Juniper Networks | Engineering Simplicity

Contrail Service Orchestration

Deployment Guide

Published
2022-01-27

RELEASE
6.2.0

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

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

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

Contrail Service Orchestration Deployment Guide
6.2.0
Copyright © 2022 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About This Guide | vii

1 Introduction

About this Deployment Guide | 2

Contrail Service Orchestration (CSO) Solutions Overview | 3

Understand CSO Versions (On-Premises and Software as a Service) | 7

Contrail Service Orchestration (CSO) GUI Portals | 9

Access the Contrail Services Orchestration (CSO) GUIs | 9

SD-WAN Overview | 11

Understand SD-WAN Sites and Devices | 17

CSO SD-WAN Topologies | 23

CSO Next-Generation Firewall Topology | 27

2 Before You Deploy SD-WAN or NGFW

Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall | 29

Log In to CSO Administration Portal | 30

Configure SMTP Settings in CSO | 31

Download the Signature Database | 33

Configuration Templates Workflow | 35

Device Templates Workflow | 36

Device Images Workflow | 37

Add a Point of Presence (POP) | 37

Add a Provider Hub Device | 38

Add an Operating Company (OpCo) | 48

Add a Tenant | 51

WAN Link Redundancy in Enterprise Hubs Using Aggregated Ethernet | 61

- Aggregated Ethernet Links in Enterprise Hubs | 62
- Configuration Guidelines for Aggregated Ethernet on WAN Links | 63
- Example: Configure Aggregated Ethernet in Enterprise Hub Devices | 64

Add CSO Licenses | 68

Assign CSO Licenses to Tenants | 71

3 SD-WAN Deployment

CSO SD-WAN Deployment Workflow | 75

Switch Scope or Log in as Tenant Administrator | 77

Add Provider Hub Sites | 77

Add Enterprise Hub Sites | 78

Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites | 115

Add and Install (Push) Device Licenses | 115

Install the Signature Database on Devices | 117

Add Path-Based Steering Profiles | 119

Add SLA-Based Steering Profiles | 120

Add and Deploy SD-WAN Policy Intents | 124

Add SD-WAN Breakout Profiles | 127

Add Cloud Breakout Settings | 129

Add SD-WAN Branch Sites | 133

Supported Devices for SD-WAN, and Ports and Protocols to Open | 168

Manually Activate a Site | 174

Monitor SD-WAN Sites and Devices | 175

4 Standalone Next-Generation Firewall Deployment

CSO Next-Generation Firewall (NFGW) Deployment Workflow | 178

Add Next-Generation Firewall (Branch) Sites | 180

Deploy a Firewall Policy | 188

Deploy a NAT Policy | 189

Configure Unified Threat Management (UTM) in CSO | 190

- Explanation of Procedure | **190**
- Configure UTM Settings for a Tenant | **191**
- Add UTM Profiles | **193**
- Add Web Filtering Profiles | **197**
- Add Antivirus Profiles | **203**
- Add Antispam Profiles | **208**
- Add Content Filtering Profiles | **209**
- Add URL Patterns | **213**
- Add URL Categories | **214**

Configure and Deploy SSL Proxy Policy in CSO | 215

- Explanation of Procedure | **216**
- Import a Certificate | **217**
- Install a Certificate | **219**
- Add SSL Forward Proxy Profiles | **219**
- Add SSL Proxy Policy Intents | **225**
- Deploy an SSL Proxy Policy | **227**

Configure Intrusion Prevention System (IPS) in CSO | 228

- Explanation of Procedure | **229**
- Add IPS Profiles | **230**
- Add IPS or Exempt Rules to IPS Profiles | **230**

Add and Deploy Firewall Policies | 238

Add and Deploy NAT Policies | 243

Supported Devices for NGFW, and Ports and Protocols to Open | 255

Monitor Next-Generation Firewall Sites and Devices | 258

5 Appendix

Designing and Publishing Network Services | 262

Use Site Templates to Add SD-WAN and NGFW Spoke Sites | 263

Add Branch or Enterprise Hub Sites Without Provisioning a Service | 264

Understand Breakout in CSO | 273

Network Function Virtualization in the Contrail Service Orchestration Deployments | 274

VNFs Supported by Contrail Service Orchestration | 277

Install Junos OS Software onto an NFX Series Device from a USB Drive | 277

About This Guide

Use this guide to understand the next steps you should take after a successful installation of CSO software (either CSO on-premises or CSO SaaS). This guide describes the solutions available in CSO and the workflows involved in their deployment.

CHAPTER 1

Introduction

About this Deployment Guide | 2

Contrail Service Orchestration (CSO) Solutions Overview | 3

Understand CSO Versions (On-Premises and Software as a Service) | 7

Contrail Service Orchestration (CSO) GUI Portals | 9

Access the Contrail Services Orchestration (CSO) GUIs | 9

SD-WAN Overview | 11

Understand SD-WAN Sites and Devices | 17

CSO SD-WAN Topologies | 23

CSO Next-Generation Firewall Topology | 27

About this Deployment Guide

The intent of this deployment guide is to provide an understanding of the available Contrail Service Orchestration (CSO) solutions by:

- Briefly discussing each of the available solutions
- Discussing the tools and devices required to implement the solutions
- Providing an end-to-end walkthrough of each of the solutions and covering the deployment specifics

This guide is available on the CSO Documentation page, which contains links to the different CSO releases and the documentation published for those releases. Table 1 on page 2 lists some of the guides that are available on the CSO Documentation page.

Table 1: Additional CSO Documentation

Documentation	Available for	Location
Administration Portal User Guide Customer Portal User Guide	CSO On-Premises CSO SaaS	**How To > User Guides** section of the CSO Documentation page
Release Notes	CSO On-Premises CSO SaaS	**Learn > Release Notes** section of the CSO Documentation page
CSO SD-WAN Design and Architecture Guide	CSO On-Premises CSO SaaS	**Set Up > Get Started** section of the CSO Documentation page
Monitoring and Troubleshooting Guide	CSO On-Premises CSO SaaS	**How To > System Admin Guides** section of the CSO Documentation page
CSO Installation and Upgrade Guide	CSO On-Premises	**Set Up > Install/Upgrade Software** section of the CSO Documentation page,
Designer Tools User Guide	CSO On-Premises	**How To > User Guides** section of the CSO Documentation page

Contrail Service Orchestration (CSO) Solutions Overview

IN THIS SECTION

- Contrail Software-Defined WAN Solution (SD-WAN) | 4
- Next Generation Firewall (NGFW) Deployment Model | 6

Juniper Networks CSO SD-WAN and NGFW management solutions offer automated branch connectivity while improving network service delivery and agility. CSO is a multi-tenant platform that manages physical and virtual network devices, creates and manages Juniper Networks and third-party virtualized network functions (VNFs), and uses those elements to deploy network solutions for both enterprises and service providers (SPs) and their customers. CSO multi-tenancy provides security and tenant isolation that keeps the objects and users belonging to one tenant or operating company (OpCo) from seeing or interacting with those of another tenant or OpCo.

CSO is available as an on-premises version (CSO on-premises) or a Software as a Service (CSO SaaS). For more information, see "Understand CSO Versions (On-Premises and Software as a Service)" on page 7.

CSO offers multiple network solutions that benefit enterprise customers and service providers and their customers:

- Provide lifecycle management for devices and services
- Automate physical and virtual device provisioning
- Provide Day 0, Day 1, and Day 2 configuration
- Monitor remote devices
- Provide full lifecycle management of firewall, NAT, and Internet breakout policies for user traffic
- Provide high-level reporting about devices and user traffic

Contrail Software-Defined WAN Solution (SD-WAN)

The Contrail SD-WAN solution offers a flexible and automated way to route traffic through the cloud using overlay networks. It is an overlay network solution that provides an enhanced application user experience. It acts as both a data controller and a management orchestrator. At its most basic, an SD-WAN solution encompasses multiple sites, multiple connections between sites, and a WAN controller as shown in Figure 1 on page 4.

Figure 1: Basic SD-WAN Concept

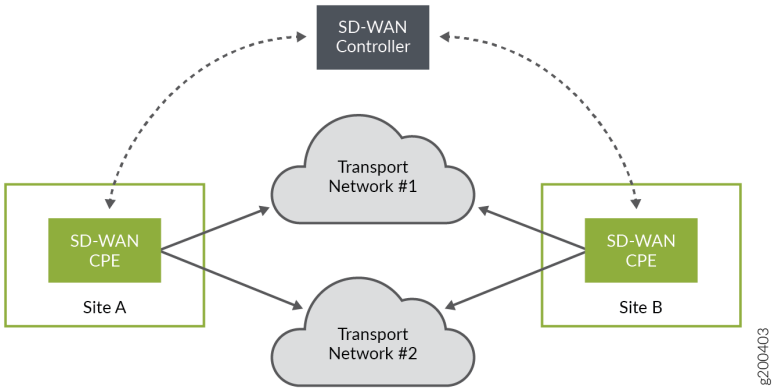

The CPE devices in a Contrail SD-WAN solution (also known as *on-premises spoke devices*) have a WAN side and a LAN side. On the WAN side, hub-and-spoke and dynamic mesh topologies are supported. The CPE devices use at least one, and up to four, WAN interfaces as connection paths to provider hub devices, enterprise hub devices, other spoke devices, and the Internet. The supported hub devices are shown in Table 2 on page 4:

Table 2: Supported Hub Devices

Hub Device	Used as
vSRX	Enterprise Hub and Provider Hub
SRX380	Enterprise Hub
SRX1500	Enterprise Hub and Provider Hub
SRX4100	Enterprise Hub and Provider Hub

Table 2: Supported Hub Devices *(Continued)*

Hub Device	Used as
SRX4200	Enterprise Hub and Provider Hub
SRX4600	Enterprise Hub and Provider Hub

The hub devices help to provide the overlay networking needed for the Contrail SD-WAN solution.

CSO allows you to give preference to one WAN path over another for any given traffic through the use of traffic steering and breakout profiles. Thus, business-critical traffic and data can be routed through the provider hub using MPLS/GRE while non-critical traffic can be routed over the Internet connection through an IPsec tunnel. Each path can have a service level agreement (SLA) profile applied. The SLA profile monitors the path for latency, congestion, and jitter while also accounting for path preference. Should the path fail to meet one or more of the required parameters, traffic is re-routed to another path automatically.

The LAN side of the CPE devices connect to the customer's LAN segments. Multiple departments at the customer site that occupy different LAN segments can have their traffic securely segregated. NFX Series spoke devices can also provide service chains of network services in addition to the routing flexibility already available.

You can use the solutions as turnkey implementations or connect to other operational support and business support systems (OSS/BSS) through northbound Representational State Transfer (REST) APIs.

Next Generation Firewall (NGFW) Deployment Model

The NGFW deployment focuses on providing remote network security through the use of SRX Series NGFW devices as CPE at the spoke site; unlike the SD-WAN deployments which focus on secure site-to-site connectivity. A high-level view of the spoke site with NGFW is shown in Figure 2 on page 6.

Figure 2: NGFW Spoke Site

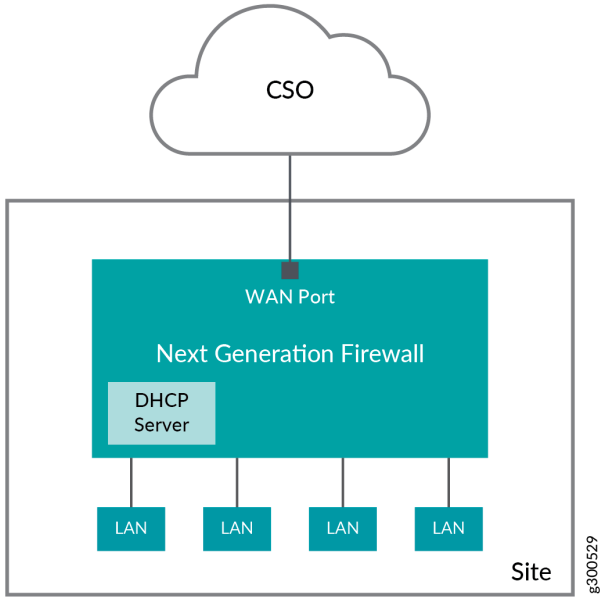

An NGFW deployment is carried out in the Customer Portal of CSO as a site deployment. The tenant under which the site is deployed must have the NGFW service available. This service is included in the tenant configuration by the tenant administrator during tenant onboarding.

RELATED DOCUMENTATION

SD-WAN Overview | **11**

CSO SD-WAN Topologies | **23**

CSO Next-Generation Firewall Topology | **27**

Understand CSO Versions (On-Premises and Software as a Service)

Juniper Networks offers Contrail Service Orchestration (CSO) software in two versions: on-premises and as a Software as a Service (SaaS). At a high-level, the difference between the two versions is as follows:

- **CSO on-premises**: The on-premises version of CSO enables you (the customer) to install CSO on your own hardware infrastructure. Therefore, you are responsible for the maintenance and administration of CSO and the underlying hardware infrastructure. In addition, you handle the upgrades to the CSO software versions, if applicable. Users access the CSO GUI by using the URL provided by your CSO administrator. An example of a CSO on-premises version is CSO Release 5.1.2.
- **CSO SaaS**: In the SaaS version of CSO, Juniper Networks handles the installation, upgrade, and maintenance and administration of CSO. Users access the CSO GUI by using the URL provided by Juniper. An example of a CSO SaaS version is CSO Release 5.4.0.

Table 3 on page 7 displays the differences between the on-premises and SaaS versions of CSO with respect to installation and upgrade tasks or requirements.

Table 4 on page 8 displays the differences in features between the on-premises and SaaS versions.

Table 3: Installation and Upgrade Differences Between CSO Versions

Task or Requirement	On-Premises	SaaS
Installed by	Installed by service provider or customer	Installed by Juniper Networks
Upgraded by	Upgraded by service provider or customer	Upgraded by Juniper Networks
Hardware Required	A minimum of three servers is needed	Not needed because CSO is hosted on public cloud infrastructure and managed by Juniper Networks
Operation, Administration, and Maintenance (OAM) Hubs	OAM hubs must be added to the CSO infrastructure by the service provider or customer	Not needed because Juniper adds the OAM hubs

Table 3: Installation and Upgrade Differences Between CSO Versions *(Continued)*

Task or Requirement	On-Premises	SaaS
Integration with local operations support system (OSS) or business support systems (BSS)	Can be done using CSO's REST APIs. Juniper Professional Services (PS) can assist with integration with third-party systems such as Service Now.	
Instance type	Private data center/ Dedicated AWS	Shared on public AWS

Table 4: Feature Differences Between CSO Versions

Feature	On-Premises	SaaS
Multitenancy	Available at three levels: • Service provider • Operating Company (OpCo) • Tenant	Available at two levels: • OpCo • Tenant
Juniper Identity Management System (JIMS)	Supported	Not qualified by Juniper
Customization of Administration and Customer Portals	Display preferences, e-mail templates, and terms of use	E-mail templates and terms of use
Designer Tools	Supported	Not supported
Download of Signature Database	Supported at the service provider level	Periodically downloaded by Juniper Networks

Contrail Service Orchestration (CSO) GUI Portals

The following sections provide a high-level overview of the CSO GUI portals:

- Administration Portal—The Administration Portal in Contrail Service Orchestration (CSO) provides a Web-based UI that service providers (SPs) and operating companies (OpCos) can use to manage physical and virtual resources, add and manage tenants, monitor system performance, perform administrative tasks (such as manage users and roles), and so on. For more information, see *Administration Portal Overview* in the *CSO Administration Portal User Guide* available on the CSO Documentation page).

- Customer Portal—The Customer Portal in CSO provides a Web-based UI that tenants of service providers (SPs) and operating companies (OpCos) can use to manage sites and devices, apply policies (firewall, NAT, SD-WAN, and so on), and perform administrative tasks. In Customer Portal, the objects that you manage and the actions that you perform are done in the context of a single tenant. Therefore, objects belonging to one tenant are isolated from objects belonging to other tenants. For more information, see *Customer Portal Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

- Designer Tools—Designer Tools are a suite of GUI-based tools that enable network designers to support virtualized network functions (VNFs) and network services in CSO. Designer Tools comprise Configuration Designer, Resource Designer, and Network Services Designer. For more information, see "Designing and Publishing Network Services" on page 262.

  NOTE: Designer Tools is available only for the CSO on-premises version and can be accessed only by users with the service provider (SP) Administrator role.

Access the Contrail Services Orchestration (CSO) GUIs

You access the CSO GUI portals using a web browser. Table 5 on page 10 provides access and login details for the different CSO GUI portals.

NOTE: We recommend that you use Google Chrome (Version 60 or later) or Firefox (Version 78 or later) to access the Contrail Service Orchestration (CSO) GUIs.

Table 5: Access Details for the GUIs

GUI	URL	Login Credentials
Administration Portal	• For CSO SaaS: Login credentials are sent to each Administration Portal user as an e-mail. The address to which the e-mail is sent is the *username* and the e-mail contains a link including an activation code. Clicking the link takes you to the CSO login page which then prompts you to create a password. Once the new password is set, the CSO login URL can be seen in your browser's address bar. • For on-premises CSO: https://central-IP-Address Where: *central-IP-Address* is the IP address of the VM that hosts the microservices for the central POP For example: https://192.0.2.1	• For CSO on-premises: Specify the OpenStack Keystone username and password. The default username is **cspadmin**. Specify the autogenerated cspadmin password that is displayed on the console after the installation is complete. After upgrade, you must specify the cspadmin password of the previously installed version.

Table 5: Access Details for the GUIs *(Continued)*

GUI	URL	Login Credentials
Customer Portal	Same as the URL used to access the Administration Portal	Login credentials are sent to each Customer Portal user as an e-mail. The address to which the e-mail is sent is the *username* and the e-mail contains a link including an activation code. Clicking the link takes you to the CSO login page which then prompts you to create a password.
Designer Tools—Log into Network Service Designer and click the menu in the top left of the page to access the other designer tools. **NOTE**: Access to Designer Tools is only available for on-premises deployments of CSO.	https:// *central-IP-Address*:83 Where: *central-IP-Address* is the IP address of the VM that hosts the microservices for the central POP For example: https://192.0.2.1:83	Specify the OpenStack Keystone username and password. The default username is **cspadmin**. Specify the autogenerated cspadmin password that is displayed on the console after the installation is complete. After the upgrade, you must specify the cspadmin password of the previously installed version.

SD-WAN Overview

IN THIS SECTION

- Branch Management Without and With SD-WAN | **13**
- SD-WAN Overlay Tunnels | **15**
- High-Level SD-WAN Architecture | **15**
- Additional Information | **17**

In simple terms, software-defined WAN (SD-WAN) refers to the application of software-defined networking (SDN) principles to the WAN. In SD-WAN, the path for the application traffic can be dynamically controlled and selected based on specified service-level agreement (SLA) criteria. Thus, SD-WAN enables you to identify the best path for an application's traffic and to then forward the traffic on that path.

According to Gartner, SD-WAN has the following characteristics:

- Support for multiple WAN connection types (such as MPLS, Internet, LTE) simultaneously.
- Ability to select the traffic path dynamically, which allows for load sharing of traffic across WAN connections.
- Ability to simplify the management and monitoring of WANs.
- Support for VPNs and other third-party services, such as gateways and firewalls.

CSO supports the following SD-WAN service types for a site:

- *Secure SD-WAN Essentials*—Provides the basic SD-WAN services. This service is ideal for small enterprises, looking for simplified management of their network and comprehensive NGFW security services at the branch sites. The SD-WAN Essentials service allows Internet traffic to breakout locally, and thus avoids the need to backhaul web traffic over costly VPN or MPLS links. This service supports features such as intent-based firewall policies, WAN link management and control, CSO-controlled routing between sites connected through the static VPN, and site to site communication through MPLS or internet links. A tenant with the SD-WAN Essentials service level can create only SD-WAN Essentials sites.

  NOTE: You can upgrade the SD-WAN service level of a tenant from SD-WAN Essentials to SD-WAN Advanced by editing the tenant information from the CSO Administration portal, provided that you have purchased the corresponding license.

- *Secure SD-WAN Advanced*—Provides the complete SD-WAN service. All sites of the tenant with Secure SD-WAN Advanced service are connected in full mesh or hub-and-spoke topology. The SD-WAN Advanced service includes SD-WAN Essentials.

  NOTE: SD-WAN sites on CSO Release 5.4 or earlier versions are treated as SD-WAN Advanced sites. You cannot downgrade the SD-WAN service level of a tenant from SD-WAN Advanced to SD-WAN Essentials.

Branch Management Without and With SD-WAN

Figure 3 on page 13 displays a topology in which a branch is managed without SD-WAN. In this scenario, the service provider (SP) maintains the quality-of-service-enabled (QoS-enabled) network and the branch, and manages the traffic (including route announcements), and VPN. In Figure 3 on page 13, the area bounded by the shaded rectangles indicates the what the service provider manages and maintains.

Figure 3: Branch Management Without SD-WAN

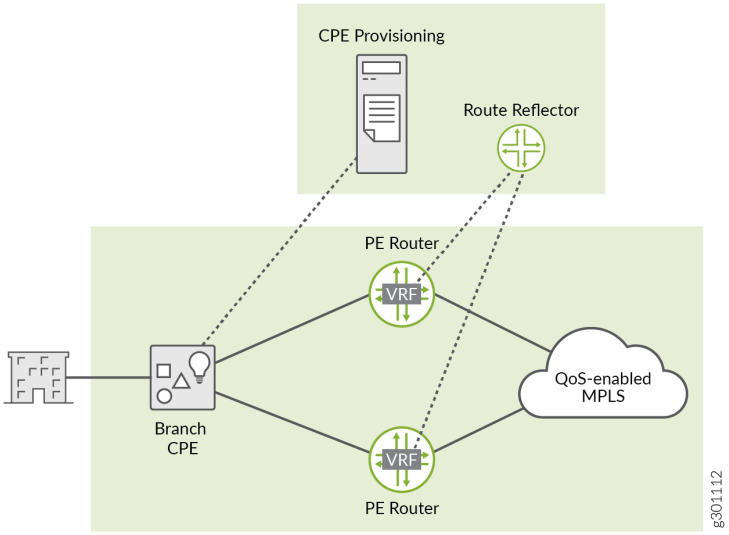

The branch customer sends traffic, which is directed over one of the two redundant links to one of the two provider edge (PE) routers, where the traffic is forwarded inside the virtual routing and forwarding (VRF) instance. Typically, the PE routers are configured in an active-backup mode (for redundancy), where traffic flows only through one router at any given time. The PE router builds queues for the traffic and the queues are respected inside the QoS-enabled MPLS network meant for that branch customer. Additionally, bandwidth might be reserved for applications that need a guaranteed bandwidth. Optionally, the service provider can provide additional value-added services, where the traffic is marked using differentiated services code point (DSCP) values and the DSCP values are adhered to downstream in the network.

Figure 4 on page 14 displays the topology for managing a branch with SD-WAN. In this scenario, the service provider provides the PE router and the MPLS network and *can* be the provider for the Internet network. However, the enterprise has an option to add a different network (for example, broadband

Internet) instead of using the service provider's network, and the enterprise can manage the customer premises equipment (CPE) device.

Figure 4: Branch Management With SD-WAN

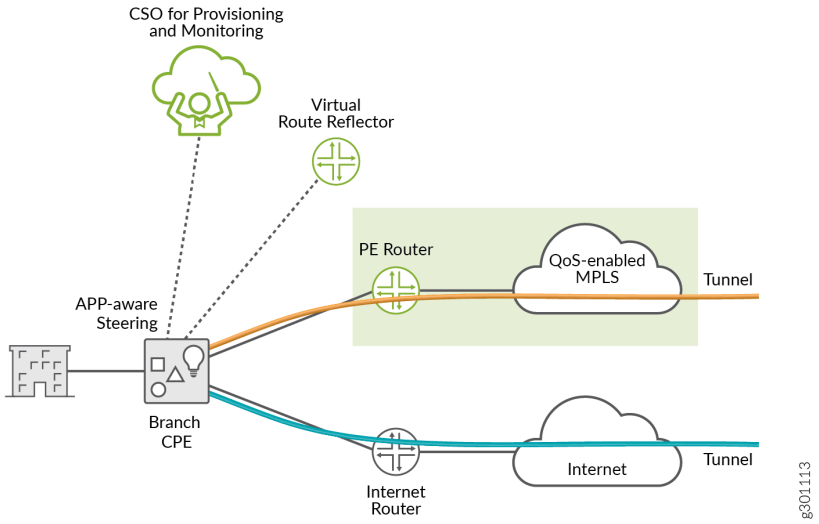

To build a VPN, the traffic must be tunneled through the different networks. So, instead of sending traffic through the underlay, you use the underlay to build tunnels through the networks to the next element (node). To dynamically select the traffic path, you need to have application-aware (also called app-aware) traffic steering that identifies the application, monitors the tunnel (path) that the traffic is on, and decides the tunnel on which to send the traffic. If a tunnel degrades, the SD-WAN controller can move the traffic to a different tunnel. In the SD-WAN scenario, both the tunnels are active simultaneously.

Therefore, in the SD-WAN scenario, you don't squeeze traffic into queues; instead, you identify the traffic and select the tunnel on which to send the traffic. Services provided throughout the network (such as route reflection) can be moved to the top as shown in Figure 4 on page 14.

NOTE: In branch management with SD-WAN, you can have redundant PE routers in the topology, if needed. (This is not shown in Figure 4 on page 14.)

SD-WAN Overlay Tunnels

In SD-WAN, the overlay tunnels (see Figure 5 on page 15) are transport-agnostic, which means that they are built independent of the underlying transport technology (such as MPLS or Internet) of the network. Tunnels are built based on the IP addresses assigned to the WAN interfaces, and can be between one spoke (branch) and another, or between a spoke and a hub (headquarters).

In CSO, you can build GRE tunnels or GRE tunnels with IPsec for additional security. When CSO identifies the application, it creates inner DSCP markings that are written to the outside tunnel so that the forwarding queues that might exist in the outside network are respected.

NOTE: In CSO, the term MPLS refers to a QoS-engineered path and is used to designate the network. Therefore, CSO doesn't create MPLS frames on the underlay and only creates Ethernet frames.

Figure 5: SD-WAN Overlay Tunnels (Transport-Agnostic)

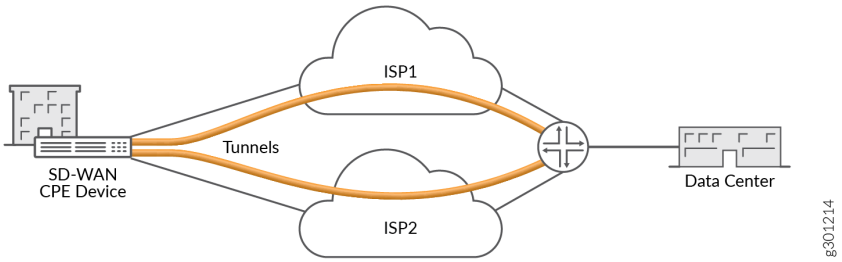

High-Level SD-WAN Architecture

Figure 6 on page 16 shows an example of a high-level SD-WAN architecture. There are two branch sites connected to SD-WAN gateways (also known as spokes or CPE devices) and one central site (headquarters) connected to another SD-WAN gateway, which could be a hub device. In addition, an

SD-WAN controller controls the SD-WAN gateways using a single UI, manages the devices, the creation of tunnels, and so on.

Figure 6: Example of SD-WAN Architecture

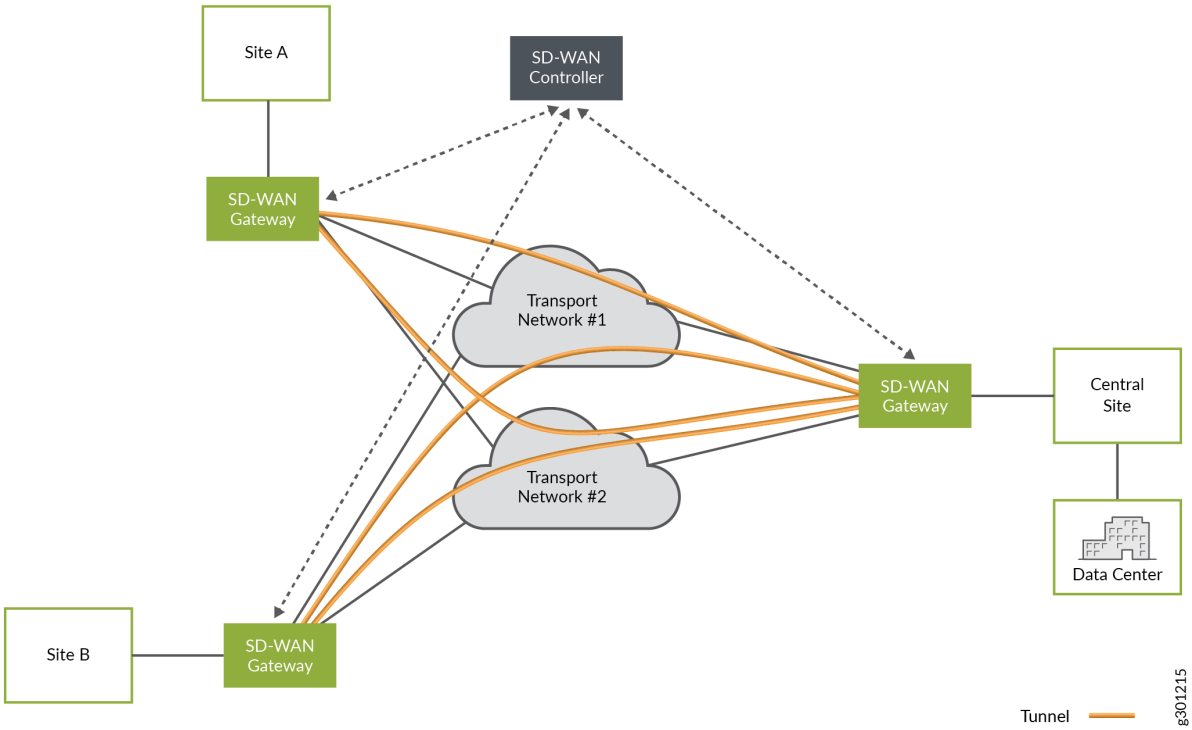

Figure 7 on page 17 shows how SD-WAN is applied using CSO in a topology that has one branch site and one hub site. CSO builds one tunnel for the WAN links going over the MPLS network and a second tunnel for the WAN links going over the Internet. When you configure SD-WAN, you can ensure that

mission-critical application data is sent over the MPLS link (reliable and secure path) and the non-mission critical application data is sent over the Internet link (best-effort, non-secure path).

Figure 7: Example of a CSO SD-WAN Topology

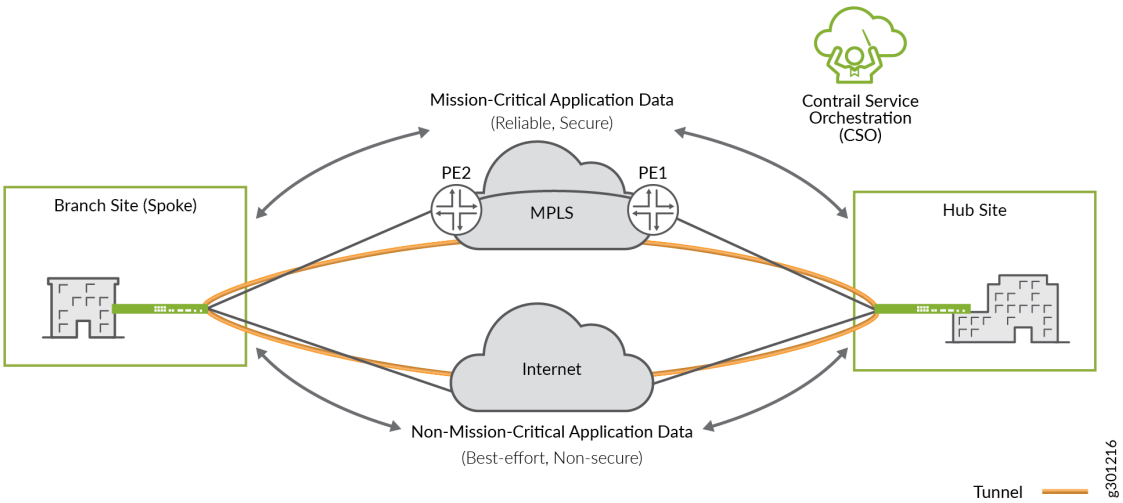

Additional Information

For more information about CSO SD-WAN, watch the Contrail SD-WAN Demos—15 Features in 15 Minutes video.

Understand SD-WAN Sites and Devices

IN THIS SECTION

- Spoke Devices | **18**
- Provider Hub Devices | **21**
- Enterprise Hub Sites and Devices | **22**

In Contrail Service Orchestration (CSO), there are two categories of SD-WAN devices: spoke devices and hub devices. These are explained in the sections below.

Spoke Devices

The CPE device at an enterprise customer's branch site acts as a spoke device in the SD-WAN model. The device also acts as a gateway router, providing connectivity from the branch site to other sites in the tenant network and to the Internet. There are two types of spoke devices: on-premises spoke and cloud spoke.

On-Premises Spoke Devices

On–premises spoke devices can be either NFX Series devices or specific SRX Series devices.

NFX Series Network Services Platform

The NFX Series Network Services Platform used as an on-premises spoke device differentiates from traditional CPE devices in that it can host a range of multivendor VNFs and support service chaining, managed by orchestration software in the cloud. NFX Series devices eliminate the operational complexities of deploying multiple physical network devices at a customer site.

A key VNF supported on the NFX Series platform is the vSRX Virtual Firewall. In the CSO SD–WAN solution, the vSRX instance performs the gateway router function, given its routing and switching capabilities. It also provides the same feature-rich security services found on a standard SRX series devices. Table 6 on page 19 shows the supported NFX hardware models.

NOTE: The NFX150 features a built–in SRX firewall in place of the vSRX functionality found on other NFX Series devices.

Table 6: NFX Series for On-Premises Spoke Devices

Platform	Models Supported
NFX150 Network Services Platform	NFX150–S1 NFX150–S1E NFX150–C–S1 NFX150–C–S1–AE/AA NFX150–C–S1E–AE/AA
NFX250 Network Services Platform	NFX250–LS1 NFX250–S1 NFX250–S2

SRX Series Devices and vSRX Virtual Firewalls

A physical SRX device can be used in place of the NFX platform to provide the gateway router function, as can a vSRX instance installed on a server. Table 7 on page 20 shows the supported SRX hardware and vSRX virtual firewalls

Table 7: SRX Series for On-Premises Spoke Devices

Platform	Models Supported
SRX Series	SRX4600 (SD-WAN spoke) SRX4200 SRX4100 SRX550M SRX380 SRX345 SRX340 SRX320 SRX300
	SRX1500
vSRX Virtual Firewalls	vSRX (standalone)
	vSRX (installed in NFX250)
	vSRX 3.0 (standalone)

Cloud Spoke Devices

A CSO SD–WAN cloud spoke device, in the form of a vSRX, can be located in an AWS virtual private cloud (VPC). The vSRX serves as a spoke device in the cloud; once the endpoint comes online, it acts like any other spoke device.

Spoke Redundancy

Two redundant CPE devices can be used at spoke sites to protect against device and link failures. For more detail, see the *Resiliency and High Availability* section of the CSO Design and Architecture Guide.

Provider Hub Devices

The CSO SD–WAN solution supports two deployment topologies: dynamic mesh and hub-and-spoke. In a dynamic mesh deployment, each site has a CPE device that connects to the other sites and the enterprise hub device. In a hub-and-spoke deployment, there is at least one provider hub device and one or more spoke devices.

The provider hub device terminates both MPLS/GRE and IPsec tunnels from spoke devices.

Provider Hubs

In a service provider (SP) environment, the service provider hosts a *provider hub* device in their network. The provider hub device acts as a point of presence (POP) or connection point. It is typically a shared device, providing hub functionality to multiple customers (tenants) through the use of virtual routing and forwarding instances (VRF). The SP administrator and the OpCo administrator can both manage the provider hub device.
In CSO SaaS, the SP administrator role is performed by Juniper Networks as the cspadmin user (or equivalent). The OpCo administrator role can be assigned to a user by the SP administrator, but the OpCo administrator does not have SP administrator privileges.
Table 8 on page 21 lists the provider hub devices supported in a CSO SD-WAN environment.

Table 8: Provider Hub Devices

Role	Supported Device Types
Provider Hub	SRX4600 SRX4200 SRX4100 SRX1500
	vSRX
	vSRX 3.0

Provider Hub Redundancy

Two redundant provider hub devices can be used at one POP to protect against device and link failures, and to provide upstream multi-homing for spoke sites. For more detail, see the Resiliency and High Availability section of the CSO SD-WAN - Design and Architecture Guide.

Enterprise Hub Sites and Devices

A special type of spoke device, called an *enterprise hub device*, can be deployed as the CPE at an on-premises site. The spoke site that functions this way, must be configured as an *enterprise hub site* during site addition. Adding an enterprise hub site opens additional functionality for the site:

- Can act as the anchor point for site–to–site communications on the customer's network.
- Can act as the central breakout node for the customer's network.
- Offers a specialized department called the *data–center department*.
- Supports dynamic LAN segments with BGP and OSPF route imports, including default routes, from the LAN–side L3 device.
- Allows for intent-based breakout profiles to create granular breakout behavior based on department, application, site, and so on.

In an enterprise environment, the enterprise hub is owned by the customer (tenant) and usually resides within an enterprise data center. Only the customer's spoke sites can connect to the enterprise hub device. OpCo administrators and tenant administrators can manage the enterprise hub. Table 9 on page 22 lists the enterprise hub devices supported in a CSO SD-WAN environment.

Table 9: Enterprise Hub Devices and Supported Software

Role	Supported Device Types
Enterprise Hub	SRX380 SRX4600 SRX4200 SRX4100
	SRX1500

Table 9: Enterprise Hub Devices and Supported Software *(Continued)*

Role	Supported Device Types
	vSRX
	vSRX 3.0

CSO SD-WAN Topologies

This topic explains two Contrail Service Orchestration (CSO) SD-WAN topologies to give you a basic understanding. You can then construct other topologies based on your network requirements.

Figure 8 on page 24 shows a simplified CSO SD-WAN topology. The SD-WAN branch site (on-premises spoke) is shown with two WAN links configured: one Internet link and one MPLS link. The WAN links are connected to a single SD-WAN hub site. In CSO, you can configure two types of hubs: enterprise hubs and provider hubs.

Provider hubs can be configured to carry only Operation, Administration, and Maintenance (OAM) traffic, data traffic, or both data and OAM traffic. On the other hand, enterprise hubs can carry only data traffic. Furthermore, provider hubs can be shared by multiple tenants, but enterprise hubs are dedicated to a single tenant.

As shown in Figure 8 on page 24, control (OAM) traffic between the sites and CSO is carried over a secure tunnel through an OAM hub. In the CSO SaaS version, the OAM functionality is provided by Juniper Networks, so the OAM hub is transparent to the CSO SaaS user. However, in the CSO on-premises version, the service provider is responsible for providing the OAM functionality.

An example of a CSO on-premises deployment is managed services provider (MSP) who deploys a minimum of two provider hubs (that are configured to carry both OAM and data traffic), which takes care of the OAM functionality. An example of a CSO SaaS deployment is an enterprise (tenant) who

wants to use only enterprise hubs. In this case, the OAM functionality is provided by the provider hubs configured and maintained by Juniper Networks.

Figure 8: Simplified CSO SD-WAN Topology

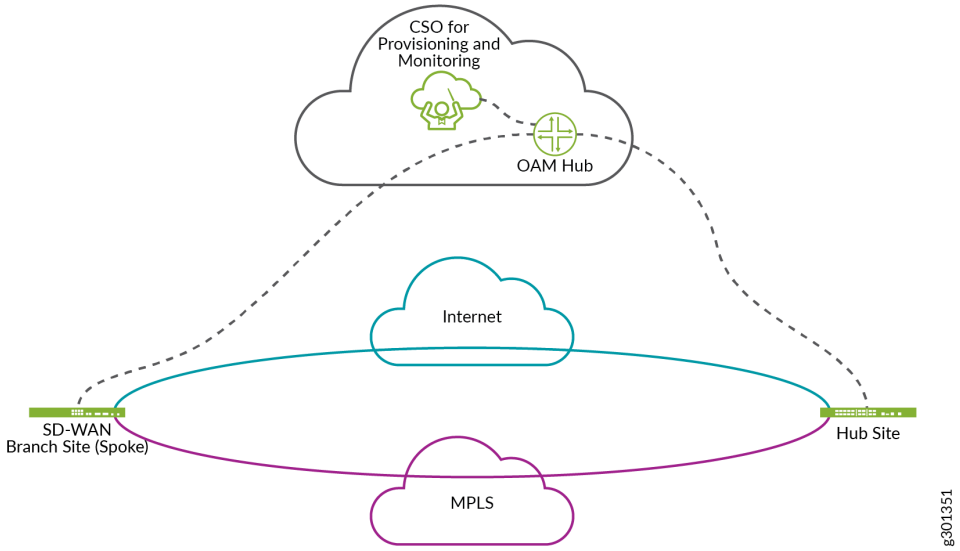

The CSO release 5.4.0 and later versions support management of OAM hub on provider hub from CSO running either on AWS via Direct Connect or in private data center with data center connectivity. In the deployment model involving AWS, CSO is connected to the provider hubs through AWS direct connect (see Figure 9 on page 25). In the deployment model involving private data center, CSO is connected to

the provider hubs within the data center (see Figure 10 on page 25). The provider hubs also support MPLS and Internet WAN links for branch site connectivity.

Figure 9: Simplified CSO SD-WAN Topology (On-Premises Deployment in which CSO is Connected to Provider Hubs via AWS Direct Connect)

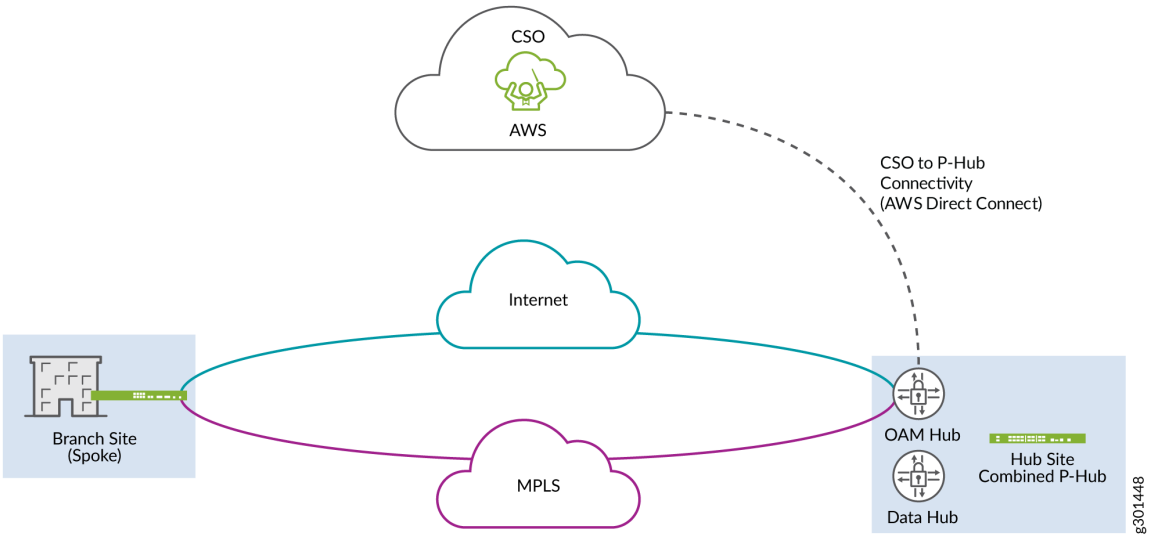

Figure 10: Simplified CSO SD-WAN Topology (On-Premises Deployment in which CSO is Connected to Provider Hubs within a Customer's Private Data Center)

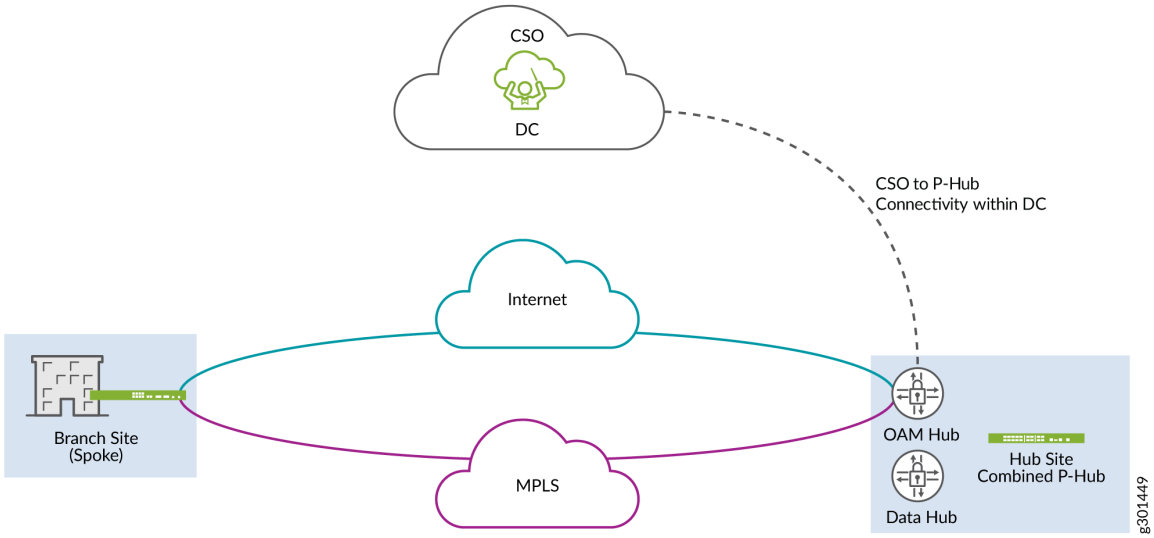

Figure 11 on page 26 shows a multihoming scenario (also called hub redundancy), where the branch site connects to two hub sites: a primary and a secondary. In this case, each WAN link has two overlay tunnels to each hub site, thereby providing redundancy.

If the primary hub site goes down, then traffic is redirected to the secondary hub site until the primary hub site comes back up. In addition, two OAM hubs are configured to provide redundancy for OAM traffic. In this case, CSO establishes secure tunnels between the:

- Branch site and the two OAM hubs.
- Two hub sites and the two OAM hubs.
- CSO and the two OAM hubs.

Therefore, if one OAM hub goes down, OAM traffic can flow through the second OAM hub, thereby providing redundancy.

Figure 11: SD-WAN Topology with Multihoming and OAM Hub Redundancy

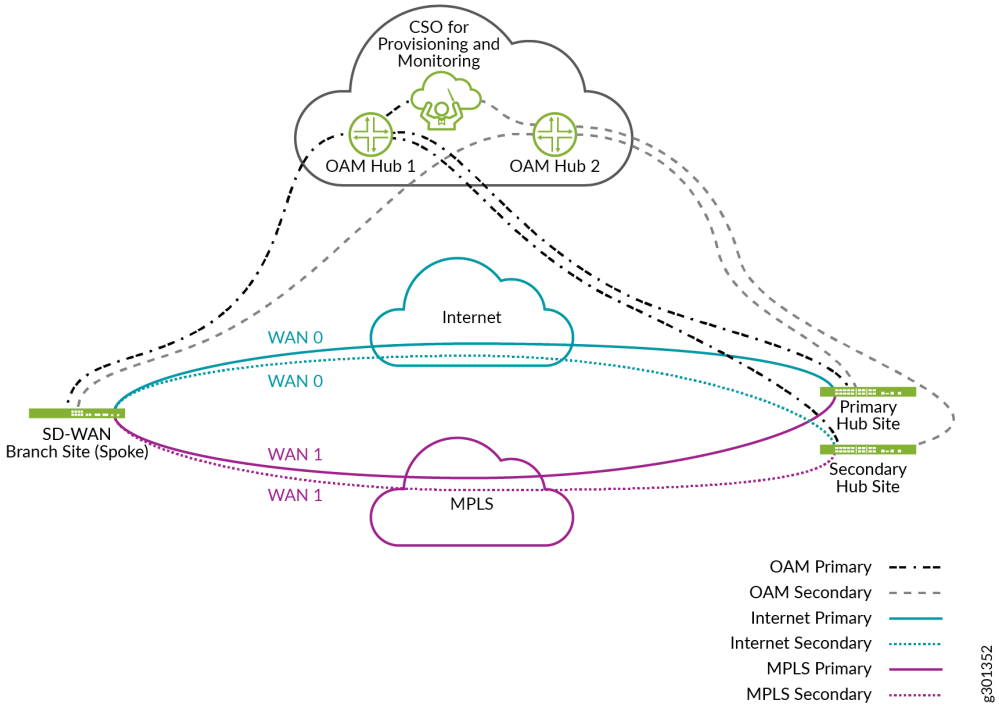

RELATED DOCUMENTATION

CSO SD-WAN Deployment Workflow | 75

CSO Next-Generation Firewall Topology

Figure 12 on page 27 shows the CSO standalone next-generation firewall (NGFW) topology. On the WAN side, the NGFW site, which is a standalone SRX Series or vSRX device, establishes the following connections:

- Data connection for Internet traffic.
- Management connection to CSO for establishing connectivity between CSO and the device, and for sending encrypted syslogs to CSO.

On the LAN side, which is not shown in the figure, the NGFW site can connect to LAN hosts. For NGFW sites, CSO allows you to provision greenfield devices or brownfield devices, with an option to import existing firewall and NAT policies into CSO for brownfield devices.

Figure 12: CSO Next-Generation Firewall Topology

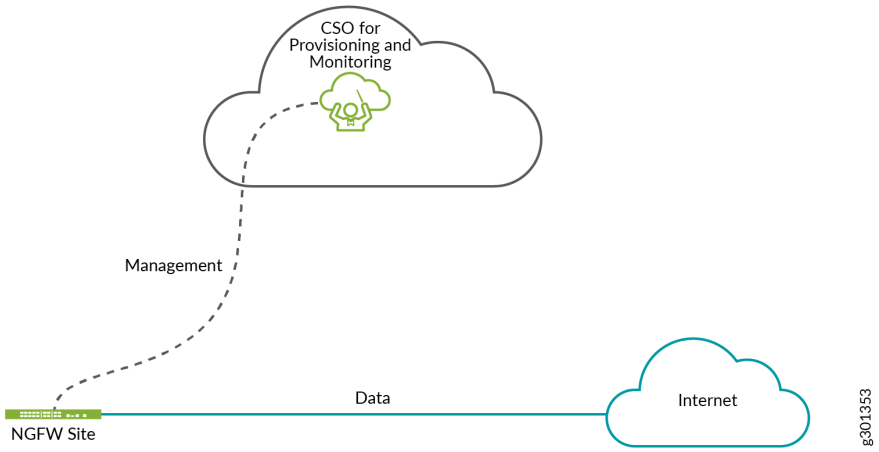

RELATED DOCUMENTATION

CSO Next-Generation Firewall (NFGW) Deployment Workflow | 178

2

CHAPTER

Before You Deploy SD-WAN or NGFW

Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall | 29

Log In to CSO Administration Portal | 30

Configure SMTP Settings in CSO | 31

Download the Signature Database | 33

Configuration Templates Workflow | 35

Device Templates Workflow | 36

Device Images Workflow | 37

Add a Point of Presence (POP) | 37

Add a Provider Hub Device | 38

Add an Operating Company (OpCo) | 48

Add a Tenant | 51

WAN Link Redundancy in Enterprise Hubs Using Aggregated Ethernet | 61

Add CSO Licenses | 68

Assign CSO Licenses to Tenants | 71

Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall

Before you deploy SD-WAN or next generation firewall (NGFW) in CSO, pre-deployment tasks must be carried out by the SP Administrator (for CSO on-premises) or the OpCo Administrator (for CSO SaaS):

1.
   - If you are an SP Administrator, go to step 2.
   - If you are an OpCo Administrator, go to step 3.
2. The following tasks are performed by an SP Administrator (CSO on-premises version) or the Juniper Networks team (CSO SaaS version)
   a. Log in to the CSO Administration Portal. See "Log In to CSO Administration Portal" on page 30.
   b. Configure SMTP settings. See "Configure SMTP Settings in CSO" on page 31.
   c. Download latest signature database. See "Download the Signature Database" on page 33.
   d. (Optional) Customize configuration templates. See "Configuration Templates Workflow" on page 35.
   e. (Optional) Customize device templates. See "Device Templates Workflow" on page 36.
   f. (Optional) Upload the latest software images to CSO. See "Device Images Workflow" on page 37.
   g. If you are deploying SD-WAN, perform the following tasks:
      i. Add one or more points of presence (POPs). See "Add a Point of Presence (POP)" on page 37.
      ii. Add Operation, Administration, and Maintenance (OAM) provider hub devices. See "Add a Provider Hub Device" on page 38.
      iii. (Optional) Add data or data and OAM provider hub devices.
   h. (Optional) Add an OpCo. See "Add an Operating Company (OpCo)" on page 48.
   i. Add a tenant with SD-WAN service, NGFW service, or both services. See "Add a Tenant" on page 51.
   j. Add CSO licenses. See "Add CSO Licenses" on page 68.

   After completing these tasks, go to step 5.
3. If you are an OpCo Administrator:
   a. Log in to the CSO Administration Portal. See "Log In to CSO Administration Portal" on page 30.

b. (Optional) Configure SMTP settings. See "Configure SMTP Settings in CSO" on page 31.

   This task is optional because the SMTP settings are configured by the service provider (for CSO on-premises) or Juniper Networks (for CSO SaaS).

c. (Optional) Customize configuration templates. See "Configuration Templates Workflow" on page 35.

d. (Optional) Customize device templates. See "Device Templates Workflow" on page 36.

e. (Optional) Upload the latest software images to CSO. See "Device Images Workflow" on page 37.

f. (Optional) If you are deploying SD-WAN and you want to add provider hubs for tenants, do the following:

   i. Add one or more points of presences (POPs). See "Add a Point of Presence (POP)" on page 37.

   ii. Add data or data and OAM provider hub devices. See "Add a Provider Hub Device" on page 38.

   > **NOTE**: For CSO SaaS, OpCo Administrators should add only provider hub devices with DATA_ONLY capability because the Juniper Networks adds the OAM-capable hubs.

g. Add a tenant with SD-WAN service, NGFW service, or both services. See "Add a Tenant" on page 51.

h. Assign CSO licenses to tenants. See "Assign CSO Licenses to Tenants" on page 71.

4. To ensure WAN redundancy in CSO, an SP Administrator or an OpCo Administrator can preconfigure aggregated Ethernet links enterprise hub devices for tenants. See "WAN Link Redundancy in Enterprise Hubs Using Aggregated Ethernet" on page 61 for more information.
5. Depending on whether you're deploying SD-WAN or NGFW, see "CSO SD-WAN Deployment Workflow" on page 75 or "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Log In to CSO Administration Portal

Review "Access the Contrail Services Orchestration (CSO) GUIs" on page 9 to get details of the portal link, username, and so on.

To log in to the CSO Administration Portal:

1. Open the link to the CSO portal in a Web browser.
2. Enter your username and password in the respective text boxes.
3. Click **Log In**.

   CSO authenticates your username and password and if the login in successful, the Welcome page appears.
4. Click the close icon (**X**) to close the Welcome page, or click **Go to Dashboard** to go the Dashboard page.

   You can now perform tasks in Administration Portal.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

Configure SMTP Settings in CSO

SMTP settings are used when CSO sends e-mails to users when their CSO account is created, when their password is reset, and so on. Therefore, we recommend that you configure the SMTP settings for the CSO on-premises version; for CSO SaaS, Juniper Networks configures SMTP settings.

Users with the service provider (SP) Administrator or Operating Company (OpCo) Administrator roles (or users with the necessary access privileges) can configure SMTP settings on CSO.

NOTE: The SMTP settings configured by the SP Administrator are applicable to the SP's tenants, OpCos, and OpCo's tenants. The SMTP settings configured by the OpCo Administrator are applicable only to the OpCo and the OpCo's tenants.

To configure SMTP settings:

1. Click **Administration > SMTP**.

   The SMTP page appears.
2. Configure the SMTP settings, as explained in Table 10 on page 32.

   NOTE: Fields marked with an asterisk (*) are mandatory.
3. Click **Save**.

   The SMTP settings are saved and a confirmation message appears at the top of the page.

   After the confirmation message appears, you can navigate away from this page.

Table 10: SMTP Settings

Field	Description
SMTP Server	
Server Address	Enter the hostname of the SMTP server.
TLS	By default, this field is enabled, which means that Transport Layer Security (TLS) protocol is used to ensure that the e-mail messages are transmitted over an encrypted channel. You can click the toggle button to disable TLS.
Port Number	By default, the port number for the SMTP server is set to 587 when TLS is enabled and to 25 when TLS is not enabled. However, you can modify the port number. The port number is typically provided by your e-mail service provider.
SMTP Authentication	
SMTP Authentication	By default, SMTP authentication is enabled, which means that you must provide authentication credentials (username and password) for the SMTP server. The **Username** and **Password** fields are displayed when you enable this option. Click the toggle button to disable SMTP authentication, which means that you don't need to provide authentication credentials for the SMTP server.
Username	If SMTP authentication is enabled, enter the username to use for authenticating with the SMTP server.
Password	If SMTP authentication is enabled, enter the password to use for authenticating with the SMTP server.
Confirm Password	If SMTP authentication is enabled, re-enter the password for confirmation.
From Name	If SMTP authentication is disabled, enter the name from which you want the email to be sent.

Table 10: SMTP Settings *(Continued)*

Field	Description
From Email Address	Enter your e-mail address from which you want the e-mails to be sent. This e-mail address appears as the sender's e-mail address to the e-mail recipient.
Test SMTP Settings	
Email Address	Though this is not mandatory, we recommend that you verify that the SMTP settings are configured correctly. To verify that the SMTP server is working, enter your e-mail address, and click **Send Test E-mail**. If you receive an e-mail at the specified e-mail address, this confirms that the SMTP settings are configured correctly.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

Download the Signature Database

NOTE: This topic is applicable only for the CSO on-premises version. For CSO SaaS, you can skip this task because the signature database is downloaded by the Juniper Networks team.

The signature database that Juniper Networks provides contains application and intrusion prevention system (IPS) signatures:

- Application signatures are definitions of predefined attacks and applications, and can be used to identify applications for tracking firewall policies and quality-of-service (QoS) prioritization.
- IPS signatures are definitions of predefined attack object and attack object groups that you can use in IDP policies to match traffic against known attacks.

To download the signature database:

1. Select **Administration > Signature Database**.

   The Signature Database page appears.
2. Click **Signature Download Settings**.

The Signature Download Settings page appears.

3. Enter the download settings according to the guidelines provided in Table 11 on page 34.
4. Click **OK** to save the changes:
   - If you specified that the signature database should be downloaded immediately, a Job Tasks page appears displaying information about the signature download job. Click **OK** to close this page and return to the Signature Database page.
   - If you scheduled the signature download for later, a job is triggered and you are returned to the Signature Database page. A confirmation message (with the job ID) is displayed at the top of the page.

   After the signature download operation is complete, predefined signatures (application and IPS) and IPS profiles are available in CSO. You cannot modify predefined signatures or IPS profiles.

Table 11: Fields on the Signature Download Settings Page

Field	Description
Download URL	Specifies the location of the Juniper hosted server from which the signature database is downloaded to the CSO server. The default download URL is https://signatures.juniper.net/. To download signatures from this location, Internet connectivity must be available from CSO.
Signature Version	**NOTE**: This field is enabled only when you change the download URL from https://signatures.juniper.net/. Enter the numeric value of the signature database version. The value must only contain numbers and not have any special characters or negative values.
Type	You can chose to download the signature database immediately or schedule the download for later. • Select **Run now** to automatically download the signature database immediately. • Select **Schedule at a later time** to download the signature database later and specify the date and time. **NOTE**: The time zone is based on the time-zone specified when CSO is installed.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

Configuration Templates Workflow

> **NOTE**: In Administration Portal, users with the SP (Service Provider) Administrator role (CSO on-premises only) or OpCo (Operating Company) Administrator role can perform the configuration template workflow tasks indicated in this topic. In Customer Portal, users with the Tenant Administrator role can perform these workflow tasks.

The high-level workflow for configuration templates is as follows:

1. You can use a pre-existing template (skip to step 2), or create a new template using one of the following methods:
   - Import a configuration template by specifying the template configuration file (Jinja syntax), Yang model file, and the Viewdef file.
   - Clone an existing configuration template and modify the cloned template.
   - Add a configuration template by specifying the template configuration and logic.
2. (Optional) Although this is an optional step, we recommend that you validate the configuration template by using the preview workflow *before* attaching the configuration template to a device template or deploying the configuration template directly on a device.
3. You can assign a configuration template to a device template from the Configuration Templates or the Device Templates pages. This enables you to deploy additional configuration on the device during zero touch provisioning (ZTP) and after the device is activated.
4. You can deploy a configuration template directly on one or more devices that were previously activated, which enables you to deploy templates that were added after a device was activated or to deploy additional configuration to devices. You can deploy configuration templates to devices from the Configuration Templates or Tenant Devices pages
5. (Optional) Dissociate or undeploy configuration templates:
   - You can dissociate a configuration template from a device, which remove the references to the configuration template from the device, but retains the configuration already deployed on the device.
   - You can undeploy the configuration template, which deletes the configuration previously deployed on the device, but retains the references to the configuration template.

For more information, see *About the Configuration Templates Page* topics in the *CSO Administration Portal User Guide* and the *CSO Customer Portal User Guide* respectively (available on the CSO Documentation page).

Depending on the workflow, see "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29, "CSO SD-WAN Deployment Workflow" on page 75, or "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178 for the next task.

Device Templates Workflow

Device templates contain configuration and provisioning settings for the different spoke and hub devices that are supported by CSO. When you add a provider hub device, an enterprise hub site, an on-premises spoke site, and a cloud spoke site, you must choose the device template to use and specify the parameters for the template. CSO then uses these values for configuring and provisioning.

The high-level workflow for device templates is as follows:

1. Select **Resources > Templates > Device Templates**.

   The Device Template page appears.
2. You can do one of the following:
   - Use predefined device templates as-is.
   - Create a customized template by cloning a predefined template and then modifying the cloned template.
   - Create a customized template by Importing a device template by using a JavaScript Object Notation (JSON) file.

   NOTE:
   - Customized device templates created by the SP Administrator are available to the service provider's tenants, OpCos, and the OpCo's tenants.
   - Customized device templates created by the OpCo Administrator are available to the OpCo and the OpCo's tenants.
   - Customized device templates created by the Tenant Administrator are available only to the tenant.
3. Provide values for the initial configuration parameters for the existing configuration templates that are associated with the device template.
4. Add one or more configuration templates to a device template, and provide initial configuration values for the parameters specified in the configuration templates.
5. Use the device template when you are adding a site.

For more information, see the *About the Device Template Page* topics in the *CSO Administration Portal User Guide* and the *CSO Customer Portal User Guide* respectively (available on the CSO Documentation page).

Depending on the workflow, see "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29, "CSO SD-WAN Deployment Workflow" on page 75, or "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178 for the next task.

Device Images Workflow

CSO's image management system allows you to upload device images, stage the image on devices, and deploy images on devices. Users with SP Administrator or OpCo Administrator roles can manage images in Administration Portal and Tenant Administrators can manage images in Customer Portal. The availability of a device image depends on the scope at which the image was uploaded. For example, a device image uploaded in the global scope is available at all scopes but an image uploaded at the tenant scope is available only to that tenant.

The device image workflow is as follows:

1. Upload one or more device images. See *Uploading a Device Image* in the *CSO Administration Portal User Guide* (available at the CSO Documentation page.
2. (Optional) Stage the device image on one or more devices. See *Staging an Image* in the *CSO Administration Portal User Guide*.

   NOTE: Staging an image is useful when you have a low bandwidth connection and you want to avoid the image deployment from timing out.

3. Deploy the device image on one or more devices. See *Deploying Device Images to Devices* in the *CSO Administration Portal User Guide*.

See "CSO SD-WAN Deployment Workflow" on page 75 for the next task.

Add a Point of Presence (POP)

In CSO, a POP refers to a location where one or more provider hub devices are located. Therefore, you must add at least one POP to which provider hub devices can then be assigned.

To add a POP:

1. Select **Resources > POPs**.

   The POPs page appears, displaying a list of existing POPs.
2. Click the Add (**+**) icon.

   The Add POP page appears.
3. Complete the configuration settings according to the guidelines provided in Table 12 on page 38.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   CSO triggers a job to add the site, and displays confirmation messages when the job is triggered and when the job is completed. You are returned to the POPs page.

   TIP: Refresh the page and verify that the POP is added.

Table 12: Add POP Settings

Field	Guideline
Region	Displays *regional* as the region. You cannot modify this field.
Name	Enter the name for the POP
Address Information	Optionally, enter the address where the POP is located in the fields provided:

See "Add a Provider Hub Device" on page 38.

Add a Provider Hub Device

A provider hub device resides in a POP within the SP or OpCo network. Provider hub devices are shared amongst multiple tenants through the use of virtual routing and forwarding (VRF) instances configured on the provider hub itself. They allow site-to-site traffic to flow in hub-and-spoke deployments, serve as OAM gateway devices for management traffic between CSO and CPE devices, and can serve as backup data hubs when an enterprise hub device is used in a tenant.

Provider hubs come in three varieties: OAM_ONLY, DATA_ONLY, or OAM_AND_DATA.

- OAM_ONLY and OAM_AND_DATA hubs pass OAM traffic between CSO and the CPE devices. CPE devices connect to these OAM-capable hubs over IPSec. In the CSO on-premises installation, the SP administrator adds the OAM-capable hubs. In CSO SaaS, the OAM-capable hubs are provided (by Juniper Networks) as part of the service.
- DATA_ONLY and OAM_AND_DATA hubs route site-to-site user traffic in a hub-and-spoke topology. These data-capable provider hubs are optional. In the CSO on-premises installation, the SP or OpCo administrator creates the data-capable hubs. In CSO SaaS, the OpCo administrator creates the data-capable hubs.

BEST PRACTICE: It is recommended that all provider hubs be clearly named for their data and OAM capabilities.

NOTE:

- For SD-WAN Advanced service, we recommend that you configure two OAM-capable provider hubs to provide redundancy in the OAM network. In CSO SaaS, Juniper Networks provisions two OAM-capable hubs by default. In the CSO on-premises version, the SP Administrator must add the OAM-capable hubs.
- Before you add the provider hub, check the cable connections, review the NAT and firewall ports and protocols, and check the Junos OS version of the enterprise hub device, as explained in "Supported Devices for SD-WAN, and Ports and Protocols to Open" on page 168.

To add a provider hub device:

1. Select **Resources > Provider Hub Devices**.

   The Provider Hub Devices page appears.
2. Click the add (**+**) icon.

   The Add Provider Hub page appears, displaying the General settings to configured.
3. Configure the General settings as explained in Table 13 on page 41, and click **Next**.

   You are taken to the WAN section of the wizard.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Configure the WAN settings as explained in Table 14 on page 42 and click **Next**.

   You are taken to the Summary section of the wizard.

5. Review the configuration in the Summary tab, and modify the settings, if required.

   You can also download the settings that you configure as a JavaScript Object Notation (JSON) file by clicking the **Download as JSON** link at the bottom of the page

6. Click **OK**.

   - If you entered a serial number during activation and automatic activation is enabled, the Site Activation Progress page appears. The site activation process proceeds through the tasks explained in Table 27 on page 110.

     Click **OK** to close the page.

     > **NOTE**: If you don't want to wait for the provider hub activation to finish, you can close the page and monitor the status of the activation from the Jobs page (**Monitor > Jobs**).
     > The time taken for provider hub activation varies depending on the device that CSO is activating.

   - If you did not enter a serial number or if automatic activation is disabled, you are returned to the Provider Hub Devices page. CSO triggers a job and displays a confirmation message with a job link. Click the link to view the status of the job.

     After the job is finished, CSO displays a confirmation message with a job link. The status of the site changes to CREATED. You must manually activate the device to finish the activation process. To manually activate the provider hub:

     a. Select the device and click **Activate Device**.

        The Activate Site page appears.

     b. If a serial number was not specified when the site was added, enter the serial number of the device in the **Serial Number** field. Serial numbers are case sensitive.

        If the serial number that you entered is already present in the system, CSO displays an error message. If the serial number is not present, then CSO displays a green check mark.

     c. If automatic activation was disabled when the site was added, enter the activation code of the device in the **Activation Code** field.

     d. Click **OK**.

        CSO triggers a job and the Site Activation Progress page appears after a few seconds. Because the site was previously modelled, the Ship Device task is the first task to be executed. The rest of the steps are as explained in Table 27 on page 110.

TIP: After you add a provider hub, you can modify certain parameters for DATA_ONLY provider hubs. For more information, see the *Edit Provider Hub Site Parameters* topic in the *CSO Administration Portal User Guide* (available on the CSO Documentation page).

Table 13: General Settings (Add Provider Hub [Device] Page)

Field	Guideline
Site Information	
Site Name	Enter a name for the provider hub device. The name can contain alphanumeric characters and hyphens (-) and must not exceed 15 characters. For example, LA-PHub-OAM
Device Host Name	The device host name is auto-generated by default. You can change the name but not the prefix. Use alphanumeric characters and hyphen (-); the maximum length allowed is 32 characters.
Management Region	Displays *regional* as the management region. You cannot modify this field.
Site Capability	Select the capability of the provider hub device: • **OAM_ONLY**—Transmits only OAM traffic. **NOTE**: This option is available only for SP Administrator users in the on-premises version of CSO. • **DATA_ONLY**—Transmits only data traffic. • **OAM_AND_DATA**—Transmits both data traffic and OAM traffic. For provider hubs added with data only capability, CSO establishes a secure OAM tunnel between the provider hub with data capability and a provider hub with OAM_ONLY or OAM AND DATA capability).
POP	Select the POP to which you want to assign the provider hub device.

Table 13: General Settings (Add Provider Hub [Device] Page) *(Continued)*

Field	Guideline
Authentication Type	Select the type of authentication to use for establishing secure IPsec tunnels: • Pre-shared key, which is the default. • Public Key Infrastructure
Advanced Configuration	
Domain Name Server	Specify the IPv4 or IPv6, or both IPv4 and IPv6 addresses of one or more Domain Name System (DNS) servers.
NTP Server	Specify the IP address or fully-qualified domain name (FQDN) of the NTP server.
Select Timezone	Select the time zone to which the provider hub device belongs.

Table 14: WAN Settings (Add Provider Hub [Device] Page)

Field	Guideline
Device Series	Displays SRX as the device series because currently only SRX Series devices are supported as provider hubs.
[Device Template]	Ensure that you select the correct device template for the provider hub device from the carousel. For example, for an SRX1500 device, you can select **SRX as SD-WAN Hub** (or a modified version of that template) as the device template. **NOTE**: Check that the interface names in the device template match the ones on the device that you're using.
Device Information	

Table 14: WAN Settings (Add Provider Hub [Device] Page) ***(Continued)***

Field	Guideline
Serial Number	If you want CSO to proceed with the provider hub activation immediately after you complete the add provider hub workflow, enter the serial number. If the serial number that you entered is already present in the system, CSO displays an error message. If the serial number is not present, then CSO displays a green check mark. If you want CSO to only model the provider hub, leave this field blank. If you don't enter a serial number, you must manually activate the provider hub later.
Auto Activate	Automatic activation is typically enabled by default (based on the setting in the device template). When automatic activation is enabled, zero-touch provisioning (ZTP) of the provider hub device is automatically triggered after the site is added to CSO. If you want the device to be activated manually, click the toggle button to disable automatic activation.
Activation Code	If you disabled automatic activation, enter the activation code that must be entered when the device is manually activated later. When you manually activate the device later, CSO checks the activation code entered against the activation code specified here and activates the device only if the activation codes match.
Device Root Password	The default root password is fetched from the ENC_ROOT_PASSWORD field in the services template. You can keep the password or change it by entering a password in plain-text format. The password is encrypted and stored.
Boot Image	If you want to upgrade the provider hub device with the latest supported Junos OS version, select the boot image from the list. The boot image is used to upgrade the device when CSO starts the zero touch provisioning (ZTP) process. If you don't specify a boot image, which is the default option (**Use Image on Device**) in the list, then the CSO skips the procedure to upgrade the device during ZTP.
Management Connectivity	
Loopback IP Prefix	By default, CSO assigns the IPv4 address prefix for the loopback interface on the device.

Table 14: WAN Settings (Add Provider Hub [Device] Page) *(Continued)*

Field	Guideline
OAM Interface	For provider hubs with OAM or OAM and data capabilities, select the interface on the provider hub device that you want to use to connect the provider hub device to CSO. This interface is used only for OAM connectivity. The interface names are listed are the names configured in device template.
OAM VLAN	For provider hubs with OAM or OAM and data capabilities, enter an OAM VLAN ID for in-band management of the hub device. If you specify an OAM VLAN ID, then in-band OAM traffic reaches the device through the selected OAM interface.
OAM IP Prefix	For provider hubs with OAM or OAM and data capabilities, enter an IPv4 address prefix for the OAM interface in the provider hub device. The prefix must be unique across the entire management network.
OAM Gateway	For provider hubs with OAM or OAM and data capabilities, enter the IP address of the next-hop through which the connectivity from the provider hub device to CSO is established.
EBGP Peer-AS	For provider hubs with OAM or OAM and data capabilities, enter the autonomous system (AS) number of the external BGP (EBGP) peer. The AS number is unique to the service provider and is needed to establish the EBGP peering session.
WAN Links	
WAN_0 (*Interface-Name*)	This field is enabled by default. Enter parameters related to WAN_0. You must configure the fields marked with an asterisk (*) to proceed.
Local Interface	Displays the interface name configured in the device template. You cannot modify this field.
Link Type	Select the underlay network type (MPLS or Internet) of the WAN link.

Table 14: WAN Settings (Add Provider Hub [Device] Page) ***(Continued)***

Field	Guideline
Public IP Address	Enter the public IPv4 address for the WAN link. This IP address should be provided only if the static IP prefix is a private address and 1:1 NAT is configured.
Data VLAN ID	Enter the VLAN ID that is associated with the WAN link.
Underlay Address Families	
IPv4	Click the toggle button to enable or disable IPv4 address assignment for the WAN link. By default, IPv4 address assignment is enabled for the WAN link. The WAN link requires an IPv4 address to connect to an IPv4 network.
Address Assignment Method	Displays the address assignment method used for the IPv4 WAN link (STATIC). You cannot modify this field.
Static IP Prefix	Enter the IPv4 address prefix of the WAN link.
Gateway IP Address	Enter the IPv4 address of the gateway of the WAN service provider.

Table 14: WAN Settings (Add Provider Hub [Device] Page) *(Continued)*

Field	Guideline
MTU	Applicable only to IPv4 addresses. Enter the maximum transmission unit (MTU) size for the media or protocol. The supported MTU range can vary depending on the device, interface type, network topology, and other individual requirements. See also: MTU Default and Maximum Values and LTE Mini Physical Interface Modules (LTE Mini-PIM). Editing the MTU values of all the OAM-enabled WAN links of a site at the same time might result in tunnel flapping. You must ensure that at least one OAM-enabled WAN link always remains undisrupted for a site. For example, if you have a site with four WAN links (including two links that support OAM traffic), you can edit the MTU values of all the WAN links except one OAM-enabled link at the same time. After the edit is complete and the changes are saved, you can edit the site again and update the remaining WAN link. **NOTE**: If you enable the PPPoE/PPP option under a WAN link, the MTU option is displayed under the PPPoE/PPP Settings section for that link.
IPv6	Click the toggle button to enable or disable IPv6 address assignment for the WAN link. By default, IPv6 address assignment is disabled for the WAN link. The WAN link requires an IPv6 address to connect to an IPv6 network.
Address Assignment Method	Displays the address assignment method used for the IPv6 WAN link (STATIC). You cannot modify this field.
Static IP Prefix	Enter the IPv6 address prefix of the WAN link.
Gateway IP Address	Enter the IPv6 address of the gateway of the WAN service provider.
WAN_1 (*Interface-Name*)	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. You must configure the fields marked with an asterisk (*) to proceed. Refer to the fields described for WAN_0 (*Interface-Name*) for an explanation of the fields

Table 14: WAN Settings (Add Provider Hub [Device] Page) ***(Continued)***

Field	Guideline
WAN_2 (*Interface-Name*)	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. You must configure the fields marked with an asterisk (*) to proceed. Refer to the fields described for WAN_0 (*Interface-Name*) for an explanation of the fields
WAN_3 (*Interface-Name*)	Click the toggle button to enable or disable the WAN link. When you enable the WAN link, fields related to the WAN link appear. You must configure the fields marked with an asterisk (*) to proceed. Refer to the fields described for WAN_0 (*Interface-Name*) for an explanation of the fields
Additional Configuration	If you want to deploy additional configuration during the ZTP process, you can select one or more configuration templates and set the parameters for each template. The configuration templates for the device family are displayed.
Configuration Templates List	For each configuration template that you select: 1. Select one or more configuration templates from the list that you want to deploy on the device during ZTP. 2. Click **Set Parameters**. The Device Configurations page appears. The names and configuration parameters of the configuration templates that you selected are displayed in the Configure tab. 3. For each configuration template, enter values for the parameters. 4. (Optional) Click the **Summary** tab to view the Junos OS configuration commands that will be deployed on the device for the different configuration templates. 5. Click **Save**. You are returned to the WAN tab. The Junos OS configuration commands will be deployed on the device during the ZTP process.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

Add an Operating Company (OpCo)

NOTE: This topic is applicable to the CSO on-premises version.

An operating company (OpCo) is a managed service provider that can add and manage its own tenants and provide services to those tenants. Only users with the SP Administrator role or equivalent permissions can add OpCos. For more information on OpCos, see *Operating Companies Overview* in the *CSO Administration Portal User Guide* (available on the CSO Documentation page).

To add an OpCo:

1. Select **Tenants > Operating Companies**.

   The Operating Companies (OpCo) page appears.
2. Click the Add icon (**+**).

   The Create Operating Company (OpCo) page appears.
3. Complete the configuration according to the guidelines provided in Table 15 on page 48.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   CSO triggers a job to add an OpCo and returns you to the Operating Companies (OpCo) page. A confirmation message with a link to the triggered job appears on the top of the page. When the job finishes, a confirmation message appears along with a link to the job. (You can click the job link to view the job details or go to the Jobs page (**Resources > Jobs**) and view the job details.)

   The new operating company is then displayed on the Operating Companies (OpCo) page.

Table 15: Fields on the Create Operating Company Page

Field	Description
Name	Enter a unique name for the operating company. The name can contain alphanumeric characters, underscores, and periods, and cannot exceed 15 characters.
Portal URL	

Table 15: Fields on the Create Operating Company Page *(Continued)*

Field	Description
Admin Portal	Enter the URL that OpCo users must use to access the CSO Administration Portal.
Tenant Portal	Enter the URL that OpCo users must use to access the CSO Customer Portal.
Authentication Method for OpCo	
SP User	Select the authentication method to authenticate OpCo users. • **Same as Global**—Select this option to use the authentication method which is used by the Global SP. • **Allow OpCo to decide**—Select this option to use OpCo's own authentication method.
Tenant User	Select the authentication method to authenticate OpCo's tenant users. The default method is local authentication. • **Same as Global**—Select this option to use the authentication method which is used by the Global SP. • **Allow OpCo to decide**—Select this option to use OpCo's own authentication method.
Admin User	
First Name	Enter the first name of the OpCo administrator user.
Last Name	Enter the last name of the OpCo administrator user.
Username (Email)	Enter the e-mail address of the OpCo administrator user. The e-mail address is the username that the OpCo Administrator user will use to log in to CSO.

Table 15: Fields on the Create Operating Company Page *(Continued)*

Field	Description
OpCos [Roles]	Select one or more OpCo roles that you want to assign to the OpCo administrator user, and click the greater than (>) icon. Both predefined and custom roles are displayed. The following are the predefined roles for OpCo users: • **OpCo Admin**—Users with the OpCo Admin role have full access to the OpCo's Administration Portal UI or API capabilities. • **OpCo Operator**—Users with the OpCo Operator role have read-only access to the OpCo's Customer Portal UI and APIs.
Tenants [Roles]	Select one or more tenant roles (predefined or custom) that you want to assign to the OpCo administrator user and click the greater than (>) icon. Both predefined and custom roles are displayed. The following are the predefined tenant roles available: • **Tenant Admin**—Users with the Tenant Admin role have full access to the Customer Portal UI or API capabilities. • **Tenant Operator**—Users with the Tenant Operator role have read-only access to the Customer Portal UI and APIs.
Password Policy	
Password Expiration Days	Specify the duration (in days) after which the password expires and must be changed. The range is from 1 through 365, and the default value is 180 days.
MAP-E Network Settings	
MAP-E for Tenants	Click the toggle button to enable the OpCo's tenants to configure the Mapping of Address and Port with Encapsulation (MAP-E) functionality for their branch sites with NFX150 as the CPE. MAP-E supports transporting IPv4 packets across an IPv6 network by using IPv4-in-IPv6 encapsulation. If you enable this toggle button, you must select a manufacturer code to be associated with the OpCo user. **NOTE**: MAP-E is compliant only with the Japan Network Enabler (JPNE) standards.

Table 15: Fields on the Create Operating Company Page *(Continued)*

Field	Description
Manufacturer Code	From the list of manufacturer codes assigned to Juniper Networks, select a manufacturer code to be associated with the OpCo user. If a tenant belonging to the OpCo enables MAP-E for a branch site with NFX150 as the CPE device, the device uses this manufacturer code to obtain the MAP-E rules from the MAP-E rule server.

See "Add a Tenant" on page 51.

Add a Tenant

In CSO, a tenant is a logical representation of a customer. Tenants enable the separation and isolation of resources (such as sites) and traffic of different customers from one another.

To add a tenant:

1. From the CSO menu, select **Tenants**.

   The Tenants page appears.
2. Click the Add (**+**) icon.

   The Add Tenants wizard appears, displaying the General settings to be configured.

   NOTE: Fields marked with an asterisk (*) are mandatory.
3. Configure the General settings as explained in Table 16 on page 52, and click **Next**.

   You are taken to the Deployment Info section of the wizard.
4. Configure the Deployment Info settings as explained in Table 17 on page 53, and click **Next**.

   You are taken to the Tenant Properties section of the wizard.
5. Configure the Tenant Properties settings as explained in Table 18 on page 54, and click **Next**.

   You are taken to the Summary section of the wizard, where a summary of the settings that you configured is listed.
6. Review the configuration in the Summary section and, if needed, modify the settings.

> **NOTE**: You can download the tenant settings that you configured as a JavaScript Object Notation (JSON) file by clicking the **Download as JSON** link at the bottom of the Summary section.

7. Click **Finish**.

   You are returned to the Tenants page, and CSO triggers a job to add the tenant and displays a confirmation message. Click the link in the message to view the details of the job. Alternatively, you can check the status of the job on the Jobs (**Resources > Jobs**) page.

   After the job finishes successfully, the tenant that you added is displayed on the Tenants page.

   If an SMTP server is configured. an e-mail is sent to the tenant administrator user that you configured, which includes a URL to access Customer Portal. The URL is active for only 24 hours and is valid only for the first login.

Table 16: General Settings (Add Tenant)

Field	Guideline
Basic Information	
Name	Enter a unique name for the tenant. The name can contain alphanumeric characters, underscores, and hyphens, and must be less than 32 characters long. For example, Ent_Tenant.
Password Policy	
Password Expiration Days	Specify the duration (in days) after which the password will expire and must be changed. Range: 1 through 365. Default: 180.
Admin User	You must add an administrator user that can perform the administration tasks for that tenant.
First Name	Enter the first name of the administrator user.
Last Name	Enter the last name of the administrator user.

Table 16: General Settings (Add Tenant) ***(Continued)***

Field	Guideline
Username (Email)	Enter the e-mail address of the administrator user. The e-mail address will be the username that the administrator user will use to log in to the CSO portal. After the tenant is added successfully, CSO sends an e-mail containing the link to the CSO portal and a link to set the password.
Roles	Select one or more roles (predefined or custom) that you want to assign to the tenant user, and click the right arrow (>) to move the selected role or roles from the Available column to the Selected column.

Table 17: Deployment Info Settings (Add Tenant)

Field	Guideline
Services	
Services for Tenant	Select the services that you want to be available for the tenant: • **SD-WAN**—If you select SD-WAN, the tenant can add on-premise spoke sites (with SD-WAN capability), enterprise hub sites, and cloud spoke sites. • **Security Services**—If you select NGFW (next-generation firewall), the tenant can add on-premise spoke sites with NGFW capability.
Service Level	**NOTE**: This field appears only if you selected the **SD-WAN** in the Services for Tenant field. Choose an SD-WAN service type for the tenant. The following options are available: • **Essential**—Provides the basic SD-WAN services. This service does not support multihoming, dynamic mesh tunnels, cloud breakout profiles, SLA-based steering profiles, or pool-based source NAT rules. • **Advanced**—Provides complete SD-WAN services. All sites of the tenant are connected in full mesh or hub-and-spoke topology. This service includes Secure SD-WAN Essential service.

Table 18: Tenant Properties Settings (Add Tenant)

Field	Guideline
SSL Settings	This setting is applicable only to tenants with SD-WAN service.
Default SSL Proxy Profile	Click the toggle button to enable a default SSL proxy profile for the tenant. This option is disabled by default. If you enable this option, you must add a root certificate. If you enable this option and add the root certificate, the following items are created when a tenant is added: • A default root certificate with the certificate content specified (in the Root Certificate field) • A default SSL proxy profile • A default SSL proxy profile intent that references the default profile **NOTE**: You use this option to add a tenant-wide default profile; enabling or disabling this option does *not* mean that SSL is enabled or disabled.

Table 18: Tenant Properties Settings (Add Tenant) ***(Continued)***

Field	Guideline
Root Certificate	You can add a root certificate (X.509 ASCII format) by importing the certificate content from a file or by pasting the certificate content: • To import the certificate content directly from a file: 1. Click **Browse**. The **File Upload** dialog box appears. 2. Select a file and click **Open**. The content of the certificate file is displayed in the Root Certificate field. • Copy the certificate content from a file and paste it in the text box. After the tenant is successfully added, a default root certificate, a default SSL proxy profile, and a default SSL proxy profile intent are added. **NOTE**: • The root certificate must contain both the certificate content and the private key. • For full-fledged certificate operations, such as certificates that need a passphrase, or that have RSA private keys, you must use the Certificates page (**Administration > Certificate Management > Certificates**) to import the certificates and install the certificates on one or more sites.
VPN Authentication	This setting is applicable only to tenants with SD-WAN service.

Table 18: Tenant Properties Settings (Add Tenant) *(Continued)*

Field	Guideline
Authentication Type	Select the VPN authentication method to establish a secure IPsec tunnel: • **Preshared Key**—Select this option if you want CSO to establish IPsec tunnels using keys. This is the default VPN authentication method. **NOTE**: When preshared key is used as the VPN authentication method, CSO generates a random preshared key for each IPsec tunnel and pushes the key to the two devices between which the IPsec tunnel is established. • **PKI Certificate**—Select this option if you want CSO to establish IPsec tunnels using public key infrastructure (PKI) certificates. If you select this option, the following fields appear: • **CA Server URL**—Specify the Certificate Authority (CA) Server URL. For example, http:// *CA-Server-IP-Address*/certsrv/mscep/ mscep.dll/pkiclient.exe. The CA server manages the life cycle of a certificate. The CA server also publishes revoked certificates to the certification revocation list (CRL) server. To obtain trusted CA certificates, CSO communicates with the CA server using the Simple Certificate Enrollment Protocol (SCEP). • **Password**—Specify the password for the CA server. This field is optional. • **CRL Server URL**—Specify the certificate revocation list (CRL) server URL. For example, http:// *Revocation-List-Server-IP-Address*/certservices/abc.crl. CSO retrieves the list of revoked certificates from the CRL server. • **Auto Renew CA Certificates**—Click the toggle button to enable automatic renewal of certificates. By default, the Auto Renew toggle button is disabled, which means that certificates must be manually renewed. If you enable the Auto Renew toggle button, certificates are automatically renewed for all sites in the tenant. **NOTE**: If the certificate is expired before the renewal, CSO might not be able to reach the device. • **Renew before expiry**—This field appears only if you enabled the automatic renewal of certificates.

Table 18: Tenant Properties Settings (Add Tenant) ***(Continued)***

Field	Guideline
	Select the period (3 days, 1 week, 2 weeks, or 1 month) before the expiration date when the certificates get automatically renewed. **NOTE**: The default value is 2 weeks. You can also change the duration in the VPN Authentication page in Customer Portal (**Administration > Certificate Management > VPN Authentication**) page.
Overlay Tunnel Encryption	This setting is applicable only to tenants with SD-WAN service (Advanced or Essential).
Encryption Type	For security reasons, all data that passes through the VPN tunnel must be encrypted. Select the type of encryption to use: • 3DES-CBC—Triple Data Encryption Standard with Cipher-Block Chaining (CBC) algorithm. • AES-128-CBC—128-bit Advanced Encryption Standard with CBC algorithm. • AES-128-GCM—128-bit Advanced Encryption Standard with Galois/Counter Mode (GCM) algorithm. • AES-256-CBC—256-bit Advanced Encryption Standard with CBC algorithm. This is the default. • AES-256-GCM—256-bit Advanced Encryption Standard with GCM algorithm.
Network Segmentation	This setting is applicable only to tenants with SD-WAN service.
Network Segmentation	In CSO, network segmentation, which is enabled by default, allows you to isolate the traffic of one department from another because CSO creates a unique Layer 3 VPN for each department. Enabling network segmentation also allows you to use overlapping IP addresses across departments. **NOTE**: • After the tenant is added, you cannot change this setting. • If you disable network segmentation, then the LAN segments (across different sites in a tenant) cannot have overlapping subnets.

Table 18: Tenant Properties Settings (Add Tenant) ***(Continued)***

Field	Guideline
Dynamic Mesh	This setting is applicable only to tenants with SD-WAN Advanced service.
Threshold for Creating a Tunnel	Set a threshold value, above which a tunnel is created between two sites.
Number of sessions	For creating dynamic tunnels, specify the threshold, which is the maximum number of sessions closed between two spoke sites in a two-minute duration. If the number of sessions closed between two spoke sites (in two minutes) exceeds the specified threshold, then a dynamic mesh tunnel is created between the spoke sites The default threshold for tunnel creation value is 5.
Threshold for Deleting a Tunnel	Set a threshold value, below which a tunnel is deleted between two sites.
Number of sessions	For deleting tunnels, specify the threshold, which is the minimum number of sessions closed between two spoke sites in a 15-minute duration. If the number of sessions closed between two spoke sites (in 15 minutes) is lesser than or equal to the specified threshold, then the dynamic mesh tunnel between two spoke sites is deleted The default threshold value for tunnel deletion) is 2.
Max Dynamic Mesh Tunnels	
Max tunnels per CSO	Displays the maximum number of dynamic mesh tunnels that can be created in CSO. The total number of dynamic mesh tunnels that can be created by all tenants in a CSO instance is to 125,000. A major alarm is raised if the number of dynamic mesh tunnels created by all tenants reaches 70 percent of the maximum value. A critical alarm is raised if the number of dynamic mesh tunnels created by all tenants reaches 90 percent of the maximum value. You can view the alarms on the Alarms page (**Monitor > Alerts & Alarms > Alarms**) in Administration Portal.

Table 18: Tenant Properties Settings (Add Tenant) ***(Continued)***

Field	Guideline
Max tunnels per tenant	Specify the maximum number of dynamic mesh tunnels that the tenant can add. Range: 1 through 50,000. A major alarm is raised if the number of dynamic mesh tunnels created by all sites in a tenant reaches 70 percent of the maximum value. A critical alarm is raised if the number of dynamic mesh tunnels created by all sites in a tenant reaches 90 percent of the maximum value. You can view alarms for the tenant on the Alarms page (**Monitor > Alerts & Alarms > Alarms**) in Customer Portal.
Dynamic Mesh	Click the toggle button to disable dynamic meshing between sites in the tenant. Dynamic meshing is enabled by default.
Cloud Breakout Settings	This setting is applicable only to tenants with SD-WAN Advanced service.
Customer Domain Name	Enter the domain name of the tenant. 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. For example, juniper.example.com.
Quality of service settings	This setting is applicable only to tenants with SD-WAN service.

Table 18: Tenant Properties Settings (Add Tenant) *(Continued)*

Field	Guideline
Class of Service	This setting is enabled by default, which means that CSO configures the class of service (CoS) parameters on an SD-WAN site (on-premise spoke, cloud spoke, or enterprise hub site) when you deploy the SD-WAN policy for the site. The CoS parameters are derived from the application traffic type profile associated with the path-based steering profile, SLA-based steering profile, or breakout profile, which is referenced in an SD-WAN policy intent. You can click the toggle button to disable this setting, which means that CSO does not configure CoS parameters for SD-WAN sites, so no CoS parameters are applied to SD-WAN traffic. If you then want to apply CoS parameters on SD-WAN traffic, you must use configuration templates to configure and deploy CoS parameters on the SD-WAN sites. Therefore, unless you want to apply customized CoS parameters by using configuration templates, we recommend that you *do not disable* this setting.
Advanced Settings (Optional)	
Tenant-Owned Public IP Pool	You can add one or more public IPv4 subnets that are part of the tenant's pool of public IPv4 addresses. The tenant IP pool addresses are assumed to be public IP addresses and represent public LAN subnets in SD-WAN on-premise spoke sites. To add an IPv4 subnet: **1.** Click the add (**+**) icon. An editable row appears inline in the table. **2.** In the Addresses field, enter a valid, public IPv4 prefix. **NOTE**: Ensure that the IP addresses configured for a tenant are unique. **3.** Click **✓** (check mark) to save your changes. The prefix that you entered is displayed in the table. You can enter more IPv4 subnets by following the preceding procedure. You can also modify subnets that you entered by selecting a row and clicking the edit (pencil) icon. To delete a subnet, select the subnet and click the delete icon. If you update the IP address pool of a tenant, CSO runs a job to automatically update and reprovision the tenant sites.

Table 18: Tenant Properties Settings (Add Tenant) ***(Continued)***

Field	Guideline
Tenant-Specific Attributes (Optional)	If you have set up a third-party provider edge (PE) device by using software other than CSO, then configure settings on that router by specifying custom properties (parameters) and its corresponding values.
Custom Properties	To add a custom property:: 1. Click the add (**+**) icon. An editable row appears inline in the table. 2. In the **Role Name** field, enter the description of the parameter (property) that you want to pass to the third-party router. 3. In the **Value** field, enter the value of the parameter that you want to pass to the third-party router. 4. Click **√** (check mark) to save your changes. The information that you entered is displayed in the table.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

WAN Link Redundancy in Enterprise Hubs Using Aggregated Ethernet

SUMMARY

Learn about aggregated Ethernet links (AE), how to manually configure LAG and LACP on an enterprise hub, and enable AE links on the enterprise hub WAN links.

IN THIS SECTION

- Aggregated Ethernet Links in Enterprise Hubs | **62**
- Configuration Guidelines for Aggregated Ethernet on WAN Links | **63**

Example: Configure Aggregated Ethernet in Enterprise Hub Devices | 64

Aggregated Ethernet Links in Enterprise Hubs

In CSO Release 6.0.0, a service provider or an OpCo Administrator can aggregate full-duplex gigabit Ethernet WAN links into a single logical aggregated Ethernet (ae*x*) link or link aggregation group (LAG) bundle, as defined by the IEEE 802.3ad standard. Aggregated Ethernet (AE) links topology (shown in Figure 13 on page 62) allows data traffic to flow between two WAN Ethernet interfaces operating at the same speed. This results in WAN redundancy and improves availability even if one physical link fails, as data traffic can flow through the alternative member in the aggregated Ethernet interface.

AE can be configured on WAN links of SRX Series enterprise hub devices. Provisioning LAG bundles in an enterprise hub involves three processes: pre-staging an SRX device, modifying the SRX device template, and enabling aggregated Ethernet on physical WAN ports. The pre-staging configuration of LAG bundle (aggregated Ethernet interface) is performed by service providers or operating companies.

Figure 13 on page 62 shows the topology with LAG bundle configurations deployed during the pre-staging of an enterprise hub. Two gigabit Ethernet interfaces — xe-0/0/0 and xe-0/0/1 — are bundled together into one aggregated Ethernet interface (such as ae0). Similarly, xe-0/0/2 and xe-0/0/3 are configured to form ae1. If xe-0/0/0 fails, data traffic is switched to the xe-0/0/1 interface in ae0. Hence, data traffic continues to flow through the same WAN_0 port configured for AE. The branch site does not have to do WAN link switchover because of hub WAN link failure.

Figure 13: Aggregated Ethernet Topology of Enterprise Hub WAN Links (Without VLAN Tagging)

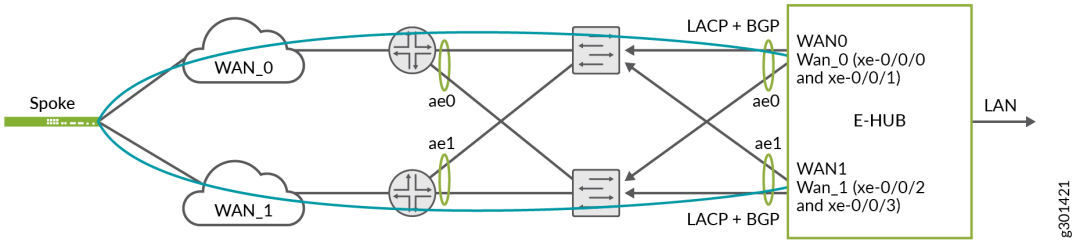

Starting in Release 6.1.0, CSO supports aggregated Ethernet interfaces with VLAN tagging on enterprise hub WAN links, thereby allowing you to ensure WAN link redundancy with lesser number of ports.

Figure 14 on page 63 illustrates a scenario where a VLAN-tagged aggregated Ethernet interface containing two member links is used to provide link redundancy for two WAN links. See .

Figure 14: Aggregated Ethernet Topology of Enterprise Hub WAN Links (With VLAN Tagging)

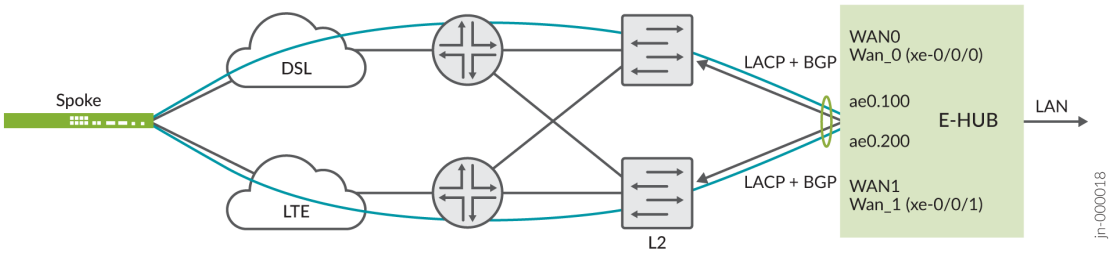

The Link Aggregation Control Protocol (LACP), the protocol defined in IEEE 802.3ad, monitors the interfaces in the aggregated Ethernet link. LACP initiates and establishes LAG connection between the WAN aggregated Ethernet interfaces in enterprise hub and the remote device, monitors the AE interfaces for link failures, and dynamically switches the traffic between member links in an AE interface. LACP flags an AE link down only if all physical member links are operationally down.

After configuring LAG and LACP on the enterprise hub, an SP or OpCo Administrator can modify the device template for enterprise hub in CSO to map physical WAN ports — WAN_0 and WAN_1 — to ae*x* links. Tenant Administrators must enable aggregated Ethernet on WAN ports (while adding an enterprise hub site in Customer Portal).

NOTE: : Links in the aggregated Ethernet bundle support MPLS and Internet data traffic with only Ethernet as the access type for the underlay. VLAN tagging is not supported on aggregated Ethernet interfaces.

Configuration Guidelines for Aggregated Ethernet on WAN Links

Note the following guidelines before you configure aggregated Ethernet or LAG bundle on enterprise hub devices.

- In CSO Release 6.0.0, you must manually configure LAG bundles on the enterprise hub device before zero touch provisioning (ZTP) is initiated to provision an enterprise hub site.

- You must configure link aggregation groups within a configuration group and not at the root level. For example, `set groups WANredundancy interfaces xe-0/0/0 gigether-options 802.3ad ae0`. In CSO, LAG configured at the root level will be removed when sites are provisioned through ZTP.
- Ensure that the LAG configuration group name is unique. The configuration group name must not be the same as groups CSO uses to configure devices. You need to also ensure that the LAG groups used in WAN links are different from LAG groups configured for LAN links.

Example: Configure Aggregated Ethernet in Enterprise Hub Devices

Table 19 on page 64 describes an example configuration snippet for aggregated Ethernet links on enterprise hub devices.

NOTE: You must execute all commands in configuration mode.

Table 19: Example Configuration for Aggregated Ethernet (Without VLAN Tagging)

Configuration Steps	Commands
Step 1: Specify the number of aggregated Ethernet interfaces you want on your device. In the topology for enterprise hub WAN redundancy, the `device-count` value supported is 2. This means, you can configure two aggregated Ethernet interfaces.	`[edit]` `user@host# set groups WANredundancy chassis aggregated-devices ethernet device-count 2`
Step 2: Specify the WAN interfaces (for example, xe-0/0/0) you want to include within the aggregated Ethernet bundle and add them individually. Also enter the interface name of the aggregate Ethernet link to which you add physical WAN member links (for example, ae0).	`[edit]` `user@host# set groups WANredundancy interfaces xe-0/0/0 gigether-options 802.3ad ae0`
Step 3: Specify the minimum number of links in the aggregated Ethernet interface (aex) so that, the ae link is labeled *up*. Only one physical link need to be up for the bundle to be labeled *up*.	`[edit]` `user@host# set groups WANredundancy interfaces ae0 aggregated-ether-options minimum-links 1`

Table 19: Example Configuration for Aggregated Ethernet (Without VLAN Tagging) *(Continued)*

Configuration Steps	Commands
Step 4: Configure LACP on the defined aggregated Ethernet link (for example, ae0) as 'active'. A port with 'active' LACP state can start negotiating an LACP connection with the remote end by sending LACP packets, even if the device at the remote end is in 'passive' state.	[edit] user@host# **set groups *WANredundancy* interfaces *ae0* aggregated-ether-options lacp active**
Step 5: Map an aggregated Ethernet link (ae0) to the IP address of the WAN interface.	[edit] user@host# **set interfaces *ae0* unit 0 family inet address 198.51.100.40/24**
Step 6: Set security zone for the defined aggregated Ethernet (for example, ae0) link and enable traffic on the interface from the defined system services available in the enterprise hub device.	[edit] user@host# **set security zones security-zone *untrust interfaces ae0*.0 host-inbound-traffic system-services all**
Step 7: Set security zone for the defined aggregated Ethernet (for example, ae0) link and enable traffic from all protocols to reach the interfaces in the specified zone.	[edit] user@host# **set security zones security-zone untrust interfaces *ae0*.0 host-inbound-traffic protocols all**
Step 8: Apply the LAG and LACP group configurations on the device.	[edit] user@host# **set apply-groups *WANredundancy***

Table 20: Example Configuration for VLAN-Tagged Aggregated Ethernet Interface

Configuration Steps	Commands
Step 1: Specify the number of aggregated Ethernet interfaces you want on your device. In the topology for enterprise hub WAN redundancy, the `device-count` value supported is 2. This means, you can configure two aggregated Ethernet interfaces.	`[edit]` `user@host# set groups WANredundancy chassis aggregated-devices ethernet device-count 2`
Step 2: Specify the WAN interfaces (for example, xe-0/0/0) you want to include within the aggregated Ethernet bundle and add them individually. Also, enter the interface name of the aggregate Ethernet link to which you add physical WAN member links (for example, ae0).	`[edit]` `user@host# set groups WANredundancy interfaces xe-0/0/0 gigether-options 802.3ad ae0`
Step 3: Specify the minimum number of links in the aggregated Ethernet interface (ae*x*) so that the ae link status is *up*. Only one physical link needs to be up for the bundle to be labeled *up*.	`[edit]` `user@host# set groups WANredundancy interfaces ae0 aggregated-ether-options minimum-links 1`
Step 4: Configure LACP on the defined aggregated Ethernet link (for example, ae0) as `active`. A port with `active` LACP state can start negotiating an LACP connection with the remote end by sending LACP packets, even if the device at the remote end is in `passive` state.	`[edit]` `user@host# set groups WANredundancy interfaces ae0 aggregated-ether-options lacp active`

Step 5: Configure VLAN tagging on the AE interface (for example, ae0) and assign a VLAN ID (for example, 100) to it.	`[edit]`
	`user@host# set interfaces ae0 vlan-tagging`
	`user@host# set interfaces ae0 unit 100 vlan-id 100`

Step 6: Map an aggregated Ethernet link (ae0) to the IP address of the WAN interface.	`[edit]` `user@host# set interfaces ae0 unit 100 family inet address 198.51.100.40/24`

Table 20: Example Configuration for VLAN-Tagged Aggregated Ethernet Interface *(Continued)*

Configuration Steps	Commands
Step 7: Configure a static route to CSO and associate a next hop (gateway) to it to ensure that the device is reachable to CSO.	[edit] user@host# **set routing-options static route 203.0.113.0/24 next-hop 198.51.100.1**
Step 8: Set security zone for the defined aggregated Ethernet (for example, ae0) link and enable traffic on the interface from the defined system services available in the enterprise hub device.	[edit] user@host# **set security zones security-zone *untrust interfaces ae0*.100 host-inbound-traffic system-services all**
Step 9: Set security zone for the defined aggregated Ethernet (for example, ae0) link and enable traffic from all protocols to reach the interfaces in the specified zone.	[edit] user@host# **set security zones security-zone untrust interfaces *ae0*.100 host-inbound-traffic protocols all**
Step 10: Apply the LAG and LACP group configurations on the device.	[edit] user@host# **set apply-groups *WANredundancy***

To verify if the configuration works as intended, enter the `show interfaces` command.

RELATED DOCUMENTATION

- *Predefined Configuration Templates*
- *Add Enterprise Hubs with SD-WAN Capability*

Add CSO Licenses

NOTE: This topic is applicable only to the CSO on-premises version.

To maintain a record of CSO licenses purchased by tenants or operating companies (OpCos), users with the SP Administrator role (or users with the necessary access privileges) can add the CSO license for a tenant or an OpCo from the CSO Licenses page.

To add a CSO license:

1. In Administration Portal, select **Administration** > **Licenses** > **CSO Licenses**.

   The CSO Licenses page appears.
2. Click the add (**+**) icon.

   The Add CSO License page appears.
3. Complete the configuration according to the guidelines in Table 21 on page 68.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the CSO Licenses page. A job is triggered to add the license and a confirmation message appears at the top of the page. After the job completes successfully, a confirmation message appears and the page refreshes to display the newly added license SKUs.

Table 21: Fields on the Add CSO License page

Setting	Guideline
Add License	Select whether you are adding the license for a tenant or for an operating company.
Tenant	If you are adding the license for a tenant, select the name of the tenant from the drop-down list.
Operating Company	If you are adding the license for an OpCo, select the name of the OpCo from the drop-down list.

Table 21: Fields on the Add CSO License page *(Continued)*

Setting	Guideline
Sales Order	Specify the sales order number; For example, 15563238.
SSRN	Specify the software support reference number (SSRN). This information is necessary to identify your sales order if you contact Juniper Networks for support.
Start Date	Specify the start date (in MM/DD/YYYY format) from which the license is effective.

Table 21: Fields on the Add CSO License page *(Continued)*

Setting	Guideline

License SKUs	Add one or more license SKUs:
	1. Click the add (**+**) icon.
	A row appears inline in the License SKU List grid.
	2. In the **License SKU** field, enter the SKU name.
	The SKU format is as follows: S-CSO-*Release-Type*-*License-Type*-*Device-Class*-*License-Period*, where:
	• S, which indicates that the SKU is for software.
	• CSO, which indicates that the SKU is for CSO.
	• *Release-Type*, which indicates whether the SKU is for a cloud release (C) or an on-premise release (P).
	• *License-Type*, which indicates whether the license is standard (S1) or advanced (A1)
	• *Device-Class*
	  • A denotes SRX300, SRX320, SRX340, SRX345, vSRX (2 vCPUs), NFX150 devices
	  • B denotes NFX250 (2 vCPUs), SRX550 High Memory Services Gateway (SRX550M), SRX1500, vSRX (5 vCPUs) devices.
	  • C denotes NFX250 (8 vCPUs), SRX4100, SRX4200, vSRX (9 or 17 vCPUs) devices.
	• *License-Period*, which indicates the term for the CSO license (1, 3, or 5 years).
	3. In the **Device Quantity** field, enter the maximum number of on-premise spoke sites that a tenant is authorized to add. You must enter a non-zero number to proceed.
	4. Click **✓** (check mark) to save your changes.
	The license SKU is saved and displayed in the grid.
	5. (Optional) Repeat the preceding steps if you want to add more license SKUs.

Table 21: Fields on the Add CSO License page *(Continued)*

Setting	Guideline
	You can modify a license SKU by selecting the corresponding row and clicking the edit (pencil) icon.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

RELATED DOCUMENTATION

WAN Link Redundancy in Enterprise Hubs Using Aggregated Ethernet | 61

Assign CSO Licenses to Tenants

Users with the Operating Company (OpCo) Administrator role (or users with the necessary access privileges) can assign the CSO licenses (that were previously assigned to them) to one or more tenants.

To assign a CSO license that is not yet assigned to a tenant:

1. Select **Administration > Licenses > CSO Licenses**.

   The CSO Licenses page appears.
2. Click the **Assign** link (in the Assigned column) corresponding to the license that you want to assign .Alternatively, select the license that you want to assign, and click the **Update Assignment** button.

   The Update License Assignment page appears.
3. Configure the fields according to the guidelines provided in Table 22 on page 72.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **Assign**.

   CSO validates the quantities that you assigned against the total quantity for the license:

   - If the sum of assigned quantities is greater than the total quantity, an error message is displayed. You must then modify the assigned quantities to proceed.

- If the sum of assigned quantities is less than or equal to the total quantity, a job is triggered. You are returned to the CSO Licenses page and a confirmation message is displayed on the top of the page. After the job finishes successfully, the CSO Licenses page displays the updated information in the Available and Assigned columns.

Table 22: Fields on the Assign CSO License page

Field	Description
License Information	Displays the following information for the license: • Sales Order • License SKU • Start Date
License Assignment	
Device Quantity	Displays the total quantity that can be assigned to tenants.
Assigned	Displays the quantity that is already assigned to tenants.
Tenants List	To assign the license to one or more tenants: 1. Click the **+** icon. A row is added in the grid and selected. 2. In the **Tenant** column, select the tenant to which you want to assign the license. 3. In the **Device Quantity** column, enter the quantity that you want to assign to the tenant. 4. Click **✓** (check mark) to save your changes. 5. (Optional) Click the pencil icon to modify the tenant name or the quantity and click **✓** (check mark) to save your changes. 6. (Optional) Repeat the steps if you want to assign the license to additional tenants.

See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29 for the next task.

RELATED DOCUMENTATION

CSO SD-WAN Deployment Workflow | 75

3
CHAPTER

SD-WAN Deployment

CSO SD-WAN Deployment Workflow | 75

Switch Scope or Log in as Tenant Administrator | 77

Add Provider Hub Sites | 77

Add Enterprise Hub Sites | 78

Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites | 115

Add and Install (Push) Device Licenses | 115

Install the Signature Database on Devices | 117

Add Path-Based Steering Profiles | 119

Add SLA-Based Steering Profiles | 120

Add and Deploy SD-WAN Policy Intents | 124

Add SD-WAN Breakout Profiles | 127

Add Cloud Breakout Settings | 129

Add SD-WAN Branch Sites | 133

Supported Devices for SD-WAN, and Ports and Protocols to Open | 168

Manually Activate a Site | 174

Monitor SD-WAN Sites and Devices | 175

CSO SD-WAN Deployment Workflow

CSO makes use of advanced features of the devices used in SD-WAN deployments. In order to use features such as link-switching based on application identification, or remote access IPsec VPNs on vSRX Series devices, you must purchase the required licenses. However, the underlay and overlay networks, and thus SD-WAN connectivity can be established without special licensing.

Starting in Release 6.0.0, CSO supports the following SD-WAN service types for a site:

- *Secure SD-WAN Essentials*—Provides the basic SD-WAN services. This service is ideal for small enterprises, looking for simplified management of their network and comprehensive NGFW security services at the branch sites. The SD-WAN Essentials service allows Internet traffic to breakout locally, and thus avoids the need to backhaul web traffic over costly VPN or MPLS links. This service supports features such as intent-based firewall policies, WAN link management and control, CSO-controlled routing between sites connected through the static VPN, and site to site communication through MPLS or internet links. A tenant with the SD-WAN Essentials service level can create only SD-WAN Essentials sites.

  NOTE: You can upgrade the SD-WAN service level of a tenant from SD-WAN Essentials to SD-WAN Advanced by editing the tenant information from the CSO Administration portal, provided that you have purchased the corresponding license.

- *Secure SD-WAN Advanced*—Provides the complete SD-WAN service. This service is ideal for enterprises with one or more data centers, requiring flexible topologies and dynamic application steering. You can establish site-to-site connectivity by using a hub in a hub-and-spoke topology or through static or dynamic full mesh VPN tunnels. Enterprise wide intent based SD-WAN policies and service-level agreement (SLA) measurements allow to differentiate and dynamically route traffic for different applications.

  NOTE: SD-WAN sites on CSO Release 5.4 or earlier versions are treated as SD-WAN Advanced sites. You cannot downgrade the SD-WAN service level of a tenant from SD-WAN Advanced to SD-WAN Essentials.

NOTE: Ensure that the pre-deployment tasks related to SD-WAN are carried out *before* you follow the procedure outlined in this topic. See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29.

The following tasks for configuring SD-WAN must be performed in the tenant scope in Customer Portal.

1. If you are a Tenant Administrator, log in to Customer Portal. If you are an SP Administrator (CSO on-premises) or OpCo Administrator (with appropriate permissions), switch scope to the tenant. See "Switch Scope or Log in as Tenant Administrator" on page 77.
2. Although the following optional tasks *can* be available in Customer Portal, these tasks are typically not performed in the tenant scope:
   - (Optional) Customize configuration templates. See "Configuration Templates Workflow" on page 35.
   - (Optional) Customize device templates. See "Device Templates Workflow" on page 36.
   - (Optional) Upload the latest software images to CSO. See "Device Images Workflow" on page 37.
3. For SD-WAN Advanced service, you can add one or more provider hub sites, one or more enterprise hub sites, or a combination of provider hub sites and enterprise hub sites. For SD-WAN Essentials service, you can add only one provider hub site, one enterprise hub site, or a combination of one provider hub site and one enterprise hub site (SD-WAN Essentials service does not support multihoming):
   a. Add provider hub sites. See "Add Provider Hub Sites" on page 77.
   b. Add enterprise hub sites. See "Add Enterprise Hub Sites" on page 78.

      Starting in CSO Release 6.0.0, the ZTP process is simplified to separate the device and service provisioning processes for faster deployment. You can add a site without applying a service and then edit the site to add the SD-WAN service later. See "Add Branch or Enterprise Hub Sites Without Provisioning a Service" on page 264.

      NOTE: Starting in CSO Release 6.0.0, adding a hub site is optional for an SD-WAN deployment scenario.
4. If you added enterprise hub sites, perform post-processing tasks for the enterprise hub sites. See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.
5. Add one or more SD-WAN branch sites. See "Add SD-WAN Branch Sites" on page 133. To add a site without applying a SD-WAN service, see "Add Branch or Enterprise Hub Sites Without Provisioning a Service" on page 264.
6. Perform post-processing tasks for the SD-WAN branch sites. See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.
7. (Optional) Configure a cloud spoke site. See *Adding Cloud Spoke Sites for SD-WAN Deployment* and *Provisioning a Cloud Spoke Site in AWS VPC* in the *CSO Administration Portal User Guide* (available on the CSO Documentation page).
8. Monitor SD-WAN sites and devices. See "Monitor SD-WAN Sites and Devices" on page 175.

Switch Scope or Log in as Tenant Administrator

NOTE: Because certain tasks in CSO are done in the tenant scope, that is, for a specific tenant, you need to switch scope to a tenant or log in as a Tenant Administrator user.

You can change the scope to the tenant as follows:

- If you are an SP Administrator or an OpCo Administrator user, you can switch the scope by doing one of the following:
  - On the Tenants page, click the ***Tenant-Name*** link.
  - Select the tenant name from the scope switcher list that is displayed on the CSO banner.
- If you are a Tenant Administrator user, log in to the CSO portal by accessing the CSO URL in a browser and entering your username and password.

The Welcome page appears. Click the close icon (**X**) or click **Go to Dashboard** to go the Dashboard page for Customer Portal.

Depending on whether you're deploying SD-WAN or NGFW, see "CSO SD-WAN Deployment Workflow" on page 75 or "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Add Provider Hub Sites

Before a Tenant Administrator user can add provider hub sites in Customer Portal, an SP Administrator or an OpCo Administrator user must create a point of presence (POP) and add the provider hub device (with DATA_ONLY or OAM_AND_DATA capabilities) to it from the Administration Portal. A provider hub device resides in a POP within the SP or OpCo network.

Provider hub sites are logical entities that connect branch sites, cloud spoke sites, or enterprise hub sites to provider hub devices through overlay tunnels in an SD-WAN deployment. Provider hub sites enable the tenant's sites to backhaul traffic to the provider hub devices and to the Internet.

To add one or more provider hub sites:

1. Select **Resources > Site Management**.
   The Site Management page appears.
2. Click **Add** and select **Add Provider Hub**.

The Add Provider Hub for *Tenant-Name* page appears.

3. Complete the configuration settings according to the guidelines provided in Table 23 on page 78.

> **NOTE**: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   CSO triggers a job and displays a job link. You are returned to the Site Management page. When the job is finished, the provider hub sites are listed, with the Site Status displaying Provisioned.

Table 23: Fields on the Provider Hub for <Tenant-Name> Page

Field	Description
Configuration	
Service POP	Select the POP from which you want to specify the provider hub device.
Hub Device Name	Select one or more provider hub devices from the list. (Provider hub devices with DATA_ONLY and OAM_AND_DATA capabilities are listed.) If you select two or more provider hubs, the CSO provisions the provider hub sites in the order in which you selected the provider hub devices.

See "CSO SD-WAN Deployment Workflow" on page 75 for the next task.

Add Enterprise Hub Sites

Unlike provider hubs, which can be shared by different tenants, an enterprise hub is available only to a single tenant. An enterprise hub is an SD-WAN site that is used to connect all the branch sites (spokes) in the hub and spoke topology and to break out backhaul (also called central breakout) traffic from branch sites. An enterprise hub typically has a data center department behind it; however, this is not enforced in CSO.

> **NOTE**: Starting in CSO Release 6.0.0, in SD-WAN deployments, using hubs to connect sites is optional.

For more information, see *Enterprise Hubs Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

NOTE: Before you add the enterprise hub site, check the cable connections, review the NAT and firewall ports and protocols, and check the Junos OS version of the enterprise hub device, as explained in "Supported Devices for SD-WAN, and Ports and Protocols to Open" on page 168.

To add an enterprise hub site:

1. Click **Resources > Site Management** in Customer Portal.

   The Sites page appears.
2. Click **Add**, and select **Add Enterprise Hub**.

   The Add Enterprise Hub wizard appears, displaying the General settings to be configured.
3. Configure the General settings as explained in Table 24 on page 80, and click **Next**.

   You are taken to the WAN section of the workflow.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Configure the WAN settings as explained in Table 25 on page 83, and click **Next**.

   NOTE: In Release 6.1.0, CSO moves a site to the PROVISIONED state when at least one of the WAN links obtains the IP address and is activated.

   You are taken to the LAN section of the workflow.
5. Add a LAN segment:
   a. Click the Add (**+**) icon.

      The Add LAN Segment page appears.

   b. Configure the LAN segment settings as explained in Table 26 on page 102

   c. Click **OK**.

      You are returned to the LAN section of the workflow, and the LAN segment that you added is displayed.
6. Click **Next**.

   You are taken to the Summary section of the workflow.
7. (Optional) Review the configuration in the Summary section and, if required, modify the settings.
8. Click **Finish**.

- If you entered a serial number during activation and automatic activation is enabled, the Site Activation Progress page appears. The site activation process proceeds through the tasks explained in Table 27 on page 110.

  Click **OK** to close the page.

  > **NOTE**: If you don't want to wait for the site activation to finish, you can close the page and monitor the status of the site activation from the Jobs page (**Monitor > Jobs**).
  >
  > The time taken for site activation varies depending on the device that CSO is activating.

- If you did not enter a serial number or if automatic activation is disabled, you are returned to the Sites page. CSO triggers a job and displays a confirmation message with a job link. Click the link to view the status of the job.

  After the job is finished, CSO displays a confirmation message with a job link. The status of the site changes to CREATED and an Activate Site link is displayed. You must manually activate the site to finish the process. For more information, see "Manually Activate a Site" on page 174.

After the site is activated, CSO applies the service provisioning configuration if you selected a service when adding the site. If you did not select a service, then the status of the site remains as Managed. You can edit the site later to add the service and provision the device.

> **TIP**: After you provision a site, you can modify (depending on the site status) certain parameters of the site. For more information, see *Edit Site Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Table 24: General Information (Add Enterprise Hub)

Field	Guideline
Site Information	
Site Name	Enter a unique name for the site. The name can contain alphanumeric characters and hyphens (-), and cannot exceed 32 characters.

Table 24: General Information (Add Enterprise Hub) *(Continued)*

Field	Guideline
Device Host Name	The device host name is auto-generated and uses the format *tenant-name.host-name*. You cannot change the tenant-name part in the device host name. Use alphanumeric characters and hyphen (-); the maximum length allowed is 32 characters.
Site Group	If you want the site to be part of a site group, select the site group. By default, **None** is selected, which means that the site doesn't belong to any site group.

Table 24: General Information (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Site Capabilities	**NOTE**: Device Management, enabled by default, allows you to create a site with only device management capability (without any services) and add services later. To add an SD-WAN capability for this site, choose one of the following SD-WAN service types: • Secure SD-WAN Essentials—(Available for tenants with SD-WAN Essentials or Advanced service level) Provides basic SD-WAN services. This service is ideal for small enterprises looking for managing simple WAN connectivity with comprehensive NGFW security services at the branch sites, using link-based application steering. The SD-WAN Essentials service does not support multihoming, dynamic mesh tunnels, cloud breakout profiles, SLA-based steering profiles, pool based source NAT rules, IPv6, MAP-E, or underlay BGP. **NOTE**: A tenant with the Advanced SD-WAN service level can create enterprise hubs only with the Advanced SD-WAN service. A Secure SD-WAN Advanced branch site connects only to secure SD-WAN Advanced enterprise hubs. • Secure SD-WAN Advanced—(Available for tenants with SD-WAN Advanced service level) Provides complete SD-WAN services. This service is ideal for enterprises with one or more data centers, requiring flexible topologies and dynamic application steering. You can establish site-to-site connectivity by using a hub in a hub-and-spoke topology or through static or dynamic full mesh VPN tunnels. Enterprise wide intent based SD-WAN policies and service-level agreement (SLA) measurements allow to differentiate and dynamically route traffic for different applications. This service includes Secure SD-WAN Essentials service.

Table 24: General Information (Add Enterprise Hub) *(Continued)*

Field	Guideline
Address and Contact Information	Enter the address and contact information in the fields provided. Although it is not mandatory, providing an address lets you visualize where the site is located on a geographical map on the Monitor Overview page.
Advanced Configuration	For the DNS and NTP servers, you can either use the defaults or specify DNS and NTP servers.
Domain Name Server	If needed, specify the IPv4 or IPv6, or both IPv4 and IPv6 addresses of one or more DNS servers.
NTP Server	If needed, specify the IP addresses of one or more NTP servers.
Select Timezone	Select a time zone for the site.

Table 25: Device Settings (Add Enterprise Hub)

Field	Guideline
Device Redundancy	Disabled by default. Enable this option only for dual CPEs.
Device Series	Displays SRX as the device series (family). You cannot modify this field because only certain SRX Series devices can be configured as enterprise hubs.
Device Model	Select the SRX model.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
[Device Template]	Ensure that you select the correct device template from the carousel; the template depends on the device that you are using as the enterprise hub. For example, for an SRX4100 device, select **SRX4x00 as SD-WAN CPE** (or a modified version of that template) as the device template.
Device Information	**NOTE**: If you enabled Device Redundancy, additional fields are displayed. For more information, see *Add Enterprise Hubs with SD-WAN Capability* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).
Serial Number	If you want CSO to proceed with the site activation immediately after you complete the site addition workflow, enter the serial number. If the serial number that you entered is already present in the system, CSO displays an error message. If the serial number is not present, then CSO displays a green check mark. If you want CSO to only model the site, leave this field blank. If you don't enter a serial number, you must manually activate the site later.
Device Root Password	The default root password is fetched from the ENC _ROOT_PASSWORD field in the device template. You can retain the password or change it by entering a password in plain-text format. The password is encrypted and stored on the device.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline

Zero Touch Provisioning	By default, Zero Touch Provisioning is enabled. If you want to disable ZTP, click the toggle button.
	
	NOTE: By default, this button is disabled for vSRX. You can enable this button, if the Junos OS version running on vSRX supports phone-home client.
	
	To use ZTP, ensure the following:
	
	• Device must have connectivity to CSO and Juniper phone-home server (https://redirect.juniper.net)
	
	Use telnet to verify connectivity:
	
	`telnet redirect.juniper.net:443`
	
	`telnet CSO Hostname/IP:443`
	
	If the connection is established, the device has connectivity to the phone-home server and CSO.
	
	• Required certificates for phone-home server and CSO must be present on the device.
	
	If ZTP is enabled, the Boot Image field is displayed and you must select an image that supports the Phone-Home client. During ZTP, the image on the firewall device is upgraded to the image that you select for the Boot Image.
	
	If you disable ZTP, you must copy the stage-1 configuration from CSO and commit it on the device. Use any of the following options to copy the stage-1 configuration:
	
	• Click the **Click to copy stage-1 config** link next to Prestage Device task in the Site Activation Progress page.
	
	If you close the Site Activation Progress page inadvertently, you can access the page from the Site Management page. Click the **View** link next to the status of the site under the Site Status column.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
	• On the Devices page (**Resources > Devices**), select the device and click **Stage1 Config**.
Is Cluster Already Formed?	**NOTE**: This field is available only for SRX dual CPE devices. Click the toggle button to specify whether the SRX cluster has been manually formed (Yes) or not (No).

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Cluster ID	**NOTE**: This field is available only for SRX dual CPE devices. If the SRX cluster hasn't been formed manually, specify a unique ID for the cluster. Range: 1 through 15 If you've enabled ZTP for the site, the cluster is automatically formed when the site is activated. If you've disabled ZTP, the following processes are displayed on the Site Activation Progress page (that appears after you've added the branch site): 1. After CSO models the site (that is, after the Model Site process completes successfully), click the **Click to copy pre script link**, which appears next to the Pre Script process. 2. Execute the commands as directed. After the Pre Script process completes successfully, the SRX cluster is formed and the **recovery.conf** file is saved on the cluster. In case you want to delete the site later, you'll need this file to remove the stage-1 configuration and other configurations pushed to the device by CSO. 3. Manually copy the stage-1 configuration (generated automatically by CSO) to the primary device in the cluster, and commit the configuration on the device. After the cluster is detected, CSO executes the bootstrap and provisioning processes and completes provisioning the cluster.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Auto Activate	Click the toggle button to specify whether the site activation requires an activation code or not: • **Enabled**—The site is activated automatically without an activation code. This is the default setting. • **Disabled**—The site activation proceeds only after you enter an activation code. If you choose this setting, enter the activation code (in the **Activation Code** field) that must be entered to activate the device.
Boot Image	If you want to upgrade the enterprise hub device with the latest supported Junos OS version, select the boot image from the list. The boot image is used to upgrade the device when CSO starts the zero touch provisioning (ZTP) process. If you don't specify a boot image, which is the default option (**Use Image on Device**) in the list, then the CSO skips the procedure to upgrade the device during ZTP.
Management Interface Family	Select the IP address type (IPv4 or IPv6) for the management interface. This field is displayed only if you have enabled **Zero Touch Provisioning**.
Management Connectivity **NOTE**: This section is displayed only if you disable Zero Touch Provisioning.	
Address Family	Select the IP address type (IPv4 or IPv6).
Interface Name	Enter the management interface.
Access Type	Select the access type for the underlay link. LTE, ADSL, and VDSL access types are supported only on Internet links. You cannot add LTE, ADSL, and VDSL access types to the same WAN link.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
Address assignment	DHCP is selected by default. If you want to provide a static IP address, select STATIC.
Management VLAN ID	Enter a VLAN ID for the WAN link.
PPPoE	Click the toggle button to enable authenticated address assignment for the WAN link by using PPPoE (Point-to-Point Protocol over Ethernet).
Hub Configuration **NOTE**: Hub selection is optional for both SD-WAN Advanced and Essentials sites. SD-WAN Essentials sites do not support multihoming.	
Primary Provider Hub	If you previously added provider hub sites (DATA or OAM and DATA capability) for the tenant and want to have a backup for the enterprise hub, select a provider hub site as the primary provider hub.
Secondary Provider Hub	**NOTE**: Not applicable to sites with SD-WAN Essentials service. If you previously added provider hub sites (DATA or OAM and DATA capability) for the tenant and want provider hub redundancy, select another provider hub as the secondary provider hub.
WAN Links	You can configure a maximum of four WAN links and must configure at least one WAN link.
WAN_0 (*WAN-Interface-Name*)	The first WAN link is enabled by default. Fields marked with an asterisk (*) must be configured to proceed.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Link Type	For the first WAN link, we use the default (**Internet**) for the underlay network type to ensure reachability to the redirect server.
Egress Bandwidth	Enter the maximum egress bandwidth (in megabits per second [Mbps]) that is allowed for the WAN link.
Underlay Address Families	
IPv4	By default, IPv4 address assignment is enabled for the WAN link. The WAN link requires an IPv4 address to connect to an IPv4 network.
Address Assignment Method	Displays the method of assigning an IPv4 address to the WAN link (STATIC). You cannot modify this field. You must provide the IPv4 address prefix and the gateway IPv4 address for the WAN link.
Static IP Prefix	Enter the IPv4 address prefix of the WAN link.
Gateway IP Address	Enter the IPv4 address of the gateway of the WAN service provider.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
MTU	Applicable only to IPv4 addresses. Enter the maximum transmission unit (MTU) size for the media or protocol. The supported MTU range can vary depending on the device, interface type, network topology, and other individual requirements. See also: MTU Default and Maximum Values and LTE Mini Physical Interface Modules (LTE Mini-PIM). Editing the MTU values of all the OAM-enabled WAN links of a site at the same time might result in tunnel flapping. You must ensure that at least one OAM-enabled WAN link always remains undisrupted for a site. For example, if you have a site with four WAN links (including two links that support OAM traffic), you can edit the MTU values of all the WAN links except one OAM-enabled link at the same time. After the edit is complete and the changes are saved, you can edit the site again and update the remaining WAN link. **NOTE**: If you enable the PPPoE/PPP option under a WAN link, the MTU option is displayed under the PPPoE/PPP Settings section for that link.
Public IP Address	**NOTE**: You should provide a public IP address only if the static IP prefix is a private IP address and 1:1 NAT is configured. Enter the public IPv4 address for the link, if needed.
Advanced Settings	
Advanced Settings	
Address Family (Tunnel Creation)	Displays the underlay address family (IPv4) that is used to establish the overlay tunnel.
Provider	Enter the name of the WAN link's service provider.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Cost/Month	Leave this as the default because this field is currently not used in CSO.
Enable Local Breakout	Click the toggle button to enable the WAN link to be used for local breakout. The toggle button is disabled by default, which means that the WAN link cannot be used for local breakout. Local breakout is an SD-WAN feature that enables Internet links to break out traffic directly from a site. For example, if you want to provide guests who visit your enterprise with Internet access, you can use local breakout to break out guest traffic locally from the site directly to the Internet. **NOTE**: If you enable local breakout, this only means that the WAN link *can* be used for local breakout. To enable traffic to break out from the site, you must also configure a breakout profile, reference that profile in an SD-WAN policy intent, and deploy the SD-WAN policy. If you enable local breakout, additional fields appear.
Breakout Options	This field is displayed only if local breakout is enabled for the WAN link. Select whether you want to use the WAN link for both breakout and WAN traffic (default) or only for breakout traffic.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
Autocreate Source NAT Rule	**NOTE**: Sites with Secure SD-WAN Essentials service support interface-based source NAT rules only. If you enable this options for an SD-WAN Essentials site, interface-based source NAT rules are automatically applied. If you enable this options for an SD-WAN Advanced site, you must select a source NAT rule from the Translation field. This field is displayed only if local breakout is enabled for the WAN link. When you enable local breakout on a link, this setting is enabled by default, which triggers automatic creation of source NAT rules for the site. You can click the toggle button to disable the automatic creation of source NAT rules. If you disable this field, then you must manually add a source NAT rule for local breakout and deploy the NAT policy on the site. **NOTE**: If NAT is not enforced by a separate device in your network (for example, an Internet gateway firewall), then we recommend that you enable this setting because it allows CSO to automatically create a NAT policy for the site. Table 28 on page 113 explains how source NAT rules are automatically created on the WAN link. The automatically-created source NAT rules are implicitly defined and applied to the site and is not visible on the NAT Policies page. **NOTE**: You can manually override automatically created NAT rules, by creating a NAT rule, which is placed at a higher priority than the automatically created NAT rule

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Translation	This field is displayed only if the automatic creation of source NAT rules is enabled for the WAN link, and the SD-WAN service used is Advanced. Sites with Secure SD-WAN Essentials service support interface-based source NAT rules only. Select the type of NAT to use for the traffic on the WAN link: • **Interface**—Use interface-based NAT, which is the default setting. • **Pool**—Use pool-based NAT. If you select this option, you must specify the IP addresses that are to be used for the NAT pool.
IP Addresses	For pool-based NAT, enter one or more IP addresses, subnets, or an IP address range. Separate multiple IP addresses by using commas and use a hyphen to denote a range; for example, 192.0.2.1-192.0.2.50. **NOTE**: No NAT is performed for tenant-owned public IP addresses that were added during the tenant addition workflow.
Preferred Breakout Link	if the WAN link is enabled for local breakout, click the toggle button to enable the WAN link as the most preferred breakout link. If you disable this option, then the breakout link is chosen using ECMP (equal-cost multipath) from the available breakout links.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
BGP Underlay Options	**NOTE**: Not applicable to sites with SD-WAN Essentials service. **NOTE**: BGP underlay routing is typically used by service providers, and can be configured only if local breakout is enabled for the WAN link. Click the toggle button to enable BGP underlay routing. When you enable BGP underlay routing, route advertisements to the primary Provider Edge (PE) node and, if configured, the secondary PE node occur as follows: • CSO advertises the WAN interface subnet. • If you configured pool-based translation, CSO advertises the NAT address pool. **NOTE**: If underlay BGP is enabled for a WAN link, then the routes learnt from BGP are installed for local breakout; CSO does not generate the static default route.
Primary Neighbor	Displays the IP address that you entered for the gateway for the WAN link.
Secondary Neighbor	If you want to provide PE resiliency, you can configure a secondary PE node. Enter the IP address of the secondary PE node. **NOTE**: If the primary PE node goes down, then the secondary PE is used as the next hop. When the primary PE comes back up, the route next hops are changed to the primary PE.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
eBGP Peer-AS-Number	Enter the autonomous system (AS) number for the external (EBGP) peer. **NOTE**: If the peer AS number is not configured or the peer AS number that is configured is the same as that of the CPE site, then the BGP type is assumed to be internal BGP (IBGP).
Local AS Number	Enter the local AS number for the WAN link. When you configure this parameter, the local AS number is used for eBGP peering instead of the global AS number configured for the device.
Authentication	Select the BGP route authentication method to be used: • **None**—Indicates that no authentication should be used. This is the default. • **Use MD5**—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.
Advertise Public LAN Prefixes	Click the toggle button to enable the advertisement of public LAN prefixes. This field is disabled by default. If the tenant has a public IP address pool configured and you enable the advertisement of public LAN prefixes, then for LAN segments that are created with a subnet that falls under the tenant public IP address pool, CSO advertises the LAN subnet to the BGP underlay. **NOTE**: When public LAN advertisement is enabled for the WAN link, public LAN prefixes are advertised through the BGP underlay towards MPLS or the Internet.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
Use for Fullmesh	Click the toggle button to enable the WAN link to be part of a full mesh topology. A site can have all WAN links enabled for meshing. **NOTE**: • You must enable at least one WAN link for full mesh. • Even if you enable this option, sites with SD-WAN Essentials service do not support creation or deletion of dynamic mesh tunnels based on a user-defined threshold for the number of sessions closed between two branch sites. However, an OpCo administrator or the Tenant administrator can create a static tunnel between a source site and destination site by using the CSO GUI in Customer Portal. Configure the two additional fields that appear:
Mesh Overlay Link Type	If the WAN link is enabled for full mesh, select the type of encapsulation to be used for the overlay tunnels in the full mesh topology: **NOTE**: For links with public IP addresses, we recommend that you use GRE over IPsec as the mesh overlay link type. • **GRE_IPSEC**—Use GRE over IPsec. • **GRE**—Use GRE. This option is available only for MPLS links.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
Mesh Tag	Select one or more mesh tags for the WAN link. **NOTE**: The tunnels between the enterprise hub site and the branch site are added based on matching mesh tags. So, if you want meshing to take place between a WAN link on the enterprise hub and a WAN link on the branch site, the mesh tags must be the same for both sites. For more information about mesh tags, see *Mesh Tags Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).
Use for OAM traffic	Click the toggle button to enable the use of the WAN link for Operation, Administration, and Maintenance (OAM) traffic. The WAN link is then used to establish an OAM tunnel for communication between the enterprise hub site and CSO. **NOTE**: To ensure redundancy, we recommend that you configure at least two WAN links that can be used for OAM traffic. In addition, for added management redundancy, use two links with different transport paths.
Connects to Hubs	**NOTE**: The **Connects to Hubs** field is available only if you have selected a provider hub. Click the toggle button to specify that the WAN link of the site connects to a hub. **NOTE**: • For sites with a single CPE, you must enable at least one WAN link to connect to the hub so that OAM traffic can be transmitted. • For sites with a dual CPE, you must enable at least one WAN link per device to connect to the hub so that OAM traffic can be transmitted.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
VLAN ID	Enter a VLAN ID for the WAN link. Range: 0 through 4049 (4050 to 4094 is reserved by CSO). **NOTE**: • If you are configuring more than one WAN link on the same physical interface, only one WAN link can be untagged; for the remaining WAN links, you must configure a VLAN ID. • A combination of tagged and untagged on the same physical interface is supported only for single CPE devices. To enable the configuration of WAN links as logical interfaces, you must modify the device template and configure the WAN ports as logical interfaces.
Backup Link	Select a backup link through which traffic can be routed when the primary (other) links are unavailable. You can select any link other than the default links or links that are configured exclusively for local breakout traffic. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, SLA data is not monitored for the backup link.
Default Link	Select one or more links that will be used for routing traffic in the absence of matching SD-WAN policy intents. A site can have multiple default links to the hub site. Default links are used primarily for overlay traffic but can also be used for local breakout traffic. However, a default link cannot be used exclusively for local breakout traffic. If you do not specify a default link, then equal-cost multipath (ECMP) is used to choose the link on which to route traffic.

Table 25: Device Settings (Add Enterprise Hub) ***(Continued)***

Field	Guideline
WAN_1 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
WAN_2 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
WAN_3 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
Advanced Configuration **NOTE**: Sites with SD-WAN Essentials service do not support creation or deletion of dynamic mesh tunnels based on a user-defined threshold for the number of sessions closed between two branch sites. However, an OpCo administrator or a tenant administrator can create a static tunnel between a source site and destination site by using the CSO GUI in Customer Portal.	
OAM IP Prefix	We recommend that you *do not* configure this setting (leave the IP Prefix field blank) because management connectivity is handled automatically by CSO.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
DVPN Threshold for Tunnel Creation	Specify the threshold for the number of sessions (flows) closed (in a two-minute duration) between the enterprise hub site and a destination site. When the number of sessions closed exceeds the specified threshold, a tunnel is created between the enterprise hub site and the destination site. For example, if you specify a threshold as 7, dynamic mesh tunnels are created if the number of sessions closed (in two minutes) between the enterprise hub site and destination site exceeds 7.
DVPN Threshold for Tunnel Deletion	Specify the threshold for the number of sessions closed (in a 15-minute duration) between the enterprise hub site and a destination site. When the number of sessions closed is lower than the specified threshold, the tunnel between the enterprise hub site and destination site is deleted. For example, if you specify the number of sessions closed as 5, dynamic mesh tunnels between the enterprise hub site and destination site are deleted if the number of sessions closed (in a 15-minute duration) is lesser than or equal to 5.
Additional Configuration	If you want to deploy additional configuration during the ZTP process, you can select one or more configuration templates and set the parameters for each template.

Table 25: Device Settings (Add Enterprise Hub) *(Continued)*

Field	Guideline
Configuration Templates List	For each configuration template that you select 1. Select one or more configuration templates from the list that you want to deploy on the device during ZTP. 2. Click **Set Parameters**. The Device Configurations page appears. The names and configuration parameters of the configuration templates that you selected are displayed in the Configure tab. 3. For each configuration template, enter values for the parameters. 4. (Optional) Click the **Summary** tab to view the Junos OS configuration commands that will be deployed on the device for the different configuration templates. 5. Click **Save**. You are returned to the WAN tab. The Junos OS configuration commands will be deployed on the device during the ZTP process.

Table 26: Fields on the Add LAN Segment page (Enterprise Hub)

Field	Description
Use for Overlay VPN	Enable the **Use for Overlay VPN** field to associate the LAN segment with the selected department (VRF + ZONE) for overlay traffic to other sites. Disable the **Use for Overlay VPN** field to associate the LAN segment with a security zone for underlay breakout. You must define zone-based security policies. **NOTE**: When adding a new site, this field is enabled by default and cannot be modified. However, when you add a new LAN Segment to a provisioned site from the **LAN** tab of the **Site-Name** page, you can enable or disable this option.

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) *(Continued)*

Field	Description
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length allowed is 15 characters.
CPE Port	**NOTE**: Applicable to SRX Series devices. Select the CPE port to be added in the LAN segment. When you add a new LAN Segment to a provisioned site from the LAN tab of the *Site-Name* page, you can select (or create) a LAG interface or a redundant Ethernet (reth) interface (for dual CPE cluster) to connect the SRX Series CPE devices to an EX series switch. To use the et interface on SRX4600 devices, you must create a LAG interface and configure the et interface as a member of the LAG (aggregated Ethernet or ae) interface. See *Create LAG Interface*. For an SRX4600 dual CPE cluster, you can use the et interface if it is configured as a member of the redundant Ethernet (reth) interface.
Add LAG Interface	**NOTE**: This option is available when you add a new LAN Segment to a provisioned site from the LAN tab of the *Site-Name* page. Click the link to create a LAG interface (ae interface) if you want to use it to connect the SRX Series CPE to the EX Series switch. See *Create LAG Interface* for details.
Create RETH Interface	**NOTE**: This option is available when you add a new LAN Segment to a provisioned site from the LAN tab of the *Site-Name* page. Click the link to create a reth interface for an SD-WAN site with a dual CPE cluster. See *Create a RETH Interface* for details.

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) ***(Continued)***

Field	Description
Type **NOTE**: This field is displayed only for LAN segments associated with enterprise hub sites.	Select the type of LAN segment: • Directly Connected (default)—Indicates that the LAN segment is directly connected to the site. • Dynamic Routed—Indicates that the LAN segment is not directly connected to the site and is reachable by using a dynamic route. If you select this option, you must specify the dynamic routing information.
VLAN ID	Enter the VLAN ID for the LAN segment. By default, VLAN ID is set to 1 and native VLAN is enabled for untagged traffic. Range: 1 to 4049 .
Use for Native VLAN	Enable this option to use the VLAN ID specified above for untagged traffic. The CPE interface is configured with a native-vlan-id, which has the same value as the VLAN ID.
Department	**NOTE**: This field is available only if the **Use for Overlay VPN** field is enabled. Select a department to which the LAN segment is assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See *Add a Department* for details. You can group LAN segments as departments for ease of management and for applying policies at the department-level. For LAN segments that are dynamically routed, you can assign only a data center department.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment. This address will be the default gateway for endpoints in this LAN segment. For example: 192.0.2.8/24.
Zone	**NOTE**: This field is available only if the **Use for Overlay VPN** field is disabled. Select a security zone to be associated with this LAN segment. Alternatively click **Create Zone** to create a new security zone and assign that to this LAN segment. See *Adding a Security Zone* for details.

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) ***(Continued)***

Field	Description
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP. You can enable DHCP if you want to assign IP addresses by using a DHCP server or disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE**: If you enable DHCP, additional fields appear on the page.
Additional fields related to DHCP	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on the DHCP server. Default: 1440 Range: 0 through 4,294,967,295 seconds.
Name Server	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address. **NOTE**: DNS servers are used to resolve hostnames into IP addresses.
CPE Ports	**NOTE**: Applicable to NFX150 and NFX250 devices. For sites with SD-WAN capability, the CPE Ports field is disabled and the CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) ***(Continued)***

Field	Description
Static Routing Use this section to configure static routing on the LAN segment. Provide the IP addresses of all the LAN routers connected to the CPE device and the static subnets behind these routers.	
Add LAN Router IP Prefix	
LAN Router IP	Enter the IP address of the LAN router that is connected to the CPE device.
Prefix	Enter the subnets that are connected to the LAN router.
BFD	Enable Bidirectional Forwarding Detection (BFD) to detect any failures on the static route.
Dynamic Routing	
Routing Protocol	Enable this toggle button to configure dynamic routing using the BGP or OSPF protocol.
BFD	Enable Bidirectional Forwarding Detection (BFD) to detect any failures in the LAN segment.
Protocol	Select either BGP or OSPF.
BGP Configuration **NOTE**: Starting in Release 6.1.0, CSO explicitly disables the long-lived graceful restart (LLGR) capability for BGP peering sessions with provider edge (PE) and data center or LAN routers. Disabling LLGR ensures that the CPE does not differentiate the route advertisements to the peering router irrespective of the peering router's LLGR capability. Prior to CSO Release 6.1.0, LLGR helper mode is enabled by default (implicit behavior of Junos OS) on the CPE for BGP peering towards PE router in IP VPN deployments, and data center or LAN routers in data center deployments.	

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) ***(Continued)***

Field	Description
Authentication	Select the BGP route authentication method to be used: • None—Indicates that no authentication should be used. This is the default. • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.
BGP Options	You can select the following options based on your requirements: • AS-OVERRIDE: Replaces all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer. • AS-PATH-PREPEND: Prepends one or more autonomous system (AS) numbers at the beginning of an AS path. Prepending an AS path makes a shorter AS path look longer and therefore it becomes less preferable to BGP. • AS-LOOP: Allows the local device's AS number to be added in the received AS paths. You can specify the number of times the detection of local AS is allowed in the AS path.
Loop Count	This field is displayed only if you select AS-LOOP. Enter the maximum number of times the detection of local AS is allowed in the AS path.
Peer IP Address	Enter the IP address of the LAN BGP peer.
Peer AS Number	Enter the autonomous system (AS) number of the LAN BGP peer. By default, CSO uses the AS number 64512. You can enter a different AS number.
Local AS Number	Enter the local AS number. When you configure this parameter, the local AS number is used for BGP peering instead of the global AS number configured for the CPE.

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) ***(Continued)***

Field	Description
OSPF Configuration	
OSPF Area ID	Specify the OSPF area identifier to be used for the dynamic route.
Authentication	Select the OSPF route authentication method to be used: • Password—Indicates that password-based authentication should be used. If you choose this option, you must specify the password. (This is the default). • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key. • None—Indicates that no authentication should be used.
Password	Enter the password to be used to verify the authenticity of OSPF packets.
Confirm Password	Retype the password for confirmation purposes.
MD5 Auth Key ID	If you specified that MD5 should be used for authentication, enter the OSPF MD5 authentication key ID. Range: 1 through 255.
Auth Key	If you specified that MD5 should be used for authentication, enter an MD5 authentication key, which is used to verify the authenticity of OSPF packets.
Route Advertisement Control	

Table 26: Fields on the Add LAN Segment page (Enterprise Hub) ***(Continued)***

Field	Description
LAN Route(s) to Overlay	When this option is enabled, LAN routes are advertised to the remote CPEs. By default, this option is enabled. Starting in CSO Release 6.2.0, you can configure export policies in conjunction with the **LAN Route(s) to Overlay** option for more granular control over routes that are advertised to the overlay network. For example, when the **LAN Route(s) to Overlay** option is enabled, you can configure policies to prevent specific routes from being advertised. Similarly, when the **LAN Route(s) to Overlay** option is disabled, you can configure policies to allow only specific routes to be advertised.
Overlay Route(s) to LAN	This option is displayed only if you enable the Routing Protocol toggle button. By default, this option is disabled. Enable this option to advertise the remote CPE routes received in a department to the LAN router. **NOTE**: In CSO Release 6.0.0 and earlier releases, this option is called Advertise LAN Prefix and is applicable only for data center departments. Starting in CSO Release 6.2.0, you can use the following policies for granular control of the route advertisements: • Import policies for granular control of the routes that a CPE device accepts from the list of routes advertised by the LAN router. • Export policies for granular control of the routes that a CPE device advertises to the LAN router.
Static/Aggr Routes to Overlay	Enable this option to allow advertisement of static or aggregate routes to the overlay network. • If a large number of LAN routes are present, then you can disable the **LAN Route(s) to Overlay** option and use this option to advertise aggregate routes. • If you want to advertise additional routes, then you can enable the **LAN Route(s) to Overlay** option and use this option to advertise additional static routes.

Table 27: Site Activation Tasks and Troubleshooting

Activation Task	Troubleshooting
Model Site—CSO first models the site to begin the activation process. If you didn't enter a serial number or disabled automatic activation, you must manually activate the site as explained in "Manually Activate a Site" on page 174.	
Prestage Device—Depending on the type of device used, you might need to copy the configuration that is generated by CSO and commit the configuration on the device. For such devices, CSO can move to the next step (detecting the device) only after the configuration is committed successfully on the device. 1. On the Devices page (**Resources > Devices**), select the device and click **Stage1 Config**. The configuration to be copied appears in a separate page. 2. Click **Copy** to copy the configuration to the clipboard 3. Log in to the device by using SSH and enter Junos OS configuration mode. 4. Paste the configuration that you copied and commit the configuration.	This step typically goes through without problems. However, if you encounter a problem, log in to the device (using a console or a management interface), access the CLI, and verify that the stage-1 configuration was committed on the device.

Table 27: Site Activation Tasks and Troubleshooting *(Continued)*

Activation Task	Troubleshooting

Detect Device—The device reaches out to CSO, and communication with CSO is established.
 	
This task typically takes a few minutes. If the status shows as Pending after about 10 minutes, try the troubleshooting steps.
	If the device is not detected:
	
	1. Check that the correct interfaces on the device are connected.
	
	2. Log in to the device, and access the CLI.
	
	3. Check the system time that is configured on the device by executing the `show system uptime` command, and ensure that the system time is accurate. A mismatch in time might mean that the device is unable to connect to the redirect server.
	
	4. **NOTE**: This step is applicable only for branch sites.
	
	Execute the `show interfaces terse` command.
	
	In the command output, verify whether the device received a DHCP IP address. If the device did not receive an IP address, try to reconnect.
	
	5. If the device has a valid IP address, then verify that the device can reach the Internet by using the `ping` command. For example, `ping www.juniper.net`.
	
	If the `ping` command executes successfully, this means that the device can reach the Internet, and DNS resolution is working.
	
	6. Verify whether the device has the permissions required for outgoing connections on port 443 by executing the `telnet redirect.juniper.net 443` command.
	
	If the device has the required permissions, you should see an output similar to the following:
	
	`Trying 192.0.2.155...`
	`Connected to telnet-host.example.com.`
	`Escape character is '^]'.`

Table 27: Site Activation Tasks and Troubleshooting *(Continued)*

Activation Task	Troubleshooting
Bootstrap Device—This task comprises the following sub-tasks: 1. A secure OAM tunnel (using IPsec) from the device to the OAM hub is established. 2. An outbound SSH connection from the device is established with CSO. 3. An Internal BGP (iBGP) peering between the device and the virtual route reflector (VRR) is established. 4. The device sends a Bootstrap Complete message to CSO, which CSO receives and marks the bootstrap as completed. CSO applies the pre-script and stage-1 (includes the device configuration) configuration. This task typically takes a few minutes to finish. If the status shows as Pending after about 10 minutes, try the troubleshooting steps.	If the bootstrap device task does not finish successfully: 1. Verify whether the stage-1 configuration was deployed on the device by executing the `show configuration | display set | match outbound-ssh | match 7804` command. If the resulting output is similar to the following sample output, it means that the stage-1 configuration was deployed successfully. `set system services outbound-ssh client CSO-xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx 192.0.2.100 port 7804` 2. Check if the secure OAM tunnels are up by executing the following commands: • `show security ike sa` command. If the State field in the output doesn't display `UP`, it means that port 500 is blocked. Ensure that you open 500 and retry the activation job (from the Jobs page). • `show security ipsec sa` command. If the State field in the output doesn't display `UP`, it means that port 4500 is blocked. Open port 4500, and retry the activation job (from the Jobs page). 3. Verify whether the device has established BGP peering with the VRR by executing the `show bgp summary` command. If the State field in the output displays `Establ`, it means that BGP peering is established successfully. 4. Verify whether the secure OAM session is established by executing the `show security flow session destination-port 7804` command.

Table 27: Site Activation Tasks and Troubleshooting *(Continued)*

Activation Task	Troubleshooting

	If the resulting output is similar to the following output, it means that the secure OAM session was established successfully.
	
	`Session ID: 430000098, Policy name: default-policy-00/2, Timeout: 1778, Valid`
	`  In: 192.0.2.10/15190 --> 192.0.2.20/23;tcp, If: ge-7/1/0.0, Pkts: 109, Bytes: 5874, CP Session ID: 430000093`
	`  Out: 192.0.2.20/23 --> 192.0.2.10/15190;tcp, If: ge-7/1/1.0, Pkts: 64, Bytes: 4015, CP Session ID: 430000093`
	`Total sessions: 1`

Manage Device—After CSO applies the configuration on the device, the status of the device changes to Managed. If the status is showing Pending after about 10 minutes, try the troubleshooting steps.	Go to the Jobs page (**Monitor > Jobs**), search for the ZTP job, and check the status. Click the *job-name* link to view the tasks associated with the job and their status. You can drill down further by clicking the *task-name* link. If the status of the job or task is **In Progress**, wait until the job or task finishes. If the job failed, you can retry the job by selecting the job, and clicking the **Retry Job** button.

Table 28: Automatic Creation of Source NAT Rules

Autocreate Source NAT Rule	Translation	NAT Rules Creation
Disabled	Not applicable (No NAT)	None.

Table 28: Automatic Creation of Source NAT Rules *(Continued)*

Autocreate Source NAT Rule	Translation	NAT Rules Creation

Enabled	Interface-Based (Default)—CSO creates interface-based NAT rules.	Source NAT rules are automatically created, with each rule from a department zone to the WAN interface, with a translation of type interface. Each pair of [zone - interface] represents a rule-set.
		
		For example, the following department zone to (WAN link) W1 interface rule-set might be created:
		
		`Dept-Zone1 --> W1: Translation=Interface`
		`Dept-Zone2 --> W1: Translation=Interface`
		`Dept-Zone3 --> W1: Translation=Interface`
		
		When traffic from a branch site breaks out at an enterprise hub, a source NAT rule is automatically created at the enterprise hub from the department routing group (also referred to as VRF group) to the WAN interface.
		
		`Dept-vrf-group --> W1: Translation=Interface`

Enabled	Pool-Based—CSO automatically creates pool-based NAT rules (Not applicable to sites with SD-WAN Essentials service).	Source NAT rules are automatically created, with each rule from a department zone to the WAN NAT pool with a translation of type pool.
		
		For example, a source NAT rule from department zone to NAT pool might be created:
		
		`Dept-Zone1 --> W1 : Translation=Pool-1`
		`Dept-Zone2 --> W1 : Translation=Pool-1`
		
		When traffic from a branch site breaks out at an enterprise hub, a source NAT rule is automatically created at the enterprise hub from the department routing group to the WAN pool.
		
		`Dept-vrf-group --> W1: Translation=Pool`

After the site is provisioned, you must perform "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.

Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites

After the enterprise hub or the SD-WAN on-premise spoke site is provisioned successfully, perform the following post-provisioning tasks:

1. Upload and install device licenses. See "Add and Install (Push) Device Licenses" on page 115.
2. Install the signature database. See "Install the Signature Database on Devices" on page 117.
3. Add and deploy a firewall policy. See "Add and Deploy Firewall Policies" on page 238.
4. If you want to perform path-based or SLA-based (dynamic) steering of the traffic:
   a. Add a path-based steering profile. See "Add Path-Based Steering Profiles" on page 119.
   b. (Not applicable to sites with SD-WAN Essential service) Add an SLA-based steering profile. See "Add SLA-Based Steering Profiles" on page 120.
   c. Add an SD-WAN policy intent that references the path-based or SLA-based steering profile and deploy the SD-WAN policy. See "Add and Deploy SD-WAN Policy Intents" on page 124.

   NOTE: The SD-WAN Essential service does not support department-level policy intents or SLA-based steering profiles.
5. To break out traffic from the site:
   a. Add a local breakout, central breakout, or a cloud breakout profile. See "Add SD-WAN Breakout Profiles" on page 127.
   b. For cloud-breakout, add cloud breakout settings and assign the settings to the site. See "Add Cloud Breakout Settings" on page 129.
   c. Add an SD-WAN policy intent that references the breakout profile and deploy the SD-WAN policy. See "Add and Deploy SD-WAN Policy Intents" on page 124.

See "CSO SD-WAN Deployment Workflow" on page 75 for the next task.

Add and Install (Push) Device Licenses

After a site is successfully provisioned, you must add the required device licenses into CSO, and then install the licenses on the device (that is associated with the site).

Starting in CSO Release 6.2.0, tenants can use a single license (golden license) to push the device licenses to all the physical SRX devices in its network.

To add and install device licenses:

1. Add the device license file:

   NOTE: Device license files can be added by the SP Administrator or OpCo Administrator (in Administration Portal) or by the Tenant Administrator (in Customer Portal).

   a. Select **Administration** > **Device Licenses**.

      The Device License Files page appears.

   b. Click the Add (**+**) icon.

      The Add License page appears.

   c. Configure the parameters as explained in Table 29 on page 117.

      NOTE: Fields marked with an asterisk (*) are mandatory.

   d. Click **OK**.

      CSO parses the license file, and verifies whether the license file format is valid. If the format is valid, CSO adds the license file, and returns you to the Device License Files page.

   If needed, you can add additional device license files.

2. Install (push) the license to the device:

   a. Select the device license file that you want to push to the device.

   b. Click **Push License**, and select **Push**.

      The Push License page appears, displaying the sites and devices to which the license can be pushed.

   c. Select one or more devices to which you want to push the license, and click **OK**.

      CSO initiates a job to push the license to the device and displays a confirmation message. After the job completes successfully, the license is pushed to the device. You can view the status of the job on the Jobs page (**Monitor > Jobs**).

Table 29: Add License Page Settings

Field	Guideline
License File	Click **Browse** to select the license file, and click **Open**. The License File field displays the license file that you selected. **NOTE**: A license file can contain only one license key.
Tenant	**NOTE**: This field is displayed only if you're adding a license in Administration Portal. • If you're an SP Administrator user, this field displays Global, which means that the license file can be used by all tenants added by the SP Administrator. • If you're an OpCo Administrator user, select the tenant with which you want to associate the license file. When a device is activated during ZTP, the license is downloaded to the device. Licenses associated with a tenant can be applied only to devices that belong to that tenant.
Description	Enter a description for the license file.

The next step after installing device licenses is to install the signature database on the device. See "Install the Signature Database on Devices" on page 117.

Install the Signature Database on Devices

Users with the Tenant Administrator role can install the active signature database on one or more devices. Signatures must be present on the device for application firewall or intrusion prevention system (IPS) features to be used. If you do not install the signature database on a device, the deployment of IPS profiles or application firewall will fail.

NOTE: Before you install the signature database on the device, ensure that the IPS license is installed on the device. If the IPS license is not installed, only the application signatures will be installed when the signature database installation is triggered.

You can install the signature database on NFX150, NFX250, SRX Series, and vSRX devices.

To install the active signature database:

1. Select **Administration > Signature Database**.

   The Signature Database page appears.

2. Click **Install Signatures**.

   The Install Signatures page appears, displaying the signature database version and the devices on which you can install the signature database.

3. Select the check boxes corresponding to the devices on which you want to install the signature database.

   You can also search for, filter, or sort the devices that are displayed.

4. From the **Type** field:

   - Select **Run now** to trigger the installation of the signature database immediately.
   - Select **Schedule at a later time** to install the signature database later, and specify a date and time at which you want the installation to be triggered.

5. Click **OK**.

   - If you specified that the database must be installed immediately, a job is triggered. In the Job Tasks page that appears, the tasks associated with the signature database installation are displayed. Click **OK** to exit and return to the Signature Database page.
   - If you specified that the database must be installed later, a job is triggered and you are returned to the Signature Database page. A confirmation message (with the job ID) is displayed at the top of the page.

NOTE: In addition to using the predefined signatures present in the database, you can add and use the following:

- Customized application signatures and signature groups on the Application Signatures page (**Configuration > Application Signatures** in Administration Portal or **Configuration > Shared Objects > Application Signatures** in Customer Portal).

  For more information, see the *About the Application Signatures Page* topics in the *CSO Administration Portal User Guide* and the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

- Customized IPS signatures, static groups, and dynamic groups on the IPS Signatures page (**Configuration > IPS > IPS Signatures** in Customer Portal).

  For more information, see *About the IPS Signatures Page* in the *CSO Customer Portal User Guide*.

Depending on whether you're deploying SD-WAN or NGFW, see "CSO SD-WAN Deployment Workflow" on page 75 or "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Add Path-Based Steering Profiles

A path-based steering profile is an SD-WAN traffic steering profile in which you specify only a path preference for the SD-WAN traffic, and, optionally, rate limiting parameters. You cannot configure service-level agreement (SLA) parameters or path failover criteria in a path-based steering profile.

To add a path-based steering profile:

1. Select **Configuration > Path Based Steering Profiles** (in Administration Portal) or **Configuration > SD-WAN > Path Based Steering Profiles** (in Customer Portal).

   The Path-Based Steering Profiles page appears.
2. Click the add (**+**) icon.

   The Add Path Profile page appears.
3. Enter the path-based steering profile information according to the guidelines provided in Table 30 on page 119.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the Path-Based Steering Profiles page and a confirmation message indicating that the path-based steering profile was added is displayed. The page refreshes to display the path-based steering profile that you added.

   NOTE: After you add a path-based steering profile, you must add an SD-WAN policy intent that references the path-based steering profile, so that profile parameters are applied to SD-WAN traffic.

Table 30: Add Path Profile Settings

Field	Guideline
Name	Enter a unique string that can contain alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Traffic Type Profile	Select a traffic type profile to apply the class-of-service configuration and priority to the path-based steering profile. You can select only traffic type profiles that are enabled.

Table 30: Add Path Profile Settings *(Continued)*

Field	Guideline
Path Preference	Select the path (**Internet** or **MPLS**) that the SD-WAN traffic should take. The path here refers to the overlay path, which means that traffic will take the overlay tunnel on the WAN link whose type is the same as the path specified.
Advanced Configuration	You can optionally configure parameters for rate limiting the SD-WAN traffic for applications associated with the path-based steering profile.
Maximum Upstream Rate	Enter the maximum upstream rate (in Kbps) for all applications associated with the path-based steering profile.
Maximum Upstream Burst Size	Enter the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the upstream rate limit for short periods.
Maximum Downstream Rate	Enter the maximum downstream rate (in Kbps) for all applications associated with the path-based steering profile.
Maximum Downstream Burst Size	Enter the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the downstream rate limit for short periods.
Loss Priority	Select a loss priority based on which packets are dropped or retained when network congestion occurs. Packet drops are most likely when the loss priority is High and least likely when the loss priority is Low.

See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.

Add SLA-Based Steering Profiles

An SLA-based steering profile is an SD-WAN traffic steering profile with service-level agreement (SLA) parameters and path failover criteria for WAN traffic, and, optionally, rate limiting parameters. CSO provides predefined SLA-based steering profiles that are tuned for specific application categories and traffic types, which you can use in an SD-WAN policy. You can also add customized SLA-based steering profiles in CSO.

NOTE: SD-WAN Essentials service does not support SLA-based steering profiles.

To add a customized SLA-based steering profile:

1. Select **Configuration > SLA Based Steering Profiles** (in Administration Portal) or **Configuration > SD-WAN > SLA Based Profiles** (in Customer Portal).

   The SLA Profiles page appears.
2. Click the add icon (**+**).

   The Add SLA Profile page appears.
3. Configure the parameters according to the guidelines provided in Table 31 on page 121.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the SLA-Based Steering Profiles page and a confirmation message indicating that the SLA-based steering profile was added is displayed. The page refreshes to display the SLA-based steering profile that you added.

   NOTE: After you add an SLA-based steering profile, you must add an SD-WAN policy intent that references the SLA-based steering profile, so that profile parameters are applied to SD-WAN traffic.

Table 31: Add SLA Profile Settings

Field	Guideline
Name	Enter the name of the SLA-based steering profile, which is a unique string that can contain alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Traffic Type Profile	Select a traffic type profile to apply the class-of-service configuration and priority to the SLA-based steering profile. You can select only traffic type profiles that are enabled.

Table 31: Add SLA Profile Settings *(Continued)*

Field	Guideline
SLA Configuration	Select whether you want to use predefined SLA thresholds or specify customized SLA parameters: • **Use Recommended**—Select this option if you want to specify an SLA threshold, which uses predefined SLA parameters for the steering profile. • **Enter Custom**—Select this option if you want to specify customized SLA parameters for the steering profile.
SLA Threshold	If you specified that an SLA threshold should be used, use the slider to select the predefined threshold to use for the SLA parameters in the steering profile: • **Liberal**: Use liberal (relaxed) SLA parameters for the steering profile. This is the default setting. • **Baseline**—Use the baseline SLA parameters for the steering profile. • **Conservative**—Use conservative (strict) SLA parameters for the steering profile.
Packet Loss	If you specified that customized SLA parameters should be used, enter the target packet loss percentage for the steering profile. If the percentage of data packets dropped by the network (to manage congestion) exceeds the specified target packet loss percentage, then an SLA violation is reported.
RTT	If you specified that customized SLA parameters should be used, enter the target round-trip time (RTT), in milliseconds, for the steering profile. If the RTT for a packet exceeds the specified target RTT, then an SLA violation is reported.
Jitter	If you specified that customized SLA parameters should be used, enter the target jitter, in milliseconds, for the steering profile. If the jitter for data packets exceeds the specified target jitter, then an SLA violation is reported.
Path Selection Criteria	

Table 31: Add SLA Profile Settings *(Continued)*

Field	Guideline
Path Preference	Displays Any, which means that either an MPLS or an Internet link can be used because the link is selected based on the SLA criteria. You cannot modify this field.
Path Failover Criteria	Select the failover criterion to use, which determines how links are switched when the active link fails to meet the SLA specified in the steering profile. When path failover occurs, the traffic is routed to a link that meets the SLA. Choose one of the following options: • **Does not meet one or more SLA parameters**—The link is switched if any of the SLA parameters is violated. • **Does not meet all SLA parameters**—The link is switched only when all the SLA parameters are violated.
Advanced Configuration	You can optionally configure parameters for rate limiting the SD-WAN traffic for applications associated with the SLA-based steering profile.
Maximum Upstream Rate	Enter the maximum upstream rate (in Kbps) for all applications associated with the SLA-based steering profile.
Maximum Upstream Burst Size	Enter the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the upstream rate limit for short periods.
Maximum Downstream Rate	Enter the maximum downstream rate (in Kbps) for all applications associated with the SLA-based steering profile.
Maximum Downstream Burst Size	Enter the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the downstream rate limit for short periods.
Loss Priority	Select a loss priority based on which packets are dropped or retained when network congestion occurs. Packet drops are most likely when the loss priority is High and least likely when the loss priority is Low. If you select **None**, which is the default, the loss priority is configured as Low.

Table 31: Add SLA Profile Settings *(Continued)*

Field	Guideline
SLA Sampling	You must specify the parameters related to SLA sampling, as explained in the following fields.
Session Sampling %	Enter the percentage of flow sessions for which the application quality of experience (AppQoE) probes should be sent. If you specify a 3% sampling rate, then 3 probes are sent for every 100 flow sessions. For information about AppQoE, see *Application Quality of Experience Overview* in the *CSO Administration Portal User Guide* (available on the CSO Documentation page).
SLA-violation-count	Enter the number of times that the SLA can be violated before CSO considers that the AppQoE is affected. For example, if you specify a violation count of 10, then CSO considers the AppQoE to be affected if the SLA is violated 10 times.
Sampling Period	Enter the period (in seconds) for which the sampling packets are sent.
Switch-cool-Off-period	Enter the period (in seconds) for which CSO should stop the AppQoE probes, when a link switch takes place. For example, if you specify a period of 60 seconds, then CSO will wait for 60 seconds after a link switch before restarting the sending of AppQoE probes. The default period is 120 seconds. This setting is used to prevent frequent switching of links.

See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.

Add and Deploy SD-WAN Policy Intents

By default, CSO provides predefined SD-WAN policy intents, which reference predefined SLA-based steering profiles. These policy intents are applicable to all sites and can be deployed to the branch sites or enterprise hub sites of a tenant. You can also modify these intents based on your network requirements or delete them if you don’t want to use the predefined intents.

CSO also allows you to add customized intents that reference path-based steering profiles, SLA-based steering profiles, and SD-WAN breakout profiles, and then deploy the intents to the branch sites or enterprise hub sites of a tenant.

For example, if you enable local breakout on a WAN link of an enterprise hub site, you can add an SD-WAN breakout profile, reference that breakout profile in an SD-WAN policy intent, and then deploy the SD-WAN policy intent, which ensures that traffic will break out locally from the WAN link that you configured for local breakout.

NOTE: SD-WAN Essentials service does not support department level policy intents or SLA-based steering profiles.

To add and deploy an SD-WAN policy intent:

1. Add the SD-WAN policy intent:
   a. Select **Configuration > SD-WAN > SD-WAN Policy**.
      The SD-WAN Policy page appears.
   b. Click the add icon (**+**).
      The parameters for an SD-WAN policy intent appear inline on the SD-WAN Policy page.
   c. Enter the policy intent information according to the guidelines provided in Table 32 on page 126.
   d. Click **Save**.
      The SD-WAN policy intent is added, and a confirmation message is displayed. The Undeployed field is incremented by one, indicating that the policy intent must be deployed.
2. Deploy the SD-WAN policy intent:
   a. Click the **Deploy** button.
      The Deploy page appears.
   b. From the **Choose Deployment Time** field:
      - Select **Run now** to deploy the policy immediately.
      - Select **Schedule at a later time** to schedule the deployment for later.

        If you schedule the deployment for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) at which you want the deployment to be triggered. You specify the time in the local time zone of the client from which you access the CSO GUI.

   You are returned to the SD-WAN Policy page, and a job to deploy the policy is triggered. You can check the status of the deployment on the Jobs page (**Monitor > Jobs**).

   After the SD-WAN policy is successfully deployed:

   - For intents that reference steering profiles, the profile parameters are applied to the traffic that matches the intent.

- For intents that reference breakout profiles, traffic can break out directly from the sites or departments that match the intent.

Table 32: SD-WAN Policy Intent Settings

Field	Guideline
Name	Enter a name for the policy intent, or use the name generated by CSO.
Description	Enter a description for the policy intent.
Source	Select one or more of the following source endpoints: • Site • Site Group • Department By default, **All Sites** is selected as the source endpoint, which means that the SD-WAN policy intent is applicable to all spoke and enterprise hub sites in the tenant.
Application	For an SD-WAN policy intent that references a steering profile, select the applications or application groups to which you want the steering profile parameters to be applied to the traffic. For an SD-WAN policy intent that references a breakout profile, select the applications or application groups for which you want to break out traffic. **NOTE**: The option **Any**, which means that this policy intent is applicable to all applications, can be used only if you specify a breakout profile. For example, if you want all guest traffic to break out to the Internet through the underlay, you can select the guest department as the source and select Any as the application.
Traffic Steering Profile	Click inside the text box, and select one of the following depending on the traffic for which you plan to apply the intent: • SLA-based steering profile • Path-based steering profile • Breakout profile

See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.

Add SD-WAN Breakout Profiles

Read the "Understand Breakout in CSO" on page 273 topic for high-level overview of breakout in Contrail Service Orchestration (CSO).

You can add SD-WAN breakout profiles either in the Administration Portal (SP Administrator or OpCo Administrator users only) or in Customer Portal. By default, for SD-WAN Advanced deployments, CSO adds a cloud breakout profile. The availability of the profile depends on the role of the user adding the breakout profile:

- Breakout profiles added by the SP Administrator are available to all tenants, OpCos, and the OpCo's tenants.
- Breakout profiles added by the OpCo Administrator are available only to that OpCo and the OpCo's tenants.
- Breakout profiles added in the tenant scope or by the Tenant Administrator are available only for the tenant.

To add an SD-WAN breakout profile:

1. Select **Configuration > SD-WAN Breakout Profiles** (in Administration Portal) or **Configuration > SD-WAN > Breakout Profiles** (in Customer Portal).

   The Breakout Profiles page appears.
2. Click the Add (**+**) icon. (If you're accessing the page from Customer Portal, ensure that you are on the Breakout Profiles tab).

   The Add Breakout Profile page appears.
3. Complete the configuration according to the guidelines provided in Table 33 on page 128.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the Breakout Profiles page, and a message confirming that the breakout profile was added is displayed. The page refreshes to display the breakout profile that you added.

After you add an SD-WAN breakout profile, you must add an SD-WAN policy intent and then deploy the SD-WAN policy to ensure that traffic breaks out locally from the WAN link that you configured for local breakout.

Table 33: Fields on the Add Breakout Profile Page

Field	Guideline
Type	Select the type of breakout profile that you want to add: • **Local Breakout (Underlay)**—Select this option if you want traffic to break out locally (on the underlay) from the site. • **Backhaul**—Select this option if you want traffic to break out through a hub or a enterprise hub (if configured). • **Local Breakout (Cloud)**—Select this option if you want to break out traffic through a cloud-based security platform. Currently, Zscaler is the only cloud-based security platform supported. **NOTE**: SD-WAN Essentials sites do not support cloud breakout profiles.
Name	Enter a unique name for the breakout profile. You can use alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Description	Enter a description for the breakout profile.
Traffic Type Profile	Select a traffic type profile to apply class of service (CoS) parameters to the breakout traffic.
Preferred Path	Select the preferred path (MPLS, Internet, or Any) to be used for breaking out the traffic. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic. If you specify that any path can be used, then there is no preference and all WAN links that are enabled for breakout are used in a load-balancing mode.
Advanced Configuration	You can optionally configure parameters for rate limiting the breakout traffic for cacheable applications.
Upstream Rate	Enter the maximum upstream rate (in Kbps) for all cacheable applications associated with the breakout profile.

Table 33: Fields on the Add Breakout Profile Page *(Continued)*

Field	Guideline
Upstream Burst Size	Enter the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the upstream rate limit for short periods.
Downstream Rate	Enter the maximum downstream rate (in Kbps) for all cacheable applications associated with the breakout profile.
Downstream Burst Size	Enter the maximum size (in bytes) of a steady stream of traffic sent at average rates that exceed the downstream rate limit for short periods.
Loss Priority	Select a loss priority based on which packets are dropped or retained when network congestion occurs. Packet drops are most likely when the loss priority is High and least likely when the loss priority is Low.

See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.

Add Cloud Breakout Settings

If you want to break out traffic to a cloud-based security platform, then you must add settings for cloud breakout and assign the settings to one or more sites. You assign cloud breakout settings to sites to enable the provisioning of the tunnels from the sites to the cloud breakout node. For traffic to break out from the site, you must reference the cloud breakout profile in an SD-WAN policy intent and then deploy the SD-WAN policy.

To add cloud breakout settings:

1. Select **Configuration > SD-WAN > Breakout Profiles**.

   The Breakout Profiles page appears.
2. On the **Cloud Breakout Settings** tab, click the add icon (**+**).

   The Add Cloud Breakout Settings page appears.
3. Complete the configuration according to the guidelines provided in Table 34 on page 130.

NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

   You are returned to the Breakout Profiles page (Cloud Breakout Settings tab) and a confirmation message indicating that the breakout settings are added is displayed.

Table 34: Fields on the Add Cloud Breakout Settings Page

Field	Description
Name	Enter a unique name for the cloud breakout settings. You can use alphanumeric characters and hyphens (-); the maximum length is 15 characters.
Tunnel Type	Select the type of overlay tunnel (IPsec or GRE) used to break out the traffic to the cloud breakout node.
IPsec Configuration Parameters	
Domain Name	Displays the domain name that is used to generate the fully qualified domain name (FQDN) for SD-WAN policies. The FQDN is used by cloud security providers to identify the IPsec tunnels. The domain name is populated based on the customer domain name that was provided when the tenant was onboarded. You can modify the domain name.
Phase 1	In Phase 1, the SD-WAN branch site and the cloud breakout node establish a secure tunnel to negotiate the IPsec security associations (SAs).
Encryption Type	Select an encryption type for IPsec proposals: • **AES-256-CBC (default)**—Advanced Encryption Standard (AES) 256-bit encryption algorithm in Cipher Block Chaining (CBC) mode. • **AES-192-CBC**—AES 192-bit encryption algorithm. • **AES-128-CBC**—AES 128-bit encryption algorithm. • **3DES-CBC**—Triple Data Encryption Algorithm (3DES) in CBC mode. Has a block size of 24 bytes; the key size is 192 bits long.

Table 34: Fields on the Add Cloud Breakout Settings Page *(Continued)*

Field	Description
Authentication Type	Select an IPsec authentication algorithm for security association: • **SHA-256 (default)**—Secure Hash Algorithm (SHA) that converts a text of any length into a string of 256 bits. • **SHA-384**—Produces a 384-bit string. • **SHA1**—Produces a 160-bit string.
DH Group	Specify the Diffie-Hellman (DH) group to match the IPsec encryption algorithm: • **GROUP2 (default)**—1024-bit Modular Exponential (MODP) algorithm. • **GROUP5**—1536-bit MODP algorithm. • **GROUP14**—2048-bit MODP algorithm.
Phase 2	In Phase 2, the SD-WAN branch site and the cloud breakout node negotiate the IPsec security associations for encrypting and authenticating the exchange of data.
Encryption Type	Select an encryption type for IPsec proposals. • **NULL**—No encryption. This is the default. • **AES-256-CBC**—AES 256-bit encryption algorithm. • **AES-192-CBC**—AES 192-bit encryption algorithm. • **AES-128-CBC**—AES 128-bit encryption algorithm.
Authentication Type	Select an IPsec authentication algorithm for security association. • **HMAC-MD5-96**—Produces a 128-bit digest. This is the default. • **HMAC-SHA-256-128**—Produces a 256-bit digest, truncated to 128 bits. • **HMAC-SHA1-96**—Produces a 160-bit digest.

Table 34: Fields on the Add Cloud Breakout Settings Page *(Continued)*

Field	Description
Protocol	This setting is enabled only if you select a non-null encryption type. Select the type of protocol to be used for authentication: • **ESP**—Encapsulating Security Payload (ESP) protocol. This is the default. • **AH**—Authentication Header (AH) Protocol.
Primary Gateway	Specify the configuration parameters for the primary cloud breakout node.
Link Type	Select the preferred type of WAN link (MPLS or Internet) to be used for breaking out the traffic to the primary cloud breakout node. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic.
IP Address/ Hostname	Enter the IPv4 address or hostname of the primary cloud breakout node. Currently, Zscaler is the only cloud-based security platform supported. CSO validates the IP address or hostname, and if the IP address or host name is not reachable, a `Host Unreachable` message is displayed.
Preshared Key	Enter the preshared key (provided by Zscaler) to be used for Internet Key Exchange (IKE) authentication with the primary cloud breakout node. The key that you enter is masked by default but you can click the eye icon to unmask the key.
Confirm Preshared Key	Re-enter the preshared key for confirmation.
Secondary Gateway	Specify the configuration parameters for the primary cloud breakout node.
Link Type	Select the preferred type of WAN link (MPLS or Internet) to be used for breaking out the traffic to the secondary cloud breakout node. If a WAN link type that matches the preferred path is enabled for breakout, then that WAN link type is used for breakout traffic.

Table 34: Fields on the Add Cloud Breakout Settings Page *(Continued)*

Field	Description
IP Address/ Hostname	Enter the IPv4 address or hostname of the secondary cloud breakout node. CSO validates the IP address or hostname, and if the IP address or host name is not reachable, a `Host Unreachable` message is displayed.
Preshared Key	Enter the preshared key (provided by Zscaler) to be used for Internet Key Exchange (IKE) authentication with the secondary cloud breakout node. The key that you enter is masked by default but you can click the eye icon to unmask the key.
Confirm Preshared Key	Reenter the preshared key for confirmation.

After you add cloud breakout settings, you can assign the settings to one or more sites, which provisions the overlay tunnels to the cloud breakout nodes. For more information, see *Assigning Cloud Breakout Settings to Sites* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

To enable the breakout settings to be applied to SD-WAN traffic of a site, you must assign the cloud breakout setting to the site, and reference a cloud breakout profile in an SD-WAN policy intent, and deploy the SD-WAN policy.

Add SD-WAN Branch Sites

An on-premise spoke (or a branch) represents an endpoint, like the customer premises equipment (CPE) device at a physical location, such as a branch office. Typically, these sites are connected using overlay connections to hub sites. Starting in CSO Release 6.0.0, in SD-WAN deployments, using hubs to connect sites is optional.

NOTE: Before you add the SD-WAN branch sites, check the cable connections, review the NAT and firewall ports and protocols, and check the Junos OS version of the SD-WAN CPE device. For details, see "Supported Devices for SD-WAN, and Ports and Protocols to Open" on page 168.

To add branch sites with SD-WAN capability:

1. Click **Resources > Site Management**.

   The Sites page appears.

2. Click **Add**, and select **Branch Site (Manual)**.

   The Add Branch Site wizard appears, displaying the General settings to be configured.

   > **NOTE**: Fields marked with an asterisk (*) are mandatory.

3. Configure the General settings as explained in Table 35 on page 135, and click **Next**.

   You are taken to the WAN section of the workflow.

4. Configure the WAN settings as explained in Table 36 on page 137, and click **Next**.

   > **NOTE**: In Release 6.1.0, CSO moves a site to the PROVISIONED state when at least one of the WAN links obtains the IP address and is activated.

   You are taken to the LAN section of the workflow.

5. You can add LAN segments when you're adding the site or after a site is provisioned. To add a LAN segment during the site addition workflow:

   a. Click the add (**+**) icon.

      The Create LAN Segment page appears.

   b. Configure the LAN segment settings as explained in Table 38 on page 160.

   c. Click **OK**.

      You are returned to the LAN section of the workflow and the LAN segment that you added is displayed.

6. Click **Next**.

   You are taken to the Summary section of the workflow.

7. Review the configuration in the Summary section and, if required, modify the settings.

8. Click **Finish**.

   - If you entered a serial number during activation and automatic activation is enabled, the Site Activation Progress page appears. The site activation process proceeds through the tasks explained in Table 27 on page 110.

     Click **OK** to close the Site Activation Progress page.

NOTE: If you don't want to wait for the site activation to finish, you can close the Site Activation Progress page and monitor the status of the site activation from the Jobs page (**Monitor > Jobs**).

The time taken for site activation varies depending on the device that CSO is activating.

- If you did not enter a serial number or if automatic activation is disabled, you are returned to the Sites page. CSO triggers a job and displays a confirmation message with a job link. Click the link to view the status of the job.

  After the job is finished, CSO displays a confirmation message with a job link. The status of the site changes to CREATED and an Activate Site link is displayed. You must manually activate the site to finish the process. For more information, see "Manually Activate a Site" on page 174.

TIP: After you add a site, you can modify (depending on the site status) certain parameters of the site. For more information, see *Edit Site Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Table 35: General Information (Add [SD-WAN] Branch)

Field	Guideline
Site Information	
Site Name	Enter a unique name for the site. The name can contain alphanumeric characters, and hyphens (-) and cannot exceed 32 characters.
Site Group	If you want the site to be part of a site group, select the site group. By default, **None** is selected, which means that the site doesn't belong to any site group.
Site Capabilities	

Table 35: General Information (Add [SD-WAN] Branch) *(Continued)*

Field	Guideline
WAN Capabilities	**NOTE**: Device Management, enabled by default, allows you to create a site with only device management capability (without any services) and add services later. To add an SD-WAN capability for this site, choose one of the following SD-WAN service types: • Secure SD-WAN Essentials—(Available for tenants with SD-WAN Essentials or Advanced service level) Provides basic SD-WAN services. This service is ideal for small enterprises looking for managing simple WAN connectivity with comprehensive NGFW security services at the branch sites, using link-based application steering. This service supports features such as intent-based firewall policies, WAN link management and control, and site to site communication through MPLS or internet links. The SD-WAN Essentials service does not support multihoming, dynamic mesh tunnels, cloud breakout profiles, SLA-based steering profiles, pool based source NAT rules, IPv6, MAP-E, or underlay BGP. • Secure SD-WAN Advanced—(Available for tenants with SD-WAN Advanced service level) Provides complete SD-WAN services. This service is ideal for enterprises with one or more data centers, requiring flexible topologies and dynamic application steering. You can establish site-to-site connectivity by using a hub in a hub-and-spoke topology or through static or dynamic full mesh VPN tunnels. Enterprise wide intent based SD-WAN policies and service-level agreement (SLA) measurements allow to differentiate and dynamically route traffic for different applications.
Address and Contact Information	Enter the address of the branch site and contact information in the fields provided. Although it is not mandatory, providing an address lets you visualize where the site is located on the geographical map on the Monitor Overview page.

Table 35: General Information (Add [SD-WAN] Branch) *(Continued)*

Field	Guideline
Advanced Configuration	For the DNS and NTP servers, you can either use the defaults or specify DNS and NTP servers.
Domain Name Server	Specify one or more IPv4 or IPv6, or both IPv4 and IPv6 addresses of the DNS server. To specify more than one DNS server address, type the address, press Enter, and then type the next address, and so on.
NTP Server	If needed, specify the IP addresses of one or more NTP servers.
Select Timezone	Select a time zone for the site.

Table 36: Device (Add Branch Site)

Field	Guideline
Device Series	Select the device series of the CPE device; for example, SRX. Based on the device series that you selected, the supported device templates are displayed. Ensure that you select the correct device template from the carousel. For example, for an SRX300 device, select **SRX as SD-WAN CPE** (or a modified version of that template) as the device template.
Device Information	**NOTE**: If you selected a dual CPE template, additional fields are displayed. For more information, see *Creating On-Premise Sites* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Serial Number	If you want CSO to proceed with the site activation immediately after you complete the site addition workflow, enter the serial number. If the serial number that you entered is already present in the system, CSO displays an error message. If the serial number is not present, then CSO displays a green check mark. If you want CSO to only model the site, leave this field blank. If you don't enter a serial number, you must manually activate the site later.
Device Root Password	The default root password is fetched from the ENC _ROOT_PASSWORD field in the device template. You can retain the password or change it by entering a password in plain-text format. The password is encrypted and stored on the device
Zero Touch Provisioning	Click the toggle button to enable or disable Zero Touch Provisioning (ZTP). This option is enabled by default. **NOTE**: By default, this button is disabled for vSRX. You can enable this button, if the Junos OS version running on vSRX supports phone-home client. If ZTP is disabled, you must manually copy the stage-1 configuration (generated automatically by CSO) to the device and commit the configuration on the device. If ZTP is enabled, the Boot Image field is displayed and you must select an image that supports the Phone-Home client. During ZTP, the image on the device is upgraded to the image that you select for the Boot Image.
Is Cluster Already Formed?	**NOTE**: This field is available only for SRX dual CPE devices. Click the toggle button to specify whether the SRX cluster has been manually formed (Yes) or not (No).

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Cluster ID	**NOTE**: This field is available only for SRX dual CPE devices. If the SRX cluster hasn't been formed manually, specify a unique ID for the cluster. Range: 1 through 15 If you've enabled ZTP for the site, the cluster is automatically formed when the site is activated. If you've disabled ZTP, the following processes are displayed on the Site Activation Progress page (that appears after you've added the branch site): 1. After CSO models the site (that is, after the Model Site process completes successfully), click the **Click to copy pre script link**, which appears next to the Pre Script process. 2. Execute the commands as directed. After the Pre Script process completes successfully, the SRX cluster is formed and the **recovery.conf** file is saved on the cluster. In case you want to delete the site later, you'll need this file to remove the stage-1 configuration and other configurations pushed to the device by CSO. 3. Manually configure the stage-1 configuration (generated automatically by CSO) to the primary device in the cluster and commit the configuration on the device. After the cluster is detected, CSO executes the bootstrap and provisioning processes and completes provisioning the cluster.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Auto Activate	Click the toggle button to specify whether the site activation requires an activation code or not: • **Enabled**—The site is activated automatically without an activation code. This is the default setting. • **Disabled**—The site activation proceeds only after you enter an activation code. If you choose this setting, enter the activation code (in the **Activation Code** field) that must be entered to activate the device.
Boot Image	If you want to upgrade the branch device with the latest supported Junos OS version, select the boot image from the list. The boot image is used to upgrade the device when CSO starts the ZTP process. If you don't specify a boot image, which is the default selection (**Use Image on Device**) in the list, then CSO skips the procedure to upgrade the device during ZTP.
Hub Configuration	**NOTE**: Hub selection is optional for both SD-WAN Advanced and Essentials sites. SD-WAN Essentials sites do not support multihoming, that is, you cannot select a secondary hub for SD-WAN Essentials branch sites. For sites with SD-WAN Advanced service, you must specify at least one hub to which the branch site must connect (in the **Primary Provider Hub**, **Secondary Provider Hub**,**Primary Enterprise Hub**, and **Secondary Enterprise Hub** fields). The combinations supported are listed in Table 37 on page 159.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Use Mesh Tags to connect EHub	This toggle button is enabled by default. If this button is enabled, CSO uses mesh tags to automatically form the overlay tunnel between the site and the enterprise hubs. Disable this toggle button if you want to manually create static tunnel (per WAN link) between the branch site and the enterprise hubs. If you disable this option, you must manually enable at least one WAN link to connect to the enterprise hub by using the **Connects to Enterprise Hubs** toggle button in the **Advanced Settings** of the WAN link.
WAN Links	You can configure a maximum of four WAN links and must configure at least one WAN link.
WAN_0 (*WAN-Interface-Name*)	The first WAN link is enabled by default. **NOTE**: Fields marked with an asterisk (*) must be configured to proceed.
Link Type	Select the type of link (**MPLS** or **Internet**) for the WAN link. For the first WAN link, we recommend that you use the default (**Internet**) for the underlay network type to ensure reachability to the redirect server.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Access Type	Select the access type for the underlay link: You can select LTE access type only for one WAN link. You can select ADSL or VDSL access type only for two WAN links. **NOTE**: • You can select the LTE, ADSL, or VDSL access type only for one WAN link. • You cannot configure LTE, ADSL, or VDSL as the access type if you are using the Dual SRX and Dual NFX device templates; Ethernet is configured as the access type for the underlay link. • SRX300 does not support LTE and ADSL access types. • On SRX300 Series devices (except SRX300 devices) and NFX150 devices, the LTE WAN link is supported through a SIM card that is inserted in the SIM slot of the Mini-Physical Interface Module (Mini-PIM). On NFX250 devices, the LTE WAN link is supported through a USB dongle (Vodafone K5160 dongle) that is plugged into the USB port of the CPE device. CSO supports the following combination of MPLS tunnels (with ADSL or VDSL access types) for a branch device: • From the branch site (with an ADSL or VDSL access type) to an enterprise hub, provider hub, or a branch site (with Ethernet link). • From the branch site (with an ADSL or VDSL access type) to another branch site (with an ADSL or VDSL link).

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
PPPoE/PPP	This field is displayed only for Internet links with Ethernet, ADSL, or VDSL access type, and for MPLS links with Ethernet or LTE access types. Click the toggle button to enable authenticated address assignment for the WAN link by using PPPoE (Point-to-Point Protocol [PPP] over Ethernet) or PPP. By default, this toggle button is disabled. PPPoE works with Ethernet, ADSL, and VDSL access types while PPP works with the LTE access type. If you've enabled this toggle button, you must specify the authentication parameters in the PPPoE/PPP Settings section of the page.You can enable PPPoE or PPP per WAN link.
ADSL/VDSL SFP Annex	Applicable only to MPLS or Internet links with ADSL or VDSL access types. Click the toggle button to enable the Annex J support through an *x*DSL SFP module. Annex J is specified in ITU-T recommendations G.992.3 and G.992.5. If you keep this option disabled, you must use a Mini-PIM module for connectivity. You can enable or disable this option only on new WAN links being added to a site. You cannot enable or disable this option for the existing WAN links by using the site edit workflow. You can enable this option along with other parameters such as PPoE, static IP address (IPv4/IPv6), DHCP, and VLAN ID.
Egress Bandwidth	This field is not available when you configure LTE as the access type. Enter the maximum egress bandwidth (in Mbps) allowed for the WAN link.

Table 36: Device (Add Branch Site) *(Continued)*

Field	Guideline
Underlay Address Families	
IPv4	Click the toggle button to enable or disable IPv4 address assignment for the WAN link. By default, IPv4 address assignment is enabled for the WAN link. The WAN link requires an IPv4 address to connect to an IPv4 network.
Address Assignment Method	This field is not available if you've enabled PPPoE/PPP. For LTE access type, only DHCP is available as the address assignment method. Select the method of assigning an IPv4 address to the WAN link—DHCP (Dynamic Host Configuration Protocol) or STATIC. • If you select **DHCP**, the IP address is provided by using the DHCP server of the WAN link's service provider. • If you select **STATIC**, you must provide the IPv4 address prefix and the gateway IPv4 address for the WAN link.
Static IP Prefix	If you've configured the address assignment method as STATIC, enter the IPv4 address prefix of the WAN link.
Gateway IP Address	If you've configured the address assignment method as STATIC, enter the IPv4 address of the gateway of the WAN service provider.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
MTU	Applicable only to IPv4 addresses. Enter the maximum transmission unit (MTU) size for the media or protocol. The supported MTU range can vary depending on the device, interface type, network topology, and other individual requirements. See also: MTU Default and Maximum Values and LTE Mini Physical Interface Modules (LTE Mini-PIM). Editing the MTU values of all the OAM-enabled WAN links of a site at the same time might result in tunnel flapping. You must ensure that at least one OAM-enabled WAN link always remains undisrupted for a site. For example, if you have a site with four WAN links (including two links that support OAM traffic), you can edit the MTU values of all the WAN links except one OAM-enabled link at the same time. After the edit is complete and the changes are saved, you can edit the site again and update the remaining WAN link. **NOTE**: If you enable the PPPoE/PPP option under a WAN link, the MTU option is displayed under the PPPoE/PPP Settings section for that link.
IPv6	Click the toggle button to enable or disable IPv6 address assignment for the WAN link. By default, IPv6 address assignment is disabled for the WAN link. The WAN link requires an IPv6 address to connect to an IPv6 network. **NOTE**: • IPv6 address assignment is supported only for sites with Secure SD-WAN Advanced service. • You cannot enable IPv6 address assignment for NFX250 devices.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Address Assignment Method	Select the method of assigning an IPv6 address to the WAN link—DHCP (Dynamic Host Configuration Protocol), STATIC, or SLAAC (Stateless Address Auto Configuration). If you select **STATIC**, you must provide the IPv6 address prefix and the gateway IPv6 address for the WAN link.
Static IP Prefix	If you've configured the address assignment method as STATIC, enter the IPv6 address prefix of the WAN link.
Gateway IP Address	If you've configured the address assignment method as STATIC, enter the IPv6 address of the gateway of the WAN service provider.
Access Point Name (APN)	This field can be configured only for MPLS links with LTE access type and PPPoE/PPP enabled. For MPLS links with LTE as the access type and PPPoE/PPP disabled, CSO uses the default APN settings that the CPE device is shipped with. The access point name (APN) determines the Packet Data Network Gateway (P-GW) that the CPE device must use to connect to the Packet Data Network (PDN) such as Internet. All CPE devices are shipped with default APN settings. However, if you choose to use a private APN with the current LTE service provider or to use a different LTE service provider, enter the APN for the CPE device (as specified by the service provider) in this field.
Advanced Settings	

Table 36: Device (Add Branch Site) *(Continued)*

Field	Guideline
Address Family (Tunnel Creation)	Select the underlay address family (IPv4 or IPv6) that is used to establish the overlay tunnel. The options on the list are populated based on the address family that you've configured for the underlay (either IPv4 or IPv6, or both).
Provider	Enter the name of the WAN link's service provider.
Cost/Month	Leave this as the default because this field is currently not used in CSO.
Enable Local Breakout	Click the toggle button to enable the WAN link to be used for local breakout. The toggle button is disabled by default, which means that the WAN link cannot be used for local breakout. Local breakout is an SD-WAN feature that enables Internet links to break out traffic directly from a site. For example, if you want to provide guests who visit your enterprise with Internet access, you can use local breakout to break out guest traffic locally from the site directly to the Internet. **NOTE**: • If you enable local breakout, this only means that the WAN link *can* be used for local breakout. To enable traffic to break out from the site, you must also configure a breakout profile, reference that profile in an SD-WAN policy intent, and deploy the SD-WAN policy. • If you do not enable local breakout on at least one WAN link for a single CPE site and at least two WAN links for a dual CPE site, then local breakout is disabled for the site. If you enable local breakout, additional fields appear.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Breakout Options	This field is displayed only if local breakout is enabled for the WAN link. Select whether you want to use the WAN link for both breakout and WAN traffic (default) or only for breakout traffic.
MAP-E	Click the toggle button to enable or disable the Mapping of Address and Port with Encapsulation (MAP-E) functionality on the IPv6 WAN link. By default, MAP-E is disabled. MAP-E supports transporting IPv4 packets across an IPv6 network by using IPv4-in-IPv6 encapsulation. For more information on MAP-E, see Mapping of Address and Port with Encapsulation on NFX Series Devices. **NOTE**: • MAP-E is compliant only with the Japan Network Enabler (JPNE) standards. • CSO supports MAP-E only on one WAN link of the branch site (Secure SD-WAN Advanced service only) with NFX150 as the CPE. IPV6 address assignment and local breakout must be enabled for the WAN link.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Autocreate Source NAT Rule	**NOTE**: Sites with Secure SD-WAN Essentials service support interface-based source NAT rules only. If you enable this options for an SD-WAN Essentials site, interface-based source NAT rules are automatically applied. If you enable this options for an SD-WAN Advanced site, you must select a source NAT rule from the **Translation** field. This field is displayed only if IPv4 address assignment and local breakout are enabled for the WAN link. When you enable local breakout on a link, this setting is enabled by default, which triggers automatic creation of source NAT rules for the site. You can click the toggle button to disable the automatic creation of source NAT rules. If you disable this field, then you must manually add a source NAT rule for local breakout and deploy the NAT policy on the site. **NOTE**: If NAT is not enforced by a separate device in your network (for example, an Internet gateway firewall), then we recommend that you enable this setting because it allows CSO to automatically create a NAT policy for the site.
Translation	This field is displayed only if the automatic creation of source NAT rules is enabled for the WAN link, and the SD-WAN service used is Advanced. Sites with Secure SD-WAN Essentials service support interface-based source NAT rules only. Select the type of NAT to use for the traffic on the WAN link: • **Interface**—Use interface-based NAT, which is the default setting. • **Pool**—Use pool-based NAT. If you select this option, you must specify the IP addresses that are to be used for the NAT pool.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
IP Addresses	For pool-based NAT, enter one or more IP addresses, subnets, or an IP address range. Separate multiple IP addresses by using commas and use a hyphen to denote a range; for example, 192.0.2.1-192.0.2.50. **NOTE**: No NAT is performed for tenant-owned public IP addresses that were added during the tenant addition workflow.
Preferred Breakout Link	if the WAN link is enabled for local breakout, click the toggle button to enable the WAN link as the most preferred breakout link. If you disable this option, then the breakout link is chosen using ECMP (equal-cost multipath) from the available breakout links.
BGP Underlay Options	**NOTE**: Not applicable to sites with SD-WAN Essentials service. **NOTE**: BGP underlay routing is typically used by service providers, and can be configured only if IPv4 address assignment (with STATIC as the address assignment method) and local breakout are enabled for the WAN link. Click the toggle button to enable BGP underlay routing. When you enable BGP underlay routing, route advertisements to the primary Provider Edge (PE) node and, if configured, the secondary PE node occur as follows: • CSO advertises the WAN interface subnet. • If you configured pool-based translation, CSO advertises the NAT address pool. **NOTE**: If underlay BGP is enabled for a WAN link, then the routes learnt from BGP are installed for local breakout; CSO does not generate the static default route.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Primary Neighbor	Displays the IP address that you entered for the gateway for the WAN link.
Secondary Neighbor	If you want to provide PE resiliency, you can configure a secondary PE node. Enter the IP address of the secondary PE node. **NOTE**: If the primary PE node goes down, then the secondary PE is used as the next hop. When the primary PE comes back up, the route next hops are changed to the primary PE.
eBGP Peer-AS-Number	Enter the autonomous system (AS) number for the external (EBGP) peer. **NOTE**: If the peer AS number is not configured or the peer AS number that is configured is the same as that of the CPE site, then the BGP type is assumed to be internal BGP (IBGP).
Local AS Number	Enter the local AS number for the WAN link. When you configure this parameter, the local AS number is used for eBGP peering instead of the global AS number configured for the device.
Authentication	Select the BGP route authentication method to be used: • **None**—Indicates that no authentication should be used. This is the default. • **Use MD5**—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.

Table 36: Device (Add Branch Site) *(Continued)*

Field	Guideline
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.
Advertise Public LAN Prefixes	Click the toggle button to enable the advertisement of public LAN prefixes. This field is disabled by default. If the tenant has a public IP address pool configured and you enable the advertisement of public LAN prefixes, then for LAN segments that are created with a subnet that falls under the tenant public IP address pool, CSO advertises the LAN subnet to the BGP underlay. **NOTE**: When public LAN advertisement is enabled for the WAN link, public LAN prefixes are advertised through the BGP underlay towards MPLS or the Internet.
Use for Fullmesh	Click the toggle button to enable the WAN link to be part of a full mesh topology. **NOTE**: Sites with SD-WAN Essentials service do not support creation or deletion of dynamic mesh tunnels based on a user-defined threshold for the number of sessions closed between two branch sites. However, an OpCo administrator or the Tenant administrator can create a static tunnel between a source site and destination site by using the CSO GUI in Customer Portal. Configure the two additional fields that appear:

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Mesh Overlay Link Type	If the WAN link is enabled for full mesh, select the type of encapsulation to be used for the overlay tunnels in the full mesh topology: **NOTE**: For links with public IP addresses, we recommend that you use GRE over IPsec as the mesh overlay link type. • **GRE_IPSEC**—Use GRE over IPsec. • **GRE**—Use GRE. This option is available only for MPLS links. **NOTE**: If you've enabled IPv6 address assignment for the WAN links, you can select only GRE-IPSEC as the type of mesh overlay link.
Mesh Tag	If the WAN link is enabled for full mesh, select the mesh tag for the WAN link. **NOTE**: The tunnels between two branch sites or an branch site and an enterprise hub site are added based on matching mesh tags. So, if you want meshing to take place between such sites, the mesh tags must be the same for both sites. For more information about mesh tags, see *Mesh Tags Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).
Connects to Enterprise Hubs	This field is displayed only if you have enabled the **Use Mesh Tags to Connect EHub** field in the **Hub Configuration** section. Enable this toggle button if you want to manually connect the site to an enterprise hub, without using mesh tags.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Primary EHub Tunnel Type	This field is displayed only if you have enabled the **Connects to Enterprise Hubs** field. Select the tunnel type to be used for the connection between the branch site and the primary enterprise hub.
Primary EHub Peer Device	This field is displayed only if you have enabled the **Connects to Enterprise Hubs** field. Displays the name of the primary enterprise hub you have selected.
Primary Ehub Peer Interface	This field is displayed only if you have enabled the **Connects to Enterprise Hubs** field. Select the primary enterprise hub WAN link that needs to be part of the tunnel. You can select multiple WAN links.
Secondary EHub Tunnel Type	This field is displayed only if you have enabled the **Connects to Enterprise Hubs** field. Select the tunnel type to be used for the connection between the branch site and the secondary enterprise hub.
Secondary EHub Peer Device	This field is displayed only if you have enabled the **Connects to Enterprise Hubs** field. Displays the name of the secondary enterprise hub you have selected.
Secondary Ehub Peer Interface	This field is displayed only if you have enabled the **Connects to Enterprise Hubs** field. Select the secondary enterprise hub WAN link that needs to be part of the tunnel. You can select multiple WAN links.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Use for OAM traffic	**NOTE**: The **Connects to Hubs** field is available only if you have selected a provider hub. Click the toggle button to enable the use of the WAN link for OAM traffic. The WAN link is then used to establish an OAM tunnel for communication between the enterprise hub site and CSO. You must configure at least one WAN link to be used for OAM traffic. To ensure redundancy, we recommend that you configure at least two WAN links that can be used for OAM traffic. In addition, for added management redundancy, use two links with different transport paths.
Backup Link	Select a backup link through which traffic can be routed when the primary (other) links are unavailable. You can select any link other than the default links or links that are configured exclusively for local breakout traffic. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, SLA data is not monitored for the backup link.
Default Link	Select one or more links that will be used for routing traffic in the absence of matching SD-WAN policy intents. A site can have multiple default links to the hub site. Default links are used primarily for overlay traffic but can also be used for local breakout traffic. However, a default link cannot be used exclusively for local breakout traffic. If you do not specify a default link, then ECMP is used to choose the link on which to route traffic.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
Data VLAN ID	Enter a VLAN ID for the WAN link. Range: 0 through 4049 (4050 to 4094 is reserved by CSO). **NOTE**: • If you are configuring more than one WAN link on the same physical interface, only one WAN link can be untagged; for the remaining WAN links, you must configure a VLAN ID. • A combination of tagged and untagged on the same physical interface is supported only for single CPE devices.
WAN_1 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
WAN_2 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields

Table 36: Device (Add Branch Site) *(Continued)*

Field	Guideline
WAN_3 (*WAN-Interface-Name*)	Click the toggle button to enable or disable (default) the WAN link. When you enable the WAN link, fields related to the WAN link appear. Fields marked with an asterisk (*) must be configured to proceed. Refer to the fields described for WAN_0 (*WAN-Interface-Name*) for an explanation of the fields
Advanced Configuration **NOTE**: Sites with SD-WAN Essentials service do not support creation or deletion of dynamic mesh tunnels based on a user-defined threshold for the number of sessions closed between two branch sites. However, an OpCo administrator or a tenant administrator can create a static tunnel between a source site and destination site by using the CSO GUI in Customer Portal.	
OAM IP Prefix	We recommend that you *do not* configure this setting (leave the IP Prefix field blank) because management connectivity is handled automatically by CSO.
DVPN Threshold for Tunnel Creation	Specify the threshold for the number of sessions (flows) closed (in a two-minute duration) between the branch site and a destination site. When the number of sessions closed exceeds the specified threshold, a tunnel is created between the branch site and the destination site. For example, if you specify a threshold of as 7, dynamic mesh tunnels are created if the number of sessions closed (in two minutes) between the branch site and destination site exceeds 7.

Table 36: Device (Add Branch Site) ***(Continued)***

Field	Guideline
DVPN Threshold for Tunnel Deletion	Specify the threshold for the number of sessions closed (in a 15-minute duration) between the branch site and a destination site. When the number of sessions closed is lower than the specified threshold, the tunnel between the branch site and destination site is deleted. For example, if you specify the number of sessions closed as 5, dynamic mesh tunnels between the branch site and destination site are deleted if the number of sessions closed (in a 15-minute duration) is lesser than or equal to 5.
Configuration Templates (Optional) If you want to deploy additional configuration during the ZTP process, you can select one or more configuration templates and set the parameters for each template.	

Table 36: Device (Add Branch Site) *(Continued)*

Field	Guideline
Configuration Templates List	For each configuration template that you select 1. Select one or more configuration templates from the list that you want to deploy on the device during ZTP. 2. Click **Set Parameters**. The Device Configurations page appears. The names and configuration parameters of the configuration templates that you selected are displayed in the Configure tab. 3. For each configuration template, enter values for the parameters. 4. (Optional) Click the **Summary** tab to view the Junos OS configuration commands that will be deployed on the device for the different configuration templates. 5. Click **Save**. You are returned to the WAN tab. The Junos OS configuration commands will be deployed on the device during the ZTP process.

Table 37: Supported Combinations of Provider and Enterprise Hubs

Provider Hubs Specified	Enterprise Hubs Specified
Primary	None
Primary	Primary
Primary	Primary and Secondary
Primary and Secondary	None

Table 37: Supported Combinations of Provider and Enterprise Hubs *(Continued)*

Provider Hubs Specified	Enterprise Hubs Specified
Primary and Secondary	Primary
Primary and Secondary	Primary and Secondary
None	Primary
None	Primary and Secondary

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke)

Field	Description
Use for Overlay VPN	Enable the **Use for Overlay VPN** field to associate the LAN segment with the selected department (VRF + ZONE) for overlay traffic to other sites. Disable the **Use for Overlay VPN** field to associate the LAN segment with a security zone for underlay breakout. You must define zone-based security policies. **NOTE**: When adding a new site, this field is enabled by default and cannot be modified. However, when you add a new LAN Segment to a provisioned site from the **LAN** tab of the **Site-Name** page, you can enable or disable this option.
Name	Enter a name for the LAN segment. The name for a LAN segment should be a unique string of alphanumeric characters and some special characters (. -). No spaces are allowed and the maximum length allowed is 15 characters.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) *(Continued)*

Field	Description
CPE Port	**NOTE**: Applicable to SRX Series devices. Select the CPE port to be added in the LAN segment. When you add a new LAN Segment to a provisioned site from the LAN tab of the *Site-Name* page, you can select (or create) a LAG interface or a redundant Ethernet (reth) interface (for dual CPE cluster) to connect the SRX Series CPE devices to an EX series switch. To use the et interface on SRX4600 devices, you must create a LAG interface and configure the et interface as a member of the LAG (aggregated Ethernet or ae) interface. See *Create LAG Interface*. For an SRX4600 dual CPE cluster, you can use the et interface if it is configured as a member of the redundant Ethernet (reth) interface.
Add LAG Interface	**NOTE**: This option is available when you add a new LAN Segment to a provisioned site from the LAN tab of the *Site-Name* page. Click the link to create a LAG interface (ae interface) if you want to use it to connect the SRX Series CPE to the EX Series switch. See *Create LAG Interface* for details.
Create RETH Interface	**NOTE**: This option is available when you add a new LAN Segment to a provisioned site from the LAN tab of the *Site-Name* page. Click the link to create a reth interface for an SD-WAN site with a dual CPE cluster. See *Create a RETH Interface* for details.
Type **NOTE**: This field is displayed only for LAN segments associated with enterprise hub sites.	Select the type of LAN segment: • Directly Connected (default)—Indicates that the LAN segment is directly connected to the site. • Dynamic Routed—Indicates that the LAN segment is not directly connected to the site and is reachable by using a dynamic route. If you select this option, you must specify the dynamic routing information.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) ***(Continued)***

Field	Description
VLAN ID	Enter the VLAN ID for the LAN segment. By default, VLAN ID is set to 1 and native VLAN is enabled for untagged traffic. Range: 1 to 4049 .
Use for Native VLAN	Enable this option to use the VLAN ID specified above for untagged traffic. The CPE interface is configured with a native-vlan-id, which has the same value as the VLAN ID.
Department	**NOTE**: This field is available only if the **Use for Overlay VPN** field is enabled. Select a department to which the LAN segment is assigned. Alternatively, click the **Create Department** link to create a new department and assign the LAN segment to it. See *Add a Department* for details. You can group LAN segments as departments for ease of management and for applying policies at the department-level. For LAN segments that are dynamically routed, you can assign only a data center department.
Gateway Address/Mask	Enter a valid gateway IP address and mask for the LAN segment. This address will be the default gateway for endpoints in this LAN segment. For example: 192.0.2.8/24.
Zone	**NOTE**: This field is available only if the **Use for Overlay VPN** field is disabled. Select a security zone to be associated with this LAN segment. Alternatively click **Create Zone** to create a new security zone and assign that to this LAN segment. See *Adding a Security Zone* for details.
DHCP	For directly connected LAN segments, click the toggle button to enable DHCP. You can enable DHCP if you want to assign IP addresses by using a DHCP server or disable DHCP if you want to assign a static IP address to the LAN segment. **NOTE**: If you enable DHCP, additional fields appear on the page.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) ***(Continued)***

Field	Description
Additional fields related to DHCP	
Address Range Low	Enter the starting IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Address Range High	Enter the ending IP address in the range of IP addresses that can be allocated by the DHCP server to the LAN segment.
Maximum Lease Time	Specify the maximum duration (in seconds) for which a client can request for and hold a lease on the DHCP server. Default: 1440 Range: 0 through 4,294,967,295 seconds.
Name Server	Specify one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address. **NOTE**: DNS servers are used to resolve hostnames into IP addresses.
CPE Ports	**NOTE**: Applicable to NFX150 and NFX250 devices. For sites with SD-WAN capability, the CPE Ports field is disabled and the CPE ports that you can include in the LAN segment are listed. Select the ports from the **Available** column and click the right-arrow to move the ports to the **Selected** column.
Static Routing Use this section to configure static routing on the LAN segment. Provide the IP addresses of all the LAN routers connected to the CPE device and the static subnets behind these routers.	
Add LAN Router IP Prefix	
LAN Router IP	Enter the IP address of the LAN router that is connected to the CPE device.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) *(Continued)*

Field	Description
Prefix	Enter the subnets that are connected to the LAN router.
BFD	Enable Bidirectional Forwarding Detection (BFD) to detect any failures on the static route.
Dynamic Routing	
Routing Protocol	Enable this toggle button to configure dynamic routing using the BGP or OSPF protocol.
BFD	Enable Bidirectional Forwarding Detection (BFD) to detect any failures in the LAN segment.
Protocol	Select either BGP or OSPF.
BGP Configuration **NOTE**: Starting in Release 6.1.0, CSO explicitly disables the long-lived graceful restart (LLGR) capability for BGP peering sessions with provider edge (PE) and data center or LAN routers. Disabling LLGR ensures that the CPE does not differentiate the route advertisements to the peering router irrespective of the peering router's LLGR capability. Prior to CSO Release 6.1.0, LLGR helper mode is enabled by default (implicit behavior of Junos OS) on the CPE for BGP peering towards PE router in IP VPN deployments, and data center or LAN routers in data center deployments.	
Authentication	Select the BGP route authentication method to be used: • None—Indicates that no authentication should be used. This is the default. • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key.
Auth Key	If you specified that MD5 should be used for authentication, specify an MD5 authentication key (password), which is used to verify the authenticity of BGP packets.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) ***(Continued)***

Field	Description
BGP Options	You can select the following options based on your requirements: • AS-OVERRIDE: Replaces all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer. • AS-PATH-PREPEND: Prepends one or more autonomous system (AS) numbers at the beginning of an AS path. Prepending an AS path makes a shorter AS path look longer and therefore it becomes less preferable to BGP. • AS-LOOP: Allows the local device's AS number to be added in the received AS paths. You can specify the number of times the detection of local AS is allowed in the AS path.
Loop Count	This field is displayed only if you select AS-LOOP. Enter the maximum number of times the detection of local AS is allowed in the AS path.
Peer IP Address	Enter the IP address of the LAN BGP peer.
Peer AS Number	Enter the autonomous system (AS) number of the LAN BGP peer. By default, CSO uses the AS number 64512. You can enter a different AS number.
Local AS Number	Enter the local AS number. When you configure this parameter, the local AS number is used for BGP peering instead of the global AS number configured for the CPE.
OSPF Configuration	
OSPF Area ID	Specify the OSPF area identifier to be used for the dynamic route.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) ***(Continued)***

Field	Description
Authentication	Select the OSPF route authentication method to be used: • Password—Indicates that password-based authentication should be used. If you choose this option, you must specify the password. (This is the default). • Use MD5—Indicates that MD5 is to be used for authentication. If you choose this option, you must specify an authentication key. • None—Indicates that no authentication should be used.
Password	Enter the password to be used to verify the authenticity of OSPF packets.
Confirm Password	Retype the password for confirmation purposes.
MD5 Auth Key ID	If you specified that MD5 should be used for authentication, enter the OSPF MD5 authentication key ID. Range: 1 through 255.
Auth Key	If you specified that MD5 should be used for authentication, enter an MD5 authentication key, which is used to verify the authenticity of OSPF packets.
Route Advertisement Control	
LAN Route(s) to Overlay	When this option is enabled, LAN routes are advertised to the remote CPEs. By default, this option is enabled. Starting in CSO Release 6.2.0, you can configure export policies in conjunction with the **LAN Route(s) to Overlay** option for more granular control over routes that are advertised to the overlay network. For example, when the **LAN Route(s) to Overlay** option is enabled, you can configure policies to prevent specific routes from being advertised. Similarly, when the **LAN Route(s) to Overlay** option is disabled, you can configure policies to allow only specific routes to be advertised.

Table 38: Fields on the Add LAN Segment page (Add Branch Spoke) ***(Continued)***

Field	Description
Overlay Route(s) to LAN	This option is displayed only if you enable the Routing Protocol toggle button. By default, this option is disabled. Enable this option to advertise the remote CPE routes received in a department to the LAN router. **NOTE**: In CSO Release 6.0.0 and earlier releases, this option is called Advertise LAN Prefix and is applicable only for data center departments. Starting in CSO Release 6.2.0, you can use the following policies for granular control of the route advertisements: • Import policies for granular control of the routes that a CPE device accepts from the list of routes advertised by the LAN router. • Export policies for granular control of the routes that a CPE device advertises to the LAN router.
Static/Aggr Routes to Overlay	Enable this option to allow advertisement of static or aggregate routes to the overlay network. • If a large number of LAN routes are present, then you can disable the **LAN Route(s) to Overlay** option and use this option to advertise aggregate routes. • If you want to advertise additional routes, then you can enable the **LAN Route(s) to Overlay** option and use this option to advertise additional static routes.

See "Post-Provisioning Tasks for Enterprise Hub and SD-WAN Spoke Sites" on page 115.

Supported Devices for SD-WAN, and Ports and Protocols to Open

IN THIS SECTION

- What's Next | 173

For the SD-WAN devices supported by CSO, and list of ports or protocols that must be opened for the devices, see:

- Table 39 on page 169 for enterprise hub and SD-WAN on-premise spoke devices.
- Table 40 on page 173 for provider hub devices.

NOTE: During the site activation process for SRX4100, SRX4200, SRX4600, and vSRX 3.0, you must copy the stage-1 configuration (generated automatically by CSO) to the device, and commit the configuration on the device.

Before you add a provider hub device, enterprise hub site, or an SD-WAN on-premise spoke site:

- Connect cables to the device according to your network design, and power on the device.
  - For enterprise hubs and SD-WAN on-premise spoke devices, see the hardware documentation links in Table 39 on page 169.

    NOTE: We assume that the SD-WAN on-premise spoke devices will obtain the DHCP IP address (if DHCP is configured as the address assignment method) and will have Internet connectivity along with DNS resolution, when connected according to the network design.

  - For provider hub devices, see the hardware documentation links in Table 40 on page 173
- For enterprise hubs and SD-WAN on-premise spoke devices, ensure that the NAT and firewall ports and protocols listed in Table 39 on page 169 are open on the network.
- For provider hubs, ensure that the ports and protocols listed in Table 40 on page 173 are open on the network.

- Ensure that the devices are running the recommended version of Junos OS for the CSO release that you are using. For up-to-date information about the supported Junos OS versions in a CSO release, refer to the CSO Release Notes for that release (available at the CSO Documentation page).
- Before you initiate ZTP for the enterprise hub, ensure that the hub device can connect to CSO.

Table 39: Supported Enterprise Hub and SD-WAN On-Premise Spoke Devices, and NAT and Firewall Ports to Open

Device Model	Supported Site Type	NAT and Firewall Protocols or Ports	WAN Link Ports (minimum one port; maximum four ports)	Hardware Documentation Links
NFX150	On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	heth-0-0 through heth-0-5	NFX150 Chassis
NFX250	On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 7804 TCP Port 8060	ge-0/0/10 ge-0/0/11 xe-0/0/12 xe-0/0/13	NFX250 Chassis

Table 39: Supported Enterprise Hub and SD-WAN On-Premise Spoke Devices, and NAT and Firewall Ports to Open *(Continued)*

Device Model	Supported Site Type	NAT and Firewall Protocols or Ports	WAN Link Ports (minimum one port; maximum four ports)	Hardware Documentation Links
SRX300 SRX320	On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	ge-0/0/0 through ge-0/0/7	SRX300 Chassis SRX320 Chassis
SRX340 SRX345	On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	ge-0/0/0 through ge-0/0/15	SRX340 Chassis SRX345 Chassis
SRX380	Enterprise hub On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	ge-0/0/0 ge-0/0/2 to ge-0/0/15 xe-0/0/16 to xe-0/0/19	SRX380 Chassis

Table 39: Supported Enterprise Hub and SD-WAN On-Premise Spoke Devices, and NAT and Firewall Ports to Open *(Continued)*

Device Model	Supported Site Type	NAT and Firewall Protocols or Ports	WAN Link Ports (minimum one port; maximum four ports)	Hardware Documentation Links
SRX550M	On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	ge-0/0/0 through ge-0/0/9	SRX550 HM Chassis
SRX1500	Enterprise hub On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	ge-0/0/7 ge-0/0/8 xe-0/0/18 xe-0/0/19	SRX1500 Chassis
SRX4100 SRX4200	Enterprise hub On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 TCP Port 500 UDP Port 4500 TCP Port 8060	xe-0/0/0 xe-0/0/1 xe-0/0/2 xe-0/0/3	SRX4100 Chassis SRX4200 Chassis

Table 39: Supported Enterprise Hub and SD-WAN On-Premise Spoke Devices, and NAT and Firewall Ports to Open *(Continued)*

Device Model	Supported Site Type	NAT and Firewall Protocols or Ports	WAN Link Ports (minimum one port; maximum four ports)	Hardware Documentation Links
SRX4600	Enterprise hub On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 TCP Port 500 UDP Port 4500 TCP Port 8060	xe-1/1/0 xe-1/1/1 xe-1/1/2 xe-1/1/3	SRX4600 Chassis
vSRX	Enterprise hub On-premise (SD-WAN) spoke	IP Protocol 50 IP Protocol 51 TCP Port 443 UDP Port 500 UDP Port 4500 TCP Port 8060	ge-0/0/0 ge-0/0/1 ge-0/0/2 ge-0/0/3	vSRX Deployment Guides

Table 40: Provider Hub Devices Supported, and Ports and Protocols to Open

Device Model	Ports and Protocols	Hardware Documentation Links
SRX1500	IP Protocol 50 IP Protocol 51 TCP and UDP Ports 53 (for DNS) UDP Port 123 (for NTP) TCP Port 443 UDP Port 500 UDP Port 4500	SRX1500 Chassis
SRX4100 SRX4200 SRX4600	IP Protocol 50 IP Protocol 51 TCP and UDP Ports 53 (for DNS) UDP Port 123 (for NTP) TCP Port 443 UDP Port 500 UDP Port 4500	SRX4100 Chassis SRX4200 Chassis SRX4600 Chassis
vSRX	IP Protocol 50 IP Protocol 51 TCP and UDP Ports 53 (for DNS) UDP Port 123 (for NTP) TCP Port 443 UDP Port 500 UDP Port 4500	vSRX Deployment Guides

WHAT'S NEXT

See "CSO SD-WAN Deployment Workflow" | 75.

Manually Activate a Site

If you did not enter a serial number for a device during the site addition workflow or if you required an activation code to be entered before a site could be activated, you must manually activate the site.

NOTE: The following procedure is for a site associated with a single device. For sites with dual CPE devices, the workflow remains the same except that the serial number, activation code, or both parameters must be specified for two devices.

To manually activate a site:

1. Click **Resources > Site Management**.

   The Sites page appears.
2. Click the **Activate Site** link (in the Site Status column) for the site, or select the site and select **More > Activate Site**.

   The Activate Site page appears.
3. If a serial number was not specified when the site was added, enter the serial number of the device in the **Serial Number** field. Serial numbers are case sensitive.

   If the serial number that you entered is already present in the system, CSO displays an error message. If the serial number is not present, then CSO displays a green check mark.
4. If automatic activation was disabled when the site was added, enter the activation code of the device in the **Activation Code** field.
5. Click **OK**.

   If ZTP was enabled for the site, CSO triggers a job and the Site Activation Progress page appears after a few seconds. Because the site was previously modelled, the Ship Device task is the first task to be executed. The rest of the steps are as explained in Table 27 on page 110.

   NOTE: If you don't want to wait for the site activation to finish, you can close the Site Activation Progress page and monitor the status of the site activation from the Jobs page (**Monitor > Jobs**). The time taken for site activation varies depending on the device that CSO is activating.

Depending on the workflow, see "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29, "CSO SD-WAN Deployment Workflow" on page 75, or "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Monitor SD-WAN Sites and Devices

After configuring SD-WAN, you can perform the following monitoring tasks:

- On the Overview page (**Monitor > Overview**), you can view the sites that you configured on a geographical map, and the site status and connections between the sites.

  You can filter based on sites, connections, or both, and zoom in to the map.

  For more information, see *About the Monitor Overview Page* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

- On the *Site-Name* page (**Resources > Site Management > *Site-Name***), you can view general information about the site, WAN overlay and underlay links, policies, and devices.

  For more information, see *Managing a Single Site* in the *CSO Customer Portal User Guide*.

- On the *Device-Name* page (**Resources > Devices > *Device-Name***), you can view general information about the device, and view recent alerts and alarms.

  For more information, see *Managing a Single CPE Device* in the *CSO Customer Portal User Guide*.

- On the Generated Alerts page (**Monitor > Alerts**), you can view the alerts generated by the SD-WAN CPE or enterprise hub devices.

  For more information, see *About the Generated Alerts Page* in the *CSO Customer Portal User Guide*.

- On the Alarms page (**Monitor > Alarms**), you can view the alarms raised by the SD-WAN CPE or enterprise hub devices.

  For more information, see *About the Alarms Page* in the *CSO Customer Portal User Guide*.

- On the *Tenant-Name* SLA Performance page (**Monitor > Application SLA Performance**). you can view the SLA performance of the tenant's sites that have met and not met the defined SLA values.

  For more information, see *About the SLA Performance of a Single Tenant Page* and *Viewing the SLA Performance of a Site* topics in the *CSO Customer Portal User Guide*.

- On the Application Visibility page (**Monitor > Application Visibility**, you can view information about your applications such as sessions, bandwidth consumed, and risk levels.

For more information, see *About the Application Visibility Page* in the *CSO Customer Portal User Guide*.

- On the User Visibility page (**Monitor > User Visibility**), you can view information about the devices (such as top 50 devices accessing high bandwidth-consuming applications and establishing higher number of sessions) on your network.

  For more information, see *About the User Visibility Page* in the *CSO Customer Portal User Guide*.

- On the Traffic Logs page (**Monitor > Traffic Logs**), you can view the traffic logs from different sites.

  For more information, see *About the Traffic Logs Page* in the *CSO Customer Portal User Guide*.

- On the SD-WAN Report Definitions page (**Reports > Report Definitions**), you can use predefined report definitions or create custom report definitions to generate SD-WAN performance, tenant performance, and site performance reports.

  For more information, see *About the SD-WAN Report Definitions Page* in the *CSO Customer Portal User Guide*.

CHAPTER 4

Standalone Next-Generation Firewall Deployment

CSO Next-Generation Firewall (NFGW) Deployment Workflow | 178

Add Next-Generation Firewall (Branch) Sites | 180

Deploy a Firewall Policy | 188

Deploy a NAT Policy | 189

Configure Unified Threat Management (UTM) in CSO | 190

Configure and Deploy SSL Proxy Policy in CSO | 215

Configure Intrusion Prevention System (IPS) in CSO | 228

Add and Deploy Firewall Policies | 238

Add and Deploy NAT Policies | 243

Supported Devices for NGFW, and Ports and Protocols to Open | 255

Monitor Next-Generation Firewall Sites and Devices | 258

CSO Next-Generation Firewall (NFGW) Deployment Workflow

The Contrail Service Orchestration (CSO) next generation firewall (NGFW) deployment focuses on providing remote network security through the use of SRX Series NGFW devices as the customer premises equipment (CPE) at the branch site. In CSO, you can add two types of NGFW devices:

- Greenfield—Greenfield devices are generally devices on which you've not deployed any configuration. When you add a greenfield NGFW site, CSO provisions the device by using Zero Touch Provisioning (ZTP). You can then configure and use the NGFW as needed.
- Brownfield—Brownfield devices are generally devices that are already configured and operational. When you add a brownfield NGFW site, CSO *does not* provision the device by using ZTP. This allows you to import existing policies on the device into CSO and deploy the policies. You can then manage the NGFW by using CSO.

NOTE: Ensure that the pre-deployment tasks related to NGFW are carried out *before* you follow the procedure outlined in this topic. See "Pre-Deployment Tasks for CSO SD-WAN and Next-Generation Firewall" on page 29.

The following tasks must be performed in the tenant scope in Customer Portal:

1. If you are a Tenant Administrator, log in to Customer Portal. If you are an SP Administrator (CSO on-premises) or OpCo Administrator (with appropriate permissions), switch scope to the tenant. See "Switch Scope or Log in as Tenant Administrator" on page 77.
2. (Optional) Customize configuration templates. See "Configuration Templates Workflow" on page 35.
3. (Optional) Customize device templates. See "Device Templates Workflow" on page 36.
4. Depending on whether you're using a greenfield or a brownfield device:
   - Starting in CSO Release 6.0.0, the ZTP process is simplified to separate the device and service provisioning processes for faster deployment. You can add a greenfield site without applying a service and then edit the site to add the NGFW service later. See "Add Branch or Enterprise Hub Sites Without Provisioning a Service" on page 264.
   - To add a greenfield next-generation firewall site, select the SRX as Security Services CPE (or a modified version) as the device template. See "Add Next-Generation Firewall (Branch) Sites" on page 180.
   - To add a brownfield next-generation firewall site, select SRX_Standalone_Pre_Staged_NonZTP (or a modified version) as the device template. CSO generates a stage-1 configuration that you

must commit on the device, so that CSO can take over the management of the device. See "Add Next-Generation Firewall (Branch) Sites" on page 180.

5. Upload and install device licenses. See "Add and Install (Push) Device Licenses" on page 115.
6. Install the signature database. See "Install the Signature Database on Devices" on page 117.
7. (Greenfield only) Before you add firewall and NAT policies, you must add interfaces (physical and logical), routing instances, and zones for the device. You can do this on the Configuration tab of the *Device-Name* page (**Resource > Devices > *Device-Name***). See *Configuring the Firewall Device* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).
8. (Brownfield only) If you specified that policies should be imported during the activation process, you must deploy the imported policies in CSO:
   - If a firewall policy was imported, deploy the firewall policy. See "Deploy a Firewall Policy" on page 188.
   - If a NAT policy was imported, deploy the NAT policy. See "Deploy a NAT Policy" on page 189.

   NOTE: CSO imports the existing routing instances, interfaces, and zones on the brownfield device into CSO.
9. (Brownfield only) If you did not import the policies as part of the site activation, you can import the policies manually and deploy the policies:
   a. To import firewall policies, go to the Firewall Policy page (**Configuration > Firewall > Firewall Policy**) and click **Import**. For more information, see *Importing Firewall Policies* in the *CSO Customer Portal User Guide*.
   b. To import NAT policies, go to the NAT Policy page (**Configuration > NAT> NAT Policy**) and click **Import**. For more information, see *Importing NAT Policies* in the *CSO Customer Portal User Guide*.
   c. Deploy the manually imported policies, as explained in step 8.
10. (Optional) Configure unified threat management (UTM) on the next-generation firewall. See "Configure Unified Threat Management (UTM) in CSO" on page 190.
11. (Optional) Configure SSL proxy on the next-generation firewall site. See "Configure and Deploy SSL Proxy Policy in CSO" on page 215.
12. (Optional) Configure intrusion prevention system (IPS) on the next-generation firewall. See "Configure Intrusion Prevention System (IPS) in CSO" on page 228.
13. Add a firewall policy and zone-based intents and deploy the firewall policy. See "Add and Deploy Firewall Policies" on page 238.

> **NOTE**: You can also use the default firewall policy in CSO by either deploying the policy as-is or modifying the intents as required and deploying the policy.
>
> This step is optional for the brownfield device if you've already imported the firewall policies previously configured on the device.

14. (Optional) Add a NAT policy and rules and deploy the NAT policy. See "Add and Deploy NAT Policies" on page 243.

> **NOTE**: You can also use the default NAT policy in CSO by either deploying the policy as-is or modifying the rules as required and deploying the policy.
>
> This step is optional for the brownfield device if you've already imported the NAT policies previously configured on the device.

15. Monitor the NGFW sites and devices. See "Monitor Next-Generation Firewall Sites and Devices" on page 258.

Add Next-Generation Firewall (Branch) Sites

> **NOTE**: Before you add the next generation firewall (NGFW) branch site, check the cable connections, review the NAT and firewall ports and protocols, and check the Junos OS version of the NGFW device. For details, see "Supported Devices for NGFW, and Ports and Protocols to Open" on page 255.

To add a NGFW branch site:

1. Click **Resources > Site Management**.

   The Sites page appears.
2. Click **Add**, and select **Branch (Manual)**.

   The Branch Site wizard appears, displaying the General settings to be configured.

> **NOTE**: Fields marked with an asterisk (*) are mandatory.

3. Configure the General settings as explained in Table 42 on page 182, and click **Next**.

   You are taken to the Device section of the workflow.

4. Configure the Device settings as explained in Table 43 on page 183, and click **Next**.

   You are taken to the Configuration (Templates) section of the workflow.

5. (Optional) Configure the templates as explained in Table 44 on page 188, and click **Next**.

   You are taken to the Summary section of the workflow.

6. Review the configuration in the Summary section and, if required, modify the settings.

7. Click **Finish**.

   CSO triggers the activation of the site. See Table 41 on page 181 for how the site activation proceeds for greenfield and brownfield sites.

Table 41: Site Activation Process for Greenfield and Brownfield NGFW Sites

Type of NGFW	Serial Number	Auto-Activate	Site Activation Process
Greenfield or Brownfield	Not specified	Disabled	You are returned to the Sites page. CSO triggers a job and displays a confirmation message with a job link. Click the link to view the status of the job. After the job is finished, CSO displays a confirmation message and the status of the site changes to CREATED and an Activate Site link is displayed. You must manually activate the site to finish the activation process. For more information, see "Manually Activate a Site" on page 174.
Greenfield or Brownfield	Not specified	Enabled	
Greenfield or Brownfield	Specified	Disabled	
Brownfield	Specified	Enabled	CSO triggers the site activation and the Site Activation: Progress page appears. The site activation process proceeds through the tasks explained in Table 27 on page 110. **NOTE**: Because you're adding a brownfield NGFW site, you must copy the stage-1 configuration that CSO generates, paste it, and commit it on the NGFW device for the activation to proceed.
Greenfield	Specified	Enabled	CSO triggers the site activation and the Site Activation Progress page appears. The site activation process proceeds through the tasks explained in Table 27 on page 110. If you don't want to wait for the site activation to finish, you can close the Site Activation Progress page and monitor the status of the site activation from the Jobs page (**Monitor > Jobs**). The time taken for site activation varies depending on the device that CSO is activating.

Table 42: General Settings (Branch Site Page)

Field	Guideline
Site Information	
Site Name	Enter a unique name for the site. The name can contain alphanumeric characters, and hyphens (-) and cannot exceed 32 characters.
Device Host Name	The device host name is auto-generated and uses the format *tenant-name.host-name*. You cannot change the tenant-name part in the device host name. Use alphanumeric characters and hyphen (-); the maximum length allowed is 32 characters.
Site Group	If you want the site to be part of a site group, select the site group. By default, **None** is selected, which means that the site doesn't belong to any site group.
Site Capabilities	Because we're configuring a next-generation firewall site, click the **Security Services** card. By default, Device Management is selected.
Address and Contact Information	Enter the address of the branch site and contact information in the fields provided. Although it is not mandatory, providing an address lets you visualize where the site is located on the geographical map on the Monitor Overview page.
Advanced Configuration	For the DNS and NTP servers, you can either use the defaults or specify DNS and NTP servers.
Domain Name Server (DNS)	Specify the IPv4 addresses of one or more DNS servers.
NTP Server	If needed, specify the IP addresses of one or more NTP servers.
Select Timezone	Select a time zone for the site.

Table 43: Device Settings (Add Device Site)

Field	Guideline
Device Redundancy	Disabled by default. Enable this option only for dual CPEs.
Device Series	Because only SRX Series devices can be configured as NGFW sites, this field displays SRX.
Device Model	Select the SRX model.
Serial Number	If you want CSO to proceed with the site activation immediately after you complete the site addition workflow, enter the serial number. If the serial number that you entered is already present in the system, CSO displays an error message. If the serial number is not present, then CSO displays a green check mark. If you want CSO to only model the site, leave this field blank. If you don't enter a serial number, you must manually activate the site later.
Device Root Password	The default root password is fetched from the ENC _ROOT_PASSWORD field in the device template. You can retain the password or change it by entering a password in plain-text format. The password is encrypted and stored on the device

Table 43: Device Settings (Add Device Site) *(Continued)*

Field	Guideline

Zero Touch Provisioning	By default, Zero Touch Provisioning is enabled. If you want to disable ZTP, click the toggle button.
	
	NOTE: By default, this button is disabled for vSRX. You can enable this button, if the Junos OS version running on vSRX supports phone-home client.
	
	To use ZTP, ensure the following:
	
	• Device must have connectivity to CSO and Juniper phone-home server (https://redirect.juniper.net)
	
	Use telnet to verify connectivity:
	
	`telnet redirect.juniper.net:443`
	
	`telnet CSO Hostname/IP:443`
	
	If the connection is established, the device has connectivity to the phone-home server and CSO.
	
	• Required certificates for phone-home server and CSO must be present on the device.
	
	If ZTP is enabled, the Boot Image field is displayed and you must select an image that supports the Phone-Home client. During ZTP, the image on the firewall device is upgraded to the image that you select for the Boot Image.
	
	If you disable ZTP, you must copy the stage-1 configuration from CSO and commit it on the device. Use any of the following options to copy the stage-1 configuration:
	
	• Click the **Click to copy stage-1 config** link next to Prestage Device task in the Site Activation Progress page.
	
	If you close the Site Activation Progress page inadvertently, you can access the page from the Site Management page. Click the **View** link next to the status of the site under the Site Status column.
	
	• On the Devices page (**Resources > Devices**), select the device and click **Stage1 Config**.

Table 43: Device Settings (Add Device Site) *(Continued)*

Field	Guideline
Auto Activate	Click the toggle button to specify whether the site activation requires an activation code or not: • **Enabled**—The site is activated automatically without an activation code. This is the default setting. • **Disabled**—The site activation proceeds only after you enter an activation code. If you choose this setting, enter the activation code (in the **Activation Code** field) that must be entered to activate the device.
Management Interface Family	Select the IP address type (IPv4 or IPv6) for the management interface. This field is displayed only if you have enabled **Zero Touch Provisioning**.
Management Connectivity **NOTE**: This section is displayed only if you disable Zero Touch Provisioning.	
Address Family	Select the IP address type (IPv4 or IPv6).
Interface Name	Enter the management interface.
Access Type	Select the access type for the underlay link. LTE, ADSL, and VDSL access types are supported only on Internet links. You cannot add LTE, ADSL, and VDSL access types to the same WAN link.
Address assignment	DHCP is selected by default. If you want to provide a static IP address, select STATIC.
Management VLAN ID	Enter a VLAN ID for the WAN link.
PPPoE	Click the toggle button to enable authenticated address assignment for the WAN link by using PPPoE (Point-to-Point Protocol over Ethernet).

Table 43: Device Settings (Add Device Site) *(Continued)*

Field	Guideline
Boot Image	This field is displayed only if ZTP is enabled. If you want to upgrade the next-generation firewall device with the latest supported Junos OS version, select the boot image from the list. The boot image is used to upgrade the device when CSO starts the zero touch provisioning (ZTP) process. If you don't specify a boot image, which is the default option (**Use Image on Device**) in the list, then the CSO skips the procedure to upgrade the device during ZTP.
Device Template	You must choose the device template that you want to use for the site from the carousel. For NGFW, the following predefined templates are available. • **SRX_Standalone_Pre_Staged_NonZTP**—Select this template if you want to use a *brownfield* device, which is a device that already has existing firewall and NAT configurations that you want to import into CSO. If you select this template, CSO *does not* perform ZTP for the site. • **SRX as Security Services CPE**—Select this template if you're using a *greenfield* device, which means that CSO will provision the device. **NOTE**: If modified versions of these templates are available, you can choose a different template.
Device Information	
Secure Log Source Interface	This field displays the default interface to be used for in-band management of the device, If you want to use a different interface, remove the default and select a different interface from the list.

Table 43: Device Settings (Add Device Site) *(Continued)*

Field	Guideline
Firewall Policies	This field is displayed only if you enable Zero Touch Provisioning. By default, CSO applies a default firewall policy to the next-generation firewall device. If you don't want to apply the default policy, select **None**.
NAT Policies	This field is displayed only if you enable Zero Touch Provisioning. By default, CSO applies a default NAT policy to the next-generation firewall device. If you don't want to apply the default policy, select **None**.
Import Policy Configuration	This field is displayed only if you disable Zero Touch Provisioning. Click the toggle button to enable the automatic import of previously configured NAT and firewall policies from the device to CSO, after the site is provisioned. By default, the automatic import of policies is disabled. However, you can import firewall and NAT policies manually using the Import workflow. For more information, see *Importing Policies Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Table 44: Configuration Templates (Branch Site Page)

Field	Guideline
Configuration Templates (Optional)	If you want to deploy additional configuration, you can select one or more configuration templates and set the parameters for each template. For each configuration template that you select: 1. Select one or more configuration templates from the list that you want to deploy on the device. 2. Click **Set Parameters**. The Device Configurations page appears. The names and configuration parameters of the configuration templates that you selected are displayed in the Configuration tab. 3. For each configuration template, enter values for the parameters. 4. (Optional) Click the **Summary** tab to view the Junos OS configuration commands that will be deployed on the device for the different configuration templates. 5. Click **Save**. You are returned to the Configuration Templates tab. The Junos OS configuration commands will be deployed on the device.

See "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Deploy a Firewall Policy

To deploy a firewall policy:

1. Select **Configuration > Firewall > Firewall Policy**.
   The Firewall Policy page appears
2. Click the name of the firewall policy that you want to deploy.
   You are taken to the *Firewall-Policy-Name* page.
3. Click **Deploy**.
   The Deploy page appears.
4. In the **Choose Deployment Time** field, select:
   - **Run now** to trigger the deployment of the policy immediately.

- **Schedule at a later time** to schedule the deployment for later.

  If you schedule the deployment for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur. You specify the time in the local time zone of the client from which you access the CSO GUI.

5. Click **OK**

   You are returned to the Firewall Policy page and a job to deploy the policy is triggered. You can check the status of the deployment on the Jobs page (**Monitor > Jobs**). When the job completes successfully, it means that the firewall policy was deployed. The Undeployed field on the *Firewall-Policy-Name* page should be 0.

See "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Deploy a NAT Policy

To deploy a NAT policy:

1. Select **Configuration > NAT > NAT Policy**.

   The NAT Policy page appears

2. Click the name of the NAT policy that you want to deploy.

   You are taken to the *NAT-Policy-Name* page.

3. Click **Deploy**.

   The Deploy page appears.

4. In the **Choose Deployment Time** field, select:

   - **Run now** to trigger the deployment of the policy immediately.
   - **Schedule at a later time** to schedule the deployment for later.

     If you schedule the deployment for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur. You specify the time in the local time zone of the client from which you access the CSO GUI.

5. Click **OK**

   You are returned to the NAT Policy page and a job to deploy the policy is triggered. You can check the status of the deployment on the Jobs page (**Monitor > Jobs**). When the job completes successfully, it means that the NAT policy was deployed. The Undeployed field on the *NAT-Policy-Name* page should be 0.

See "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Configure Unified Threat Management (UTM) in CSO

IN THIS SECTION

- Explanation of Procedure | 190
- Configure UTM Settings for a Tenant | 191
- Add UTM Profiles | 193
- Add Web Filtering Profiles | 197
- Add Antivirus Profiles | 203
- Add Antispam Profiles | 208
- Add Content Filtering Profiles | 209
- Add URL Patterns | 213
- Add URL Categories | 214

Unified threat management (UTM) consolidates several security features to protect against multiple threat types. CSO allows you to configure antispam, antivirus, Web filtering, and content filtering profiles as part of a single UTM profile. You can then reference the UTM profile in a firewall policy intent and deploy the firewall policy to apply UTM, thereby protecting the site from multiple threat types. For more information about UTM, see *UTM Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Explanation of Procedure

The high-level workflow for UTM in CSO is as follows:

1. (Optional) Configure UTM settings that are applicable to all sites in a tenant. See "Configure UTM Settings for a Tenant" on page 191.
2. CSO provides predefined Web filtering, antivirus, antispam, and content filtering profiles that you can use in the UTM profile. If you want to use customized profiles:
   - (Optional) Add a Web filtering profile. See "Add Web Filtering Profiles" on page 197.
   - (Optional) Add an antivirus profile. See "Add Antivirus Profiles" on page 203.

- (Optional) Add an antispam profile. See "Add Antispam Profiles" on page 208.
- (Optional) Add a content filtering profile. See "Add Content Filtering Profiles" on page 209.

3. CSO also provides predefined UTM profiles that you can use. You can also add customized UTM profiles. See "Add UTM Profiles" on page 193.
4. Add a firewall policy and reference the UTM profile in a firewall policy intent, and deploy the firewall policy. See "Add and Deploy Firewall Policies" on page 238.

Configure UTM Settings for a Tenant

You can configure unified threat management (UTM) antispam, antivirus, and Web filtering settings for a tenant that are applicable to all sites belonging to a tenant. The settings are pushed to all those sites to which a firewall policy intent with UTM enabled is applicable.

To configure UTM settings for a tenant:

1. Select **Configuration > UTM > UTM Settings**.

   The Edit UTM Settings page appears.
2. Complete the configuration according to the guidelines provided in Table 45 on page 191.
3. Do one of the following:
   - Click **Reset** to reset the settings to the previously saved configuration.
   - Click **OK** to save the changes.

     The settings are saved and a confirmation message is displayed. You can now navigate away from this page.

Table 45: Edit UTM Settings

Setting	Guideline
Antispam Settings	Specify the antispam settings for the tenant.
Address Whitelist	Select the URL pattern to be used as the antispam allow list. Alternatively, click **Create a New URL Pattern** to add a new URL pattern to use as an allowlist. For more information, see "Add URL Patterns" on page 213.

Table 45: Edit UTM Settings *(Continued)*

Setting	Guideline
Address Blacklist	Select the URL pattern to be used as the antispam block list. Alternatively, click **Create a New URL Pattern** to add a new URL pattern to use as a block list.
Antivirus Settings	Specify the antispam settings for the tenant.
MIME Whitelist	Enter one or more MIME types (separated by commas) to include as part of the MIME allow list; these MIME types are excluded from antivirus scanning.
Exception MIME Whitelist	Enter one or more MIME types (separated by commas) that are to be excluded from the list of MIME types specified as part of the MIME allow list. This list is a subset of the MIME types that you specified in the MIME allow list. For example, if you specify **video/** in the allow list and **video/x-shockwave-flash** in the exception allow list, all objects of MIME type **video/** except MIME type **video/x-shockwave-flash** are excluded from antivirus scanning.
URL Whitelist	Select a URL category (that contains one or more URLs) that you want the antivirus to allow or select **None** if you don't want to add any URLs to the allow list.. Alternatively, click **Create a New URL Category** to add a new URL category to use as an allow list. For more information, see "Add URL Categories" on page 214
Web Filtering Settings	Specify the Web filtering settings for the tenant.
URL Whitelist	Select a URL category (that contains one or more URLs) that you want the Web filtering system to allow or select **None** if you don't want to add any URLs to the allow list. Alternatively, click **Create a New URL Category** to add a new URL category to use as an allow list.

Table 45: Edit UTM Settings *(Continued)*

Setting	Guideline
URL Blacklist	Select a URL category (that contains one or more URLs) that you want the Web filtering system to add to the block list or select **None** if you don't want to add any URLs to the block list. Alternatively, click **Create a New URL Category** to add a new URL category to use as a block list.
Site Reputation	Use the slider to specify the site reputation ranges (for the tenant) for different site categories: • Harmful • Suspicious • Fairly-safe • Moderately-safe • Very-safe

Add UTM Profiles

To add a UTM profile:

1. Select **Configuration > UTM > UTM Profiles** in Customer Portal.
   The UTM Profiles page appears.
2. Click the add icon (**+**).
   The Create UTM Profiles wizard appears.
3. Complete the configuration according to the guidelines provided in Table 46 on page 194.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. (Optional) Review the settings in the Summary section, and modify the settings, if required.
5. Click **Finish**.

You are returned to the UTM Profiles page, and a confirmation message is displayed when the UTM profile is added successfully. After you add a UTM profile, you can assign it to a firewall policy intent.

Table 46: Create UTM Profiles Settings

Setting	Guideline
General	
Name	Enter a unique name for the UTM profile. The name can contain alphanumeric characters, hyphens, or underscores and cannot exceed 29 characters.
Description	Enter a description for the UTM profile.
Traffic Options	In an attempt to consume all available resources, a malicious user might generate a large amount of traffic all at once. To prevent such activity from succeeding, you can impose traffic options.
Connection Limit per Client	For client connections on the device, enter the connection limit per client. The default is 2000; enter 0 to indicate that there is no connection limit.
Action when connection limit is reached	Specify the action that must be taken when the connection limit is reached: • **No action**—Don't take any action. This is the default setting. • **Log and permit**—Log the event and permit the connection. • **Block**—Block the connection. Click **Next** to continue.
Web Filtering	
Web Filtering By Traffic Protocol	You can click **Create Another Profile** to add a Web filtering profile that you can then assign. See "Add Web Filtering Profiles" on page 197.
HTTP	Select the Web filtering profile to be applied for HTTP traffic, or select **None** if you don't want to apply a Web filtering profile. Click **Next** to continue.

Table 46: Create UTM Profiles Settings *(Continued)*

Setting	Guideline
Antivirus	
Antivirus Profiles by Traffic Protocol	You can click **Create Another Profile** to add an antivirus profile that you can then assign. See "Add Antivirus Profiles" on page 203.
Apply to all protocols	Click the toggle button to enable the application of a single antivirus profile to all traffic protocols. You must then specify the profile in the **Default Profile** field. If you disable this toggle button, which is the default, you can specify antivirus profiles for each traffic type.
Default Profile	If you specified that a single antivirus profile should be applied to all traffic protocols, select the antivirus profile. Click **Next** to continue.
HTTP	Select the antivirus profile to be applied to HTTP traffic.
FTP Upload	Select the antivirus profile to be applied to FTP upload traffic.
FTP Download	Select the antivirus profile to be applied to FTP download traffic.
IMAP	Select the antivirus profile to be applied to Internet Message Access Protocol (IMAP)traffic.
SMTP	Select the antivirus profile to be applied to SMTP traffic.
POP3	Select the antivirus profile to be applied to Post Office Protocol 3 (POP3) traffic. Click **Next** to continue.
Antispam	

Table 46: Create UTM Profiles Settings *(Continued)*

Setting	Guideline
Antispam Profiles by Traffic Protocol:	You can click **Create Another Profile** to add an antispam profile that you can then assign. See "Add Antispam Profiles" on page 208.
SMTP	Select the antispam profile to be applied for SMTP traffic. Click **Next** to continue.
Content Filtering	
Content Filtering Profiles by Traffic Protocol:	You can click **Create Another Profile** to add a content filtering profile that you can then assign. See "Add Content Filtering Profiles" on page 209.
Apply to all protocols	Click the toggle button to enable the application of a single content filtering profile to all traffic protocols. You must then specify the profile in the **Default Profile** field. If you disable this toggle button, which is the default, you can specify content filtering profiles for each traffic type.
Default Profile	If you specified that a single antivirus profile should be applied to all traffic protocols, select the antivirus profile. Click **Next** to continue.
HTTP	Select the content filtering profile to be applied to HTTP traffic.
FTP Upload	Select the content filtering profile to be applied to FTP upload traffic.
FTP Download	Select the content filtering profile to be applied to FTP download traffic.
IMAP	Select the content filtering profile to be applied to IMAP traffic.
SMTP	Select the content filtering profile to be applied to SMTP traffic.

Table 46: Create UTM Profiles Settings *(Continued)*

Setting	Guideline
POP3	Select the content filtering profile to be applied to POP3 traffic. Click **Next** to continue.

Add Web Filtering Profiles

Web filtering profiles enable you to manage Internet usage by preventing access to inappropriate Web content over HTTP.

To add a Web filtering profile:

1. Select **Configuration > UTM > Web Filtering Profiles** in Customer Portal.

   The Web Filtering Profiles page appears.
2. Click the add icon (**+**).

   The Create Web Filtering Profiles wizard appears.
3. Complete the configuration according to the guidelines provided in Table 47 on page 197.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. (Optional) Review the settings in the Summary section, and modify the settings, if required.
5. Click **Finish**.

   You are returned to the Web Filtering Profiles page, and a confirmation message is displayed when the Web filtering profile is added successfully. After you add a Web filtering profile, you can associate with a UTM profile.

Table 47: Create Web Filtering Profiles Settings

Setting	Guideline
General	
General Information	

Table 47: Create Web Filtering Profiles Settings *(Continued)*

Setting	Guideline
Name	Enter a unique name for the Web filtering profile. The name can contain alphanumeric characters, hyphens, or underscores and cannot exceed 29 characters.
Description	Enter a description for the Web filtering profile.
Timeout	Enter the time (in seconds) to wait for a response from the Websense server. The default is 15 seconds and the maximum is 1800 seconds.
Engine Type	Select an engine type for Web filtering: • **Juniper Enhanced**—UTM-enhanced Web filtering engine. This is the default engine. • **Websense Redirect**—Use the Websense Redirect Web filtering engine. • **Local**—Use the local Web filtering engine For more information, see Web Filtering Overview.
Safe Search	Safe search ensures that embedded objects, such as images on the URLs received from the search engines, are safe and that undesirable content is not returned to the client. This setting is available only for the Juniper Enhanced engine type and is enabled by default. Click the toggle button to disable safe search redirects. **NOTE**: Safe search redirect supports only HTTP and you cannot extract the URL for HTTPS. Therefore, it is not possible to generate a redirect response for HTTPS search URLs.

Table 47: Create Web Filtering Profiles Settings *(Continued)*

Setting	Guideline
Custom Block Message/URL	Specify the redirect URL or a custom message to be sent when HTTP requests are blocked. The maximum length is 512 characters. **NOTE**: If a message begins with http: or https:, the message is considered a block message URL. Messages that begin with values other than http: or https: are considered custom block messages.
Custom Quarantine Message	For Juniper Enhanced or local engine types, define a custom message to allow or deny access to a blocked site based on a user's response to the message. The maximum length is 512 characters. The quarantine message contains the following information: • URL name • Quarantine name • Category (if available) • Site reputation (if available) Click **Next** to continue.
Account	Specify the user account associated with the Websense Redirect engine.
Server	Specify the hostname or IP address for the Websense server.
Port	Specify the port number to use to communicate with the Websense server. The default port value is 15,868.

Table 47: Create Web Filtering Profiles Settings *(Continued)*

Setting	Guideline
Sockets	Enter the number of sockets used for communication between the client and the Websense server. The default value is 8. Click **Next** to continue.
URL Categories	
Deny Action List	Click the **Add URL Categories** link (next to the text box) to specify a list of URL categories that should be denied access. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 48 on page 203. The list of URL categories selected is displayed in the text box.
Log & Permit Action List	Click the **Add URL Categories** link (next to the text box) to specify a list of URL categories that are logged and then permitted. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 48 on page 203. The list of URL categories selected is displayed in the text box.
Permit Action List	Click the **Add URL Categories** link (next to the text box) to specify a list of URL categories that should be permitted access. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 48 on page 203. The list of URL categories selected is displayed in the text box.

Table 47: Create Web Filtering Profiles Settings *(Continued)*

Setting	Guideline
Quarantine Action List	Click the **Add URL Categories** link (next to the text box) to specify a list of URL categories that should be quarantined. The Select URL Categories page appears. Complete the configuration according to the guidelines provided in Table 48 on page 203. The list of URL categories selected is displayed in the text box. Click **Next** to continue.
Fallback Options	

Table 47: Create Web Filtering Profiles Settings *(Continued)*

Setting	Guideline
Global Reputation Actions	Enhanced Web filtering intercepts HTTP and HTTPS requests and sends the HTTP URL or the HTTPS source IP to the Websense ThreatSeeker Cloud (TSC). The TSC categorizes the URL into one of the predefined categories and also provides site reputation information for the URL to the device. The device determines if it can permit or block the request based on the information provided by the TSC. By default, URLs can be processed using their reputation score if there is no URL category available. You can click the toggle button to disable global reputation actions or select the action to take for the uncategorized URLs based on their reputation score: • **Very Safe**—Permit, log and permit, block, or quarantine a request if a site reputation of 90 through 100 is returned. By default, Permit is selected. • **Moderately Safe**—Permit, log and permit, block, or quarantine a request if a site reputation of 80 through 89 is returned. By default, Log and Permit is selected. • **Fairly Safe**—Permit, log and permit, block or quarantine a request if a site-reputation of 70 through 79 is returned. By default, Log and Permit is selected. • **Suspicious**—Permit, log and permit, block, or quarantine a request if a site reputation of 60 through 69 is returned. By default, Quarantine is selected. • **Harmful**—Permit, log and permit, block, or quarantine a request if a site reputation of zero through 59 is returned. By default, Block is selected.

Table 47: Create Web Filtering Profiles Settings *(Continued)*

Setting	Guideline
Default Action	Choose the actions to be taken for URL categories with no assigned action and for uncategorized URLs. This is used only if no reputation action is assigned.
Fallback Action	Select the fallback action, which is used when: • The ThreatSeeker Websense Cloud servers are unreachable. • A timeout occurs for requests to ThreatSeeker Cloud. • There are too many requests to be handled by the device. Click **Next** to continue.

Table 48: Select URL Categories Settings

Setting	Guideline
Show	Choose which URL categories should be displayed for selection: **All categories**, **Custom URL categories**, or **Websense URL categories**. The first column of the URL Categories field displays URL categories based on your selection.
URL Categories	Select one or more URL categories in the first column and click the forward arrow to confirm your selection. The selected URL categories are displayed in the second column. Click **OK**. You are returned to the Create Web Filtering Profiles page.

Add Antivirus Profiles

The *antivirus* profile defines the content to scan for any malware and the action to be taken when malware is detected. You can add an antivirus profile and then assign it to a UTM profile.

To add an antivirus profile:

1. Select **Configuration > UTM > Antivirus Profiles** in Customer Portal.

   The Antivirus Profiles page appears.

2. Click the add icon (**+**).

   The Create Antivirus Profiles wizard appears.

3. Complete the configuration according to the guidelines provided in Table 49 on page 204.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. (Optional) Review the settings in the Summary section, and modify the settings, if required.

5. Click **Finish**.

   You are returned to the Antivirus Profiles page, and a confirmation message is displayed when the antivirus profile is added successfully. After you add an antivirus profile, you can associate with a UTM profile.

Table 49: Create Antivirus Profile Settings

Setting	Guideline
General	
General Information	
Name	Enter a unique name for the antivirus profile. The name can contain alphanumeric characters, hyphens, or underscores and cannot exceed 29 characters.
Description	Enter a description for the antivirus profile.
Engine Type	Displays the engine type used for scanning. Currently, **Sophos** is the only antivirus engine supported. Sophos antivirus is an in-the-cloud antivirus solution. The virus and malware database is located on external servers maintained by Sophos (Sophos Extensible List) servers, thus there is no need to download and maintain large pattern databases on the Juniper Networks device. Click **Next** to continue.

Table 49: Create Antivirus Profile Settings *(Continued)*

Setting	Guideline
Fallback Options	Fallback options are used when the antivirus system experiences errors and must fall back to one of the previously configured actions to either deny (block) or permit the object. You can specify the fallback options to use when there is a failure, or select the default action if no specific options are to be configured:
Content Size	Select the action to be taken on the content(**None**, **Log and Permit**, or , **Block** [default]) if the content size exceeds the defined limit.
Content Size Limit	Enter the content size limit, in kilobytes (KB), based on which action is taken. 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.
Engine Error	Select the action to take (**None**, **Log and Permit**, or , **Block** [default]) when an engine error occurs. The term *engine error* refers all engine errors, including engine not ready, timeout, too many requests, password protected, corrupt file, decompress layer, and out of resources.
Default Action	Select the default action to take (**None**, **Log and Permit**, or , **Block** [default]) when an engine error occurs. Click **Next** to continue.
Notification Options	Use the notification options to configure a method of notifying the user when a fallback occurs (block or non-block) or when a virus is detected:

Table 49: Create Antivirus Profile Settings *(Continued)*

Setting	Guideline

Fallback Deny	Click the toggle button to enable fallback notifications to e-mail senders when their messages are blocked. By default, fallback block notifications are disabled. If you enable notifications, you can configure the following additional parameters:
	• **Notification Type**—Select the type of notification to be sent:
	  • **None**—Don't sent notifications.
	  • **Protocol**—Send a protocol-specific notification. With protocol-only notifications, a protocol-specific error code might be sent.
	  • **Message**—Send a generic notification.
	• **Custom Message Subject**—Enter the subject line that you want to send for the e-mail notification when a block occurs.
	• **Custom Message**—Enter the text of the e-mail to be sent for the e-mail notification when a block occurs.
	• **Display Hostname**—Click the toggle button to enable the display of the computer hostname in the notification e-mail sent to the administrator when a block occurs.
	• **Allow E-Mail**—Click the toggle button to enable e-mail notification to notify a specified administrator when a block occurs.
	• **Administrator E-Mail Address**—Enter the administrator e-mail address that will be notified when a block occurs.

Table 49: Create Antivirus Profile Settings *(Continued)*

Setting	Guideline
Fallback Non-Deny	Click the toggle button to enable fallback notifications to e-mail senders when their messages are not blocked. By default, fallback unblock notifications are disabled. If you enable notifications, you can configure the following additional parameters: • **Custom Message Subject**—Enter the subject line that you want to send for the e-mail notification. • **Custom Message**—Enter the text of the e-mail to be sent for the e-mail notification.

Virus Detected	Click the toggle button to enable notifications to e-mail senders when a virus is detected. By default, notifications are disabled. If you enable notifications, you can configure the following additional parameters:
	• **Notification Type**—Select the type of notification to be sent:
	  • **None**—Don't sent notifications.
	  • **Protocol**—Send a protocol-specific notification. With protocol-only notifications, a protocol-specific error code might be sent.
	  • **Message**—Send a generic notification.
	• **Custom Message Subject**—Enter the subject line that you want to send for the e-mail notification when a virus is detected.
	• **Custom Message**—Enter the text of the e-mail to be sent for the e-mail notification when a virus is detected.
	Click **Next** to continue.

Add Antispam Profiles

E-mail spam consists of unwanted e-mail messages usually sent by commercial, malicious, or fraudulent entities. When the device detects an e-mail message deemed to be spam, it either blocks the message or tags the message header or subject field with a preprogrammed string. *Antispam* filtering allows you to use a third-party server-based spam block list (SBL) and to optionally add your own local allow lists (benign) and block lists (malicious) for filtering against e-mail messages.

NOTE: Sophos updates and maintains the IP-based SBL. Antispam is a separately licensed subscription service.

To add an antispam profile:

1. Select **Configuration > UTM > Antispam Profiles** in Customer Portal.

   The Antispam Profiles page appears.
2. Click the add icon (**+**).

   The Create Antispam Profiles wizard appears.
3. Complete the configuration according to the guidelines provided in Table 50 on page 208.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the Antispam Profiles page, and a confirmation message is displayed when the antispam profile is added successfully. After you add an antispam profile, you can associate with a UTM profile.

Table 50: Create Antispam Profile Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the antispam profile. The name can contain alphanumeric characters, hyphens, or underscores and cannot exceed 29 characters.
Description	Enter a description for the antispam profile.

Table 50: Create Antispam Profile Settings *(Continued)*

Setting	Guideline
Sophos Blacklist	Click the toggle button to enable the use of server-based spam filtering. If the toggle button is disabled, which is the default, local spam filtering is used. Server-based antispam filtering requires Internet connectivity with the SBL server. Domain Name Service (DNS) is required to access the SBL server. The firewall performs SBL lookups using the DNS protocol. **NOTE**: Server-based spam filtering supports only IP-based spam block list blocklist lookup. Sophos updates and maintains the IP-based spam block list. Server-based antispam filtering is a separately licensed subscription service.
Action	
Default Action	Select the action to be taken when spam is detected: • **Tag Email Subject Line**—Tag the subject line of the e-mail with the configured custom tag. • **Tag SMTP Header**—Tag the SMTP header of the e-mail with the configured custom tag. • **Block Email**—Block the e-mail. • **None**—Don't take any action.
Custom Tag	Enter the tag to use for identifying a message as spam. The maximum length is 512 characters, and the default is ***SPAM***.

Add Content Filtering Profiles

Content filtering blocks or permits certain types of traffic based on the MIME type, file extension, and protocol command. The content filter controls file transfers across the device by checking traffic against configured filter lists. Table 51 on page 210 displays the types of content filters that you can configure as part of a content filtering profile.

NOTE: The content filtering profile evaluates traffic before all other UTM profiles. Therefore, if traffic meets criteria configured in the content filter, the content filter acts first upon this traffic.

Table 51: Supported Content Filter Types

Type	Description
MIME pattern filter	MIME patterns are used to identify the type of traffic in HTTP and MAIL protocols. There are two lists of MIME patterns that are used by the content filter to determine the action to be taken. The block MIME list contains a list of MIME type traffic that is to be blocked. The MIME exception list contains MIME patterns that are not to be blocked by the content filter and are generally subsets of items on the block list. **NOTE**: The exception list has a higher priority than the block list.
Block Extension List	Because the name of a file is available during the transfers, using file extensions is a highly practical way to block or allow file transfers. All protocols support the use of the block extension list.
Protocol Command Block and Permit Lists	Different protocols use different commands to communicate between servers and clients. By blocking or allowing certain commands, traffic can be controlled on the protocol command level. The block or permit command lists are intended to be used in combination, with the permit list acting as an exception list to the block list. **NOTE**: If a protocol command appears on both the permit list and the block list, the command is permitted.

To add a content filtering profile:

1. Select **Configuration > UTM > Content Filtering Profiles** in Customer Portal.

   The Content Filtering Profiles page appears.
2. Click the add icon (**+**).

   The Create Content Filtering Profiles wizard appears.
3. Complete the configuration according to the guidelines provided in Table 52 on page 211.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. (Optional) Review the settings in the Summary section, and modify the settings, if required.
5. Click **Finish**.

You are returned to the Content Filtering Profiles page, and a confirmation message is displayed when the content filtering s profile is added successfully. After you add a content filtering profile, you can associate with a UTM profile.

Table 52: Create Content Filtering Profiles Settings

Setting	Guideline
General	
General Information	
Name	Enter a unique name for the content filtering profile. The name can contain alphanumeric characters, hyphens, or underscores and cannot exceed 29 characters.
Description	Enter a description for the content filtering profile.
Notification Options	
Notify Mail Sender	Click the toggle button to enable a notification when a content filter is matched. Notifications are disabled by default.
Notification Type	Select the type of notification to send: • **None**—Don't sent notifications. • **Protocol**—Send a protocol-specific notification. With protocol-only notifications, a protocol-specific error code might be sent. • **Message**—Send a generic notification.
Custom Notification Message	Enter a custom notification message. The maximum length is 512 characters. Click **Next** to continue.
Filter Settings	
Protocol Commands	

Table 52: Create Content Filtering Profiles Settings *(Continued)*

Setting	Guideline
Command Block List	Enter the protocol commands to be blocked for the HTTP, FTP, SMTP, IMAP, and POP3 protocols. Use commas to separate each command. Protocol commands allow you to control traffic at the protocol-command level.
Command Permit List	Enter specific commands to be permitted for the HTTP, FTP, SMTP, IMAP, and POP3 protocols. Use commas to separate each command.
Block Content Type	Use the content filter to block other types of harmful files that the MIME type or the file extension cannot control. Select one or more of the following types of content blocking (supported only for HTTP): • **ActiveX** • **Windows executables (.exe)** • **HTTP cookie** • **Java applet** • **ZIP files**
Extension Block List	You use a file extension list to define a set of file extensions to block over HTTP, FTP, SMTP, IMAP, and POP3. Enter file extensions to block separated by commas. For example, exe, pdf, js, and so on.
MIME Block List	Enter the MIME types that you want to block over HTTP, FTP, SMTP, IMAP, and POP3 connections. Use commas to separate each MIME type.
MIME Permit List	Enter the MIME types that you want to permit over HTTP, FTP, SMTP, IMAP, and POP3 connections. Use commas to separate each MIME type. Click **Next** to continue.

Add URL Patterns

You can add URL patterns, and, optionally, assign URL patterns to a URL category.

To add a URL pattern:

1. Select **Configuration > UTM > URL Patterns** in Customer Portal.

   The URL Patterns page appears.
2. Click the add icon (**+**).

   The Create URL Patterns page is displayed.
3. Complete the configuration according to the guidelines provided in Table 53 on page 213.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the URL Patterns page, and a confirmation message is displayed when the URL pattern is added.

Table 53: Create URL Patterns Settings

Settings	Guidelines
Name	Enter a unique name for the URL pattern. The name must begin with a letter or an underscore (_) and can contain alphanumeric character, hyphens, and underscores. The maximum length is 29 characters.
Description	Enter a description for the URL pattern. The maximum length is 255 characters.
URL Category	Select the URL category to which you want to assign the URL pattern. Alternatively, click **Create New URL Category** to add a URL category, enter the URL category name in the text box, and click **Save** to assign the URL pattern to the new URL category.

Table 53: Create URL Patterns Settings *(Continued)*

Settings	Guidelines
[Add URLs]	Click the add (**+**) icon, enter the URL in the inline text box that appears in the table, and click **✓** (check mark) to save the URL. You can enter additional URLs if needed. **NOTE**: • The following wildcard characters are supported: asterisk (*), period (.), square brackets ([]), question mark (?) • Precede all wildcard characters with http://. • The asterisk (*) can only be used at the beginning of a URL and must be followed by a period (.). • The question mark (?) can only be used at the end of a URL. • The following are examples of wildcard syntaxes that are supported: http://*.example.net, http://www.example.ne?, and http://www.example.n??. • The following are examples of wildcard syntaxes that are not supported: *.example.???, http://*example.net, http://?, and www.example.ne?.

Add URL Categories

A URL category is a list of URL patterns grouped under a single title.

To add a URL category:

1. Select **Configuration > UTM > URL Categories** in Customer Portal.
   The URL Categories page appears.
2. Click the add icon (**+**).
   The Create URL Categories page is displayed.
3. Complete the configuration according to the guidelines provided in Table 54 on page 215.

   NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **OK**.

You are returned to the URL Categories page, and a confirmation message is displayed when the URL category is added.

Table 54: Create URL Categories Settings

Settings	Guidelines
Name	Enter a unique name for the URL category. The name must begin with a letter or an underscore (_) and can contain alphanumeric characters, hyphens, and underscores. The maximum length is 59 characters.
Description	Enter a description for the URL category.
URL Patterns	Select one or more URL patterns and click the forward arrow (**>**) to confirm your selection. The selected URL patterns are displayed in the column on the right. Alternatively, click **Create a New Pattern** to add a URL pattern and assign it to the URL category. For more information, see "Add URL Patterns" on page 213.

RELATED DOCUMENTATION

CSO Next-Generation Firewall (NFGW) Deployment Workflow | 178

Configure and Deploy SSL Proxy Policy in CSO

IN THIS SECTION

- Explanation of Procedure | 216
- Import a Certificate | 217
- Install a Certificate | 219
- Add SSL Forward Proxy Profiles | 219
- Add SSL Proxy Policy Intents | 225
- Deploy an SSL Proxy Policy | 227

SSL proxy is enabled as an application service within a security policy. You specify the traffic that you want the SSL proxy enabled on as match criteria and then specify the SSL proxy profile to be applied to the traffic. For more information, see *SSL Forward Proxy Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Explanation of Procedure

The following is the workflow to configure and deploy an intent-based SSL forward proxy policy in CSO:

1. Obtain the root certificate and private key from your trusted certificate authority (CA).
2. Combine the root certificate and private key into a single file.
3. Import the certificate and private key file. See "Import a Certificate" on page 217
4. (Optional) Install the imported certificate on one or more sites. See "Install a Certificate" on page 219.
5. By default, Juniper Networks ships trusted certificates for sites that use HTTPS. These certificates are installed automatically by CSO when the site is successfully provisioned.

   If you want to use additional trusted certificates, import and install the certificates as explained in Steps 3 and 4.
6. Add an SSL proxy profile. See "Add SSL Forward Proxy Profiles" on page 219.

   NOTE:
   - Use the imported root certificate when you add the SSL proxy profile.
   - For trusted certificates, specify that all trusted certificates on the device are used.
7. Add an SSL proxy policy intent that uses the SSL proxy profile that you added. See "Add SSL Proxy Policy Intents" on page 225.
8. Deploy the SSL proxy policy. See "Deploy an SSL Proxy Policy" on page 227.

   NOTE:
   - Understanding How SSL Proxy Policy Intents Are AppliedEnsure that the root and trusted certificates are imported into CSO before the policy is deployed.
   - If you have not installed the certificates referenced in the SSL proxy profile, then they are automatically installed when the SSL proxy policy is deployed.
9. For Internet access from an SRX Series device by using the SSL proxy, ensure that you import the root certificate (obtained in Step 1) into the browsers of the clients accessing the Internet.

NOTE: If you do not import the certificate, the traffic does not go through for clients in the LAN segments.

For examples of how SSL proxy policy intents are applied, see *Understanding How SSL Proxy Policy Intents Are Applied* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

Import a Certificate

NOTE: If you want to use the SSL proxy feature in CSO, you must import at least one root certificate for a tenant. The certificate can then be installed in one or more sites.

To import a certificate:

1. Select **Administration > Certificate Management > Certificates** in Customer Portal.

   The Certificates page appears.
2. Select **More > Import Certificate**.

   The Import Certificate page appears.
3. Complete the configuration according to the guidelines provided in Table 55 on page 217.

   NOTE: Fields marked with * are mandatory.
4. Click **OK**.

   You are returned to the Certificates page. If the certificate content that you imported is validated successfully, a confirmation message is displayed; if not, an error message is displayed.

After importing a certificate, you can use it when you add an SSL proxy profile.

Table 55: Import Certificate Settings

Setting	Guideline
Certificate Name	Enter the certificate name, which must be a unique string of alphanumeric characters and some special characters (_ -). No spaces are allowed and the maximum length is 32 characters.

Table 55: Import Certificate Settings *(Continued)*

Setting	Guideline
Certificate Type	Select an option to specify whether the certificate that you are importing is a root certificate (**Root CA**) or a trusted certificate (**Trusted CA**).
Passphrase	Enter the passphrase to protect the private key or key pair of the Privacy-Enhanced Mail (PEM) certificate file.
Description	Enter a description for the certificate.
Certificate Content	Select whether you want to import the certificate content from a file or if you want to paste the certificate content. **NOTE**: • The following certificate file extensions are supported: **.cert**, **.pem**, and **.txt**. • The certificate content must be in the X.509 ASCII format. • If you're importing a root certificate, then the both the certificate content and private key must be specified.
File Path for Certificate	To import the certificate content from a file, click **Browse**. In the File Upload dialog that appears, select the certificate file and click **Open**. The filename of the file that you uploaded is displayed.
Paste Certificate Content	To paste the certificate content directly from a file, open the certificate file in a text editor, copy the certificate content, and paste it in the text box.

The following is an example of root certificate content.

```
-----BEGIN PRIVATE KEY-----
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123AbcXyz123A
-----END PRIVATE KEY-----
```

```
-----BEGIN CERTIFICATE-----
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456DefQrs456A
-----END CERTIFICATE-----
```

Install a Certificate

After you import a certificate into CSO, you can install the certificates on one or more sites.

To install a certificate:

1. Select **Administration > Certificate Management > Certificates** in Customer Portal.

   The Certificates page appears.
2. Select the certificate that you want to install, and then select **More > Install Certificate**. Alternatively, right-click a certificate and select **Install Certificate**.

   The Install Certificate page appears, displaying a list of sites.
3. Select the sites on which you want to install the certificate.
4. Click **Install**.

   You are returned to the Certificates page. A job is triggered and a confirmation message appears with the ID of the job. Click the job ID to go to the Jobs page, where you can view the status of the job.
5. (Optional) After the job completes successfully. you can verify that the certificate was installed on the sites. On the Certificates page, select the certificate and select **More > View Installed Sites**.

   The View Installed Sites page appears listing the sites on which the certificate was installed.

Add SSL Forward Proxy Profiles

To add an SSL forward proxy profile:

NOTE: Ensure that you have a root certificate imported for the tenant before you add an SSL forward proxy profile. You can import SSL certificates (root and trusted) from the Certificates page (**Administration > Certificate Management > Certificates**) and associate the certificates with SSL forward proxy profiles.

1. Select **Configuration > SSL Proxy > Profiles** in Customer Portal.

   The SSL Proxy Profiles page appears.
2. Click the add icon (**+**).

   The Create SSL Proxy Profiles page appears.
3. Complete the configuration according to the guidelines provided in Table 56 on page 220.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the SSL Proxy Profiles page, and a confirmation message is displayed when the SSL proxy profile is added.

   The SSL forward proxy profile can be used in an SSL proxy policy intent.

Table 56: Create SSL Proxy Profile Settings

Setting	Guideline
General Information	
Name	Enter a unique name for the profile, which can contain alphanumeric characters, hyphens, and underscores. No spaces are allowed and the maximum length is 63 characters.
Description	Enter a description for the profile. The maximum length is 255 characters.

Table 56: Create SSL Proxy Profile Settings *(Continued)*

Setting	Guideline
Preferred Cipher	Select a preferred cipher, which enables you to define an SSL cipher that can be used with acceptable key strength: • **None** (Default)—Do not specify a preferred cipher. • **Medium**—Use ciphers with key strength of 128 bits or greater. • **Strong**—Use ciphers with key strength of 168 bits or greater. • **Weak**—Use ciphers with key strength of 40 bits or greater. • **Custom**—Configure a custom cipher suite.

Table 56: Create SSL Proxy Profile Settings *(Continued)*

Setting	Guideline
Custom Ciphers	If you specified a custom preferred cipher, you can define a custom cipher list by selecting one or more ciphers that the SSH server can use to perform encryption and decryption functions: • **None**—No encryption. • **rsa-with-RC4-128-md5**—RSA, 128- bit RC4, MD5 hash • **rsa-with-RC4-128-sha**—RSA, 128-bit RC4, SHA hash • **rsa-with-des-cbc-sha**—RSA, DES/CBC, SHA hash • **rsa-with-3DES-ede-cbc-sha**—RSA, 3DES EDE/CBC, SHA hash • **rsa-with-aes-128-cbc-sha**—RSA, 128-bit AES/CBC, SHA hash • **rsa-with-aes-256-cbc-sha**—RSA, 256 bit AES/CBC, SHA hash • **rsa-export-with-rc4-40-md5**—RSA-export, 40 bit RC4, MD5 hash • rsa-export-with-des40-cbc-sha—RSA-export, 40 bit DES/CBC, SHA hash • **rsa-export1024-with-des-cbc-sha**—RSA 1024 bit export, DES/CBC, SHA hash • **rsa-export1024-with-rc4-56-md5**—RSA 1024 bit export, 56 bit RC4, MD5 hash • **rsa-export1024-with-rc4-56-sha**—RSA 1024 bit export, 56 bit RC4, SHA hash • **rsa-with-aes-256-gcm-sha384**—RSA, 256 bit AES/GCM, SHA384 hash • **rsa-with-aes-256-cbc-sha256**—RSA, 256 bit AES/CBC, SHA256 hash • **rsa-with-aes-128-gcm-sha256**—RSA, 128 bit AES/GCM, SHA256 hash • **rsa-with-aes-128-cbc-sha256**—RSA, 256 bit AES/CBC, SHA256 hash • **ecdhe-rsa-with-aes-256-gcm-sha384**—ECDHE, RSA, 256 bit AES/GCM, SHA384 hash • **ecdhe-rsa-with-aes-256-cbc-sha384**—ECDHE, RSA, 256 bit AES/CBC, SHA384 hash • **ecdhe-rsa-with-aes-256-cbc-sha**—ECDHE, RSA, 256 bit AES/CBC, SHA hash • **ecdhe-rsa-with-aes-3des-ede-cbc-sha**—ECDHE, RSA, 3DES, EDE/CBC, SHA hash

Table 56: Create SSL Proxy Profile Settings *(Continued)*

Setting	Guideline
	• **ecdhe-rsa-with-aes-128-gcm-sha256**—ECDHE, RSA, 128 bit AES/GCM, SHA256 hash • **ecdhe-rsa-with-aes-128-cbc-sha256**—ECDHE, RSA, 128 bit AES/CBC, SHA256 hash • **ecdhe-rsa-with-aes-128-cbc-sha**—ECDHE, RSA, 128 bit AES/CBC, SHA hash
Flow Trace	Click the toggle button to enable flow tracing to enable the troubleshooting of policy-related issues. Flow tracing is disabled by default.
Root Certificate	Select a *root certificate* from the list or click **Add Root Certificate** to import a root certificate. In a public key infrastructure (PKI) hierarchy, the root certificate authority (CA) is at the top of the trust path.
Trusted Certificate Authorities	Choose whether you want to add all trusted certificates present on the device (**All**) or select specific trusted certificates (**Select Specific**). Before establishing a secure connection, the SSL proxy checks CA certificates to verify signatures on server certificates. If you chose to add selected trusted certificates, the existing trusted certificates are displayed. Select one or more certificates by clicking the check boxes, and click the **>** icon. The selected certificates are displayed in the column on the right. Optionally, click **Add Trusted Certificates** to import a trusted certificate. See "Import a Certificate" on page 217. **NOTE**: • Specifying that all trusted certificates should be used means that all trusted certificates on a particular device (site) will be used during SSL policy deployment. • If you specify that all trusted certificates should be used in an SSL forward proxy profile, you must ensure that at least one trusted certificate is installed on the device.

Table 56: Create SSL Proxy Profile Settings *(Continued)*

Setting	Guideline
Exempted Addresses	Exempted addresses include addresses that you want to exempt from undergoing SSL proxy processing. To specify exempted addressees, select one or more addresses in the left column and click the **>** icon to confirm your selection. The selected addresses are then displayed in the right column. These addresses are used to create allow lists that bypass SSL forward proxy processing. Because SSL encryption and decryption are complicated and expensive procedures, network administrators can selectively bypass SSL proxy processing for some sessions. Such sessions typically include connections and transactions with trusted servers or domains with which network administrators are very familiar. There are also legal requirements to exempt financial and banking sites. Such exemptions are achieved by configuring the IP addresses or domain names of the servers under allow lists. **NOTE**: You can also add addresses by clicking **Add New Address**. For more information, see *Creating Addresses or Address Groups* in the *CSO Customer Portal User Guide* (available at the CSO Documentation page).
Exempted URL Categories	Select one or more previously defined URL categories in the left column and click the **>** icon to confirm your selection. The selected addresses are then displayed in the right column. These URL categories are used to create allow lists that bypass SSL forward proxy processing. The selected URL categories are exempted during SSL inspection.
Actions	
Server Auth Failure	Click the toggle button to enable CSO to ignore errors encountered during the server certificate verification process, such as CA signature verification failure, self-signed certificates, and certificate expiry. This toggle button is disabled by default, which means that server authentication errors are not ignored. We do not recommend that you ignore authentication errors because it results in websites not being authenticated at all. However, you can use this option to effectively identify the root cause for dropped SSL sessions.

Table 56: Create SSL Proxy Profile Settings *(Continued)*

Setting	Guideline
Session Resumption	Click the toggle button to enable session resumption. Session resumption is disabled by default. To improve throughput and still maintain an appropriate level of security, SSL session resumption provides a session-caching mechanism so that session information, such as the pre-master secret key and agreed-upon ciphers, can be cached for both the client and server.
Logging	Select one or more events to be logged. You can choose to log all events, warnings, general information, errors, or different sessions (allowed, dropped, or ignored). By default, no events are logged.
Renegotiation	Select one of the following options if a change in SSL parameters requires renegotiation: • **None**—Renegotiation is not required. This is the default setting. • **Allow**—Allow secure and nonsecure renegotiation. • **Allow Secure**—Allow secure negotiation only. • **Drop**—Drop session on renegotiation request. After a session is created and SSL tunnel transport has been established, a change in SSL parameters requires renegotiation. SSL forward proxy supports both secure (RFC 5746) and nonsecure (TLS v1.0 and SSL v3) renegotiation. When session resumption is enabled, session renegotiation is useful in the following situations: • Cipher keys need to be refreshed after a prolonged SSL session. • Stronger ciphers need to be applied for a more secure connection.

Add SSL Proxy Policy Intents

An SSL proxy policy intent enables you to configure an SSL proxy between source and destination endpoints by associating the latter with an SSL proxy profile. You can add an SSL proxy policy intent inline on the SSL Proxy Policy page.

To add an SSL proxy policy intent:

1. Select **Configuration > SSL Proxy > Policy** in Customer Portal.

   The SSL Proxy Policy page appears.
2. Click the add icon (**+**).

   The options to add a policy intent appears inline on the SSL Proxy Policy page.
3. Enter the policy intent information according to the guidelines provided in Table 57 on page 226
4. Click **Save**.

   The SSL proxy policy intent is saved and a confirmation message is displayed. When an SSL proxy policy intent is added, the Undeployed field is incremented by one indicating that intents are pending deployment.

> **NOTE**: After the policy intent is added, you must deploy the policy to ensure that the changes take effect

Table 57: Add SSL Proxy Policy Intent Settings

Setting	Guideline
[Name]	Enter the name of the SSL proxy policy intent in the first text box. If you do not enter a name, the system-generated name is used. The name that you enter must begin with an alphanumeric character and can contain alphanumeric characters and some special characters (- _). The maximum length is 63 characters.
[Description]	Enter the description of the SSL proxy policy intent in the second text box.
Source	Select one or more of the following source endpoints: • IP address or IP address group • Site • Site group • Department The default source for an SSL proxy policy intent is **All Sites**. If you don't add a source, then the default is used. **NOTE**: A source IP address value of **Any** signifies any IP address from any site.

Table 57: Add SSL Proxy Policy Intent Settings *(Continued)*

Setting	Guideline
Destination	Select one or more of the following destination endpoints: • IP address or address group • Site • Site group • Department The default destination for an SSL proxy policy intent is **Internet**. If you don't add a destination, then the default is used. **NOTE**: A destination IP address value of **Any** signifies traffic going to the Internet (any address). Traffic within sites (internal traffic) is not covered by the destination IP address value of **Any**. If you want to cover traffic between two sites, ensure that the sites are included in both the source and destination endpoints.
SSL Proxy Profile	Specify an SSL proxy profile to associate with the SSL proxy policy intent in one of the following ways: • Click the add icon (**+**) and select the SSL proxy profile from the list of previously configured profiles. • Filter the profiles by entering a search term in the **SSL Proxy Profile** field and select a profile. • Add a SSL proxy profile—Click the **Add New Profile** link. See "Add SSL Forward Proxy Profiles" on page 219. • Click the **View more results** link to view additional configured profiles. The list of SSL proxy profiles is displayed in the **End Points** panel on the right. To add a profile, select it and click the check mark icon (✓) that appears when you hover over the profile.

Deploy an SSL Proxy Policy

After you add one or more SSL proxy policy intents, you must deploy the SSL proxy policy.

To deploy an SSL proxy policy:

1. Select **Configuration > SSL Proxy > Policy**.

   The SSL Proxy Policy page appears

2. Click **Deploy**.

   The Deploy page appears.

3. In the **Choose Deployment Time** field, select:

   - **Run now** to trigger the deployment of the policy immediately.
   - **Schedule at a later time** to schedule the deployment for later.

     If you schedule the deployment for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur. You specify the time in the local time zone of the client from which you access the CSO GUI.

4. Click **OK**

   You are returned to the SSL Proxy Policy page and a job to deploy the policy is triggered. You can check the status of the deployment on the Jobs page (**Monitor > Jobs**). When the job completes successfully, it means that the SSL proxy policy was deployed. The Undeployed field on the SSL Proxy Policy page should be 0.

RELATED DOCUMENTATION

CSO Next-Generation Firewall (NFGW) Deployment Workflow | 178

Configure Intrusion Prevention System (IPS) in CSO

IN THIS SECTION

- Explanation of Procedure | 229
- Add IPS Profiles | 230
- Add IPS or Exempt Rules to IPS Profiles | 230

Intrusion prevention system (IPS) signatures are used to monitor and prevent intrusions. IPS compares traffic against signatures of known threats and blocks traffic when a threat is detected.

CSO provides predefined IPS signatures, IPS signature static groups, and IPS signature dynamic groups that you can use in IPS or exempt rules in an IPS profile. However, you cannot modify the predefined signatures and groups. CSO also lets you add customized IPS signatures, static groups, and dynamic groups

CSO also provides predefined IPS profiles that contain predefined IPS rules, both of which can't be modified. You can add customized profiles and add IPS or exempt rules to the profiles. You enable intrusion detection by referencing an IPS profile in a firewall policy intent and deploying the firewall policy.

Explanation of Procedure

The high-level workflow to configure IPS is as follows:

1. On to the IPS Signatures page (**Configuration > IPS > IPS Signatures** in Customer Portal), and review the predefined IPS signatures, signature static groups, and signature dynamic groups to determine if you need to use customized signatures, static groups, or dynamic groups. You can create customized signatures, static groups, or dynamic groups in two ways:
   - Clone a predefined IPS signature, static group, or dynamic group and then modify the cloned signature, static group, or dynamic group.
   - Add a customized signature, static group, or dynamic group by specifying the parameters from scratch.

   For more information, see*About the IPS Signatures Page* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).
2. Go to the IPS Profiles page (**Configuration > IPS > IPS Profiles** in Customer Portal), and review the predefined IPS profiles to determine if you need to use customized IPS profiles and customized rules. You can create customized IPS profiles and rules in two ways:
   - Clone a predefined IPS profile and then modify the cloned profile and the rules within the profile.

     For more information, see *About the IPS Profiles Page* in the *CSO Customer Portal User Guide*.
   - Add a customized IPS profile and then add IPS or exempt rules to that profile. See "Add IPS Profiles" on page 230 and "Add IPS or Exempt Rules to IPS Profiles" on page 230.
3. Use the IPS profile in a firewall policy intent and deploy the firewall policy. See "Add and Deploy Firewall Policies" on page 238.

Add IPS Profiles

Contrail Service Orchestration (CSO) contains predefined intrusion prevention system (IPS) profiles that you can use. You can also add customized IPS profiles from the Create IPS Profile page.

To add a customized IPS profile:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.
2. Click the add (**+**) icon.

   The Create IPS Profile page appears.
3. Complete the configuration according to the guidelines in Table 58 on page 230.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **OK**.

   You are returned to the IPS Profiles page and a confirmation message is displayed indicating that the IPS profile is added.

After you add an IPS profile, you can add one or more IPS or exempt rules to the profile, and then use the IPS profile in a firewall policy intent.

Table 58: Create IPS Profile Settings

Setting	Guideline
Name	Enter a unique name for the IPS profile that is a string of alphanumeric characters and some special characters (colon, hyphen, period, and underscore). No spaces are allowed and the maximum length is 255 characters.
Description	Enter a description for the IPS profile; the maximum length is 255 characters.

Add IPS or Exempt Rules to IPS Profiles

An IPS rule is used to protect your network from attacks by using attack objects to detect known and unknown attacks, based on stateful signature and protocol anomalies. In contrast, an exempt rule works in conjunction with an IPS rule to prevent unnecessary alarms from being generated. If traffic matches

an IPS rule, the system attempts to match the traffic against the exempt rules before performing the action specified.

You can add intrusion prevention system (IPS) rules or exempt rules only to customized IPS profiles.

To add an IPS rule or an exempt rule to a customized IPS profile:

1. Select **Configuration > IPS > IPS Profiles**.

   The IPS Profiles page appears.
2. Click ***IPS-Profile-Name*** for the profile for which you want to add a rule.

   The *IPS-Profile-Name* page appears.
3. You can add IPS rules and exempt rules from this page:
   - To add an IPS rule:
     a. Select **Create > IPS Rule**.

        The parameters for an IPS rule appear inline at the top of the page.

     b. Complete the configuration according to the guidelines in Table 59 on page 232.

        NOTE: Fields marked with an asterisk (*) are mandatory.

     c. Click **Save**.

        The IPS rule is added and a confirmation message appears at the top of the page.
   - To add an exempt rule:
     a. Select **Create > Exempt Rule**.

        The parameters for an exempt rule appear inline at the top of the page.

     b. For exempt rules, you can configure only the following fields:
        - Rule Name
        - Description
        - IPS Signatures

        See Table 59 on page 232 for an explanation of these fields.

     c. Click **Save**.

        The exempt rule is added and a confirmation message appears at the top of the page.

After adding IPS and exempt rules, you can use the IPS profile in a firewall policy intent and deploy the firewall policy, which deploys the IPS and exempt rules associated with the IPS profile.

Table 59: Add IPS Rule Settings

Setting	Guideline
[Name]	CSO generates a unique IPS rule name by default. You can modify the name if needed. The name must begin with an alphanumeric character and can contain alphanumeric characters and some special characters (colons, hyphens, forward slashes, periods, and underscores); 63-character maximum.
[Description]	Enter a description for the IPS rule.
IPS Signatures	You can add one or more IPS signatures and IPS signature static and dynamic groups to be associated with the rule: 1. Click inside the text box with the **+** icon. A list of IPS signatures and IPS signature static and dynamic groups appears. 2. (Optional) Enter a search term and press Enter to filter the list of items displayed. 3. Click a list item to add it to the IPS signatures and IPS signature static or dynamic groups associated with the rule. 4. (Optional) Repeat the preceding step to add more signatures, static groups, and dynamic groups. 5. Click the **View more results** link to view the full list of IPS signatures and IPS signature static and dynamic groups. The full list is displayed in the End Points panel on the right. To add one or more signatures, static groups, or dynamic groups: a. Mouse over a list item and select the check box that appears. b. Repeat the preceding step for the other signatures, static groups, or dynamic groups that you want to add. c. Click the check mark icon (✓) at the top of the End Points panel, and select **Signatures**. The signatures, static groups, or dynamic groups that you selected are added and displayed in the IPS Signatures field.

Table 59: Add IPS Rule Settings *(Continued)*

Setting	Guideline
Actions	Select the action to be taken when the monitored traffic matches the attack objects specified in the rules: • **No Action**—No action is taken. Use this action to only generate logs for some traffic. • **Ignore**—Stops scanning traffic for the rest of the connection if an attack match is found. IPS disables the rulebase for the specific connection. • **Drop Connection**—Drops all packets associated with the connection, preventing traffic for the connection from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing. • **Drop Packet**—Drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a denial of service that prevents you from receiving traffic from a legitimate source-IP address. • **Close Client and Server**—Closes the connection and sends a TCP reset (RST) packet to both the client and the server. • **Close Client**—Closes the connection and sends an RST packet to the client, but not to the server. • **Close Server**—Closes the connection and sends an RST packet to the server, but not to the client. • **Recommended**—Uses the action that Juniper Networks recommends when that attack is detected. All predefined attack objects have a default action associated with them. • **DiffServ Marking**—Assigns the specified differentiated services code point (DSCP) value to the packet in an attack and pass the packet on normally. When you select DiffServ Marking, you must enter a DSCP value as follows: 1. Click the **Code Point: *Vaule*** hyperlink. The Code Point page appears. 2. In the **Code Point** field, enter a DSCP value from 0 through 63. 3. Click **OK**. You are returned to the previous page; the value that you entered is displayed

Table 59: Add IPS Rule Settings *(Continued)*

Setting	Guideline
Additional Actions	In addition to the IPS action, you can configure one or more additional actions.
Notifications	When attacks are detected, you can choose to log the attack, create log records with attack information, and send that information to the log server. To configure notifications: 1. Click the **Notification** link. The Notification page appears. 2. Complete the configuration according to the guidelines shown in Table 60 on page 235. 3. Click **OK**. You are returned to the previous page. A gear icon next to the Notification link indicates that you have configured notification settings.
IP Action	When attacks are detected, you can configure actions that you want IPS to take against future connections that use the same IP address. To configure IP actions: 1. Click the **IP Action** link. The IP Action page appears. 2. Complete the configuration according to the guidelines shown in Table 61 on page 236. 3. Click **OK**. You are returned to the previous page. A gear icon next to the IP Action link indicates that you have configured IP action settings.

Table 59: Add IPS Rule Settings *(Continued)*

Setting	Guideline
[Additional actions]	When attacks are detected, you can configure additional actions that you want CSO to take. To configure additional actions: 1. Click the **Additional** link. The Additional page appears. 2. Complete the configuration according to the guidelines shown in Table 62 on page 238. 3. Click **OK**. You are returned to the previous page. A gear icon next to the Additional link indicates that you have configured additional settings.

Table 60: Notification Settings

Setting	Guideline
Attack Logging	Click the toggle button to enable an attack to be logged when it is detected. By default, attack logging is disabled.
Alert Flag	If you enabled attack logging, click the toggle button to enable an alert flag to be set in the attack log. This field is disabled by default.
Log Packets	Click the toggle button to enable the logging of packets when an attack is detected. When you enable this field, the **Packets Before**, **Packets After**, or **Post Window Timeout** fields appear and you must specify at least one field. By default, packets are not logged when an attack is detected. In response to a rule match, you can capture the packets received before and after the attack for further offline analysis of attacker behavior. You can configure the number of pre-attack and post-attack packets to be captured for this attack, and limit the duration of post-attack packet capture by specifying a timeout value.
Packets Before	Specify the number of packets received before an attack that should be captured for further analysis of the behavior of the attack. Range: 1 through 255.

Table 60: Notification Settings *(Continued)*

Setting	Guideline
Packets After	Specify the number of packets received after an attack that should be captured for further analysis of attacker behavior. Range: 1 through 255.
Post Window Timeout	Specify a time limit (in seconds) for capturing packets received after an attack. No packets are captured after the specified timeout has elapsed. Range: 1 through 1800.

Table 61: IP Action Settings

Setting	Guideline
IP Action	Select the action to be taken on future connections that use the same IP address: **NOTE**: If there is an IP action match with more than one rule, then the most severe IP action of all the matched rules is applied. In decreasing order of severity, the actions are block, close, and notify. • **None**—Do not take any action, which is the default seting. This is similar to if you did not configure the IP action. • **IP Notify**—Don't take any action on future traffic but log the event. • **IP Close**—Close future connections of new sessions that match the IP address by sending RST packets to the client and server. • **IP Block**—Block future connections of any session that matches the IP address.

Table 61: IP Action Settings *(Continued)*

Setting	Guideline
IP Target	Specify how the traffic should be matched for the configured IP actions: • **None**—Do not match any traffic. • **Destination Address**—Match traffic based on the destination IP address of the attack traffic. • **Service**—For TCP and UDP, match traffic based on the source IP address, source port, destination IP address, and destination port of the attack traffic. • **Source Address**—Match traffic based on the source IP address of the attack traffic. • **Source Zone**—Match traffic based on the source zone of the attack traffic. • **Source Zone Address**—Match traffic based on the source zone and source IP address of the attack traffic. • **Zone Service**—Match traffic based on the source zone, destination IP address, destination port, and protocol of the attack traffic.
Refresh Timeout	Click the toggle button to enable the refresh of the IP action timeout (that you specify in the **Timeout Value** field) if future traffic matches the IP actions configured. This setting is disabled by default.
Timeout Value	Configure the time (in seconds) that you want the IP action to remain in effect. For example, if you configure a timeout of 3600 seconds (1 hour) and traffic matches the IP actions configured, the IP action remains in effect for 1 hour. Range: 0 through 64,800 seconds.
Log Taken	Click the toggle button to enable the logging of information about the IP action against the traffic that matches a rule. This setting is disabled by default.
Log Creation	Click the toggle button to enable the generation of an event when the IP action filter is triggered. This setting is disabled by default.

Table 62: Additional Settings

Setting	Guideline
Severity	Select a severity level *__None__, __Critical__, __Info__, __Major__, __Minor__, __Warning__) to override the inherited attack severity in the rules. The most dangerous level is critical, which attempts to crash your server or gain control of your network. Informational is the least dangerous level and is used by network administrators to discover holes in their security systems.
Terminal	Click the toggle button to enable the marking of the IPS rule as terminal. When a terminal rule is matched, the device stops matching for the rest of the rules in that IPS profile. the generation of an event when the IP action filter is triggered. This setting is disabled by default.

RELATED DOCUMENTATION

CSO Next-Generation Firewall (NFGW) Deployment Workflow | 178

Add and Deploy Firewall Policies

NOTE: For SD-WAN deployments, because Juniper's SD-WAN devices are tightly integrated with security, you must configure a firewall policy to allow traffic that traverses zones. By default, traffic between one site and another, and traffic from a site to the Internet is *not allowed* and must be explicitly allowed by using a firewall policy. CSO supports intent-based policies, which makes it simple to configure firewall policies.

To add a firewall policy and then deploy the policy:

1. Add a firewall policy:
   a. Select **Configuration > Firewall > Firewall Policy**.
      The Firewall Policy page appears.
   b. Click the Add (**+**) icon.
      The Add Firewall Policy page appears.

c. Complete the configuration according to the guidelines provided in Table 63 on page 240.

> **NOTE**: Fields marked with an asterisk (*) are mandatory.

d. Click **OK**.

You are returned to the Firewall Policy page. A confirmation message appears when the firewall policy is added.

2. Add one or more firewall policy intents to the policy:

a. Click the **Firewall-Policy-Name** link.

The *Firewall-Policy-Name* page appears.

b. Click the add (**+**) icon.

The fields to add an intent are displayed inline.

c. Complete the configuration according to the guidelines provided in Table 64 on page 240.

d. Click **Save**.

The intent is saved and a confirmation message is displayed. The CSO classifies intents as zone-based and enterprise-based intents. For more information, see *Firewall Policy Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

e. (Optional) Add more intents by following the same procedure.

After intents are added, you must deploy the policy to ensure that the changes take effect on the applicable sites, departments, or applications. When each firewall policy intent is added, the Undeployed field is incremented by one indicating that intents are pending deployment.

3. Deploy the firewall policy:

a. Click the **Deploy** button.

The Deploy page appears.

b. From the **Choose Deployment Time** field, select:

- **Run now** to trigger the deployment of the policy immediately.
- **Schedule at a later time** to schedule the deployment for later.

  If you schedule the deployment for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur. You specify the time in the local time zone of the client from which you access the CSO GUI.

You are returned to the Firewall Policy page and a job to deploy the policy is triggered. You can check the status of the deployment on the Jobs page (**Monitor > Jobs**). When the job completes successfully, it means that the firewall policy was deployed.

Table 63: Add Firewall Policy Settings

Field	Guideline
Name	Enter a unique name for the firewall policy. The name con contain alphanumeric characters, hyphens, and underscores, and cannot exceed 255 characters.
Description	Enter a description for the firewall policy.
All Sites	Click the toggle button to enable the firewall policy to be applied to all sites. By default, a firewall policy is not applied to all sites.
Select Sites	To apply the firewall policy only to specific sites, select the sites from the left column and click the **>** icon. The sites that you selected are displayed in the right column.

Table 64: Add Firewall Policy Intent Settings

Field	Guideline
[Name]	Enter a name for the policy intent or use the one generated by CSO. The name must start with an alphanumeric characters, can contain alphanumeric characters, hyphens, and underscores, and cannot exceed 63 characters.
[Description]	Enter a description for the policy intent.
[Select Schedule]	Policy schedules enable you to define when a policy is active, and thus are an implicit match criterion. Click inside the text box to select a pre-existing schedule or click **Add schedule** to add a new schedule. For more information on adding schedules, see *Creating Schedules* in the *CSO Customer Portal User Guide* (available at the CSO Documentation page).
Logging	Click the toggle button to enable logging. By default, logging is disabled. You can see the logged firewall events in the Firewall Events page (**Monitor > Security Events > Firewall**). For more information, see *About the Firewall Events Page* in the *CSO Customer Portal User Guide*.

Table 64: Add Firewall Policy Intent Settings *(Continued)*

Field	Guideline
Source	Select one or more of the following source endpoints: • IP address or IP address group • User • Site • Site group • Department • Zone If you don't select a source, the default source used is **All Sites**.
Action	Click the add icon (**+**) and select the action to take on the traffic between the specified source and destination endpoints: • **Allow**—Permit traffic between the source and the destination. • **Deny**—Silently drop all packets for the session and do not send any active control messages, such as TCP Reset or ICMP unreachable. • **Reject**—Drop the packets and send a TCP reset (for TCP protocol) or an ICMP reset (for UDP, ICMP, or any other IP protocol) message. This option is useful when you're dealing with trusted resources, so that applications don't have to wait for a timeout but receive an active message.

Table 64: Add Firewall Policy Intent Settings *(Continued)*

Field	Guideline
Destination	Select one or more of the following destination endpoints: • IP address or IP address group • Site • Site group • Department • Service • Application signature or application signature group • Zone If you don't select a destination, the default destination used is **Internet**. **NOTE**: The address endpoint **Any** refers to any address on the Internet and not to any IP address. So, if you want to enable site-to-site traffic, you must explicitly add intents to allow the traffic. For example, if you want traffic from Site A to Site B to be allowed in both directions (A to B and B to A), you must add 2 intents, one allowing traffic from Site A to Site B and another allowing traffic from Site B to Site A.

Table 64: Add Firewall Policy Intent Settings *(Continued)*

Field	Guideline

Advanced Security	**NOTE**: This field is enabled only if you select **Allow** for the action, or if you select a zone as a source and destination.
	
	• When you set the action to **Allow**:
	
	  • You can specify a UTM profile by selecting a profile from the list (under **UTM Profiles [UTM]**).
	
	  You specify a UTM profile for protection against multiple threat types including spam and malware, and control access to unapproved websites and content.
	
	  You can add a new UTM profile by clicking **+** in the End Points pane and selecting **UTM Profiles**. See "Add UTM Profiles" on page 193.
	
	  • You can specify an IPS profile by selecting a profile from the list (under **IPS Profiles [IPS]**).
	
	  You specify an IPS profile to monitor and prevent intrusions.
	
	• When you configure a zone as part of the source and the destination, you can specify an SSL proxy profile by selecting a profile from the list (under **SSL Profiles [SSLP]**).
	
	  You add an SSL proxy profile to ensure the secure transmission of data between a client and a server through a combination of privacy, authentication, confidentiality, and data integrity.
	
	  You can also add a new SSL proxy profile by clicking **+** in the End Points pane and selecting **SSL Proxy Profile**. See "Add SSL Forward Proxy Profiles" on page 219.

See "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Add and Deploy NAT Policies

CSO supports source NAT, destination NAT, and static NAT. In addition, CSO supports persistent NAT depending on the type of source and destination address. In addition, during the addition of an SD-WAN on-premise spoke site and an enterprise hub site, you can trigger the automatic creation of source NAT rules for local breakout traffic. For more information about NAT in CSO, see *NAT Policies Overview* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

To add and deploy a NAT policy:

1. Add the source NAT policy:
   a. Select **Configuration > NAT > NAT Policies**.
      The NAT Policies page appears.
   b. Click the Add (**+**) icon.
      The Add NAT Policy page appears.
   c. Configure the NAT policy according to the guidelines in Table 65 on page 246.

      NOTE: Fields marked with an asterisk (*) are mandatory.

   d. Click **OK**.
      You are returned to the NAT Policies page. After the NAT policy is added, a confirmation message is displayed.

      After you add the NAT policy, you can add one or more rules.
2. You can add three types of NAT rules in CSO: source NAT, static NAT, destination NAT.

- To add a source NAT rule:

  NOTE: If you don't have a separate NAT device in your network and want traffic to break out directly from either an enterprise hub or an on-premise spoke site, you must have a source NAT policy for the hub site or the on-premise spoke site.
  If you enabled the automatic creation of source NAT rules during the addition of the site, CSO automatically creates the source NAT rules.

  a. Click the name of that NAT policy that you added.
     The *NAT-Policy-Name* page appears.
  b. Click **Create > Source**.
     The fields to be configured appear inline on the page.
  c. Configure the source NAT rule according to the guidelines in Table 66 on page 247.

     NOTE: Fields marked with an asterisk (*) are mandatory.

  d. Click **Save**.

A confirmation message appears at the top of the page when the source NAT rule is added successfully. The Undeployed field is incremented by one indicating that intents are pending deployment.

- To add a static NAT rule:

  a. Click the name of that NAT policy that you added.

     The *NAT-Policy-Name* page appears.

  b. Click **Create > Static**.

     The fields to be configured appear inline on the page.

  c. Configure the static NAT rule according to the guidelines in Table 68 on page 252.

     NOTE: Fields marked with an asterisk (*) are mandatory.

  d. Click **Save**.

     A confirmation message appears at the top of the page when the static NAT rule is added successfully. The Undeployed field is incremented by one indicating that intents are pending deployment.

- To add a destination NAT rule:

  a. Click the name of that NAT policy that you added.

     The *NAT-Policy-Name* page appears.

  b. Click **Create > Destination**.

     The fields to be configured appear inline on the page.

  c. Configure the destination NAT rule according to the guidelines in Table 70 on page 253.

     NOTE: Fields marked with an asterisk (*) are mandatory.

  d. Click **Save**.

     A confirmation message appears at the top of the page when the destination NAT rule is added successfully. The Undeployed field is incremented by one indicating that intents are pending deployment.

After adding the NAT rules, you must deploy the rules to the sites with which the NAT policy is associated.

3. Click **Deploy**. (Alternatively, you can trigger the deployment from the NAT Policies page by selecting the policy and clicking **Deploy**).

   The Deploy page appears displaying the name of the policy to be deployed.

4. From the **Choose Deployment Time** field, select:

   - **Run now** to deploy the policy immediately.
   - **Schedule at a later time** to schedule the deployment for later.

     If you schedule the deployment for later, enter the date (in MM/DD/YYYY format) and time (in HH:MM:SS 24-hour or AM/PM format) that you want the deployment to occur. You specify the time in the local time zone of the client from which you access the CSO GUI.

   You are returned to the previous page and a job to deploy the policy is triggered. You can check the status of the deployment on the Jobs page (**Monitor > Jobs)**. When the job completes successfully, it means that the NAT policy was deployed.

Table 65: Add NAT Policy Settings

Field	Guideline
Name	Enter the name of NAT policy. The name can contain alphanumeric characters, colons, periods, hyphens, and underscores. No spaces are allowed and the maximum length is 255 characters.
Description	Enter a description for the NAT policy.

Table 65: Add NAT Policy Settings *(Continued)*

Field	Guideline
Manage Auto-Proxy ARP	Click the toggle button to enable or disable automatic proxy Address Resolution Protocol (ARP). This field is disabled by default. Typically, when an interface receives an ARP request, it responds with its MAC address only then the ARP request corresponds to the interface's IP address. However, when you enable this field, the interface also acts as a proxy and responds to ARP requests for IP addresses other than its own. **NOTE**: Proxy ARP management applies to translated addresses in a source NAT rule or to a destination address in a destination NAT rule: • When you add a source NAT rule with pool-based translation, the address pool assigned must be in the same subnet as the outgoing interface selected. • When you add a destination NAT rule, the external WAN interface can be a proxy for another IP address in the same subnet as the original IP address of the interface.
Sites Applied On	Select the sites on which you want to apply the NAT policy and click the right arrow (**>**).
Sequence No.	Click **Select Policy Sequence** link if you want to reorder this NAT policy among the existing NAT policies. If you deploy more than one NAT policy on a site, the policy sequence number determines the order in which the policies (and therefore the NAT rules) are deployed. The Select Policy Sequence page appears, displaying all NAT policies. Select the policy you want to reorder and click **Move Policy Up** or **Move Policy Down** to reorder your NAT policy among the existing policies.

Table 66: Add Source NAT Rule

Field	Guideline
Name	You can use the default name (that CSO generates automatically) for the NAT rule or enter a unique name.
Description	Enter a description for the NAT rule.

Table 66: Add Source NAT Rule *(Continued)*

Field	Guideline
Source	Specify one or more of the following source endpoints: • Address • Port: To specify a port, type **Port** and press Tab, enter the port number, and press Enter. • Zone • Routing instance • Protocols • Interface • VRF Group **NOTE**: You must specify at least one zone, interface, or VRF group as a source endpoint and specify at least one address for the source or destination endpoints.
Destination	Specify one or more of the following destination endpoints: • Address • Service • Zone • Routing instance • Protocols • Interface • VRF Group **NOTE**: You must specify at least one zone, interface, or VRF group as a destination endpoint and specify at least one address for the source or destination endpoints.

Table 66: Add Source NAT Rule *(Continued)*

Field	Guideline
Translation	Select the type of translation to apply to the traffic: • **None**—Don't perform any translation. • **Interface**—Perform interface-based translation. • **Pool**—Perform pool-based translation. If you select this option, you must specify an address pool by clicking inside the text box adjacent to the list and selecting a NAT pool.
[Advanced Settings]	If you selected interface or pool as the translation type, you can specify additional settings by clicking the gear icon. The Advanced Settings page appears. See Table 67 on page 249for an explanation of the fields.

Table 67: Advance Settings for Source NAT Rule

Field	Description	Translation Type

Persistent	Click the toggle button to enable persistence, which ensures that all requests from the same internal transport address are mapped to the same reflexive transport address.
	NOTE: For persistence to be applicable for the NAT policy, ensure that port overloading is turned off for the device to which the NAT policy is applicable. Use the following command to turn off port overloading for a device:
	`[Edit mode]`
	`set security nat source interface port-overloading off`
		Interface
		Pool

Table 67: Advance Settings for Source NAT Rule *(Continued)*

Field	Description	Translation Type
Persistent NAT Type	Select the type of persistent NAT mapping to use: • **Permit any remote host**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. (The reflexive transport address is the public IP address and port created by the NAT device closest to the STUN server.) Any external host can send a packet to the internal host by sending the packet to the reflexive transport address. • **Permit target host**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address. • **Permit target host port**—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address and port.	Interface Pool
Inactivity Timeout	Enter the period (in seconds) for which the persistent NAT binding remains in the device's memory when all the sessions of the binding entry have ended. When the configured timeout is reached, the binding is removed from memory. Range: 60 through 7,200 seconds. Default: 60 seconds.	Interface Pool
Maximum Session Number	Enter the maximum number of sessions with which a persistent NAT binding can be associated. For example, if the maximum session number of the persistent NAT rule is 2000, then a 2001st session cannot be established if that session uses the persistent NAT binding created from the persistent NAT rule. Range: 8 through 65,536	Interface Pool

Table 67: Advance Settings for Source NAT Rule *(Continued)*

Field	Description	Translation Type
Address Mapping	Allows requests from a specific internal IP address to be mapped to the same reflexive IP address (the public IP address created by the NAT device closest to the STUN server); internal and external ports can be any ports. An external host using any port can send a packet to the internal host by sending the packet to the reflexive IP address (with a configured incoming policy that allows external to internal traffic). If this option is not configured, the persistent NAT binding is for specific internal and reflexive transport addresses.	Pool
Pool Address	Displays the name of the NAT pool that you previously added. You cannot modify this field.	Pool
Host Address Base	Displays the base address of the original source IP address range for the NAT pool that you previously added. The host address base is used for IP address shifting. You cannot modify this field.	Pool
Port Translation	Displays whether port translation is enabled or disabled for the NAT pool that you previously added. You cannot modify this field.	Pool
Overflow Pool Type	Displays the source pool to be used when the address pool is exhausted. You cannot modify this field.	Pool
Overflow Pool Name	Displays the name of the overflow pool. You cannot modify this field.	Pool

Table 68: Add Static NAT Rule

Field	Guideline
Name	You can use the default name (that CSO generates automatically) for the NAT rule or enter a unique name.
Description	Enter a description for the NAT rule.
Source	Specify one or more of the following source endpoints: • Address • Zone • Routing instance • Interface • VRF Group **NOTE**: You must specify at least one zone, interface, or VRF group as a source endpoint.
Destination	Specify one or more of the following destination endpoints: • Address **NOTE**: You must specify at least one address as a destination endpoint. • Port: To specify a port, type **Port** and press Tab, enter the port number, and press Enter.
Translation	Select the type of translation to apply to the traffic: • **Address**—Perform address-based translations on the source or destination packet. If you choose this option, click inside the text box to specify the translation address. • **Corresponding IPv4**—Perform translation using the corresponding IPv4 address.
[Advanced Settings]	You can specify additional settings by clicking the gear icon. The Advanced Settings page appears. See Table 69 on page 253 for an explanation of the fields.

Table 69: Advance Settings for Static NAT Rule

Field	Description	Translation Type
Mapped Port Type	Specify the type of port mapping to use: • **Any**—Allow any port with the translated address. • **Port**—Map to the port specified in the **Port** field. • **Range**—Map to the port range specified in the **Start** and **End** fields.	Address
Routing Instance	Select the routing instance to use for NAT or select **None** not to use a routing instance. **NOTE**: If you're configuring the NAT policy for a site with SD-WAN capability, then you must select the routing instance corresponding to the translation address	Address Overlapping IPv4 Address
Port	Enter the port number to be used for port mapping. Range: 0 through 65,535.	Address
Start	Enter the starting port number of the port range to be used for port mapping. Range: 0 through 65,535.	Address
End	Enter the ending port number of the port range to be used for port mapping. Range: 0 through 65,535.	Address

Table 70: Add Destination NAT Rule

Field	Guideline
Name	You can use the default name (that CSO generates automatically) for the NAT rule or enter a unique name.
Description	Enter a description for the NAT rule.

Table 70: Add Destination NAT Rule *(Continued)*

Field	Guideline
Source	Specify one or more of the following source endpoints: • Address • Zone • Routing instance • Interface • VRF Group **NOTE**: You must specify at least one zone, interface, or VRF group as a source endpoint.

Destination	Specify one or more of the following destination endpoints:
	• Address
	NOTE: You must specify at least one address as a destination endpoint.
	• Port: To specify a port, type **Port** and press Tab, enter the port number, and press Enter.
	• Service
	NOTE: When you add a destination NAT rule for traffic arriving on an interface that terminates a VPN link, the translation process might break the VPN link if the destination addressis specified only as the WAN-facing IP address of the interface.
	For example, in the following NAT rule, any traffic destined to WAN IP address is translated to the destination pool, which breaks the functionality of the VPN link packets terminating on the interface.
	`[Any.Address] --> [Wan.IP] :: [Dest-Pool-1]`
	Therefore, we recommend that you specify both the address and port number as the destination endpoint:
	`[Any.Address] --> [Wan.IP + Port] :: [Dest-Pool-1]`

Table 70: Add Destination NAT Rule *(Continued)*

Field	Guideline
Translation	Select the type of translation to apply to the traffic: • **None**—Don't apply translation. • **Pool**—Perform pool-based translation. If you choose this option, click inside the text box and specify the NAT pool to use. **NOTE**: For sites with SD-WAN capability, the destination NAT pool selected must be configured with a site and a routing instance corresponding to the pool address. For example, if a webserver with IP address IP-Addr-1 is running in the HR department of a site called Site-A. To add a destination NAT pool corresponding to this webserver IP address, you must specify the following mandatory fields while adding the NAT pool: • Address—IP-Addr-1. • Site: Site-A. • Routing Instance: natVR_HR.

See "CSO Next-Generation Firewall (NFGW) Deployment Workflow" on page 178.

Supported Devices for NGFW, and Ports and Protocols to Open

IN THIS SECTION

- What's Next | **257**

Table 71 on page 256 lists the Next-Generation Firewall (NGFW) devices that are supported by CSO and the list of ports or protocols that must be opened for these devices.

NOTE: During the site activation process for SRX4100, SRX4200, and vSRX 3.0, you must copy the stage-1 configuration (generated automatically by CSO) to the device, and commit the configuration on the device.

Before you add a NGFW spoke site:

- Connect cables to the device according to your network design, and power on the device. For more information, see the hardware documentation links in Table 71 on page 256.

  NOTE: We assume that the NGFW device will obtain the DHCP IP address and will have Internet connectivity along with DNS resolution when connected according to the network design.

- Ensure that the ports and protocols listed in Table 71 on page 256 are open on the network.
- Ensure that the devices are running the recommended version of Junos OS. For information about the supported Junos OS versions in a CSO release, refer to the CSO Release Notes for that release (available at the CSO Documentation page).
- If you are using an SRX Series device as the NGFW, ensure that you configure either the first port (ge-0/0/0) or the last port (ge-0/0/7 or ge-0/0/15 based on the SRX model) for Internet connectivity.

Table 71: NGFW Devices Supported

Device Model	Protocols or Ports	Hardware Documentation Links
SRX300	TCP Port 443	SRX300 Chassis
SRX320	TCP Port 514	SRX320 Chassis
SRX340	TCP Port 6514	SRX340 Chassis
SRX345	TCP Port 7804	SRX345 Chassis
SRX380	TCP Port 8060 (only if using you are using PKI authentication to validate the certificate revocation list [CRL])	SRX380 Chassis

Table 71: NGFW Devices Supported *(Continued)*

Device Model	Protocols or Ports	Hardware Documentation Links
SRX550M	TCP Port 443 TCP Port 514 TCP Port 6514 TCP Port 7804 TCP Port 8060 (only if using you are using PKI authentication to validate the certificate revocation list [CRL])	SRX550 HM Chassis
SRX1500	TCP Port 443 TCP Port 514 TCP Port 6514 TCP Port 7804 TCP Port 8060 (only if using you are using PKI authentication to validate the certificate revocation list [CRL])	SRX1500 Chassis
SRX4100 SRX4200	TCP Port 443 TCP Port 514 TCP Port 6514 TCP Port 7804 TCP Port 8060 (only if using you are using PKI authentication to validate the certificate revocation list [CRL])	SRX4100 Chassis SRX4200 Chassis

WHAT'S NEXT

See "CSO Next-Generation Firewall (NFGW) Deployment Workflow" | **178**.

Monitor Next-Generation Firewall Sites and Devices

After configuring a next-generation firewall site, you can perform the following monitoring tasks:

- On the *Site-Name* page (**Resources > Site Management > *Site-Name***), you can view general information about the site, policies, and devices.

  For more information, see *Managing a Single Site* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).

- On the *Device-Name* page (**Resources > Devices > *Device-Name***), you can view general information about the device, view recent alerts and alarms, and view and manage interfaces, routing instances, and zones.

  For more information, see *Managing a Single CPE Device* in the *CSO Customer Portal User Guide*.

- On the Generated Alerts page (**Monitor > Alerts**), you can view the alerts generated by the next-generation firewall devices.

  For more information, see *About the Generated Alerts Page* in the *CSO Customer Portal User Guide*.

- On the Alarms page (**Monitor > Alarms**), you can view the alarms raised by the next-generation firewall devices.

  For more information, see *About the Alarms Page* in the *CSO Customer Portal User Guide*.

- On the Traffic Logs page (**Monitoring > Security Events > Traffic Logs**), you can view the traffic logs generated by next-generation firewall devices.

  For more information, see *About the Traffic Logs Page* in the *CSO Customer Portal User Guide*.

- On the All Security Events page (**Monitor > Security Events > All Events**), you can view a summary and detailed view of the security events in your network.

  For more information, see *About the All Security Events Page* in the *CSO Customer Portal User Guide*.

- On the Firewall Events page (**Monitor > Security Events > Firewall**), you can view a summary and detailed view of the firewall-related security events.

  For more information, see *About the Firewall Events Page* in the *CSO Customer Portal User Guide*.

- On the Web Filtering Events page (**Monitor > Security Events > Web Filtering**), you can view a summary and detailed view of the security events related to Web filtering.

  For more information, see *About the Web Filtering Events Page* in the *CSO Customer Portal User Guide*.

- On the IPsec VPN Events page (**Monitor > Security Events > IPsec VPNs**), you can view a summary and detailed view of the security events related to IPsec VPNs.

  For more information, see *About the IPsec VPNs Events Page* in the *CSO Customer Portal User Guide*.

- On the Content Filtering Events page (**Monitor > Security Events > Content Filtering**), you can view a summary and detailed view of the security events related to content filtering.

  For more information, see *About the Content Filtering Events Page* in the *CSO Customer Portal User Guide*.

- On the Antispam Events page (**Monitor > Security Events > Antispam**), you can view a summary and detailed view of the security events related to spam.

  For more information, see *About the Antispam Events Page* in the *CSO Customer Portal User Guide*.

- On the Antivirus Events page (**Monitor > Security Events > Antivirus**), you can view a summary and detailed view of the security events related to viruses.

  For more information, see *About the Antivirus Events Page* in the *CSO Customer Portal User Guide*.

- On the IPS Events page (**Monitor > Security Events > IPS**), you can view a summary and detailed view of the security events related to IPS.

  For more information, see *About the IPS Events Page* in the *CSO Customer Portal User Guide*.

- On the Screen Events page (**Monitor > Security Events > Screen**), you can view a summary and detailed view screen events that occur as a result of the screen options configured on next-generation firewall devices.

  For more information, see *About the Screen Events Page* in the *CSO Customer Portal User Guide*.

- On the Application Visibility page (**Monitor > Application Visibility**, you can security management information such as the type, bandwidth consumption, and behavior of applications running on your network, which you can use to identify application-level threats to your network.

  For more information, see *About the Application Visibility Page* in the *CSO Customer Portal User Guide*.

- On the User Visibility page (**Monitor > User Visibility**), you can view information about the devices (such as top 50 devices accessing high bandwidth-consuming applications and establishing higher number of sessions) on your network.

  For more information, see *About the User Visibility Page* in the *CSO Customer Portal User Guide*.

- On the Threat Map (Live) page (**Monitor > Threat Map (Live)**), you can visualize incoming and outgoing threats between geographic regions, view blocked and allowed threat events and so on.

For more information, see *About the Threats Map (Live) Page* in the *CSO Customer Portal User Guide*.

- On the Security Report Definitions page (**Reports > Report Definitions**), you can create custom report definitions or use predefined report definitions to generate log, bandwidth, and application and network risk (ANR) reports.

  For more information, see *About the Report Definitions Page* in the *CSO Customer Portal User Guide*.

5

CHAPTER

Appendix

Designing and Publishing Network Services | 262

Use Site Templates to Add SD-WAN and NGFW Spoke Sites | 263

Add Branch or Enterprise Hub Sites Without Provisioning a Service | 264

Understand Breakout in CSO | 273

Network Function Virtualization in the Contrail Service Orchestration Deployments | 274

VNFs Supported by Contrail Service Orchestration | 277

Install Junos OS Software onto an NFX Series Device from a USB Drive | 277

Designing and Publishing Network Services

NOTE: This topic is only relevant for the CSO on-premises version.

The CSO Designer Tools consist of three tools that you use to create VNF templates, packages, and service chains that can be deployed as network services for the CSO solutions. CSO Designer Tools are not available for CSO SaaS subscribers. You access the CSO Designer Tools at the same URL as the CSO Administration Portal, but on port 83. For example, if the IP address of the Administration Portal is 192..0.2.12, then the URL for Designer Tools would be: `https://192..0.2.12:83`. Figure 15 on page 262 shows an overview of the workflow used within the Designer Tools application.

Figure 15: Designer Tools Overview

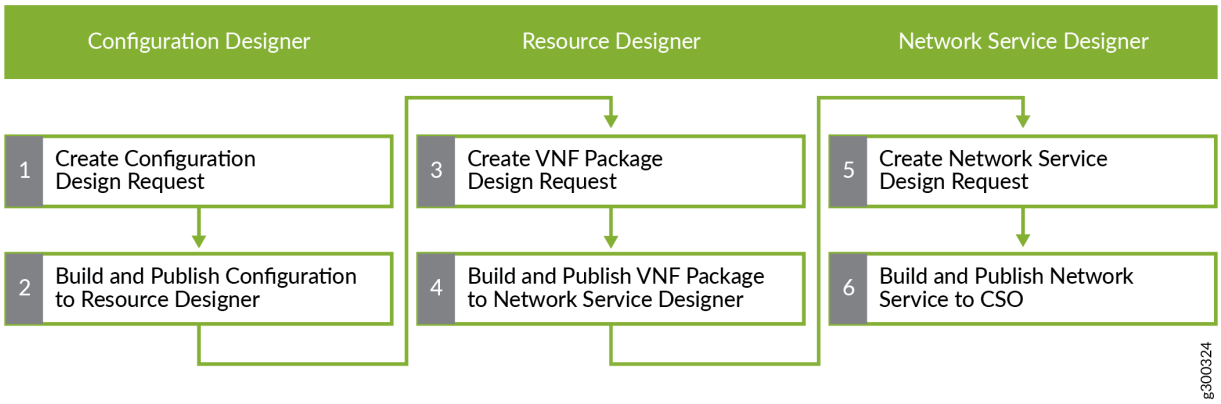

- First, you use the *Configuration Designer* to create configuration templates for virtualized network functions (VNFs). The configuration templates specify the parameters that the customer can configure for a network service.
- Then, you use the *Resource Designer* to create VNF packages. A VNF package is based on a VNF template and specifies the network functions, function chains, and performance of the package.
- Finally, you use the *Network Service Designer* to:
  - Design service chains for network services using the VNF packages that you created with the Resource Designer.
  - Configure the network services.
  - Publish network services to the network service catalog.

You use the same process to create network services for SD-WAN deployments. The same network service can not be shared between an on-premises site and the service provider's POP.

NOTE: Currently, SD-WAN deployments support only Layer 2 service chains.

For more information on Designer Tools, see *Designer Tools Overview* in the *Designer Tools User Guide* (available on the CSO Documentation page).

Use Site Templates to Add SD-WAN and NGFW Spoke Sites

If you need to add multiple branch sites that share some common attributes and some site-specific attributes, then you can use a site template to add the sites. When you add a site template, you can specify the common attributes for the site only once, which means that you only need to specify the site-specific attributes for each site that you add by using site template. In CSO, you can add site templates for branch sites with SD-WAN capability or NGFW capability.

The high-level workflow to add one or more branch sites using a site template is as follows:

1. Add a site template for an SD-WAN branch site (WAN capability SD-WAN) or a NGFW branch site (WAN capability NGFW) using the Add Site Template page (**Resources > Templates > Site Templates > +** or **Resources > Templates > Site Templates > Add Site Template**). For more information, see *Adding a Site Template* in the *CSO Customer Portal User Guide* (available on the CSO Documentation page).
2. Add the sites using the site template from the Add Branch Site page (**Resources > Site Management > Add > Branch Site (Use Site Template)**). For more information, see *Add Branch Sites by Using a Site Template* in the *CSO Customer Portal User Guide*.

TIP: You can enter site-specific attributes manually or use a JavaScript Object Notation (JSON) file to add the site-specific attributes for multiple sites in one go.

RELATED DOCUMENTATION

Add SD-WAN Branch Sites | **133**

Add Next-Generation Firewall (Branch) Sites | **180**

Add Branch or Enterprise Hub Sites Without Provisioning a Service

Starting from CSO Release 6.0.0, you can use the Device Management option to add a branch or enterprise hub site without specifying a service.

After you add the site, the status of the site changes to MANAGED. The site can remain in this state for any duration. You can perform the following tasks when the device is in the MANAGED state:

- Apply stage-2 configuration or configuration templates
- Access the device console
- Reboot the device
- Install licenses and certificates on the device
- Install application signatures
- Initiate RMA

You can deploy either a single or dual SRX CPE without adding a service. CSO Release 6.0.0 supports automatic cluster formation on SRX devices.

NOTE: You cannot add a cloud spoke site with only device management capability. You must select a service for a cloud spoke site.

To configure SD-WAN or security features, you must assign a service to the device. You can edit the site to assign the services. After the service is assigned, the status of the device changes to PROVISIONED.

To add a site with only device management capability:

1. Select **Resources > Site Management**.

   The Site Management page appears.
2. Click **Add** and select **Branch Site (Manual)** or **Enterprise Hub**.

   The Add Branch Site or Add Enterprise Hub page appears.
3. Complete the configuration settings according to the guidelines provided in Table 72 on page 266.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **Next**.

A summary page is displayed.

5. Review the configuration and modify the settings, if needed, from the Summary tab.
6. If you did not enter serial number while creating the site, you must manually enter the serial number after adding the site, in order to activate the site.

   To manually activate the site:

   a. Click **Activate Site** link that appears next to Site Status.

      The **Activate Site** page appears.

   b. Enter the serial number of the device associated with the site.

   c. Click **OK**.

   The **Site Activation Progress** page appears displaying the progress of steps executed for activating the CPE device.
7. If you enabled the **Zero Touch Provisioning** field, CSO pushes the prescript and stage-1 configurations, and the site status changes to **MANAGED** in the Sites page.

   If you disabled the **Zero Touch Provisioning** field for the device, you must copy the stage-1 configuration from CSO and commit it on the device.

   a. Click the **Click to copy stage-1 config** link next to the Prestage Device task in the Site Activation Progress page. If you close the Site Activation Progress page inadvertently, you can access the page from the Site Management page. Click the **View** link next to the status of the site, under the Site Status column.

      NOTE: You can also copy the configuration from the Devices page (Resources > Devices). Select the device and click **Stage1 Config**.

      The Stage-1 Configuration page appears displaying the stage-1 configuration.

   b. Copy the stage-1 configuration.

   c. Log in to the device and enter Junos OS configuration mode.

   d. Paste the configuration that you copied and commit the configuration.

      CSO applies the prescript and stage-1 configuration (includes the device configuration). The status of the site changes to **MANAGED** on the Sites page.

NOTE: You can also add a site using the site templates. For more information, see *Add Branch Sites by Using a Site Template*.

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability)

Field	Description
General	
Site Information	
Site Name	Enter a unique name for the firewall site. You can use alphanumeric characters and hyphen (-); the maximum length allowed is 32 characters.
Device Host Name	The device host name is auto-generated and uses the format *tenant-name.host-name*. You cannot change the *tenant-name* part in the device host name. Use alphanumeric characters and hyphen (-); the maximum length allowed is 32 characters.
Site Group	Select a site group to assign the site.
Site Capabilities	Device Management is selected by default. You need not select the service.
Address and Contact Information	
Street Address	Enter the street address of the site.
City	Enter the name of the city where the site is located.
State/Province	Select the state or province where the site is located.
ZIP/Postal Code	Enter the postal code for the site.

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) ***(Continued)***

Field	Description
Country	Select the country where the site is located. Click the **Validate** button to verify the address that you specified. • The **Address verification successful** message is displayed if the address is valid. You can click the **View location on the map** link to see the address location. • If the address is invalid, the **Site address could not be validated** message is displayed.
Contact Name	Enter the name of the contact person for the site.
Email	Enter the e-mail address of the contact person for the site.
Phone	Enter the phone number of the contact person for the site.
Advanced Configuration	
Domain Name Server (DNS)	Enter one or more IPv4 addresses of the DNS server. To enter more than one DNS server address, type the address, press Enter, and then type the next address, and so on. DNS servers are used to resolve hostnames into IP addresses.
NTP Server	Enter the fully qualified domain names (FQDNs) or IP addresses of one or more NTP servers. Example: ntp.example.net. The site must have DNS reachability to resolve the FQDN during site configuration.
Select Timezone	Select the time zone for the site.
Device **NOTE**: Some fields in this section are displayed only if you enable the Device Redundancy option.	

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) *(Continued)*

Field	Description
Device Redundancy	Disabled by default. Enable this option for dual CPEs. The following prerequisites are necessary for enabling device redundancy: • Ensure that the control and fabric ports between both the nodes are connected. • Ensure that the device is preconfigured for management connectivity (factory-default or prestaged). Do not configure the control, fabric, and data (reth) ports as these ports will be reconfigured. To identify the control, fabric, management, and data ports for each SRX model, refer to the SRX High Availability Configurator tool. **NOTE**: Do not generate the configuration in the tool as CSO generates and applies the cluster configuration automatically. • If you are using ZTP on SRX300 and SRX320 devices, use ge-0/0/7 as the predefined DHCP port instead of ge-0/0/0. • Provide the fabric and data (reth) port information in the device template. The control and fxp0 ports are predefined. To change the control port, change it in the platform device template. To change the data (reth) port, change it in the SDWAN device template.
Device Series	Select the device series. Based on the device series that you select, the supported device templates (containing information for configuring devices) are listed. Select a device template for the selected device series.
Device Model	Select the device model.
Device Root Password	The default root password is fetched from the ENC_ROOT_PASSWORD field in the device template. You can retain the password or change it by entering a password in plain-text format. The password is encrypted and stored on the device.

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) ***(Continued)***

Field	Description
Serial Number	Enter the serial number of the device. Note that the serial numbers are case-sensitive. If you do not enter the serial number, the branch site is created but not activated. See 6 to enter the serial number and activate the branch site later.
Node 0 Serial Number	For dual CPEs, enter the serial number of the primary CPE device. The serial number is case sensitive.
Node 1 Serial Number	For dual CPEs, enter the serial number of the secondary CPE device. The serial number is case sensitive.

Zero Touch Provisioning	Click the toggle button to enable or disable Zero Touch Provisioning (ZTP). This option is enabled by default.
	
	NOTE: By default, this button is disabled for vSRX. You can enable this button, if the Junos OS version running on vSRX supports phone-home client.
	
	To use ZTP, ensure the following:
	
	• Device must have connectivity to CSO and Juniper phone-home server (https://redirect.juniper.net)
	
	Use telnet to verify connectivity:
	
	`telnet redirect.juniper.net:443`
	
	`telnet CSO Hostname/IP:443`
	
	If the connection is established, the device has connectivity to the phone-home server and CSO.
	
	• Required certificates for phone-home server and CSO are present on the device.
	
	If ZTP is enabled, the Boot Image field is displayed and you must select an image that supports the Phone-Home client. During ZTP, the image on the device is upgraded to the image that you select for the Boot Image.

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) *(Continued)*

Field	Description
	If you disable ZTP, ensure that the device has connectivity to CSO. If the device is not prestaged or preconfigured, then you must provide the details under the Management Connectivity section so that CSO can generate the configuration as part of the stage-1 configuration. You can skip the Management Connectivity section if the device has connectivity to CSO. If you disable ZTP, you must copy the stage-1 configuration from CSO and commit it on the device to start the onboarding process. Use any of the following options to copy the stage-1 configuration: • Click the **Click to copy stage-1 config** link next to the Prestage Device task on the Site Activation Progress page. If you close the Site Activation Progress page inadvertently, you can access the page from the Site Management page. Click the **View** link next to the status of the site under the Site Status column. • On the Devices page (**Resources > Devices**), select the device and click **Stage1 Config**.
Is Cluster Already Formed?	Select No if the cluster is not configured.
Cluster ID	Enter the device Cluster ID. The value is ignored if the cluster is already formed on the device. Cluster ID should be unique in case more than one cluster is connected through the same Ethernet switch.
Auto Activate	Click the toggle button to enable or disable automatic activation of the device. This option is enabled by default.
Activation Code	If the automatic activation of the device is disabled, enter the activation code to manually activate the device. The activation code is provided by the administrator who adds the site.
Node 0 Activation Code	If the automatic activation of dual CPEs is disabled, enter the activation code to manually activate the primary CPE device.

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) ***(Continued)***

Field	Description
Node 1 Activation Code	If the automatic activation of dual CPEs is disabled, enter the activation code to manually activate the secondary CPE device.
Management Interface Family	Select the IP address type (IPv4 or IPv6) for the management interface. This field is displayed only if you have enabled **Zero Touch Provisioning**.
Boot Image	When the **Zero Touch Provisioning** field is enabled, select the boot image from the drop-down list to upgrade the image on the firewall device to a version that supports Phone-Home client. The boot image is the device image that was previously uploaded to the image management system. The boot image is used to upgrade the device when the CSO starts the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure. The boot image is populated based on the device template that you selected while creating a site. By default, the **Use Image on Device** option is selected.
(Device Template)	Select a device template, which contains information for configuring a device.
Management Connectivity **NOTE**: This section is displayed only when Zero Touch Provisioning is disabled. If you are adding a chassis cluster, then you must provide the interface details for both the nodes.	
Address Family	Select the IP address type (IPv4 or IPv6).
Interface Name	This is a WAN interface that the device uses to connect to CSO.
Access Type	Select the access type for the underlay link. LTE, ADSL, and VDSL access types are supported only on Internet links. You cannot add LTE, ADSL, and VDSL access types to the same WAN link.

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) ***(Continued)***

Field	Description
Address assignment	DHCP is selected by default. If you want to provide a static IP address, select STATIC.
Management VLAN ID	Enter a VLAN ID for the WAN link. Range: 0 through 4094
PPPoE	Click the toggle button to enable authenticated address assignment for the WAN link by using PPPoE (Point-to-Point Protocol over Ethernet).
ADSL/VDSL SFP Annex	Applicable only to MPLS or Internet links with ADSL or VDSL access types. Enable Annex J support with *x*DSL SFP module for ADSL and VDSL access types. Annex J is specified in ITU-T recommendations G.992.3 and G.992.5. If you keep this option disabled, you must use a Mini-PIM module for connectivity.
Configuration Templates (Optional)	

Table 72: Fields on the Add Branch Site or Add Enterprise Hub Page (Only Device Management Capability) ***(Continued)***

Field	Description
Configuration Templates List	(Optional) Select one or more configuration templates from the list. This list is filtered based on the device that you select. Configuration templates are stage-2 templates that are added by your OpCo administrators, or SP administrators, or Tenant administrators. To set the parameters for the selected configuration templates: 1. After you select one or more configuration templates, click **Set Parameters**. The **Device Configurations** page appears. This page consists of two tabs—CONFIGURATION and SUMMARY. 2. In the CONFIGURATION tab, enter the attributes for each of the configuration templates. (Optional) View the CLI commands in the Summary tab. 3. Click **Save**. You have added and set the parameters for the configuration templates that are part of the site template that you are creating.

Understand Breakout in CSO

Breakout is an SD-WAN feature that enables Internet links to break out traffic directly from a site. For example, if you want to provide guests who visit your enterprise with Internet access, you can use local breakout to break out guest traffic locally from the site directly to the Internet.

In CSO, site-to-site traffic between spoke sites of a tenant is sent (on overlay tunnels) directly from one site to another depending on the tenant topology or through the provider hub or enterprise hub associated with the spoke sites. However, for Internet-bound or Software as a Service (SaaS) traffic, you can break out the traffic in different ways:

- Local breakout—The traffic exits the VPN directly at the site and goes to the destination.

- Backhaul or central breakout—The traffic exits the VPN at the provider hub or at the enterprise hub (based on the hub associated with the spoke site) and then goes to the destination.
- Cloud breakout—The traffic is sent from the site to a designated cloud-based security platform instead of traffic being sent over an underlay.

> **NOTE**: Currently, Zscaler is the only cloud-based security platform supported.

In CSO, to configure breakout on an on-premise spoke site, cloud spoke site, or enterprise hub site, you must do the following:

1. Enable local breakout on one or more Internet WAN links of a site.
2. Add a breakout profile.
3. For cloud breakout, you must add settings for cloud breakout and apply the settings on the site.
4. Add an SD-WAN policy intent that references the breakout profile.
5. Deploy the SD-WAN policy.

To learn more about breakout and breakout profiles in CSO, see *Breakout and Breakout Profiles Overview* in the *CSO Customer Portal User Guide* (available at the CSO Documentation page).

RELATED DOCUMENTATION

Add SD-WAN Breakout Profiles | 127

Add Cloud Breakout Settings | 129

Network Function Virtualization in the Contrail Service Orchestration Deployments

Network Function Virtualization (NFV) is a concept in which network functions traditionally performed by dedicated hardware devices are performed by software that runs on virtual machines in various network locations. The virtual machines run software that performs traditional functions like routing, firewall, or network address translation (NAT). These functions are known as virtual network functions (VNFs). In Figure 16 on page 275, the upper part of the diagram shows conventional physical network

devices chained together to provide network services. The lower part of the diagram shows how the same service chain can be created from a pool of VNFs available on an NFX Series device.

Figure 16: Network Function Virtualization

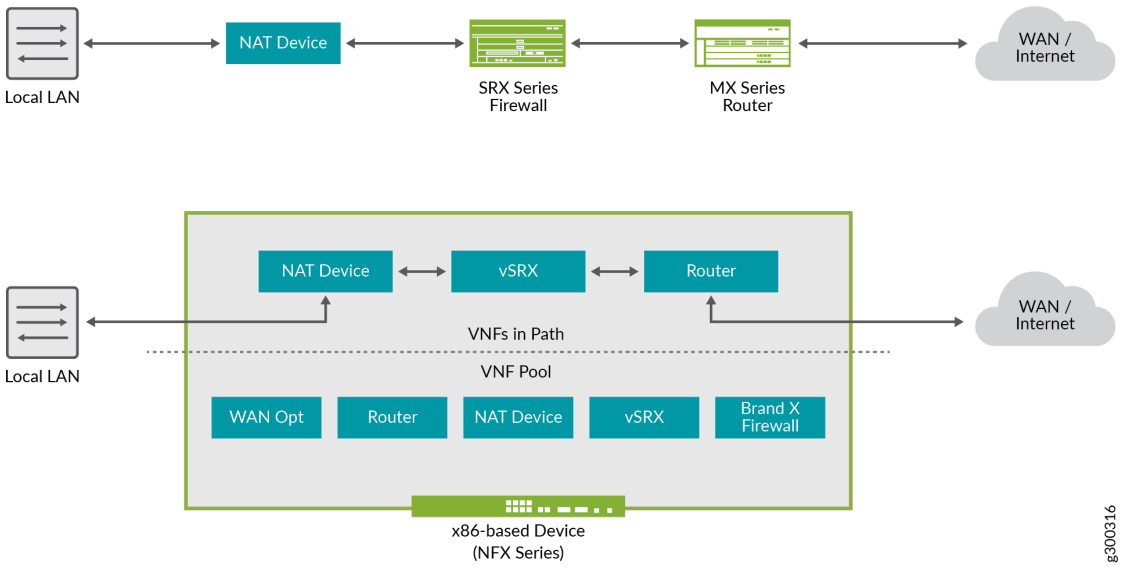

Juniper's CSO solutions comply with European Telecommunications Standards Institute (ETSI) standards for lifecycle management of network service instances.

The Contrail SD-WAN Solution uses the following components for the Network Functions Virtualization (NFV) environment:

- For SD-WAN deployments:
  - The Network Service Orchestrator, together with the Network Service Controller, provides ETSI-compliant management of the life cycle of network service instances.
  - The Network Service Controller provides service-chaining and the VIM.
  - The CPE device provides the NFV infrastructure (NFVI).

Other CSO components connect to the Network Service Orchestrator through its REST API:

- Administration Portal, which you use to set up and manage your virtual network and customers through a graphical user interface (GUI).

  Administration Portal offers role-based access control for administrators and operators. Administrators can make changes; however, operators can only view the portal.

- Customer Portal, a GUI that your customers use to manage sites, CPE devices, and network services for their organizations.

  Customer Portal offers role-based access control for administrators and operators. Administrators can make changes; however, operators can only view the portal.

- Designer Tools:

  - Configuration Designer, which you use to create configuration templates for virtualized network functions (VNFs). When you publish a configuration template, it is available for use in Resource Designer.

  - Resource Designer, which you use to create VNF packages. A VNF package consists of a configuration template and specifications for resources. You use configuration templates that you create with Configuration Designer to design VNF packages. When you publish a VNF package, it is available for use in Network Service Designer.

  - Network Service Designer, which you use to create a network service package. The package offers a specified performance and provides one or more specific network functions, such as a firewall or NAT, through one or more specific VNFs.

> **NOTE**: In Administration and Customer Portals, you can add configuration templates only for Juniper Networks devices. In Configuration Designer, you can create configuration templates for both Juniper and non-Juniper devices.

CSO solutions extend the NFV model through the support of physical network elements (PNEs). A *PNE* is a networking device in the deployment that you can configure through CSO, but not use in a service chain. Configuration of the PNE through CSO as opposed to other software, such as Contrail or Junos OS, simplifies provisioning of the physical device through automation. Combining provisioning and configuration for PNEs and VNFs provides end-to-end automation in network configuration workflows. An example of a PNE is a vSRX device serving as a provider hub for the termination of IPsec and GRE data tunnels.

OSS/BSS applications and CSO components with OSS/BSS capabilities send requests to Network Service Orchestrator through its northbound REST API. Network Service Orchestrator then communicates through its southbound API to the northbound API of the appropriate, directly connected, component. Subsequently, each component in the deployment communicates through its southbound API to the northbound API of the next component in the hierarchy. Components send responses in the reverse direction.

VNFs Supported by Contrail Service Orchestration

Contrail Service Orchestration (CSO) supports Juniper Networks and third-party virtualized network functions (VNFs) listed in the *VNFs Supported* section of the CSO Release Notes, available on the Contrail Service Orchestration (CSO) Documentation page.

An on-premises version of CSO is not shipped with any VNFs. Immediately after installation, you have to upload any desired VNFs to the CSO using the Administration Portal.

You can use VNFs in service chains and configure some settings for them in Network Service Designer. You can then view those network service configuration settings in the Administration Portal. Customers can also configure some settings for the VNFs in their network services through Customer Portal. VNF configuration settings that customers specify in the Customer Portal override VNF configuration settings specified in Network Service Designer, which is not available for CSO SaaS subscribers.

NOTE: Currently, SD-WAN deployments support only Layer 2 service chains.

In CSO SaaS, CSO contains only the VNFs installed by Juniper Networks administrators. Requests for additional VNFs must be made through your account manager and Professional Services.

Install Junos OS Software onto an NFX Series Device from a USB Drive

This section details how to install Junos OS software onto an NFX Series device from a USB drive. Doing this sets the device to the factory-default state. We also perform some confirmation steps and obtain the device's serial number.

Before You Begin

In order for this procedure to succeed, be sure that you have the following:

- Physical access to the USB port of the NFX Series device
- A USB drive of at least 4GB containing the Junos OS software image inserted into the USB port of the NFX Series device
- Access to the console port of the NFX Series device (This can be physical access or access over a terminal server.)

- A DHCP server that is reachable from the `ge-0/0/11` interface of the NFX Series device. This DHCP server must be able to provide IP address, name server, and default gateway to the NFX Series device upon request.

To install Junos OS software onto an NFX Series device by using a USB drive:

1. Ensure that the USB drive containing the Junos OS software image is inserted in the USB port of the NFX Series device.

   This allows you to boot the NFX Series device from the USB drive.
2. Access the NFX Series device console either directly or using a terminal server.

   You do not need to login; just ensure that you are actively connected.
3. Power off the NFX Series device.
4. Power on the NFX Series device.
5. Immediately return to the session that you have open to the console port of the nfx1 device.

   From the console of the nfx1 device, press the ESC key every second until the following message appears: `Esc is pressed. Go to boot options.`

   > **NOTE**: If you do not see this message in the console and the NFX appears to be booting normally, you need to wait for the boot to complete and then go back to step 1.
6. A menu appears after a brief time. Use the down arrow key to select **Boot Manager**, then press **Enter**.
7. When the **Boot Manager** menu appears, press **Enter** to boot from the **USB00** drive.
8. When the **GNU GRUB** menu appears, use the up or down arrow keys to select **Install Juniper Linux with secure boot support** and then press **Enter**.

At this point, the NFX Series device installs the software contained on the USB drive. Installation takes some time. You can keep your console connection active to watch the installation process.

The NFX Series device is made up of multiple components that load and boot in a specific order. See NFX250 Overview for details. The PFE of the NFX Series device may take a few minutes to complete the boot and allow the `jsxe0` interface to obtain its address from DHCP.

Log in to the console of the NFX Series device as the `root` user and confirm that the `jsxe0` interface has received its address using the following procedure:

1. Press **Enter** to refresh the login prompt.
2. At the `jdm login` prompt, type **root** and press **Enter**.

> **NOTE**: There is no password assigned to the root user at this point. For the purposes of this deployment exercise, do not set a root password at this time.

3. At the `root@jdm:~#` prompt, type **cli** and press **Enter**.

4. Type **show interfaces jsxe0** and press **Enter**.

The `jsxe0` interface has a number of logical interfaces used internally by the NFX Series device for different purposes. Look for the `jsxe0.0` logical interface. Confirm that the DHCP server has provided an address in the proper range before continuing.

```
root@jdm:~# show interfaces jsxe0
Logical interface jsxe0.1 (Index 4)
        Flags: Up
        Input packets : 0
        Output packets: 252
        Protocol inet, MTU: 1500

Logical interface jsxe0.2 (Index 5)
        Flags: Up
        Input packets : 3
        Output packets: 274
        Protocol inet, MTU: 1500

Logical interface jsxe0.0 (Index 3)
        Flags: Up
        Input packets : 7097
        Output packets: 8722
        Protocol inet, MTU: 1500
            Destination: 172.26.133.0/24, Local: 172.26.133.106,
            Broadcast: 172.26.133.255
```

At this point, you can confirm that the DNS name server and default gateway are working by issuing the **ping** command to some host on the Internet.

```
root@jdm:~ # cli
root@jdm:~ > ping www.juniper.net count 1
PING e1824.dscb.akamaiedge.net (23.223.165.73) 56(84) bytes of data.
64 bytes from a23-223-165-73.deploy.static.akamaitechnologies.com (23.223.165.73): icmp_seq=1
ttl=56 time=2.67 ms
```

```
--- e1824.dscb.akamaiedge.net ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 2.670/2.670/2.670/0.000 ms
```

The last part of this procedure is to login to the Junos Control Plane (jcp) to obtain the device serial number which will be used later in the SD-WAN deployment.

```
root@jdm:~ > ssh vjunos0
Last login: Tue Jan 22 06:28:51 2019
--- JUNOS 15.1X53-D40.3 Kernel 32-bit FLEX JNPR-10.1-20160217.114153_fbsd-builder_stable_10
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@:~ #cli
root> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DXXXXXXXXXX3
Pseudo CB 0
Routing Engine 0          BUILTIN      BUILTIN           RE-NFX250-S2
FPC 0            REV 04   650-066113   DXXXXXXXXXX3
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0          REV 04   BUILTIN      BUILTIN           10x10/100/1000 Base-T-2x1G SFP-
Power Supply 0
Fan Tray 0                                               fan-ctrl-0 0, Front to Back Airflow -
AFO
Fan Tray 1                                               fan-ctrl-0 1, Front to Back Airflow -
AFO
```

The device serial number is listed on the `Chassis` line of the output. In this example, it is partly obscured for security purposes. Make note of the serial number for later use.