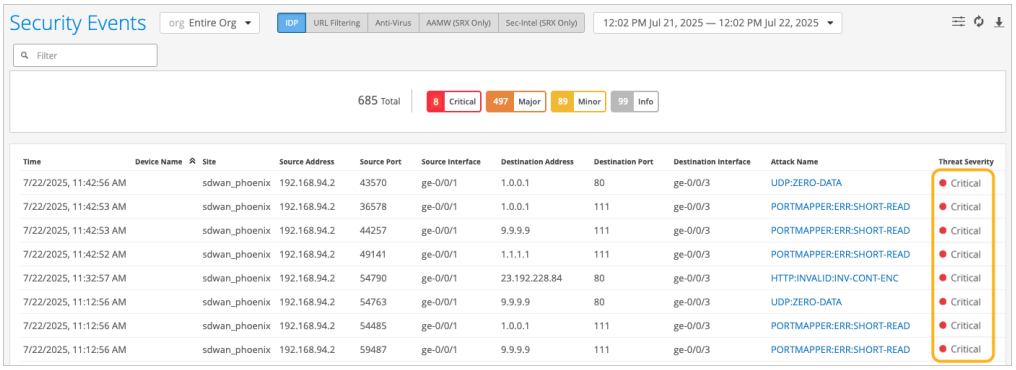

Customize IDP with Bypass Profiles

If you're seeing unnecessary alarms ("false positives"), you can create IDP bypass profiles as a counter-measure. For example, exclude a specific destination or attack type from IDP.

An IDP profile can have multiple bypass profiles, each with multiple bypass rules.

To create an IDP bypass profile:

1. From the left menu, select **Organization** > **WAN** > **Application Policy**.

2. In the Profiles section, click the IDP Bypass tab, and then click **Add Bypass Profile**.
3. In the Create Bypass Profile pop-up window, enter the information for this profile:

Table 33: Settings

Field	Description
Name	Enter a unique name for this profile. It can include letters, numbers, underscores, and dashes. It can contain up to 63 characters maximum.
Base Profile	You need a base IDP profile to create an IDP bypass profile. The supported types are: • **Standard**—Standard profile is the default profile and represents the set of IDP signatures and rules recommended by Juniper Networks. The actions include: Close the client and server TCP connection. Drop current packet and all subsequent packets. • **Strict**—Strict profile contains a similar set of IDP signatures and rules as the standard profile. However, when the system detects an attack, profile actively blocks any malicious traffic or other attacks detected in the network. • **Critical Only (SRX)**—The Critical-Only profile is suitable for critical-severity attacks. When the system detects a critical attack, this profile takes appropriate action. We recommend the Critical – Only SRX profile for SRX300 line of firewalls.

4. Click **Next** at the bottom of the Create Bypass Profile pop-up window.
The IDP Bypass Profiles page appears, showing your profile's name at the top of the page.

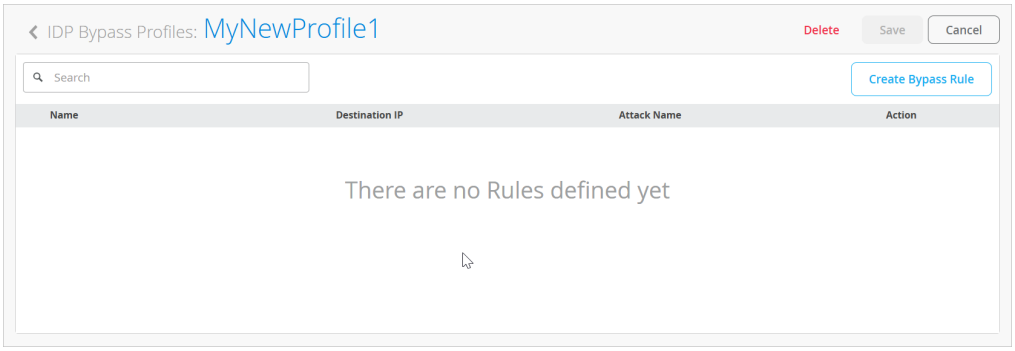

5. Click **Create Bypass Rule**.
6. Enter the settings for your rule in the side panel.

   Table 34: Settings

Field	Description
Name	Enter a descriptive name for this rule. Can contain letters, numbers, underscores, and dashes. Must start and end with a letter or number. Cannot exceed 32 characters.
Action	Select a traffic action: • **Alert**—Sends an alert about the event but doesn't discard • **Drop**—Discards packets without sending a response • **Close**—Discards packets and sends a response • **None**—Takes no action
Destination IP	IP address of the destination for traffic you want to exempt. You can select one or more destination IP addresses from the populated list or click **Add Destination IP**. Format must follow: <IP Address>/<Network Mask>. After adding an IP, click the check mark in the Add Destination IP title bar to save it.
Attack Name	Click **Add Attack Name** to specify the attacks to exempt for the specified destination addresses. The attack you enter must be of type supported by Juniper Networks IPS Signature. After adding an attack name, click the check mark in the Attack Name title bar to save it.

7. When you're done adding IP addresses and attacks, click **Add** at the bottom of the Create Bypass Rule window.
8. Repeat the above steps if you want to add more bypass rules to the profile.
9. When you're done configuring the profile, click **Save** at the top-right corner of the IDP Bypass Profiles page.
10. Apply the custom IDP profile to your application policies as needed.

For help, see "Add an IDP Profile to an Application Policy" on page 159 (earlier in this topic). Your custom profiles appear in the same IDP drop-down menu where you can select built-in profiles.

Internet Backhaul Through an SSR Hub

SUMMARY

This document describes how to use SSR Traffic Steering to originate and inject routes into an overlay network without relying on dynamic routing protocols. It applies to environments where dynamic routing is undesirable or unavailable.

The configuration uses an SSR hub providing Internet backhaul and demonstrates how the SSR can inject a default route into the overlay. To enable route injection, create an Application Policy that satisfies two requirements:

1. **Apply Traffic Steering toward a LAN or WAN interface.**

   Traffic Steering that targets only the overlay does not cause the SSR to originate routes.

2. **Apply the policy only to non-local networks.**

   Ensure that none of the networks in the policy are reachable from any LAN interface on the SSR.

When these conditions are met, the SSR performs protocol redistribution for the addresses defined in the policy. If the policy includes an application such as “any” (0.0.0.0/0), the SSR originates a default route.

In networks where the SSR should learn routes from external devices, avoid using Traffic Steering in application policies except when steering traffic to the overlay. In these situations, the SSR follows the routes in its routing table without requiring policy-driven injection.

Routing Policies

SUMMARY

Use this information to add a routing policy to a WAN Edge template, hub profile, or device configuration.

IN THIS SECTION

Navigating to Routing Policies

You'll find routing policies on the configuration page for your WAN Edge templates (for spokes), hub profiles, and individually managed devices.

- For a WAN Edge template—From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create a new one. Scroll down to the **Routing Policies** section.
- For a hub profile—From the left menu, select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create a new one. Scroll down to the **Routing Policies** section.
- For an individually managed WAN Edge device—From the left menu, select **WAN Edges** > **WAN Edges**. Click a device. Scroll down to the **Routing Policies** section.

Routing Policy Settings

SUMMARY

On the configuration page, the Routing Policy section lists the Routing Policies that you've defined. To add a Routing Policy click **Add Routing Policy**, enter the settings in the side panel, and then click **Add** at the bottom of the panel.

Table 35: Settings for the Routing Policy Panel

Field	Description
Name	Enter a name to identify this policy.

Table 35: Settings for the Routing Policy Panel *(Continued)*

Field	Description
Terms	This section of the policy lists all terms that you've defined. You can add multiple terms, as needed. You must add at least one term. Click an existing term to modify, or click **Add Terms**, enter the settings, and then click the check mark in the Add Term title bar. Settings for Terms include: • Prefix—Enter a comma-separated list of IP addresses, or an IP address range. Format: *x.x.x.x/y* for a single IP address or *x.x.x.x/y-z* for a range . • AS Path—Enter a regular expression or a number in the range 1-4294967294. • Protocol—Select a protocol from the list. • Community—Enter a regular expression or a number in the range 1-4294967294. • Then—Use the drop-down menu to select the action to apply when the condition is met. If you choose Accept, enter the Append Community and the Set Community. • Add Action—Click the link to add an action, then click the action in the list. • **Prepend AS Path**—Prepend a AS number to the start of a BGP AS path. • **Exclude AS Path**—Exclude a AS number from the start of a BGP AS path. • **Set Local Preference**—Set preference to assign to routes that are advertised to the group or peer.

Table 35: Settings for the Routing Policy Panel *(Continued)*

Field	Description
	• **Add Target VRs**— Add virtual Routing and Forwarding (VRF) instances for the intentional sharing of route information across VRF instances. Optionally, add overlay paths and conditions, as described in the remaining rows of this table.
Overlay Path Preference	This section of the term configuration lists all overlay paths that you've defined. Click an existing path to modify, or click **Add Paths**, enter the settings, and then you can add multiple paths, as needed. After you click **Add Paths**, select a path from the drop-down menu, and then click the check mark in the Add Path title bar. The menu includes all overlays that you've defined in the WAN section.
Conditions	This section of the term configuration lists all conditions that you've defined. Click an existing condition to modify, or click **Add Condition**, enter the settings, and then click the check mark in the Add Condition title bar. You can add multiple conditions, as needed. Settings for conditions include: • Prefix • Custom VR

OSPF

SUMMARY

Add OSPF to your WAN Edge template, hub profile, or standalone device configuration.

IN THIS SECTION

Configure OSPF in a WAN Edge Template

Open Shortest Path First (OSPF) is a link-state routing protocol used in IP networks to determine the best path for forwarding IP packets. OSPF divides a network into areas to improve scalability and control the flow of routing information. Follow these steps to configure OSPF for your WAN Edge device.

Before You Begin: Create your WAN Edge template, hub profile, or standalone device configuration. For help, see "WAN Assurance Configuration Overview" on page 75.

To configure OSPF:

1. Navigate to your WAN Edge template, hub profile, or standalone WAN Edge device.
2. In the **Routing** section, from the **OSPF Areas** section, click **Add OSPF Area**.
3. In the **Add OSPF Area** panel, add the following information:

   Table 36: Add OSPF Area Options

Field	Description
Area	This number indicates the identification area that your OSPF network or WAN Edge device belongs to.

Table 36: Add OSPF Area Options *(Continued)*

Field	Description
Type	This is the OSPF Area type. Select one of the following options: • Default (Area 0) — This represents the core of an OSPF network. • Stub — Using this OSPF area type blocks external routes. • Not So Stubby Area (NSSA) — Using this OSPF area type allows redistribution of some external routes and not others. For a more in depth explanation of the different area types, see Configuring OSPF Areas.

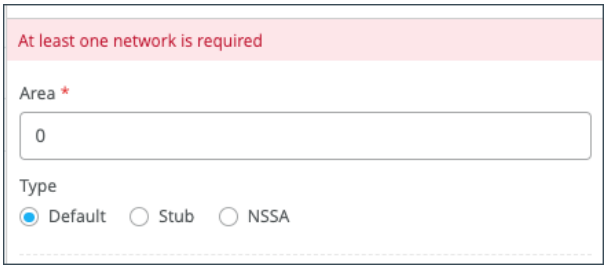

4. Click **Add OSPF Network**, then, in the **Add OSPF Network** section of the panel, enter the following information:

Table 37: Add OSPF Network Options

Field	Description
Network	This is the name of your OSPF network. **NOTE**: Select the **Passive** check box if you do not want OSPF to send Hello packets on an interface. This prevents the interface from forming unnecessary neighbor relationships, which reduces overhead on routers and ensures that only the crucial connections are being made.

Table 37: Add OSPF Network Options *(Continued)*

Field	Description
Interface Type	• Broadcast — The default interface type for an OSPF Ethernet interface. • p2p (point to point) —A connection between two OSPF routers. • p2mp (SRX Only)—A connection between multiple OSPF routers.
BFD Interval	This value determines how frequently (in milliseconds) BFD packets will be sent to BFD peer.
Metric	This is the cost metric used by OSPF to determine the best path between two OSPF-enabled devices.
Hello Interval	This interval specifies the length of time, in seconds, before the routing device sends a hello packet out an interface. By default, the routing device sends Hello packets every 10 seconds.
Dead Interval	This interval specifies the length of time, in seconds, that the routing device waits before declaring a neighboring routing device as unavailable. By default, the routing device waits 40 seconds (four times the Hello interval).
Auth Type	• None — Selecting this means you are selecting no authentication to be done. • md5 (message-digest algorithm) —A hashing algorithm that uses a one-way cryptographic function, which takes a message of any length and returns a fixed-length output value. • password—Requires a password for authentication.

Table 37: Add OSPF Network Options *(Continued)*

Field	Description
Export or Import *(SRX Only)*	Select an existing routing policy or click the **Create Policy** link below the drop-down menu. If you click the create button, then the Add Routing Policy panel appears. For help with policy settings, see "Routing Policies" on page 168.

5. Click the check mark at the top of the Add OSPF Network section.

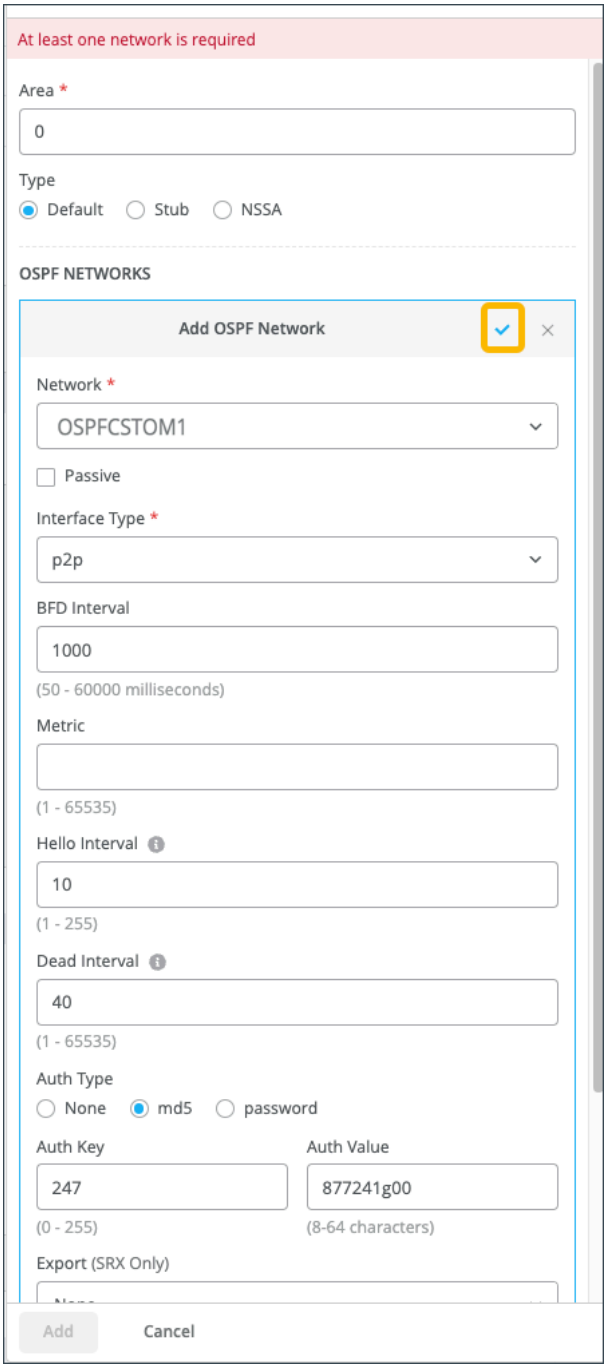

6. Click **Add** at the bottom of the panel.
   Your OSPF area appears in the **OSPF Areas** section.
7. In the **OSPF Configuration** section, select the **Enabled** check box.
   This causes the **Enable OSPF Areas** button to appear.
8. Click **Enable OSPF Areas**.

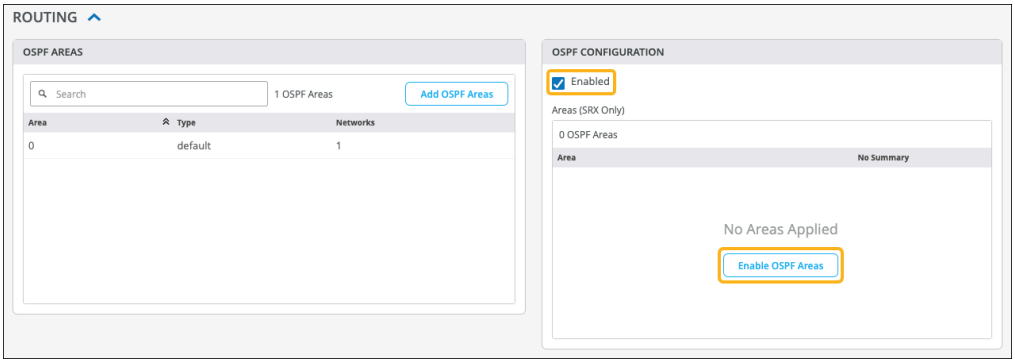

9. In the **Enable OSPF Area** panel, select the area that you just created.
10. (Optional) For SRX Series Firewalls, select the **No Summary** check box if you do not want your SRX to advertise summary routes.
11. Click **Add** at the bottom of the panel.

You will see your area listed in the **OSPF CONFIGURATION** section.

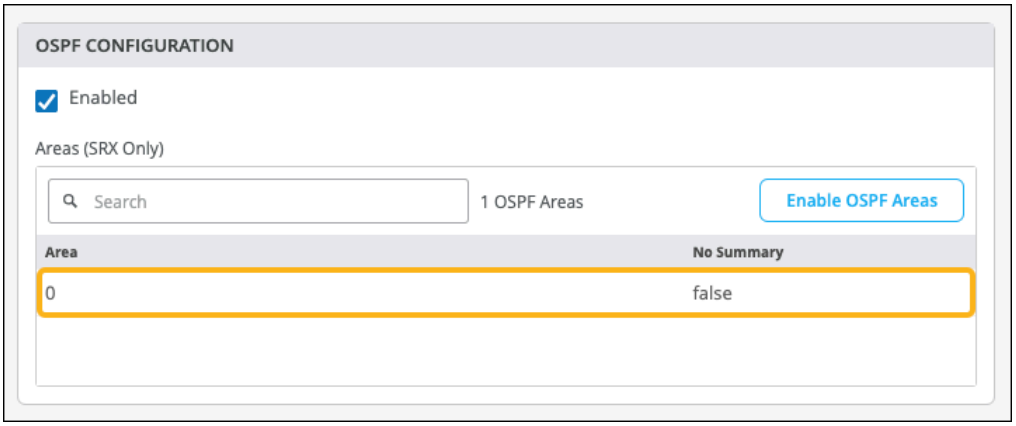

BGP

SUMMARY

When configuring WAN links, create BGP groups to integrate your WAN with your data center or data providers.

IN THIS SECTION

Before You Begin

- Create the networks that you want to reference in your routing configurations. For help, see "Network Settings" on page 94.
- Configure any Secure Edge Connectors that you want to reference in your routing configurations. For help, see "Juniper Mist Secure Edge Connector Overview" on page 222.
- Create your WAN Edge template, hub profile, or standalone device configuration. For help, see "WAN Assurance Configuration Overview" on page 75.

Create a BGP Group

The BGP Neighbors configuration is where you define the other BGP routers that your WAN Edge device is going to talk with. Follow these steps to configure a BGP (Border Gateway Protocol) group. Specify the policies and protocols, and add the neighbors to include in the group.

To create a BGP group:

1. Navigate to your WAN Edge template, hub profile, or standalone WAN Edge device.
2. In the **BGP** section, click **Add BGP Groups**.
3. In the Add BGP Group configuration panel, enter the settings.

   Refer to the following descriptions of the general settings (in the main Add BGP Group panel) and additional settings that are hidden until you click specific links and buttons.

Table 38: Settings for BGP Group—Major Fields

Field	Description
Name	Enter a unique, descriptive name to identify this group. You can enter letters, numbers, underscores, or dashes. The name must start and end with a letter or number. Maximum length: 32 characters.
Peering Network	Peering network is the network where you'll establish your BGP neighbor. Select the appropriate WAN, LAN, SEC tunnel. Tips: • For the SEC Tunnel option, you can use Secure Edge Connectors with Custom or Prisma Access as the provider. JSE and Zscaler do not support dynamic routing protocols like BGP. • If you're configuring BGP in a WAN Edge template or WAN Edge device configuration, you can select **Overlay** to specify a preferred path to use for the traffic traversing from a spoke device to the BGP-learned prefixes. If you select this option, also specify an Export Policy. For help with the policy settings, see "Routing Policies" on page 168.
Remove Private AS	Select this check box if you want the WAN Edge device to remove private AS numbers from an AS path. When a router sends route information to a BGP neighbor in a different AS, the private numbers will not appear. This setting is useful because some ISPs automatically reject routes that contain private AS numbers.

Table 38: Settings for BGP Group—Major Fields *(Continued)*

Field	Description
BFD	The Bidirectional Forwarding Detection (BFD) protocol is a simple hello mechanism that detects failures or faults between network forwarding elements that share a link. Hello packets are sent at a specified, regular interval. If it doesn't get a reply within the expected time period, it considers this as a neighbor failure. Enabling BFD can be useful because the BFD timers are shorter, providing faster failure detection than with the default settings for BGP.
Type	Select **Internal** unless you have an existing External BGP (EBGP) network; in that case, select **External**.
Local AS	Enter your WAN Edge device's AS number. It's the default autonomous system number for when the Session Smart Router advertises BGP. You can configure your Session Smart Router to be in one AS for one neighbor, and another AS for different neighbors.
Graceful Restart Time	Graceful restart allows a routing device to inform its adjacent neighbors when it is restarting. Enter a value from 1 to 495 seconds.
Authentication Key	For added security, you can add an MD5 password. Neighbors in this group will use this password to verify the authenticity of packets sent from this system.
Export *or* Import Policy	Select an existing routing policy or click the **Create Policy** link below the drop-down menu. If you click the create button, then the Add Routing Policy panel appears. For help with policy settings, see "Routing Policies" on page 168.

Table 38: Settings for BGP Group—Major Fields *(Continued)*

Field	Description
Neighbors	In the Neighbors section of the Add BGP panel, you manage the devices that are included in this BGP group. These are the other BGP routers the session smart router is going to talk with. You can edit a neighbor by clicking it. For new entries, click **Add Neighbor**. Enter the settings, and click the check mark at the top of the Neighbors section. For help, see the Table 39 on page 181 table.

Table 39: Settings for Neighbors

Field	Description
Enabled *or* Disabled	Administratively enable or disable a BGP neighbor.
IP Address	The IP address of the device, or a variable representing the address.
Neighbor AS	Enter the autonomous system number for this neighbor device.
Multihop TTL	Enter the number of hops that BGP packets can traverse to reach a remote BGP peer. This setting is necessary if the peers are separated by non-BGP routers. Valid entries: 0 (no multihopping) to 64
Export *or* Import Policy	Select a routing policy from the drop-down menu, or click the **Create Policy** link below the menu. If you click the create button, then the Add Routing Policy panel appears. For help with policy settings, see "Routing Policies" on page 168.

NOTE: After entering the required details for a neighbor, click the check mark in the Add Neighbor title bar. (The check mark appears only after the required information is entered).

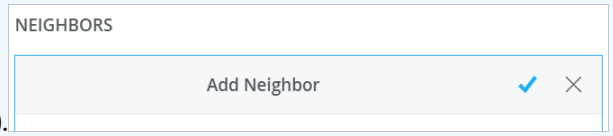

Static Routes (SRX Only)

SUMMARY

When configuring an SRX, as your WAN Edge device, you can add static routes to your hub, spoke, or standalone configuration.

IN THIS SECTION

Add a Static Route

1. Navigate to the hub profile, WAN Edge template, or device.
2. In the **Static Routes** section, click **Add Static Route**.
3. Select **Network** or **Prefix** to define this route.
   - **Network**—Define the route by selecting a network and entering the gateway IP address.
   - **Prefix**—Define the route by entering the prefix and the gateway IP address.
4. Click **Add**.
5. Click **Save** at the top-right corner of the configuration page.

5

CHAPTER

Threat Detection and Prevention

IN THIS CHAPTER

Intrusion Detection and Prevention

SUMMARY

Enhance network security by applying Intrusion Detection and Prevention (IDP) profiles.

IN THIS SECTION

Getting Started with IDP Profiles

Intrusion *detection* involves monitoring network events and identifying incidents and threats. Intrusion *prevention* follows up on the findings of the detection process by taking countermeasures. IDP profiles consist of many attack signatures, each with their own severity and recommended actions. When you apply a profile to an application policy, you enable that appropriate safeguards.

With the required licenses, IDP is available for all WAN Edge devices that are configured as spokes. IDP is also supported on SRX devices that are configured as hubs (SRX only).

NOTE:

- IDP is a calculation-heavy feature. You will likely see performance degradation on entry-level WAN Edge devices if you enable IDP in your policies.
- IDP requires licenses. For help, see IDP Basic Configuration (in the Junos OS Intrusion Detection and Prevention User Guide).

WARNING: When you activate the IDP feature for the first time on a spoke-device, we recommend you to plan it in a maintenance window. The start of the IDP engine and inclusion into the path from LAN to WAN (that is, service-chaining) might take a few minutes and might also interrupt ongoing communications.

Add an IDP Profile to an Application Policy

To enable intrusion detection and prevention, apply an IDP profile to an application policy.

Before You Begin: Configure your networks, templates, profiles, devices, and application policies. For help, see "WAN Assurance Configuration Overview" on page 75.

To add an IDP profile to an application policy:

1. Navigate to your WAN Edge template, hub profile, or standalone WAN Edge device.
2. In the **Applications Policies** section, click **Add Application Policy** (for a new policy), or scroll to an existing policy.

   NOTE: For help with a new application policy, see "Application Policies" on page 144.

3. In the **IDP** column, select an IDP profile.

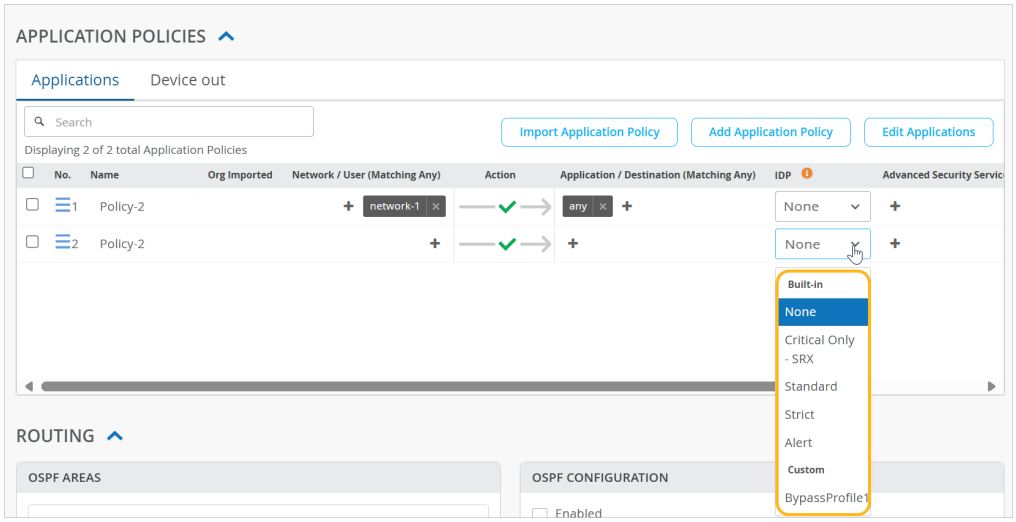

Options include:

Table 40: IDP Options

IDP Option	Description
None	No IDP profile applied.

Table 40: IDP Options *(Continued)*

IDP Option	Description
Standard	Standard profile is the default profile and represents the set of IDP signatures and rules recommended by Juniper Networks. The actions include: • Close the client and server TCP connection. • Drop current packet and all subsequent packets.
Strict	Strict profile contains a similar set of IDP signatures and rules as the standard profile does. However, when the system detects an attack, the strict profile actively blocks any malicious traffic or other attacks detected in the network.
Alert	Alert profile generates alerts only and does not take any additional action. Alert profiles are suitable only for low severity attacks. The IDP signature and rules are the same as in the standard profile.
Critical Only - SRX	The Critical-Only profile is suitable for critical-severity attacks. When the system detects a critical attack, this profile takes appropriate action. We recommend the Critical – Only SRX profile for SRX300 line of firewalls.
Recommended - SRX	Contains only the attack objects tagged as recommended by Juniper Networks. All rules have their Actions column set to take the recommended action for each attack object.
Server-Protection - SRX	Designed to protect servers. To be used on high memory devices with 2 GB or more of memory.
Client-Protection - SRX	Designed to protect clients. To be used on high memory devices with 2 GB or more of memory.
Client-And-Server-Protection - SRX	Designed to protect both clients and servers. To be used on high memory devices with 2 GB or more of memory.
Custom	Select one of your custom profiles, as listed in the drop-down menu. For help creating profiles, see Configure Advanced Threat Prevention Features (SRX Only).

NOTE: If you've created bypass profiles, they also appear in the Custom Profiles section of the IDP drop-down list. For help, see "Customize IDP with Bypass Profiles" on page 164(later in this topic).

4. Click **Save** at the top-right corner of the template.

The selected IDP profile is applied.

Running a Simulator to Test Your IDP Settings

You can test the effects of your IDP-based policies by launching sample attacks. You can use tools such as Nikto in Kali Linux, which has a variety of options available for security-penetration testing.

Before You Begin: Set up a virtual machine (VM) desktop (desktop1) in a sandbox or lab environment, and install a simple security scanner for web servers, such as Nikto. Nikto is an open-source webserver and web application scanner. For example, you can run Nikto against an unhardened Apache Tomcat webserver (or its equivalent) that is local to your lab. In this test, you can send plain or unencrypted HTTP requests for IDP inspection.

The following sample shows a process where you install the tool, check the presence of the HTTP server, and then launch the attacks.

```
virsh console desktop1
apt-get update
apt-get install -y nikto
# Using your individual Lab-Access-IP we test if the labinternal
# Apache Tomcat Server of the Apache guacamole container is avail
wget http://172.16.77.155:8080
-2022-09-16 15:47:32- http://172.16.77.155:8080/
Connecting to 172.16.77.155:8080… connected.
HTTP request sent, awaiting response… 200
Length: unspecified [text/html]
Saving to: ‘index.html’

index.html [ <=> ] 10.92K -.-KB/s in 0s

2022-09-16 15:47:32 (85.3 MB/s) - ‘index.html’ saved [11184]

# Now start our security scanner for the first time
nikto -h http://172.16.77.155:8080
- Nikto v2.1.5
---------------------------
+ Target IP: 172.16.77.155
+ Target Hostname: 172.16.77.155
+ Target Port: 8080
```

```
+ Start Time: 2022-09-16 15:48:22 (GMT0)
---------------------------
+ Server: No banner retrieved
+ The anti-clickjacking X-Frame-Options header is not present.
+ No CGI Directories found (use ‘-C all’ to force check all possible dirs)
+ Server leaks inodes via ETags, header found with file /favicon.ico, fields: 0xW/21630
0x1556961512000
+ OSVDB-39272: favicon.ico file identifies this server as: Apache Tomcat
+ Allowed HTTP Methods: GET, HEAD, POST, PUT, DELETE, OPTIONS
+ OSVDB-397: HTTP method (‘Allow’ Header): ‘PUT’ method could allow clients to save files on the
web server.
+ OSVDB-5646: HTTP method (‘Allow’ Header): ‘DELETE’ may allow clients to remove files on the
web server.
+ /examples/servlets/index.html: Apache Tomcat default JSP pages present.
+ Cookie JSESSIONID created without the httponly flag
+ OSVDB-3720: /examples/jsp/snp/snoop.jsp: Displays information about page retrievals, including
other users.
+ OSVDB-3233: /manager/manager-howto.html: Tomcat documentation found.
+ /manager/html: Default Tomcat Manager interface found
+ 6544 items checked: 1 error(s) and 10 item(s) reported on remote host
+ End Time: 2022-09-16 15:50:03 (GMT0) (101 seconds)
---------------------------
+ 1 host(s) tested
```

Run the security scanner. You'll notice that the scanner takes longer to run because it detects more errors and less events.

```
nikto -h http://172.16.77.155:8080
- Nikto v2.1.5
---------------------------
+ Target IP: 172.16.77.155
+ Target Hostname: 172.16.77.155
+ Target Port: 8080
+ Start Time: 2022-09-16 16:01:51 (GMT0)
---------------------------
+ Server: No banner retrieved
+ The anti-clickjacking X-Frame-Options header is not present.
+ No CGI Directories found (use ‘-C all’ to force check all possible dirs)
+ Server leaks inodes via ETags, header found with file /favicon.ico, fields: 0xW/21630
0x1556961512000
+ OSVDB-39272: favicon.ico file identifies this server as: Apache Tomcat
+ Allowed HTTP Methods: GET, HEAD, POST, PUT, DELETE, OPTIONS
```

```
+ OSVDB-397: HTTP method (‘Allow’ Header): ‘PUT’ method could allow clients to save files on the
web server.
+ OSVDB-5646: HTTP method (‘Allow’ Header): ‘DELETE’ may allow clients to remove files on the
web server.
+ /examples/servlets/index.html: Apache Tomcat default JSP pages present.
+ 6544 items checked: 5657 error(s) and 6 item(s) reported on remote host
+ End Time: 2022-09-16 16:05:27 (GMT0) (216 seconds)
---------------------------
+ 1 host(s) tested
```

Viewing Event Information

You can view the generated events by navigating to **Site** > **WAN** > **Security Events**.

On the Security Events page, you'll see all generated events, as shown in this example.

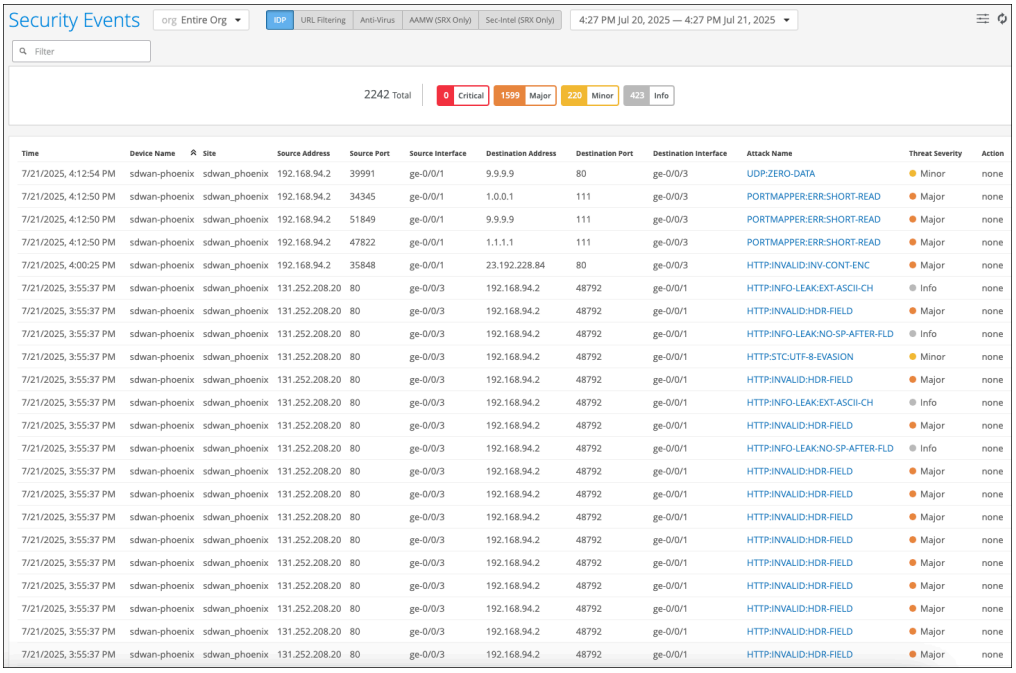

The Security Events page includes these helpful features:

- To filter the table—Above the table, click a button to filter the list by severity level, such as Critical or Minor.
- To view attack details—Click the hyperlink in the Attack Name column to see information about the attack and the default actions.

You can view more information about an event and the IDP actions by clicking the hyperlink.

In the previous example, you used passive logging for the events by using IDP profile type Alerts. Next, use IDP profile type Strict to stop or mitigate the events. When you use the Strict profile, the IDP engine closes TCP connections against the detected attacks.

You can follow the same process as shown in the sample. However, this time you change the spoke device template and change the IDP profile from **Alert** to **Strict**, as shown in this example.

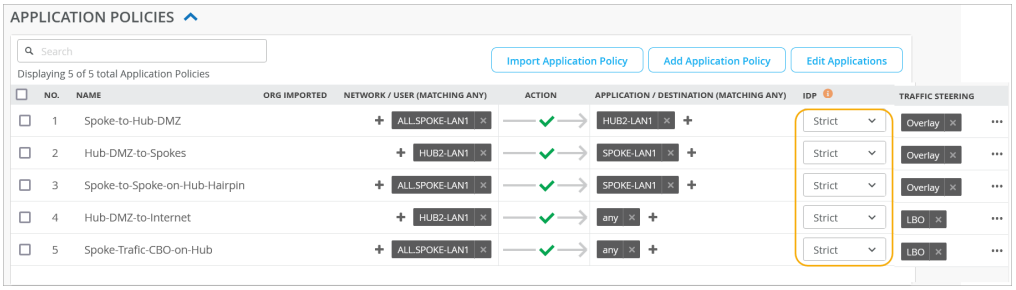

This example shows that for some events, the action is to close the session to mitigate the threats.

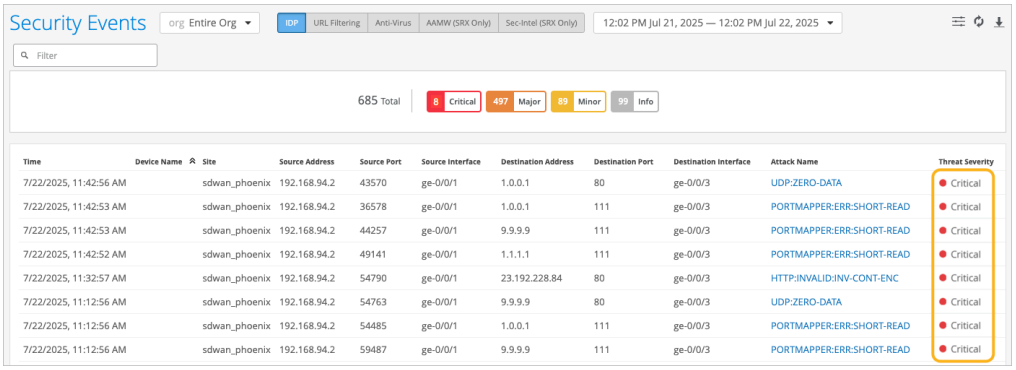

Customize IDP with Bypass Profiles

If you're seeing unnecessary alarms ("false positives"), you can create IDP bypass profiles as a counter-measure. For example, exclude a specific destination or attack type from IDP.

An IDP profile can have multiple bypass profiles, each with multiple bypass rules.

To create an IDP bypass profile:

1. From the left menu, select **Organization** > **WAN** > **Application Policy**.
2. In the Profiles section, click the IDP Bypass tab, and then click **Add Bypass Profile**.
3. In the Create Bypass Profile pop-up window, enter the information for this profile:

   Table 41: Settings

Field	Description
Name	Enter a unique name for this profile. It can include letters, numbers, underscores, and dashes. It can contain up to 63 characters maximum.
Base Profile	You need a base IDP profile to create an IDP bypass profile. The supported types are: • **Standard**—Standard profile is the default profile and represents the set of IDP signatures and rules recommended by Juniper Networks. The actions include: Close the client and server TCP connection. Drop current packet and all subsequent packets. • **Strict**—Strict profile contains a similar set of IDP signatures and rules as the standard profile. However, when the system detects an attack, profile actively blocks any malicious traffic or other attacks detected in the network. • **Critical Only (SRX)**—The Critical-Only profile is suitable for critical-severity attacks. When the system detects a critical attack, this profile takes appropriate action. We recommend the Critical – Only SRX profile for SRX300 line of firewalls.

4. Click **Next** at the bottom of the Create Bypass Profile pop-up window.

The IDP Bypass Profiles page appears, showing your profile's name at the top of the page.

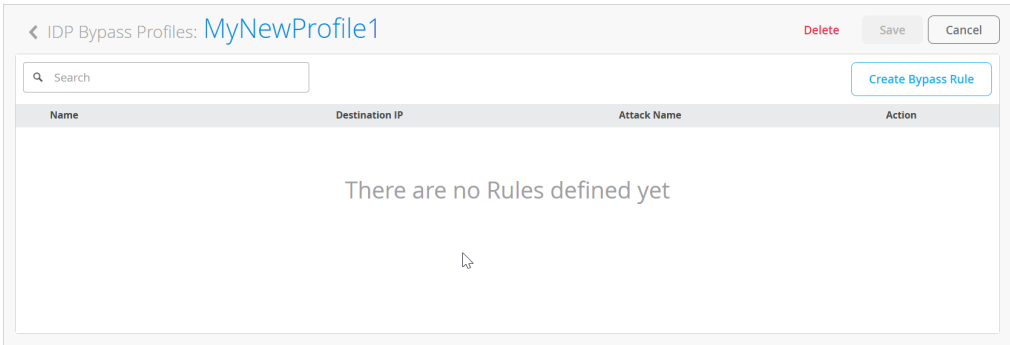

5. Click **Create Bypass Rule**.
6. Enter the settings for your rule in the side panel.

Table 42: Settings

Field	Description
Name	Enter a descriptive name for this rule. Can contain letters, numbers, underscores, and dashes. Must start and end with a letter or number. Cannot exceed 32 characters.
Action	Select a traffic action: • **Alert**—Sends an alert about the event but doesn't discard • **Drop**—Discards packets without sending a response • **Close**—Discards packets and sends a response • **None**—Takes no action
Destination IP	IP address of the destination for traffic you want to exempt. You can select one or more destination IP addresses from the populated list or click **Add Destination IP**. Format must follow: <IP Address>/<Network Mask>. After adding an IP, click the check mark in the Add Destination IP title bar to save it.

Table 42: Settings *(Continued)*

Field	Description
Attack Name	Click **Add Attack Name** to specify the attacks to exempt for the specified destination addresses. The attack you enter must be of type supported by Juniper Networks IPS Signature. After adding an attack name, click the check mark in the Attack Name title bar to save it.

7. When you're done adding IP addresses and attacks, click **Add** at the bottom of the Create Bypass Rule window.
8. Repeat the above steps if you want to add more bypass rules to the profile.
9. When you're done configuring the profile, click **Save** at the top-right corner of the IDP Bypass Profiles page.
10. Apply the custom IDP profile to your application policies as needed.

    For help, see "Add an IDP Profile to an Application Policy" on page 159 (earlier in this topic). Your custom profiles appear in the same IDP drop-down menu where you can select built-in profiles.

Anti-Virus Profiles for WAN Edge Devices

SUMMARY

Read this topic to understand how to create anti-virus profiles and apply them in application policies on WAN Edge devices.

IN THIS SECTION

When you create an anti-virus profile, you enable Juniper Mist™ to inspect files for known malicious content. You can create different profiles to define different types of content to scan and different actions to take.

To implement a profile, you assign it to an application policy. By doing so, you integrate inline malware scanning directly into your traffic control rules. With this approach, you gain effective protection against viruses and other malicious content.

This feature requires relevant anti-virus license on the WAN Edge device.

When configuring antivirus profiles for Session Smart Routers, ensure your device is running version 6.3.5 or higher.

Create an Anti-Virus Profile

Before You Begin: Create your application policies. For help, see "Application Policies" on page 144.
To create an anti-virus profile:

1. From the left menu, select **Organization** > **WAN** > **Application Policy**.
2. Under **Profiles**, click the **Anti-Virus** tab. The page displays any anti-virus profiles defined (if available).
3. Click **Add Anti-Virus Profile** and enter the following details:

   Table 43: How to Configure an Anti-Virus Profile

Field	How to Configure
Name	Enter a name for the anti-virus profile.
Max. File Size	Enter the content size limit in kilobytes (KB). The range is 20 through 40,000 KB. The content size limit check occurs before the scan request is sent. The content size refers to accumulated TCP payload size.
Protocols	Select one more more protocols to include in this anti-virus profile.
URL White List	Enter a list of trusted websites or URLs to exclude from anti-virus scans.
Mime White List	Enter a list of specific file types, identified by their MIME headers, to exclude from anti-virus scans. Example: image/gif, audio/mp3, video/avi, application/zip, application/pdf, and so on.

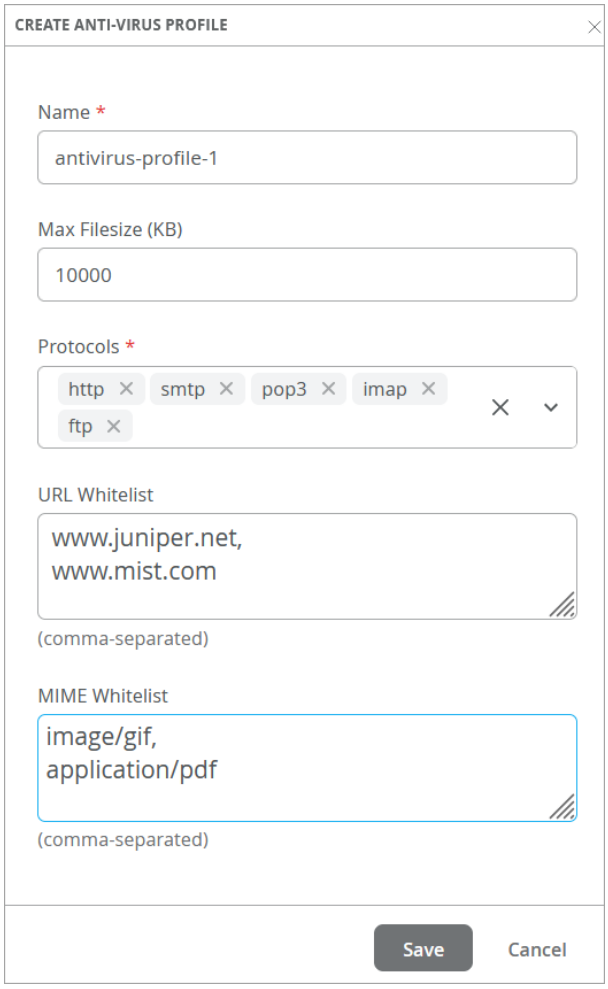

4. In the list of application policies, find the one that you want to apply your anti-virus policy to.

   For help creating an application policy, see "Application Policies" on page 144.

5. In the **IDP** column, select an anti-virus policy.

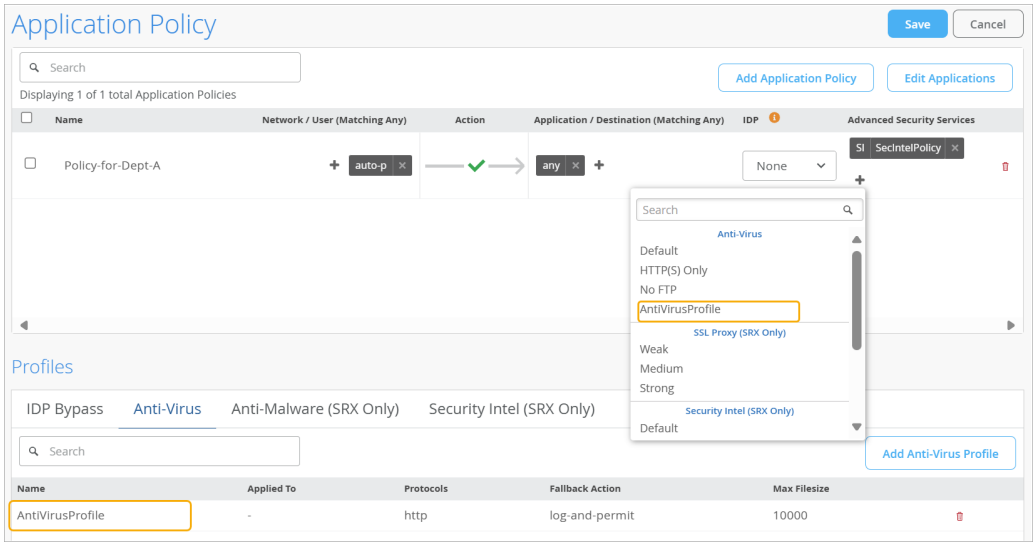

6. Optionally, also select available profiles:
   - Default—Scans files sent across HTTP, FTP, SMTP, POP3, and IMAP protocols.
   - HTTP(S)-only—Scans files sent across HTTP or HTTPS.
   - No-FTP—Excludes files sent across FTP from anti-virus scanning.
7. Save the configuration changes.

View WAN Edge Device Status

In the Juniper Mist portal, select **WAN Edges** > **WAN Edges** to view basic device monitoring information.

The **Advanced Security** section, located below the device ports, shows the status of security services. A green check mark indicates that the service is active on the device.

Figure 43: Advanced Security Status Details

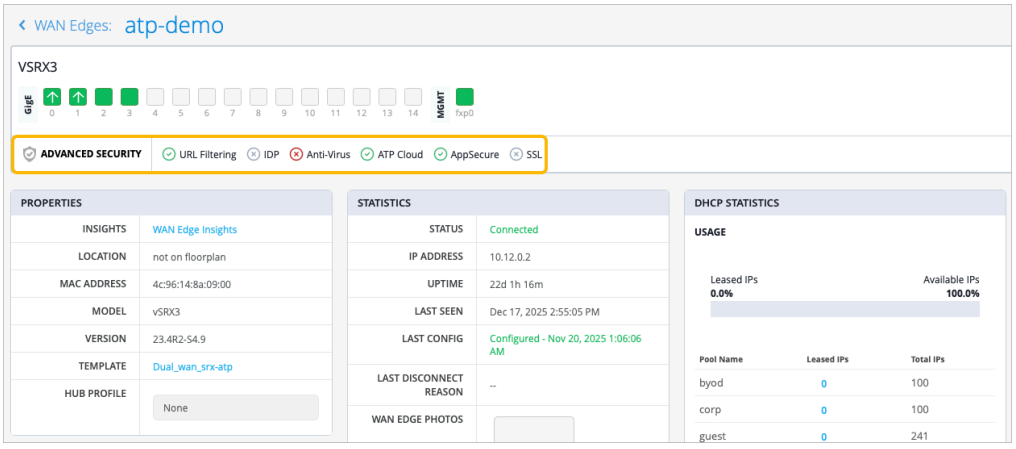

Below the Advanced Security section, you'll find **Properties** section that contains generalized platform-related information.

Click **WAN Edge Events** or navigate through **Monitor** > **Insights** and select the site and the WAN Edge that you want to view.

Click an event to see a summary on the right side of the page.

Configure Advanced Threat Prevention Features (SRX Only)

SUMMARY

Create Security Intel (SecIntel) profiles and advanced anti-malware profiles and apply them in application policies on WAN Edge devices.

IN THIS SECTION

Juniper Networks' Advanced Threat Prevention (ATP) for SRX Series Firewalls offers a comprehensive suite of features designed to detect, analyze, and prevent advanced cyber threats. Juniper Mist supports the following features:

- SecIntel Threat Intelligence Feeds—Juniper's curated security intelligence feeds, SecIntel, provide dynamic and automatic updates to identify and block malicious domains, URLs, and IP addresses. In Juniper Mist, SecIntel profiles enable you to block malicious and unwanted traffic such as Command and Control (C&C) communications, compromised IP addresses or IP subnets, and domains connected to malicious activities.
- Advanced Anti-Malware (AAMW)—Juniper Networks Anti-Malware is a security solution that uses cloud-sourced data to protect against advanced cybersecurity threats. This feature detects and blocks malware and unwanted files on the network before they reach an endpoint. In Juniper Mist, you can create anti-malware profiles for WAN Edge devices, detailing which files need cloud analysis and the steps to take when malware is detected.
- Third-party threat feeds—Sky ATP allows you to enable additional threat intelligence feeds (known malicious IPs, domains and URLs) from external vendors. Once enabled, these feeds will be a part of the Security Intel CC Category on the device. Supported feed types include:
  - IP-based—Threatfox IP, Feodo Tracker, DShield, Tor, Blocklist

  - URL-based—Threatfox URL, URLHaus, OpenPhish
  - Domain-based—Threatfox Domain
- SecIntel custom allowlist and blocklist—You can define customs IP addresses and domains under two categories - Allowlist and Blocklist. These user-defined IP addresses or domains are included under the Security Intel CC Category on the device.

The Advanced Threat Prevention feature is supported on SRX Series Firewalls.

Before You Begin

Ensure you have following available:

- Juniper Advanced Threat Prevention Cloud account and an ATP Cloud Realm created from your account. See Enroll an SRX Series Firewall Using Juniper ATP Cloud Web Portal.
- Your ATP Cloud account associated with a license. For more information, see Software Licenses for ATP Cloud.

Enrollment of a WAN Edge device in ATP Cloud occurs once a realm is created and either a SecIntel or an AAMW profile is associated with a security policy for that device.

Add Your ATP Credential Details to Your Organization and Specify the Features

Juniper Mist automatically enrolls devices in Cloud ATP Services as required. To integrate Juniper Mist Cloud with ATP Cloud, you need to provide ATP credential details in the Juniper Mist portal.

1. From the left menu of the Juniper Mist portal, select **Organization** > **Admin** > **Settings**.
2. Scroll-down to the **Secure WAN Edge Integration** section.
3. Add the credentials for your ATP Cloud account:
   a. Click **Add Credentials**.
   b. In the **Add Credentials** pop-up window, enter the details.
      - **Provider**—Select **ATP Cloud**.
      - **Email Address**—Enter the username for your ATP account.
      - **Password**—Enter the password for your ATP account.

- **Realm**—Enter the realm name for your ATP account.

> **NOTE**: Only the Global instance of ATP is supported.

c. Click **Save** at the bottom of the Add Credentials pop-up window.

Your new provider now appears in the providers list.

4. Set up the threat feeds and domains for the newly added provider:

a. In the providers list, click the pencil button for the provider that you just added.

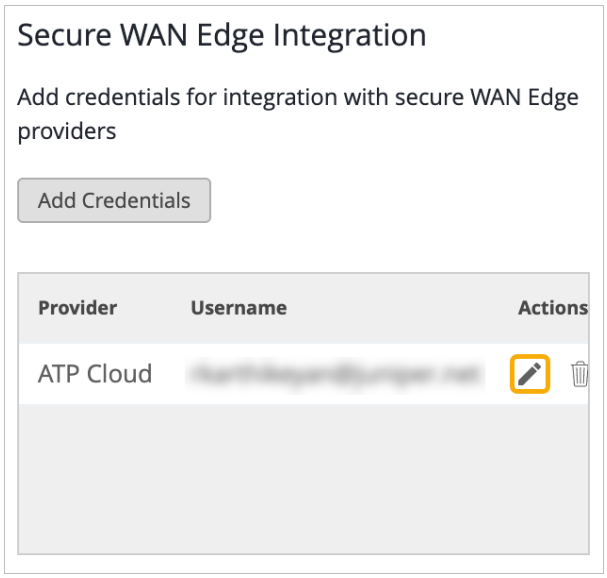

b. In the Configure Security Intelligence pop-up window, select the check box for each feature to enable.

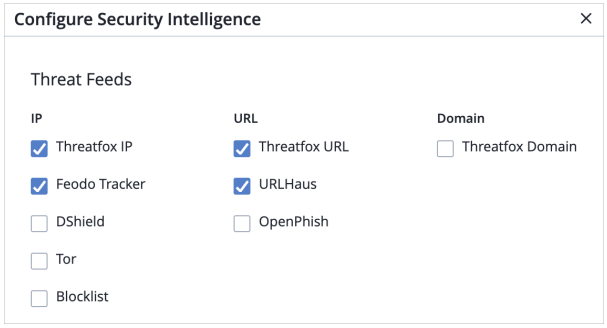

c. In the C&C IP/Domain section of the window, click **Add C&C IP/Domain**, enter the details to identify the domain, and then click the check mark in the C&C IP/Domain title bar.

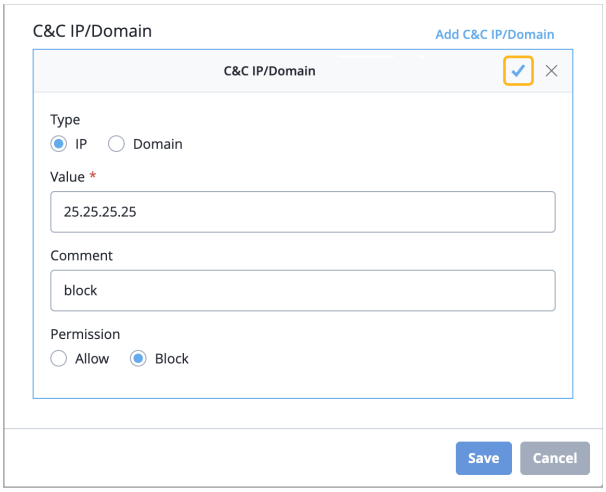

d. After adding all needed domains, click **Save** at the bottom of the Configure Security Intelligence window.

5. Click **Save** at the top of the Organization Settings page.

You'll need to reboot individual or clustered devices enrolled in Cloud ATP Services in order to activate Enhanced Services Mode. This increases the maximum number of services for L7 service processing.

Create Security Intelligence (SecIntel) Profiles

SecIntel offers meticulously curated and verified threat intelligence sourced from Juniper Networks' Advanced Threat Prevention (ATP) Cloud. This intelligence is delivered to WAN Edge device for effectively blocking Command and Control (C&C) communications at line rate. By enabling automatic and responsive traffic filtering, SecIntel provides real-time threat intelligence.

Many of the feeds include an associated threat score, allowing customers to define security rules and controls that are applied to traffic passing through their devices. The SecIntel security service integrates Juniper threat feeds, including those for C&C communications, malicious domains, and infected hosts. See also: SecIntel Feeds Overview and Benefits.

SecIntel profiles, which can be incorporated into application policies, enable the blocking of malicious and unwanted traffic such as C&C communications, compromised IP addresses or subnets, and domains linked to malicious activities.

To create a SecIntel profile:

1. From the left menu, select **Organization** > **WAN** > **Application Policy**.
2. Scroll down to the Profiles section.
3. Click the **Security Intel (SRX Only)** tab.
4. Click **Add Security Intel Profile**.

5. In the Create Security Intel Profile pop-up window, enter the details for this profile.

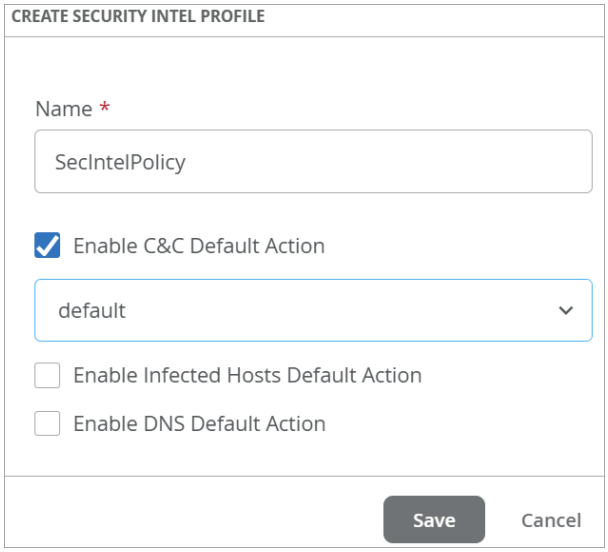

- Name—Enter a name to identify this profile.
- Check boxes—Select the actions for this profile.
  - Enable C&C Default Action—Actions against C&C servers that have attempted to contact and compromise hosts on your network.
  - Enable Infected Host Default Action—Actions against infected hosts, which are local devices that are potentially compromised because they appear to be part of a C&C network or exhibit other symptoms.
  - Enable DNS Default Action—Actions against the domains that are known to be associated with malicious activities.

  After selecting a check box, also use the drop-down menu to select the response level.
  - Default—The least aggressive response. Monitors and logs events with a threat score of 1-8; blocks events scoring 9-10.
  - Standard—The mid-level response. Monitors and logs events with a threat score of 1-5; blocks events scoring 6-10.
  - Strict—The most aggressive response. Monitors and logs events with a threat score of 1-2; blocks events scoring 3-10.

6. Click **Save** at the bottom of the Create Security Intel Profile window.
   The profile you created appears under **Security Intel (SRX Only)** pane.
7. Repeat the above steps as needed to create additional profiles.

You can now apply your new profile(s) to your application policies.

Create Advanced Anti-Malware Profiles

This feature detects and blocks malware and unwanted files on the network before they reach an endpoint. Like SecIntel, anti-malware profiles can be created from the application policy screen and included in an application policy.

To create an Anti-Malware profile:

1. (Skip this step if you're continuing from the previous procedure.) From the left menu, select **Organization** > **WAN** > **Application Policy**.
2. In the Profiles section, click the **Anti-Malware (SRX Only)** tab.
3. Click **Add Anti-Malware Profile**, and then enter the details.

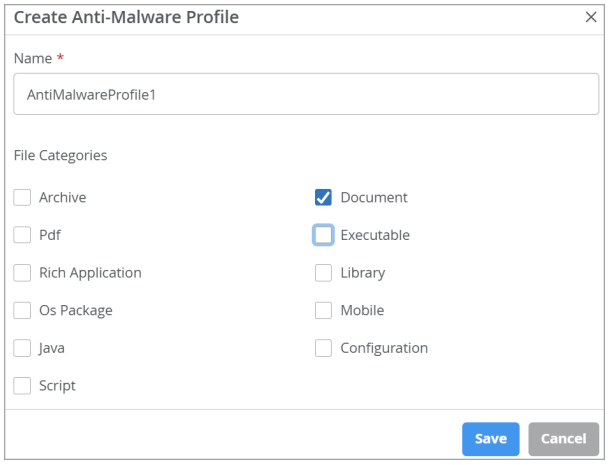

   a. Enter a name for the profile.

   b. Select one or more file categories as provided in the table below:
   Table 44: File Category Contents

Category	Description	File Types
Archive	Archive files	.zip, .rar, .tar, .gzip
PDF	PDF, e-mail, and MBOX files	.email, .mbox, .pdf, .pdfa

Table 44: File Category Contents *(Continued)*

Category	Description	File Types
Rich Application	Installable Internet Applications such as Adobe Flash, JavaFX, Microsoft Silverlight	.swf, .xap, .xbap
OS package	OS-specific update applications	.deb, .dmg
Java	Java applications, archives, and libraries	.class, .ear, .jar, .war
Script	Scripting files	.bat, .js, .pl, .ps1, .py, .sct, .sh, .tcl, .vbs, plsm, pyc, pyo
Document	All document types except PDFs	.chm, .doc, .docx, .dotx, .hta, .html, .pot, .ppa, .pps, .ppt, .pptsm, .pptx, .ps, .rtf, .txt, .xlsx, .xml,
Executable	Executable binaries	.bin, .com, .dat, .exe, .msi, .msm, .mst
Library	Dynamic and static libraries and kernel modules	.a, .dll, .kext, .ko, .o, .so, .ocx
Mobile	Mobile formats	.apk, .ipa
Configuration	Configuration files	.inf, .ini, .lnk, .reg, .plist

4. Click **Save** at the bottom of the pop-up window.

Your new profile appears in the list of anti-malware profiles.

Apply a Profile to an Application Policy

Before You Begin

In this procedure, you'll apply security profiles to application policies. First, you must complete the following tasks:

- Create your application policies. For help, see "Application Policies" on page 144.
- Create your security profiles. For help, see the earlier sections of this topic.

1. From the left menu, select **Organization** > **WAN** > **Application Policy**.
2. In the list of policies, find the one that you want to modify.
3. In the **Advanced Security Services** column, click **+**, and then click the profile to apply.
   In the drop-down menu, search by the name of the profile, or scroll through the list. Refer to the section headings to find profiles by type, such as Security Intel or Anti-Malware.

   For example, this image shows a profile in the Security Intel section of the drop-down menu.

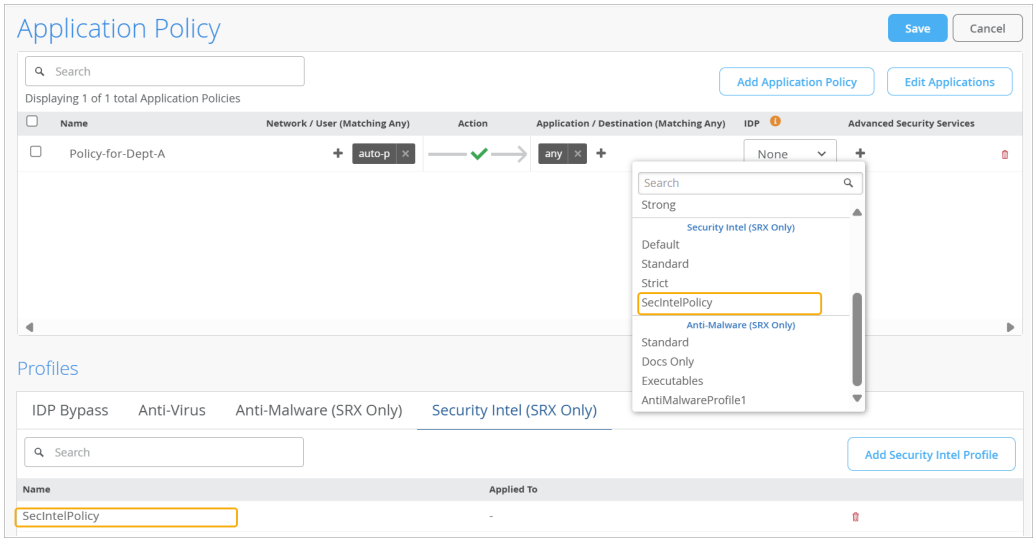

4. When you're done applying profiles to policies, click **Save** at the top-right corner of the Application Policy page.

Status Information

You can view basic device monitoring information on the WAN Edges page. From the left menu,navigate to **WAN Edges** > **WAN Edges**.

After you click a device on the WAN Edges page, the device details appear. On the details page, you'll see:

- **Advanced Security**—The status of security services. A green check mark indicates that the service is active on this device.

Figure 44: Advanced Security Status Details

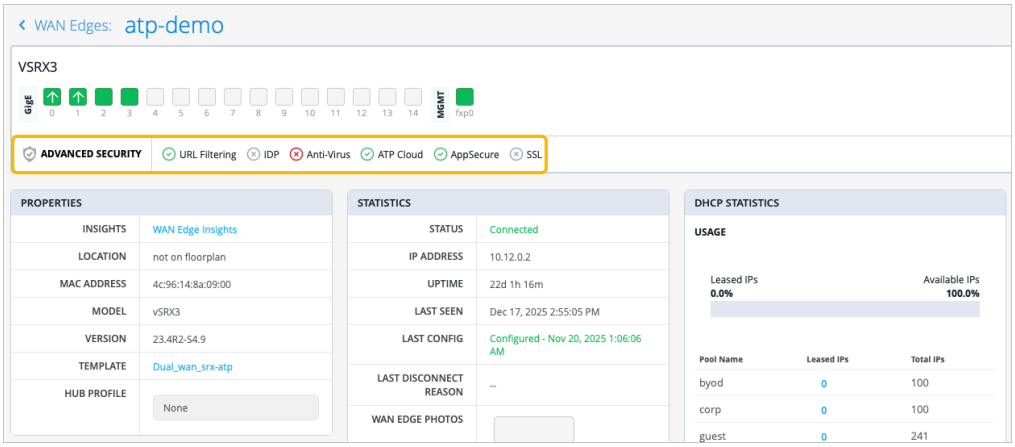

- **Properties**—Click **WAN Edge Insights** to view recent events and other information on the Insights page.

Figure 45: WAN Edge Events Details

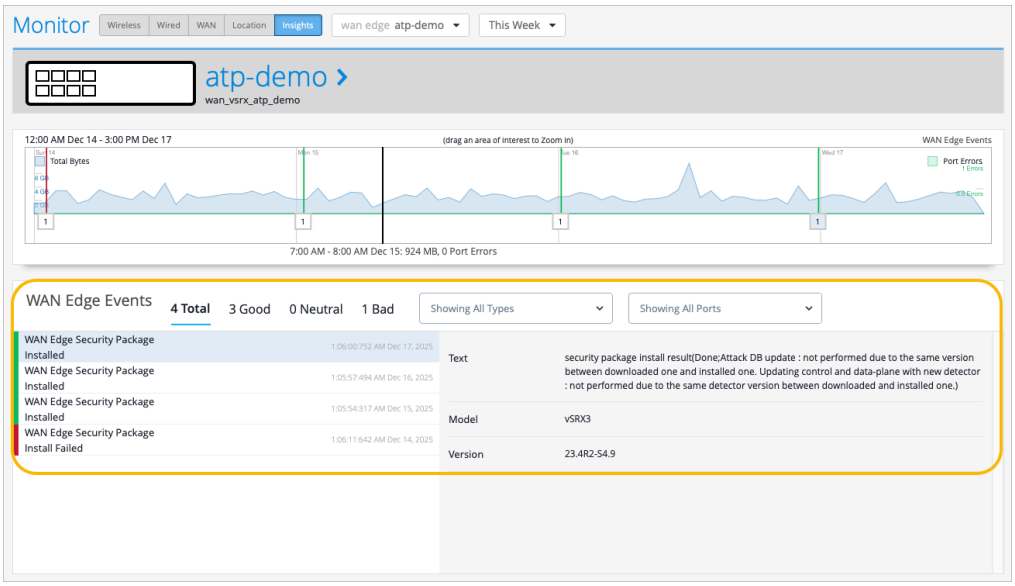

TIP:

- By default, the Insights page shows only today's events, but you can select a time period from the drop-down menu near the top-right corner of the page.
- In the WAN Edge Events list, you can click an event to see a summary on the right side of the page.

View Security Events

The Juniper Mist Security Events page, accessible through **Site > WAN Edge > Security Events**, provides a centralized view of security-related events. It displays a log of security events detected by Juniper Mist to monitor the security posture of the network. Users can filter and view details allowing for proactive security response and analysis.

Figure 46: Security Events

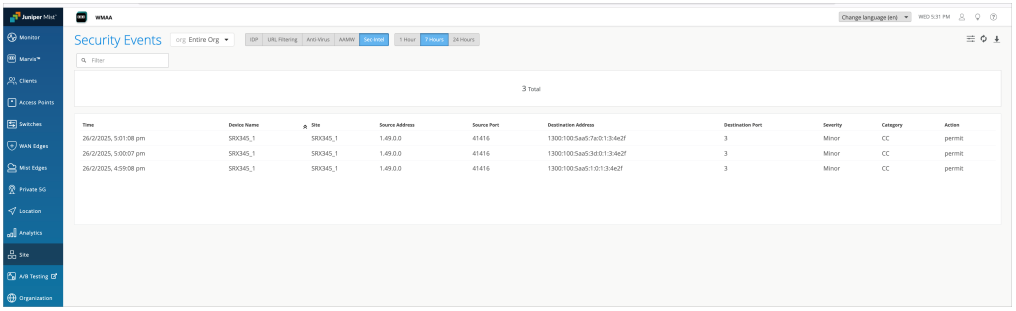

Click one of the tabs **AAMW** (Advanced Anti-Malware) or **SecIntel** to see the related security event details. In the example above, the page shows incident details for Command and Control (C&C) with a severity level of Minor. It also indicates the action taken, which is Permit in this case. Additionally, you can view other information such as the device name, site, source and destination addresses, and source and destination ports information.

CHAPTER 6

WAN Edge Device Options

IN THIS CHAPTER

Onboard Session Smart Routers with Static IP Address

SUMMARY

Follow these steps if your routers need static IP addresses.

IN THIS SECTION

Dynamically assigned IP addresses are recommended, but in some cases you might need static addresses instead. For example, consider a situation when you're deploying sites that are connected solely through WAN circuits.

To onboard your WAN Edge devices, you can take various approaches:

- Lightweight Field ZTP
- One-Time Static Device Provisioning

Before You Begin

- Get your WAN Edge device up and running in the Mist cloud. See SSR Series Devices.
- Set up your WAN Edge templates. In the WAN section of the template, add a WAN, select the static IP configuration option, and enter the IP address, prefix length, and gateway.

  Here's an example using variables to represent the information. You'd then enter the static IP addresses in the variable definitions in the site configuration.

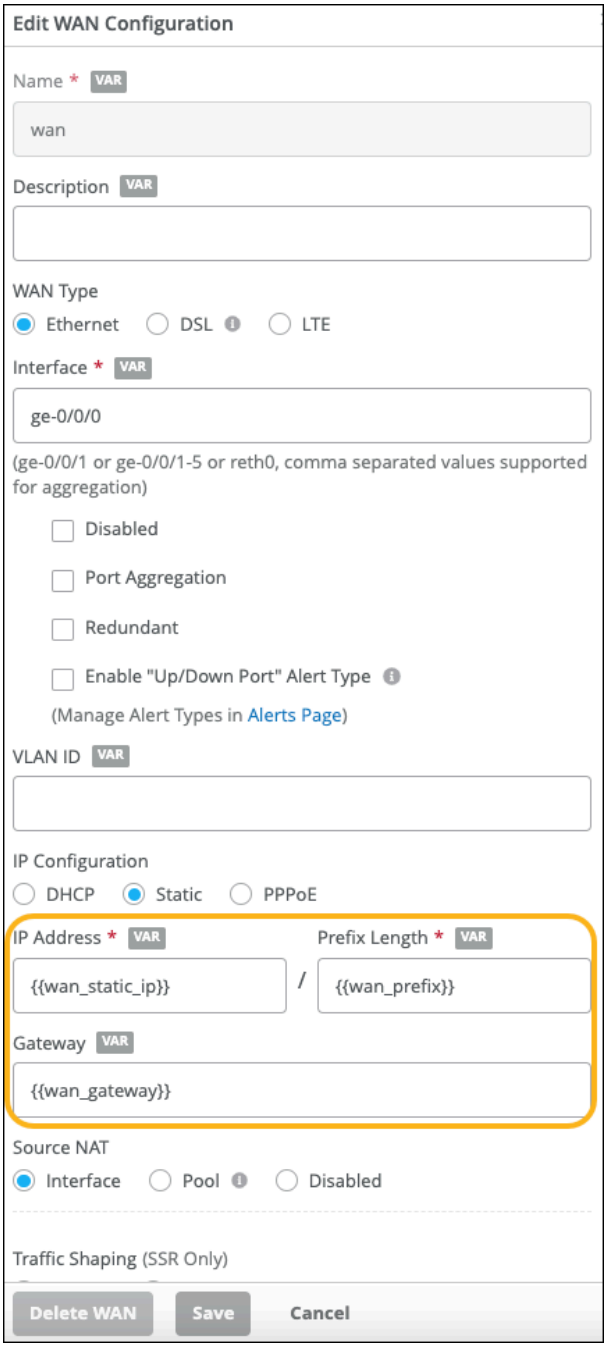

For help with templates, see "Configure WAN Edge Templates." on page 87

Onboard Devices with Lightweight Field ZTP

The lightweight field ZTP method for device onboarding in a static WAN scenario requires temporarily connecting the Session Smart Router device to onsite connectivity. This connection provides dynamic address assignment through DHCP and temporary connectivity to Mist cloud. For example, connect the

WAN port of the Session Smart Router device to a mobile wireless router carried by the installer. Alternatively, connect it to the LAN port of a site router about to be decommissioned.

To onboard a WAN Edge device with lightweight field ZTP:

1. Claim the device to your Mist org inventory (if not done already), and power it on. For details, see SSR Series Devices.
2. Connect the ge-0-0 WAN port to an onsite device that can provide temporary connectivity to the Mist cloud.
3. From the left menu of the Juniper Mist portal, select **Organization** > **Admin** > **Inventory**.
4. Refresh your browser.
5. From the left menu, select **WAN Edges** > **WAN Edges**.
6. In the inventory list, find your newly added device, and select its check box.
7. Click the **More** menu near the top-right corner of the page, and then click **Assign to Site**.

   NOTE: The menu appears only if you've selected a check box in the list of devices.

8. In the Assign WAN Edges pop-up window, select a site and keep the default Device Configuration options.

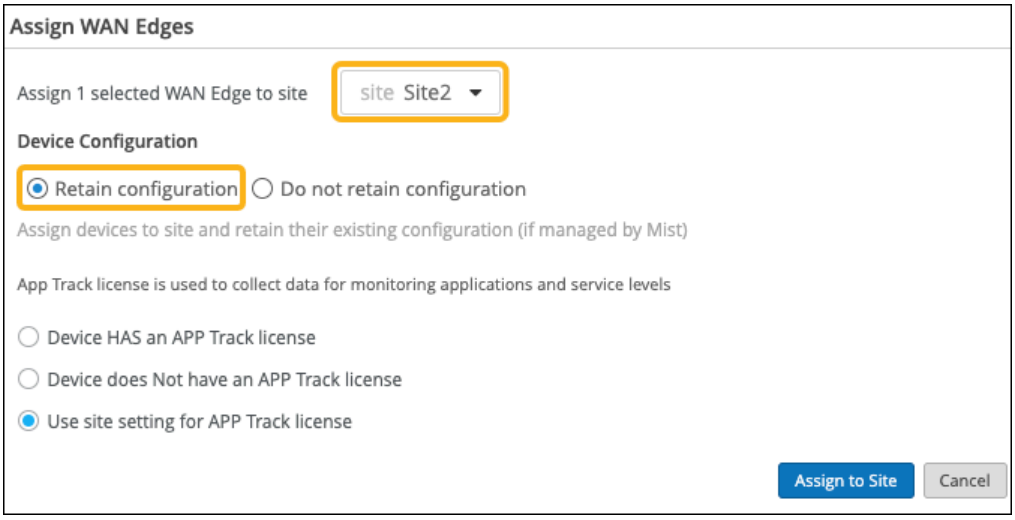

9. Upgrade the device to fully apply and run the configuration from Mist cloud.
10. Monitor WAN Edge events to confirm that the device receives the configuration.
    a. From the left menu, select **Monitor** > **Service Levels**.
    b. Click the **Insights** button.
    c. Select the WAN Edge device by using the shortcut menu next to the Insights button. Details appear in the WAN Edge Events section of the page.

d. Wait until the WAN Edge device receives the configuration and then disconnects from the Mist cloud.

11. Disconnect the device from the temporary connection that you established earlier.
12. Connect the WAN Edge device to the static WAN link.
13. Verify that the device uses the static WAN link to connect to the Mist cloud.

Onboard Devices with One-Time Static Device Provisioning

This method involves connecting to the local console of the device, and pasting in some one-time provisioning information including the static address.

1. Claim the device to your Mist org inventory (if not done already), and power it on. For details, see SSR Series Devices.
2. In the Juniper Mist portal, click **Organization** > **Admin** > **Inventory**.
3. Refresh your browser.
4. Check under **WAN Edges** to confirm that your device is in the inventory.
5. Assign Session Smart Router to an individual site using the **Assign to Site** option.

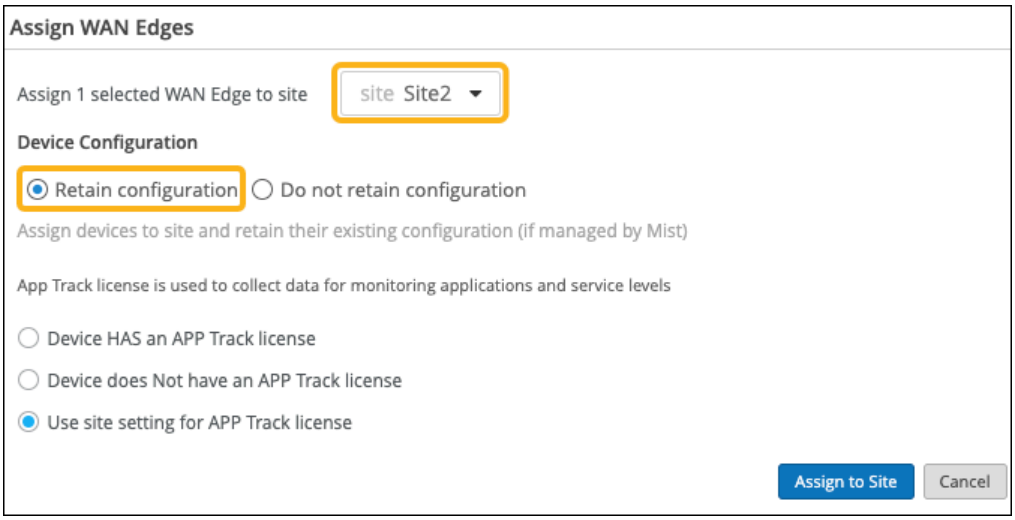

6. Connect a laptop or PC to the local serial console port of the Session Smart Router on the site.
7. Enter the default login credentials.
8. Use the following codeblock to enter the configuration.

   When entering the configuration, replace {{wan0_static_ip}}, {{wan0_prefix}}, and {{wan0_gateway}} with the correct details for the WAN link.

```
configure authority router router node node device-interface ge-0-0 network-interface
ge-0-0-intf dhcp disabled
```

```
configure authority router router node node device-interface ge-0-0 network-interface
ge-0-0-intf address {{wan0_static_ip}} prefix-length {{wan0_prefix}}

configure authority router router node node device-interface ge-0-0 network-interface
ge-0-0-intf address {{wan0_static_ip}} gateway {{wan0_gateway}}

commit force
```

9. Connect the device to the appropriate static WAN link.
10. Locally on the device, verify the local link state and connectivity by using commands such as `show arp`, `show network-interface`, and `ping`.

```
admin@node.router# show arp

Thu 2023-08-03 03:04:33 UTC

Node: node.router Page 1

========== ====== ================ ================== =======

Dev Name   VLAN   IP               Dest MAC           State

========== ====== ================ ================== =======

ge-0-0          0   96.93.100.38      80:b2:34:5a:b2:43   Valid

lte-0         0   169.254.128.131   92:3d:72:f7:cb:66   Valid

kni254        0   169.254.127.127   fe:d4:a2:e9:17:ec   Valid
```

11. Upgrade the device to fully apply and run the configuration from Mist cloud.

On the WAN Edges page, you can view the status of the onboarded device.

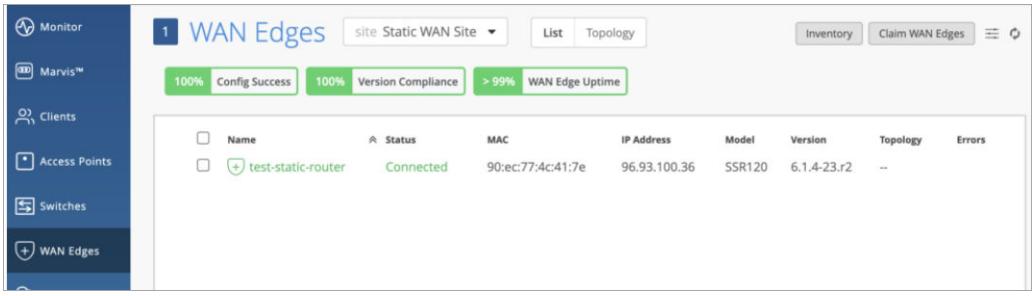

On the Insights page, you can see more information under WAN Edge Events.

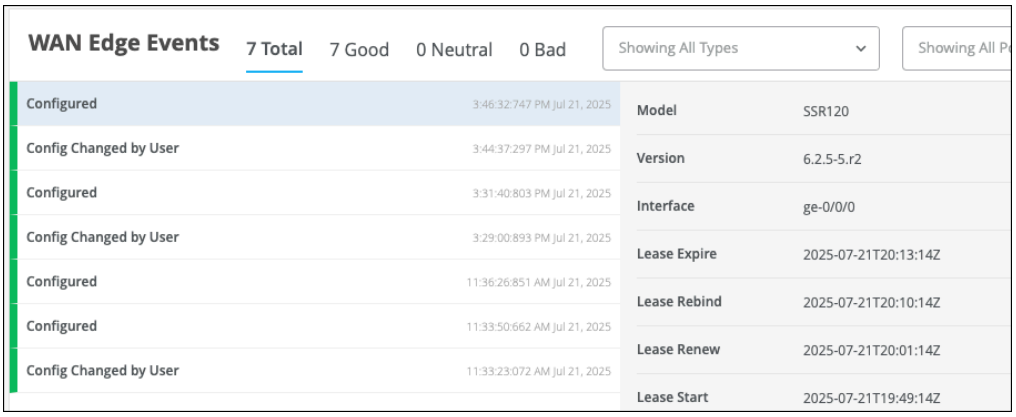

Upgrade a WAN Edge Device

SUMMARY

Follow these steps to keep your WAN Edge device up to date with the latest features and fixes.

Regularly upgrade your device to stay current with new features, enhancements, bug fixes, and compatibility improvements.

1. From the left menu, select **WAN Edges** and then select the device.
2. Click the **Utilities** button at the top-right corner of the page, and then click **Upgrade Firmware**.

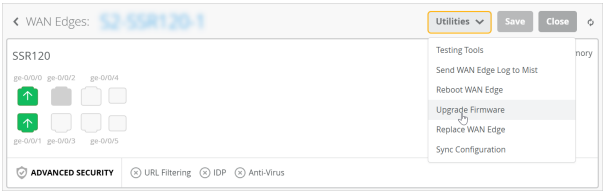

3. In the pop-up window, enter the settings for this upgrade.

Table 45: Upgrade Settings

Field	Description
Selected Channel	• Beta—Non-production lab environment • Production—Production environment
Upgrade to Version	Select the version to install.
Schedule Download Time	Start the download immediately or at a later time. • Now—Download the firmware immediately when you click the button at the end of this procedure. • Later—Download the firmware at the date and time that you specify (for example,during off-peak hours for the site).
Upgrade device after download *(Available for SSR only)*	Start the upgrade immediately after download, or at a later time. • Upgrade Immediately After Download—Start the upgrade after the download is completed. • Later—Start the upgrade at the date and time that you specify (for example, during off-peak hours for the site).
Reboot device after upgrade *(Available for SRX only)*	Select this option if you want the device to reboot automatically after the upgrade. If you don't select this option, you'll need to manually reboot the device to complete the upgrade process.
Create a recovery snapshot post upgrade *(Available for SRX only)*	If you select this option, a recovery snapshot is stored in Junos OAM (Operations, Administration, and Maintenance). OAM holds a full backup that can be used to restore the configuration if something goes wrong with the Junos volume.

4. Read and accept the End User License Agreement.
5. Click the button at the bottom of the window to continue.

   The name of the button depends on the device and the options that you selected earlier in the process: **Start Upgrade**, **Schedule Download**, or **Schedule Download and Upgrade**.

To monitor the upgrade progress, you can go to **WAN Edges** > **WAN Edges** and view the progress in the Status column next to the device.

Figure 47: View Upgrade Status from the WAN Edges page

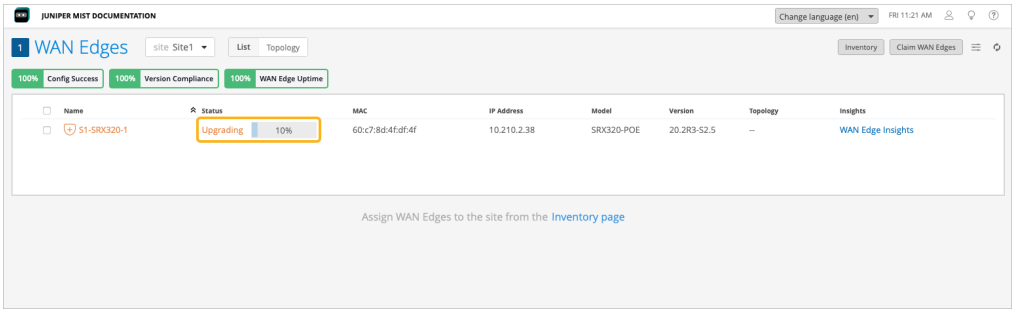

You can also click the WAN Edge name from the WAN Edges page and view the upgrade progress in the Status row in the Statistics section.

You can see the following stages of firmware upgrade:

- Initiation of the upgrade operation
- Completion of firmware download
- Device reboot to complete the upgrade process
- Completion of firmware upgrade

Replace a Standalone WAN Edge Device

SUMMARY

You can replace a connected or disconnected WAN Edge device with another device of the same model.

To replace a WAN Edge device with another device of the same model:

1. From the left menu of the Juniper Mist portal, select **Organization** > **Admin** >**Inventory** and select the **WAN Edges** tab.

The page displays a list of WAN Edge devices. You can set the view as **org (Entire Org)** or **Site** in the Inventory page.

2. Click the device that you want to replace.
3. Select **Replace WAN Edge** from Utilities drop-down.
4. On the Replace WAN Edge window, select the new replacement device's MAC address from the **MAC Address of unassigned WAN Edge** drop-down list.

Figure 48: Replace a Standalone WAN Edge device

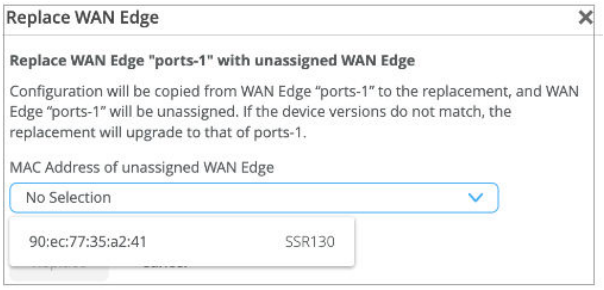

The Juniper Mist portal displays a list of supported models available in the Inventory page in an unassigned state.

After you click **Replace**, allow about 15 minutes to complete the replacement procedure. Mist copies the configuration of the replaced WAN Edge device to the new device. The device you replaced continues to be part of the site in an unassigned state.

NOTE: Replacing a device causes minimal impact on network services. Therefore, we recommend that you plan to do this task during a maintenance window.

Refresh your browser and check under WAN Edges to find out if your WAN Edge device is available as a part of the inventory.

Revoke DHCP Lease on a WAN Edge Device

SUMMARY

Follow these steps to release a client device from its current DHCP lease.

To revoke a DHCP lease on a WAN Edge device:

1. From the left menu, select **WAN Edges** > **WAN Edges**.
2. Click the device responsible for this DHCP lease.
3. Scroll down to the **DHCP Statistics** section of the device details page.
4. Under **Leased IPs** , click the hyperlink for the pool that includes this lease.

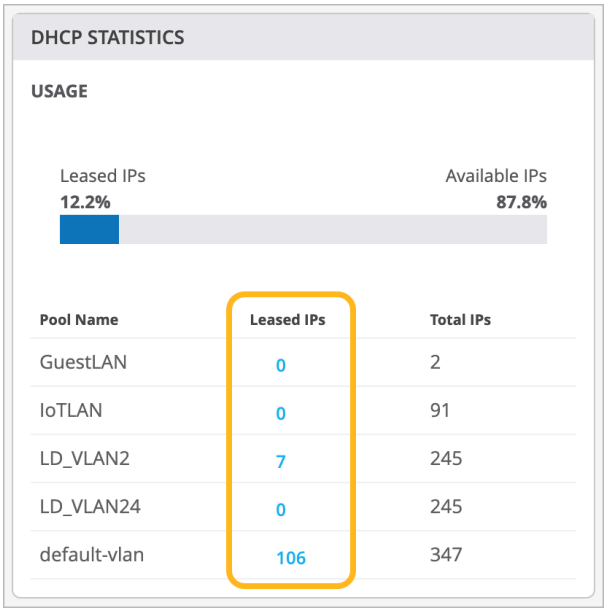

On the Leased IPs window, you can view the client devices (MAC addresses or hostnames) along with the leased IP addresses and the expiration dates and times.

5. Select the checkbox for each device that you want to revoke, and then click **Revoke**.

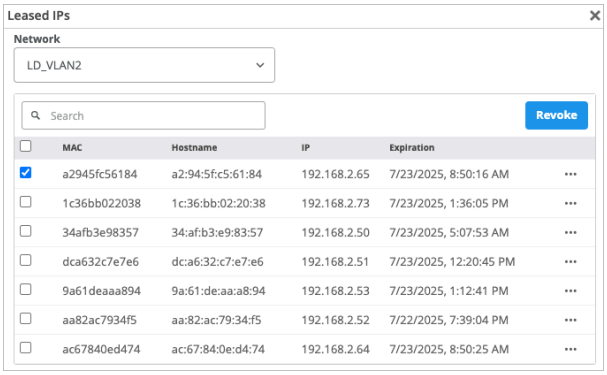

TIP: To manage leases for a different pool, select it from the **Network** list at the top of the Leased IPs pop-up window.

6. When finished, click the **X** in the top-right corner of the Leased IPs pop-up window.

Reserve a Static IP Address for a Device

SUMMARY

Follow these steps to reserve an IP address for exclusive use by a single device.

To reserve a static IP address for a device:

1. Navigate to the WAN Edge template, hub profile, or device configuration.
2. Scroll down to the **LAN** configuration section.
3. In the **DHCP Config** section, click the configuration that you want to modify.

   Or, to create a DHCP configuration, click **Add DHCP Config**. For help with the settings, see "LAN Interfaces" on page 125.
4. In the **Static Reservations** section of the side panel, click **Add Reservation**.

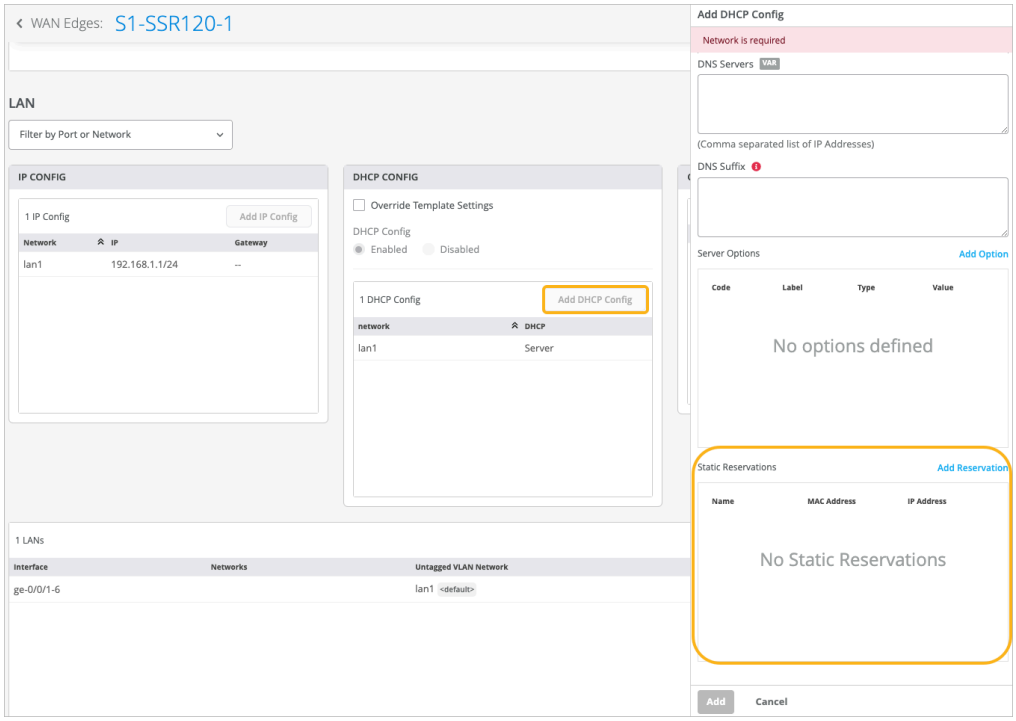

5. Enter the settings:

Table 46: Reservation Settings

Field	Description
Name	Enter a descriptive name to identify the use or purpose of this static IP address.
MAC Address	Enter the MAC address of the device that needs the static IP address.
IP Address	Enter an address or a variable that represents an address within the DHCP range. Refer to the **IP Start** and **IP End** fields in this DHCP Config.

6. Click the blue checkmark in the Add Static Reservation title bar.
7. Click **Add** at the bottom of the Add DHCP Config side panel.
8. Click **Save** at the top-right corner of the WAN Edge Template page.

CHAPTER 7

Secure Edge Connector

IN THIS CHAPTER

Juniper Mist Secure Edge Connector Overview

SUMMARY

Before integrating Secure Service Edge (SSE), get familiar with key concepts, deployment options, and the high-level workflow.

IN THIS SECTION

Juniper Mist provides pre-built connectors specifically designed for the Juniper Networks® SRX Series Firewalls and Juniper® Session Smart™ Routers deployed as WAN edge devices. These connectors facilitate seamless integration with your Secure Service Edge (SSE) deployments. With minimal configuration, you can integrate the SSE into the Juniper Mist portal. As a result, your WAN Edge device establishes connections to the SSE using either IPsec or GRE protocols.

Figure 49: Traffic Inspection by Juniper Secure Edge

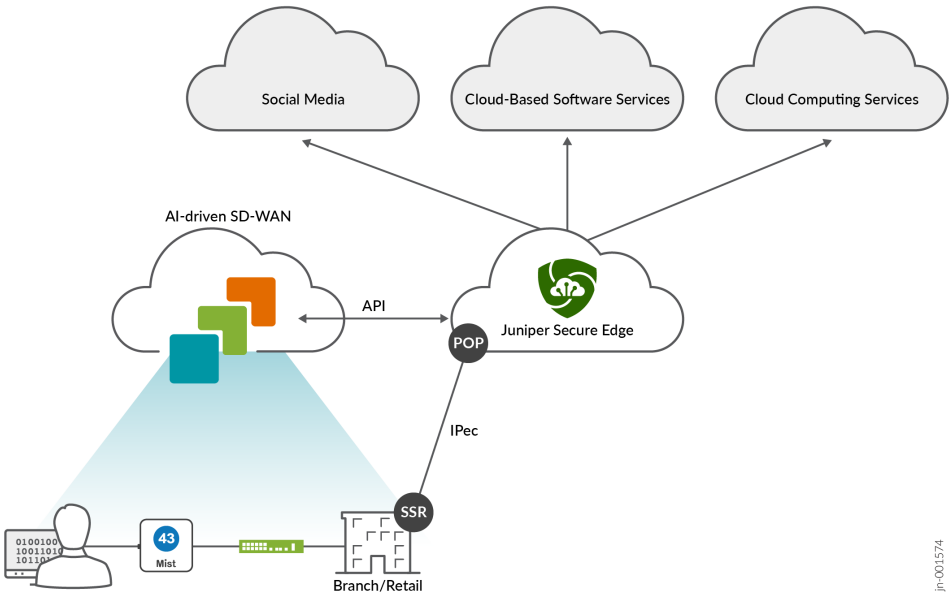

In this solution, an IPsec tunnel is configured between the WAN Edge device and SSE using the Secure Edge Connector (SEC) within the WAN Edge template. Additionally, a BGP over IPsec connection is configured to dynamically learn routing destinations from the SSE device.

Juniper Mist supports these SEC types:

- Juniper Secure Edge (manual provisioning and auto provisioning)
- Zscaler (manual provisioning and auto provisioning)
- Custom

High-level workflow for setting up secure edge connectors with Juniper Secure Edge, custom, or Zscaler deployment to offload traffic from your WAN edge device (SSR Series Routers or SRX Series Firewalls):

1. Create and deploy a basic branch template for device connectivity.
2. Optionally configure a remote network in SSE. This step defines a remote source for inbound connectivity through the tunnel.
3. Configure an SEC and SSE provider in the device template. This step creates a custom IPsec tunnel to the remote location and define encryption parameters.
4. Optionally configure a BGP peer to learn routes dynamically.
5. Configure an application that allows traffic to be steered toward the IPsec tunnel. This application will be used in application policy to allow client networks to access the BGP learned routes.
6. Configure a traffic steering policy that steers the Internet-bound traffic from the LAN side of a spoke or hub device to the SEC.

Application Policies for SEC

An application policy defines which networks and users (traffic source) can access which applications (traffic destination). You also can use Traffic Steering to define which path to use.

To set up these policies, you need to create networks, applications, and traffic-steering profiles.

- For outbound traffic—Include the SEC in the traffic steering profile.
- For inbound traffic through the SEC—Include the remote network in the SEC, and define an application policy that allows inbound access. With this feature, you can securely connect to cloud-hosted services that initiate inbound traffic to a site.

Traffic Steering Profiles for Secure Edge Connector

Traffic steering is required for SEC on both SRX Series Firewalls and Session Smart Routers before Juniper Mist creates the tunnels.

This requirement remains unless:

- A remote network is assigned to an SEC
- A BGP peer is assigned to an SEC

Dynamic Routing for SEC

You can configure BGP peering over your SEC. This configuration leverages BGP for dynamic routing and uses BGP path selection to install routes in the route table. High-Level steps include:

- Verify that your SEC is established and is configured using the custom Secure Edge provider.
- Configure BGP import and export policies.
- Configure BGP neighbor options.
- Select the SEC for this BGP neighbor.
- Assign import and export policies.
- Verify that the BGP peers are exchanging routes over the tunnel interface.

Set Up Secure Edge Connector with Juniper Secure Edge (Manual Provisioning)

SUMMARY

Follow this workflow for help with every step of the workflow to manually provision Juniper® Secure Edge.

IN THIS SECTION

The Juniper Mist™ cloud works with Juniper® Secure Edge to perform traffic inspection from edge devices by using the Secure Edge connector feature. This feature allows the Juniper® Session Smart™ Routers, deployed as WAN edge device, to send a portion of traffic to Juniper Secure Edge for an inspection.

Secure Edge capabilities are all managed by Juniper Security Director Cloud, Juniper's simple and seamless management experience delivered in a single user interface (UI).

For more information, see Juniper Secure Edge.

Configuration Overview

In this task, you send the Internet-bound traffic from the LAN side of a spoke or hub device to Secure Edge for an inspection before the traffic reaches Internet.

To perform traffic inspection by Secure Edge:

- In Juniper Security Director Cloud, create and configure the service locations, IPsec profiles, sites, and policies for Secure Edge. These are the cloud-based resources that provide security services and connectivity for the WAN edge devices.
- In Mist Cloud, create and configure the WAN edge devices (Session Smart Routers or SRX Series Firewalls), that connect to the LAN networks. These are the physical devices that provide routing, switching, and SD-WAN capabilities for the branches or campuses.
- In Mist WAN-Edge, create and configure the Secure Edge tunnels that connect the WAN edge devices to the service locations. These are the IPsec tunnels that provide secure and reliable transport for the traffic that needs to be inspected by Secure Edge.
- In Mist Cloud, assign the Secure Edge tunnels to the sites or device profiles that correspond to the WAN edge devices. This enables the traffic steering from the LAN networks to the Secure Edge cloud based on the defined data policies and other match criteria.

Topics in the following table present the overview information you need to use the cloud-based security of Secure Edge with the Juniper Mist™ cloud.

Table 47: Secure Edge Connector Configuration Workflow

Step	Task	Description
1	"Access Juniper Security Director Cloud and Check Active Subscriptions " on page 228	Access Juniper Security Director Cloud, go to your organization account, and check Secure Edge subscriptions. The subscription entitles you to configure Secure Edge services for your deployments.
2	"Configure a Service Location in Juniper Security Director Cloud " on page 229	Create service locations. This is where the vSRX-based VPN gateways creates secure connections between different networks.
3	"Generate Device Certificates in Juniper Security Director Cloud" on page 230	Generate digital certificates for Juniper Secure Edge to establish secure communications between Secure Edge and user endpoints.
4	"Create an IPsec Profile in Juniper Security Director Cloud " on page 233	Create IPsec profiles to establish IPsec tunnels for communication between the WAN edge devices on your Juniper Mist cloud network with Secure Edge instance.

5	"Create a Site in Juniper Security Director Cloud " on page 233	Create a site that hosts a WAN edge device (Session Smart Router or SRX Series Firewall). The traffic from the device is forwarded to the Secure Edge instance through a secure tunnel for an inspection.
6	"Deploy a Secure Edge Policy in Juniper Security Director Cloud " on page 240	Configure policies that define the security rules and actions for the traffic originating from or destined to the site.
7	"Get IPsec Tunnel Configuration Parameters to Apply in Juniper Security Director Cloud " on page 241	Note down the details such as service location IP or hostname, the IPsec profile name, and the pre-shared key. You need these details to set up IPsec tunnels from Juniper Mist side.
8	"Create Secure Edge Connectors in the Juniper Mist Cloud Portal " on page 244	Create Secure Edge connectors in the Juniper Mist cloud portal. This task completes the configuration on the Mist cloud side of the tunnels to establish an IPsec tunnel between WAN edge device managed by Mist and the Secure Edge instance.
9	"Modify an Application Policy " on page 250	Create a new or change an existing application policy to direct the traffic from WAN edge device to the Internet through Juniper Security Director Cloud instead of going through a hub for centralized access.
10	"Verify the Configuration" on page 250	Confirm if your configuration is working by checking the established IPsec tunnels in: • WAN Insights in Mist portal • Security Director Cloud dashboard • Tunnel traffic flow on the WAN edge device CLI.

Before You Begin

- Read about the Juniper® Secure Edge subscription requirements. See Juniper Secure Edge Subscriptions Overview.

- Ensure that you have completed the prerequisites to access the Juniper Security Director Cloud Portal . See Prerequisites.
- Create Your Secure Edge Tenant. See Create Your Secure Edge Tenant.
- We assume that you have adopted and configured the Session Smart Router or SRX Series Firewall deployed as WAN edge device in Juniper Mist Cloud.

Access Juniper Security Director Cloud and Check Active Subscriptions

A tenant in Juniper Secure Edge is an organization account that you create to access the Juniper Security Director Cloud portal and manage your Secure Edge services. A tenant is associated with a unique e-mail address and a subscription plan. A tenant can have multiple service locations, which are vSRX based VPN gateways hosted in a public cloud for your organization.

A tenant can have one or more service locations, which are the connection points for end users. To create a tenant, you need to have an account on Juniper Security Director Cloud. See Create Your Secure Edge Tenant for details.

After you create your Secure Edge tenant in the Juniper Security Director Cloud portal, access the portal and check your subscriptions.

To access Juniper Security Director Cloud and check active subscriptions:

1. Open the URL to the Juniper Security Director Cloud. Enter your e-mail address and password to log in and start using the Juniper Security Director Cloud portal.

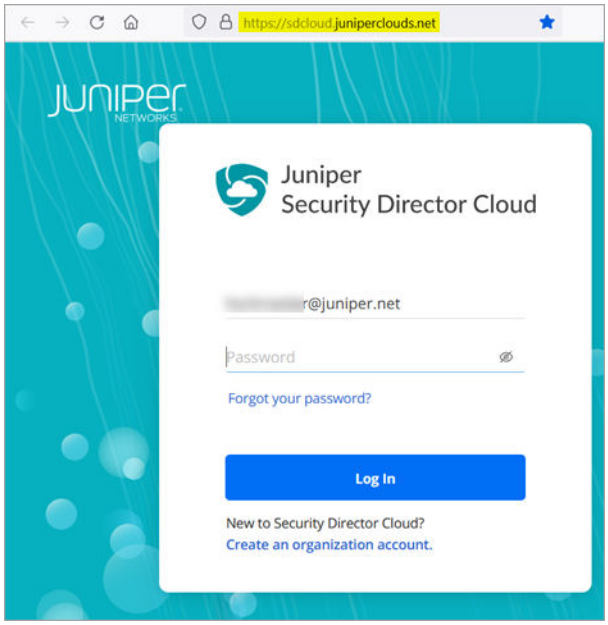

2. Select the required tenant in the upper right corner of the portal to continue.

3. Select **Administration** > **Subscriptions** to access the Juniper Security Director Cloud subscriptions page.

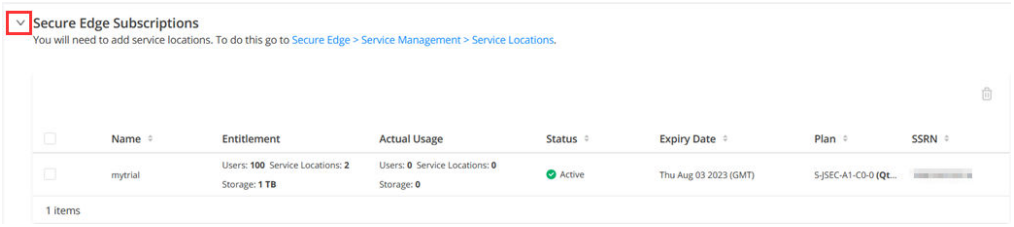

4. Scroll to the **Secure Edge Subscriptions** section to check whether you have an active subscription.

> **NOTE**: You do not need to click the **SRX Management Subscription** tab, even if you are using a Juniper Networks® SRX Series Firewall. In this task, you are not using Juniper Security Director Cloud for managing WAN edge devices.

For details, see About the Subscriptions Page.

Assuming that you have active subscriptions, continue with next steps.

Configure Service Locations

After ensuring that you have an active license to Juniper Security Director Cloud, you configure a service location. This step is your first main task in setting up a Secure Edge connector for Session Smart Routers.

A service location in Juniper Security Director Cloud is also known as POP (point of presence) and represents a Juniper® Secure Edge instance in a cloud location. The service location is the connection (access) point for both on-premises and roaming users.

Service locations are places where vSRX creates secure connections between different networks using a public cloud service. The public IP address (unique per tenant and service location) is used to:

- Set up an IPsec tunnel between the branch device and the Juniper Security Director Cloud.
- Centrally distribute the traffic when the destination is on the Internet.

To configure a service location in Juniper Security Director Cloud:

1. In Juniper Security Director Cloud menu, select **Secure Edge**>**Service Management**>**Service Locations**.
   The Service Locations page appears.
2. Click the Add (**+**) icon to create a new service location.

Enter the details for the following fields:

- **Region**—Choose the geographic region where you want to create a Secure Edge instance.
- **PoP**—Select the location for the Secure Edge in the region.
- **Number of Users**—Enter the total possible number of users this service location may need to serve.

Table 48: Service Location Fields

Field	Service Location	Service Location
Region	North America	North America
PoP	Ohio	Oregon
Number of Users	50 (we split the exiting equally)	50

3. Click **OK**.

Security Director Cloud creates a service location and lists it on the Service Locations page.

The status of the service location shows **In Progress** until the Secure Edge instance is fully deployed.

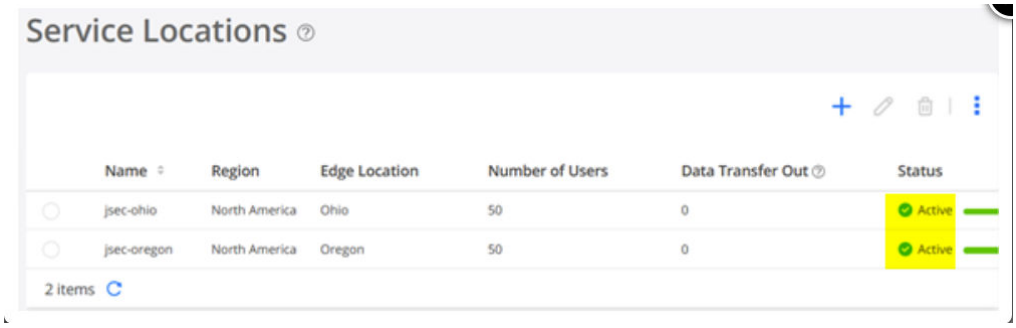

When you create a service location, the system starts the deployment of two vSRX instances as VPN gateways for your tenant system. In this deployment, vSRX instances are not shared with other tenants.

Generate Device Certificates in Juniper Security Director Cloud

Now that you have configured service locations in Juniper Security Director Cloud, you generate device certificates to secure network traffic.

You use a Transport Layer Security/Secure Sockets Layer (TLS/SSL) certificate to establish secure communications between Secure Edge and WAN edge devices. All the client browsers on your network must trust the certificates signed by the Juniper Networks and SRX Series Firewalls to use an SSL proxy.

In Juniper Security Director Cloud, you have the following choices for generating certificates:

- Create a new certificate signing request (CSR), and your own certificate authority (CA) can use the CSR to generate a new certificate.
- Select the option to have Juniper Networks create a certificate.

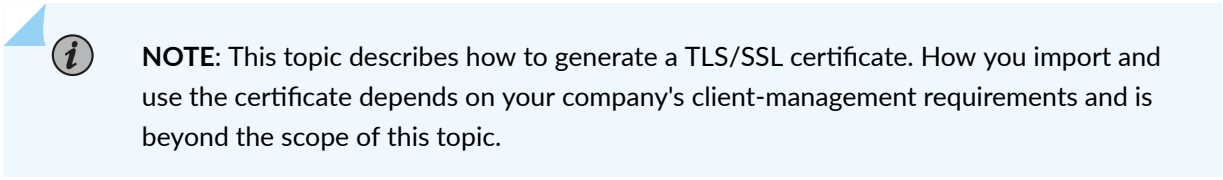

NOTE: This topic describes how to generate a TLS/SSL certificate. How you import and use the certificate depends on your company's client-management requirements and is beyond the scope of this topic.

To generate device certificates in Juniper Security Director Cloud:

1. Select **Secure Edge**>**Service Administration**>**Certificate Management**.
   The Certificate Management page appears.

   From the **Generate** list, you can generate either a new **Certificate signing request** (CSR) or a **Juniper issued certificate**.

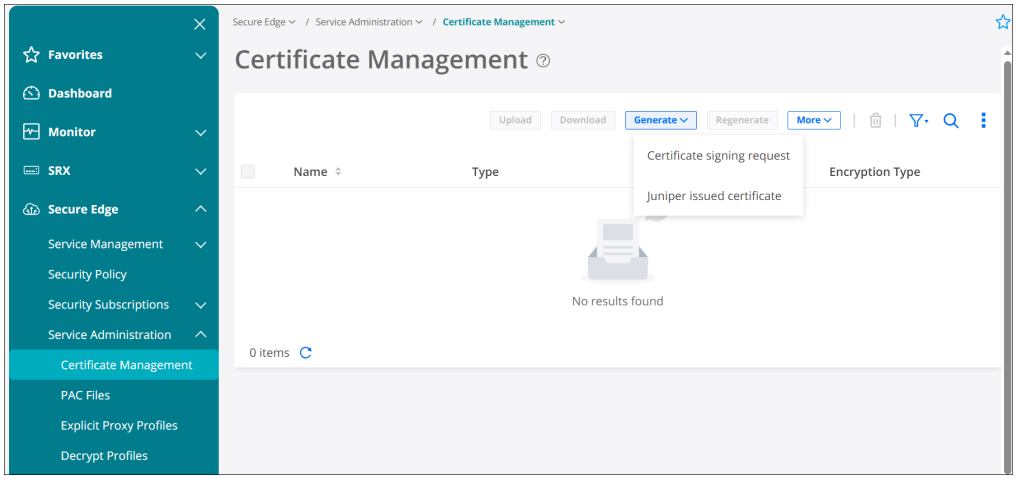

2. Select the relevant option:
   a. If your company has its own CA, and you want to generate a CSR, click **Certificate signing request**.
      After Juniper Secure Edge generates CSR, download the CSR and submit it to your CA to generate a new certificate. Once generated, click Upload to upload the certificate on the Certificate Management page.
   b. If your company does not have its own CA, click **Juniper issued certificate**, and then click **Generate** to generate the certificate. Juniper Networks will generate and keep the certificate on the system.

In this task, select **Juniper issued certificate** and continue with next step.

3. Enter the certificate details. In the **Common name** field, use the certificate's fully qualified domain name (FQDN).

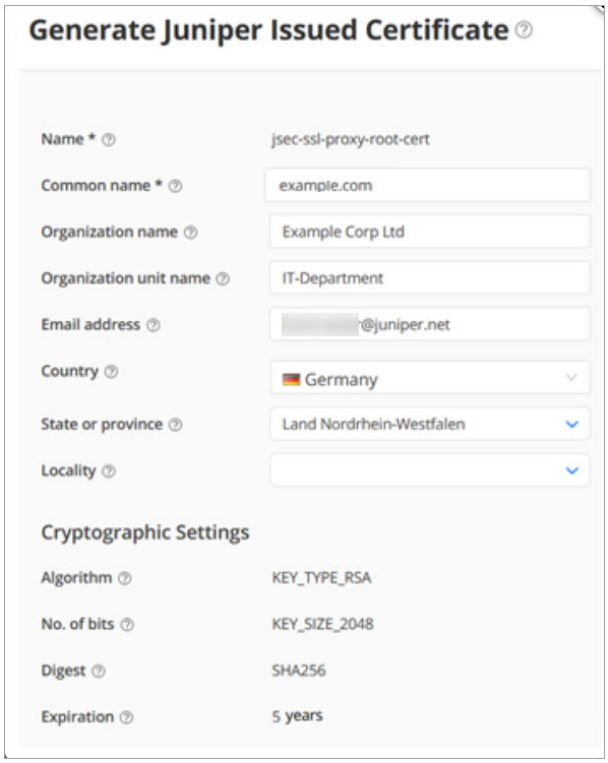

The Certificate Management page opens with a message indicating that the certificate is created successfully.

4. Download the generated certificate.

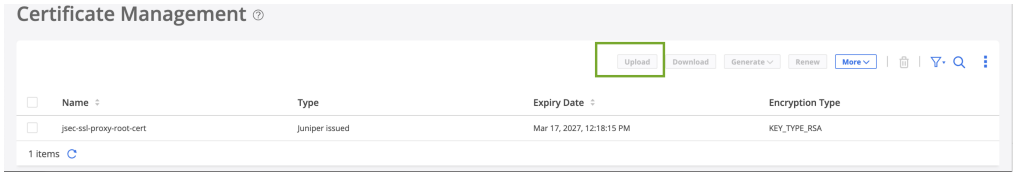

The following sample shows the downloaded certificate:

```
-----BEGIN CERTIFICATE-----
$ABC123#1$ABC123#1 $ABC123#1$ABC123#1 $ABC123#1$ABC123#1 $ABC123#1$ABC123#1.
.
J$ABC123#1$XYZ123#1 $ABCBVFCC123#1$ABC123#1 $XYZ123#1$GTF123#1 $XZY123#1$BVFD123#1=
$ABC123#1$XYZ123#1 $ABCBVFCC123#1$ABC123#1 $XYZ123#1$GTF123#1 $XZY123#1$BVFD123#1=
-----END CERTIFICATE-----
```

After you download the certificate to your system, add the certificate to client browsers.

Create an IPsec Profile in Juniper Security Director Cloud

After you generate the certificates to establish secure communications between Secure Edge and WAN edge devices, you're ready to create IPsec profiles.

IPsec profiles define the parameters with which an IPsec tunnel is established when the WAN edge devices on your Juniper Mist™ cloud network start communicating with your Secure Edge instance.

To create an IPsec profile in Juniper Security Director Cloud:

1. In Juniper Security Director Cloud portal, select **Secure Edge** > **Service Management** > **IPsec Profiles** .
2. Click the Add (**+**) icon to create an IPsec profile.
   The Create IPsec Profile page appears.
3. For the profile name, use **default-ipsec**. Retain all default values for Internet Key Exchange (IKE) and IPsec; currently, they are not configurable on the Juniper Mist cloud portal.

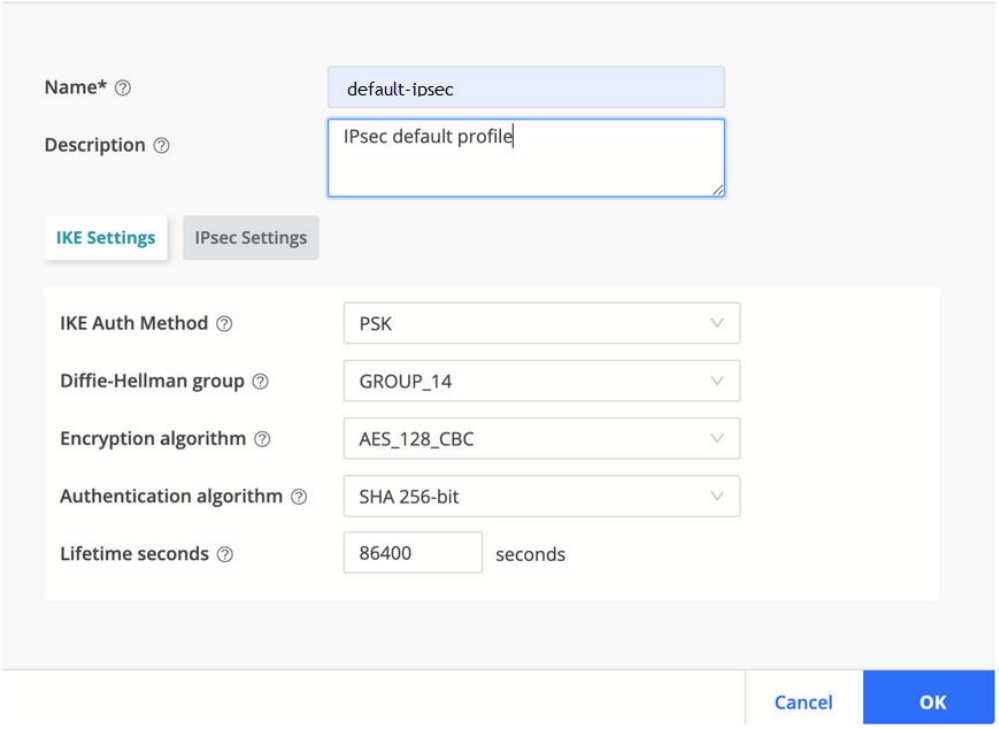

You use this IPsec profile to create a site in the next task. On the Create Site page, if you select IPsec as the tunnel type on the Traffic Forwarding tab, you will attach the IPsec profile.

Create a Site in Juniper Secure Edge Cloud

You have now created IPsec profiles. These profiles define the parameters for the IPsec tunnel between WAN edge devices on your Juniper Mist™ cloud network and your Secure Edge instance.

At this point, you need to create a site in Juniper Security Director Cloud. A site represents a location that hosts a WAN edge device. The traffic from the WAN edge device is forwarded to the Secure Edge instance through a secure tunnel, and then inspected and enforced by the Secure Edge cloud services.

You can configure to forward some or all of the Internet-bound traffic from customer sites to the Juniper Secure Edge cloud through generic routing encapsulation (GRE) or IPsec tunnels from the WAN edge devices at the site.

NOTE: Overlapping branch addresses are not supported to the same POP within Secure Edge when using a stateful firewall at branch locations. Reverse path traffic to these overlapping IPs will be routed using equal-cost multipath (ECMP) across all connections. Traffic is routed using ECMP rather than per-session routing to the interface from which traffic originated. Consider reverse path traffic through ECMP when you configure the protected networks for a site.

To create a site in Juniper Security Director Cloud:

1. In Juniper Security Director Cloud portal, select **Secure Edge** >**Service Management** > **Sites**.
2. Click the plus (**+**) icon to create a site.
3. Enter the settings:
   - Enter a unique site name and a description.
   - Select the corresponding country from the list where the site is located.
   - (Optional) Enter the zip code where the customer branch is located.
   - (Optional) Enter the location (street address) of the site.
   - Select the number of users who can use the network at the site.
   - In the Protected networks field, click the Add (**+**) icon to add the private IP address range of the interface to be used for traffic flow through the tunnel.

   This example shows a site in Juniper Secure Edge Cloud.

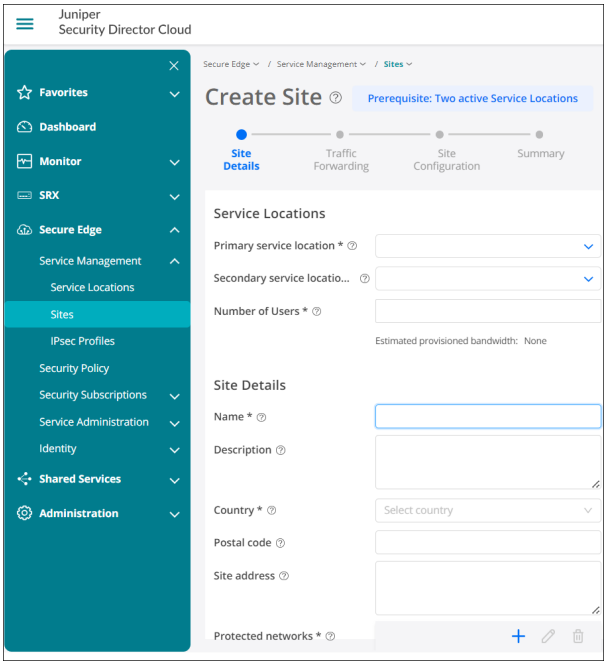

This table provides more information.

Table 49: Site-Creation Details

Fields	Values
Primary service location	jsec-oregon
Secondary service location	jsec-ohio
Number of Users	10
Name	spoke1-site
Country	Germany
Protected networks	10.99.99.0/24 (LAN network)

4. Click **Next**.
5. On the Traffic Forwarding page, enter the settings.

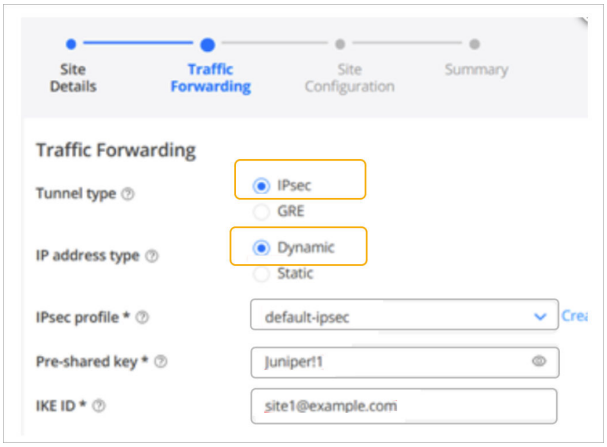

Table 50: Details for Traffic Forwarding Policy

Field	Value
Tunnel type	IPsec
IP address type	Dynamic For the Static IP address type, you need to provide the device IP address in the Site IP address field.
IPsec profile	default-ipsec If you do not have a preconfigured IPsec profile, click Create IPsec Profile to create an IPsec profile.
Pre-shared key	Define a unique PSK for each site. Example: Juniper!1
IKE ID	site1@example.com (resembles an email address and must be a unique value for each site).

6. Click **Next**.
7. On the Site Configuration page, for the **Device Type** select **Non-Juniper Device**.

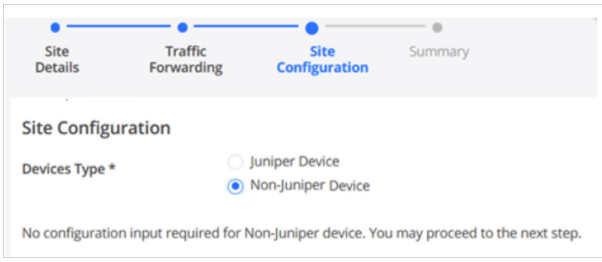

You must select this option because the devices that the Juniper Mist cloud portal manages do not have their configuration pushed through Juniper Security Director Cloud.

8. Click **Next**.
9. On the Summary page, review the configuration.

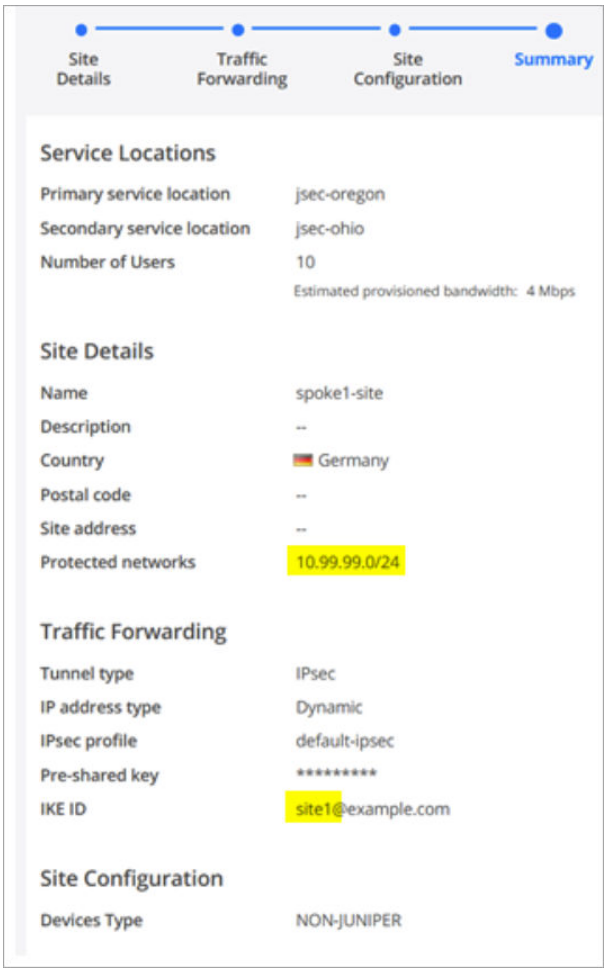

10. Click **Back** to edit any fields or **Finish** to create the new site.
11. Add two more sites using the same procedure. The following paragraphs describe the details to include in each site.

    a. Create a second site with the details provided below.

    Table 51: Site Creation for Second Site

Fields	Value
Primary service location	jsec-oregon
Secondary service location	jsec-oregon

Table 51: Site Creation for Second Site *(Continued)*

Fields	Value
Number of Users	10
Name	spoke2-site
Country	Germany
Protected networks	10.88.88.0/24 (LAN network)

Table 52: Traffic Forwarding for Second Site

Field	Value
Tunnel type	IPsec
IP address type	Dynamic
IPsec profile	default-ipsec
Pre-shared key	Define a unique PSK for each site. Example: Juniper!1
IKE ID	site2@example.com (resembles an email address and must be a unique value for each site).

b. Select **Devices Type**=**Non-Juniper Device** .

c. Create a third site with details provided below.

Table 53: Create a Third Site: Site Details

Fields	Value
Primary service location	jsec-oregon
Secondary service location	jsec-ohio
Number of Users	10

Table 53: Create a Third Site: Site Details *(Continued)*

Fields	Value
Name	spoke3-site
Country	Germany
Protected networks	10.77.77.0/24 (LAN network)

Table 54: Create a Third Site: Traffic-Forwarding Details

Field	Value
Tunnel type	IPsec
IP address type	Dynamic
IPsec profile	default-ipsec
Pre-shared key	Define a unique PSK for each site. Example: Juniper!1
IKE ID	site3@example.com (Resembles an email address and must be a unique value for each site).

d. Select **Devices Type**=**Non-Juniper Device** .

12. Review the **Summary** page. Modify any incorrect entries.

Figure 50 on page 239 displays the list of sites you created.

Figure 50: Summary of Created Sites

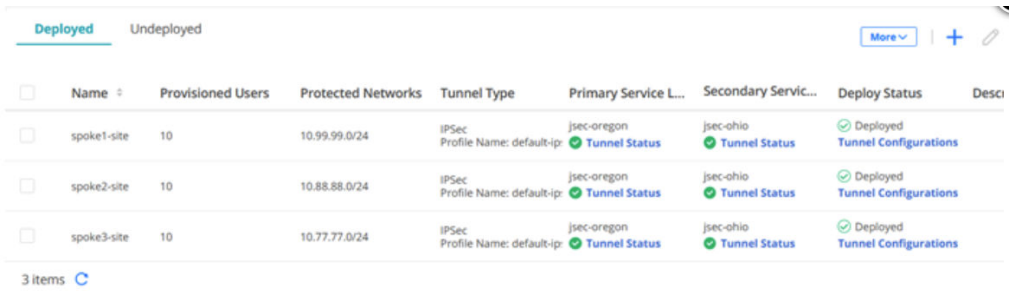

Deploy a Secure Edge Policy in Juniper Security Director Cloud

Now that you have created sites in Juniper Security Director Cloud, its time to deploy one or more Juniper® Secure Edge policies.

Secure Edge policies specify how the network routes traffic. By default, when you create a new tenant, the Security Director Cloud creates a Secure Edge policy rule set with predefined rules.

NOTE: Even if you do not change the default rule set, you must use the **Deploy** option to load the rules in your service locations.

To deploy a Secure Edge policy in Juniper Security Director Cloud:

1. In Juniper Security Director Cloud portal, click **Secure Edge** > **Security Policies**.

   A Secure Edge Policy page with default rules appears. You modify the default security policy set for better debugging. The default rule set does not allow ICMP pings to the outside (Internet), preventing you from pinging anything through the cloud.

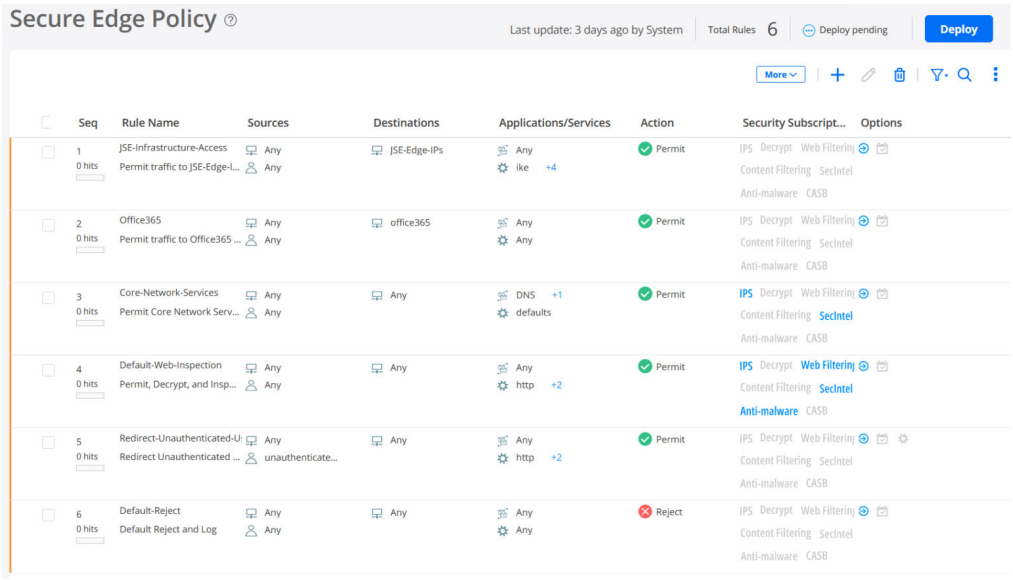

2. Click the Add (**+**) icon to create a rule, or select the existing rule and click the pencil icon to edit the rule.
3. Give the new rule the **Rule Name**=Allow-ICMP.
4. Click Add (**+**) to add sources.

   Under **Sources**, use the following default values:

   - **Addresses**=Any
   - **User Groups**=Any
5. Click Add (**+**) to add destinations.

Under **Destinations**, for **Addresses**, use the default value =Any.

6. Under **Applications/Services**, configure the following values:

   - **Applications**=Any
   - **Services**=Specific (via search)
   - **Specific Service**=icmp-all

   Using the Right Arrow (>), move **specific service=icmp-all** to the right pane to activate it before you click **OK**.

7. Configure **Action**=Permit, and retain the default values for the remaining fields.

   The system places the new rule at the bottom of the rules list and treats this rule as the last rule in the rule set. If the rule is placed after a global rule (that denies all traffic), it will never get applied, because the global rule stops all further traffic. Therefore, for this example you change the position of the rule by selecting the rule. Then, use the **Move** > **Move** > **Move Top** options to move the selected rule to the top of the rule set. Moving the rule to the top of the rule set ensures that the system applies this rule first.

   

   NOTE: Whenever you modify a rule set, ensure that you use the **Deploy** button to complete the task. Otherwise, service locations continue to use the outdated rule sets.

8. Click **Deploy**.
9. On the **Deploy** page, check the **Run now** option and click **OK**.

   Service locations get the updated rule set after few minutes.

10. Select **Administration** > **Jobs** to view the status and progress of the deployed job.

Get IPsec Tunnel Configuration Parameters to Apply in Juniper Security Director Cloud

In the preceding tasks, you completed several actions to set up an IPsec tunnels in Juniper Secure Edge and have deployed the Secure Edge policy in Juniper Security Director Cloud. The final step in Security Director Cloud is to collect configuration data for each site. You'll need these details to complete the secure edge connector configuration ("Create Secure Edge Connectors in the Juniper Mist Portal" on page 244) in the Juniper Mist™ cloud to set up an IPsec tunnel. In this step, you'll note down the details of the sites you created.

NOTE: An automated configuration push to synchronize between Juniper Security Director Cloud and Juniper Mist cloud option not available.

To get IPsec tunnel configuration parameters to apply in Juniper Security Director Cloud:

1. In Juniper Security Director Cloud portal, select **Secure Edge** >**Service Management** > **Sites**.

   The Site page opens, displaying deployed site details.

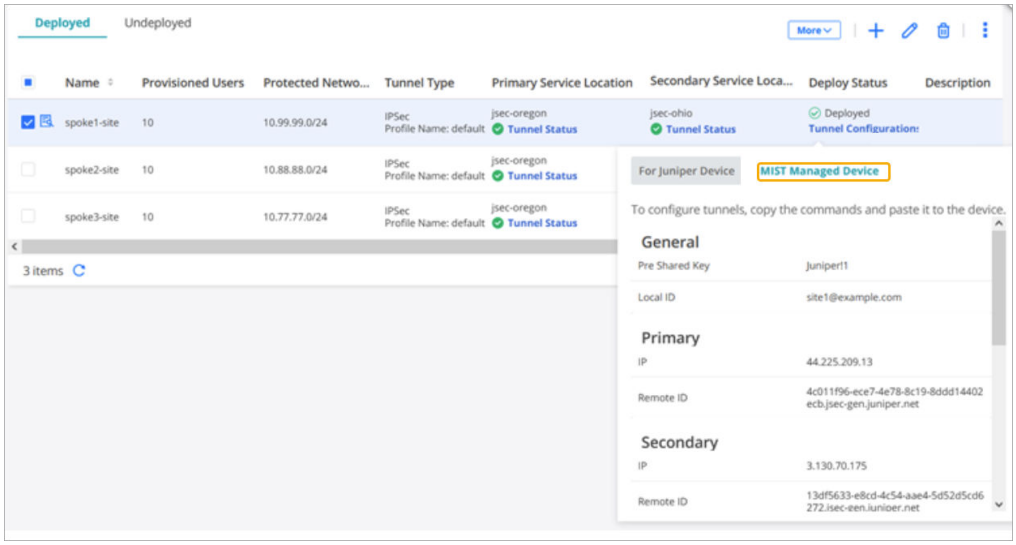

2. For each spoke site, click the **Tunnel Configuration** option under **Deployed Status**, and then check the **MIST Managed Device** tab for information.

   Note down the following details, which you will use in "Create Secure Edge Connectors in the Juniper Mist Portal" on page 244:

   - Pre-Shared Key
   - Local ID
   - IP address and remote ID of each service location tunnel

   The following samples show extracted information for all three sites you created in "Create a Site in Juniper Secure Edge Cloud" on page 233:

```
General spoke1-site

Pre Shared Key    abc!1
Local ID  site1@example.com

Primary
IP        44.225.209.13
```

```
Remote ID xxx123-exx7-xxyys-8c19-abcd123.123.juniper.net

Secondary
IP          3.130.70.175
Remote ID abc123-exx7-xxyys-8c19-abcd123.123.juniper.net


IKE Proposals
Authentication Method pre-shared-keys
DH group  group14
Encryption algorithm  aes-128-gcm
Lifetime  86400

IPsec Proposals
Protocol  esp
Encryption algorithm  aes-128-gcm
Lifetime  3600
PFS Group group14
```

The following sample is of the extracted information for site2:

```
General spoke2-site

Pre Shared Key abc!1
Local ID site2@example.com

Primary
IP 3.130.70.175
Remote ID xxx123-exx7-xxyys-8c19-abcd123.123.juniper.net


Secondary
IP 44.225.209.13
Remote ID abc123-exx7-xxyys-8c19-abcd123.123.juniper.net

IKE Proposals
Authentication Method pre-shared-keys
DH group group14
Encryption algorithm aes-128-gcm
Lifetime 86400

IPsec Proposals
```

```
Protocol esp
Encryption algorithm aes-128-gcm
Lifetime 3600
PFS Group group14
```

The following sample is of the extracted information for site3:

```
General spoke3-site

Pre Shared Key abc!1
Local ID site3@example.com

Primary
IP 44.225.209.13
Remote ID xxx123-exx7-xxyys-8c19-abcd123.123.juniper.net


Secondary
IP 3.130.70.175
Remote ID abc123-exx7-xxyys-8c19-abcd123.123.juniper.net


IKE Proposals
Authentication Method pre-shared-keys
DH group group14
Encryption algorithm aes-128-gcm
Lifetime 86400

IPsec Proposals
Protocol esp
Encryption algorithm aes-128-gcm
Lifetime 3600
PFS Group group14
```

You need these site details when you configure tunnels in the Mist cloud portal.

Create Secure Edge Connectors in the Juniper Mist Portal

You are about halfway to your ultimate goal of setting up a Secure Edge connector for the Session Smart Routers or SRX Series Firewalls in Juniper Mist™.

You create Secure Edge connectors in the Juniper Mist portal. This task completes the configuration on the Mist cloud side of the tunnels to establish an IPsec tunnel between WAN edge devices managed by Mist and Security Director Cloud. Before you create the connectors, ensure that your site has a deployed WAN edge device.

To create Secure Edge connectors in the Juniper Mist portal:

1. From the left menu, click **WAN Edges**.

   The WAN Edges page displays site details.

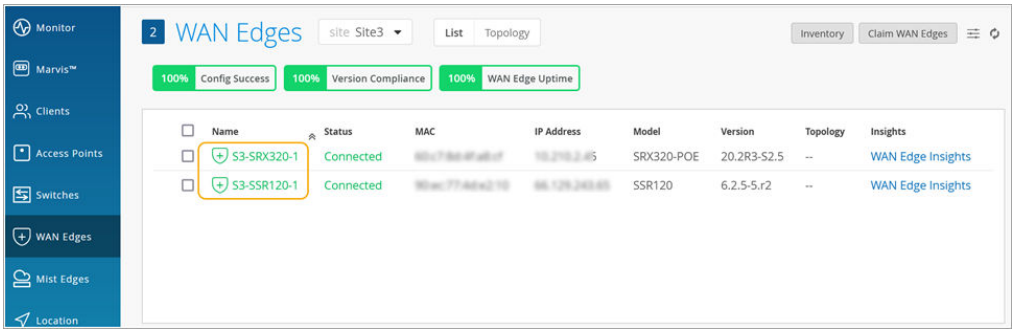

2. Click the device and scroll down to Secure Edge Connectors.

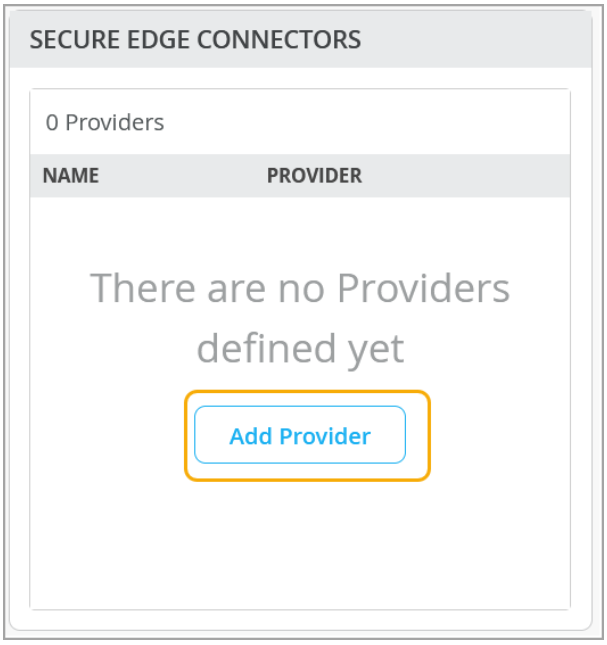

3. In the **Secure Edge Connectors** pane, click **Add Provider**.
4. Enter Secure Edge connector details.

> **NOTE**: Remember that these are same the details you gathered in "Get IPsec Tunnel Configuration Parameters to Apply in Juniper Security Director Cloud" on page 241.

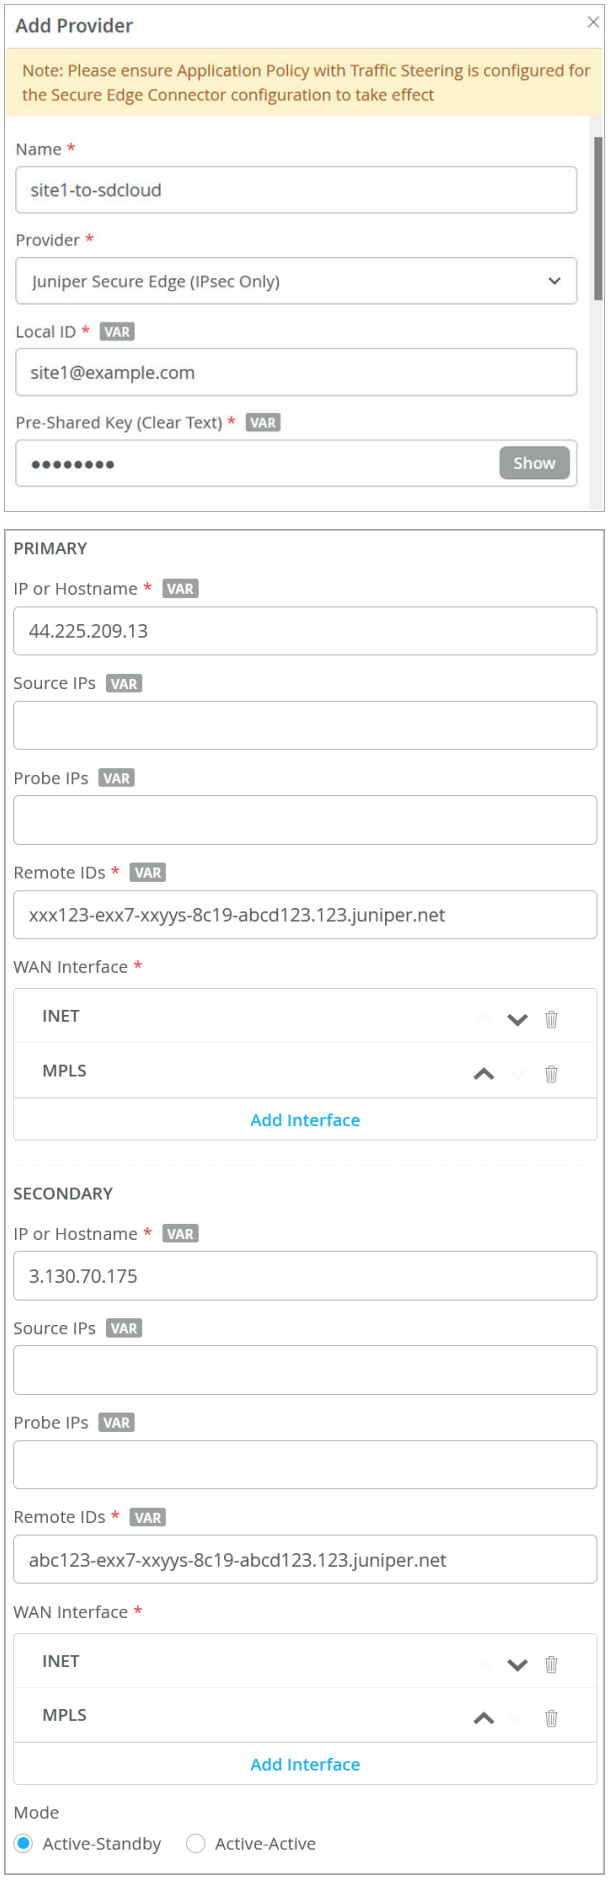
Add Provider
Note: Please ensure Application Policy with Traffic Steering is configured for the Secure Edge Connector configuration to take effect
Name *
site1-to-sdcloud
Provider *
Juniper Secure Edge (IPsec Only)
Local ID * VAR
site1@example.com
Pre-Shared Key (Clear Text) * VAR
Show
PRIMARY
IP or Hostname * VAR
44.225.209.13
Source IPs VAR
Probe IPs VAR
Remote IDs * VAR
xxx123-exx7-xxyys-8c19-abcd123.123.juniper.net
WAN Interface *
INET
MPLS
Add Interface
SECONDARY
IP or Hostname * VAR
3.130.70.175
Source IPs VAR
Probe IPs VAR
Remote IDs * VAR
abc123-exx7-xxyys-8c19-abcd123.123.juniper.net
WAN Interface *
INET
MPLS
Add Interface
Mode
Active-Standby
Active-Active

Table 55: Secure Edge Connector Details

Field	Value
Name	site1-to-sdcloud
Provider	Juniper Secure Edge
Local ID	site1@example.com
Pre-Shared Key	Juniper!1 (example)
Primary	
IP or Hostname	<IP address> (from Juniper Security Director Cloud tunnel configuration)
Probe IPs	-
Remote ID	<UUID>.jsec-gen.juniper.net (from Juniper Security Director Cloud tunnel configuration)
WAN Interface	• WAN0=INET • WAN1=MPLS
Secondary	
IP or Hostname	<IP address> from (From Juniper Security Director Cloud tunnel configuration)
Probe IPs	-
Remote ID	<UUID>.jsec-gen.juniper.net (from Juniper Security Director Cloud tunnel configuration)
WAN Interface	• WAN0=INET • WAN1=MPLS

Table 55: Secure Edge Connector Details *(Continued)*

Field	Value
Mode	Active-standby

NOTE: You don't need to enter the probe IP values. IPsec tunnels do not need additional monitoring like GRE needs.

NOTE: Do not enable ICMP **Probe IPs** for Session Smart Router-based Secure Edge configuration. ICMP probes will be sourced from a nonroutable IP address toward the Secure Edge and dropped due to policy. In addition, if the source addresses are overlapping at all branches, routing to more than one branch with a probe IP address is not supported.

NOTE: The system generates text, application, and email descriptions automatically.

5. Verify that the Mist cloud portal has added the Secure Edge connector you just configured.

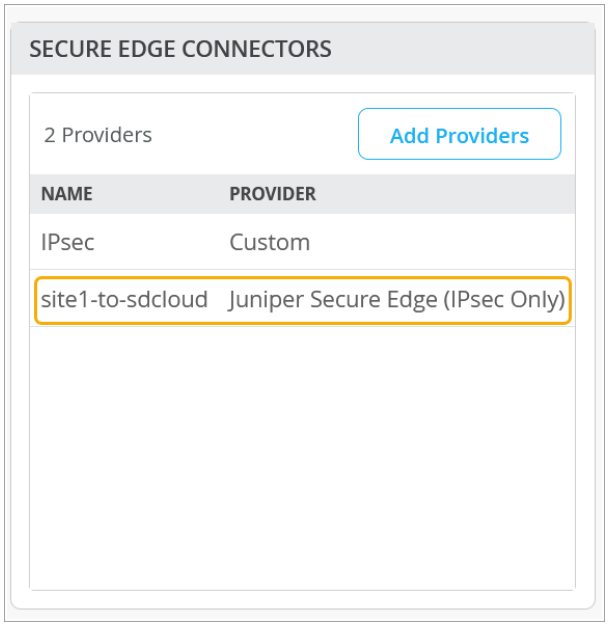

6. Add a new traffic-steering path on the WAN edge template or WAN edge device.

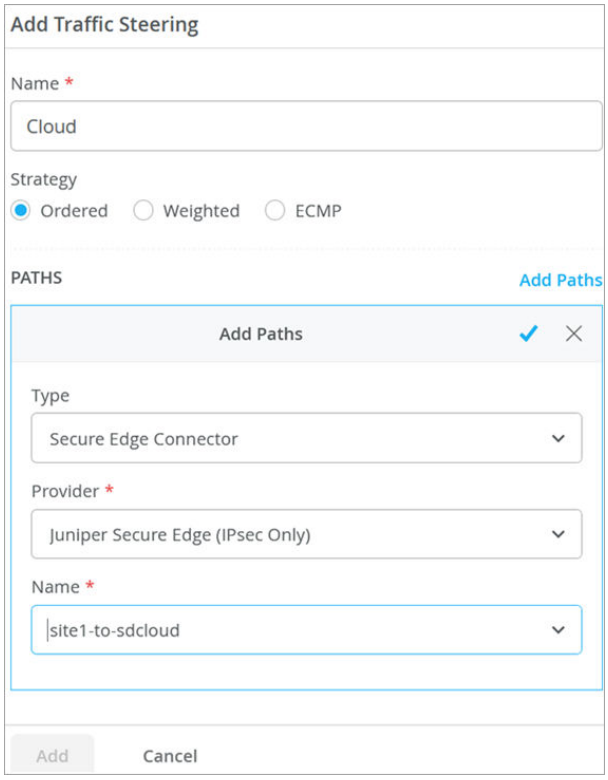

Table 56: Traffic-Steering Path Configuration

Fields	Value
Name	Cloud
Strategy	Ordered
Paths	Select Type and Destination
Type	Secure Edge Connector
Provider	Juniper Secure Edge
Name	site1-to-sdcloud

Modify an Application Policy

After you create Secure Edge connectors in the Juniper Mist™ cloud portal, next step is to modify application policies on the branch device. For example, you can allow traffic from a spoke device to a hub device. You can also allow traffic from a spoke device to another spoke device in the VPN tunnel. After that, you can send traffic from spokes to the Internet through Juniper Security Director Cloud instead of sending traffic from spokes to a hub for central breakout.

Use the following steps to confirm if the configuration is working:

1. Add or edit an Application Policy on the WAN edge template or WAN edge device page.
2. Select the policy that you want to modify, and apply the changes.

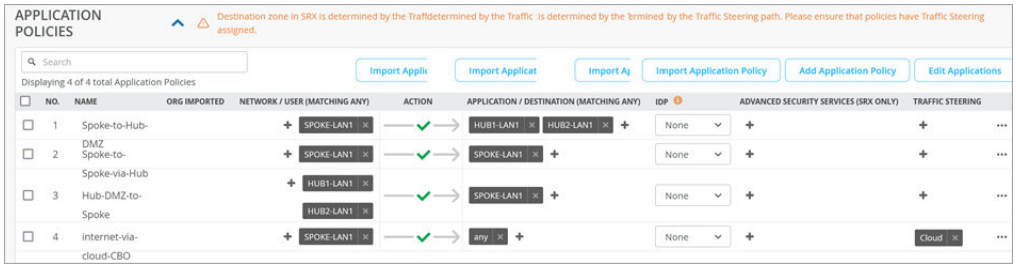

If you are creating policies from the WAN edge device page, you ay want to select the **Override Template Settings** option as per requirement.

- In the application policy, you can include the traffic steering you have created in the previous step. In this example, change the **Traffic Steering** to **Cloud** in the last rule (Internet-via-Cloud-CBO).

3. Save the changes.

   Juniper Mist cloud builds new tunnels to Juniper Security Director Cloud.

Verify the Configuration

After you modify the application policy, now it is time to confirm that your configuration is working as expected. With the desired configuration saved, you can verify if Juniper Mist cloud routes the Internet-bound traffic from spokes to Juniper Security Director Cloud instead of routing it to a hub for central breakout.

To verify the configuration:

1. (Optional) Depending on your environment, you can see the communication of the IPsec tunnel towards Juniper Security Director Cloud in CLI.

```
user@host:~# tcpdump -eni fabric6 port 4500
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
```

```
listening on fabric6, link-type EN10MB (Ethernet), capture size 262144 bytes
18:43:46.835469 52:54:00:f4:02:77 > 52:54:00:14:07:6c, ethertype IPv4 (0x0800), length 317:
192.168.173.191.16534 > 44.225.209.13.4500: NONESP-encap: isakmp: child_sa ikev2_auth[I]
18:43:46.879282 52:54:00:f4:02:77 > 52:54:00:14:07:6c, ethertype IPv4 (0x0800), length 317:
192.168.173.191.16535 > 3.130.70.175.4500: NONESP-encap: isakmp: child_sa ikev2_auth[I]
18:43:46.884834 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 292:
44.225.209.13.4500 > 192.168.173.191.16534: NONESP-encap: isakmp: child_sa ikev2_auth[R]
18:43:46.974426 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 292:
3.130.70.175.4500 > 192.168.173.191.16535: NONESP-encap: isakmp: child_sa ikev2_auth[R]
18:43:58.001576 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 103:
44.225.209.13.4500 > 192.168.173.191.16534: NONESP-encap: isakmp: parent_sa inf2
18:43:58.002603 52:54:00:f4:02:77 > 52:54:00:14:07:6c, ethertype IPv4 (0x0800), length 103:
192.168.173.191.16534 > 44.225.209.13.4500: NONESP-encap: isakmp: parent_sa inf2[IR]
18:44:06.111512 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 103:
3.130.70.175.4500 > 192.168.173.191.16535: NONESP-encap: isakmp: parent_sa inf2
18:44:06.112368 52:54:00:f4:02:77 > 52:54:00:14:07:6c, ethertype IPv4 (0x0800), length 103:
192.168.173.191.16535 > 3.130.70.175.4500: NONESP-encap: isakmp: parent_sa inf2[IR]
18:44:06.896312 52:54:00:f4:02:77 > 52:54:00:14:07:6c, ethertype IPv4 (0x0800), length 103:
192.168.173.191.16534 > 44.225.209.13.4500: NONESP-encap: isakmp: child_sa inf2[I]
18:44:06.922069 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 103:
44.225.209.13.4500 > 192.168.173.191.16534: NONESP-encap: isakmp: child_sa inf2[R]
18:44:07.022463 52:54:00:f4:02:77 > 52:54:00:14:07:6c, ethertype IPv4 (0x0800), length 103:
192.168.173.191.16535 > 3.130.70.175.4500: NONESP-encap: isakmp: child_sa inf2[I]
18:44:07.022502 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 43:
44.225.209.13.4500 > 192.168.173.191.16534: isakmp-nat-keep-alive
18:44:07.097695 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 103:
3.130.70.175.4500 > 192.168.173.191.16535: NONESP-encap: isakmp: child_sa inf2[R]
18:44:07.113678 52:54:00:14:07:6c > 52:54:00:f4:02:77, ethertype IPv4 (0x0800), length 43:
3.130.70.175.4500 > 192.168.173.191.16535: isakmp-nat-keep-alive
```

Verify the established tunnels details WAN Insights of the device in Juniper Mist cloud portal.

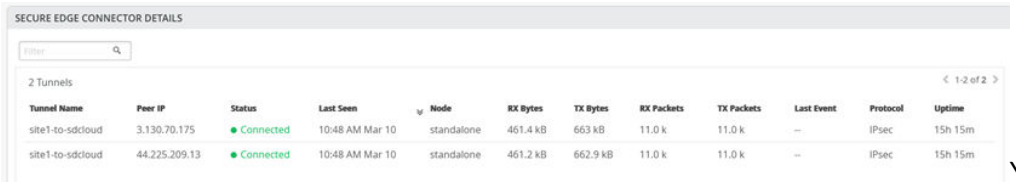

You can also check the established tunnels in the Juniper Security Director Cloud dashboard and in the service location.

2. Check the new traffic flow using a VM desktop connected to the branch device. You can verify the traffic flow by using pings to the Internet.

```
root@desktop1:~# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=100 time=40.5 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=100 time=39.2 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=100 time=38.0 ms
64 bytes from 8.8.8.8: icmp_seq=4 ttl=100 time=37.9 ms
^C
--- 8.8.8.8 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3004ms
rtt min/avg/max/mdev = 37.942/38.914/40.507/1.059 ms

root@desktop1:~# ping 9.9.9.9
PING 9.9.9.9 (9.9.9.9) 56(84) bytes of data.
64 bytes from 9.9.9.9: icmp_seq=1 ttl=41 time=35.8 ms
64 bytes from 9.9.9.9: icmp_seq=2 ttl=41 time=34.3 ms
64 bytes from 9.9.9.9: icmp_seq=3 ttl=41 time=34.9 ms
64 bytes from 9.9.9.9: icmp_seq=4 ttl=41 time=33.7 ms
^C
--- 9.9.9.9 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3001ms
rtt min/avg/max/mdev = 33.722/34.686/35.824/0.782 ms
```

NOTE: You may experience latency depending on the physical distance between your WAN edge device and Juniper Secure Edge service location.

3. Open a browser on a VM desktop and navigate to https://whatismyipaddress.com/ to view details about the source IP address used to route the Juniper Mist network traffic from a service location towards the Internet.

   This example shows traffic from a primary service location.

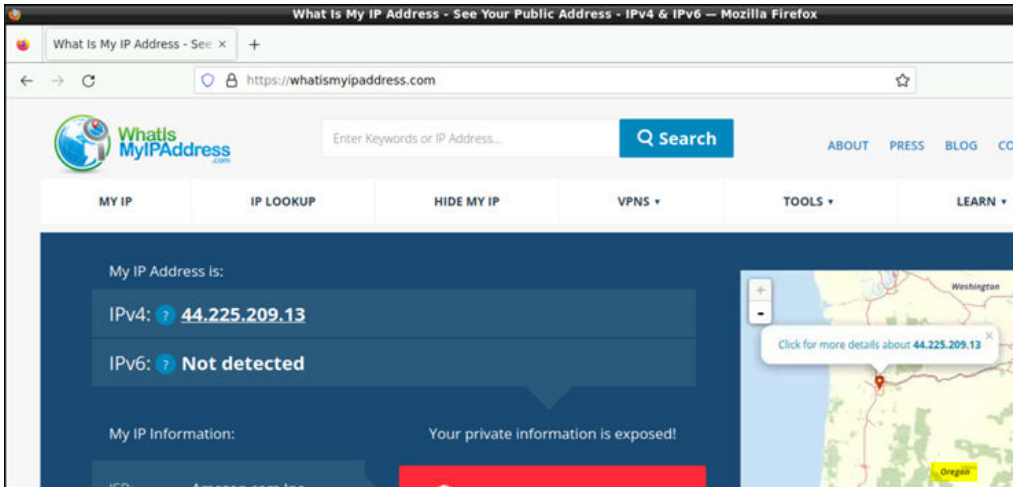

This example shows traffic from a secondary service location.

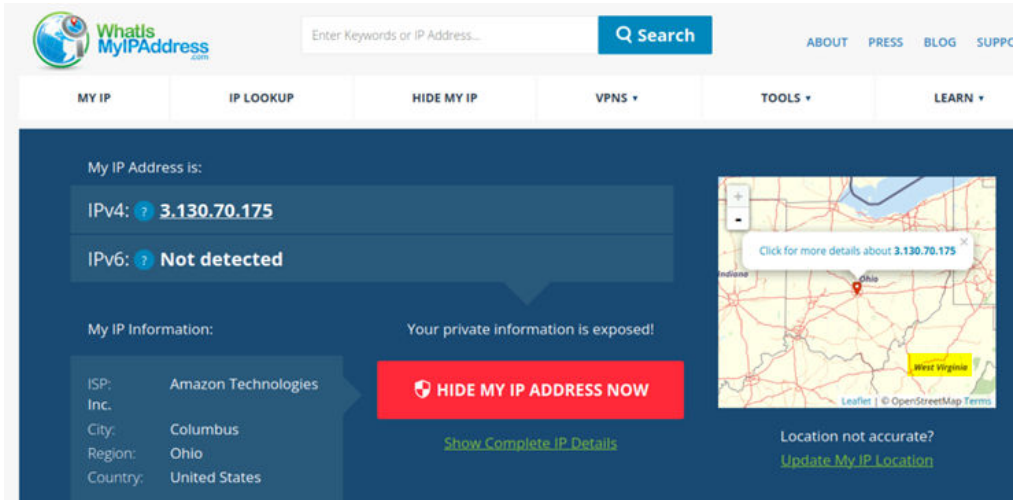

One of the two IP addresses of the service location is a public IP address and serves two purposes:

- Terminates the IPsec tunnel
- Routes traffic from branch devices to the Internet through Juniper Security Director Cloud

You can view this same public IP address in the packet captures showing established tunnel to the service location using Juniper Security Director Cloud. See "Verify the Configuration" on page 250.

Remember that a service location in Juniper Security Director Cloud is also known as POP and represents a Juniper® Secure Edge instance in a cloud location. The service location is the connection (access) point for both on-premises and roaming users.

Setup Secure Edge Connector with Juniper Secure Edge (Auto-Provisioning)

SUMMARY

Follow this workflow to set up and verify auto-provisioning for Juniper® Secure Edge.

IN THIS SECTION

Mist now offers automated Juniper Secure Edge connector tunnel provisioning. This feature allows you to effortlessly establish connections using predefined settings.

Prerequisites

- Activate Juniper Secure Edge account and check licenses, subscriptions, certificates. See "Access Juniper Security Director Cloud and Check Active Subscriptions" on page 228.
- Launch the required number of service locations (with required capacity). See "Configure Service Locations" on page 229.

Configure Secure Edge Connector Auto-Provisioning

IN THIS SECTION

Watch the following video to understand how to setup Secure Edge Connector auto provisioning:

Video: Auto Provision JSE for Session Smart Routers

Add Juniper Secure Edge Connector Credentials in Juniper Mist Portal

1. Provide Juniper Secure Edge credential details in Juniper Mist portal.
   - On Juniper Mist portal, select **Organization > Settings**.
   - Scroll-down to **Secure WAN Edge Integration** pane and click **Add Credentials**.
   - In **Add Provider** window, enter the details.

     Figure 51: Add Credentials for Juniper Secure Edge

     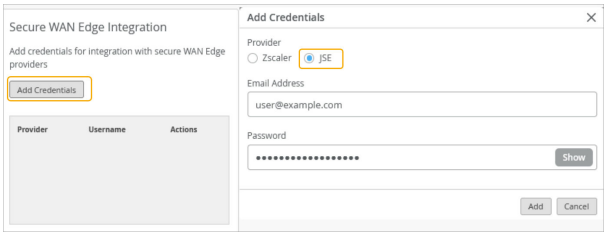

     - Provider—Select JSE.
     - Email Address—Enter user name (email address) (Credentials of the user created on the Juniper Secure Edge portal)
     - Password—Enter password for the user name.
   - Click **Add** to continue.

Configure Juniper Secure Edge Tunnel Auto-Provisioning

1. On Juniper Mist portal, go to Organization > WAN Edge Templates and click an existing template.
2. Scroll-down to **Secure Edge Connector**.
3. Click **Add Providers**

Figure 52: Add Provider

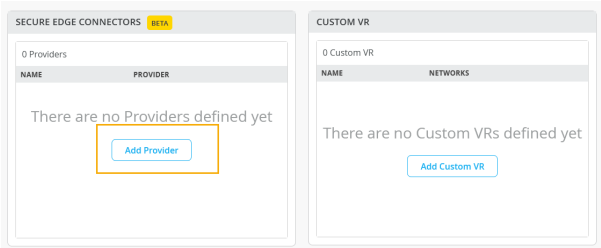

4. In **Add Provider** window, select **Juniper Secure Edge (Auto)** for automatic provisioning.

Figure 53: Select Juniper Secure Edge as Provider

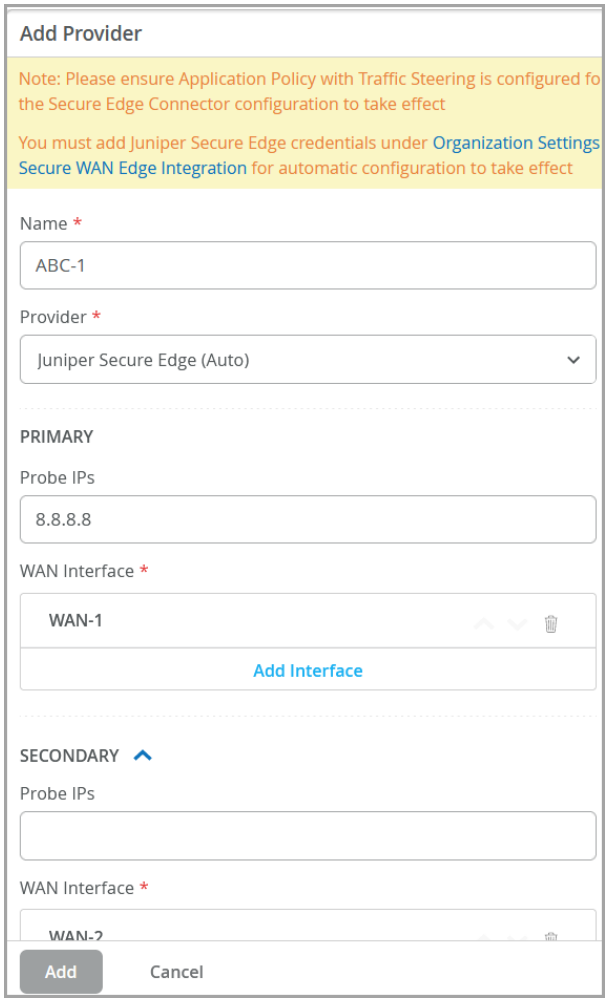

Enter the following details:

- Name—Enter a name for the JSE tunnel.

- Provider—Select Juniper Secure Edge (Auto).
- Probe IP—Enter probe IPs (primary and secondary). Enter probe IP 8.8.8.8 or any other well-known probe IP address.
- WAN Interface—Assign WAN interfaces under primary and secondary tunnel details for provisioning of primary and secondary tunnels.

5. Click **Add**.
6. In the **Secure Edge Connector Auto Provision Settings**enter the details. This option is available only if you have configured Juniper Secure Edge as provider in the previous step.

Figure 54: Secure Edge Connector Auto Provision Settings

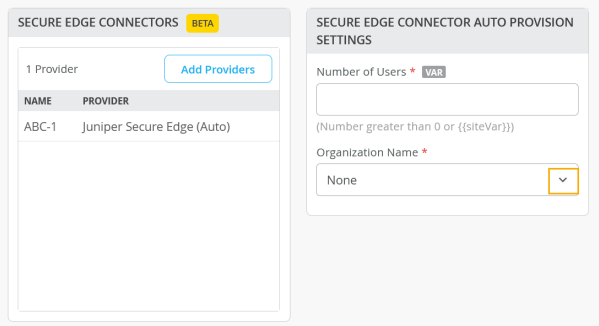

- Number of Users—Enter the maximum number of users supported by the JSE tunnel
- Organization Name—Enter the organization name. The drop-down box displays all organizations associated with the user name in Juniper Secure Edge account. This is the same user name that you have entered in Juniper Secure Edge credential in **Organization > Settings**. See step 1 for details.

7. Click **Add** to continue.

When you assign a template enabled with the Juniper Secure Edge (Auto) option to a site, an associated JSE site (location object) is automatically created and a tunnel from the device to the closest network point of presence (POP) is brought up.

For the Secure Edge Connector configuration to take effect, you must create an application policy with Mist Secure Edge Connector-to-Juniper Secure Edge traffic steering.

SEE ALSO

No Link Title

Verify Juniper Secure Edge Tunnels

On Juniper Mist portal, you can verify the established tunnels details in WAN Insights of the device once **WAN Edge Tunnel Auto Provision Succeeded** event appears under WAN Edge Events.

Figure 55: WAN Edge Events

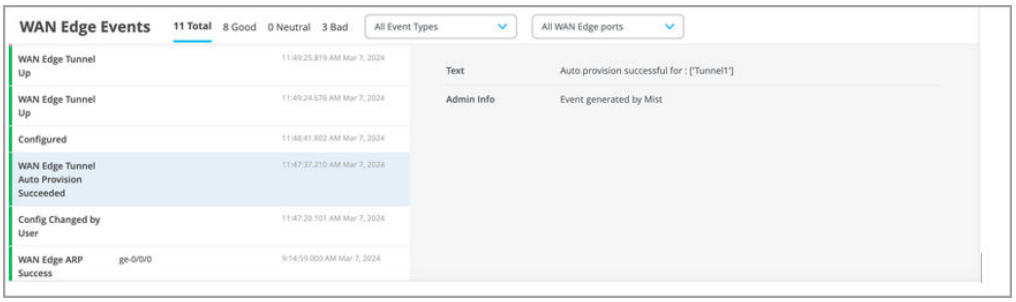

Get the established tunnels status details in **WAN Edges > WAN Edge Insights** page Juniper Mist cloud portal.

Figure 56: Established Secure Edge Tunnels

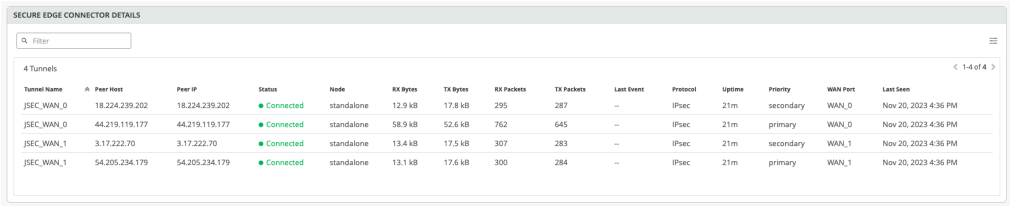

You can check the established tunnels in the Juniper Security Director Cloud dashboard and in the service location.

Use Custom Options to Configure Secure Edge Connector

SUMMARY

Use custom options to configure tunnel provisioning or to support site-to-site VPN.

IN THIS SECTION

Juniper Mist™ offers custom option for tunnel provisioning. With minimal configuration, your WAN Edge device can establish connections to the SSE using either IPsec or GRE protocols.

Configure Tunnel Provisioning

Before You Begin: Ensure you have the local and remote network account details on hand.

To configure tunnel:

1. From the left menu, navigate to a WAN Edge template, hub profile, or device.
2. Scroll to the **Secure Edge Connector** section, click **Add Provider**, and enter the seetings:
   - **Name**—Enter the name of the service.
   - **Provider**—Select **Custom**.
   - **Remote Network**—Select an existing Network or create a network.
   - **Protocol**—Select **IPsec** or **GRE**. Then enter the settings for the selected protocol.

Table 57: Settings

Field	Description
Local ID *(IPsec only)*	Enter the login ID for the local account.

Table 57: Settings *(Continued)*

Field	Description
Pre-Shared Key (Clear Text) *(IPsec only)*	Enter the preshared key (PSK) for the local account. The length of the PSK must be between 6-255 characters.
IP or Hostname	Enter the IP address or hostname.
Source IP	Enter the Source IP address of the tunnel.
Probe IPs	Enter probe IP address. You can use any well-known IP (Example: 8.8.8.8).
Remote ID *(IPsec only)*	Provide login ID of the remote account.
WAN Interface	Add one or more WAN interfaces to provision of primary and secondary tunnels. If you add multiple WAN interfaces, the first interface takes the priority. If first interface is down, then system uses the second interface to establish the tunnel. When you click Add Interface, choos from the list of WANs that have been configured for the selected template, hub profile, or device.
IKEv2 proposal *(IPsec only)*	Retain default values or click Add Proposal. Then select enter the settings.
Lifetime	Enter a value between 180 to 86400 seconds.

3. Click **Add** at the bottom of the provider panel.
4. Scroll down to the Routing section, click **Add BGP Group**, and enter the settings.

   Tips:

   - For the Peering Network, select the same SEC provider that you created in previous steps.
   - For Local AS, enter AS number or non-default AS for WAN Edge.
   - If you selected the GRE protocol, configure the BGP group as follows:
   - • Name: Give the BGP group a name, such as **BGP-over-GRE**

- Peering Network: Choose **SEC Tunnel**, and then select the tunnel you configured in step iv, above.
- Select **Advertise to the Overlay**.
- BDF: Choose **Disabled**.
- Type: Choose **External**.
- Local AS: Type the number of the AS you are using, for example, 65000.
- Hold Time: Specify a time, in seconds such as 90.
- Graceful Restart Time: Specify a time, in seconds such as 120.

- In the Neighbors section, click **Add Neighbors**. Add BGP peer IP address of SSE and AS value.
- Optionally, you can add BGP policy for import or export of routes.

For help with other BGP settings, see "BGP" on page 178.

5. Save the BGP group.
6. Scroll to the **Traffic Steering** section, click **Add Traffic Steering**, and enter the settings.
   - **Name**—Enter a name for the traffic-steering profile.
   - **Strategy**—Select a strategy. You can configure the traffic steering profile with any strategy (**Ordered** or **Weighted** or **ECMP**), based on your topology and configuration.
   - **Path**—Click **Add Paths** and enter the following details.
     - **Type**—Select **Secure Edge Connector**.
     - **Provider**—Select **Custom**.
     - **Name**—Select the custom connector's name you have created in previous step.
7. Save the traffic steering policy.
8. Click **Save** at the top-right corner of the page to save the entire configuration.
9. Add an application policy.

   Application policy allows the desired network to reach the more specific application using the route table. In the application policy, you can include the remote network you have created in the previous step. Use that network in an application policy to allow inbound access from the Secure Edge Connector. To create the application policy, in the Juniper Mist cloud portal, go to **Organization** > **WAN** > **Application Policy**. For help, see "Application Policies" on page 144.

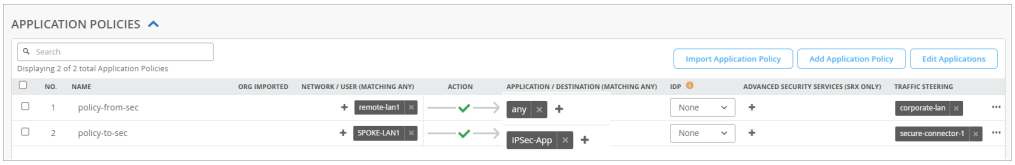

Configure a Site-to-Site VPN

IN THIS SECTION

Support for Site-to-Site VPN

You can set up site-to-site VPN using custom option for tunnel provisioning.

A site-to-site VPN is a secure, software-defined network connection that links two or more remote sites over the internet. This type of VPN is crucial for enterprises looking to connect branch offices, data centers, or other remote locations securely and efficiently.

1. Go to the Juniper Mist portal and navigate to the Secure Edge Connector section at the WAN Edge Templates level, hub profile, or site level.
2. Click **Add Providers** and select the **Custom** option.
3. Enter the necessary details for tunnel provisioning, such as local and remote network account details, IP addresses, and preshared keys
4. Define the IKEv2 and IPsec proposals, including encryption and authentication algorithms, Diffie-Hellman groups, and lifetimes. You must select IKE and IPsec values which match the device on another end of the tunnel.
5. Assign WAN interfaces for primary and secondary tunnels.
6. Create a Traffic Steering profile. This profile defines how traffic is routed through the VPN tunnel. This profile is then used in an Application Policy to apply these settings to specific types of traffic.
7. Create inbound or outbound Application Policies. If you want to allow traffic from the remote network to enter your local network, you need to create a Network representing the remote

network. Attach this network to the Custom Secure Edge Connector (SEC) Provider and use it as the source in an Application Policy.

Verification

On Juniper Mist portal, you can verify the established tunnels details in WAN Insights of the device once **WAN Edge Tunnel Auto Provision Succeeded** event appears under WAN Edge Events.

Once you update the template, the IPsec configuration will be pushed to the WAN Edge device. For first time IPSec deployment, the system takes time to download the software and configuration.

Once the IPSec configuration has been deployed, you can see the IPsec status by navigating to **WAN Edge** > ***WAN Edge Name*** > Secure Edge Connector Details.

You can view BGP neighbor status by navigating to **Monitor** > **Insights** > **WAN Edge**.

To verify the BGP over GRE session you created, you can use the WAN Edge testing tools:

- **WAN Edge** > **Utilities** > **Testing Tools**.

  Open the **BGP > Summary** tab, or **Routes** > **Show Routes**.

Palo Alto Prisma Access Integration for SASE Health (BETA)

SUMMARY

You can integrate a Palo Alto Prisma Access account with Juniper Mist™ WAN Assurance for single-pane-of-glass troubleshooting directly from the Mist portal.

IN THIS SECTION

You can integrate a Palo Alto Prisma Access account with Juniper Mist™ WAN Assurance to get a single-pane-of-glass view for troubleshooting and debugging. This integration enables you to view Prisma Access events and tunnel statistics in real-time directly from the Mist portal, so you don't have to navigate between platforms while troubleshooting.

Palo Alto Prisma Access is a cloud-based security platform that ensures secure and fast access to applications and data for users regardless of their location. If you use both Prisma Access and Mist WAN Assurance, an IPsec tunnel is used for end-to-end communication between the two. For example, if standalone WAN Edge devices need to reach specific applications, they can do so securely through the IPsec tunnel between the devices and the Prisma cloud.

Mist automatically provisions these IPsec tunnels, which significantly reduces the amount of manual configuration required by you and also reduces risk of human error. This automatic tunnel provisioning means that you do not need to configure on both sides, as the configuration you do in Mist will be replicated in Strata Cloud Manager.

This integration delivers critical insights in real-time for your WAN Edge devices, Prisma Access, and applications. You can use these insights to quickly and efficiently diagnose and resolve site issues.

Figure 57: Traffic Inspection by Juniper Secure Edge

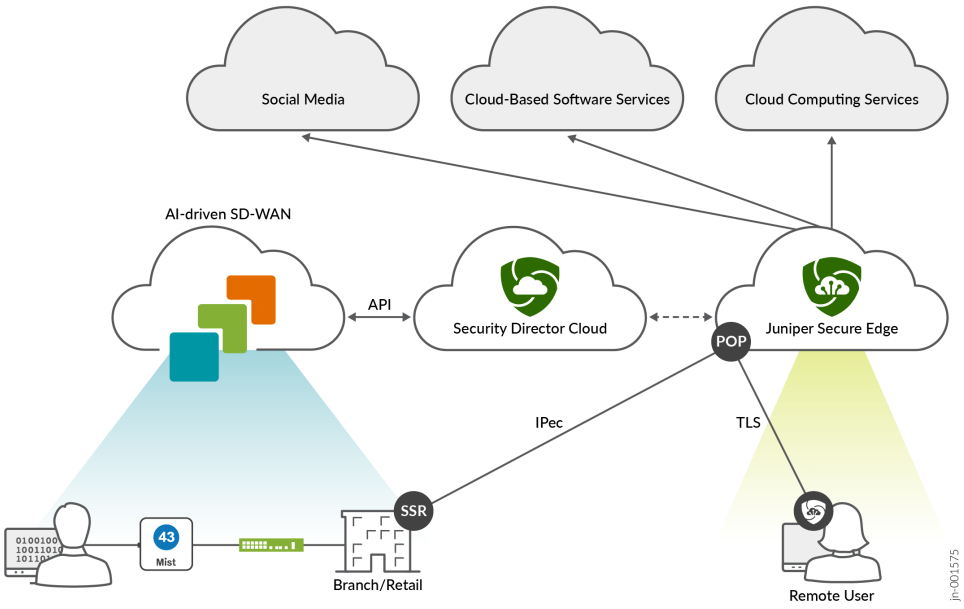

Benefits and Requirements

Table 58: Benefits and Requirements

Benefits	• For Juniper Mist™ WAN Assurance customers:
	  • No additional cost required
	  • No additional subscription required
	  • Minimal configuration required for automatic IPsec tunnel provisioning
	  • The IPsec tunnel configuration in Mist is automatically replicated in the Strata Cloud Manager
	• Single-pane-of-glass for troubleshooting:
	  • View alerts generated by Prisma Access
	  • View tunnel statistics
	  • Root cause analysis
	• Marvis Insights into network health

Requirements	• Juniper Mist™ SD-WAN devices • Juniper Mist™ WAN Assurance subscription • Palo Alto Prisma Access managed by Strata Cloud Manager (SCM) • Palo Alto Prisma Access license • API Keys generated from Palo Alto

To integrate a Prisma Access account with Mist WAN Assurance, you must complete the following steps.

Configure User Access Role in the Prisma Strata Cloud Manager and Obtain Credentials for Mist

To control access to your applications and services, configure a user Access Role in Prisma. You'll set up the role through the Strata Cloud Manager. Then store the user credentials so that they can be entered into the Juniper Mist portal for "account linking" on page 268.

1. In the Strata Cloud Manager, navigate to **Identity & Access**.

   For help, see the identify and access help topics in your Strata Cloud documentation on paloaltonetworks.com.

2. Add a new identity as described in the Strata Cloud documentation.
3. When the client credentials are displayed on the screen, copy the following values and store this information somewhere safe:
   - Client ID
   - Client Secret
   - TSG-ID—The Tenant/Service Group ID (TSG-ID) is the series of numbers directly after the `@` symbol in the Client ID. You can also view the TSG-ID from the lefthand panel of the Strata Cloud Manager.

     Refer to the following example, but consult your Strata Cloud documentation for the latest information.

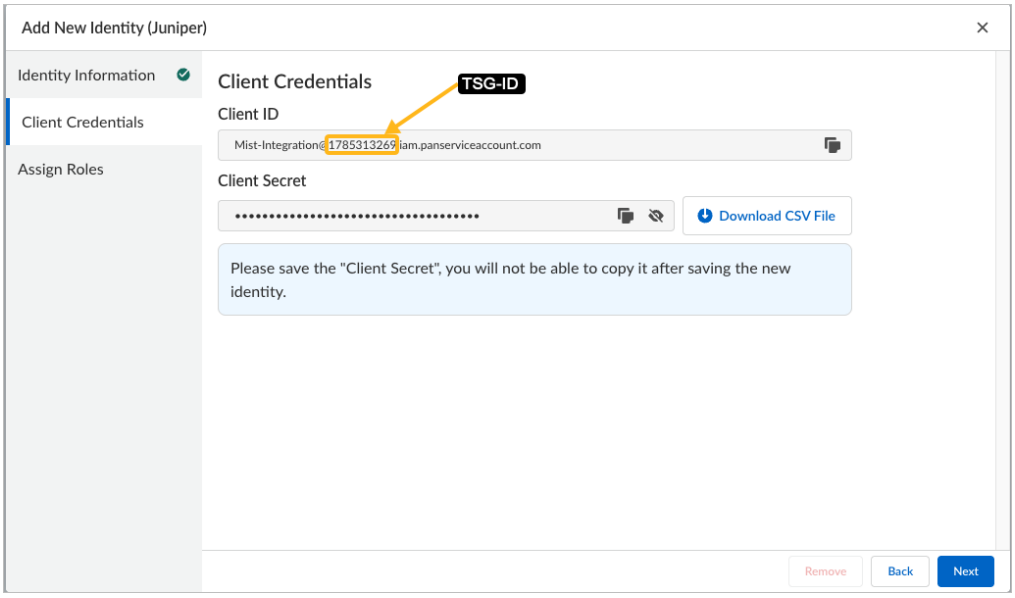

   Later, you'll enter these values in the Mist portal to "link a Prisma Access account" on page 268.
4. In Strata Cloud, assign a predefined role with the following information:
   - Apps & Services—Select **Prisma Access & NGFW Configuration**.
   - Role—Select **Network Administrator**.

> **NOTE**: The Network Administrator role is the minimum required access role to enable the link between Mist and Prisma APIs. With this role, Mist can access Prisma APIs for IPsec tunnel orchestration and can provide visibility into tunnel status, alerts, and incident notifications. If this role is not assigned, you must use the two separate dashboards (the Mist portal for Mist SD-WAN device troubleshooting and the Strata Cloud Manager for Prisma tunnel troubleshooting).

Refer to the following example, but consult your Strata Cloud documentation for the latest information.

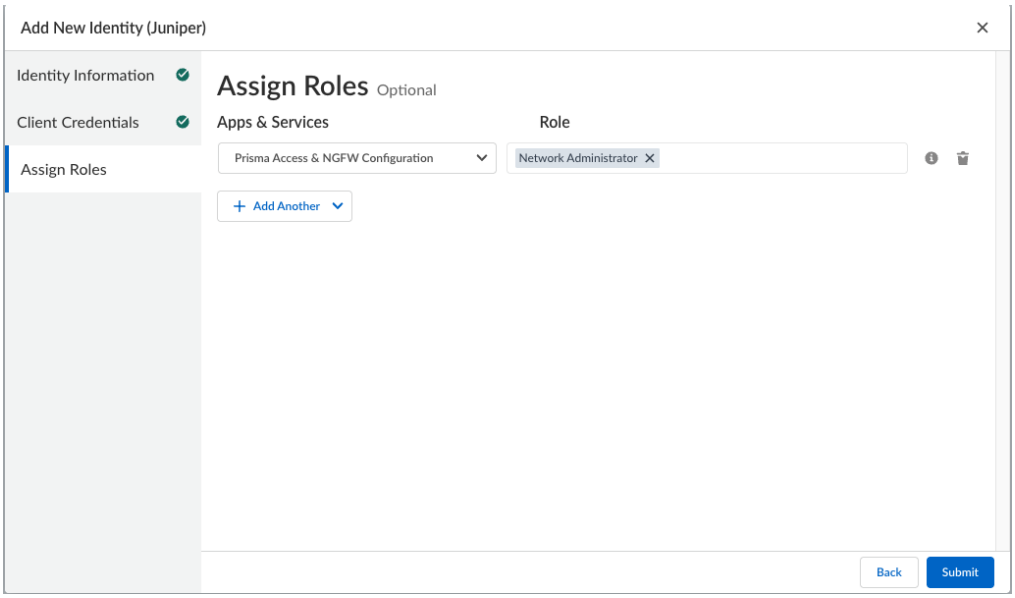

Add Prisma Access Account as a Secure Edge Connector

Follow these steps to add the Prisma Access account to your Mist Organization.

1. From the left menu of the Juniper Mist portal, select **Organization** > **Admin** > **Settings**.
2. In the **Secure WAN Edge Integration** section, click **Add Credentials**.
3. Enter the credentials you copied from the Prisma Strata Cloud Manager.
   a. Provider—Select **Prisma Access**.
   b. Add the **Service Account Identity Address**, **Client Secret**, and **TSG-ID**.
   c. Select the **I Agree** check box to consent to the terms.
   d. **Add the Tunnel Probe Source IP Address Range**—A randomly selected IP address from the specified range will be used to run a probe within the tunnel to measure key performance indicators (KPIs). Ensure that this IP is whitelisted in Prisma Access to allow the probe to pass through successfully.
   e. Click **Save**.

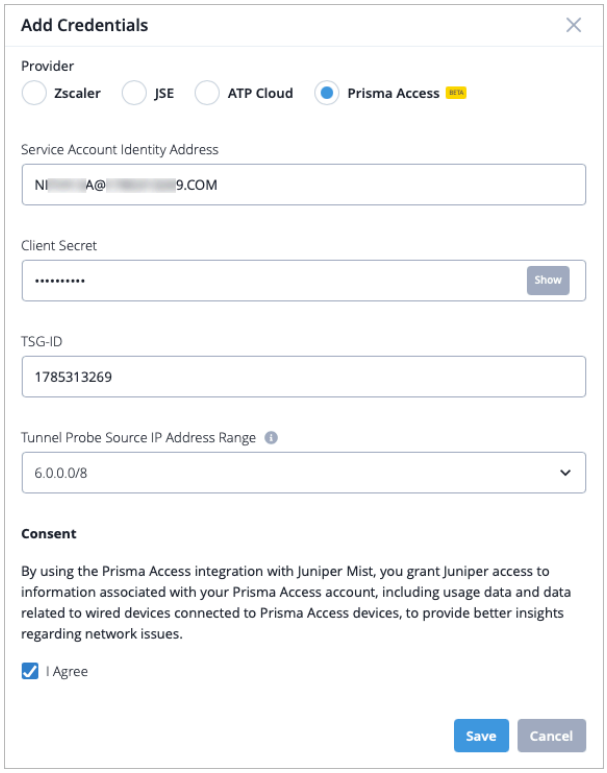

NOTE: If any of the credentials you entered are incorrect or expired, you will receive an error message and will not be able to save the settings.

The Prisma Access Provider and Username is then listed in the Secure WAN Edge Integration tile.

Auto Provision IPsec Tunnels

Mist's automatic Prisma tunnel provisioning requires very minimal configuration. You only need to configure in Mist, and that configuration will be replicated in the Strata Cloud Manager automatically.

Most aspects of the configuration are automatically set for you. All you have to do to configure your tunnel is enter a name, provider, and WAN interface. Some of the fields that are automatically configured for you include the Region, where Mist automatically finds your region for the tunnel based on your Prisma tenant information, and Data Center, where Mist automatically finds the closest geographically located Point of Presence (POP) to your service connection. However, you can configure these fields with other values if necessary.

NOTE: All configuration elements for auto tunnel creation can be manually overriden if needed by advanced users.

1. Navigate to the WAN Edge template (**Organization** > **WAN Edge Templates**). In the SECURE EDGE CONNECTOR AUTO PROVISION SETTINGS tile, select the Prisma Access Account you just created.

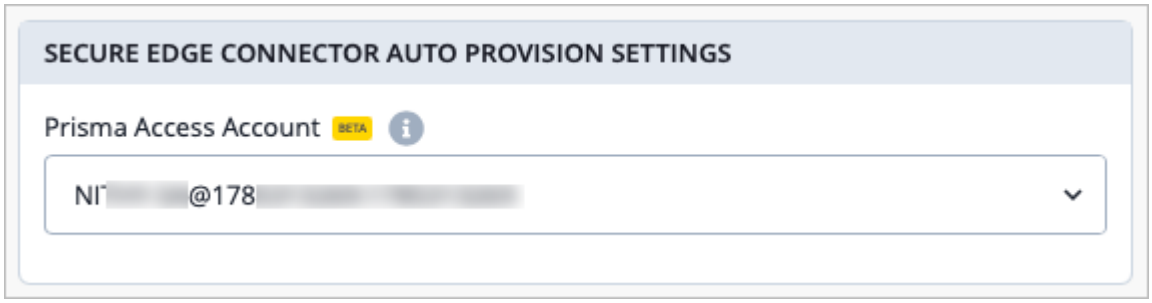

2. Configure an IPsec tunnel from the WAN Edge device to the Prisma Access cloud. To do this, you must first add a Provider.
   a. From the **Secure Edge Connectors** tile in the WAN Edge Template, select the **Add Providers** button.
   b. In the Add Provider window, enter the following information necessary for provisioning the tunnel:
      i. **Name**—Enter the name of the service.
      ii. **Provider**—Select **Prisma Access**.
      iii. **Remote Networks**—Select an existing network or create a new one.
      iv. **Probe IPs**—Enter the destination probe IP address. You can use any well-known IP (Example: 8.8.8.8). Probes are used to send information such as jitter, latency, and roundtrip time to Prisma, and are used to construct the Peer Path Statistics that display on the WAN Edge Insights page.
      v. **WAN Interfaces**—Assign WAN interfaces for provisioning of primary and secondary tunnels. You can add multiple WAN interfaces, and the first interface in the list has priority. If the first interface is down, then the second interface is used to establish the tunnel.

      In the OVERRIDE AUTO PROVISION OPTIONS section, default parameters are automatically configured for you as part of this automatic Provider configuration process. Use the fields in this section if you need to change any of the parameters that would otherwise be automatically chosen for you:

      - **Region**—This indicates the geographic region for the Secure Edge instance. When the default "Auto" is selected, the nearest Prisma cloud region is automatically chosen for you.

- **IKE v2 Proposals**—The encryption and authentication settings to be used for internet key exhange security association are automatically set for you, but you can change them if needed.
- **DH Group**—The size of the keys to be used in the IKE negotiation to establish the tunnel is automatically set for you, but you can change it if needed.
- **IPsec Proposals**—The encryption and authentication settings to be used for IPsec tunnels are automatically set for you, but you can change them if needed.
- **Data Center**—This is set to "none" by default. As part of the automatic provisioning, Mist automatically selects the location of the nearest point of presence (POP) for you. In other words, Mist automatically selects the nearest POP for your application. However, you have the option to specify a particular data center if needed.

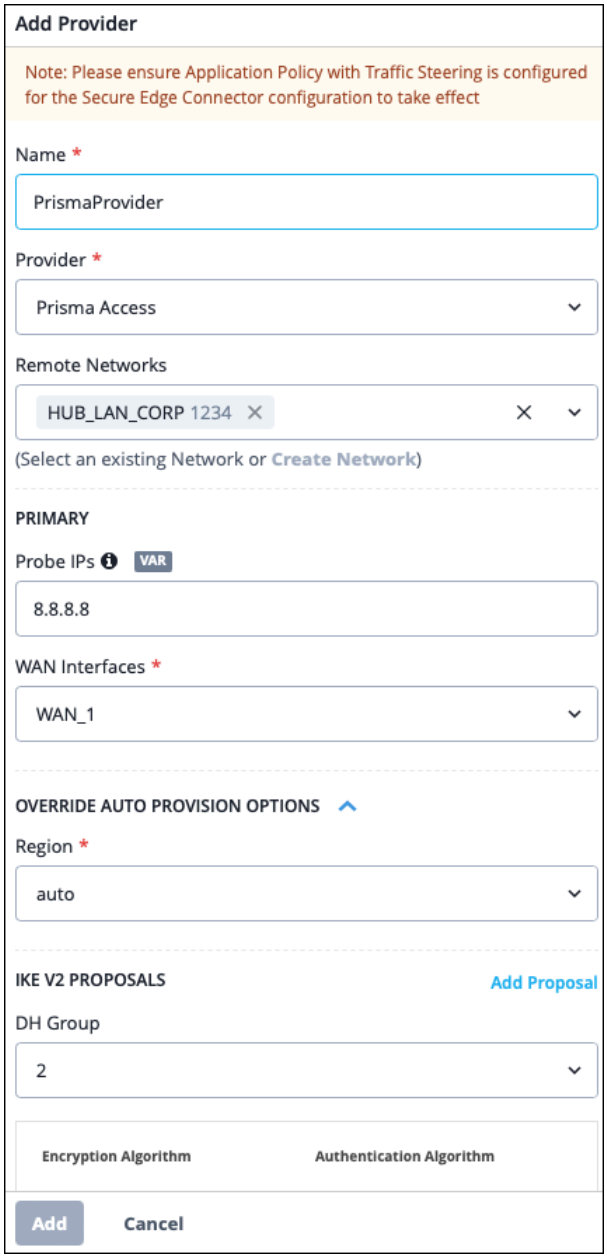

c. Finally, select **Add**.
Your Prisma Access provider is now listed in the **Secure Edge Connectors** tile.

NOTE: The provider information you configured is automatically carried over to Prisma once the WAN Edge template is updated, so there is no need for manual configuration on the Prisma side. This is possible due to the "account linking" on page 268 step.

In addition to configuring a provider, you must also complete the steps below to set up an IPsec tunnel.

Configure Traffic Steering and Application Policies

1. Navigate to the WAN Edge template (**Organization** > **WAN** > **WAN Edge Templates**).
2. Scroll down to the Traffic Steering section and click **Add Traffic Steering**.

   Later, you'll add this traffic steering profile to an application policy to specify the path traffic can take to the destination.

   a. Enter the details for the traffic-steering path:

   - **Name**—Enter a name for the traffic-steering profile.
   - **Strategy**—Select a strategy. You can configure the traffic steering profile with any strategy (**Ordered**, **Weighted**, or **ECMP**), based on your topology and configuration.
   - **PATHS**—Click **Add Paths** and enter the following details.

     i. **Type**—Select **Secure Edge Connector**.

     ii. **Provider**—Select **Custom**.

     iii. **Name**—Select the Prisma Provider's name you created in step 2.

   b. Select the blue check mark to save the changes.

   c. Select **Add**.

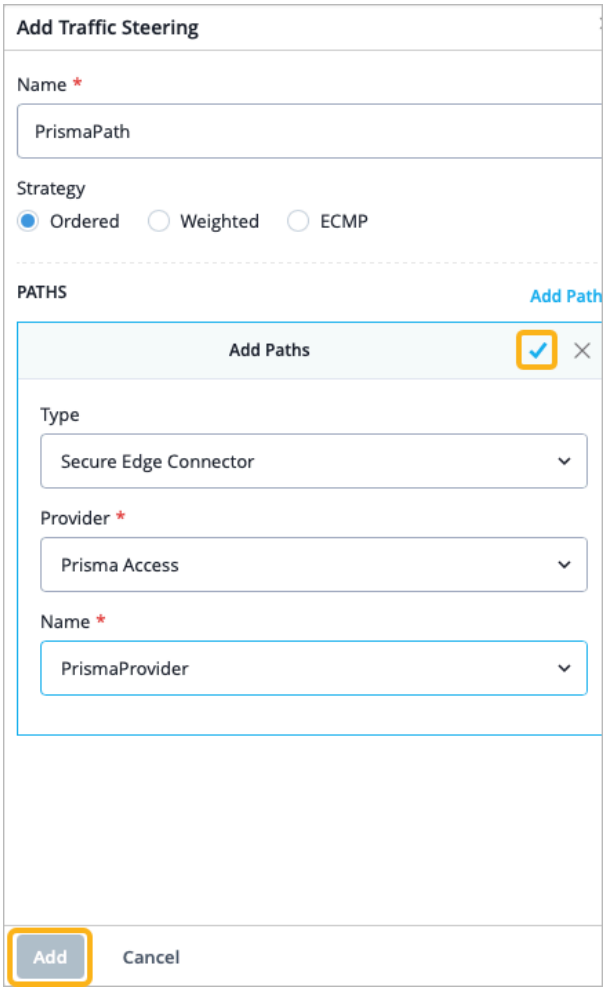

3. Next, scroll down to the APPLICATION POLICIES tile on the WAN Edge Template and click **Add Application Policy**. An application policy defines the networks and users that can access an application and which path traffic takes to the destination.

   a. Give your application policy a name.

   b. **Network/User**—This is the LAN user that needs secure access to applications through the Prisma cloud.

   c. **Action**—Select an action of **Allow** for the traffic.

   d. **Application/Destination**—Select the applications that you want the Network/User to have access to.

   e. **Traffic Steering**—Select the traffic steering profile you created in step 3. This specifies the path that traffic is allowed to take to reach its destination.

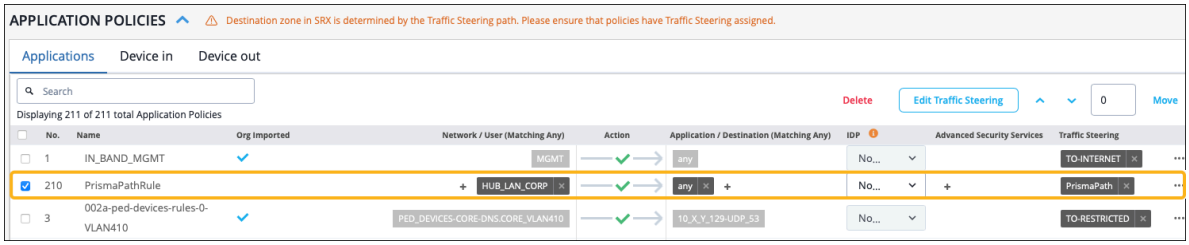

For more information about how to create Application Policies, see Configure Application Policies.

Tunnel Verification:

The IPsec configuration is pushed to any WAN Edge devices that belong to a site that has the template assigned, and a tunnel from the device to the closest Prisma cloud is brought up. To see the IPsec status, select **WAN Edges** > **WAN Edges** from the left menu, then click the WAN Edge device, and finally, click **WAN Edge Insights**.

You can verify the established tunnel's details on the WAN Edge Insights page of the device once the **WAN Edge Tunnel Auto Provision Succeeded** event appears under **WAN Edge Events**.

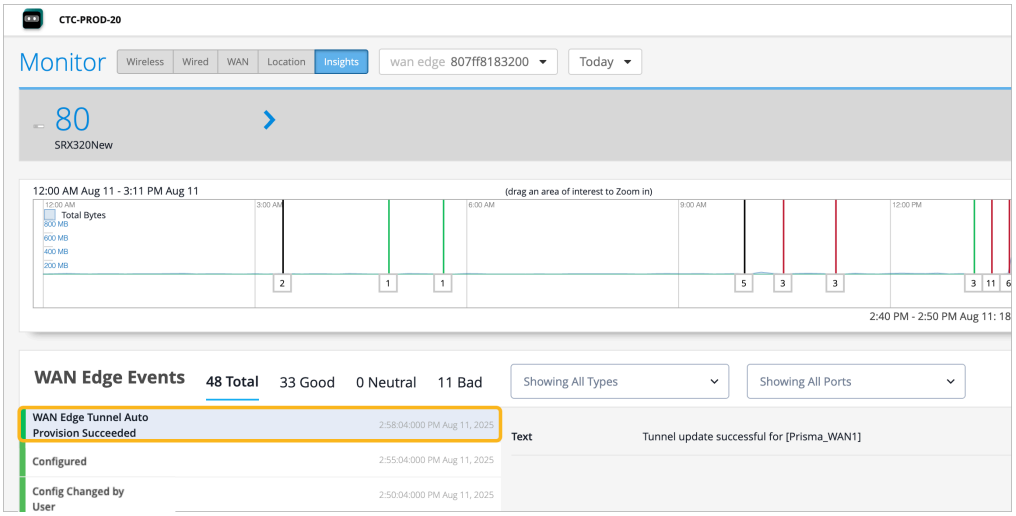

Prisma Events

You can view any WAN Edge Events including Prisma Access events under **WAN Edge Events** on the **WAN Edge Insights** page for the device that has the IPsec tunnel configured to the Prisma Access cloud.

NOTE: You must reach out to your account team in order to enable Prisma Access events. These events are disabled by default due to a notification profile issue at Prisma Access.

- Juniper Mist supports the following Prisma Access Events:

Remote Networks:

- Prisma RN ECMP BGP Down
- Prisma RN ECMP BGP Flap
- Prisma RN ECMP Proxy Tunnel Down
- Prisma RN ECMP Proxy Tunnel Flap
- Prisma RN Primary WAN BGP Down
- Prisma RN Primary WAN BGP Flap
- Prisma RN Primary WAN BGP Up
- Prisma RN Primary WAN Proxy Tunnel Down
- Prisma RN Primary WAN Proxy Tunnel Flap
- Prisma RN Primary WAN Tunnel Down
- Prisma RN Primary WAN Tunnel Flap
- Prisma RN Primary WAN Tunnel Up
- Prisma RN Secondary WAN BGP Down
- Prisma RN Secondary WAN BGP Flap
- Prisma RN Secondary WAN BGP Up
- Prisma RN Secondary WAN Proxy Tunnel Down
- Prisma RN Secondary WAN Proxy Tunnel Flap
- Prisma RN Secondary WAN Tunnel Down
- Prisma RN Secondary WAN Tunnel Flap
- Prisma RN Secondary WAN Tunnel Up

Service Connection:

- Prisma Service Connection Primary WAN BGP Down
- Prisma Service Connection Primary WAN BGP Flap
- Prisma Service Connection Primary WAN Proxy Tunnel Down
- Prisma Service Connection Primary WAN Proxy Tunnel Flap

- Prisma Service Connection Primary WAN Tunnel Down
- Prisma Service Connection Primary WAN Tunnel Flap
- Prisma Service Connection Secondary WAN Proxy Tunnel Down
- Prisma Service Connection Secondary WAN Proxy Tunnel Flap
- Prisma Service Connection Secondary WAN Tunnel Down
- Prisma Service Connection Secondary WAN Tunnel Flap
- Prisma Service Connection WAN BGP Down
- Prisma Service Connection WAN BGP Flap

NOTE: A Prisma Access icon displays next to any Prisma Access events received from Strata Cloud Manager. The information in the Prisma Access event comes directly from Strata Cloud Manager.

- When you select a Prisma Access event, you have the **View Prisma SCM Incidents** button available to you. Select the button if you need to see more details about the Prisma incident from Prisma SCM.

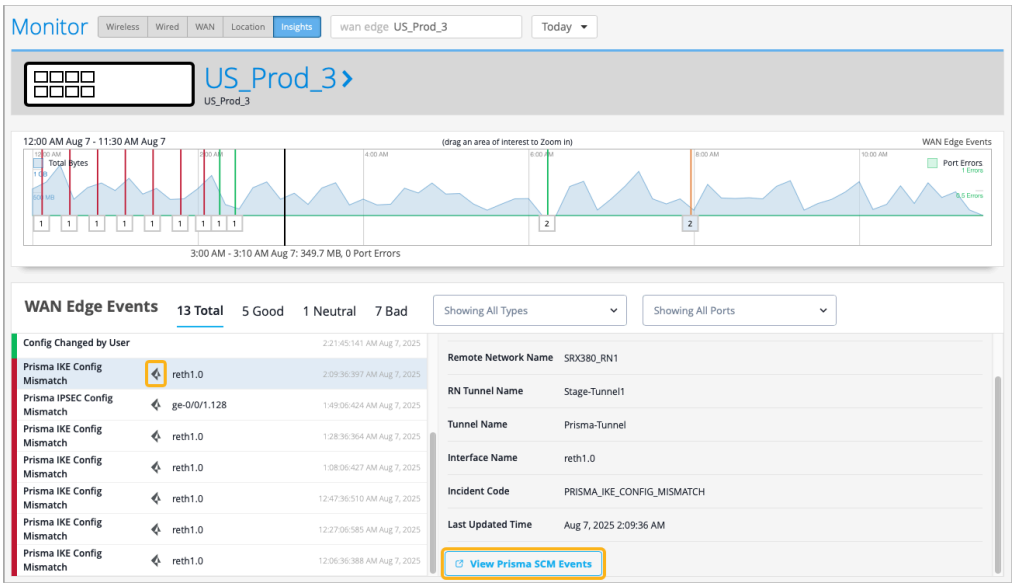

Configuration Mismatch

Marvis constantly compares the site configuration in Mist to the configuration in Prisma. If any deviations are detected, a Prisma Access event is generated and appears in the WAN Edge Events.

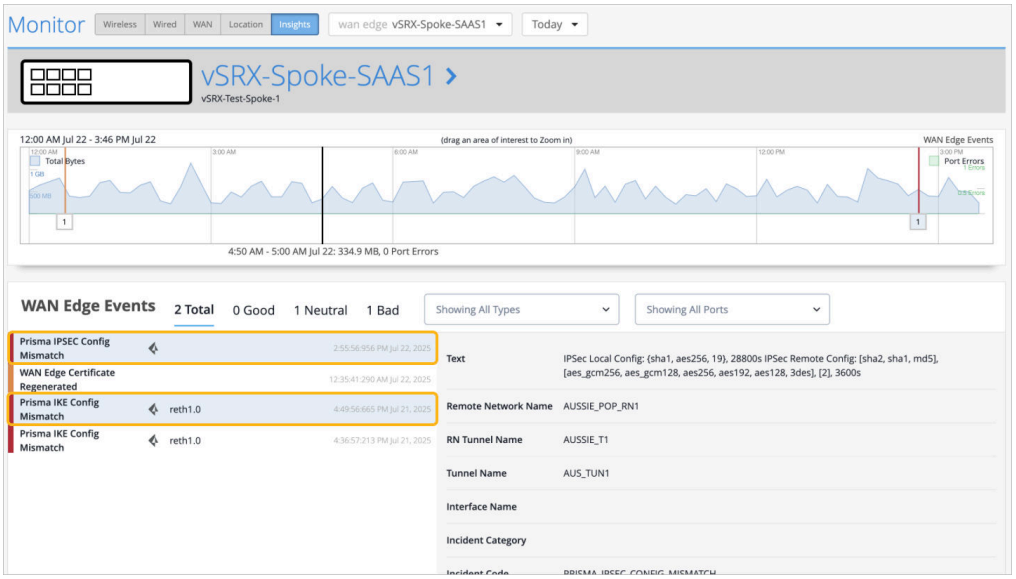

Configuration Difference Alerts:

- Prisma IKE Config Mismatch
- Prisma IPsec Config Mismatch

Tunnel Statistics

You can view the Prisma tunnel statistics for both Session Smart Router (SSRs) and SRX Series Firewalls under **Peer Path Stats** on the **WAN Edge Insights** page.

Tunnel statistics use intelligent probing to a northbound resource and generate the near-real time Key Performance Indicator (KPI). You can see the customer's experience as they access an application through the tunnel in near-real time.

NOTE: Make sure you have the appropriate Probe IP and Source IP addresses configured in the "provider configuration" on page 273 to ensure that data populates the chart.

When you hover over the Peer Path Stats chart, you get information about the path in the string that is displayed to you. You can decipher the string as follows:

egress device interface <--> Prisma tunnel name.

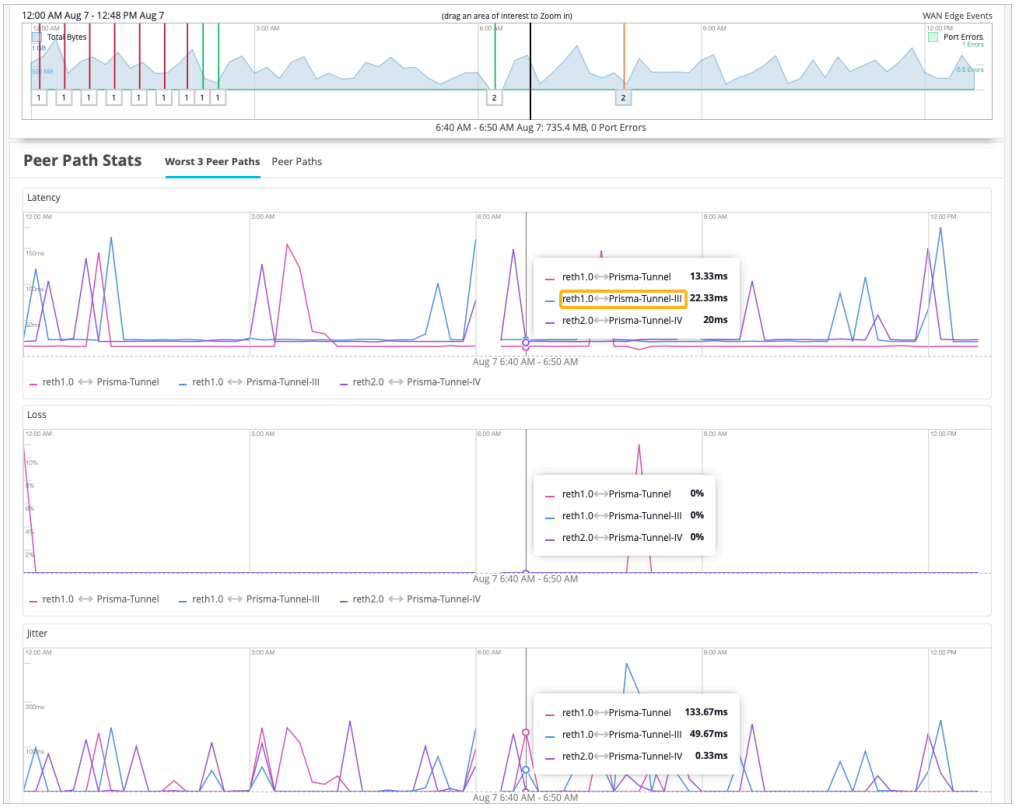

SEE ALSO

> https://docs.paloaltonetworks.com/prisma-access/integration/juniper-mist-integration-for-sase-health

Zscaler Integration

SUMMARY

Follow this workflow to use Zscaler as your Secure Edge provider.

IN THIS SECTION

Juniper Mist™ provides pre-built connectors specifically designed for the Juniper Networks® SRX Series Firewalls and Juniper® Session Smart™ Routers deployed as WAN edge devices. These connectors facilitate seamless integration with your Secure Edge (SSE) deployments. With minimal configuration, you can integrate the SSE into the Juniper Mist portal. As a result, your WAN Edge device establishes connections to the SSE using either IPsec or GRE protocols.

Overview

Juniper SD-WAN with Mist AI enhances security, and integration with Zscaler offers a top-notch SASE solution for secure access across locations. Zscaler provides comprehensive cyber-security and zero trust connectivity, all integrated seamlessly with Juniper Mist Cloud.

Figure 58: Juniper Mist Integration with Zscaler

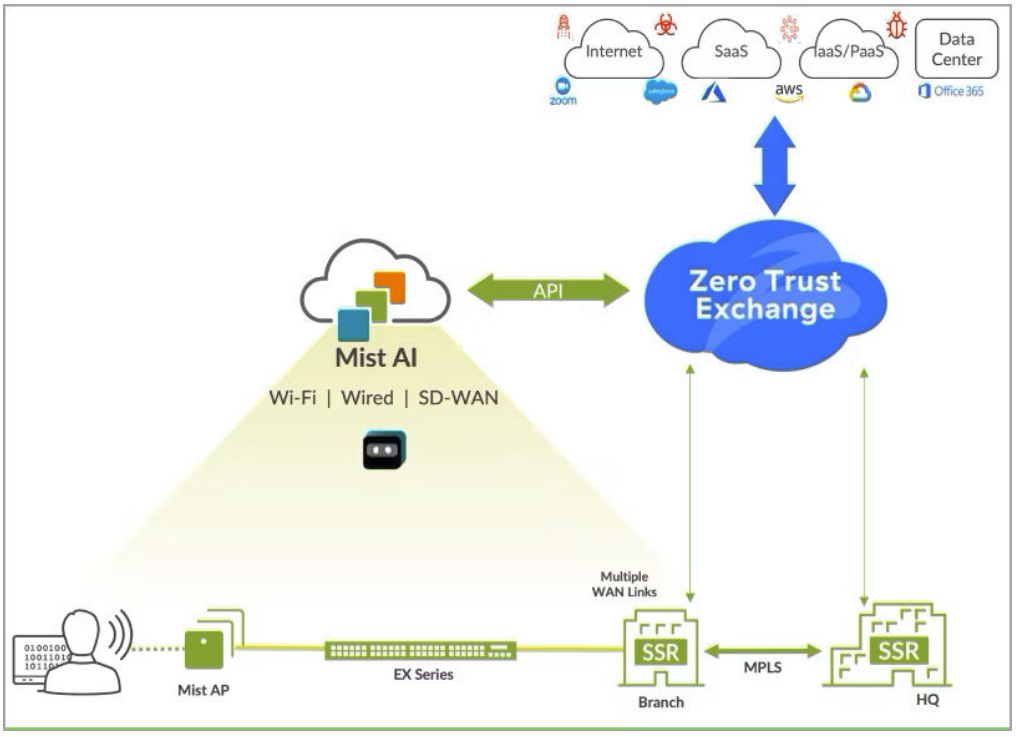

To improve cloud-to-cloud connectivity, Mist now offers automated Zscaler tunnel provisioning. WAN Assurance delivers an easy-to-use workflow for default Zscaler connections. This integration allows you to effortlessly establish connections from your devices using predefined settings. Juniper Mist Cloud portal allows simplified configuration and management of Zscaler Security Service Edge (SSE) connections—eliminating the need for additional platform logins.

Note the following about Zscaler tunnel provisioning:

- Automatic provisioning supports only IPsec tunnels.
- The tunnel auto provision process includes the creation of a new Zscaler location and VPN credential objects in the Zscaler cloud. These Zscaler resources will be deleted when the associated tunnel provider is removed from the Mist.
- Each auto orchestration requires the creation of one primary tunnel and one secondary tunnel.

Read the following topics to understand how to set up Zscaler tunnel configuration from the Juniper Mist Cloud.

Zscaler Manual Provisioning of Tunnels

IN THIS SECTION

Prerequisites

For IPsec Tunnels

- Zscaler cloud account
- Local ID and pre-shared key (PSK) configured from the Zscaler account. Ensure that you set the length of PSK between 6 to 255 characters.
- IP addresses or hostnames of the Zscaler public service edges tunnels (primary and secondary). See Locating the Hostnames and IP Addresses of Zscaler Enforcement Nodes (ZENs).
- Check Configuring an IPSec VPN Tunnel.

For GRE Tunnels

- Zscaler cloud account
- IP addresses or hostnames of the Zscaler public service edges tunnels (primary and secondary). See Locating the Hostnames and IP Addresses of Zscaler Enforcement Nodes (ZENs). See Configuring GRE Tunnels.
- Static IP address. See About Static IP

Configure Zscaler Tunnel Provisioning

1. On Juniper Mist portal, go to **Secure Edge Connector** at WAN Edge Templates-level, hub profile, or at Site-level.
2. Click **Add Providers**.
3. In **Add Provider** window, select **Zscaler** for manual provisioning.
4. Enter the details for Zscaler manual provisioning of tunnels.
   a. For IPsec Tunnels

Figure 59: Add Details for Zscaler IPsec Tunnels

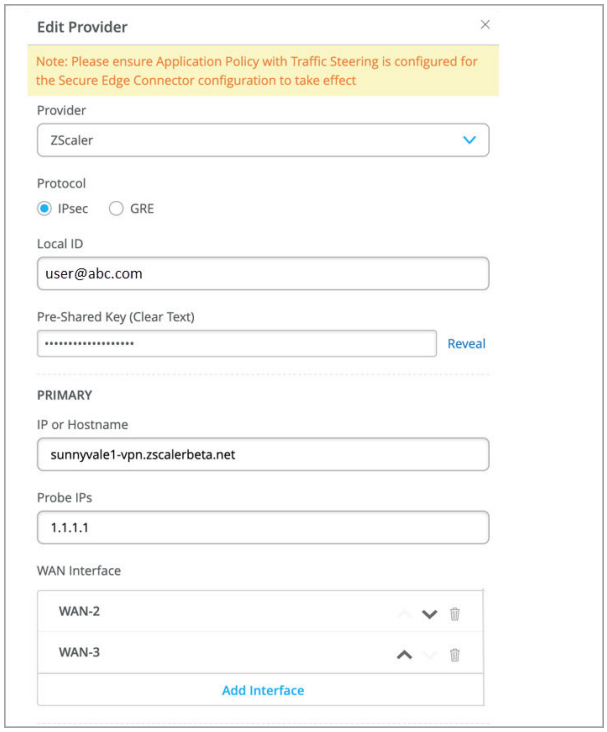

i. **Name**—Enter the name of the service.

ii. **Provider**—Select Zscaler.

iii. **Protocol**—Select the protocol as IPsec.

iv. **Local ID**—Provide login ID of the Zscaler account.

v. **Preshared Key**—Provide preshared key (PSK) created with Zscaler account. The length of the PSK must be between 6-255 characters.

vi. **IP or Hostname**—IP addresses or hostname of the Zscaler DC. (See https://config.zscaler.com/zscalerbeta.net/cenr). We recommend to add IP address of the nearest DC to your device location.

vii. **Probe IPs**—Enter probe IP address. You can use any well-known IP (Example: 8.8.8.8). You need probe IP address to monitor the status of the Zscaler IPsec tunnel using probes.

viii. **WAN Interface**—Assign WAN interfaces for provisioning of primary and secondary tunnels. You can add multiple WAN interfaces and the first interface takes the priority. If first interface is down, then system uses the second interface to establish the tunnel.

ix. **Mode**—Select active-standby option.

Configure secondary tunnel options (optional). Enter IP or Hostname, Probe IPs, WAN Interface, and Mode for the secondary tunnel.

b. **For GRE Tunnels:**

Figure 60: Add Details for Zscaler GRE Tunnels

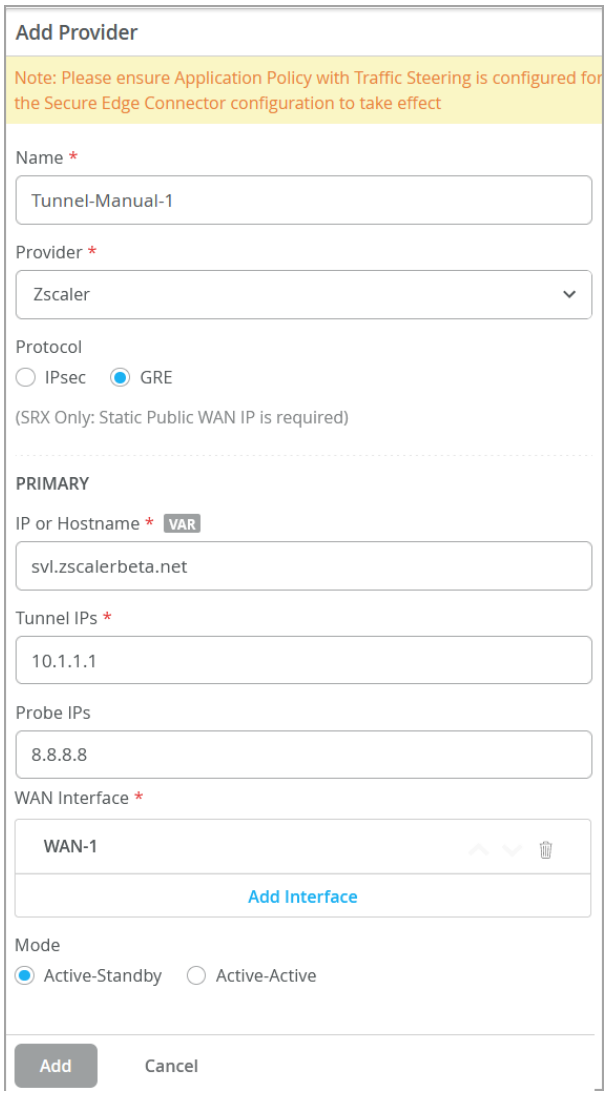

i. **Name**—Enter the name of the service.

ii. **Provider**—Select Zscaler.

iii. **Protocol**—Select the protocol as GRE.

iv. **IP or Hostname**—IP addresses or hostname of the Zscaler DC. (See https://config.zscaler.com/zscalerbeta.net/cenr). We recommend to add IP address of the nearest DC to your device location.

v. **Tunnel IP**—Static IP address created for GRE tunnel. See Configuring GRE Tunnels.

vi. **Probe IPs**—Enter probe IP address. You can use any well-known IP (Example: 8.8.8.8). You need probe IP address to monitor the status of the Zscaler IPsec tunnel using probes.

vii. **WAN Interface**—Assign WAN interfaces for provisioning of primary and secondary tunnels. You can add multiple WAN interfaces and the first interface takes the priority. If first interface is down, then system uses the second interface to establish the tunnel.

viii. **Mode**—Select active-standby option.

Enter IP or Hostname, Tunnel IP, Probe IPs, and WAN Interface for the secondary tunnel.

5. Click **Add** to continue.
6. Add a traffic steering profile on the WAN Edge Templates page or on WAN Edge Device page.

Figure 61: Traffic Steering Path

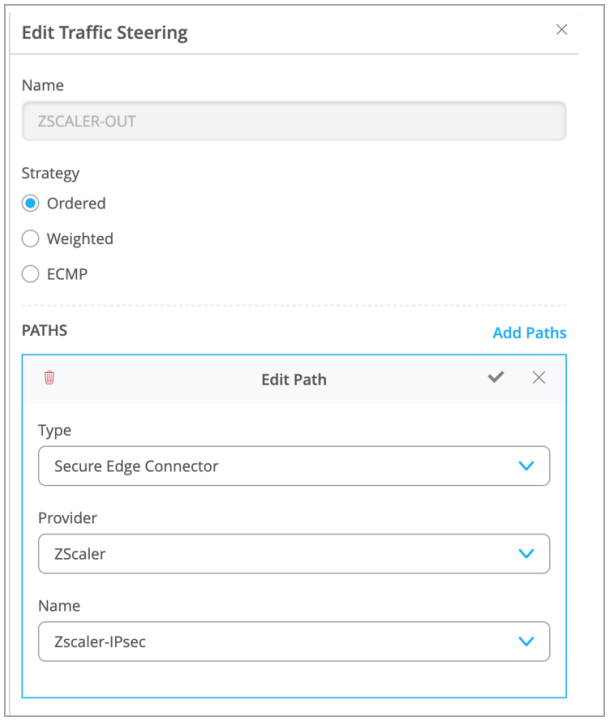

- Enter the details for the traffic-steering path:
  - **Name**—Enter a name for the traffic-steering profile.
  - **Strategy**—Select a strategy. You can configure the traffic steering profile with any strategy (Ordered/Weighted/ECMP), based on your topology and configuration.
  - **Type**—Select **Secure Edge Connector**.
  - **Provider**—Select **Zscaler**.

7. Add an application policy to refer the traffic steering profile you created. This step is required for provider tunnels to take effect. To create the application policy, in the Juniper Mist cloud portal, go to **Organization** > **WAN** > **Application Policy**

   Following image shows an example of application policy with traffic steering configured in previous step.

Figure 62: Application Policies

Zscaler Auto-Provisioning of Tunnels

IN THIS SECTION

Prerequisites

You need a partner key and the partner admin login credentials in Zscaler portal for auto-provisioning of Zscaler tunnels.

1. Add a partner API key.

a. Use your Zscaler account to login admin portal.

b. Select **Administration** > **Partner Integrations**.

c. Select the **SD-WAN** tab.

d. Click **Add Partner Key**.

e. In the **Add Partner Key** window, choose the partner name from a drop-down menu and click **Generate**.

Figure 63: Selecting SD-WAN Partner Integration

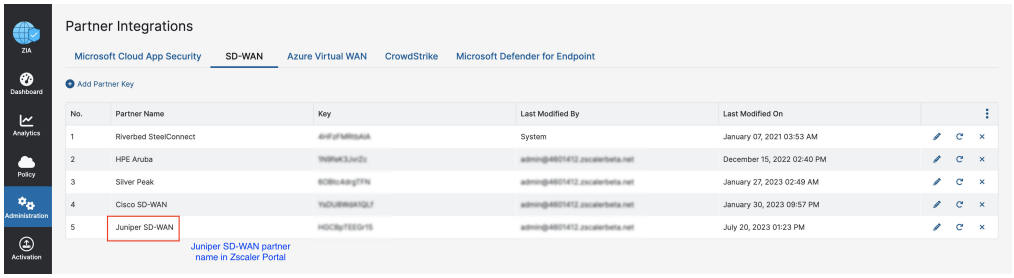

2. Provide credentials for the API access.

a. Go to **Administration** > **Role Management** > **Add Partner Administrator Role**.

b. Configure the following options:

- **Name**—Name of the administrator role.
- **Access Control**—Select the access control type as **Full**.
- **Partner Access**—Select the type of access categories for the access control (SD- WAN API Partner access permission).

Check Adding Partner Admin Roles for details.

3. Create a partner account for the SD-WAN Orchestrator. This admin user (username/password) is specifically used for SD-WAN partner authentication and it is different from Zscaler account admin user.

a. Select **Administration** > **Administrator Management** > **Add Partner Administrator**.

b. Enter the details such as login ID, email address, name, partner role, and password. See Adding Partner Admins for details.

4. Find Zscaler Cloud name using What is My Cloud Name.

Add Zscaler Credentials in Juniper Mist Portal

1. Provide Zscaler credential details in Juniper Mist portal to integrate Juniper Mist cloud with Zscaler.
   - On Juniper Mist portal, select **Organization > Settings**.
   - Scroll-down to **Secure WAN Edge Integration** pane and click **Add Credentials**.
   - In **Add Provider** window, enter the details.

     Figure 64: Add Credentials for Zscaler

     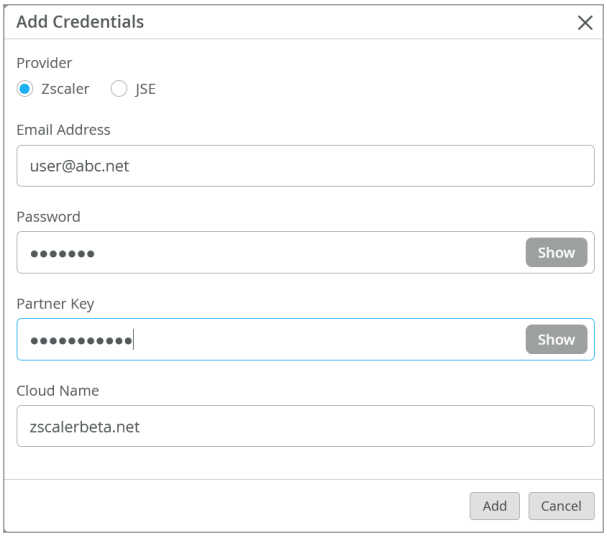

     - **Provider**—Select Zscaler.
     - **Email Address**—Enter username (email address) (SD WAN partner user credentials)
     - **Password**—Enter password for the username.
     - **Partner Key**—Input partner key you created when configuring your Zscaler account.
     - **Cloud Name**—Zscaler cloud URL. For example, zscalerbeta.net.
   - Click **Add** to continue.

   This procedure is one-time configurations at the organization level. To automatically provision the Zscaler tunnels across several sites, Mist Cloud uses the above-mentioned credentials for the given Organization.

Configure Zscaler Auto Tunnel Provisioning

1. On Juniper Mist portal, go to **Secure Edge Connector** at WAN Edge Templates-level or at Site-level.

2. Click **Add Providers**.
3. In **Add Provider** window, select **Zscaler (Auto)** for auto provisioning.

Figure 65: Add Details for Zscaler Tunnel Auto Provisioning

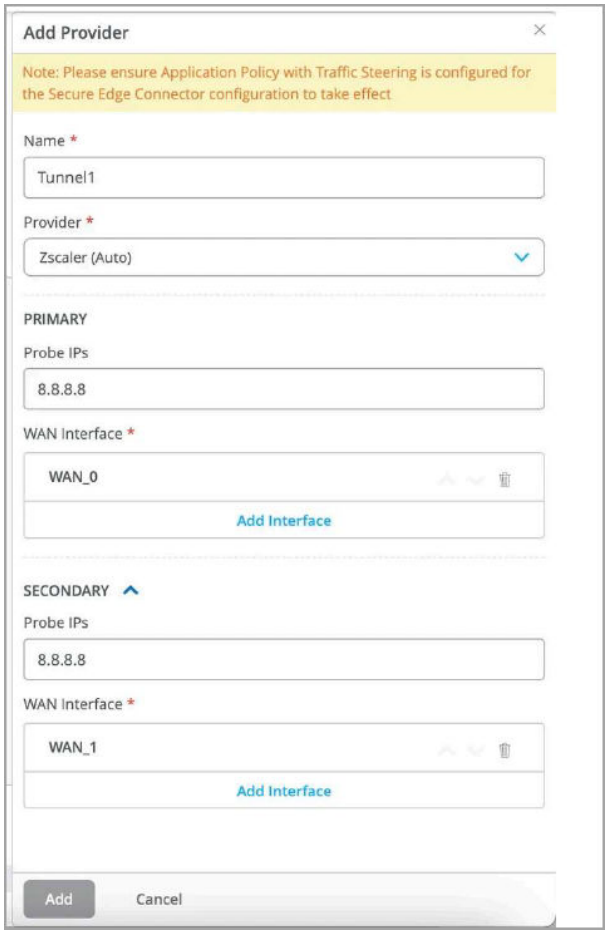

Enter the following details for Zscaler auto-provisioning:

- **Name**—Enter the username of the Zscaler account.
- **Provider**—Select Zscaler.
- **Probe IPs**—Enter probe IP address (primary and secondary). Enter any well-known IP address as probe IP. Example: 8.8.8.8.
- **WAN Interface**—Assign WAN interfaces under primary and secondary tunnel details for provisioning of primary and secondary tunnels.

4. Click **Add** to continue.

5. Add a traffic steering profile on the WAN Edge Templates page or on WAN Edge Device page.

Figure 66: Traffic Steering Path

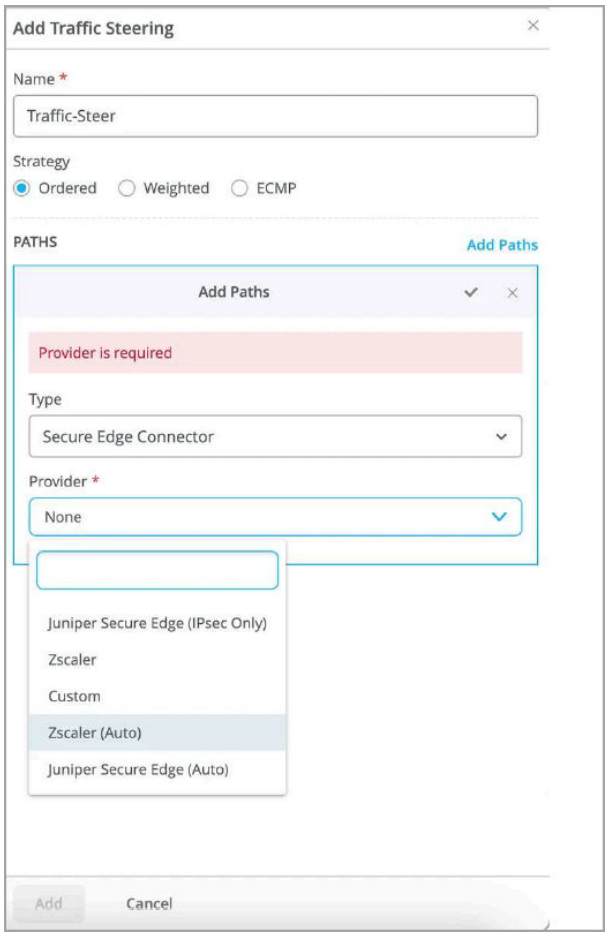

- You can add a new traffic-steering path by selecting:
  - **Name**—Enter a name for the traffic-steering profile.
  - **Strategy**—Select a strategy. You can configure the traffic steering profile with any strategy (Ordered/Weighted/ECMP), based on your topology and configuration.
  - **Type**—Select **Secure Edge Connector**.
  - **Provider**—Select **Zscaler**.

6. Add an application policy to refer the traffic steering profile you created. This is required for provisioning of provider tunnels to take effect. To create the application policy, in the Juniper Mist cloud portal, go to **Organization** > **WAN** > **Application Policy**

   Following image shows an example of application policy with traffic steering configured in previous step.

Figure 67: Application Policies

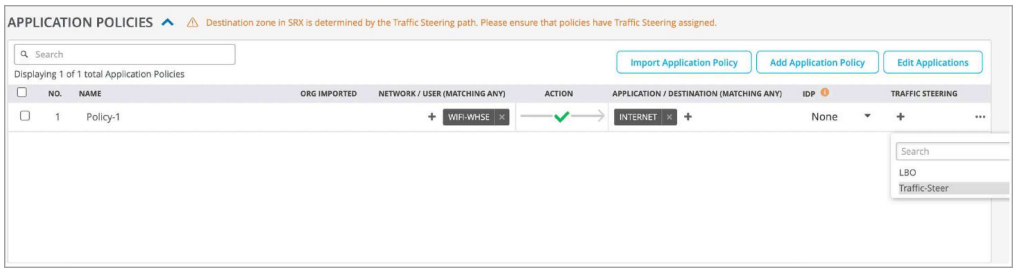

When you assign a template that is enabled with the Zscaler (Auto) option to a site, the following operations take place:

- An associated Zscaler site (location object) is automatically created. You can view the location object in **Administration** > **Location Management** on Zscaler portal.

Figure 68: Location Created in Zscaler Portal

- The details required for tunnel creation are exchanged between Juniper Mist cloud and Zscaler.
- Tunnels are powered up from the device to the closest network point-of-presence (POP).

Verify Juniper Secure Edge Tunnels

On Juniper Mist portal, you can verify the established tunnels details in WAN Insights of the device once **WAN Edge Tunnel Auto Provision Succeeded** event appears under WAN Edge Events.

Figure 69: WAN Edge Events

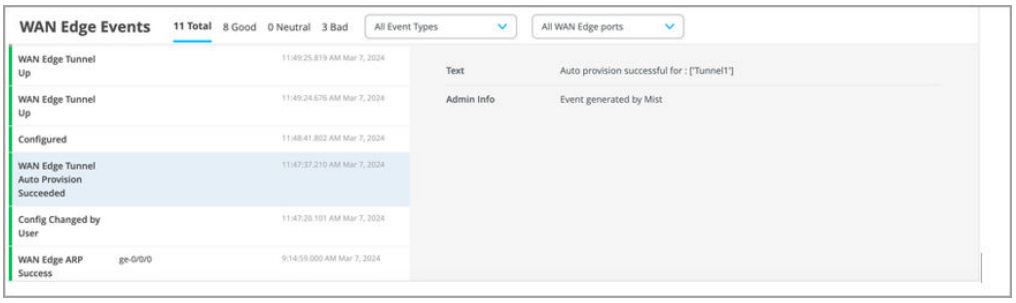

Get the established tunnels status details in **WAN Edges** > **WAN Edge Insights** page Juniper Mist cloud portal.

Figure 70: Status of Tunnels in WAN Edge Insights

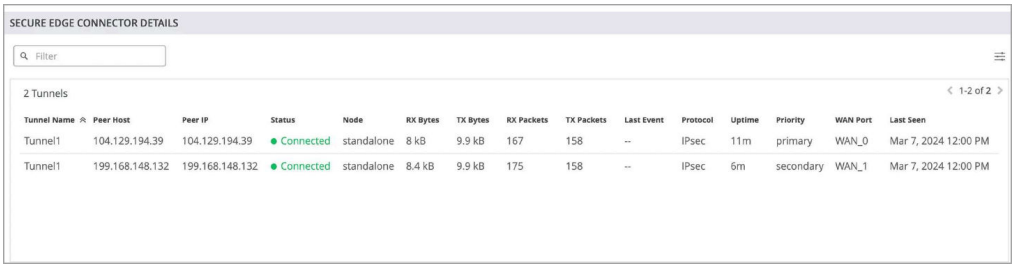

Figure 71: Established Secure Edge Tunnels

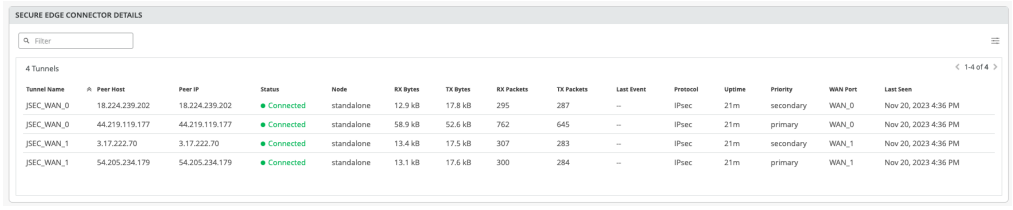

You can check the established tunnels in the Juniper Security Director Cloud dashboard and in the service location.

Secure Edge Connector Auto Provision Settings

You can configure gateway options per Zscaler location in **Secure Edge Connector Auto Provision Settings** on Juniper Mist portal. These settings offer additional control for configuring various traffic rules and policies, and they are optional parameters.

1. On Juniper Mist portal, select **Organization** > **WAN Edge Templates** or select **WAN Edges** > **WAN Edges** > *WAN Edge Name*.
2. Scroll-down to Secure Edge Connectors.
3. Select Zscaler (Auto) as the Secure Edge Connector. In the **Secure Edge Connector Auto Provision Settings** pane, you can define the gateway options.

Figure 72: Secure Edge Connector Auto Provision Settings

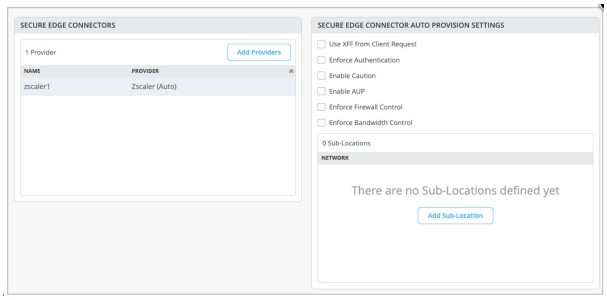

- **Use XFF from Client Request**—Enable this option if this location uses proxy chaining to forward traffic to the Zscaler service.
- **Enforce Authentication**—Enable this option if you want to authenticate users from this location.
- **Enable Caution**— Set the caution interval for more than one minute to display a caution notification for unauthenticated users. Use this option if you have disabled **Enforce Authentication** option.
- **Enable AUP**—Enable to display Acceptable Use Policy (AUP) for unauthenticated traffic and mandate it for the users to accept it. Use this option if you have disabled **Enforce Authentication**.
- **Enforce Firewall Control**—Enable the firewall control options.
- **Bandwidth Control**—Enforce bandwidth control for the location. You can specify the maximum bandwidth limits for upload and downloads.

Zscaler Auto-Tunnel Provisioning with Sub-Locations

Juniper Mist supports configuration and provisioning of Zscaler sub-locations. A Zscaler sub-location is a child entity of the location object. Locations identify the various networks from which your organization sends its Internet traffic. Sub-locations can be used for specific uses cases. For example, an organization can define a Zscaler sub-location for its corporate network, and another sub-location for its guest network, even if their traffic goes through the same IPsec tunnel.

The organization uses the sub-locations to:

- Implement different policies based on IP addresses.
- Enforce authentication for the internal corporate network, while disabling it for the guest network.
- Enforce bandwidth control for sub-locations while ensuring that unused bandwidth remains available to the parent location.

Juniper Mist supports sub-location provisioning as part of the tunnel orchestration process, and you can define the sub-location options using the Mist portal.

To configure Zscaler sub-locations:

1. On Juniper Mist portal, select **Organization** > **WAN Edge Templates** or select **WAN Edges** > **WAN Edges** > *WAN Edge Name*.
2. Scroll-down to **Secure Edge Connectors**.
3. Select Zscaler as the Secure Edge Connector. You'll see sub-location for the site below **Secure Edge Connector Auto Provision Settings** option.

Figure 73: Add Sub-Location

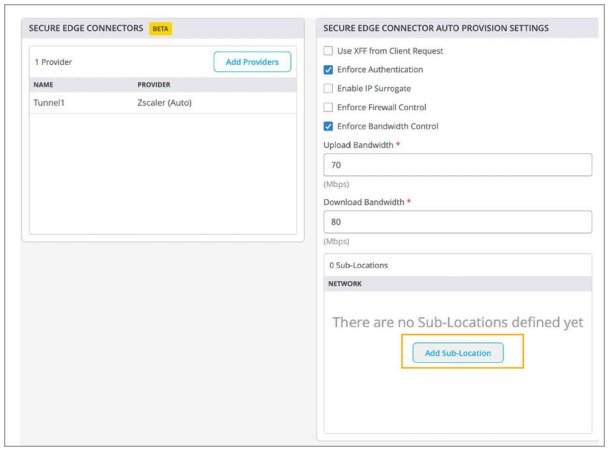

4. Click **Add Sub-Location** to define the new sublocation.

   The **Add Sub-Location** option is available only when you select Zscaler as the Secure Edge Connector.

5. In the **Add Sub-Location** window, define settings for the sub-location.

Figure 74: Sub-Locations Settings

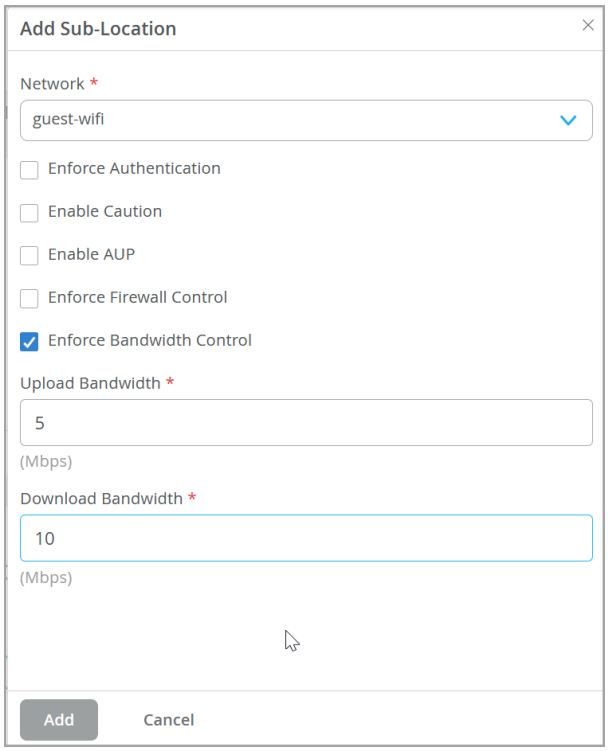

- **Network**—Select an existing network from the drop-down box.
- **Enforce Authentication**—To authenticate users from this location.
- **Enable Caution**— Display a caution notification for unauthenticated users. Use this option if you have disabled **Enforce Authentication** option.
- **Enable AUP**—Enable to display Acceptable Use Policy (AUP) for unauthenticated traffic and mandate it for the users to accept it. Use this option if you have disabled **Enforce Authentication**. The custom AUP Frequency must be a number between 1 and 180.
- **Enforce Firewall Control**—Enable the firewall control options.
- **Bandwidth Control**—Enforce bandwidth control for the location. You can specify the maximum bandwidth limits for upload and downloads. Upload and download bandwidth must be a number between 0.1 and 99999.

View Sub-Locations in Zscaler Portal

You can view the newly created sub-location in **Administration** > **Location Management** on Zscaler portal. On the Locations page, when you click the sublocation's number within the table, the sub-locations for the location appear.

Figure 75: View Configured Sub-Location in Zscaler Portal

Figure 76: View Other Sub-Location in Zscaler Portal

Verification and Troubleshooting

On Juniper Mist portal, you can verify the established tunnels details in WAN Insights of the device once **WAN Edge Tunnel Auto Provision Succeeded** event appears under WAN Edge Events.

Figure 77: WAN Edge Events

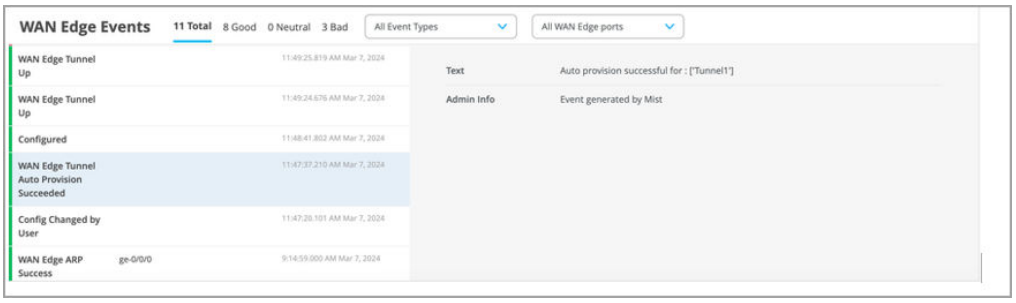

Figure 78: View Established Tunnels

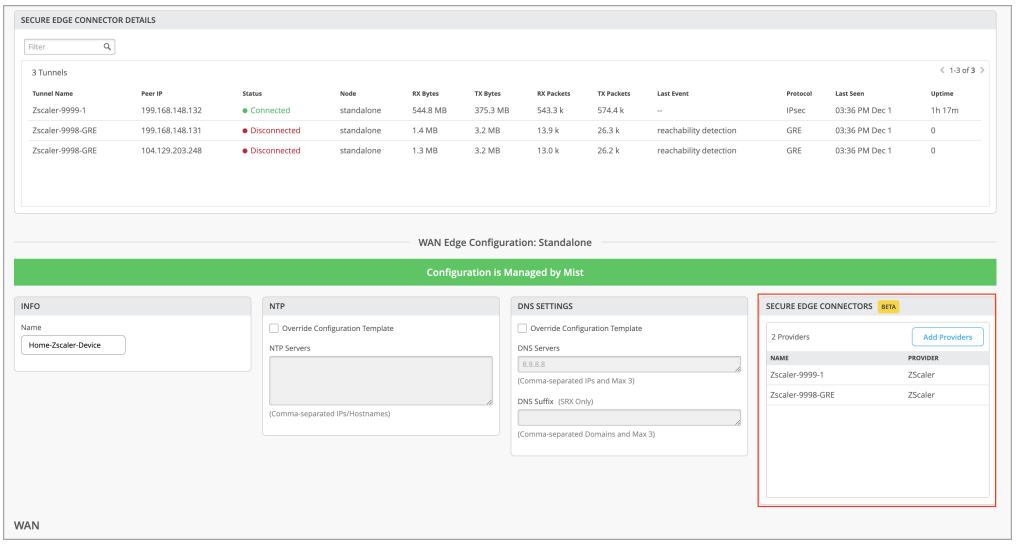

If you are not able to establish the tunnel, the possible cause could be tunnel configuration issues or reachability issues from your device. You can use the following options to troubleshoot the issue:

- Check the Zscaler IP address or hostname configured on the Juniper Mist portal. Ensure that local ID and PSK configured on Zscaler account match with those configured on Juniper Mist portal.
- Ping the Zscaler public IP address from the WAN interface on your device and check if there are responses.
- Using the packet capture tool on the Mist portal. Run a PCAP on the WAN interface with Zscaler Public IP address as filter and check see if there is bidirectional packets. See "Troubleshoot SRX Series Firewalls Using Packet Captures" on page 416.
- If your security device is sending packets to Zscaler and there is no response from the Zscaler, do following:
  - Check if the Zscaler IP address is active
  - Check if any uplink router is blocking the traffic flow
  - Check NAT configuration if applicable

  Note that the packet capture results help you to detect the above issues.

Related Topics

SEE ALSO

- No Link Title
- Configure Hub Profiles for SRX Series Firewalls
- No Link Title

8
CHAPTER

Cellular Edges

IN THIS CHAPTER

Cradlepoint Integration

SUMMARY

Follow this workflow to integrate with your Cradlepoint NetCloud account and to provision your Cradlepoint devices.

IN THIS SECTION

Juniper Networks Mist Cloud supports Cradlepoint 5G cellular adapters, adding to Juniper's wired, wireless and SD-WAN portfolio of supported products.

Video: Now in 60: Cradlepoint integration with SD-WAN

Juniper Mist WAN Assurance supports the following Cradlepoint 5G adapters:

- W1850 Series 5G Wideband Adapter
- W1855 Series 5G Wideband Adapter
- W2000 Series 5G Wideband Adapter
- W2005 Series 5G Wideband Adapter
- W4005 Series 5G Wideband Adapter

Juniper Mist WAN Assurance supports the following Cradlepoint devices:

- E300 Series Enterprise Router

- E3000 Series Enterprise Router
- R1900 router
- AER2200 router
- CBA850 LTE adapter and its sub models

The Cradlepoint 5G adapters provide an LTE WAN backhaul mechanism across many verticals such as retail, warehousing, logistics. Juniper Mist integration with Cradlepoint enables you to use Cradlepoint 5G cellular adapters with Juniper's wired, wireless, and SD-WAN solutions driven by Mist AI.

You can now integrate your Cradlepoint NetCloud Manager account with the Juniper Mist cloud portal. The integration enables you to:

- Manage Cradlepoint devices from the Mist portal including onboarding, assigning devices to a site, and view device inventory details.
- Get visibility into the health, SLE, and Insights into Cradlepoint devices.
- Leverage Marvis, Juniper's virtual network assistant, to get proactive recommendations and self-driving network actions.

The integration enhances Juniper Mist's client-to-cloud user experience by additionally providing insights into the branch WAN adapters, helping the network admins reduce Mean Time to Identify (MTTI).

Prerequisite for Onboarding Cradlepoint Devices

To onboard Cradlepoint devices into Juniper Mist portal, you must:

1. Link Cradlepoint NetCloud account to the Juniper Mist portal.
2. Add account to your Mist organization as a Cellular Edge token.

On adding the Cellular Edge token to your organization in Juniper Mist, supported Cradlepoint devices that are managed by NetCloud are automatically onboarded to the Mist portal.

To integrate your NetCloud account with the Mist cloud, you need:

- Active Marvis subscription
- Deployed Mist device type on the site (minimum one)
- Cradlepoint devices managed from NetCloud

Integrate Cradlepoint NetCloud Account to the Juniper Mist Portal

Use the following procedure to integrate your NetCloud account with the Mist cloud:

1. Generate the required API token (X-CP-API-ID, X-CP-API-KEY, X-ECM-API-ID, and X-ECM-API-KEY) from NetCloud dashboard with following steps:
   a. Log in to the NetCloud Manager.
   b. Click **Tools** in the left-side navigation panel, and then click the **NetCloud API** tab.

   Figure 79: Generating API Keys in NetCloud Dashboard

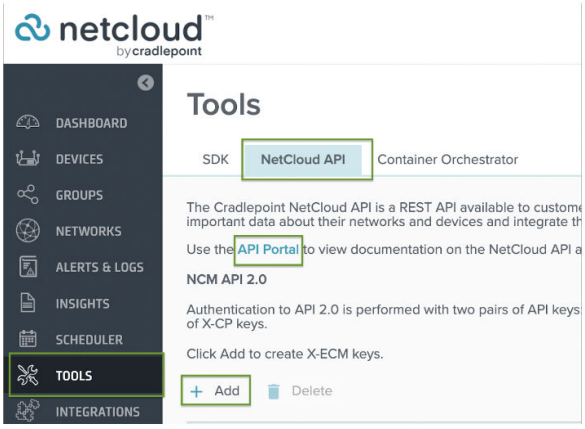

   c. Click the **API Portal link** on the Tools page.
   On the API portal, you can view information about your API keys. Copy **X-CP-API-ID** and **X-CP-API-KEY** and save them for further reference.
   d. Navigate back to **Tools> NetCloud API > Create API Key**
   e. Click **Add** and select a role to associate with the keys and click **OK.**.
   The screen displays the key values. Copy them and save them for further reference.

   Keep ready all key values—**X-CP-API-ID**, **X-CP-API-KEY**, **X-ECM-API-ID**, and **X-ECM-API-KEY**
2. Log on to Juniper Mist portal and complete the following steps:
   a. On Juniper Mist portal, go to **Organization** > **Settings**.
   b. Scroll to **Third Party Token** pane and click **Create**.
   c. Select **Cellular Edge** and enter the keys you saved in previous procedure.

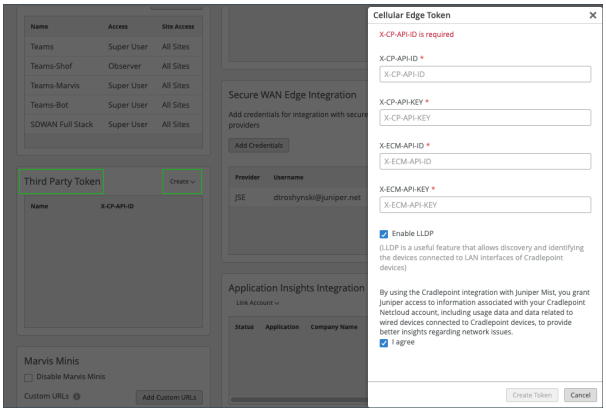

By default, the LLDP option is enabled and is required for automatic site assignment. LLDP allows discovery and identification of devices connected to the LAN ports on Cradlepoint devices.

> **NOTE**: The Cradlepoint R1900 router does not support LLDP today. Please contact your Cradlepoint account manager to request further information about future plans to add LLDP to this device.
>
> For LLDP functionality on W1850 and W1855 Cradlepoint devices, a Netcloud OS version newer than August 2024 is recommended.

d. Click **Create Token**.
You can see the entry for CradlePoint in the **Third Party Token** table after a successful integration.

View Cradlepoint Inventory

To view the inventory details of the Cradlepoint devices onboarded to the Mist portal:

1. On Juniper Mist portal, go to **Organization > Inventory > Cellular Edges**.
The Mist Dashboard's Inventory displays the CradlePoint Inventory for the supported devices that are synchronized from NetCloud.
2. Click **Sync Cellular Edges** on the Inventory page if the Cellular Edge page does not display any data.

Figure 80: View CradlePoint Devices in Cellular Edge Inventory

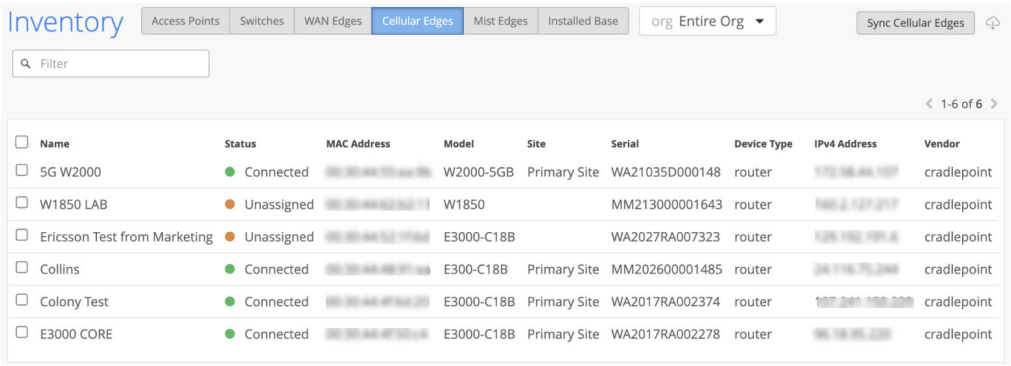

Assign Cradlepoint Devices to a Site

To assign a Cradlepoint device to site:

1. On Juniper Mist portal, go to **Organization > Inventory > Cellular Edges.**
2. Select the device and click **More > Assign to Site** on the page

Figure 81: Assign Cradlepoint Devices to a Site

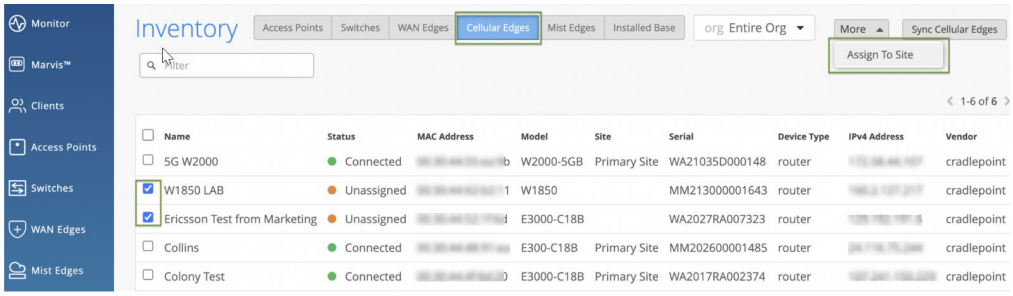

Auto-Provisioning

You can use auto-provisioning to automatically assign Cradlepoint devices to sites. Auto-provisioning allows you to define rules for assigning a device to a site. Mist automatically assigns the device to a site based on the auto-provisioning rules that you configure.

You can set up auto-provisioning from the **Organization** > **Settings** page in the Juniper Mist™ portal. See Automatically Assign Devices to a Site.

Juniper Mist™ can automatically assign a site based on:

- Device name—You can configure rules to derive the site name based on the device name (Cellular Edge Name).

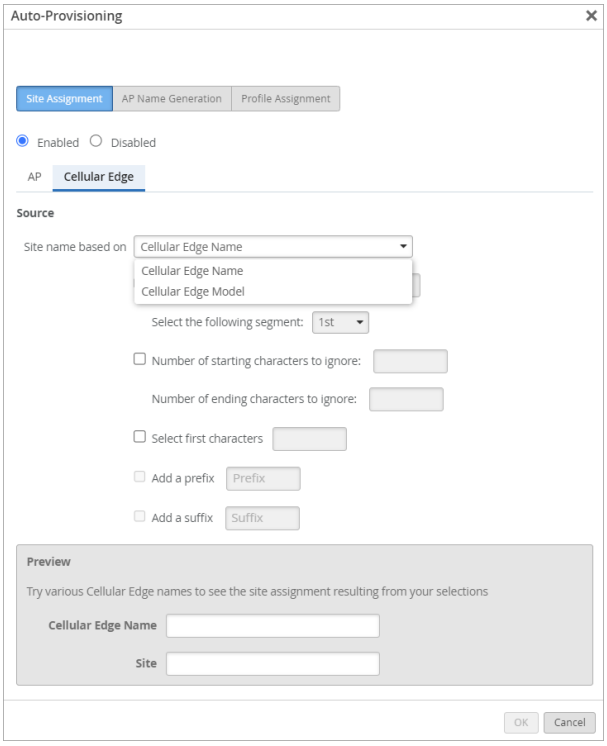

- Device model—You can assign devices to sites based on the device model (Cellular Edge Model).

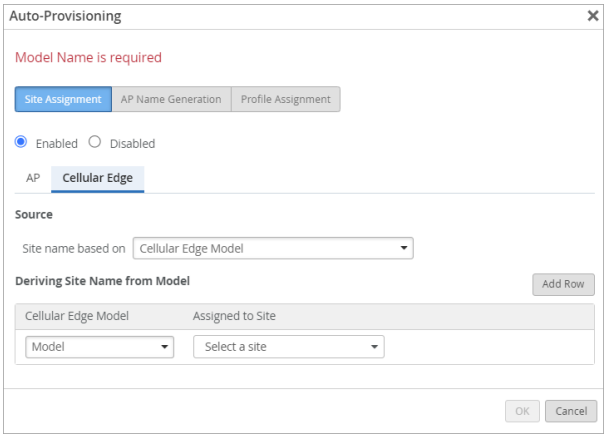

If you do not configure auto-provisioning, Mist uses the LLDP information about the Cradlepoint device for site assignments. Mist first identifies a Juniper device that is connected to a port on the Cradlepoint device and is also in the same organization. Mist then assigns the Cradlepoint device to the same site that the Juniper device is assigned to.

For the LLDP-based automatic site assignment to work, you must ensure that LLDP is enabled on both the Cradlepoint and Juniper devices. Automatic site assignment is applicable only for devices connecting to the Juniper Mist dashboard for the first time.

View Cradlepoint Device Details

To view Cradlepoint device details:

1. On Juniper Mist portal left-navigation bar, go to **WAN Edges > Cellular.**
2. Select the required site from drop-down menu to view Cradlepoint devices assigned to that site.

Figure 82: View Cradlepoint Device Details

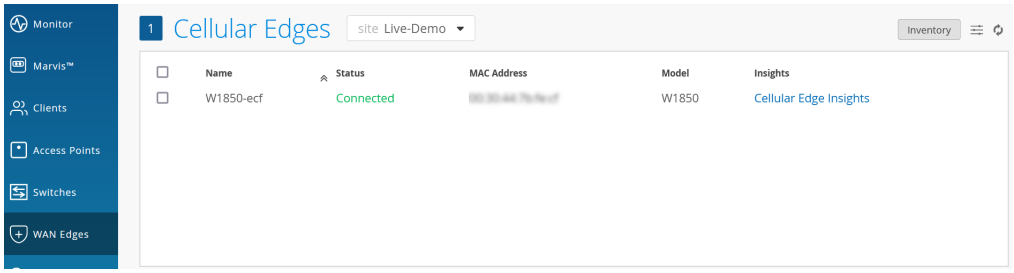

Click the device name to open the device details page.

Figure 83: View Cradlepoint Device-Specific Details

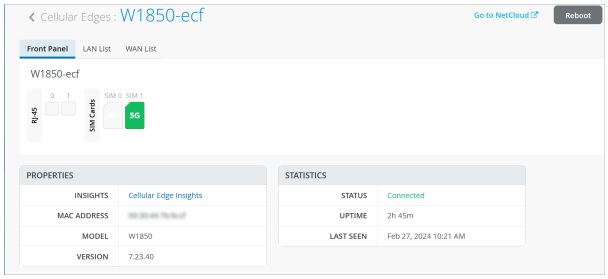

The page lists device-specific information such as properties, statistics, SIM list, and front panel details. You have an option to navigate to the NetCloud dashboard by clicking **Go to NetCloud**.

Get Cradlepoint Device Insights

To get details of events related to cellular edge devices, select **Monitor** > **Service Levels** > **Insights**, and then select **Cellular Edge** as the context.

Figure 84: Select Cellular Edge to Display Insights

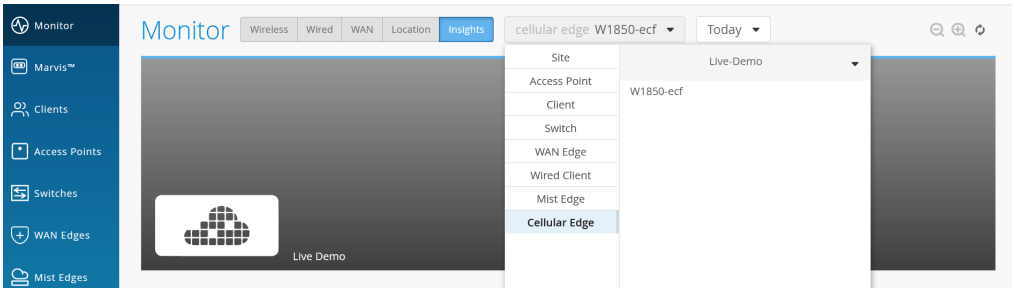

The page displays Cellular Edge device events, properties, and port details. For Cradlepoint devices, the events are ingested from NetCloud.

Figure 85: Cellular Edge Device Events

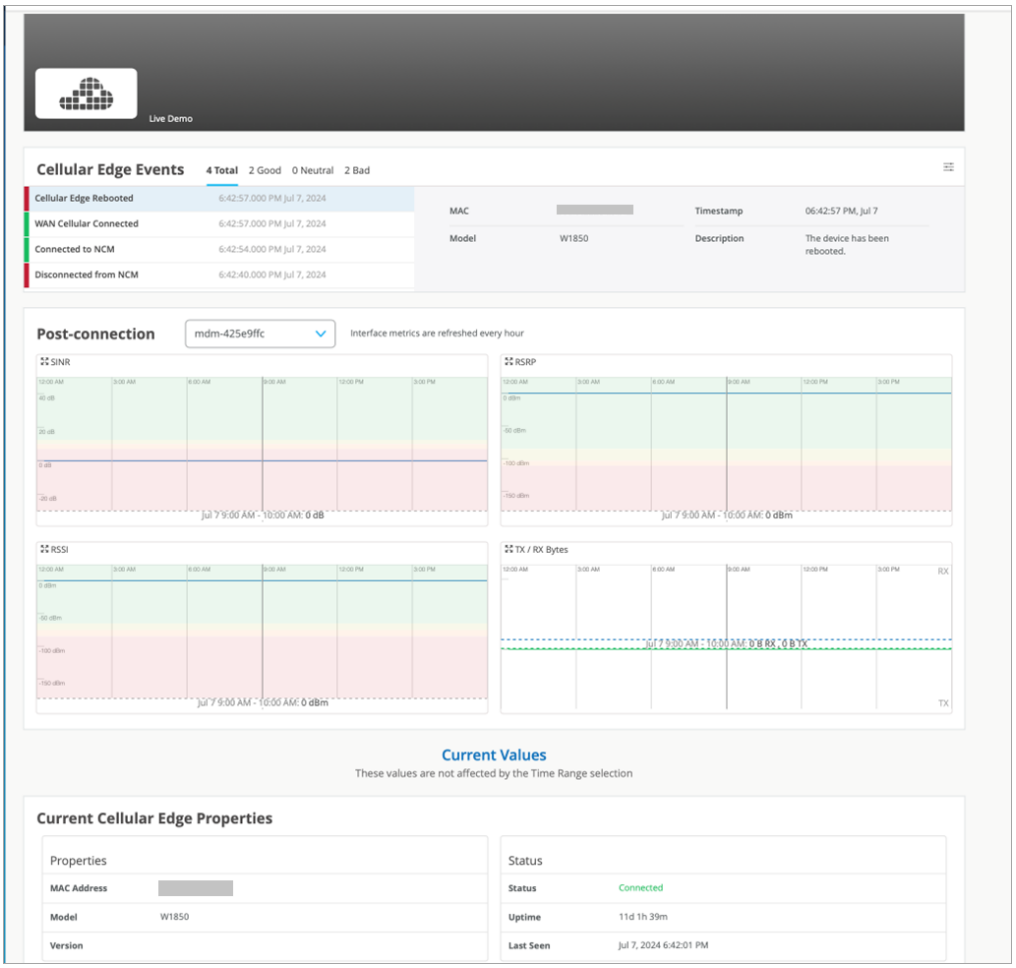

The time series data charts include:

- SINR—Signal-to-Interference-plus-Noise Ratio (SINR) graph compares the level of the received signal to the level of background noise and interference.
- RSRP—Reference Signal Received Power (RSRP) represents a measure of the received power level in an LTE network. Supported range: -200 through 10 dBm.
- RSSI—Received signal strength indicator (RSSI) is a measurement of the AP radio signal and is typically measured by the client. The scale runs from -100 dBm (weakest) to 0 dBm (strongest).
- RSRQ—Reference Signal Received Quality (RSRQ) measures the quality of the reference signal received by the Cellular Edge device. RSRQ measures the strength of the signal received, by assessing the interference and noise level in the signal.

You can find the supported device events currently ingested at List Other Device Events Definitions.

Configure Alerts for Cellular Edge Devices

The Alerts Dashboard gives you visibility into issues with Cradlepoint devices. The dashboard provides information about all alerts that you enable on the Alerts Configuration page. You can also enable e-mail notifications for issues that you want to monitor closely. For information about configuring alerts, see Configure Alerts and Email Notifications.

You can configure alerts for issues such as

- Device disconnected from NCM
- Firmware upgrade
- Login failure
- Change of WAN Cellular service type
- Device reboot

For a list of alerts that you can enable for your Cradlepoint device, see Juniper Mist Alert Types.

Here is a sample screenshot that shows the alerts for a Cradlepoint device:

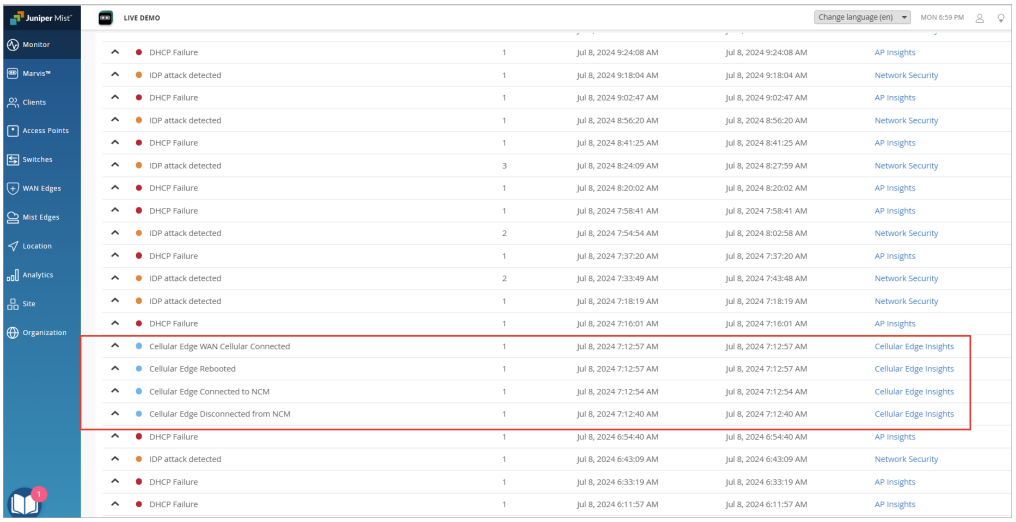

The alerts are based on the events generated by NetCloud Manager (NCM). When you integrate the Cradlepoint NetCloud Account with Juniper Mist, Mist configures the list of events that will be sent as webhooks to the Mist cloud. The ingest service consumes these webhooks and generates events and alerts. These alerts can be forwarded as e-mails or webhooks from the Mist cloud to an external destination.

You can view the list of events in the **Alert & Logs**>**Set Up Alerts** page in NetCloud Manager as shown in the following example. You can also see the webhook listed on the right.

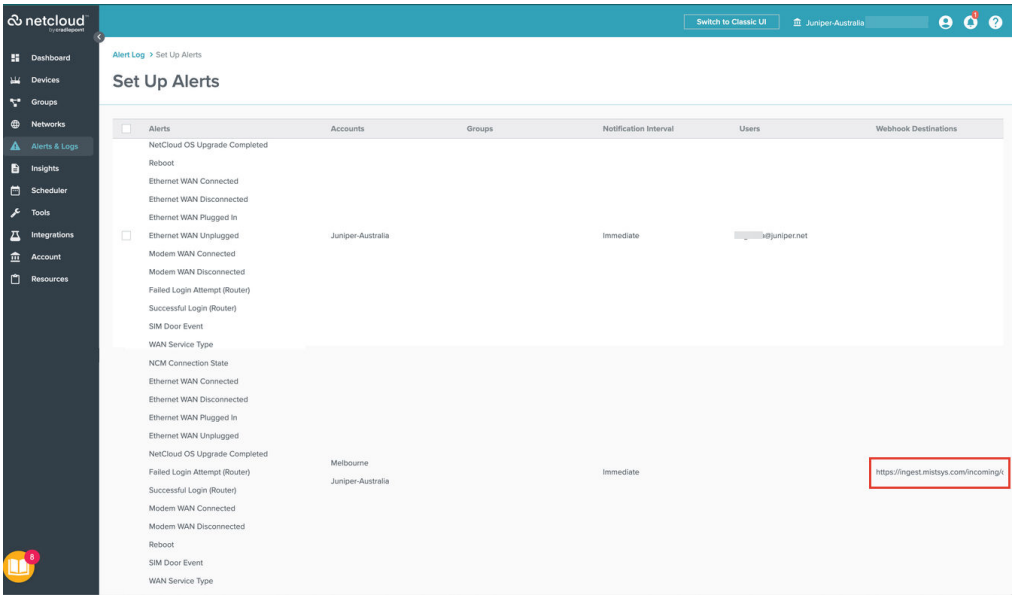

You can view the complete webhook URL in the **Alert & Logs**>**Set Up Alerts**>**Manage Destinations** page.

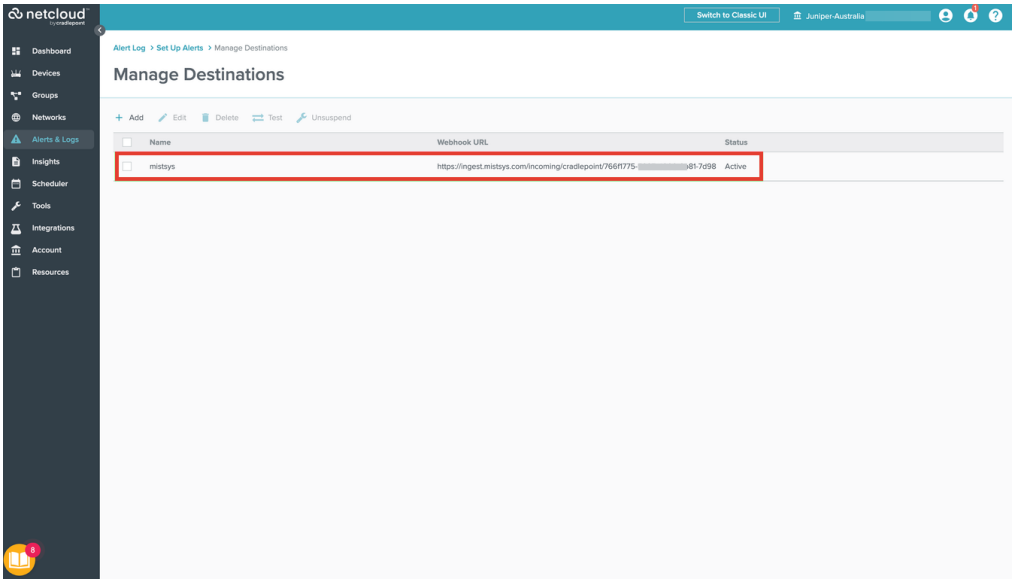

Unlink Cradlepoint Account from Juniper Mist Portal

1. On Juniper Mist portal, go to **Organization > Admin> Settings**.
2. Scroll-down to **Third Party Token** pane.
3. Click the Cradlepoint entry in the table.

   The Cellular Edge Token window opens.

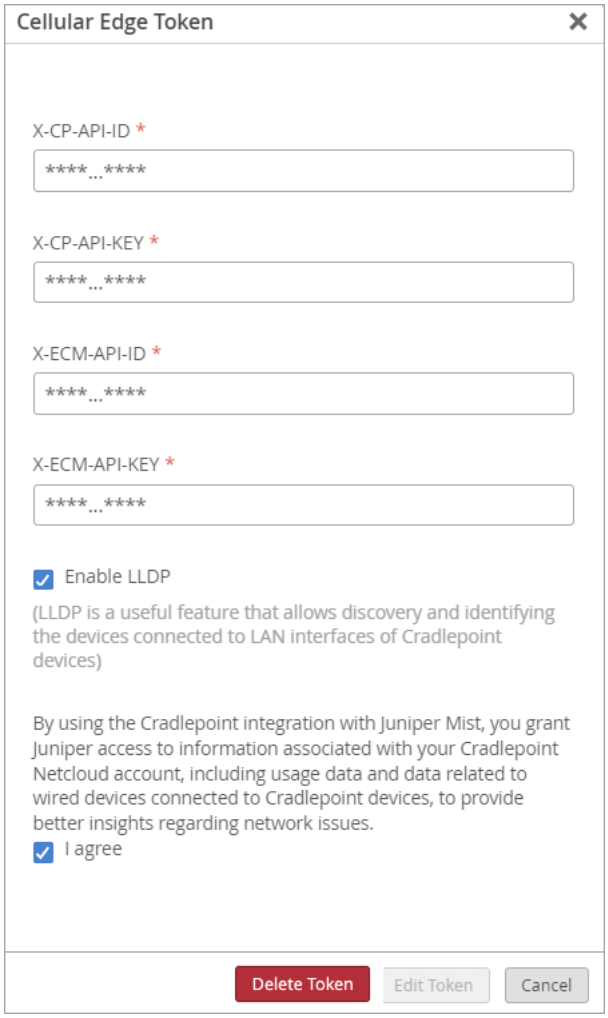

4. Click **Delete Token**.
   Table will not have any key entry post delete and Cellular Edge inventory also does not display any Cradlepoint devices.

9

CHAPTER

Monitor and Troubleshoot

SUMMARY

Use the information in this chapter to monitor your WAN and troubleshoot issues.

IN THIS CHAPTER

What Do You Want to Do?

Table 59: Top Tasks

If you want to...	Use these resources:
Learn about WAN monitoring and troubleshooting. *Get started with WAN monitoring, Service Level Expectations, and more.*	"WAN Assurance Monitoring, SLE, and Troubleshooting Overview" on page 313
Investigate specific issues. *Use the other topics in this chapter to get help with other issues.*	Troubleshooting topics in this chapter
Explore other aspects of troubleshooting and AI-driven operations. *Take a deeper dive into the various dashboards, tools, and Marvis features that support AI-driven operations.*	Juniper Mist AI-Driven Operations Guide

WAN Assurance Monitoring, SLE, and Troubleshooting Overview

SUMMARY

Get familiar with various Juniper Mist features that you can use to investigate and resolve issues and to optimize performance on your network.

IN THIS SECTION

The Juniper® Mist™ WAN Assurance Monitoring, SLE, and Troubleshooting chapter is for system administrators and technical support who maintain enterprise SD-WAN networks.

Driven by Mist AI, the Juniper SD-WAN solution simplifies the monitoring and troubleshooting of the WAN edge. Juniper Mist WAN Assurance does this by consistently and proactively monitoring numerous variables that impact a user's network experience. Both the Juniper® Session Smart™ Router and Juniper® SRX Series Firewalls WAN Assurance platforms have unique solutions to the WAN edge and overlay that impact provisioning and deployment. Because of this, those monitoring and troubleshooting actions are platform-specific, and understanding your WAN Assurance platform is crucial.

You'll find separate topics for the Juniper® Session Smart™ Router and Juniper® SRX Series Firewalls WAN edge Monitoring, SLE, and Troubleshooting.

WAN Edge Monitoring

In the WAN edge monitoring guides, you'll explore the most efficient ways to monitor your WAN edge device in the Mist UI following your initial deployment phase. The Session Smart Secure Vector Routing WAN Assurance solution monitors liveness, jitter, latency, loss, and mean opinion score (MOS) to inform those user minutes. The SRX Series Firewall monitors the utilization of the IPsec tunnels that make up the overlay on the SRX Series WAN Assurance solution.

- "Troubleshoot Session Smart Router Deployed as WAN Edge " on page 403
- "Troubleshoot Disconnected SRX Series Firewalls" on page 411

WAN Edge SLEs

The Juniper Mist WAN Assurance solution simplifies the entire network diagnosis process. WAN edge devices track metrics to assess the bandwidth usage and the health of the device, links, and applications. The result is service-level experiences (SLEs), which show the impact that issues have on the user experience. From the SLE dashboard, you can follow hyperlinks to view root causes and additional metrics. For more information, see "WAN SLEs" on page 372./

Video: NOW in 60: Mist SLEs for WAN Assurance

WAN Edge Troubleshooting

Juniper Mist WAN Assurance troubleshooting tools give system admins proactive insights alongside traditional tools for diagnosing WAN edge issues and identifying things on a per-application level down

to granular interface and port metrics for both Session Smart Routers and SRX Series Firewall WAN edge devices. Users will see the most difference in their guides depending on their platform when WAN edge troubleshooting. The Session Smart Secure Vector Routing creates the overlay, while a unique implementation of Bidirectional Forwarding Detection between Session Smart devices helps troubleshooting. You'll use Mist UI dashboard tools and leverage Juniper Mist Application Visibility, Application SLE, and Marvis to troubleshoot your Session Smart WAN edge.

For example, at times, within an overlay tunnel connecting a Hub to a spoke, not all the advertised routes may be visible on the overlay, which can lead to traffic being unable to traverse the tunnel. In such cases, we recommend following actions:

- Ensure the overlay advertisement is enabled for the networks that need to be advertised. Go to **Organization > WAN > Networks** and check if **Advertised via Overlay** is enabled for the specific network.
- Verify that the overlay tunnel is configured properly and that the correct routes are being advertised.
- Check the configuration on both the hub and spoke devices and ensure that the overlay tunnel is properly configured, including the correct IP addresses and route advertisements.
- Check the routing configuration on both the Hub and spoke devices. Verify that the routing tables on both devices include the necessary routes for the overlay tunnel, including routes to each other's network segments.
- Check the firewall configuration on both the hub and spoke devices. Ensure that the firewall rules are properly configured to allow traffic to pass across the overlay tunnel.
- Check the MTU (Maximum Transmission Unit) settings on both the hub and spoke devices. Verify that the MTU is set to the same value on both devices and that it is not set too high, which could cause fragmentation and slow down traffic.
- Check the connectivity between the hub and spoke devices. Verify that there are no network connectivity issues between the two devices, including issues with firewalls or NAT (Network Address Translation) devices in between.
- Check the logs on both the Hub and spoke devices for any error messages related to the overlay tunnel. Use the **show log** command to view the system logs and look for any errors related to the overlay tunnel.
- Check the firewall filters to make sure that traffic is allowed to pass through the tunnel.
- If you are still unable to resolve the issue, try restarting both the Hub and spoke devices

Troubleshooting gateways on your SRX Series Firewall uses the powerful and versatile Junos OS with the same Application Visibility, Application SLE, and Marvis to diagnose WAN edge connectivity.

- "Troubleshoot Session Smart Router Deployed as WAN Edge " on page 403

- "Troubleshoot Disconnected SRX Series Firewalls" on page 411

Monitor SRX Series Firewall Deployed as WAN Edge

SUMMARY

Use the WAN Edges page, the Insights page, and the Alerts page to quickly find device information, event details, and alerts for your SRX Series Firewalls.

IN THIS SECTION

In monitoring a Juniper® SRX Series Firewall deployed as a WAN edge device, you'll explore the most efficient ways to monitor your WAN edge device in the Juniper Mist™ portal following your initial deployment phase.

Monitoring WAN Edges

From the left menu, select **WAN Edges WAN Edges**to view basic device monitoring information. Notice the organization name at the top of the portal. This is the largest container and represents your entire organization. Beneath the organization name, you can see your site devices in either a List format or a graphical Topology format.

In this example, Live Demo is your organization, and LD_CUP_SRX_1 is the selected WAN edge device.

Figure 86: Accessing WAN Edges Page

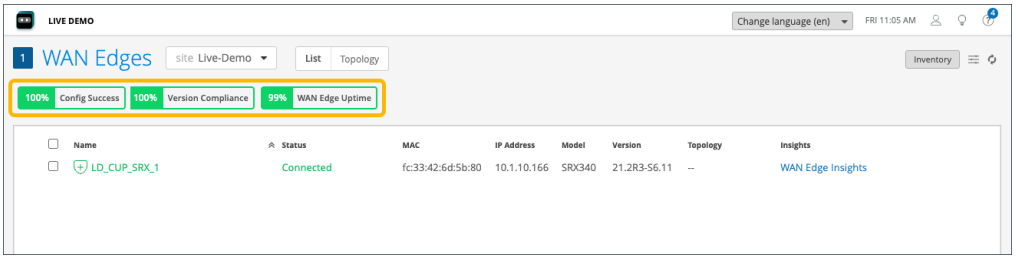

The tiles across the top provide high-level information:

- **Config Success**—Percentage of online WAN edges with successful configuration.
- **Version Compliance**—Percentage of WAN edges that have the same software version per model.
- **WAN Edge Uptime**—Percentage of time a WAN edge was up during the past seven days, averaged across all WAN edges.

Figure 87: WAN Edges List View

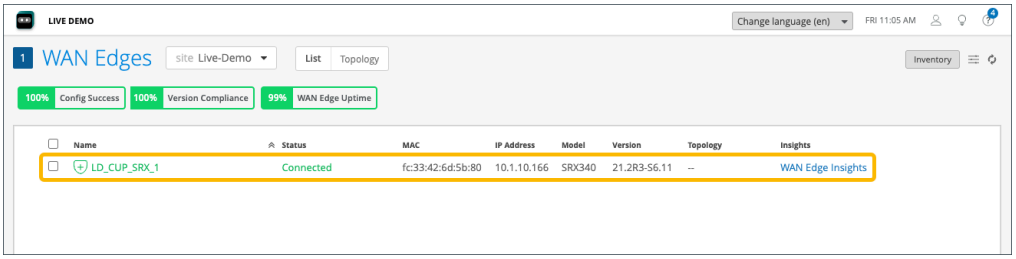

Beneath, you'll find WAN edge device details as shown in Table 60 on page 317.

Table 60: WAN Edge Device Details

Fields	Description
Name	Name
Status	Connected or disconnected
MAC	MAC address
IP Address	IP address
Model	Juniper Networks® SRX Series Firewalls or Juniper® Session Smart™ Routers.
Version	SRX Software Version
Topology	Hub or Spoke
Insights	Provides a direct link to the WAN Edge Insights page.

The **Topology** format presents the same information as the **List** view. For example, if you hover over the LD_CUP_SRX_1 device, you'll see the same information that displayed in the List view.

Figure 88: WAN Edges Topology View

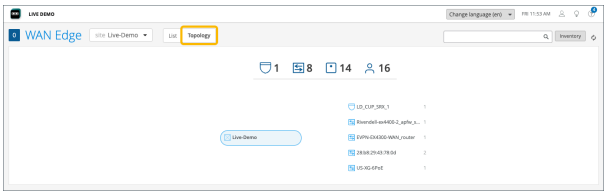

On both the List and Topology view, selecting your WAN edge device (LD_CUP_SRX_1 in this example) brings you to its Device Information page. The Device Information page provides different categories of monitoring information for your WAN edge device.

Figure 89: WAN Edges Device Information Page

The first thing you'll notice on the Device Information Page are details about the WAN edge device you selected, (LD_CUP_SRX_1 in our figure). The information includes a graphical front view of the device ports and baseline status information such as CPU and memory utilization.

Figure 90: WAN Edges Device Information - Interfaces

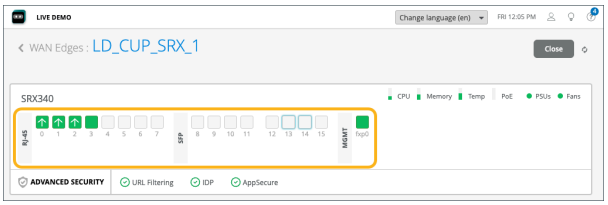

For each Gigabit Ethernet interface, you'll find link information.

Figure 91: WAN Edge Device Information Page - Details

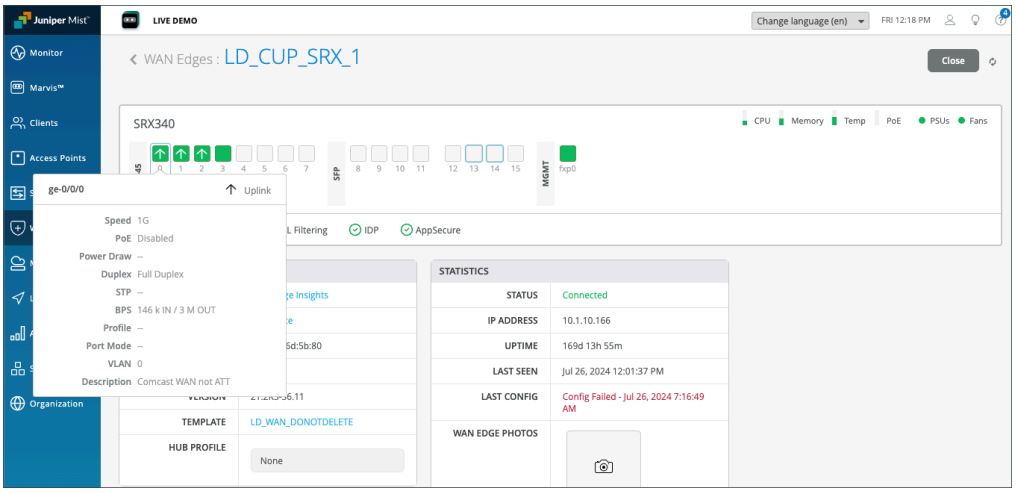

Table 61: Link Information for Gigabit Ethernet Interface

Fields	Description
Speed	Rated speed
PoE	Enabled or disabled
Power Draw	Measured PoE power draw
Duplex	Full or half
STP	True or false
BPS	Bits/second
Profile	The name of the Port profile assigned to the port
Port Mode	The mode of the port profile configuration (Trunk, Access, Port Network, or VoIP Network)
VLAN	VLAN tag
Description	Interface description

The CPU, Memory, and other status icons indicate how your device behaves. Hover over each status icon for deeper insights.

Figure 92: WAN Edges - CPU, Memory, and other status icons

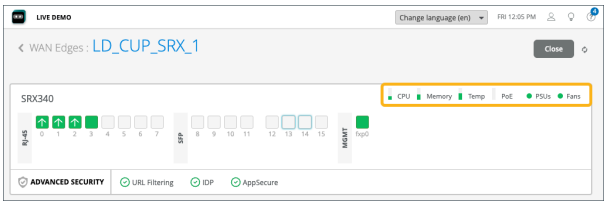

Advanced Security information is listed below the device ports with a check mark or an X, indicating whether URL filtering, intrusion detection and prevention (IDP), or AppSecure (for application visibility) is active on this device. Here, URL filtering, IDP, and AppSecure are active with the green check mark.

Figure 93: Advanced Security Details

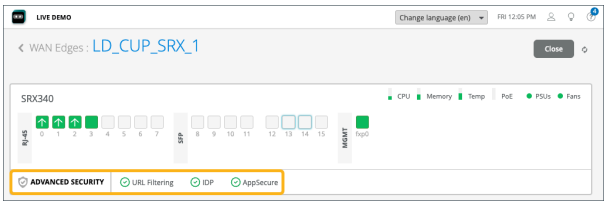

Below the port information and security section, you'll find generalized data for your WAN edge device, including:

Figure 94: WAN Edge Device Properties

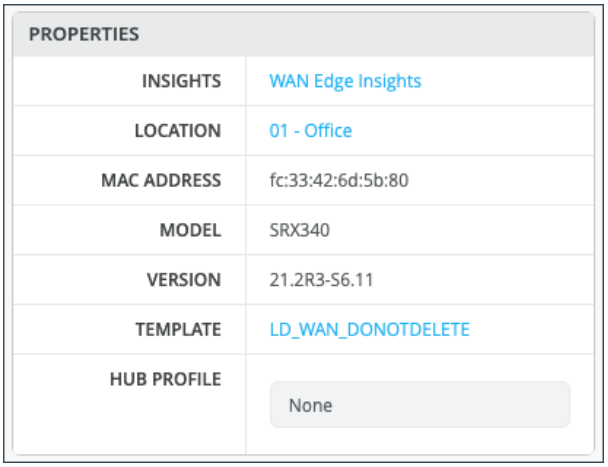

Properties contains generalized platform-related information.

Table 62: WAN Edge Platform-Related Details

Field	Description
Insights	Provides a direct link to WAN Edge Insights.
Location	Provides floorplan information.
MAC Address	MAC Address for the SRX device.
Model	Indicates if model type is SSR or SRX.
Version	Version of SRX Software.
Template	The WAN edge template applied to the device.
Hub Profile	The Hub Profile applied to the device.

Statistics displays action information about your platform.

Figure 95: WAN Edge Device Statistics

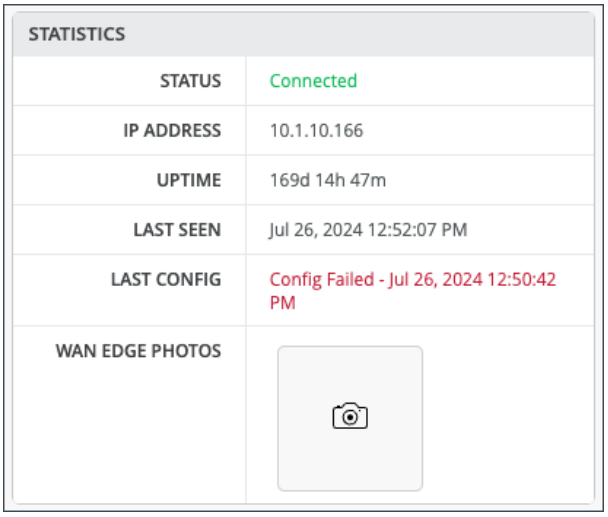

Table 63: WAN Edge Device Statistics

Field	Description
Status	Connected/Disconnected

IP Address	The IP address of the WAN edge device
Uptime	Day/Hour/Min uptime information
Last Seen	Last login
Last Config	Last Commit
WAN Edge Photos	Photos of the WAN edge device

If you configured DHCP servers on the WAN router itself, there will also be a DHCP Statistics pane with information about the leased IPs.

- **DHCP Statistics** presents IP information related to dynamic distributed IP addresses.

Figure 96: WAN Edge Device DHCP Statistics

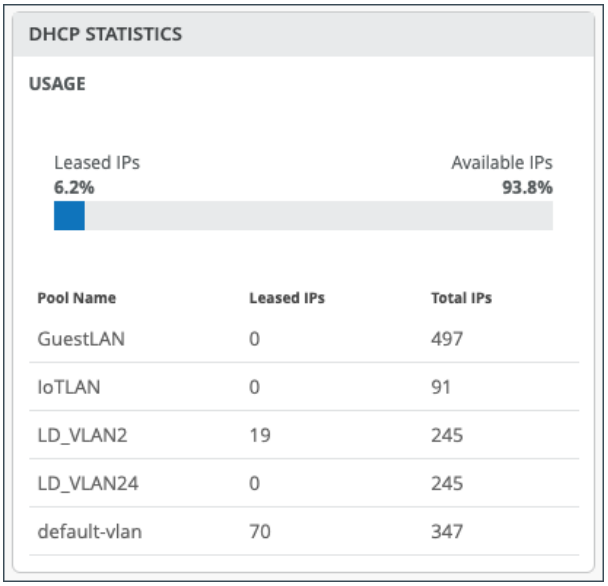

Table 64: WAN Edge Device DHCP Statistics

Field	Description
Usage	The total figure presented as a percentage of Leased and Available IPs.
Pool Name	The name for given pool of addresses.

Leased IPs	Number of used IP addresses in each pool.
Total IPs	Total number of available IP addresses in each pool.

Scrolling down the Device Information page, you'll find configuration information for your WAN edge. Usually, WAN edges inherit templates or profiles. However, you can make individual changes to the configuration to be pushed to the device. In this example, a standalone WAN Edge Template was used.

Figure 97: WAN Edge Configuration: Standalone

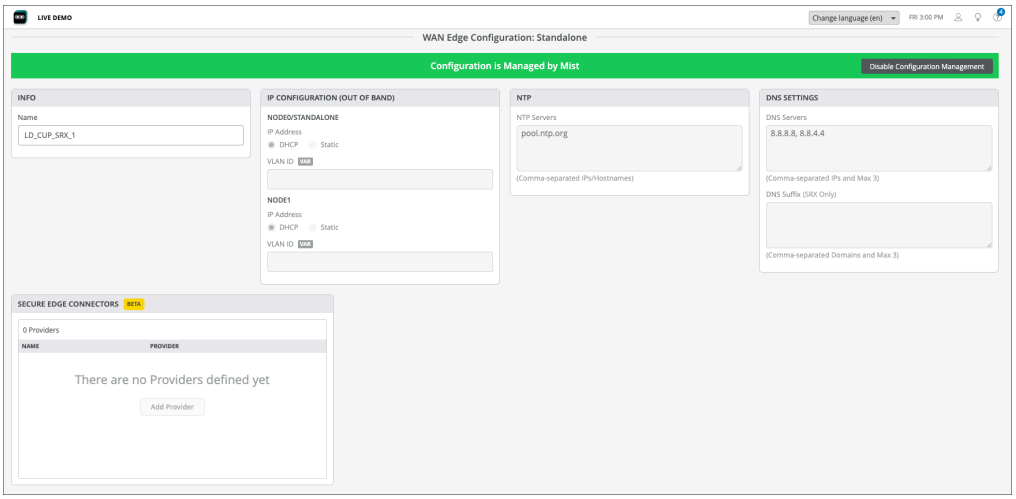

Table 65: WAN Edge Configuration: Standalone

Field	Description
Info	The name of the SRX device.
IP Configuration	Node0/standalone DHCP/Static, VLAN ID, node 1 DHCP/Static, VLAN ID.
NTP	Time Servers IP/Hostnames.
DNS Settings	DNS Servers, DNS Suffix (SRX only DNS suffix info).
Secure Edge Connectors (BETA)	Provider for the Secure Edge Connector.

Scrolling past the configuration, you'll find information for your connected WANs and LANs.

Figure 98: WAN Details

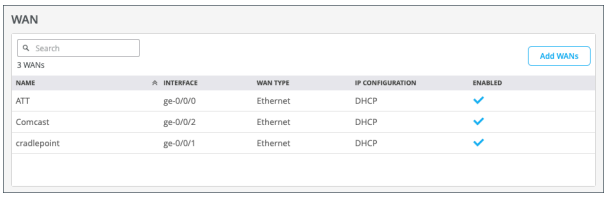

Table 66: WAN Details

Field	Description
Name	Selected WAN Interface Name
Interface	Supports one of these interfaces for aggregation: ge-0/0/1, ge-0/0/1-5, or reth0.
WAN Type	Ethernet, DSL (SRX Only), or LTE
IP Configuration	DHCP, Static, or PPPoE
Enabled	Check mark indicates that the interface is enabled.

Figure 99: LAN Details

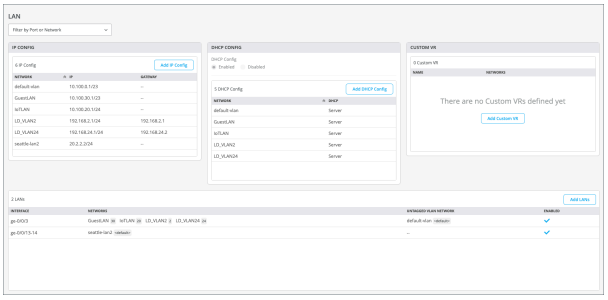

Table 67: LAN Details

Field	Description
IP Config	Network, IP Address, Prefix Length.
DHCP Config	Server or Relay.

Custom VR	A virtual router that you can configure to be used in automatic route leaking.
LANs	• Interface—Supports one of these interfaces for aggregation: ge-0/0/1, ge-0/0/1-5, or reth0. • Networks—Networks that participate in the LAN. • Untagged VLAN Network—Untagged VLAN networks that participate in the LAN (SRX only). • Enabled—Check mark indicates that the interface is enabled.

Scrolling down, you have sections for Traffic Steering, Application Policies, and Routing (OSPF, BGP, Routing Policies, and Static Routes).

The Traffic Steering and Application Policies sections show how you use the SRX Series Firewall to create rules for path preference and routing behavior. Note that on the SRX Series Firewall deployed as a WAN edge, the Application Policy and Traffic Steering path determine destination zones and must be assigned.

Traffic Steering enables you to define different paths that traffic can take to reach its destination. Traffic Steering policies allow you to specify the paths for traffic to traverse, as well as the strategies for utilizing those paths.

Figure 100: Traffic Steering

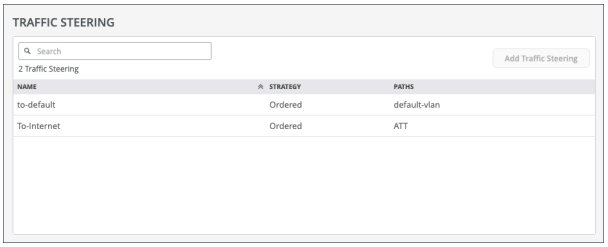

Table 68: Traffic Steering

Field	Description
Name	The name of the Traffic Steering policy.
Strategy	Ordered, weighted, ECMP.
Paths	LAN, WAN, Untagged VLAN (SRX only).

Application Policies are security policies in the Juniper WAN Assurance design, where you define which network and users can access which applications, and according to which traffic steering policy. You must create Networks, Applications, and establish Traffic Steering profiles to define an Application Policy. These elements become matching criteria to allow access to or block access from applications or destinations.

Figure 101: Application Policies

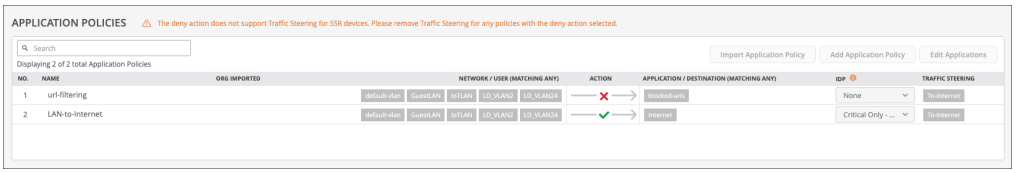

In the Juniper Mist™ cloud portal, the Networks or Users setting determines the source zone. The Applications and Traffic Steering settings determine the destination zone. Traffic Steering paths determine the destination zone in Juniper Networks® SRX Series Firewalls, so ensure that you assign Traffic Steering profiles to the Application Policies.

Table 69: Application Policies

Field	Description
Number	Ordered Policy Number
Name	Application policy name
Org Imported	Indicates if the policy was pushed down from the Organization level to the Site.
Network/User (Matching Any)	The “source” of your traffic
Action	Allow/Block
Application/Destination (Matching Any)	The “destination” for your traffic.
IDP	Indicates IDP/URL filtering (requires separate license)
Traffic Steering	Indicates path for traffic

Open Shortest Path First (OSPF) is used to determine the best path for forwarding IP packets. OSPF segments a network to improve scalability and control the flow of routing information. See "OSPF" on page 172 .

Figure 102: Routing: OSPF Areas and OSPF Configuration

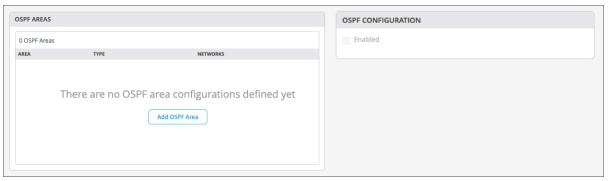

Table 70: Routing: OSPF Areas and OSPF Configuration

Field	Description
Area	The identification area that your OSPF network or SRX Series Firewall belongs to.
Type	This is the OSPF Area type. Select Default (Area 0), Stub, or Not So Stubby Area (NSSA).
Networks	The name of your OSPF network.
Enabled	Selecting this check box causes the **Enable OSPF Areas** button to become selectable.

You can configure Border Gateway Protocol (BGP) for your SRX Series Firewall deployed as a WAN edge device. You can also manually add a BGP Group here.

Figure 103: Routing: BGP

Table 71: Routing: BGP

Field	Description
Name	BGP Name
Peering Network	The type of network being used for your BGP peering (WAN or LAN).
Type	Type of BGP Route (Internal or External)
Local AS	Autonomous System Number
Export	Export Route
Import	Import Route
Neighbors	Neighbor Route
Neighbor AS	Autonomous System Number for Neighbor Route

The Routing Policies section enables you to configure path preference and allows you to determine traffic behavior.

Figure 104: Routing: Routing Policies

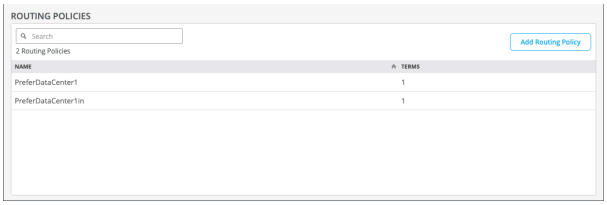

Table 72: Routing: Routing Policies

Field	Description
Name	The name of your routing policy.
Terms	These are the policy conditions such as prefix, routing protocol, and actions.

Static routes allow you to manually define the routes that your SRX Series Firewall deployed as a WAN edge device will use.

Figure 105: Routing: Static Routes

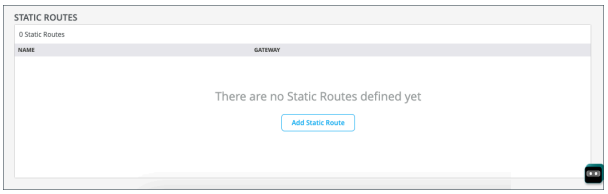

Table 73: Routing: Static Routes

Field	Description
Name	The name of your static route.
Gateway	The gateway that your static route will use when routing traffic.

View Device Information and WAN Edge Insights

IN THIS SECTION

WAN Edge Insights

The Properties pane for your selected WAN edge links to **WAN Edge Insights**. Click **WAN Edge Insights** for the next level of information about your WAN edge device.

Figure 106: WAN Edge Insights

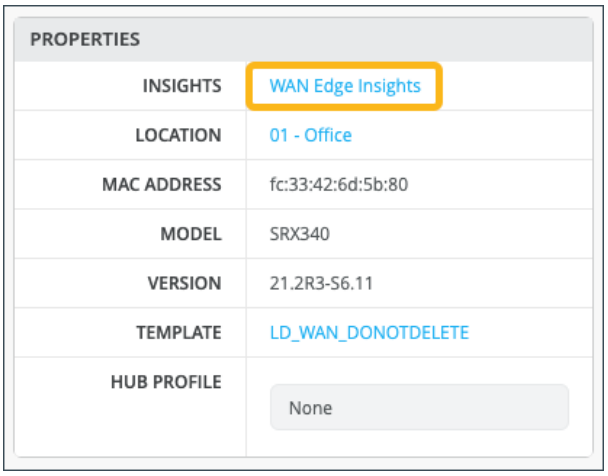

Next to the selected WAN edge (LD_CUP_SRX_1) on the Insights page, you can select a time frame for selected information. The default view is **Today,** but this can be set to a customized date or range of dates. Below this, you find (when the site location information is configured) where this WAN edge is configured via a street map.

Figure 107: WAN Edge Insights-Select Time Duration

With your time frame selected, **WAN Edge Events** displays a time line of the traffic through the WAN edge during your specified time, and also displays a list of events.

Select a specific event in the listed WAN Edge Events for greater detail of the **Good**, **Neutral**, and **Bad** events.

Figure 108: WAN Edge Events Timeline

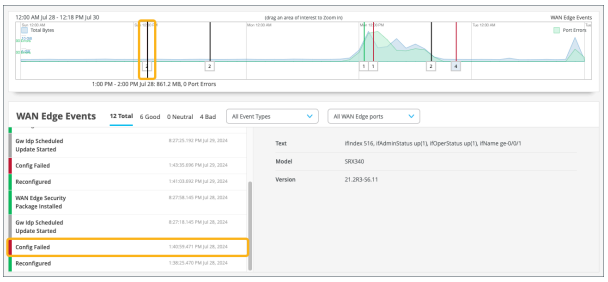

Your selection expands and displays detailed information about the selected time.

For a detailed portion of time, select a window of time with the mouse cursor. By doing this, you're able to adjust the window of events and isolate specific **Good**, **Neutral**, and **Bad** occurrences that happened on your network. With a smaller section, you'll get a more detailed view of that period.

Figure 109: WAN Edge Events Timeline Details View

Scroll down on the WAN Edge Events page for deeper insights within your selected period.

Figure 110: WAN Edge Events page

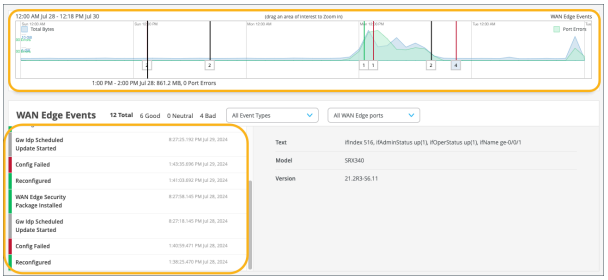

In the WAN Edge Events, you can narrow down the type of event by selecting a modifier in the Event Type drop-down menu. You can also filter your search by limiting the event types to a specific port.

Figure 111: WAN Edge Events Page

On the WAN Edge Events page, you can also view reports on applications on the Applications pane. On this pane:

- You can use the applications pane to monitor and troubleshoot specific application behavior.
- You can hover over the App Name to see more details about the services.
- You can view a client's use of a particular application by clicking the Clients tab.

NOTE: Ensure you've had a few hours for these metrics to be populated following initial deployment.

Figure 112: Table Capacity

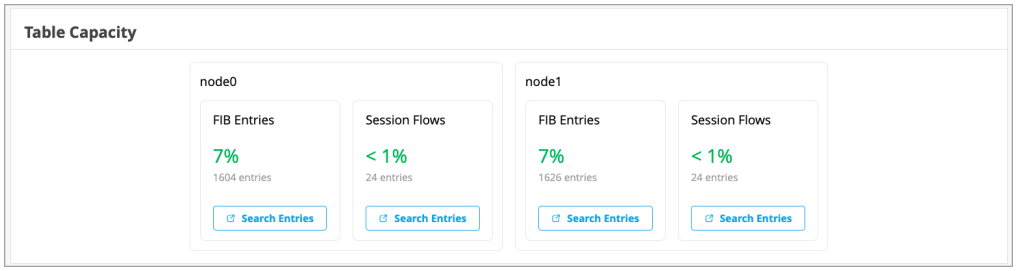

The WAN Edge Insights page provides the following indicators in the Table Capacity section:

- FIB Entries: Displays the current number of FIB entries and the percentage of utilization; essentially showing how much of the available FIB space is currently being used.
- Session Flows: Displays the current number of active sessions and the percentage of session flow utilization based on the device's capacity.

In the case of a high availability cluster, Table Capacity indicators are displayed for each node.

You can also click the **Search Entries** button under each metric to open a shell view in a new window where you can search for entries after specifying filters.

Figure 113: Applications

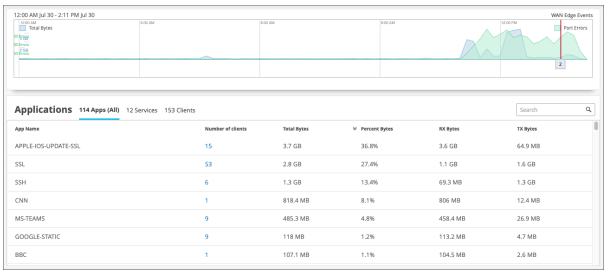

In the **Number of Clients** column, you can click on the number to see more information about the clients using the application such as the Client name, MAC Address, IP Address, Device Type, and Bytes being used.

Figure 114: Clients Using Application

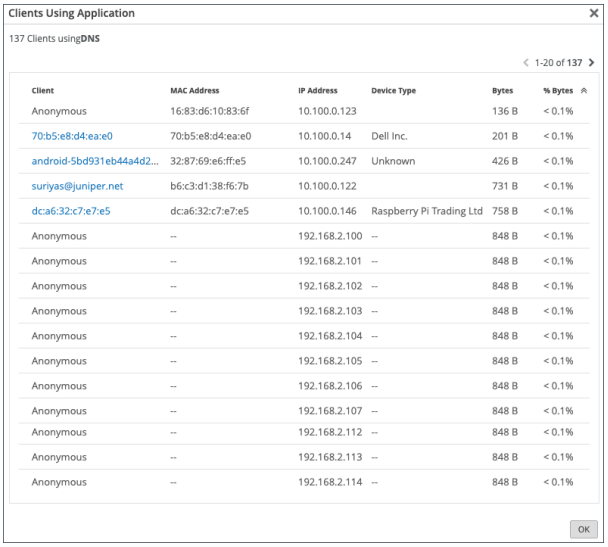

NOTE: For SRX Series Firewalls deployed as a WAN edge running a DHCP server, clients using that application will display a HostName in the Client column if available. Otherwise, the MAC address will be displayed. Device Type and MAC Address columns will be populated as well.

Back on the Applications pane, you can click the **Clients** tab to see how much bandwidth a particular client is using. You can click the number in the **Number of applications** column to see more information.

Figure 115: Applications for Client

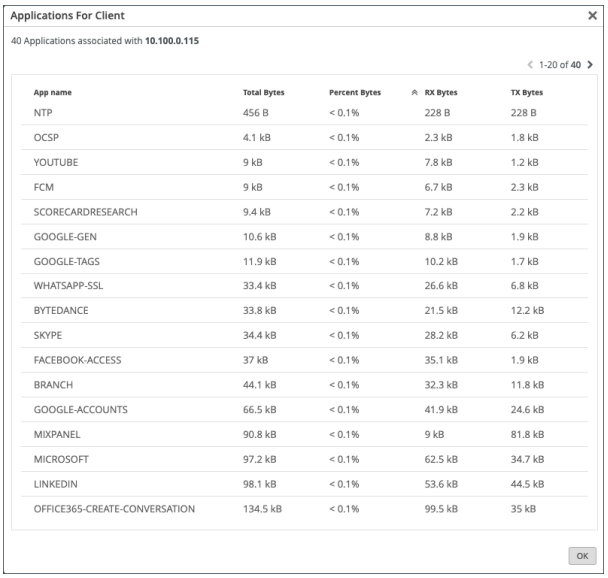

The Application Path Insights (BETA) section shows you which applications are using the most bandwidth according to the selected Application Policy and Network. It displays the effective application flow over the path for the selected Application Policy. You can also change the Data Type to Sessions to see the number of sessions occurring per application. Hover over a section of the graph to view the bandwidth or sessions per application as well as jitter, loss, and latency.

NOTE: The Application Path Insights visualization data is available only if the configuration is managed by Juniper Mist.

Video: NOW in 60: WAN Assurance Application Insights Dashboard

Figure 116: Application Path Insights (BETA)

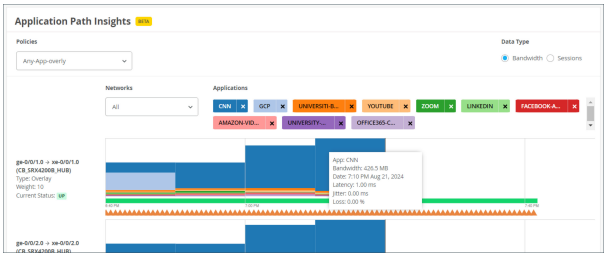

The path state bar shows path state information over a timeline, and path state events are indicated by segments highlighted in different colors. For example, Path Up events are shown in green and Path Down events are shown in red.

You can hover over the highlighted portions of the path state bar to view a summary of path state events.

The Application Path Insights section also includes a summary view on the lefthand side that displays recent path state events.

Figure 117: Application Path Insights (BETA) continued

If you click on the bar, you will get a pop-up window where you can view more detailed information about the path state events. The list of events displays on the left and when you select an event, the reason for the event displays on the right.

Path state events include:

- Path Update
- Port Up
- Port Down
- Path Up
- Path Down

Path state reasons include:

- Probe Down
- Peer Path Up
- Peer Path Down
- Config Change
- Best Path Selected
- SLA Metric Violation

Figure 118: Path State Events and Reasons

The **WAN Edge Device Charts** include Control Plane CPU, Data Plane CPU, Memory Utilization, and Power Draw.

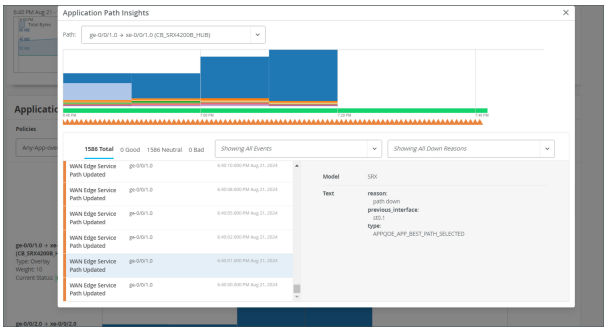

The **Control Plane CPU** and **Data Plane CPU** charts show you the percentage of CPU utilization for both max and average.

Figure 119: Control Plane CPU and Data Plane CPU

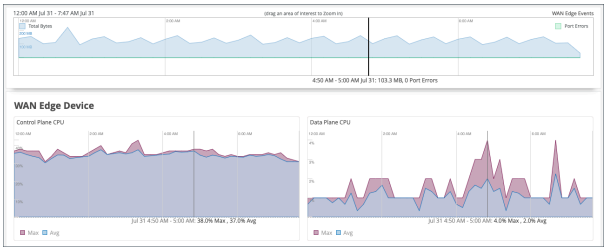

Memory Utilization and **Power Draw** shows you the percentage for both max and average.

Figure 120: Memory Utilization and Power Draw

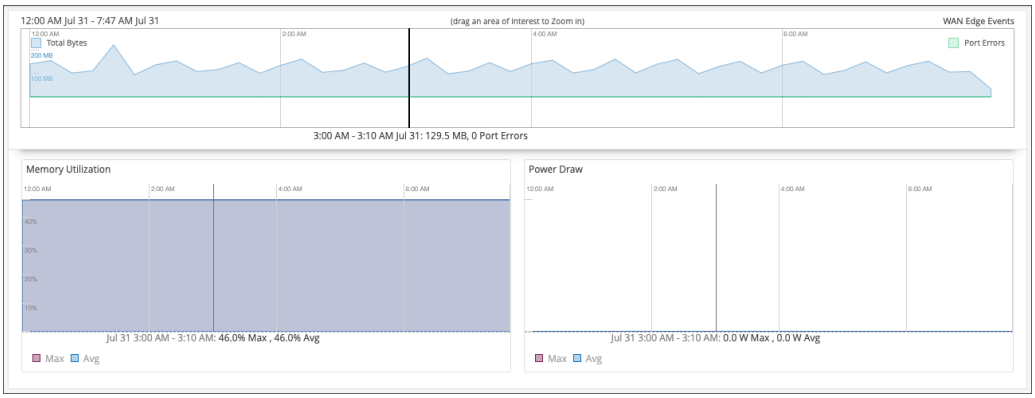

The **WAN Edge Ports** charts include Bandwidth, Max Bandwidth, Applications TX + RX Bytes, Port Errors, and IPsec Traffic. From the drop down list at the top, you can select All ports to see utilization metrics in the charts for all interfaces, or you can select an interface to see the utilization metrics for that particular interface.

Figure 121: Bandwidth and Max Bandwidth

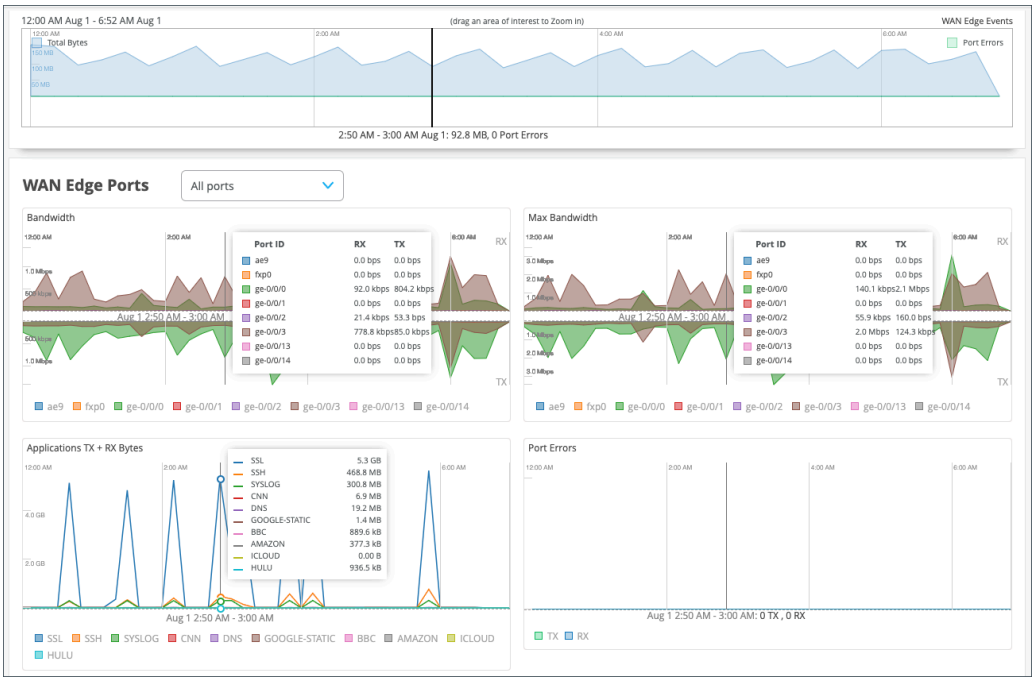

- In the **Bandwidth** chart, you will see the bandwidth utilization metrics in megabits per second (Mbps) for that particular interface.
- The **Max Bandwidth** chart displays insights into the highest point of link utilization recorded for received power signal (RX) and transmitted power signal (TX) packets on each port during the day. The data is shown in Mbps.

In the last three WAN Edge Ports charts, you'll find **Applications TX + RX Bytes**, **Port Errors**, and **IPsec Traffic**. Hover over the charts to find out more information.

- The **Applications TX + RX Bytes** chart outlines transmit and receive data information, which can be isolated at an application level by clicking on the application name at the bottom of the chart to see Client, MAC address, IP address, device type, and bytes for bandwidth utilization.
- The **IPsec Traffic** chart displays IPsec traffic for transmit and receive packets during the day in kilobytes or megabytes.
- The **Port Errors** graph displays port errors detected on the WAN Edge device over a period of time. Port errors are ethernet data link error counts that include all possible ethernet errors reported by the port device driver. Exact types of errors vary by device driver, and the total may include but is not

limited to CRC errors, collisions, etc. Errors are counted in both the transmit (TX) and receive (RX) direction. The graph displays the total for all ports, or for a particular port based on the WAN Edge Ports selection.

Peer Path Statistics

This applies only to Session Smart Routers deployed as WAN edge devices in Juniper Mist™ WAN Assurance. Therefore, no data will be populated in this section for SRX Series Firewalls deployed as a WAN edge device.

The final section of your WAN Edge Insights page is Current WAN Edge Properties. Time range selections do not impact information in the Current Values pane.

Figure 122: Current WAN Edge Properties

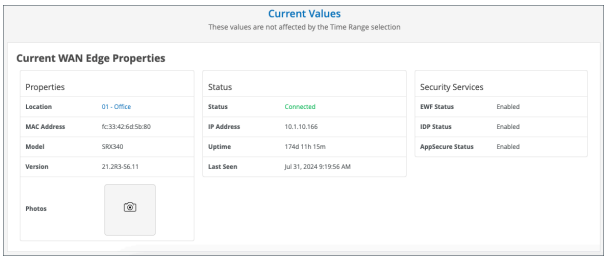

View Alerts for Interfaces Status

In Juniper Mist, alerts present network and device issues that are ongoing. You can view alerts on the Juniper Mist portal by selecting Monitor > Alerts.

You can set up alerts and email updates for when certain ports on a WAN Edge device go online or offline. To configure alerts for specific ports, you need to label these ports in the LAN or WAN settings of a WAN Edge device.

To configure the alerts and notifications for specific port, you must:

- Change the WAN or LAN settings to label the specified ports in the WAN Edge template or at device-level configuration page.

  1. In the Juniper Mist portal, select **Organization > WAN > WAN Edge Templates** and select the WAN or LAN configuration that you want to update (or add a new configuration).

     To configure this at the device-level, select **WAN Edges** on the left-navigation bar and select the WAN or LAN configuration of the selected device.

2. Under Interface, enter the port or ports, and then select the **Enable "Up/Down Port" Alert Type** check-box.

Figure 123: Marking LAN Port or WAN Interface as Critical Interface

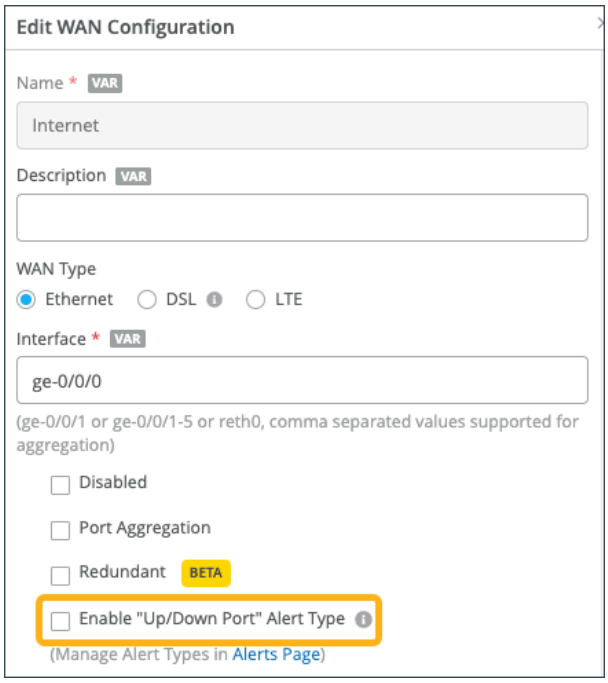

Repeat these steps for all critical ports.

- Configure alerts and e-mail notifications for the specified ports on the Alerts page.
  1. Go to **Monitor > Alerts > Alerts Configuration** and use the following check-boxes to enable alerts for the selected port:
     - Critical WAN Edge Port Up
     - Critical WAN Edge Port Down

Figure 124: Alerts Configuration for Critical Ports

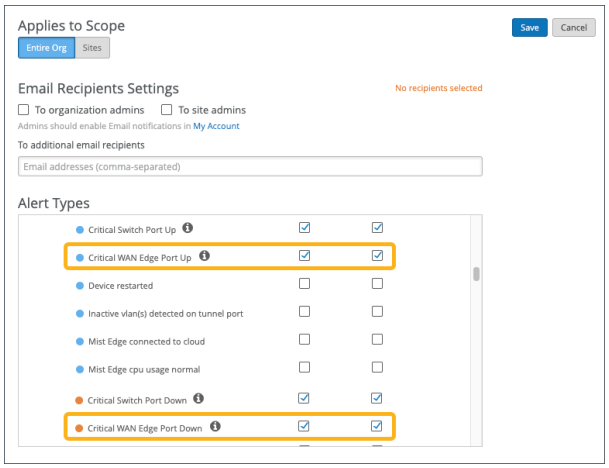

See Alert Configuration for details.

When you enable alerts and notifications:

- You'll receive an e-mail notification whenever a port transitions from one state to another.
- You can delay alerts about when the WAN edge gateway goes offline to prevent repeated alerts in the case of connectivity flaps by clicking the pencil icon and setting a time threshold.
- You can view the status in Monitor > Alerts page. Figure 125 on page 340 shows an example of the critical port status on the Juniper Mist Alerts dashboard.

Figure 125: Critical WAN Edge Port Status

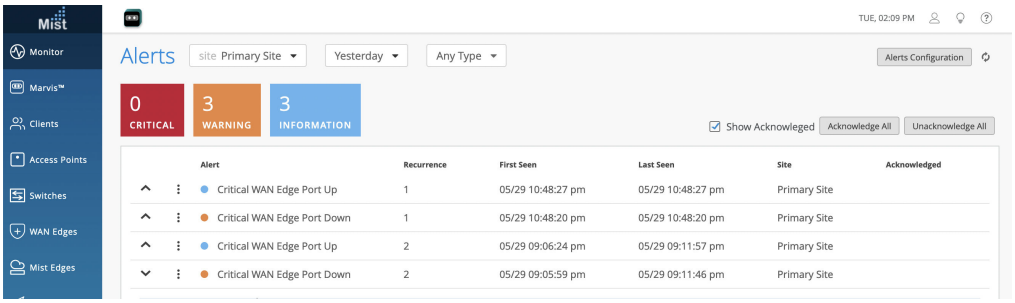

Monitor Session Smart Router Deployed as WAN Edge

SUMMARY

Use the WAN Edges page, the Insights page, and the Alerts page to quickly find device information, event details, peer path statistics, and alerts for your Session Smart Routers.

IN THIS SECTION

In monitoring Juniper® Session Smart™ Routers deployed as WAN edge device, you'll explore the most efficient ways to monitor your WAN edge device in the Juniper Mist™ portal following your initial deployment phase.

Monitor WAN Edges

In the Juniper Mist Portal, select **WAN Edges > WAN Edges** to view basic device monitoring. Notice the Organization name at the top of the GUI, AI-DRIVEN SDWAN AND SASE FULL STACK. This is the largest container and represents your entire organization. Beneath the organization name, you can see your site devices in either a List format or a graphical Topology format. Here, you see **Dallas-FullStack** is your site, and **lab1-dallas** is your WAN edge device.

Figure 126: Accessing WAN Edges Page

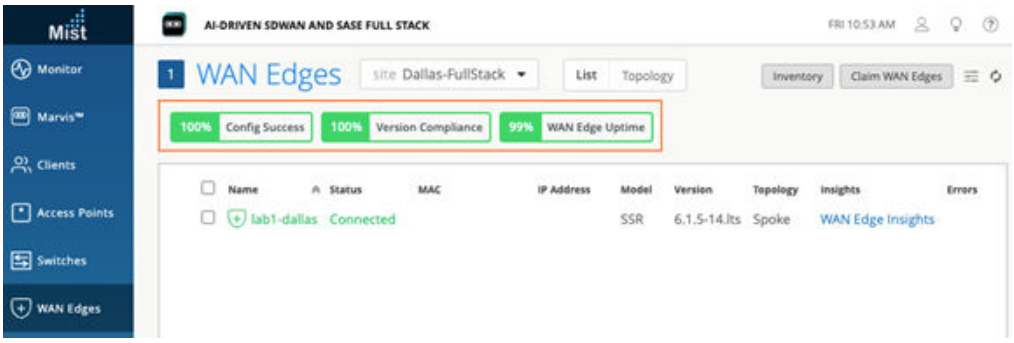

The List view outlines the following information:

- Config Success—Percentage of online WAN edges with successful configuration
- Version Compliance—Percentage of WAN edges that have the same software version per model.
- WAN Edge Uptime—Percentage of time a WAN edge was up during the past seven days, averaged across all EAN edges.

Figure 127: WAN Edges List View

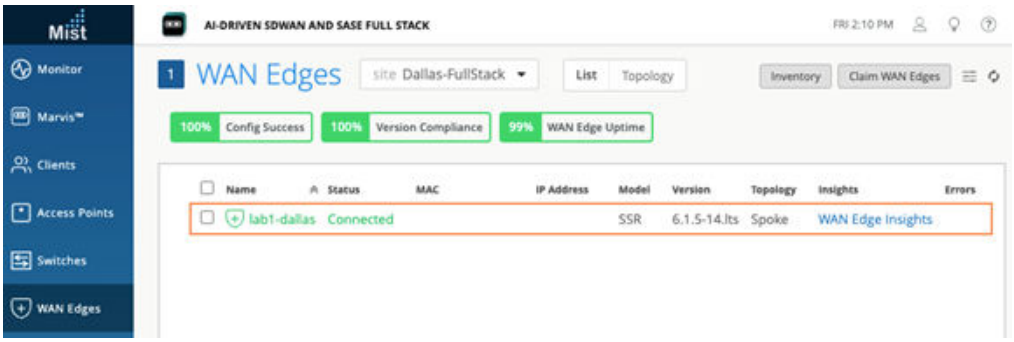

Beneath, you'll find WAN edge device details as shown in Table 74 on page 342.

Table 74: WAN Edge Device Details

Fields	Description
Name	Name
Status	Connected or disconnected
MAC	MAC address

IP Address	IP address
Model	Juniper® Session Smart™ Routers or Juniper Networks® SRX Series Firewalls
Version	SSR Software Version
Topology	Hub or Spoke
Errors	Error state

The **Topology** format presents the same information as the **List** view. For example, if you hover over the node0.lab1-dallas device, you'll see the same information as that displayed in the List view as you'll see in the figure below.

Figure 128: WAN Edges Topology View

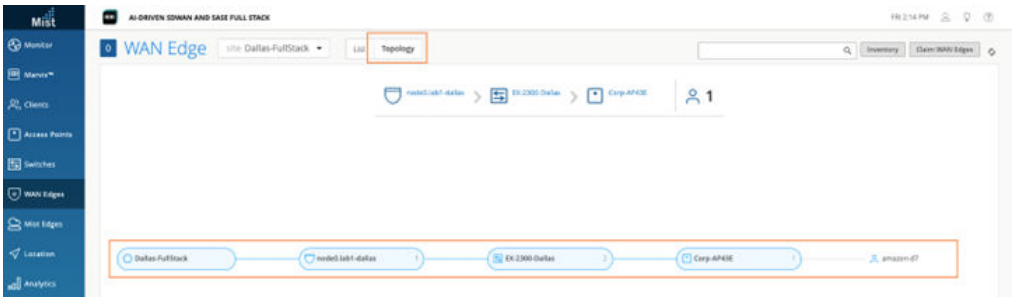

On both the List and Topology view, selecting your WAN edge device (lab1-dallas in this example) brings you to its Device Information page. The Device Information page provides different categories of monitoring information for your WAN edge device.

Figure 129: WAN Edges Device Information Page

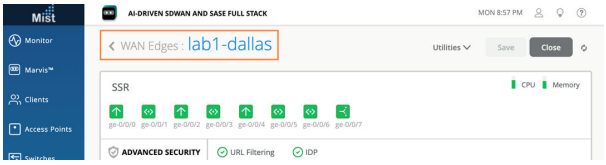

The first thing you'll notice on the Device Information Page is details about the WAN edge device you selected, (lab1-dallas in our figure). The information includes a graphical front view of the device ports and baseline status information such as CPU and memory utilization.

Figure 130: WAN Edges Device Information - Interfaces

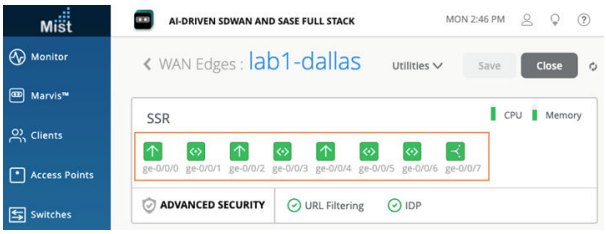

For each Gigabit Ethernet interface you'll find link information.

Figure 131: WAN Edge Device Information Page- Details

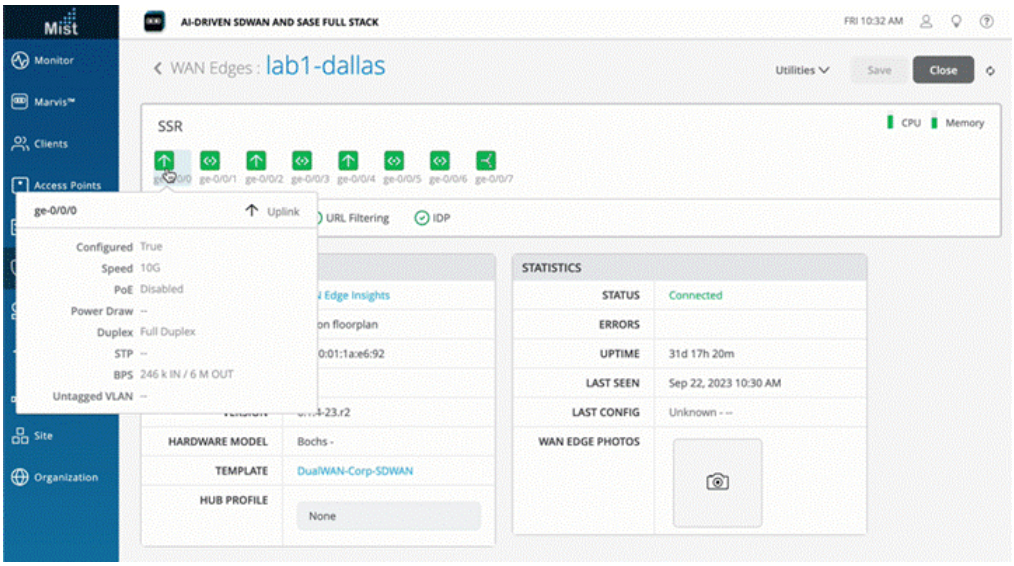

Table 75: Link Information for Gigabit Ethernet Interface

Fields	Description
Configured	True or false
Speed	Rated speed
PoE	Enabled or disbaled
Power Draw	Measured PoE power draw
Duplex	Full or half

STP	True or false
BPS	Bits/second
Untagged VLAN	-

When hovering over Wired Clients, you'll get similar information with additional information.

Table 76: Wired Clients Details

Fields	Description
Hostname	Name of the device
Username	User name
MAC	MAC address of the device
IP Address	IP address of the device
Manufacturer	Type of device- SSR / SRX

The CPU and Memory status icon indicates how your device behaves. Hover over each interface icon for deeper insights.

Figure 132: WAN Edges - CPU and Memory Status

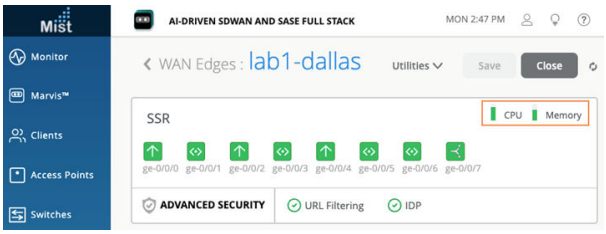

Advanced Security information is listed below the device ports with a checkmark or an X, indicating whether URL filtering or intrusion detection and prevention (**IDP**) is active on this device. Here, both URL filtering and IDP are active with the green checkmark.

Figure 133: Advanced Security Details

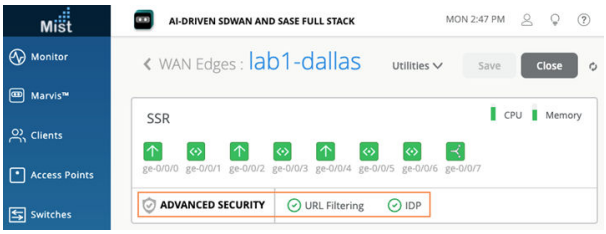

Below our port information and security section, you'll find generalized data for your WAN edge device, including:

Figure 134: WAN Edge Device Properties

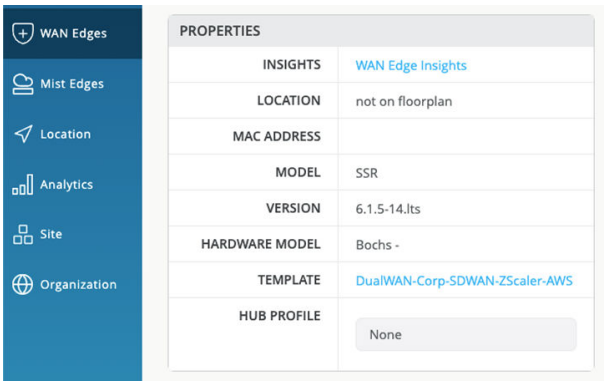

Properties contains generalized platform-related information.

Table 77: WAN Edge Platform-Related Details

Field	Description
Insights	Provides a direct link to WAN Edge Insights.
Location	Provides floorplan information
MAC Address	MAC Address for the SSR device
Model	Indicates if model type is SSR or SRX
Version	Version of the Session Smart Software
Hardware Model	Lists the Whitebox or Juniper Networks device model name and number.

Template	The applied WAN edge template to the device.
Hub Profile	The applied Hub Profile to the device.

Statistics displays action information about your platform.

Figure 135: WAN Edge Device Statistics

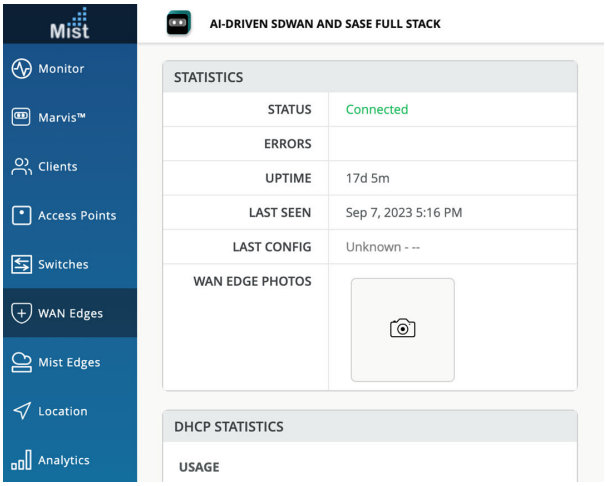

Table 78: WAN Edge Device Statistics

Field	Description
Status	Connected /Disconnected
Errors	Any commit errors
Uptime	Day/Hour/Min uptime information
Last Seen	Last login
Last Config	Last Commit
WAN Edge Photos	Photos of the WAN edge device

If you configured DHCP servers on the WAN router itself, there will also be a DHCP Statistics pane with information about the leased IPs.

- **DHCP Statistics** presents IP information related to dynamic distributed IP addresses.

Figure 136: WAN Edge Device DHCP Statistics

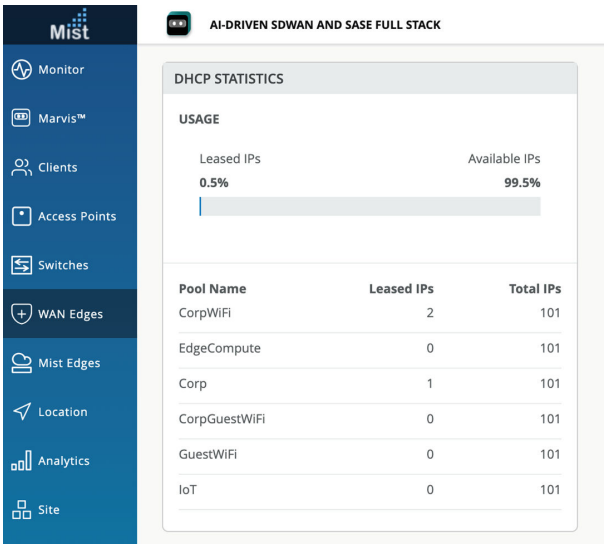

Table 79: WAN Edge Device DHCP Statistics

Field	Description
Usage	The total figure presented as a percentage of Leased and Available IPs
Pool Name	The name for given pool of addresses
Leased IPs	Number of used IP addresses in each pool.
Total IPs	Total number available of IP addresses in each pool.

As you scroll down the device information page, you'll find Secure Vector Routing (SVR)-based Paths between devices that provide information about connectivity through WAN interfaces to the hubs. Here, you can review your WAN edge device configuration. Usually, WAN edges inherit templates or profiles. However, you can make individual changes to the configuration to be pushed to the device.

Topology Details displays Peer Path information. Remember that a Session Smart SD-WAN network overlay is generated through Secure Vector Routing Peer connections between Session Smart devices.

Figure 137: Topology Details

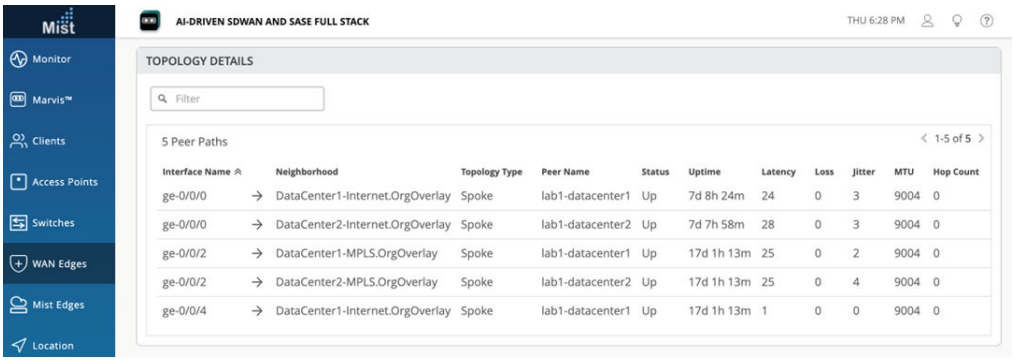

Table 80: Topology Details

Field	Description
Interface Name	Lists the name of the interface
Neighborhood	The shared layer 3 connection between Peers
Topology Type	Indicates Hub/Spoke
Status	Indicates up/down
Peer Name	Peer SVR device
Uptime	Time up and live
Latency	Measured in Milliseconds
Loss	Packet loss
Jitter	Measured in Milliseconds
MTU	Max Transmission Unit
Hop Count	Number of Hops

Secure Edge Connector Details include tunnel information from your WAN edge connection to the Secure Edge cloud.

Figure 138: Secure Edge Connector Details

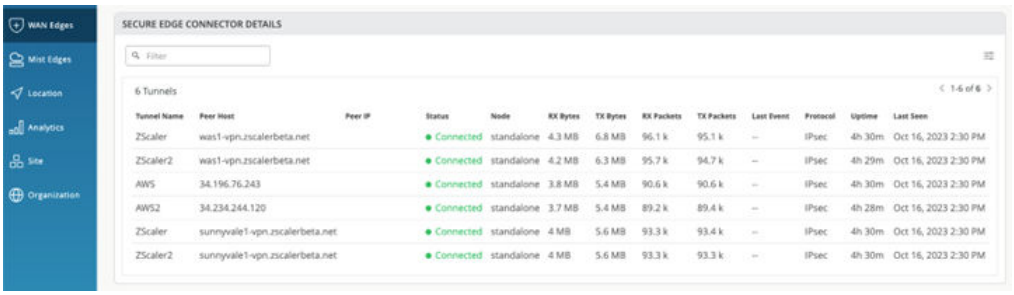

Table 81: Secure Edge Connector Details

Fields	Description
Tunnel Name	Name
Peer Host	Peer Host IP Address
Peer IP	Peer IP
Status	Connected/Disconnected
Node	Standalone/HA
RX Bytes	Volume of data, in bytes, received by the interface.
TX Bytes	Volume of data, in bytes, transmitted by the interface.
RX Packets	Packets received by the interface.
TX Packets	Packets transmitted by the interface.
Last Event	System events
Protocol	Protocol
Uptime	time live
Last Seen	Last login

Scrolling down the device information page, you'll find configuration information for your WAN edge. First, it'll indicate hub or spoke with relevant information about your **WAN Edge Configuration**.

Figure 139: WAN Edge Configuration: Spoke

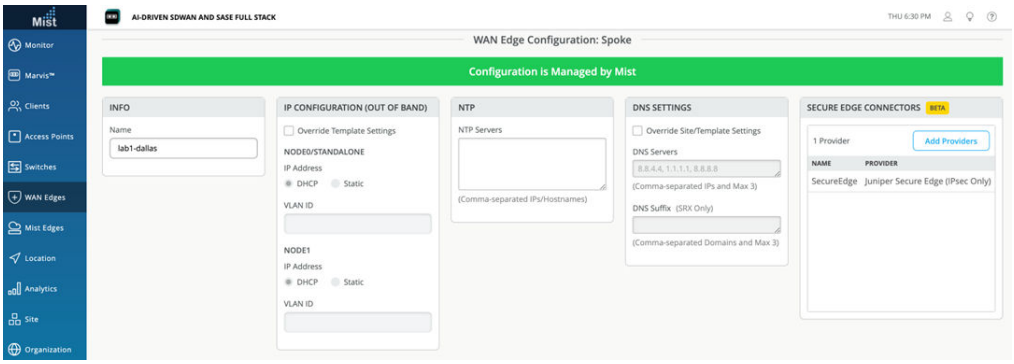

Table 82: WAN Edge Configuration: Spoke

Field	Description
Info	Name
IP Configuration	Override Template Settings, node1 DHCP/Static, VLAN ID, node 2 DHCP/Static, VLAN ID
NTP	Time Servers IP/Hostnames
DNS	Override Template Settings, DNS Servers, (SRX only DNS suffix info)
Secure Edge Connector	Provider for the Secure Edge Connector.

Scrolling past the configuration, you'll find information for your connected WANs and LANs.

Figure 140: WAN Details

Table 83: WAN Details

Field	Description
Name	Selected WAN Interface Name

Interface	Supports one of these interfaces for aggregation: ge-0/0/1, ge-0/0/1-5, or reth0.
WAN Type	Ethernet, DSL (SRX Only) LTE
IP Configuration	DHCP, Static, or PPPoE
Overlay Hub Endpoints	SVR Peer connections to the Hub

Figure 141: LAN Details

Table 84: LAN Details

Field	Description
Network	Selected LAN name.
Interface	Supports one of these interfaces for aggregation: ge-0/0/1, ge-0/0/1-5, or reth0.
Untagged	Untagged VLAN (SRX only)
VLAN ID	DHCP, Static, or PPPoE
IP Configuration	SVR Peer connections to the Hub
DHCP	Relay, Server, none.

The Traffic Steering and Application Policy sections show how you use the Session Smart Secure Vector Routing process to create rules for path choice and routing behavior. Note that on the SRX Series deployed as a WAN edge, the Application Policy and Traffic Steering path determine destination zones and must be assigned. The Session Smart router is first and foremost, a router and will use the closest match for the address.

Figure 142: Traffic Steering

Table 85: Traffic Steering

Field	Description
Name	Selected Traffic Steering name.
Strategy	Ordered, weighted, ECMP
Paths	Untagged VLAN (SRX only)

Application Policies are the heart of Juniper's AI-Driven SD-WAN. Remember that Application Policies are security policies in Juniper WAN Assurance design, where you define which network and users can access which applications, and according to which traffic steering policy. You must create Networks, Applications, and establish Traffic Steering profiles to define an Application Policy. These elements become matching criteria to allow access to or block access from applications or destinations.

Figure 143: Application Policies

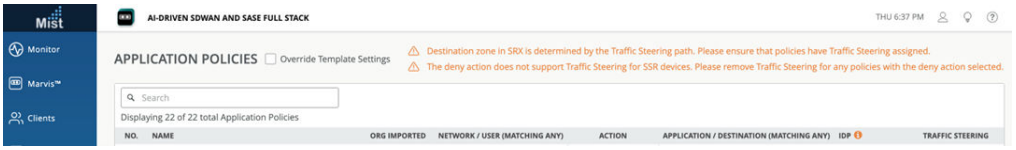

In the Juniper Mist™ cloud portal, the Networks or Users setting determines the source zone. The Applications and Traffic Steering settings determine the destination zone. Traffic Steering paths determine the destination zone in Juniper Networks® SRX Series Firewalls, so ensure that you assign Traffic Steering profiles to the Application Policies.

Table 86: Application Policies Details

Field	Description
Number	Ordered Policy Number
Name	Selected name

Org Imported	Indicates if the policy was pushed down from the Organization level to the Site.
Network/User (Matching Any)	The “source” of your traffic
Action	Allow/Block
Application/Destination (Matching Any)	The “destination” for your traffic.
IDP	Indicates IDP/URL filtering (requires separate license)
Traffic Steering	Indicate path for traffic

The bottom of the Device Information page has tables for routing properties such as BGP and static routes connected to your WAN edge device. You can also manually add a BGP Group here.

Figure 144: Routing Details

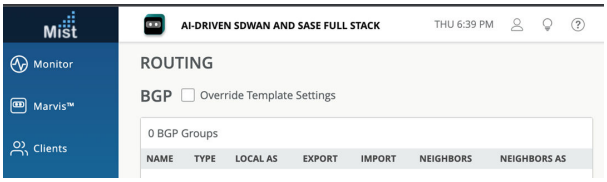

Table 87: Routing Details

Field	Description
Name	BGP Name
Type	Type of BGP Route
Local AS	Autonomous System Number
Export	Exported Route
Import	Imported Route
Neighbors	Neighbor Route

Neighbor AS	Autonomous System Number for Neighbor Route

Figure 145: Static Routes

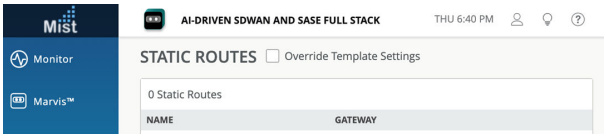

Static Routes display name and gateway information.

WAN Edge Insights

The Properties pane for your selected WAN edge links to **WAN Edge Insights**. Click **WAN Edge Insights** for the next level of information about your WAN edge device.

Figure 146: WAN Edge Insights

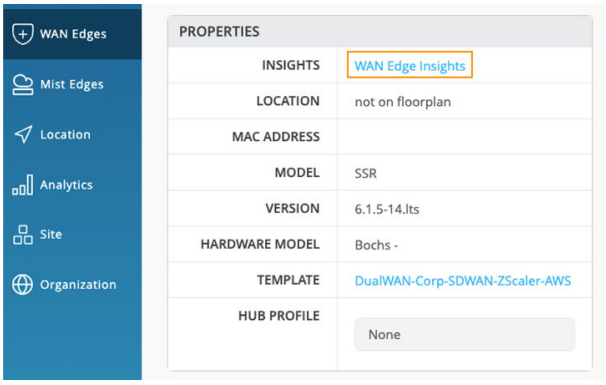

Next to the selected WAN edge (lab1-dallas) on the Insights page, you can select a timeframe for selected information. The default view is **Today,** but this can be set to a customized date or range of dates. Below this, you find (when the site location information is configured) where this WAN edge is configured via a street map.

Figure 147: WAN Edge Insights-Select Time Duration

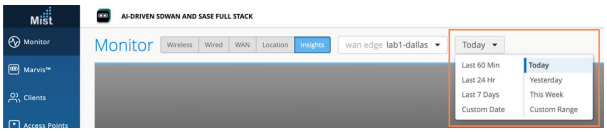

NOTE: Ensure you've had a few hours for these metrics to be populated following initial deployment.

WAN Edge Insights includes the following sections:

- WAN Edge Events
- Table Capacity
- Applications
- Applications Path Insight (Beta)
- WAN Edge Device
- WAN Edge Ports
- Peer Path Stats

At the end, you can see **Current WAN Edge Properties** for your device location, MAC address, model, version. You can see the status, uptime, and last seen, and security services status.

WAN Edge Events

With your timeframe selected, **WAN Edge Events** displays a timeline of the traffic through the WAN edge during your specified time, and a list of events in the same window.

Select a specific event in the listed WAN Edge Events for greater detail of the **Good**, **Neutral**, and **Bad** events.

Figure 148: WAN Edge Events Timeline

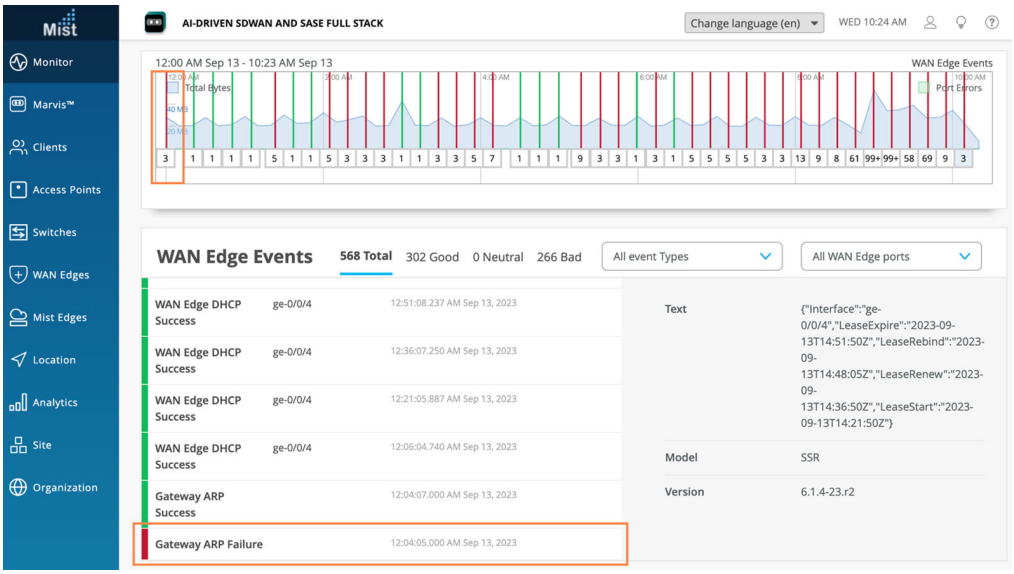

Your selection expands and displays detailed information about the selected time.

For a detailed portion of time, select a window of time with the mouse cursor. By doing this, you're able to adjust the window of events and isolate specific **Good**, **Neutral**, and **Bad** things that happened on your network. With a smaller section you'll get a more detailed view of that period.

Figure 149: WAN Edge Events Timeline Details View

Drill down the WAN Edge Events page for deeper insights within your selected period.

Figure 150: WAN Edge Events page

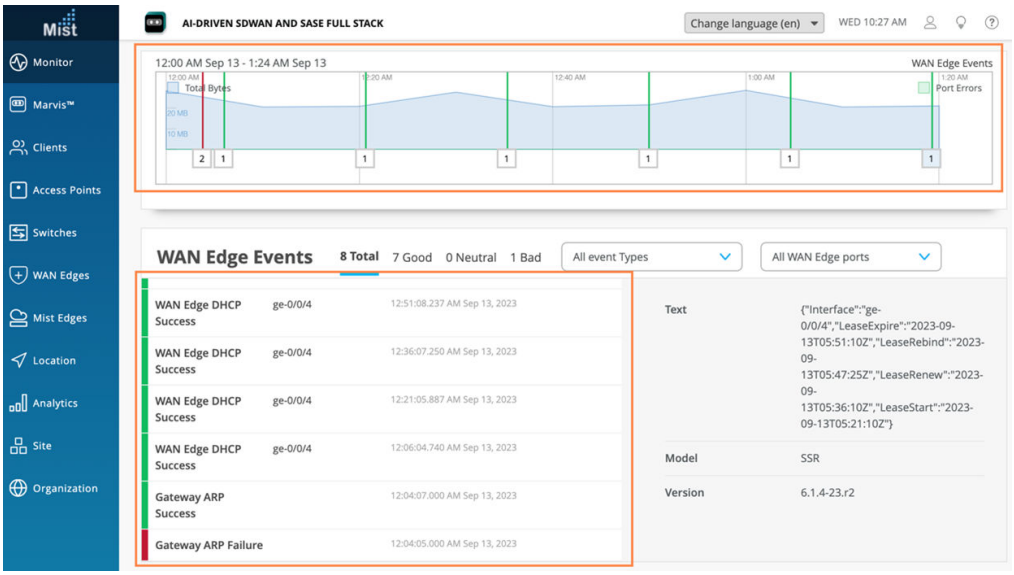

We can continue that way: You can narrow down on the type of event by selecting a modifier in the Event Type drop-down menu. You can also filter your search by limiting the event types to a specific port

Figure 151: WAN Edge Events Page

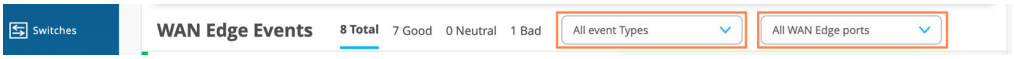

Table Capacity

The WAN Edge Insights page provides the following indicators in the Table Capacity section:

Figure 152: Table Capacity

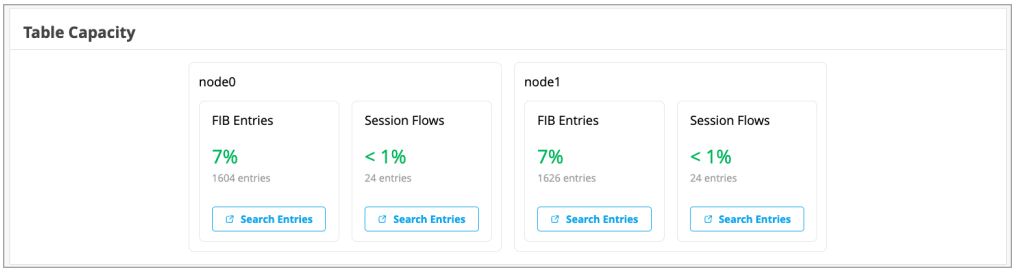

- FIB Entries: Displays the current number of FIB entries and the percentage of utilization; essentially showing how much of the available FIB space is currently being used.
- Session Flows: Displays the current number of active sessions and the percentage of session flow utilization based on the device's capacity.

In the case of a high availability cluster, Table Capacity indicators are displayed for each node.

You can also click the **Search Entries** button under each metric to open a shell view in a new window where you can search for entries after specifying filters.

Applications

The Applications section on WAN Edge Insights page allows you to:

- Ue categorized applications to monitor and troubleshoot specific application behavior.
- View a client's use of a particular application by clicking the Clients tab.

The chart provides the following details about applications usage:

- Service—Category of the applications. Expand the categories to view individual applications details.
- Total Bytes — Traffic volume that each application or website receives and transmits.
- Percent Bytes — Percentages of traffic volume that each application or website receives and transmits.
- RX Bytes —Traffic that applications or websites receive.
- TX Bytes—Traffic that applications or websites transmit.

Figure 153: Applications

The Application and Client statistics shown on the WAN Edge insights page come from the WAN Edge device. The tabs show the same data, only it is indexed according the different sort fields. The *Number of Applications* field on the client's tab is the applications reported through the SSR/SRX.

The statistics shown on the Client and Sites Insights pages, as well as that available from the Clients tab on the WAN Edge Insights page, provide data from the AP about the wireless clients. Individual client names are linked to their respective Client's Insights page, which likewise shows data for that particular client from the AP.

NOTE: For Session Smart Router devices running a DHCP server, clients using that application will display a HostName in the Client column if available. Otherwise, the MAC address will be displayed. Device Type and MAC Address columns will be populated as well.

Figure 154: Clients Using Application

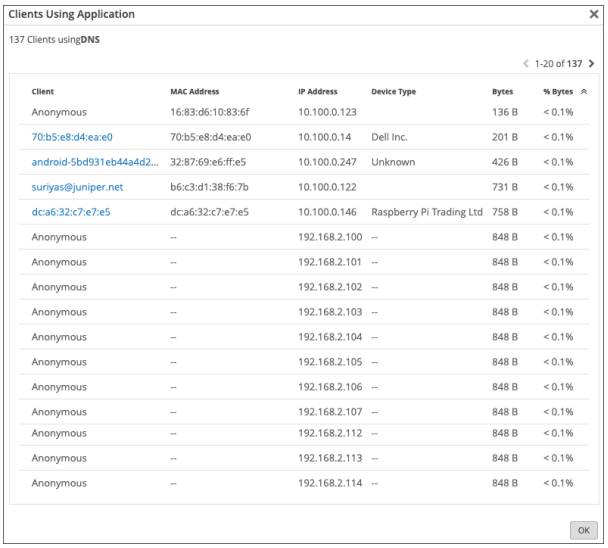

Application Path Insights

The Application Path Insights (BETA) section shows you which applications are using the most bandwidth according to the selected Application Policy and Network. You can also change the Data Type to Sessions to see the number of sessions occurring per application. Hover over a section of the graph to view the bandwidth or sessions per application as well as jitter, loss, and latency.

This tab helps you to visualize how different applications are utilizing WAN links and make informed decisions about policy tuning and path optimization.

See "Application Settings" on page 98 and "Application Policies" on page 144 for details on applications and application policies.

Video: NOW in 60: WAN Assurance Application Insights Dashboard

Figure 155: Application Path Insights (BETA)

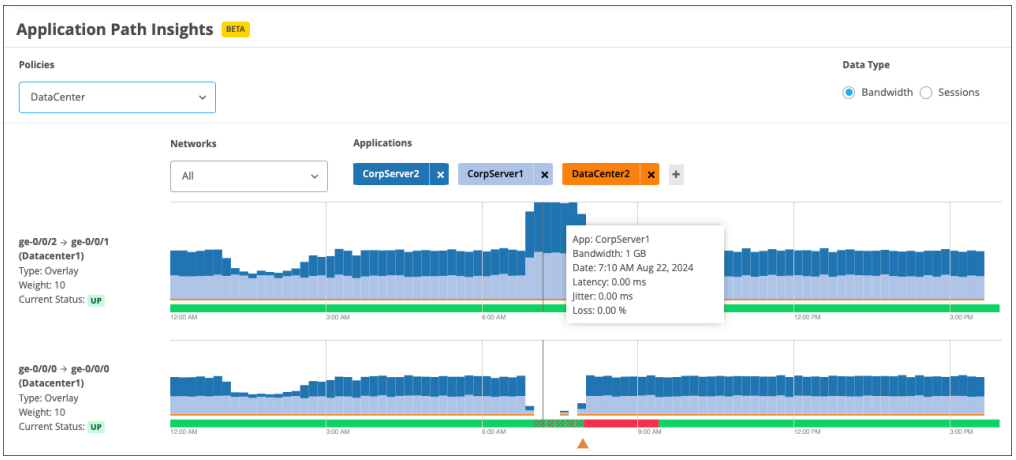

This tab provides the following configuration options:

- **Policies**: You can filter by policy type.
- **Data Type**: Toggle between **Bandwidth** and **Sessions** to analyze traffic volume or session count.
- **Networks**: Filter by specific networks.

A list of applications is shown with color-coded tags.

Figure 156: Application Path Insights with Service Path Update Event

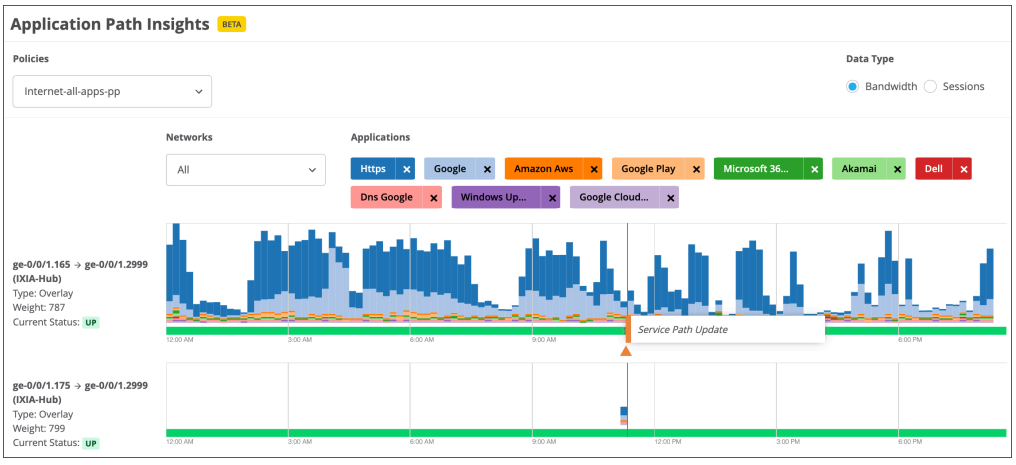

The path state bar shows path state information over a timeline, and path state events are indicated by segments highlighted in different colors. For example, Path Up events are shown in green and Path Down events are shown in red.

If you see an orange triangle below the path state bar, this indicates that a Service Path Update event occurred. You can hover over the triangle to see the details.

The Application Path Insights section also includes a summary view on the lefthand side that displays recent path state events.

You can also hover over the highlighted portions of the path state bar to view a summary of the path state event.

Figure 157: Application Path Insights (BETA) continued

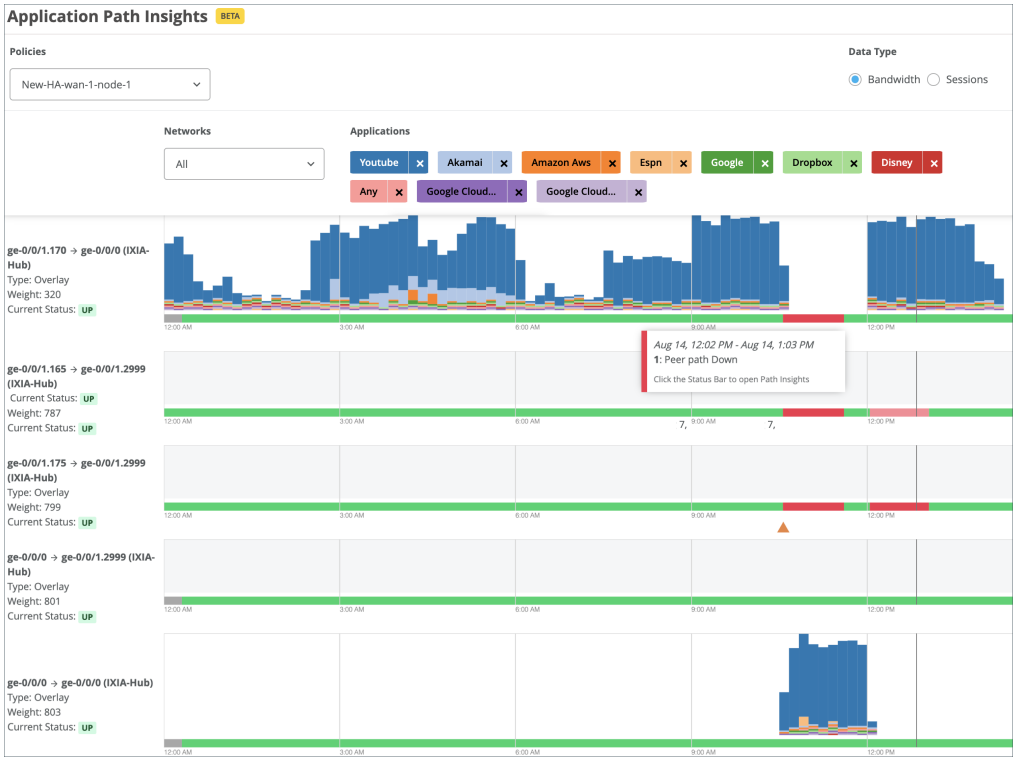

If you click on the bar, you will get a pop-up window where you can view more detailed information about the path state events. The list of events displays on the left, and when you select an event, the reason for the event displays on the right.

Path state events include:

- Path Add
- Path Remove
- Path Update
- Port Down

- Path Up
- Path Down

  Path Down Reasons include:

  - Probe Down
  - Peer Path Down
  - ARP Unresolved
  - DHCP Failure

Figure 158: Path State Events and Reasons - Example 1

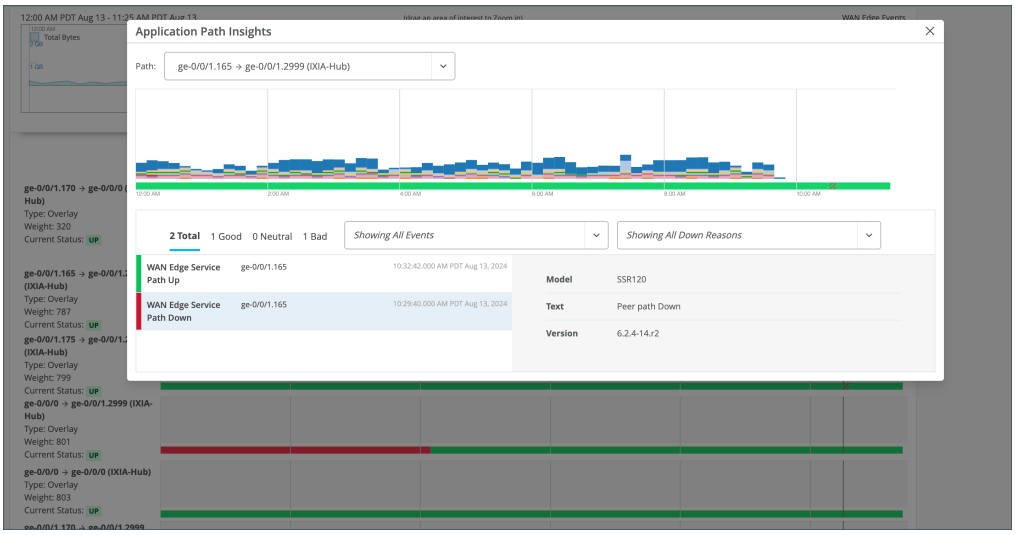

Figure 159: Path State Events and Reasons - Example 2

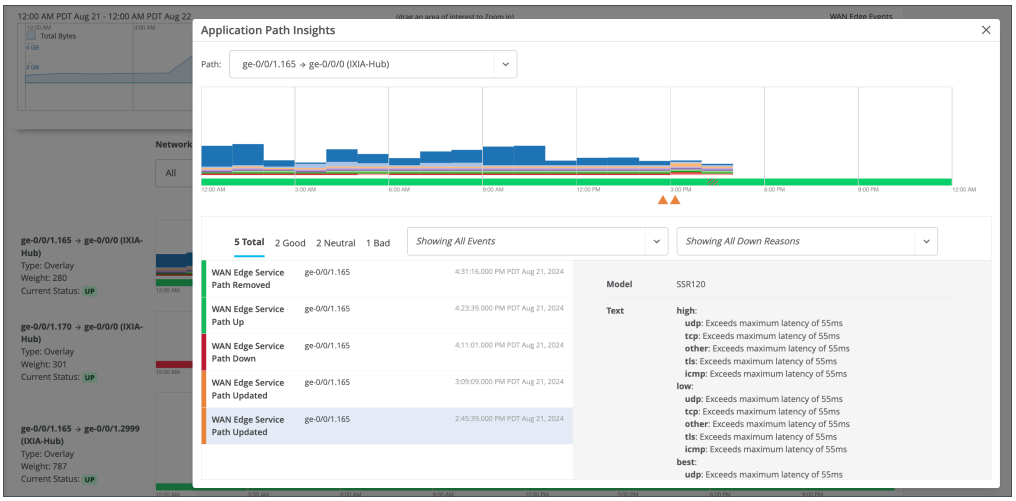

WAN Edge Device Chart

WAN Edge Device charts include Control Plane CPU, Data Plane CPU, and Memory Utilization.

Control Plane CPU shows CPU utilization for both max and average. The Data Plane CPU chart displays the CPU utilization for both max and average.

Figure 160: Control Plane CPU and Data Plane CPU

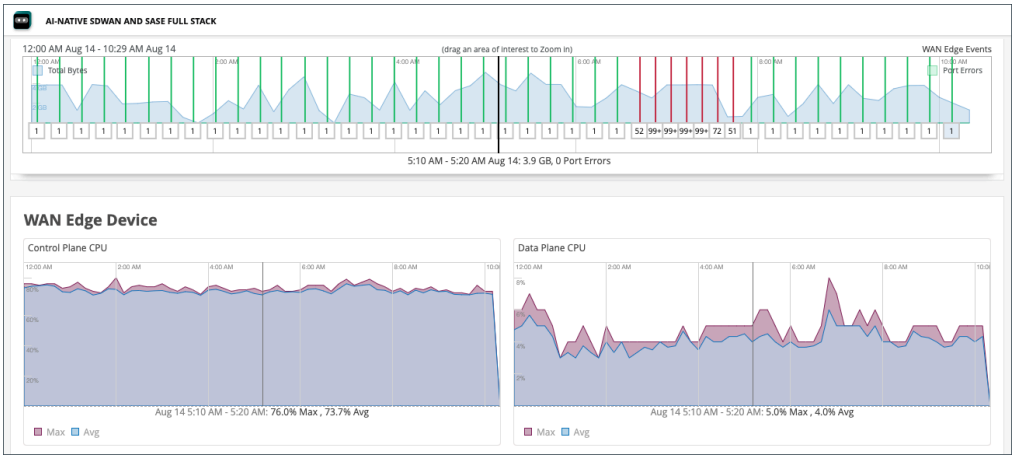

Memory Utilization displays the max and average memory utilization.

Figure 161: Memory Utilization

WAN Edge Ports Charts

The **WAN Edge Ports** charts include Bandwidth, Max Bandwidth, Applications TX + RX Bytes, and Port Errors. From the drop down list at the top, you can select All ports to see utilization metrics in the charts for all interfaces, or you can select an interface to see the utilization metrics for that particular interface.

In the **Bandwidth** chart, you will see the bandwidth utilization metrics in megabits per second (Mbps) for that particular interface.

The **Max Bandwidth** chart displays insights into the highest point of link utilization recorded for received power signal (RX) and transmitted power signal (TX) packets on each port during the day. The data is shown in Mbps.

Figure 162: Bandwidth and Max Bandwidth

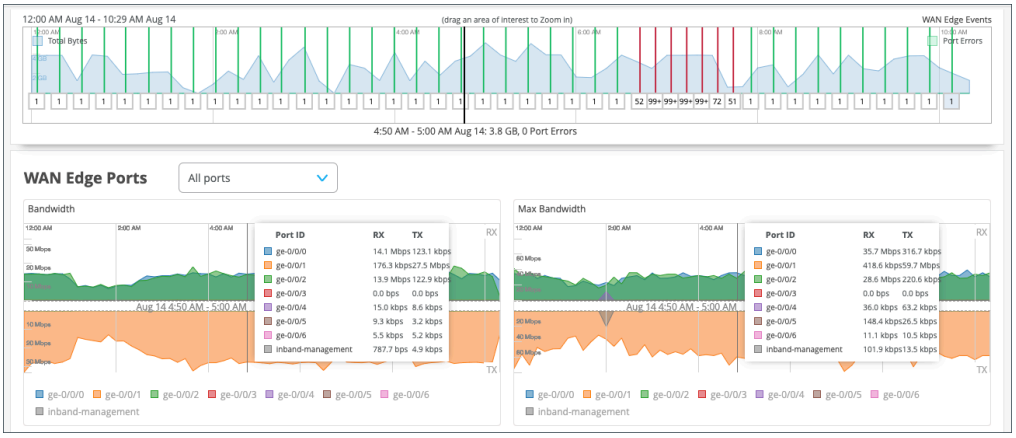

In the last two WAN Edge Ports charts, you'll find **Applications TX + RX Bytes** and **Port Errors**. Hover over the charts to find out more information.

Figure 163: Applications TX + RX Bytes and Port Errors

The **Applications TX + RX Bytes** chart outlines transmit and receive data information, which can be isolated at an application level by clicking on the application name at the bottom of the chart to see Client, MAC address, IP address, device type, and bytes for bandwidth utilization.

Figure 164: Port Errors

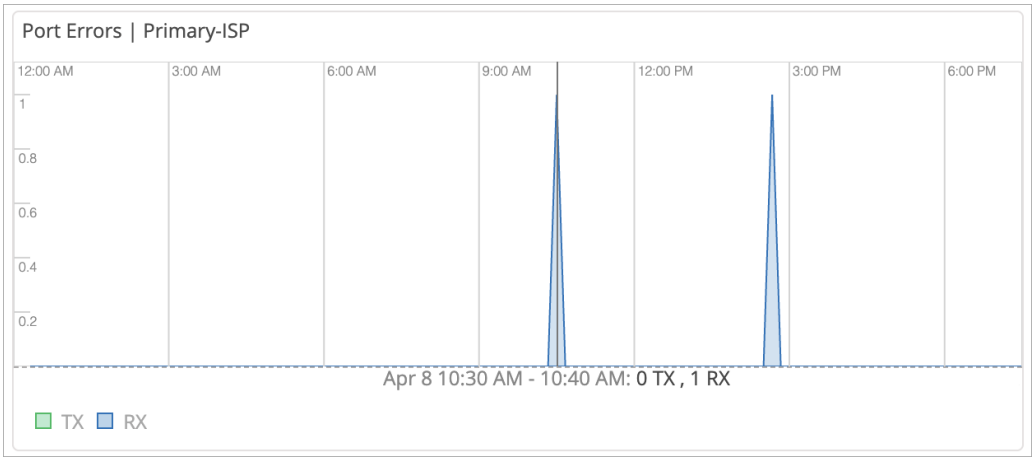

The **Port Errors** graph displays port errors detected on the WAN Edge device over a period of time. Port errors are ethernet data link error counts that include all possible ethernet errors reported by the port device driver. Exact types of errors vary by device driver, and the total may include but is not limited to

CRC errors, collisions, etc. Errors are counted in both the transmit (TX) and receive (RX) direction. The graph displays the total for all ports, or for a particular port based on the WAN Edge Ports selection.

Peer Path Statistics

The Session Smart WAN edge devices deployed in Juniper Mist™ WAN Assurance provide insights for liveness and path quality through Session Smart, Secure Vector Routing. The Session Smart use of the Bidirectional Forwarding Detection (BFD) signal on port 1280 checks with the downstream Session Smart Routers for liveness and monitors jitter, latency, loss, and mean opinion score (MOS). This insight works only with our Session Smart devices.

Figure 165: Peer Path Statistics

We return to WAN Edge Insights to find the Session Smart Peering metrics on your Mist dashboard. These graphs are at the bottom of the page, with a default view showing the worst three peer connections: jitter, latency, loss, and MOS. Drill down into the data here, using the same time ranges for the WAN Edge Charts. This also means that the graphs are interrelated and cross referenced.

You can also drill down and select a specific peer path to view statistics.

Figure 166: Peer Path Statistics for Specific Peer

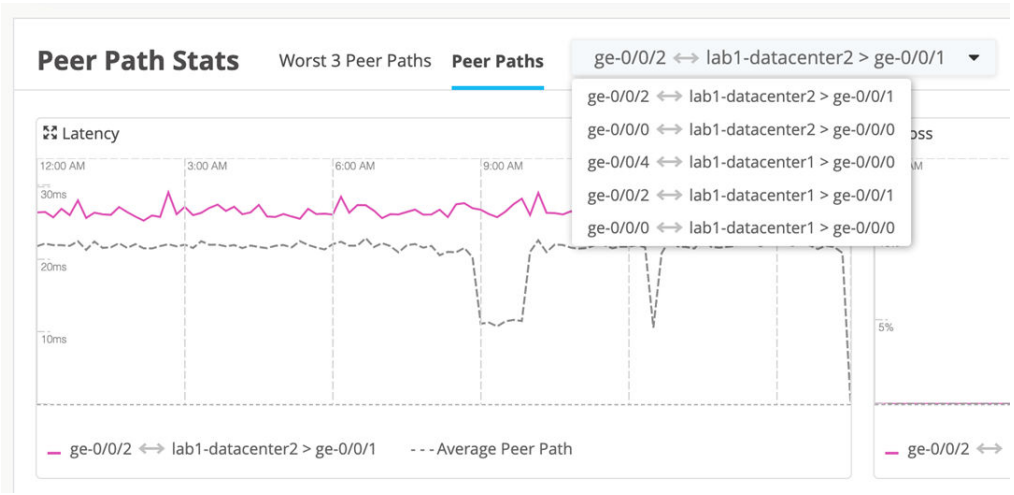

The final information on your WAN Edge Insights page is Current WAN Edge Properties. Time range selections do not impact information in the Current Values pane.

Figure 167: Current WAN Edge Properties

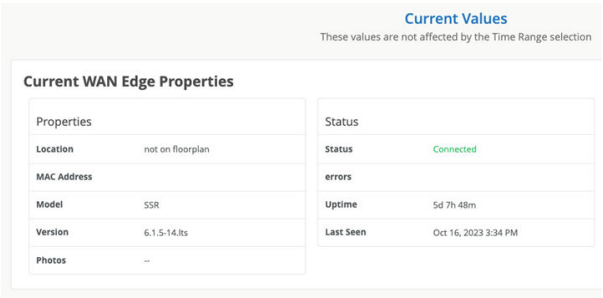

Alerts for Interfaces Status

In Juniper Mist, alerts present network and device issues that are ongoing. You can view alerts on Juniper Mist portal by selecting Monitor > Alerts.

You can set up alerts and email updates for when certain ports on a WAN Edge device go online or offline. To configure alerts for specific ports, you need to label these ports in LAN or WAN settings of WAN Edge device. You can also configure alerts to signal if the WAN edge gateway goes offline (these can be immediate or delayed, for example, if you want to prevent repeated alerts in the case of connectivity flaps).

To configure the alerts and notifications for specific port, you must:

- Change the WAN or LAN settings to label the specified ports in WAN Edge template or at device-level configuration page.

  1. In the Juniper Mist cloud portal, select **Organization > WAN > WAN Edge Templates** and select the WAN or LAN configuration that you want to update. (Or add a new configuration.)

     To configure at the device-level, select **WAN Edges** on the left-navigation bar and select WAN or LAN configuration of the selected device.

  2. Under Interface, enter the port or ports, and then select **Enable "Up/Down Port" Alert Type** check-box.

     Figure 168: Marking LAN Port or WAN Interface as Critical Interface

     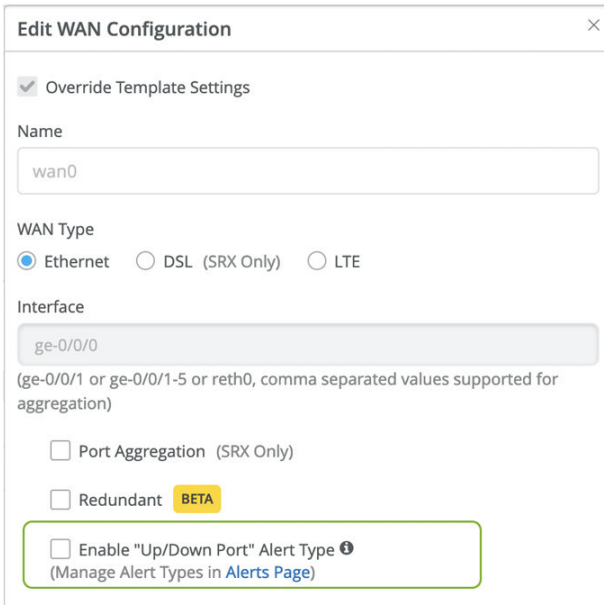

     Repeat these steps for all critical ports.

- Configure alerts and e-mail notifications for the specified ports in Alerts page.

  1. Go to **Monitor > Alerts > Alerts Configuration** and use the following check-boxes to enable alerts for the selected port:

     - Critical WAN Edge Port Up
     - Critical WAN Edge Port Down

Figure 169: Alerts Configuration for Critical Ports

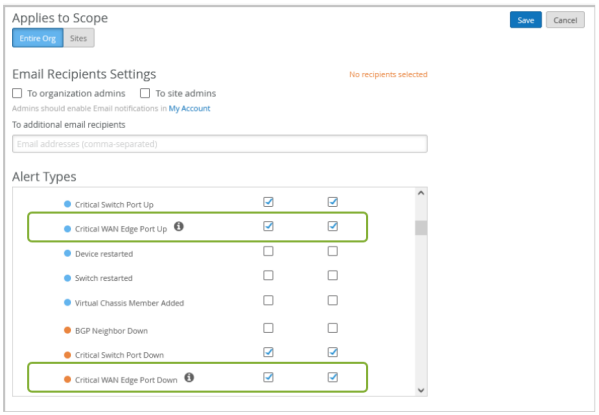

See Alert Configuration for details.

When you enable alerts and notifications:

- You'll receive an e-mail notification whenever a port transitions from one state to another.
- You can delay alerts about when the WAN edge gateway goes offline to prevent repeated alerts in the case of connectivity flaps by clicking the pencil icon and setting a time threshold.
- You can view the status in Monitor > Alerts page. Figure 170 on page 371 shows an example of the critical port status on Juniper Mist Alerts dashboard.

Figure 170: Critical WAN Edge Port Status

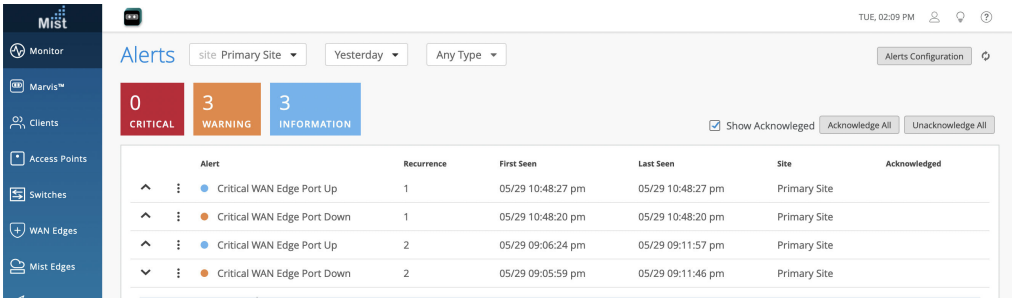

WAN SLEs

SUMMARY

Use the WAN Service-Level Experiences (SLEs) to assess user-impacting factors such as WAN Edge health, WAN link health, and application health.

IN THIS SECTION

Overview

IN THIS SECTION

Finding the WAN SLEs Dashboard

To find the WAN SLEs dashboard, select **Monitor** > **Service Levels** from the left menu, and then click the **WAN** button.

NOTE: The buttons appear only if you have the required subscriptions. See Requirements.

SLE Filter Buttons

- Use the buttons on the left to show **Success Rate** or **Values**.

- Use the **Show Custom Apps**button to show or hide your custom applications.

  In the example below, the button is in the Off position, so all applications are included. If you drag the button to the **On** position, you'll see only your custom applications.

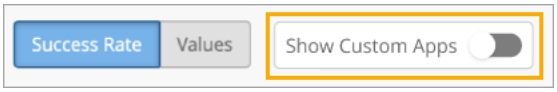

Video: WAN Assurance Overview

Video: WAN Assurance Video Overview

Video: Troubleshoot WAN Issues with SLEs

Video: SLE Example

WAN SLE Blocks

As shown in the following example, each SLE block provides valuable information.

- At the left, you see that this SLE has an 85 percent success rate. If you select the Value filter button, you'll see a number instead.
- At the center, the timeline shows variations across the time period. You can hover your mouse pointer over any point to see the exact time and SLE outcome.

  At the right, the classifiers show the percentage of the issues that were attributed to each root cause. In this example, 100 percent of the issues were attributed to Jitter.

- If you click a classifier, you'll see more information on the Root Cause Analysis page. Most classifiers have sub-classifiers for greater insight into the exact causes. The Root Cause Analysis page also provides additional details about the scope and impact of the issues.

See the following table for more information about the WAN SLEs and classifiers.

Table 88: WAN SLE Descriptions

SLEs	SLE Descriptions	Classifiers	Classifier Descriptions
WAN Edge Health	Juniper Mist monitors the user minutes when the health or performance of the WAN edge device is not optimal. Suboptimal health lowers the device's ability to pass traffic, directly affecting any clients connected to the device.	WAN Edge Disconnected	Lost connectivity to the Juniper Mist cloud
		System	High system usage relative to capacity Sub-Classifiers: • Memory—Memory utilization above 80 percent • Power—Power consumption above 90 percent • Temp CPU—CPU temperature outside the prescribed threshold range • Temp Chassis—Chassis temperature outside the prescribed threshold range • **CPU Data Plane**—CPU Data Plane utilization above 90 percent • CPU Control Plane—CPU control plane utilization above 90 percent

Table 88: WAN SLE Descriptions *(Continued)*

SLEs	SLE Descriptions	Classifiers	Classifier Descriptions
		Table Capacity	High number of table entries relative to capacity Sub-Classifiers: • Flow—Session flow table utilization • FIB—Forwarding Information Base (FIB) table utilization
		DHCP Pool	High DHCP utilization relative to pool size Sub-Classifiers: • DHCP Denied • DHCP Headroom

Table 88: WAN SLE Descriptions *(Continued)*

SLEs	SLE Descriptions	Classifiers	Classifier Descriptions
WAN Link Health	Juniper Mist monitors the user minutes when the WAN link's health meets or fails to meet the SLE threshold. Poor WAN link health lowers the device's ability to pass traffic, thus directly affecting any clients using that link.	Network	Network issues Sub-Classifiers: • Latency—Juniper Mist calculates latency by using the average value of round-trip time (RTT) for traffic over a period of time. • IPSec Tunnel Down • Jitter—Juniper Mist calculates jitter by using the variation (standard deviation) of RTT within a period of 5 to 10 minutes for a particular WAN link. We compare the calculated value with the average deviation of RTT over a day or a week. • Loss—Lost packets

Table 88: WAN SLE Descriptions *(Continued)*

SLEs	SLE Descriptions	Classifiers	Classifier Descriptions
		Interface	Interface issues Sub-Classifiers: • Congestion—High number of output packet drops. When packets enter an interface, they go in a queue for buffering. When the buffer becomes full it starts to drop packets (TxDrops). • Cable Issues • VPN • Port Down • Negotiation Incomplete (SRX only)
WAN Application Health	Juniper Mist monitors the latency of WAN applications to identify applications that are performing sub-optimally. This SLE can help you to understand the end users' experiences when accessing applications. For example, a weak network connection might give good user experiences for FTP or SMTP-based applications, but bad user experiences for VoIP applications.	Jitter	Inconsistent packet transmit times
		Latency	Slow response time (lag)
		Loss	Packet loss

Table 88: WAN SLE Descriptions *(Continued)*

SLEs	SLE Descriptions	Classifiers	Classifier Descriptions
	Performance metrics vary by device: • SSR—Values come from the Session Smart bidirectional forward detection between peers • SRX—Values come from variation detection against RTT, lost packets, and latency above the RTT For fine-tuning, you can click the **Settings** button to select individual applications to include or exclude.	Application Services (SSR only)	Issues such as slow responses to application requests, recurring disconnects, and insufficient bandwidth Sub-Classifiers: • Slow Application • Application Bandwidth • Application Disconnects

Table 88: WAN SLE Descriptions *(Continued)*

SLEs	SLE Descriptions	Classifiers	Classifier Descriptions
Gateway Bandwidth	Juniper Mist evaluates the IPsec overlay that constitutes the SD-WAN. Use this SLE to determine if you need more WAN bandwidth on your site.	Bandwidth Headroom	Current usage exceeding the baseline, which is determined by the highest usage over the past 14 days If you've enabled automatic speed tests, these results also are incorporated into the Bandwidth Headroom classifier. In this case, the headroom threshold is based on maximum usage and the speed test results, if available. Speed tests occur if configured in your organization settings and if enabled in the WAN settings for the WAN Edge template, hub profile, or WAN Edge device.
		Congestion Uplink (SRX Only)	High ratio of total transmitted bytes dropped (TX drops) to total packets transmitted (TX packets).

WAN Edge Testing Tools

SUMMARY

Follow these steps to add Juniper Mist™ testing tools to your troubleshooting process.

IN THIS SECTION

The Juniper Mist portal provides testing tools that play crucial roles in maintaining network health and diagnosing issues. You can use the tools to:

- Identify network bottlenecks, latency, and packet loss.
- Check connectivity and measuring round-trip time.
- Identify the path taken by network traffic from source to destination for a specific application.
- Monitor BGP peering status and troubleshoot connectivity issues.

To use WAN Edge testing tools:

Finding the WAN Edge Testing Tools

After you've selected a WAN Edge device on the WAN Edges page, you'll find Testing Tools in the Utilties menu.

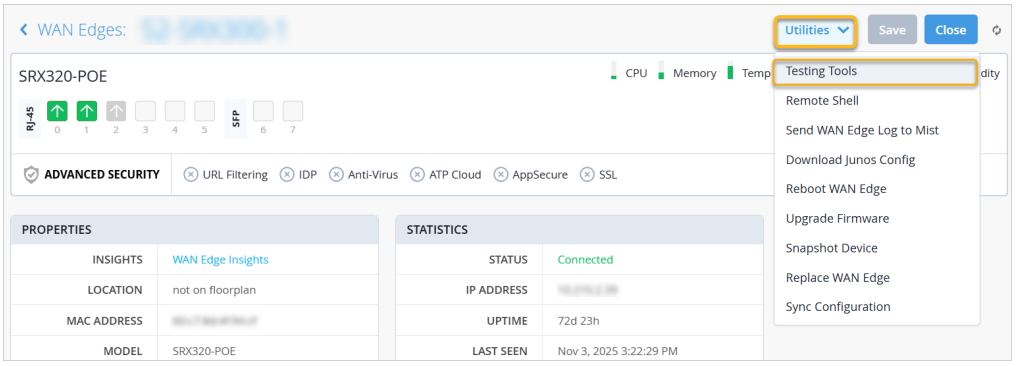

Testing Tools

On the WAN Edge Testing Tools page, use the buttons and tabs to select a testing tool. The buttons represent testing categories, and the tabs represent specific testing tools.

Table 89: Testing Tools

Category (button)	Tool (tab)	Description
Applications (SSR Only	Path	Enter the application name, and then click **Show Path**. You'll see details for each path.
	Sessions	See the path information for a specific application. You can also delete sessions. Deleting a session might be useful in situations where a stuck session was created and traffic is not sending due to upstream problems. Enter the application name, and then click **Show Sessions**. You'll see details for each session. To delete sessions: • One or More Sessions—Select the check box, and then click **Delete Selected**. • All—Click **Delete All Sessions**.
ARP (SSR Only)	Refresh ARP	During troubleshooting scenarios, remove Address Resolution Protocol (ARP) entries from the WAN Edge device or the node's ARP cache. Enter the port name (required) and other details (optional), and then click **Refresh ARP**. By using the optional details, you specify which entries to remove.

Table 89: Testing Tools *(Continued)*

Category (button)	Tool (tab)	Description
	Table	Click **Show ARP** to see all interface details including current state.
BGP	Clear BGP	Clear sessions with all BGP neighbors or with specific BGP peer. Enter the neighbor ID (required) and other details (optional), and then click **Clear BGP**.
	Advertised Routes	Display all routes that have been advertised to the specified neighbor. Enter a neighbor IP (required) and VRF (optional). Click **Show Routes**. You'll see the routes that your device has advertised to the peer router during the current BGP session.
	Received Routes	See the routes that your device has advertised to the peer router during the current BGP session. Enter a neighbor IP (required) and VRF (optional). Click **Show Routes**. You'll see the routing information as it was received through the specified neighbor.

Table 89: Testing Tools *(Continued)*

Category (button)	Tool (tab)	Description
	Routes	Debug the BGP routing table. Click **Show Routes**. Optionally, you can specify a prefix and VRF to narrow down the list of routes. You'll see how the sent or received prefixes from various neighbors are being handled and processed in the BGP table.
	Summary	Click **Show Summary**. You'll see details for all BGP entries.
FIB (SSR Only)	FIB Lookup	View the forwarding information base (FIB) data associated with the WAN Edge device. Enter the network, IP, protocol (all required) and destination port (optional). Then click **Show FIB**. You'll see information for each FIB.
	FIB By Application	Enter the application (required) and other optional information. . Then click **Show FIB**. You'll see information for each application.
OSPF	Database	Click **Show Database**. Optionally, to narrow down the results, first select a Self Originate option and enter a VRF. You'll see area, LSA, and routingdetails for each VRF.

Table 89: Testing Tools *(Continued)*

Category (button)	Tool (tab)	Description
	Interfaces	Click **Show Interfaces**. Optionally, to narrow down the results, first enter a VRF. You'll see interface details for each VRF.
	Neighbors	Click **Show Neighbors**. Optionally, to narrow down the results, first enter IP, interface, or VRF. You'll see neighbor information for each VRF.
	Routes	Click **Show Routes**. Optionally, to narrow down the results, enter a route prefix or VRF. You'll see route details for each VRF.
	Summary	Click **Show Summary**. Optionally, to narrow down the results, enter a VRF. You'll see area details for each VRF.
Utility	Bounce Port	Provisionally take down the port and bring it up without causing the external physical link to change. The connected devices will not see a link state change. Enter the port, and then click **Soft Bounce Port**.

Table 89: Testing Tools *(Continued)*

Category (button)	Tool (tab)	Description
	Ping	Use the Ping tool to send Internet Control Message Protocol (ICMP) echo requests to a specified host. You can use this command to check if a particular host is reachable at a particular IP address. Enter the IP address and the port name. Click **Ping**. Optionally, you can specify the count (the number of requests to send, 1-100) or the ping packet size in bytes (64-65,535.
	Show FIB (SRX Only)	
	Traceroute	Traceroute is a network diagnostic tool that traces an IP packet's path to a host, revealing the sequence of routers (hops) it takes along the way. Enter a host, and then click **Traceroute**. Optionally, also enter a network.
	WAN DHCP Release (SSR Only)	Release client devices from their current DHCP lease. Enter the port name, and then click **Release**.

Speed Tests for WAN Edge Devices

SUMMARY

Follow these steps to do speed edge tests on your WAN Edge devices.

IN THIS SECTION

Service Providers (SPs) as well as their end customers install and deploy telecommunication circuits (or paths) to offices, branches, and so on. As Session Smart Routers or SRX Series Firewalls are deployed at the edge of the customer premise, SPs and customers need to generate traffic to test the speed and performance of these circuits to ensure the quality is being maintained.

You can test the speed of WAN links on WAN Edge devices to ensure optimal performance. This feature allows for new link qualification and on-demand speed tests if a low link speed is suspected.

NOTE: In addition to Mist-managed WAN Edge devices, you can now perform Speed Tests on conductor-managed Session Smart Routers (SSRs) from the Juniper Mist™ portal.

From the Juniper Mist™ portal, you can use a speed test to:

- Test the speed and performance of the circuit being delivered to the customer.
- Perform new link qualification to verify that speeds are what the service provider and customer have agreed upon.
- Perform on-demand speed tests when you suspect a low link speed is causing link issues.
- Run scheduled speed tests to re-test link speeds and ensure performance continues to meet expectations on an ongoing basis.

NOTE: The WAN Edge Speed Test tool can reliably validate circuits speeds of 1 megabit per second (Mbps) to 1 gigabit per second (Gbps). Circuits exceeding 1Gbps must rely on other tools for validation.

The WAN Edge Speed Test tool does not measure or validate jitter or loss.

Perform Speed Test on WAN Edge Device

Use the following steps to perform the on-demand Speed Test:

1. Navigate to the Juniper Mist portal, then click **WAN Edges** > **WAN Edges**.
2. Select the WAN Edge name.

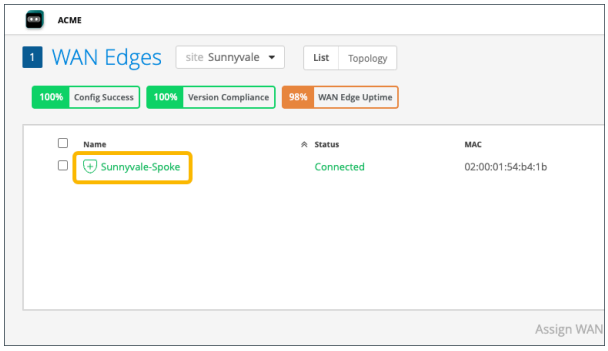

3. Select the WAN port from the Port Panel.

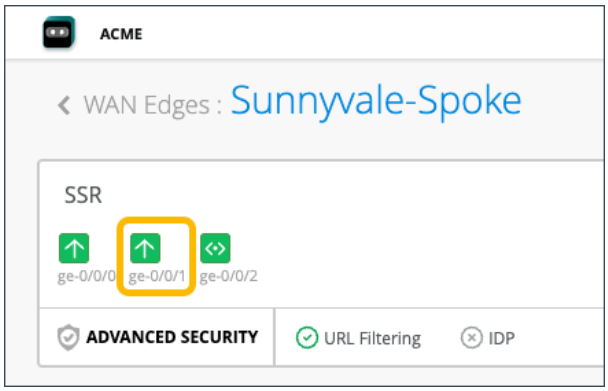

4. The Networks section then appears. From the **Networks** section, under the Speed Test column, click **Run Speed Test**.
5. In the Speed Tests section, the test you are currently running is **In Progress**. After a minute or two, the progress of the speed test changes to **Succeeded**, and the results populate in the columns, displaying information such as download speed, upload speed, latency, the interface the test was run on, and the VLAN number.

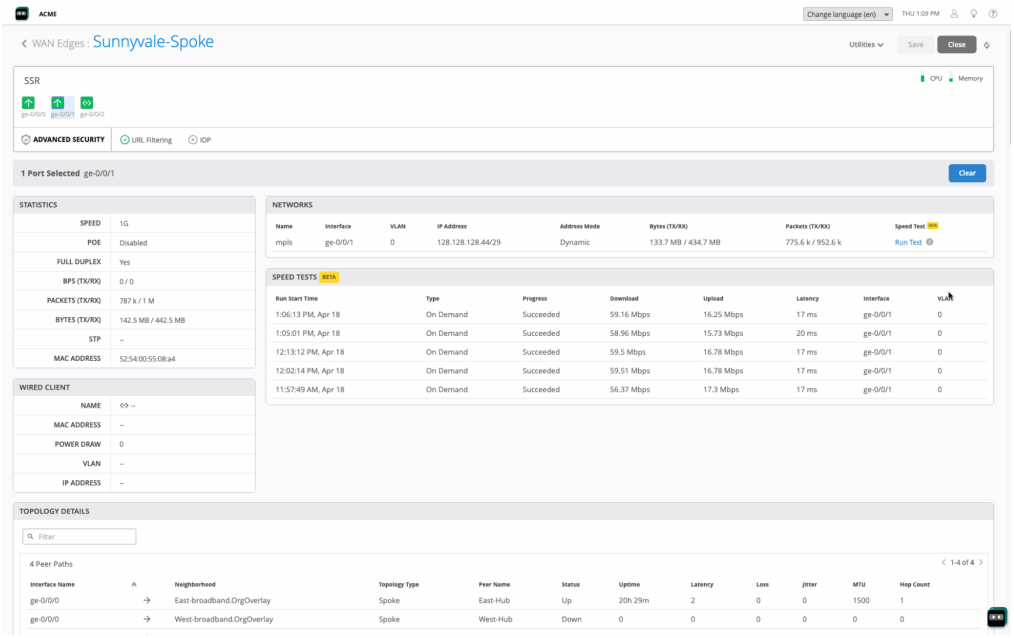

Schedule a Speed Test

If you need to run speed tests on a recurring basis, you can use the WAN Speed Test Scheduler.

1. Navigate to **Organization** > **Admin** > **Settings**.
2. In the **WAN Speed Test Scheduler** section, select **Enabled**.
3. Uncheck the **Enable Automatic Speed Test** checkbox.
4. Select the **Time of Day** and **Day of Week** that you want the speed tests to occur.
5. Select whether you want to run the tests on all WAN interfaces or select WAN interfaces.

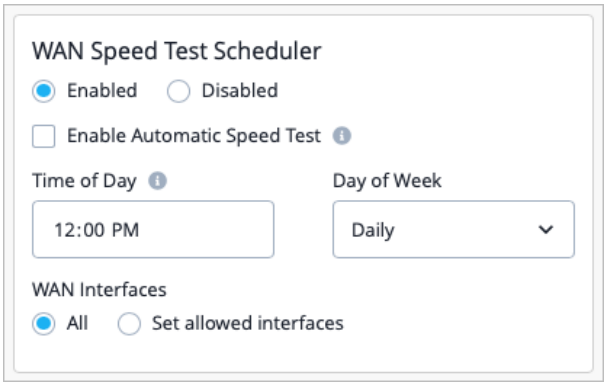

6. Save your settings.

Enable Automatic Speed Tests (BETA)

To have Marvis automatically run self-driving speed tests on selected WAN interfaces during times of low activity, you can enable automatic WAN speed tests. You can configure this in the WAN interface configuration at the WAN Edge template-level or the device-level. For Ethernet interfaces, the Enable Scheduled Speed Test checkbox is selected by default. The checkbox applies to both automatic and scheduled speed tests.

1. Navigate to **Organization** > **Admin** > **Settings**.
2. In the **WAN Speed Test Scheduler** section, select **Enabled**.
3. Select the **Enable Automatic Speed Test** checkbox.

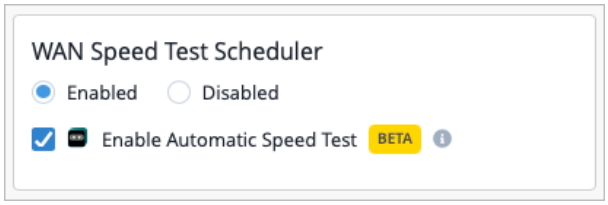

4. Now enable automatic speed tests at the WAN Edge template-level by navigating to **Organization** > **WAN** > **WAN Edge Templates** and select the template, or at the WAN Edge device-level by navigating to **WAN Edges** > **WAN Edges**, and then select the device.
5. Scroll down to the WAN section, then select the appropriate WAN interface.
6. Select or deselect the **Enable Scheduled Speed Test** checkbox as needed for the necessary interfaces.

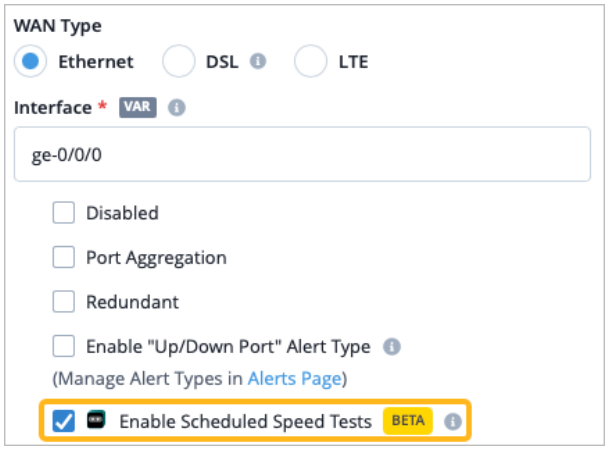

NOTE: For Ethernet interfaces, the Enable Scheduled Speed Tests checkbox is selected by default. For devices with built-in LTE interfaces, the checkbox is unchecked by

default. Automatic speed tests for LTE devices are not enabled by default , as it can result in costliness if you pay for bandwidth usage through a service provider. However, if you have LTE interfaces connected to your Ethernet interfaces, such as on your Cradlepoint devices, you must deselect the checkbox from the Ethernet interface configuration.

7. Save your settings.

8. You can view the Speed Test results from the WAN Edge page or by navigating to **Marvis** > **Marvis Minis**, then select WAN Speed Test from the drop-down menu at the top of the page. You will then see a list of completed speed tests per site and device along with the results.

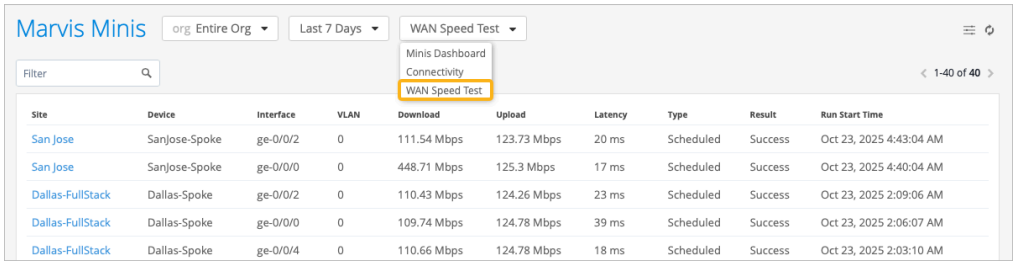

Review Speed Tests in Marvis Minis

The Marvis Minis dashboard provides visibility into the validation results. You can schedule speed test and view the results in Marvis Minis. This proactive approach helps in identifying and resolving network issues before they affect users.

Video: Marvis Minis Speed Test

To view the speed test results in Marvis Minis, you must:

- "Schedule the Speed Test" on page 388
- Enable Marvis Minis at the organization level (Organization> Settings) or at site level (Organization > Site Configuration).

1. To view the Marvis Minis dashboard, select **Marvis** > **Marvis Minis** from the left menu.

   In the Marvis Minis page, you'll see the major elements of the dashboard including the summary of speed test.

2. At the top of the page, use filters to select **WAN Speed Test** from the drop-down box. You can also select to display results at organization-level or at site-level and set the time period for the results.

Figure 171: Marvis Minis Dashboard

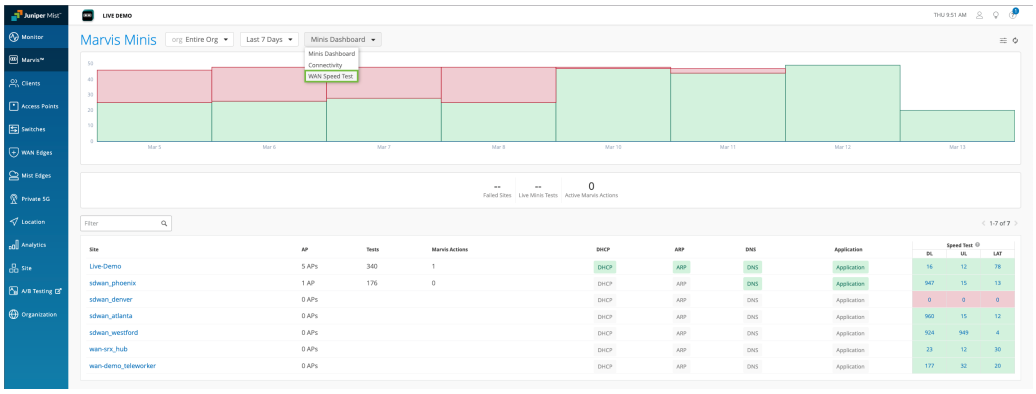

3. You can view the details of speed test results. Results populate in the columns, displaying information such as site, device, interface, download speed, upload speed, latency, type, result of the test, and start time of the test.

Figure 172: Marvis Minis WAN Test Speed Results

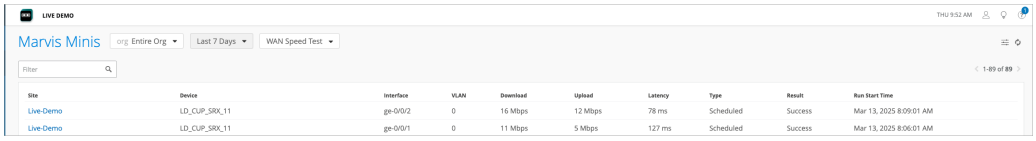

4. You can click a site to display the test results at that site-level.

Configure System Logging

SUMMARY

Send log messages to a named file, a remote location, a user, or the console. Configure the settings to receive the messages that you want to see.

IN THIS SECTION

You can enable system logging in a hub profile, WAN Edge template, or device configuration.

1. Navigate to the hub profile, WAN Edge template, or device configuration.

- For a WAN Edge template—From the left menu, select **Organization** > **WAN** > **WAN Edge Templates**. Click a template, or create a new one. Scroll down to the **WAN** section.
- For a hub profile—From the left menu, select **Organization** > **WAN** > **Hub Profiles**. Click a profile, or create a new one. Scroll down to the **WAN** section.
- For an individually managed WAN Edge device—From the left menu, select **WAN Edges** > **WAN Edges**. Click a device. Scroll down to the **WAN** section.

2. Scroll to the **Syslog** section.
3. Select **Enabled**.
4. On the first four tabs, identify the resources/recipients that will receive the logs, and specify the contents to send.

   These tabs include:

   - **Files**—Send log messages to a named file.
   - **Hosts**—Send log messages to a remote location. This could be an IP address or hostname of a device that will be notified whenever those log messages are generated.
   - **Users**—Notify a specific user of the log event.
   - **Console**—Send log messages of a specified class and severity to the console. Log messages include priority information, which provides details about the facility and severity levels of the log messages.

   To add/edit/remove files, hosts, users, or consoles:

   - Add—Click the **Add** button. Enter the settings, and then click **Add Contents** (near the bottom of the Add window) to specify the log contents. You can repeat this step to add additional content types. When finished, click **Add** at the bottom of the window.
   - Modify—Click the item to modify, and then make your changes in the Edit window. Be sure to save your changes.
   - Remove—Click the item to remove, and then click the **Delete** button at the bottom of the Edit window.
5. On the **Archive** tab, define parameters for archiving log messages.
6. On the **General** tab, specify the settings to use for the log messages.

   For more information about the various options, see "Syslog Options" on page 113.
7. Configure other settings as needed, and then click **Save** at the top-right corner of the page.

Syslog Options

SUMMARY

Use this information to send system log messages to local files, external hosts, users, or the console. Also configure content types, archive settings, and more. (Available logging options vary, depending on which settings you're configuring, such as syslog for a switch, a WAN Edge device, or other.)

Table 90: Files

Field	Description
Add File	Use the **Add File** link to specify a named file in the local file system. Syslog messages that match the specified criteria will be sent to this file. Cycle through this process as needed to specifying multiple files for different matching criteria.
File Name	Enter a name for the file.
Match	To filter messages, enter a text string. Only log messages that contain this string will be saved to the specified file.
Explicit Priority	Select this option to include priority information in the file.
Structured Data	Select this option to use structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.
Archive	Set the maximum number of files to store, and the maximum file size in megabytes (m) or bytes.

Table 90: Files *(Continued)*

Field	Description
Contents	To specify the types of log information to capture, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be saved to the specified file. For more information about the content options, see Table 93 on page 396 later in this topic.

Table 91: Hosts

Field	Description
Add Host	Use the **Add Host** link to specify an external server that is configured as a system log message host. Syslog messages that match the specified criteria will be sent to this server. Cycle through this process as needed to specifying multiple hosts for different matching criteria.
Host	Enter the host.
Port	Enter the port number, from 1 through 65535.
Match	To filter messages, enter a text string. Only log messages that contain this string will be saved to the specified file.
Explicit Priority	Select this option to include priority information in the file.
Structured Data	Select this option to use structured-data format instead of the standard Junos OS format. Structured-data format provides more information without adding significant length, and makes it easier for automated applications to extract information from a message.

Table 91: Hosts *(Continued)*

Field	Description
Archive	Set the maximum number of files to store, and the maximum file size in megabytes (m) or bytes.
Contents	To specify the types of log information to capture, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be saved to the specified file. For more information about the content options, see Table 93 on page 396 later in this topic.

Table 92: Users

Field	Description
Add User (link)	On the Users tab, you're sending log messages to the terminal session of one or more users. To add a user, click this link, and the enter the user information. Cycle through this process as needed to specify multiple users for different matching criteria.
User	Enter one or more usernames, separated by spaces, or enter an asterisk (*) to include all users.
Match	To filter messages, enter a text string. Only log messages that contain this string will be sent to the specified users.
Contents	To specify the log information to capture for this user, click **Add Content**, select the **Facility** and the **Level**, and then click the check mark in the Add Content title bar. Repeat as needed until you've specified all the content that you want to capture. Only logs that meet these conditions will be sent to the specified users. For more information about the content options, see Table 93 on page 396 later in this topic.

Table 93: Console Options and Content Types

Field	Description
Add Content (link)	Use the **Add Content** link to specify the types of log messages to capture. You'll cycle through this process multiple times to add different content types. For example, let's say you want to capture (1) critical authorization events, (2) warnings for authorization events, (3) critical change logs, (4) change log errors, and (5) ftp errors. For this example, you'd create five content types. **NOTE**: You can add content types in a few different places. For example, if you're working on the Console tab, you're specifying the content to send to the console. If you're working on the User tab, you're specifying the content to capture for that user.
Facility	Select a log event that you want to capture in this content type.
Level	For the specified log event, select the severity level to capture in this content type.

Table 94: Archive Tab

Field	Description
Files	Specify the maximum number of files (1-1000) to store in the system log. After this maximum is reached, log messages are archived.
Size	Specify the maximum size (65536 - 1073741824) of the system log. Select either **m** (megabytes) or **bytes**. After this maximum is reached, log messages are archived.

Table 95: General Tab

Field	Description
Time Format	Use this option to include the milliseconds, the year, or both in the syslog timestamps. • To add an option, click the plus (+) button, then click the option in the list. • To remove an option, click the X.
Routing Instance	By default, system logging traffic is sent from the management interface on your device and its associated routing instance. However, the Junos syslog client is completely VRF aware. If a server is reachable through a virtual routing and forwarding (VRF) instance, the syslog client can send log messages to the server. Use this option to specify a routing instance to use.
Network	Use this option to specify a network instance to use.

Dynamic and Manual Packet Captures

SUMMARY

When investigating communication failures between the client and the access point (AP), you can use the Juniper Mist™ portal to get dynamic and manual packet captures.

IN THIS SECTION

NOTE: Mist does not collect or store any payload data from packets capture. Only transmission and connection data are used.

Dynamic Packet Captures

IN THIS SECTION

Which Events Trigger Dynamic Packet Captures?

Whenever a connection failure occurs between the wireless client and an AP (AP), it automatically triggers a short-term dynamic packet capture.

These events include:

- DHCP Timeout—When the client sends a broadcast discover packet but does not receive an offer packet from server.
- DHCP Denied—When the server sends a DHCP NAK, indicating that the IP address might already be in use.
- DHCP Terminated—When the Client does not proceed with DHCP request for the offer provided by the server.
- Authorization Failure—This type of failure could be caused due to various reasons. Examples include MIC failure, RADIUS server not responding, Access-Reject from RADIUS server, client failing to complete the auth process, and so on.
- 11r FBT Failure—This type of failure is caused due to client failing 11r roam.
- OKC Auth Failure—This type of failure is caused due to client failing OKC roam.
- Association Failure—This type of failure could be caused due to transmission failures or an invalid PMKID included by the client during association request.

Finding the Packet Captures

Dynamic packet captures are saved to the cloud. You can download these files from the Insights page.

Video Demo

Video: NOW in 60: WAN Assurance - Dynamic Packet Capture

Example

This example shows how easily you can find dynamic packet captures on the Insights page.

1. From the left menu, select **Monitor** > **Service Levels**.
2. Click the **Insights** button to view the Insights page.
3. Scroll down to the **Client Events** section.

   Paperclip icons indicate the events with dynamic packet captures.
4. Click an event to see more details on the right side of the screen.
5. Below the details, click **Download Packet Capture**.

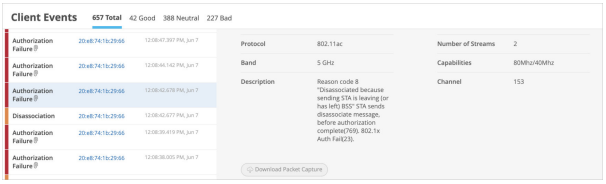

Manual Packet Captures

For manual packet captures, go to **Site** > **Packet Captures**, where you can:

- Choose which network type to capture packets from: wired, wireless, or WAN.

  NOTE: Wired packet capture applies to the wired ports of APs (not the switch ports). The switch must be running a CloudX version of Junos for it to appear in the **Add Switch +** selection window. WAN packet captures support Session Smart Router and SRX WAN edge device ports.

- Restrict the packet capture to specific clients, WLANs, APs, or wireless bands.

- Configure the number of packets captured, packet size in bytes, and the duration of the capture session.
- Configure other capture parameters such as header inclusion and capture filters. See Table 96 on page 402 for details.

After downloading the packet capture to your computer, follow the steps below to view the data in Wireshark.

Configure IEEE 802.11 on Wireshark

Packet inspection requires Wireshark. See https://www.wireshark.org for the download file and related information.

To configure Wireshark to view packets captured from the Juniper Mist portal, follow the steps below:

1. Open the Wireshark application on your computer.
2. Open the Wireshark Preferences window:

   On a Windows computer, navigate to **Edit > Preferences**.

   On a Mac computer, navigate to **Wireshark > Preferences**.
3. In the Preferences window, expand the **Protocols** menu option and scroll down to **IEEE 802.11**.
   a. Select **Yes - with IV** and then click **OK**, as shown in the following image:

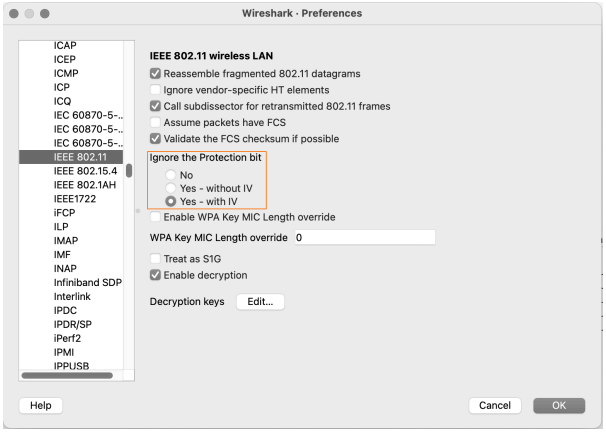

View Wireless Packet Captures in Wireshark

You can capture packets from both your wired and wireless networks. The following configuration regards wireless packet, for which you can see:

- Wireless channel information
- Wireless data rate
- Received signal strength indicator (RSSI)

To accomplish this task, you must download and install the Wireshark application on your computer. In a browser, navigate to https://www.wireshark.org for Wireshark application downloads and detailed information about Wireshark. For additional information about Wireshark, see https://www.wireshark.org/docs/.

This topic provides minimal guidance about how to configure Wireshark for use in examining wireless packet captures gathered from the Juniper Mist portal.

1. Open the Wireshark application on your computer.
2. Open the Wireshark Preferences window:

   On a Windows computer, navigate to **Edit > Preferences**.

   On a Mac computer, navigate to **Wireshark > Preferences**.
3. In the Preferences window, navigate to **Appearance > Columns**.
4. Click the **Add (+)** button to add a new radiotap column to the Wireshark display.

   Wireshark adds a new line called New Column, and the type Number.

   Radiotap headers include wireless packet frames that would otherwise not be displayed. See: https://www.wireshark.org/docs/dfref/r/radiotap.html.

   a. Double-click the **New Column** title and type Channel as the title.

   b. Double-click the **Type** column and select Frequency/Channel from the drop-down menu.

   c. Leave the **Displayed** column selected.
5. Repeat Step 4 two times

   a. The first time, use **Data Rate** for the column title and **IEEE 802.11 TX Rate** for the type.

   b. The second time, use **RSSI** as the column title and **IEEE 802.11 RSSI** for the type.
6. Click **OK** to save your changes.

   Wireshark will display the new columns when you open a packet capture (.pcap) file for viewing.

Manual Packet Capture Options

By default, Juniper Mist streams the packet capture session data, including beacon frames, to the Mist portal. The following table describes the packet capture options that you can use when you create a packet capture session.

Table 96: Packet Capture Options

Option Name	Option Function	Usage Notes	Firmware Notes
Include Network Headers	This feature includes packet headers with the packet data.	Packet capture works by buffering packets locally on the device, which has limited space available. By default, Mist truncates header data from the captured packets to reduce the size of capture files while still providing the most relevant information.	–
Local Capture	This capture is local only and is not streamed to the Mist portal.	Earlier AP firmware did not support live streaming packet captures to the Juniper Mist portal.	Required for AP firmware versions before 0.10.x
Canned Filters	These filters are based on the type of packet capture that you're performing.	The filters available in the list change depending on whether you're capturing wireless, wired, or WAN packets. For example, beacon frames are only available for wireless packet captures.	–
Advanced Filters	Use this option to apply your own filters by using `tcpdump` syntax.		0.10.x or later
Expression Builder	This interactive tool builds custom filters in `tcpdump` syntax for use in the capture session.	You can let the builder start the filter entry and then add to or delete from the entry manually.	0.10.x or later

Troubleshoot Session Smart Router Deployed as WAN Edge

SUMMARY

Follow these steps to enable alerts, use testing tools, and get packet captures.

IN THIS SECTION

Discover efficient methods for troubleshooting your WAN edge device in the Juniper Mist™ portal after the initial deployment phase. This topic provides guidance for system administrators and technical support responsible for maintaining enterprise SD-WAN networks.

Troubleshooting Session Smart WAN Edge Alarms

In this section you'll configure WAN Edge alarms and e-mail alerts to the administrator.

1. From the Mist dashboard, select **Monitor** > **Alerts**.

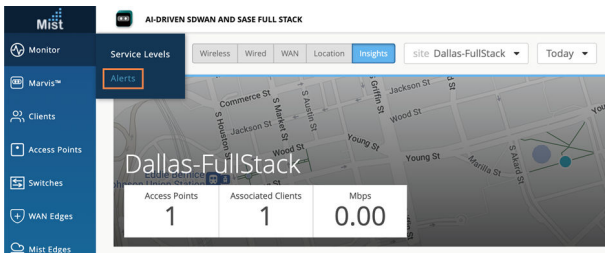

2. Filter the alarms for our lab Site, **Dallas-FullStack**.

3. Set the **Time Range** for **This Week** so we can see some Alerts. Next, we'll select the **Alerts Configuration** to set up e-mail notifications.

> **NOTE**: Under Alerts Configuration, it's best practice for **Applies to Scope** alerts to be set for the Entire Org (as shown in the screenshot below). You can also specify that alerts be sent to organization administrators and site administrators by selecting the relevant check boxes in the Email Recipient Settings section. Enter emails for additional recipients in the **To additional email recipients** field.

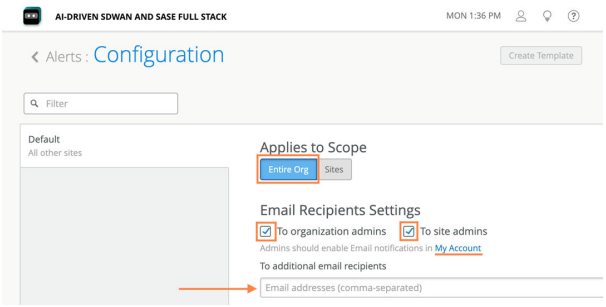

4. Use **ENABLE** button on the Alerts Configuration in the **Email Notification** section. It's important to note that even for an administrator, e-mails are disabled by default.

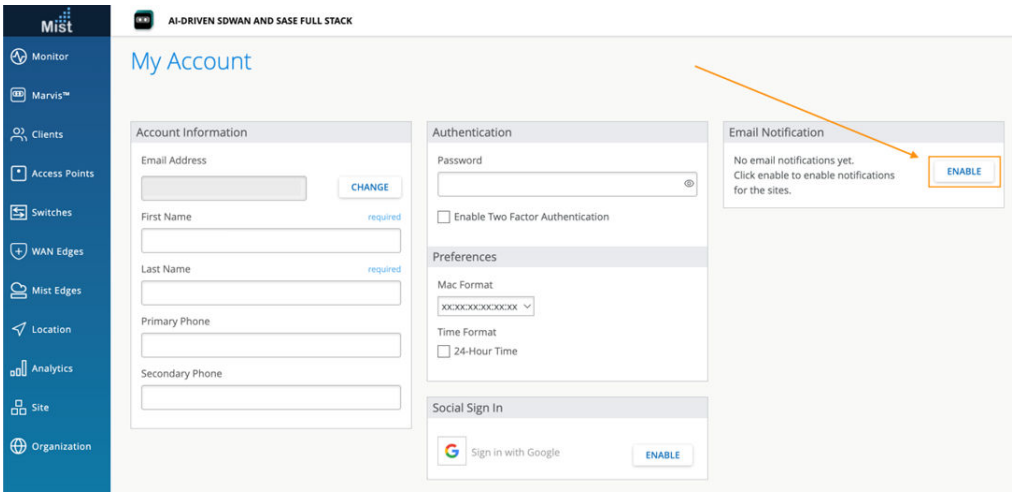

You're directed to the Enable Email Notifications window where you enable e-mail alerts by selecting the toggle for Enable Org Notifications, and selecting specific sites, in the example below you can see we're getting notifications for our Dallas-FullStack site.

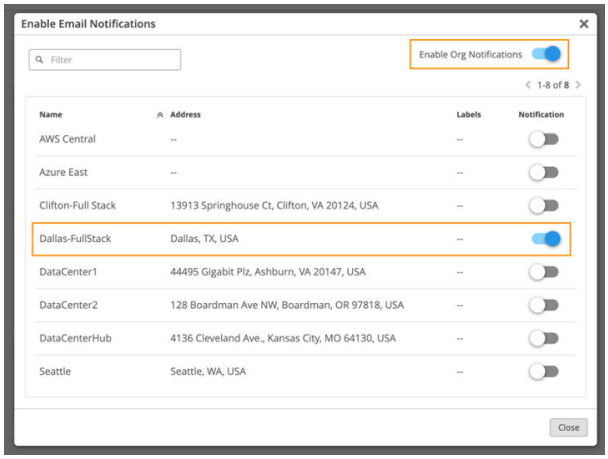

5. Now enable the gateway alerts and e-mail notifications for Infrastructure as shown in the options below:

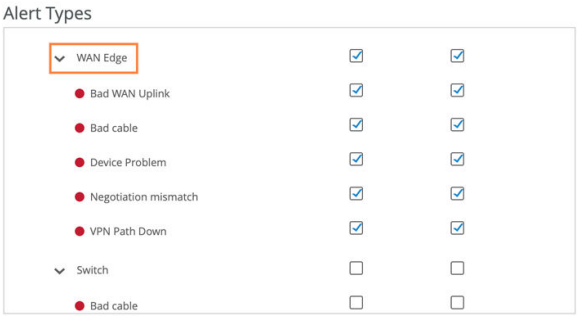

We recommend enabling the WAN edge alerts and e-mail notifications as well.

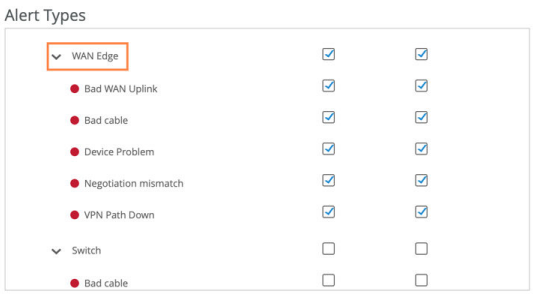

When a device loses connection to the Juniper Mist cloud, your administrator will receive an e-mail about the event after the issue is resolved. The Alert Details section in the e-mail will direct you to the Alerts page, but you can also look directly at it by seeing the Event reported below:

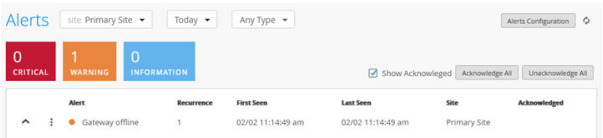

After the connection to the Juniper Mist cloud has been restored, another e-mail will indicate the status change with a link to the Alerts page, with the second event reported.

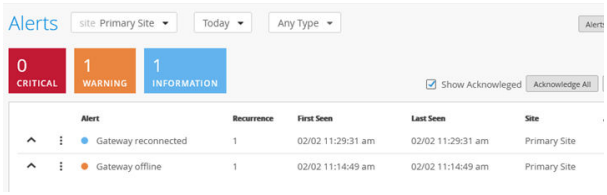

Session Smart Router Testing Tools for Troubleshooting

The Session Smart Router in WAN Assurance relies on SLE and Marvis insights for all your troubleshooting actions. An administrator's ability to see beneath the surface of the GUI is limited to the utilities we'll cover in this section. If you are familiar with the Session Smart Router device and its Programmable CLI (PCLI), which runs the Conductor-based deployments, you'll have a simplified and streamlined experience diagnosing your WAN Edge device using the Mist UI.

First, we'll dive into the **Utilities** button on your selected WAN Edge. Navigate to **WAN Edges** > **WAN Edges**, and we're selecting our Dallas-lab WAN Edge in this example.

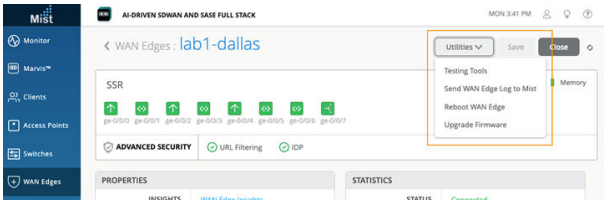

In the **Utilities** menu, you'll see the options **Testing Tools**, **Send WAN Edge Log to Mist**, **Reboot WAN Edge**, and **Upgrade Firmware**. Sending the WAN edge device logs to the Juniper Mist cloud and rebooting your selected WAN edge device are straightforward activities. For upgrading firmware, you'll have a few options for scheduling and version, as you can see in the screenshot below:

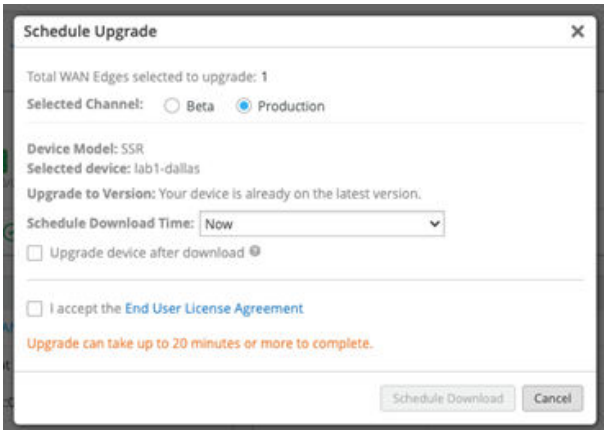

For troubleshooting, let's explore the **Testing Tools** option in the **Utilities** menu.

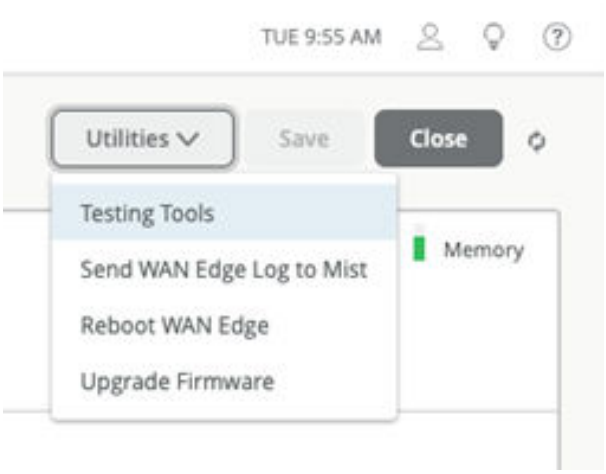

The **Testing Tools** option allows you to issue ICMP-Ping, identify BGP status reviews, sift through Application and Session knowledge, and review Address Resolution Protocol (ARP) tools for your WAN edge. The Mist UI enables quick troubleshooting for device interfaces and connectivity with pings. In this example, we send an ICMP ping to the public DNS 8.8.8.8. After you select Ping from the Utility section, enter an IP address, and specify the port name and the ping count and size, Mist opens an instance of the Session Smart PCLI for you.

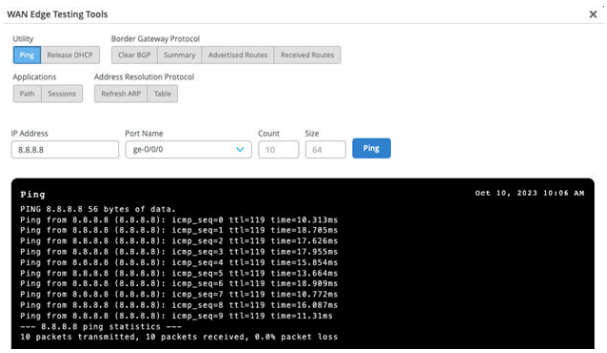

When you select any of the WAN Edge Testing tools for BGP, Mist simplifies the Session Smart PCLI and gathers the information for you. In this example, we see a summary of BGP neighbors and their relevant information by quickly selecting **Summary** under Border Gateway Protocol and then clicking **Show Summary**.

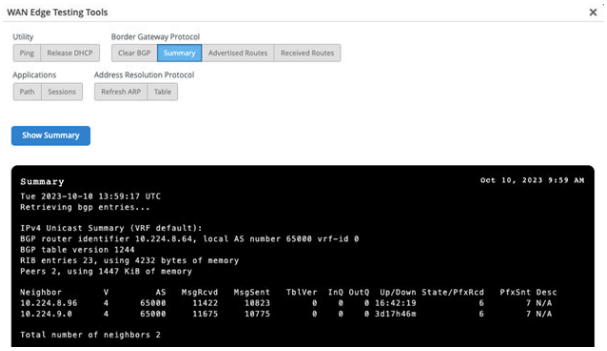

Remember that Mist is intent-driven, and the Session Smart Secure Vector Routing (SVR) protocol leverages the sources and destinations in Mist, the networks, and applications to create point-to-point SVR paths. You can diagnose those paths between Session Smart devices in the **WAN Edge Testing Tools** utility. In this example, we're looking at the Corp Application. In this example, you'll also notice that the actual BGP advertisements are sent through SVR.

NOTE: For SVR, Mist applications are called *services* in the Session Smart paradigm. For more information about the Session Smart platform, refer to the platform considerations, https://www.juniper.net/documentation/us/en/software/mist/mist-wan/topics/concept/mist-wan-platform-consideration.html.

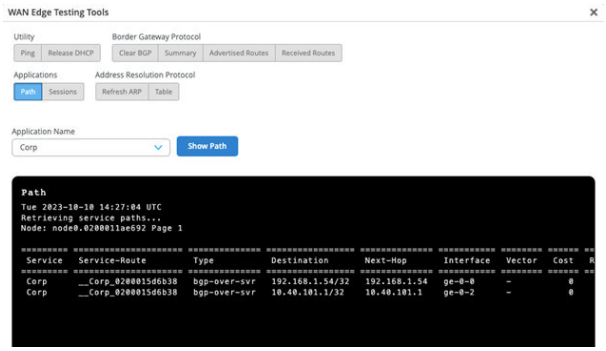

Finally, let's explore ARP tools in **WAN Edge Testing Tools**. In our example, we've selected a the table format for displaying the ARP information, and again, the Mist dashboard leverages the Session Smart PCLI for you and generates and generates ARP information for interfaces on node.0.

NOTE: For deeper insights not shown on the output text of the window, select the output from the Utility in your browser and paste it to an ASCII editor for further review.

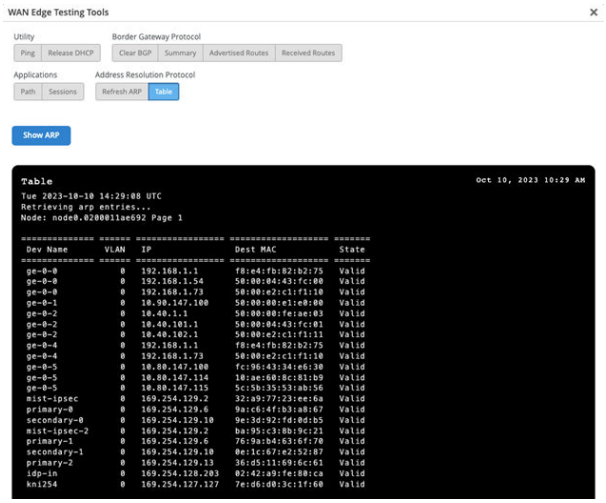

The combination of **Marvis insights**, **WAN SLEs**, **WAN edge insights**, and the **WAN Edge Testing Tools** utility make a complete suite for diagnosis at the WAN edge.

SEE ALSO

Monitor Session Smart Router Deployed as WAN Edge | **341**

Troubleshoot Session Smart Router Deployed as WAN Edge | **403**

Troubleshoot Session Smart Routers Using Packet Captures

IN THIS SECTION

Juniper Session Smart Router (SSR) ports support manual and dynamic packet captures (PCAP). Packet capture is a tool that helps you to analyze network traffic and troubleshoot network problems. It captures real-time data packets traveling over the network for monitoring and logging.

For more information, see "Dynamic and Manual Packet Captures" on page 397.

Initiate a Manual Packet Capture

Manual packet captures are initiated by users from the WAN Edge Packet capture page.

To initiate manual PCAP for an SSR device:

1. Go to **Site** > **WAN Edge Packet Captures**.
2. On the WAN tab, click **Add WAN Edge +** and select an SSR device.

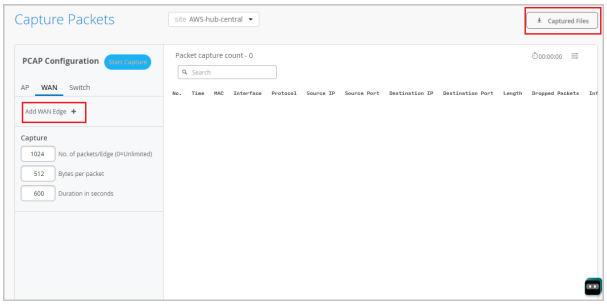

3. Specify the number of packets captured, packet size in bytes, and the duration of the capture session.

You can click **Captured Files** to access a manual PCAP file and analyze it.

Access Dynamic Packet Captures

Dynamic packet captures are short-term packet captures automatically triggered by a service impacting event. They are saved to the Mist cloud and are available for you to download in the WAN Edge Events section on the WAN Edge Insights page.

Here are the events which can generate a dynamic packet capture for an SSR device:

- Failure of the ARP request to next-hop gateway
- DHCP address resolution failure
- WAN Edge failure to establish BGP peering based on configuration
- WAN Edge failure to establish SVR peering based on configuration

To download and analyze a dynamic packet capture file:

1. Go to **Monitor** > **Service Levels**. By default, you will land on the Insights tab.
2. Scroll down to WAN Edge Events section. Each event with a dynamic packet capture file available is indicated by a paperclip icon next to the event name.

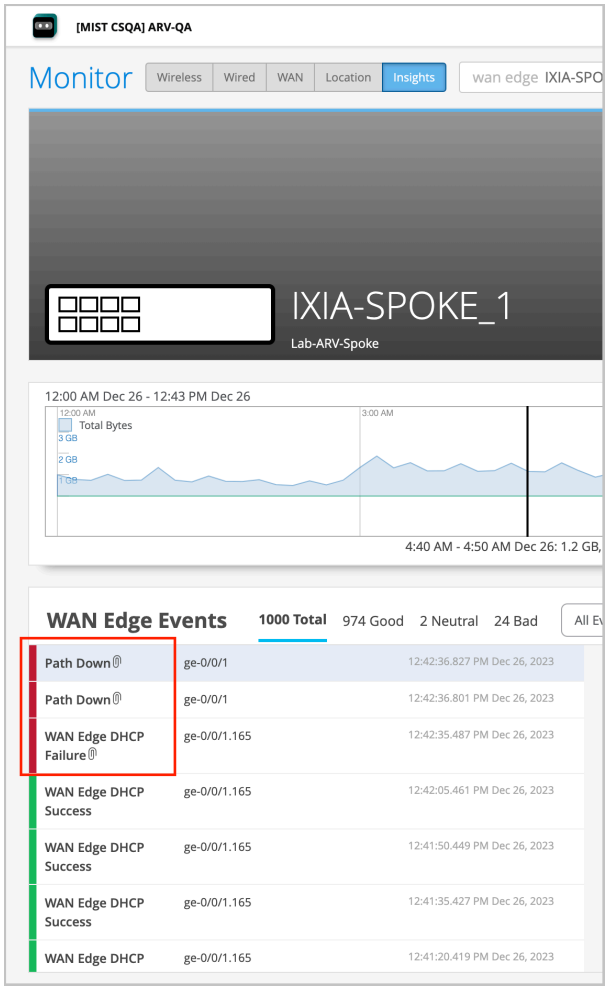

3. Click an event dynamic packet capture and then click the **Download Packet Capture** button available in the Events Details section on the right.

Troubleshoot Disconnected SRX Series Firewalls

SUMMARY

Troubleshoot disconnected and get packet captures (PCAP) for additional insights.

IN THIS SECTION

Troubleshoot SRX Series Firewalls Using Packet Captures | 416

Troubleshoot SRX Series Firewalls Shown as Disconnected

If the Juniper Mist™ portal shows a Juniper Networks® SRX Series Firewall as disconnected when it is online and reachable locally, you can troubleshoot the issue using the steps listed in this topic. You need console access or SSH access to the firewall to perform the troubleshooting steps.

1. Check if the SRX Series Firewall is running on the supported Junos OS version.

   For WAN Assurance, you need Junos OS version 19.4 and later for SRX300, SRX320, SRX340, SRX345, SRX380, SRX550M, and SRX1500. For the SRX1600, SRX2300, and SRX4300, devices must run Junos OS Release 24.2R1.17 and later. For the SRX4700, you need a Junos OS version of 24.4R1-S2 and later.

   You can use the `show version` CLI command to check the version.

2. Check if the SRX Series Firewall has a valid IP address.

   Use the `show interfaces terse` command.

```
user@host >  show interfaces terse 1 match ge-0/0/0
ge-0/0/0      up    up
ge-0/0/0.0    up    up    inet    10.0.0.51/24

user@host > show interfaces terse I match    irb
irb        up    up
irb.0         up    down
irb.2         up    up    inet    192.168.2.1/24
irb.8         up    up    inet    192.168.8.1/24
irb.10        up    up    inet    192.168.10.1/24
irb.24        up    up    inet    192.168.24.1/24
```

   You should see the integrated routing and bridging (IRB) interface (irb.0) with an IP address. You might see multiple IRB interfaces, depending on the SRX Series model (or in the case of a chassis cluster high availability configurations).

At least one IRB interface needs to have a valid IP address. The Firewall can also connect using a management IP address, which you can see on the fxp0 interface.

Ensure that:

- Either the irb or fxp0 interface has a valid IP address.
- The Admin and Link states are up.

3. Ensure that the firewall can reach the gateway as shown in the following sample.

```
user@host> ping inet 10.0.0.1
PING 10.0.0.1 (10.0.0.1): 56 data bytes
64 bytes from 10.0.0.1: icmp_seq=0 ttl=64 time=44.967 ms
64 bytes from 10.0.0.1: icmp_seq=l ttl=64 time=1.774 ms
64 bytes from 10.0.0.1: icmp_seq=2 ttl=64 time=41.347 ms
64 bytes from 10.0.0.1: icmp_seq=3 ttl=64 time=1.731 ms
64 bytes from 10.0.0.1: icmp_seq=4 ttl=64 time=1.674 ms
^C
---10.0.0.1 ping statistics---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 1.674/18.299/44.967/20.329 ms
```

4. Check if your device can reach the Internet. Initiate a *ping* test toward any public server (for example, *8.8.8.8*).

```
user@host> ping inet 8.8.8.8
PING 8.8.8.8 (8.8.8.8): 56 data bytes
64 bytes from 8.8.8.8: icmp_seq=0 ttl=58 time=9.789 ms
64 bytes from 8.8.8.8: icmp_seq=1 ttl=58 time=5.206 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=58 time=4.679 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=58 time=4.362 ms
64 bytes from 8.8.8.8: icmp_seq=4 ttl=58 time=4.497 ms
^C
--- 8.8.8.8 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 4.362/5.707/9.789/2.061 ms
```

5. Check if the firewall can resolve `oc-term.mistsys.net`.

```
user@host> ping oc-term.mistsys.net
PING ab847c3d0fcd311e9b3ae02d80612151-659eb20beaaa3ea3.elb.us-west-1.amazonaws.com
(13.56.90.212): 56 data bytes
```

If the firewall is not resolving `oc-term.mistsys.net`, make sure that the firewall has a DNS server configured.

```
user@host> show configuration | display set | grep name-server
set system name-server 8.8.8.8
set system name-server 8.8.4.4
```

If the firewall doesn't have a DNS server, configure the server as shown in the following example:

```
user@host# set system name-server 8.8.8.8
```

6. Ensure firewall ports are open (for example: tcp port 2200 for oc-term.mistsys.net).

See the following table to determine which port to enable, depending on your cloud environment:

Table 97: Ports to Enable in Different Juniper Mist Clouds

Service Type	Global 01	Global 02	Europe 01
SRX Series	redirect.juniper.net (TCP 443)	redirect.juniper.net (TCP 443)	redirect.juniper.net (TCP 443)
	ztp.mist.com (TCP 443)	ztp.gc1.mist.com (TCP 443)	ztp.eu.mist.com (TCP 443)
	oc-term.mistsys.net (TCP 2200)	oc-term.gc1.mist.com (TCP 2200)	oc-term.eu.mist.com (TCP 2200)

You can check the connections using the following command:

```
user@host> show system connections | grep 2200
tcp4       0      0      10.0.0.51.49981   54.83.93.93.2200    ESTABLISHED
```

7. Check the system time on the firewall to make sure the time is correct.

```
user@host> show system uptime
Current time: 2021-08-23 19:39:17 UTC
Time Source: LOCAL CLOCK
System booted: 2021-07-14 22:40:20 UTC (5w4d 20:58 ago)
Protocols started: 2021-07-14 22:45:39 UTC (5w4d 20:53 ago)
Last configured: 2021-08-23 19:34:05 UTC (00:05:12 ago) by root
7:39PM up 39 days, 20:59, 2 users, load averages: 0.66, 1.07, 0.92
```

If the system time is not correct, configure it. For more information, see Configure Date and Time Locally.

8. Check `device-id` to make sure it is in the format `<org_id>.<mac_addr>`, as shown below:

```
user@host# show system services outbound-ssh
traceoptions {
   file outbound-ssh.log size 64k files 5;
   flag all;
}
client mist {
   device-id abcd123445-1234-12xx-x1y2-ab1234xyz123.<mac>;
   secret "$abc123"; ## SECRET-DATA
   keep-alive {
      retry 12;
      timeout 5;
   }
   services netconf;
   oc-term-staging.mistsys.net {
      port 2200;
      retry 1000;
      timeout 60;
   }
}
```

See outbound-ssh for more information.

You can also examine the log messages by using the command `show log messages`.

9. Deactivate and then reactivate the outbound SSH, as shown below:

- To deactivate:

```
user@host# deactivate system services outbound-ssh client mist
user@host# commit
```

- To activate again:

```
user@host# activate system services outbound-ssh client mist
user@host# commit
```

10. If you are adding the SRX Series Firewall for the first time, do the following:
    - Delete the present Juniper Mist configuration from the firewall using the delete command.
    - Onboard the firewall again. For details on getting your SRX Series Firewall up and running in the Mist cloud, see Cloud-Ready SRX Firewalls .
    - Verify system service outbound-ssh and system connections using the following commands:
        - `show system services outbound-ssh`
        - `show system connections | grep 2200`

Troubleshoot SRX Series Firewalls Using Packet Captures

SRX Series Firewalls support manual packet captures (PCAP). Packet capture is a tool that helps you to analyze network traffic and troubleshoot network problems. It captures real-time data packets traveling over the network for monitoring and logging.

NOTE: SRX Series Firewalls do not support dynamic packet capture.

Manual packet captures are initiated by users from the WAN Edge Packet capture page.

To initiate manual PCAP for an SRX Series Firewall:

1. Go to **Site** > **WAN Edge Packet Captures**.
2. On the **WAN** tab, click **Add WAN Edge +** and select an SRX Series Firewall.

Figure 173: WAN Edge Packet Capture

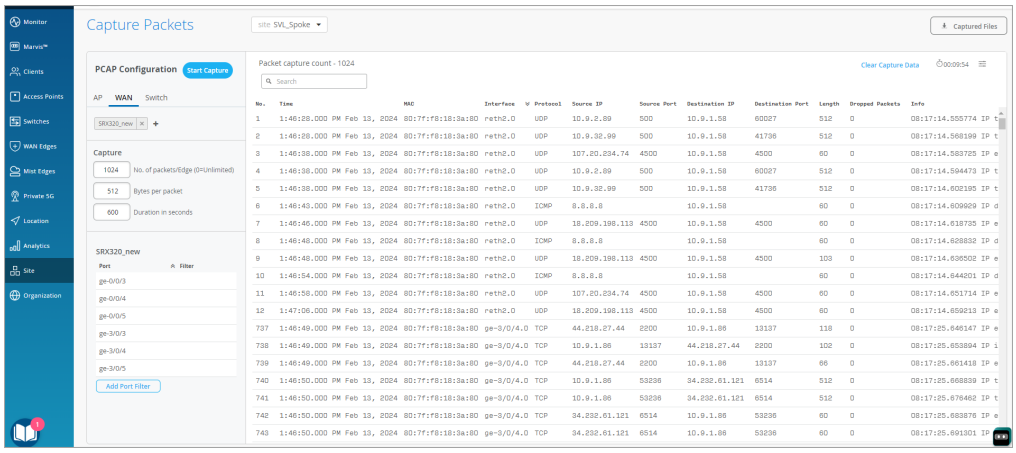

3. Specify the number of packets captured, packet size in bytes, and the duration of the capture session.

4. Use the **Add Port Filter** option to specify the port. In this pane, you can also enter filters in the TCPDUMP Expression text box.

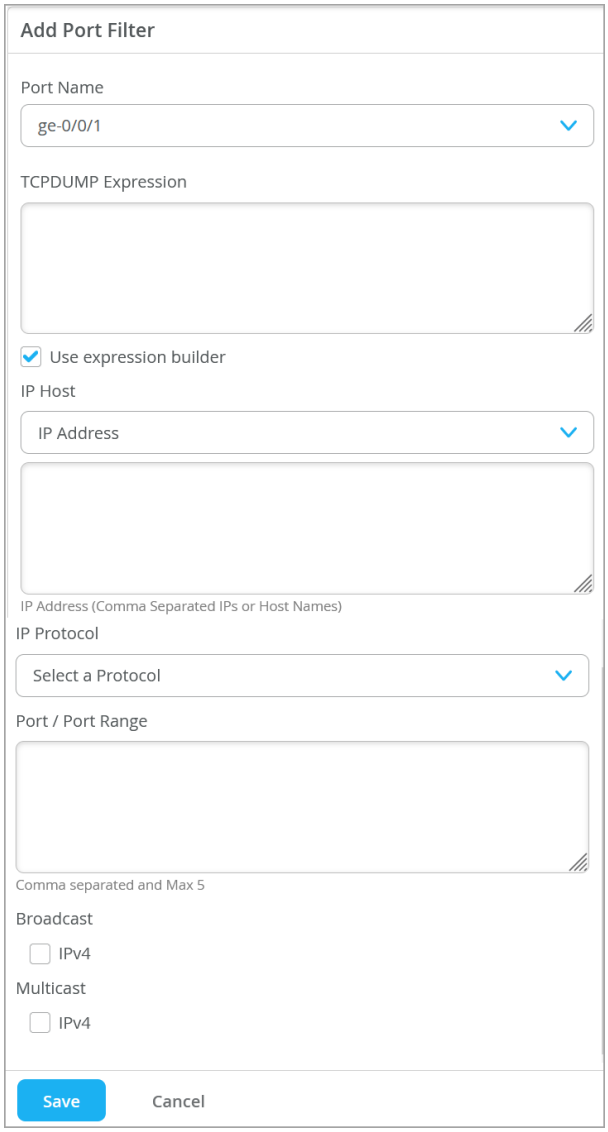

5. Optionally, select **Use Expression builder** to build the expression for packet capture. Expression builder is an interactive GUI tool to build custom filters in tcpdump syntax for use in the capture session. You can let the builder start the filter entry and then add to or delete from the entry manually. You can specify the following options:
   - IP host
   - Protocols
   - Port and port ranges
   - IP broadcast
   - IP multicast

When you enter addresses and protocols into the expression builder, the portal automatically generates the tcpdump expression on the page. You can edit the expression if needed.

6. Click **Start Capture**. The packet capture content is streamed on the page.
7. You can download the file for offline analysis by clicking **Captured Files** on the top right side of the page.

See also: "Dynamic and Manual Packet Captures" on page 397.

SEE ALSO

Juniper Mist WAN Assurance Platform Considerations | 23

No Link Title

Juniper Mist WAN Assurance Configuration Hierarchy | 6

Using the Root Cause Analysis to Troubleshoot Application Health

SUMMARY

Application Health is an important SLE to use when troubleshooting issues at the WAN edge. Use the Root Cause Analysis page to gain insights into lower than expected service levels.

Remember, you need to define application probes on SRX Series devices, but not on Session Smart Routers. However, for both, you will need traffic that can be sampled and reported to the Juniper Mist cloud.

Mist AI begins collecting and analyzing data right away, but you'll see more meaningful data after at least one week.

When you click Application Health on the WAN SLE dashboard, you'll see the Root Cause Analysis. By default, it shows the Statistics page, which displays the generalized Success Rate. In this example, it's only 79 percent.

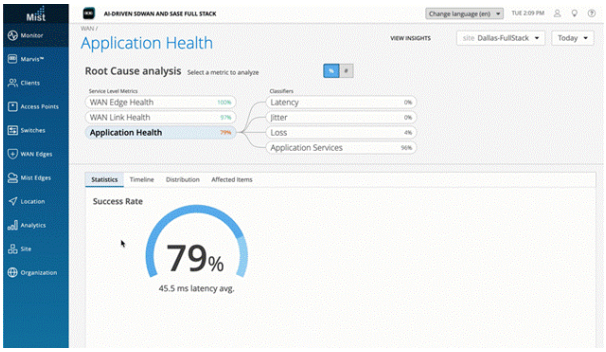

What's going on to generate such a low rate? Explore the Root Cause Analysis for the various classifiers (latency, jitter, loss, and application failures) that contribute to the overall percentage.

After clicking a classifier in the top half of the page, dig into the data in the lower half of the page.

Use these tabs:

- **Statistics**—Displays the generalized percentage of all the classifiers.
- **Timeline**—Shows a timeline of the events comprising the classifiers, listing their failures, connected clients, and system changes. Use this information to pinpoint the events that need further investigation.

  The time frame is influenced by the value at the top of the page, with options for **Today**, **Yesterday**, **This Week**, or a **Custom Range**. On the timeline, you can select a specific time, zoom in for details, or select a range. Hover over any point in time to see a pop-up message with more information about the classifier that failed to meet service expectations.
- **Distribution**—Analyzes service-level failures by attribute and is sorted by the most disruptive attribute. You can drill down into categories of traffic classes, peer paths (those Session Smart connections between Session Smart Routers), physical interfaces, WAN edges, and zones. Note that WAN Assurance deployments with Session Smart WAN edge devices leverage the **peer path** information, while the SRX Series devices leverage destination **zones** for deeper insights in the **Distribution** tab.
- **Affected Items**—Categorizes the specific items that failed to meet service-level goals. Here, you'll find numerical values for the various failures.

Replace a WAN Edge Device

SUMMARY

Follow these steps to replace a WAN Edge.

You can replace an existing WAN Edge with a new one of the same model and Junos or SSR software version. The existing WAN Edge should already be claimed/adopted to the site, and the new WAN Edge should exist in your organization's inventory so it can be adopted into the site.

The configuration from the existing WAN edge will be cloned to the replacement device, and the replaced device will become unassigned. Because the replacement process may interrupt the network for a few minutes, even if the node is in a cluster, we recommend that you schedule a maintenance window for the operation.

Figure 174: Replace a WAN Edge

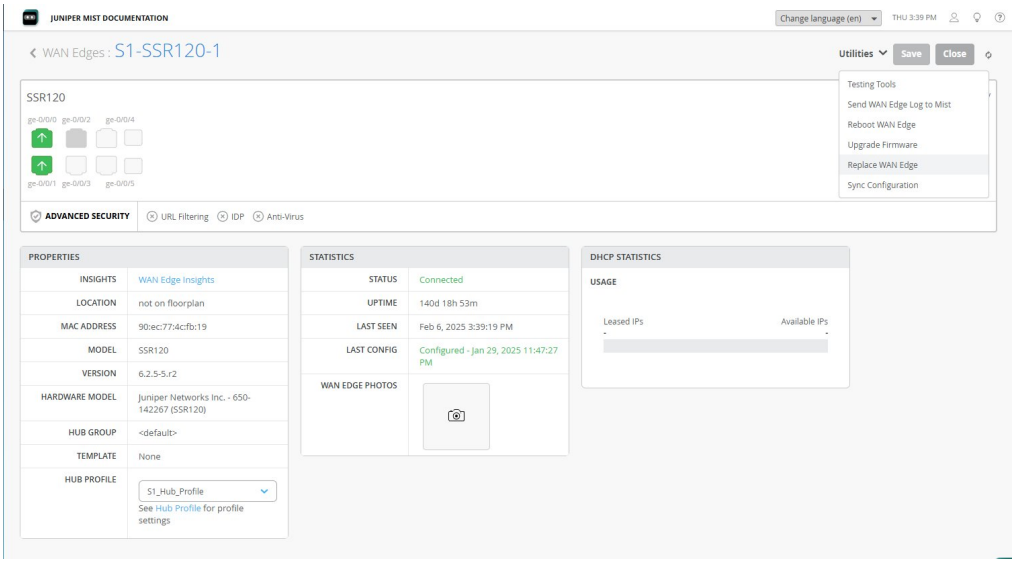

To replace a WAN Edge:

1. From the Mist menu, select **WAN Edges** > **WAN Edges** and, from the list that appears, click the name of the WAN Edge you want to replace. The WAN Edge details page opens.
2. Click the **Utilities** button, and then choose **Replace WAN Edge** from the drop-down list.

3. Choose the new WAN Edge you want to use and click the **Replace** button.
4. In the **Replace WAN Edge** window, select the MAC address of the unassigned WAN Edge (the old device) from the **MAC address of unassigned device** drop-down list.
5. Click **Replace**. The configuration of the old WAN edge will be copied to the new device, and the old device will become **unassigned** in the Inventory.