JUNIPER NETWORKS | Engineering Simplicity

Contrail Service Orchestration

Contrail Service Orchestration (CSO) Deployment Guide

Published
2020-11-05

Release
5.1.1

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 Contrail Service Orchestration (CSO) Deployment Guide
5.1.1
Copyright © 2020 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 the Documentation | vii

Documentation and Release Notes | **vii**

Documentation Conventions | **vii**

Documentation Feedback | **x**

Requesting Technical Support | **x**

- Self-Help Online Tools and Resources | **xi**
- Creating a Service Request with JTAC | **xi**

1 **Solutions Overview**

About this Deployment Guide | 13

Contrail Service Orchestration (CSO) Solutions Overview | 13

- Contrail SD-WAN Solution | **15**
- Contrail Managed LAN Solution (SD-LAN) | **16**
- Next Generation Firewall (NGFW) Deployment Model | **17**
- Hybrid WAN (Distributed CPE) Deployment Model | **18**

Building Blocks Used for Contrail Service Orchestration Deployments | 20

- Administrators | **20**
- Portals | **21**
- Tenants | **22**
- Topologies | **22**
- Points of Presence (POPs) | **25**
- Sites | **26**
- Customer Premises Equipment (CPE) | **30**
- Standalone Next-Generation Firewall (NGFW) | **31**
- SD-LAN Devices | **31**
- Virtual Route Reflector (VRR) | **31**
- SLA-Based Steering Profiles and Policies | **32**
- Path Based Steering Profiles | **33**
- Intent-based Firewall Policies | **33**
- Software Image Management | **33**

2 Deployment Tools

Contrail Service Orchestration (CSO) Deployment Tools | 35

Contrail Services Orchestration (CSO) GUIs | 35

Designing and Publishing Network Services | 37

Contrail Service Orchestration (CSO) License Tool | 38

- Overview of the License Pages | 38

3 SD-WAN Deployment

SD-WAN Deployment Overview | 42

Contrail SD-WAN Deployment Architectures | 42

- Contrail SD-WAN Reference Architecture | 43
- Spoke Devices | 44
  - On-Premises Spoke Devices | 44
  - Cloud Spoke Devices | 46
  - Spoke Redundancy | 46
- Provider Hub Devices | 47
  - Provider Hubs | 47
  - Provider Hub Redundancy | 48
- Enterprise Hub Sites and Devices | 48
- Underlay (Physical) Network | 49
  - WAN Access Options | 50
  - WAN Interface Types - Data and OAM | 51
- Overlay (Tunnels) Network | 52
  - Overlay Deployment Topologies | 53
- Orchestration and Control | 55
- Secure OAM Network | 56
  - Integration with Deployment Topologies | 57
  - OAM Hub Design Options | 58
  - Usage Notes on Provider Hub Design Options | 59
- Zero Touch Provisioning | 60
  - Usage Notes for ZTP | 60

Redirect Server | **61**

Design Considerations for CSO and Redirect Server | **61**

Bypassing the Redirect Server | **62**

Service Chaining in Contrail SD-WAN | **62**

Three Planes, Four Layers | **63**

Initial SD-WAN Deployment | 64

Before You Begin | **65**

Download Intrusion Protection System (IPS) and Application Signatures | **65**

Upload Licenses | **68**

Add a New Tenant | **68**

Modify Device Templates | **70**

Choose a Point of Presence (POP) for the Hub | **74**

Add a Provider Hub Device to Your Tenant | **74**

Add an Enterprise Hub to Your Tenant | **78**

Add an On-Premises Spoke for the Tenant | **82**

Install a License on a Device | **88**

Install an Application Signature on a Device | **88**

Add Firewall and NAT Policies to the Topology | **89**

Add SD-WAN SLA-Based Steering Profiles and Policy | **91**

4 **SD-LAN Deployment**

SD-LAN Deployment | 94

SD-LAN Deployment Overview | **94**

SD-LAN Deployment | **98**

5 **Standalone Next-Generation Firewall (NGFW) Deployment**

Next-Generation Firewall (NGFW) Deployment | 104

NGFW Deployment Overview | **104**

NGFW Deployment Architecture | **105**

NGFW Deployment | **106**

6 Hybrid WAN Deployment (uCPE)

Hybrid WAN Deployment Overview | 111

Hybrid WAN Deployment Architecture | 112

Hybrid WAN Deployment | 114

- Modify Device Templates | 114
- Add and Configure a New Tenant | 115
- Add and Configure a Site for the Tenant | 116

7 Appendix A - Network Function Virtualization in Contrail Service Orchestration

Network Function Virtualization in the Contrail Service Orchestration Deployments | 119

VNFs Supported by the Contrail Service Orchestration Solutions | 121

8 Appendix B - Manual Staging of NFX

Install Junos OS Software onto an NFX Series Device from a USB Drive | 124

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | vii
- Documentation Conventions | vii
- Documentation Feedback | x
- Requesting Technical Support | x

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

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at https://www.juniper.net/documentation/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at https://www.juniper.net/books.

Documentation Conventions

Table 1 on page viii defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page viii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		

Table 2: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback so that we can improve our documentation. You can use either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

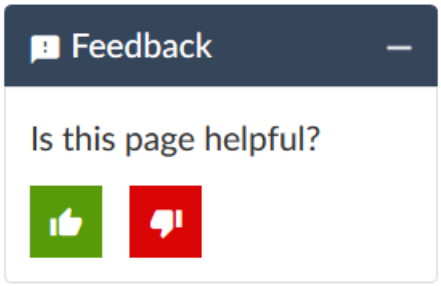

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active Juniper Care or Partner Support Services support contract, or are

covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://myjuniper.juniper.net

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see https://support.juniper.net/support/requesting-support/.

CHAPTER 1

Solutions Overview

About this Deployment Guide | 13

Contrail Service Orchestration (CSO) Solutions Overview | 13

Building Blocks Used for Contrail Service Orchestration Deployments | 20

About this Deployment Guide

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

- Briefly discussing each of the available solutions
- Discussing the building blocks used in every deployment
- Discussing the tools used to put the blocks together
- Providing an end-to-end walkthrough of each of the solutions and covering the deployment specifics

This guide is available on the Contrail Service Orchestration Documentation page, which includes several other helpful guides:

- *Quick Start Guide for Contrail Service Orchestration, Release 5.1.1*
- *Contrail Service Orchestration Installation and Upgrade Guide*
- *Contrail Service Orchestration Administration Portal User Guide*
- *Contrail Service Orchestration Customer Portal User Guide*
- *Contrail Service Orchestration Monitoring and Troubleshooting Guide*
- Contrail SD-WAN and SD-LAN Design and Architecture Guide
- And more!

Contrail Service Orchestration (CSO) Solutions Overview

IN THIS SECTION

- Contrail SD-WAN Solution | **15**
- Contrail Managed LAN Solution (SD-LAN) | **16**
- Next Generation Firewall (NGFW) Deployment Model | **17**
- Hybrid WAN (Distributed CPE) Deployment Model | **18**

Juniper Networks Contrail SD-WAN, SD-LAN, 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.

The CSO platform itself can be deployed in one of two ways:

- As a downloadable, on-premises platform in which you (or your company) become the SP administrator (cspadmin user). In an on-premises deployment, the cspadmin user has complete read-write management access and responsibility for the CSO micro-services platforms, orchestration and management infrastructure, and all underlay networks needed to allow access to CSO and its solutions.
- As a software-as-a-service (SaaS) platform, hosted in a public cloud, to which tenants and OpCos subscribe. In an SaaS deployment, Juniper Networks manages the necessary microservices infrastructure, the secure orchestration and management (OAM) infrastructure, and underlay networks needed to allow access to CSO and its solutions.

CSO offers multiple network solutions that benefit enterprise customers and service providers and their customers. The solutions are split into two overall groups, WAN solutions and LAN solutions as shown in Figure 1 on page 14.

Figure 1: WAN and LAN Solutions

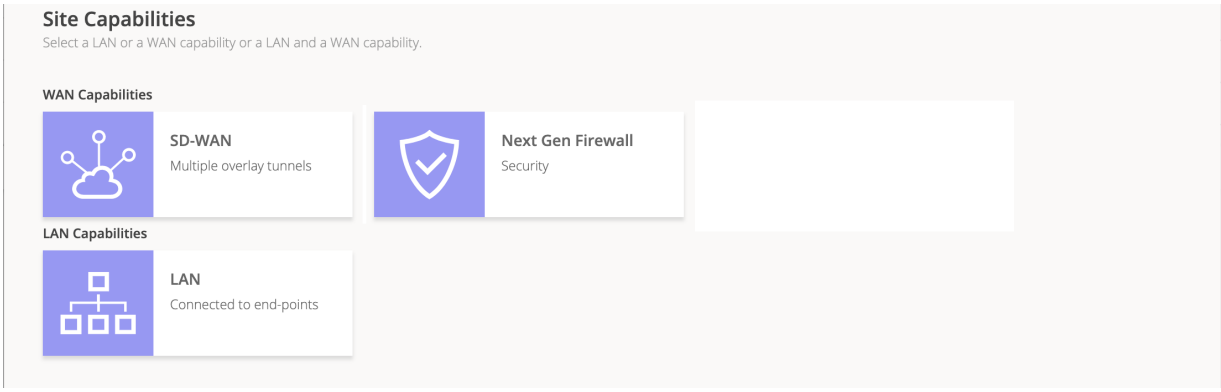

These solutions allow CSO to do the following:

- 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 SD-WAN Solution

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 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 2 on page 15.

Figure 2: Basic SD-WAN Concept

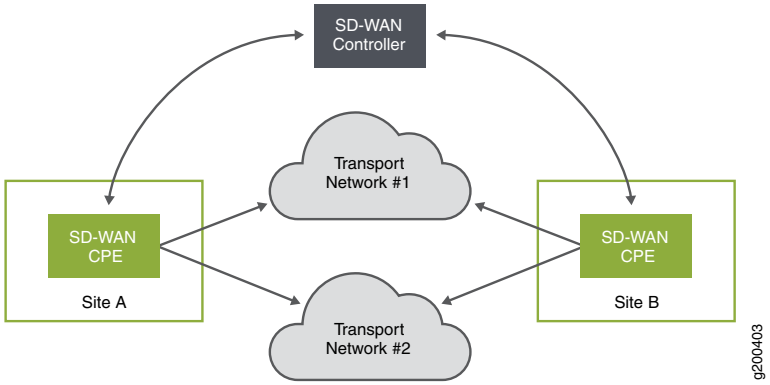

The CPE devices in a Contrail SD-WAN solution (also known as *on-premises spoke devices*) have a WAN side and a LAN side. On the WAN side, hub-and-spoke and dynamic mesh topologies are supported. The CPE devices use at least one, and up to four, WAN interfaces as connection paths to provider hub devices, enterprise hub devices, other spoke devices, and the Internet. The supported hub devices are shown in Table 3 on page 15:

Table 3: Supported Hub Devices

Hub Device	Used as
vSRX	Enterprise Hub and Provider Hub
SRX1500	Enterprise Hub and Provider Hub
SRX4100	Enterprise Hub and Provider Hub
SRX4200	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 with the use of dedicated IPSec tunnels. 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.

Contrail Managed LAN Solution (SD-LAN)

The SD-LAN solution allows CSO to manage and monitor remote LAN devices like certain EX Series LAN switches, Mist WiFi access points, and certain SRX Series next generation firewall (NGFW) devices. This extends the SD-WAN solution to provide visibility into the LANs of remote networks. At its most basic, a managed LAN implementation is as simple as connecting a supported EX switch or SRX firewall at the remote site through an Internet gateway device as shown in Figure 3 on page 17.

Figure 3: Simple SD-LAN Solution

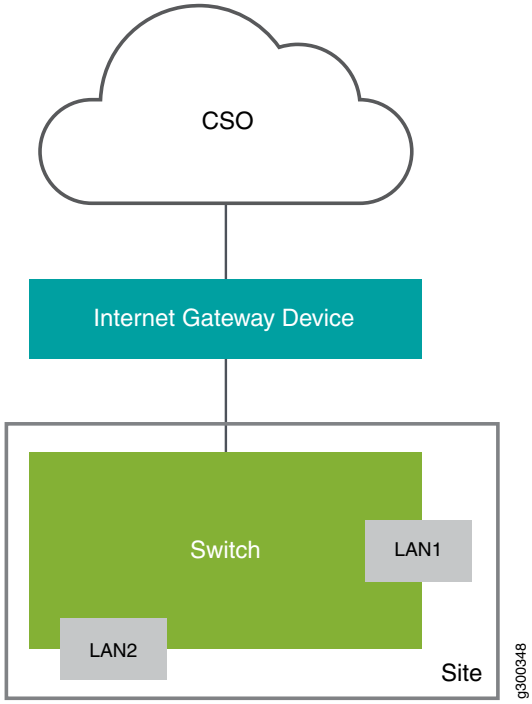

While Figure 3 on page 17 shows a single switch connected behind an Internet gateway device, there are several other deployment options available within the solution. For example, an EX switch can be attached to an existing managed CPE device, or it can be added to CSO as a standalone LAN switch. Similar deployment options are available for the NGFW solution. For more details about switch deployment in a managed LAN solution, see the *Adding and Provisioning Switches to Provide LAN Capability to a Site Overview* and the CSO Design and Architecture Guide.

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 and Hybrid WAN deployments which focus on secure site-to-site connectivity and remote VNF deployment. A high-level view of the spoke site with NGFW is shown in Figure 4 on page 18.

Figure 4: NGFW Spoke Site

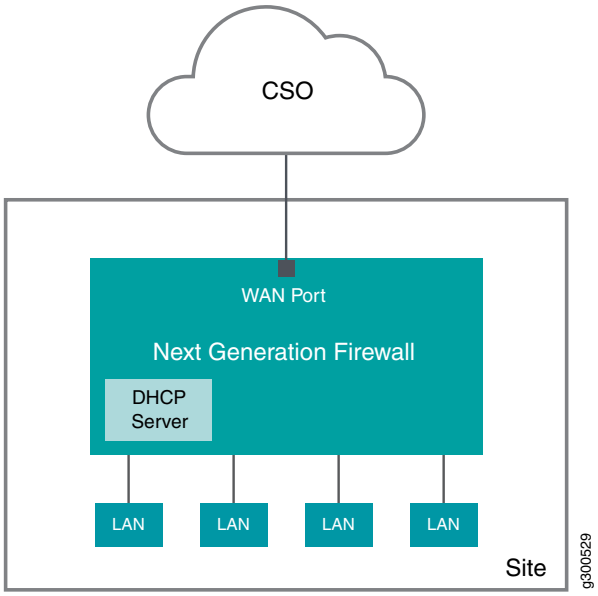

An NGFW deployment is carried out in the Customer Portal of CSO as a site deployment. The tenant under which the site is deployed must have the NGFW service available. This service is included in the tenant configuration by the tenant administrator during tenant onboarding. The remainder of this document provides a brief discussion of the architecture, and the steps that you need to perform in order to complete a NGFW deployment in CSO.

Hybrid WAN (Distributed CPE) Deployment Model

In a Hybrid WAN deployment, customers access network services from a CPE device located at the customer's site. These sites are called *on-premises sites* or *spokes*. In the workflows used in the CSO GUI, this deployment style is known as Hybrid WAN. Figure 5 on page 19 illustrates a simplified Hybrid WAN deployment.

Figure 5: Hybrid WAN Deployment

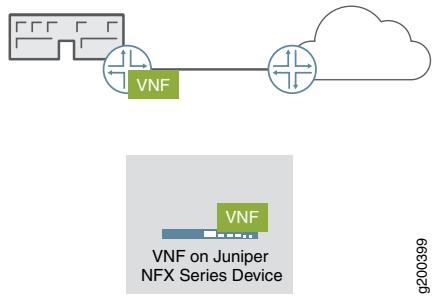

Initial configuration of the CPE device at the site can be automated through the use of zero touch provisioning (ZTP) that is orchestrated through CSO. CSO also monitors the CPE device and its services, and can push software and configuration updates to the devices remotely, reducing operating expenses. This deployment model is useful in environments where service delivery from the service provider's cloud is costly.

In fact, CSO has been designed to require only modest bandwidth, needing as little as 30 kbps for probe and secure OAM traffic over Hybrid WAN connections where there are only a few sessions active. When AppQoE is involved, the bandwidth requirement increases to somewhere between 105 kbps and 2 Mbps, depending on the number of sessions. During ZTP operations, if new device images are needed, they can be downloaded as part of the ZTP process, or pre-staged on the device. In those circumstances, the bandwidth requirement increases to a maximum of 5 Mbps only when device image download is needed. This makes these solutions applicable even in cases where connection bandwidth is limited or noisy.

The Hybrid WAN deployment uses a CPE device such as an NFX Series Network Services platform or SRX Series Services Gateway at the customer site and thus supports private hosting of network services at a site. The distributed deployment can be extended to offer SD-WAN capabilities.

NOTE: If an SRX Series device is used as the CPE device at the customer site, it cannot host VNFs.

Building Blocks Used for Contrail Service Orchestration Deployments

IN THIS SECTION

- Administrators | 20
- Portals | 21
- Tenants | 22
- Topologies | 22
- Points of Presence (POPs) | 25
- Sites | 26
- Customer Premises Equipment (CPE) | 30
- Standalone Next-Generation Firewall (NGFW) | 31
- SD-LAN Devices | 31
- Virtual Route Reflector (VRR) | 31
- SLA-Based Steering Profiles and Policies | 32
- Path Based Steering Profiles | 33
- Intent-based Firewall Policies | 33
- Software Image Management | 33

Contrail Service Orchestration (CSO) uses conceptual and logical elements as building blocks to complete deployments in the GUI. This document provides some discussion about those elements and their use in CSO. For more detailed discussions regarding these elements, see the *Contrail Service Orchestration Administration Portal User Guide* and *Contrail Service Orchestration Customer Portal User Guide*.

Administrators

CSO uses a hierarchical, domain-based administration framework. After CSO installation, the first administrator is named *cspadmin* by default. This administrator is also known as the global service provider (SP) administrator. This SP administrator has full read and write access to all of the CSO platform from the global domain. In a cloud-hosted CSO deployment, the cspadmin role is reserved for Juniper Networks.

The SP administrator can create, edit, and delete other administrators and operators who are subject to role-based access controls (RBAC) that assign them privileges to the rest of the objects in CSO.

In an on-premises CSO deployment, the next level of administrator is the Operating Company or OpCo administrator. In a cloud-hosted CSO deployment, the OpCo admin is the highest level of administrator available to customers. In this case, the first administrator for any given OpCo is created by the SP administrator. This user has full administrative privileges within an OpCo domain. An OpCo can be thought of as a region-specific service provider within the global service provider (such as Juniper Networks). The OpCo administrator can create other administrators and operators within the OpCo domain and its tenants, but can not affect elements of the global domain. Successful login by the OpCo administrator places them into the Administration Portal of their OpCo and they can switch into the Customer Portals of any Tenant of the OpCo.

The other level of administrator is the Tenant administrator. This administrator has full access to all objects within a single tenant and can create other administrator and operator users within that tenant. The tenant administrator's login places them into the Customer Portal for that Tenant.

There are also operator users at both levels, OpCo, and Tenant. While operator users are not administrators, they can be created by administrators at each level. By default, operators have read-only access to the elements in their domain.

Portals

Portals in CSO help to separate the administrators from the customers. CSO has an Administration Portal and a Customer Portal available. Access to any given portal is controlled by a user's login privileges. If your login does not grant access to the Administration Portal, then you cannot see or access any of the elements of this portal.

Administration portals allow tenant creation and creation of other high-level objects that customers make use of within the customer portals. Administration portals are the highest level of portal within a domain.

Customer portals provide users access to a subset of the objects that exist in administration portals. The primary example of this is that an OpCo administrator can see the **Tenants** page in the Administration Portal. Each tenant name is a link that, when clicked, takes you to the customer portal for that tenant.

For more information about Administration and Customer Portals, see the *Contrail Service Orchestration Administration Portal User Guide* and *Contrail Service Orchestration Customer Portal User Guide*.

Tenants

CSO uses the tenant element to logically separate one customer from another. An OpCo administrator creates one tenant to represent each customer for which they will provide network services.

Using RBAC and other means such as virtual routing and forwarding (VRF) instances within the network, CSO keeps all tenant and OpCo objects walled within their own space. This ultimately includes the traffic that traverses the customer networks. No individual tenant, its administrators, operators, or customers can see or interact with the objects of another tenant or customer. Tenants can be named in whatever way makes most sense to the SP administrator.

Topologies

There are four network topologies supported in CSO. When defining a tenant, the OpCo administrator must decide which topology type to assign to the tenant:

- **Service Provider (SP) Cloud Topology**—This is generally assumed to be a traditional MPLS topology including provider edge (PE) routers, provider routers (P) and other resources that are owned and managed by the SP.

  > **NOTE:** In cloud-hosted CSO releases, the OpCo administrator may have no access or read-only access to the SP Cloud and any of its components.

- **Standalone Topology**—This topology is one in which the customers, or users of network services remain separate from each other with no means of communication amongst themselves.

  This is the topology of the Hybrid WAN, solution wherein the SP provides network services to its on-premises customers but does not allow them to communicate with one another. Figure 6 on page 23 shows an example where the virtual network functions (VNFs) are located at an on-premises site, but the site has no access to other sites belonging to the tenant.

Figure 6: Distributed CPE (or Hybrid WAN)

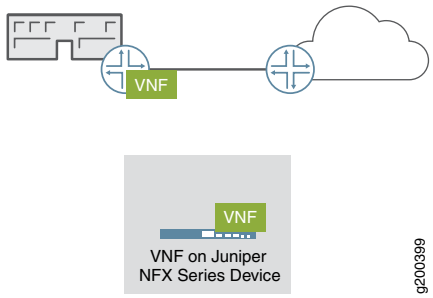

NOTE: For more information regarding network function virtualization (NFV) and VNF, see "Appendix A - Network Function Virtualization in Contrail Service Orchestration" on page 119.

It is also the topology of the NGFW and SD-LAN solutions. The NGFW solution provides for remote site security with SRX Series next-generation firewall devices. The SD-LAN solution provides for remote site LAN management with EX Series LAN access switches and Virtual Chassis. Figure 7 on page 23 and Figure 8 on page 24 below show high-level examples of these two solutions.

Figure 7: Standalone NGFW

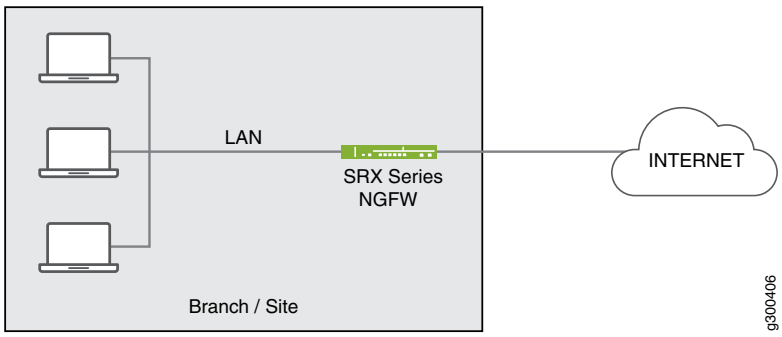

Figure 8: SD-LAN Solution with EX Switch

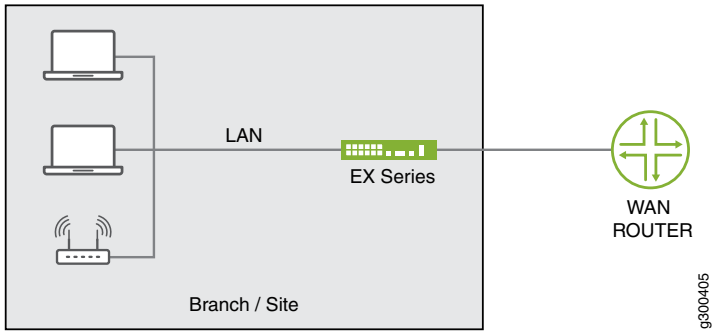

- **Hub–and–Spoke Topology**—This topology is available for SD–WAN deployments. Given that SD–WAN is intended specifically to enable and enhance the efficacy of WAN communication using network overlays, this topology does allow for communication from site to site. Specifically, if one site needs to communicate with another site, that communication goes through the hub on its way to the other site. Figure 9 on page 24 shows a very basic example of a hub–and-spoke topology. VNFs can be deployed at any of the locations shown.

Figure 9: Hub-and-Spoke Topology

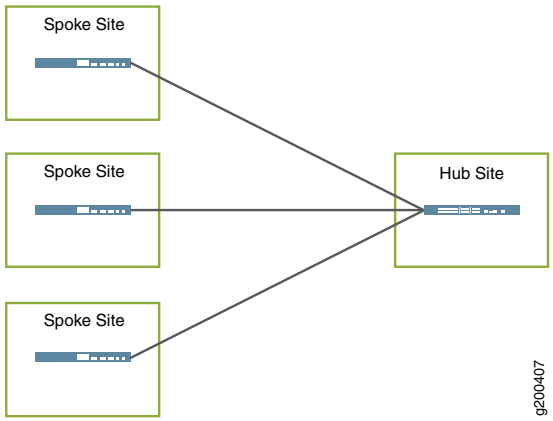

- **Dynamic Mesh Topology**—This topology is also available for SD-WAN deployments. Direct site–to–site communication is allowed and every site is considered a hub site. Figure 10 on page 25 shows a very basic example of a full mesh topology. VNFs can be deployed at any of the locations shown. This topology requires more overlay networks than the hub–and–spoke topology so CSO allows for the creation of a full mesh topology as a construct, but the tunnels from one site to another are created dynamically, (or on-demand) based on traffic thresholds thereby conserving resources and improving overall performance.

Figure 10: Dynamic Mesh Topology

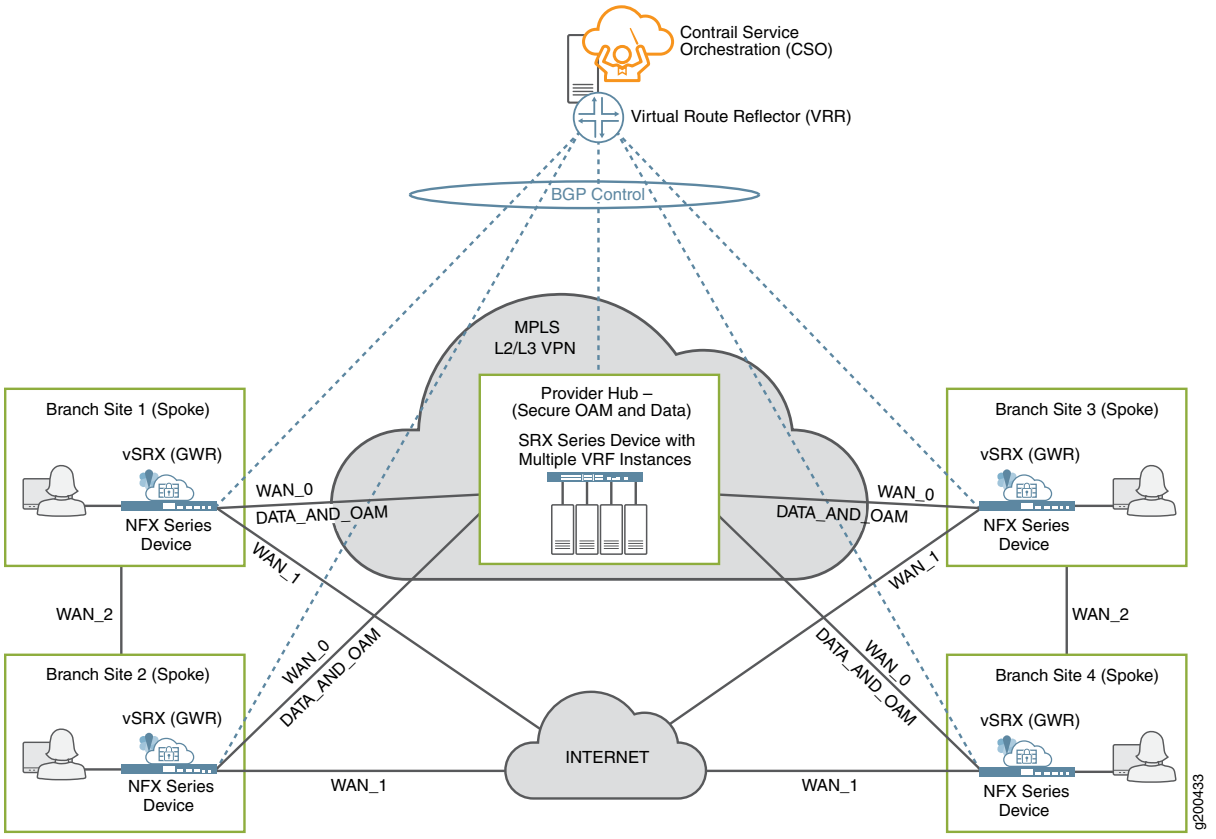

In addition, tunnelling requires the use of mesh tags. Each WAN interface on a CPE device in a dynamic mesh topology is configured with a mesh tag. Tunnels can only be formed between interfaces with matching mesh tags.

Points of Presence (POPs)

A POP is a placeholder, usually at the telco edge or enterprise datacenter, where network services can be deployed and underlay network connections are made to remote sites, as shown in Figure 11 on page 26. POPs can have PE routers and provider hubs (both data and OAM type).

Figure 11: Points of Presence (POPs)

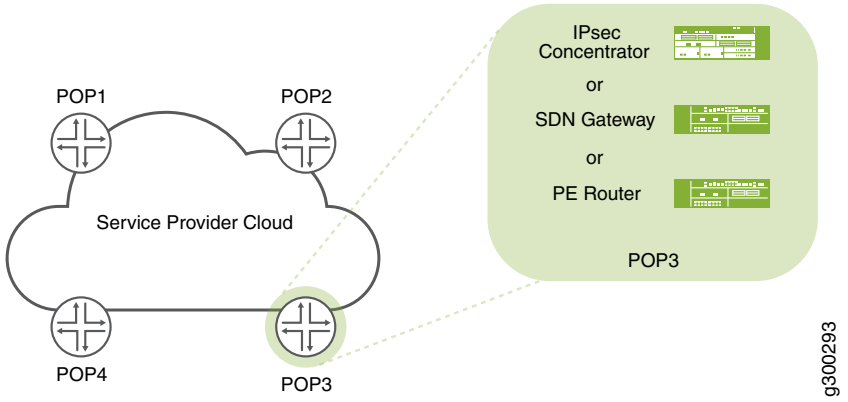

POPs are used in Hybrid WAN and SD-WAN deployments as a way to locate network access and network services closer to the users who need them. Different network services and different connection types can be offered at each POP, depending on need and availability. POPs can be named in whatever way makes the most sense to the SP administrator.

Sites

Sites are the branch offices or remote locations from which customers access the network services provided by the CSO solutions. A site is assigned to a POP and the type of sites available depend on the type of deployment you are creating: SD-WAN, Hybrid WAN, Next Gen Firewall, or SD-LAN. Sites are created by the Tenant administrator. Sites can be named whatever makes sense for the Tenant. Table 4 on page 27 lists what types of sites can be created within each deployment.

Table 4: Site Types by Deployment

Deployment	Available Site Types	Uses	Service Notes
SD-WAN	On-Premises Spoke	Use this site type for placing NFX Series or SRX Series devices at customer sites in either a hub–and–spoke or full mesh topology.	SRX Series devices deployed as on-premises spoke devices can not host VNF–based network services. NFX devices used as on–premises spoke devices can support ADSL, VDSL, and LTE access links, which can also be used for ZTP. The DSL access links allow configuration of PPPoE. Starting with CSO Release 4.0, LTE access links can be used as primary DATA, OAM, or DATA_OAM links. **NOTE:** ZTP using an xDSL interface will not work if the link is PPPoE. If the link is bridged and uses DHCP, then ZTP will work on xDSL interfaces. Local breakout is supported on this type of site when using the dynamic mesh topology.
	Cloud Spoke	This type of site is specifically for deploying a vSRX in a tenant's Amazon Web Services (AWS) Virtual Private Cloud (VPC)	Firewall and UTM services are available to protect the customer's resources in an AWS VPC. Connectivity between VPC resources and on-premises sites. WAN_0, WAN_1, and LAN interfaces need to be predefined in the VPC. Two elastic IP addresses need to be reserved in the VPC to attach to WAN interfaces later. VPC should be created and attached to an Internet gateway. Only hub-and-spoke topology is supported. The hub needs to have public IP addresses on its WAN interfaces. The hub WAN interface type should be set as **Internet** during onboarding.
	Provider Hub		

Table 4: Site Types by Deployment *(continued)*

Deployment	Available Site Types	Uses	Service Notes
		Use this type of site for placing SRX Series devices in a service provider cloud. The hub devices are used for establishment of IPSec tunnels. Provider hub devices are multi-tenant (shared amongst multiple sites) through the use of VRF instances configured on them. In a cloud-hosted CSO deployment, an OpCo or Tenant admin can create Provider Hub sites, but not the hub devices themselves. In this case, available hub devices are created by the SP administrator and made available to the lower-level administrators.	You must specify the capability of the hub devices when setting up the site. Specifying OAM capabilities (OAM Hub) allows the hub to help create secure OAM networks with the CPE devices. This option is only available for on-premises CSO deployments. For cloud-hosted CSO, data hub sites can be added only by an OpCo or tenant administrator. A hub device is required for the dynamic mesh topology. Local breakout is not supported on Hub sites.
	Enterprise Hub	Use this site type to provide additional hub-like capabilities to those of a normal spoke site.	

Table 4: Site Types by Deployment *(continued)*

Deployment	Available Site Types	Uses	Service Notes
			This type of site has the following capabilities: • Can behave as a normal spoke. • Anchor point for spokes for dynamic VPN creation. • Provides an on-premises central breakout option. • Can host a data center department. • Can import BGP and OSPF routes from the LAN–side L3 device to create a dynamic LAN segment. • Automatically meshed with other gateway sites that belong to the same tenant. • Regular spoke sites can be assigned to associate with a gateway site. • Supports local, central, and cloud breakout profiles with intent-based rules for more granular breakout control.
Hybrid WAN/Distributed CPE	On-Premises Spoke	Use this site type for locating NFX Series or SRX Series devices at customer sites.	SRX Series devices deployed as on-premises spoke devices can not host VNF–based network services. NFX devices used as on–premises spoke devices can support ADSL, VDSL, and LTE access links, which can also be used for ZTP. The DSL access links allow configuration of PPPoE. LTE access links can be used as primary DATA, OAM, or DATA_OAM links. **NOTE:** ZTP using an xDSL interface will not work if the link is PPPoE. If the link is bridged and uses DHCP, then ZTP will work on xDSL interfaces. Local breakout is supported on this type of site when using the dynamic mesh topology.

Table 4: Site Types by Deployment *(continued)*

Deployment	Available Site Types	Uses	Service Notes
SD-LAN	Switch	Use this site type to access EX Series switches in a branch location.	CSO can manage EX Series switches located behind an SD-WAN spoke site or NGFW security device.
	Access Point	CSO automatically recognizes any Mist WiFi access points attached to a switch in a branch location.	For SD-WAN and NGFW environments, CSO can detect and manage Mist access points located behind a switch.

Customer Premises Equipment (CPE)

CPE devices are those devices that are placed at remote locations in the site types mentioned previously. CPE devices serve their functions as on-premises spoke devices in Hybrid WAN or SD-WAN deployments. Figure 12 on page 30 shows available CPE device types.

Figure 12: CPE Devices

NFX150 and NFX250 Series Network Services Platforms, SRX300, SRX 550M, SRX1500, SRX4100, SRX4200, and vSRX Series Services Gateways can all be deployed as CPE devices. The NFX series devices provide the ability to host VNFs that can be deployed within the Hybrid WAN and SD-WAN solutions. The SRX Series devices cannot host VNFs but can provide their built-in security functions of firewall, UTM, and NAT as protection for the customer sites. In these cases, VNFs can still be deployed behind the SRX, but those VNFs cannot be managed by CSO.

When using SRX1500 or SRX4000 Series Services Gateways, you can create an enterprise hub site that helps implement the on-demand IPsec tunnels used in dynamic mesh topologies.

Standalone Next-Generation Firewall (NGFW)

SRX Series devices can be used as standalone firewalls, managed by CSO in the customer LAN. CSO supports the use of SRX300, SRX320, SRX345, SRX550M, SRX4100, and SRX4200 line of devices as well as the vSRX for this purpose. In this next-generation firewall (NGFW) scenario, the SRX acts as a CPE device but provides no site–to–site or site–to–hub communications as with an SD-WAN solution.

You can add LAN capabilities along with or after the deployment of an NGFW site.

SD-LAN Devices

EX Series LAN access switches can be used as CPE devices to provide managed LAN services at branch sites. This SD-LAN solution supports the use of the EX2300, EX3400, EX4300, EX4600, and EX4650 line of switches in either a standalone or Virtual Chassis configuration. These switches provide CSO-managed LAN capabilities, and you can deploy them behind an unmanaged WAN router, a CSO–managed CPE device, or NGFW device. In addition, you can add Mist WiFi access points behind the switches to provide both wired and wireless services.

In addition, CSO supports dynamic routing protocols such as BGP and OSPF in the local LAN. Therefore, when SD-LAN is configured using any of the above devices, routes to the site LANs can be updated dynamically with BGP or OSPF. For more information, see “Hybrid WAN Deployment Architecture” on page 112.

Virtual Route Reflector (VRR)

The VRR is part of CSO's SD–WAN controller. It is one of the virtual machines that get provisioned and installed during the installation process. To facilitate the routing needed in the SD–WAN deployment , the VRR forms BGP sessions with CPE spokes and hub devices using the underlay interface designated as OAM or OAM_AND_DATA during the configure site GUI workflow for site onboarding. The OAM interfaces can be implemented using dedicated IPSec tunnels which allows CPE and hub devices to be behind NAT. Figure 13 on page 32 illustrates the concept of the VRR

Figure 13: VRR Overview

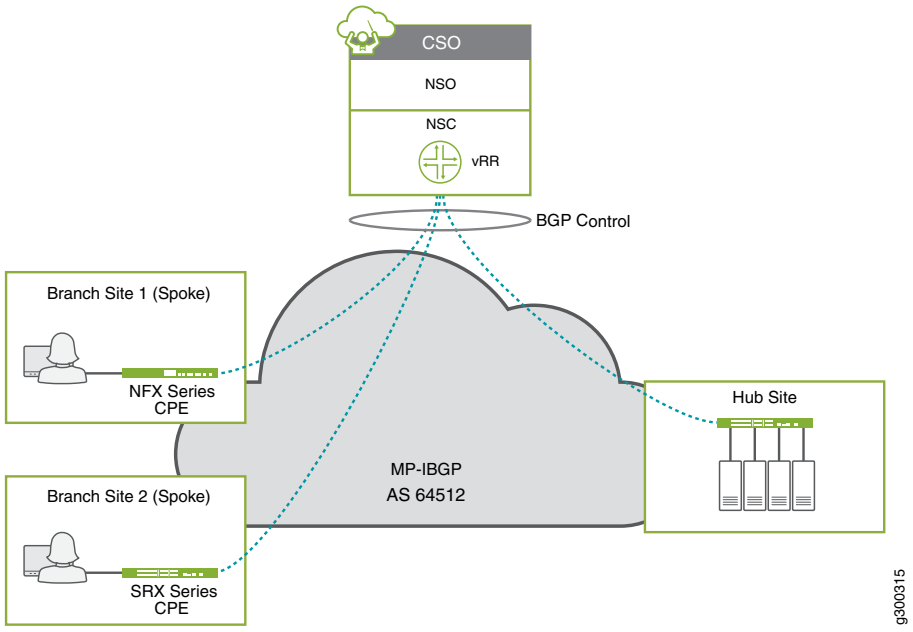

SLA-Based Steering Profiles and Policies

CSO allows for the creation of SLA–Based steering profiles that can be mapped to SD–WAN policy intents for traffic management in an SD–WAN deployment. The profiles are designed to steer traffic to a specific WAN link based on SLA parameters such as packet loss, round trip time (rtt), and jitter thresholds. SLA steering profiles are created for global application traffic types for all tenants. An SLA profile consists of a set of configurable constraints that can be defined in the Administration Portal.

You can set:

- Path preference for each of the connection paths from site–to–site
- Path preference for each of the connection paths from site–to–hub
- Threshold parameters for throughput
- Threshold parameters for packet loss
- Threshold parameters for latency
- Threshold parameters for jitter
- Class of service for various types of traffic
- Rate limiters to control upstream and downstream traffic rates and burst sizes

Once the steering profile exists, an intent-based SD–WAN policy can be created that applies that profile to specific sites or departments and against specific types of application traffic such as SSH and HTTP.

> **NOTE:** When creating an SLA profile, you must set either path preference or one of the SLA parameters. Both fields cannot be left blank at the same time.

See *SLA Profiles and SD-WAN Policies Overview* for more details.

Path Based Steering Profiles

Path based steering profiles are a simplified way to steer global application traffic types onto a specific WAN path. With these profiles, you do not need to configure any SLA parameters. All you need to do is specify which available path you want a specific traffic type to take. Just as with SLA steering profiles, you can set rate limiting parameters for these profiles. You must also assign these profiles to an SLA policy before they take effect.

Intent-based Firewall Policies

Accessed through the Customer Portal, CSO presents firewall policies as *intent-based* policies. Firewall policies provide security functionality by enforcing intents on traffic that passes through a device. Traffic is permitted or denied based on the action defined as the firewall policy intent. If your intention is to block HTTP-based traffic from social media sites, but allow HTTP-based traffic from Microsoft Outlook, you can create an intent policy to do that.

See *Firewall Policy Overview* for more information.

Software Image Management

The CSO Administration Portal allows SP administrators (cspadmin) to upload device software images and VNF images on the **Resources > Images** page. The cspadmin user in an on-premises CSO deployment can upload device images for supported SRX Series devices (including vSRX), NFX Series devices, and EX Series devices. He or she can also upload VNF images created in the Designer Tools applications.

For cloud-hosted versions of CSO, an OpCo administrator can see the images that have been uploaded to CSO by Juniper Networks. He or she can also stage and deploy uploaded device images to CPE devices and EX Series access switches.

2

CHAPTER

Deployment Tools

Contrail Service Orchestration (CSO) Deployment Tools | 35

Contrail Services Orchestration (CSO) GUIs | 35

Designing and Publishing Network Services | 37

Contrail Service Orchestration (CSO) License Tool | 38

Contrail Service Orchestration (CSO) Deployment Tools

The following sections describe the deployment tools used by CSO. While these tools are used for deployments, they are also used for other purposes in CSO.

These sections discuss:

- **Administration and Customer Portals**

  These are web-accessible portals and provide work spaces in which CSO administrators and customers can create, view, or change the tenants, sites, devices, policies, and other objects used in CSO deployments.

- **CSO Designer Tools**

  These are tools with which you can create, modify, and deploy network services into CSO. The designer tools allow you to create custom services based on Juniper or third-party virtual network functions (VNFs).

  > **NOTE:** CSO Designer Tools are only available for on-premises CSO deployments.

- **CSO License Tool**

  The license tool allows you to install and maintain software licenses on deployed devices and to track CSO license installation.

Contrail Services Orchestration (CSO) GUIs

Access to CSO's GUI interfaces is achieved using a web browser. This document briefly describes how to access the various CSO GUI interfaces.

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

See Table 5 on page 36 for information about logging into the Contrail Service Orchestration GUIs.

Table 5: Access Details for the GUIs

GUI	URL	Login Credentials
Administration Portal	• For cloud-hosted CSO: 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 on-premises CSO: 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.
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.

Designing and Publishing Network Services

NOTE: This section is only relevant for on-premises deployments of Contrail Service Orchestration (CSO).

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 cloud-hosted deployments of CSO. 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 10.2.2.12, then the URL for Designer Tools would be: **https://10.2.2.12:83**. Figure 14 on page 37 shows an overview of the workflow used within the Designer Tools application.

Figure 14: Designer Tools Overview

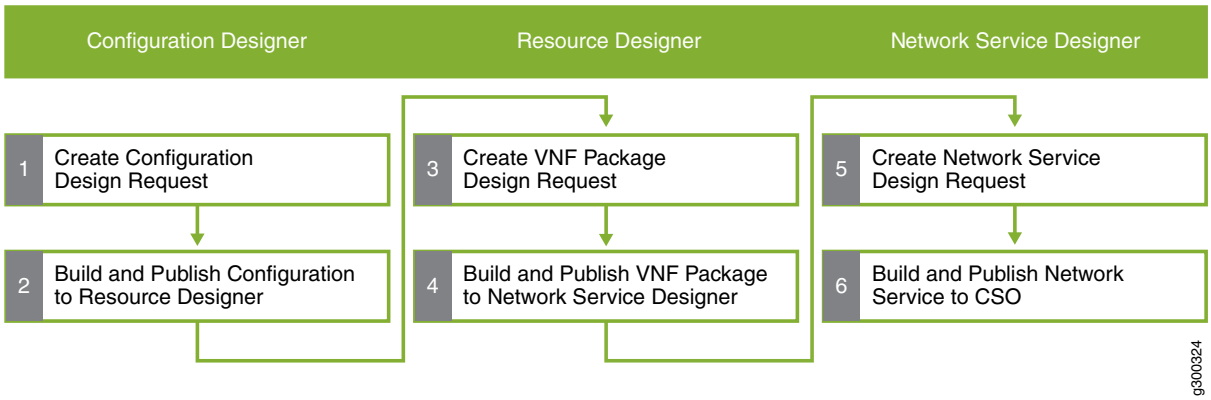

- First, you use the *Configuration Designer* to create configuration templates for virtualized network functions (VNFs). The configuration templates specify the parameters that the customer can configure for a network service.
- Then, you use the *Resource Designer* to create VNF packages. A VNF package 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 Hybrid WAN, and 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 (L2) service chains while Hybrid WAN deployments can support L2 and L3 service chains.

Contrail Service Orchestration (CSO) License Tool

IN THIS SECTION

- Overview of the License Pages | 38

Overview of the License Pages

CSO licenses come in two types: CSO software licenses and CPE platform licenses for hardware and Junos. CSO allows you to manage both the CSO licenses and any devices licenses that you use on your CPE devices. The following sections describe each of the license management pages.

> **NOTE:** For cloud-based CSO, Juniper Networks adds the licenses you have purchased on your behalf. For on-premises CSO, the SP Administrator adds the licenses.

SRX and vSRX Series devices can be used in both the Hybrid WAN and SD–WAN solutions as CPE devices or as provider or enterprise hubs. These devices require licensing in order to perform the functions needed for those solutions. Contrail Solutions Orchestration (CSO) provides a GUI-based method for loading licenses into CSO and installing them on the devices. The device licensing page is available in the Administration Portal or the Customer Portal by navigating to **Administration > Licenses > Device Licenses**. Licenses must first be purchased through your Juniper Networks account team or reseller. Once purchased, the text of the license is emailed to you.

The license page can be used to push licenses to the following devices.

- The following items in a Hybrid WAN solution:
  - vSRX gateway router on an NFX Series device

  - vSRX or SRX Series CPE devices
- vSRX, SRX Series, or NFX Series CPE devices in an SD–WAN solution
- SRX Series or EX Series CPE devices in an SD–LAN (managed branch) solution.

To upload a license to CSO for later push to an SRX, NFX, or EX device:

1. Login to CSO as an authorized user

   License management is available to tenant administrators. Operators can view, but cannot upload licenses to CSO or push them to devices.

2. Navigate to the **Administration > Licenses > Device Licenses** page.

   Here you can see a list of license files that have been uploaded to CSO. The list is empty if there have been no licenses uploaded.

3. Click the Add icon (**+**) at the top-right part of the list.

   The Add License page appears.

4. Click the **Browse** button to locate the license file that was e-mailed to you.

   Each file uploaded should be for one feature only. License files are generally named as the device serial number for which they are intended and have a **.txt** file extension.

5. (Optional) Enter a description of the license file.

   If uploading multiple licenses for a single device, a description can help you know which is which in the license list.

6. Click **OK** once you have filled in the required data.

   The license file will appear in the list along with the upload date, and your login under the **Uploaded By** column.

To install, or push, an uploaded license onto a device:

1. Click on the line or in the checkbox next to the appropriate license file.

2. Click the **Push License** pull-down menu and select **Push**. A pop-up window will appear.

   If you are logged in as a tenant administrator, you will see a list of sites and their assigned devices for your tenant.

3. Select the appropriate device, and click **Push Licenses**.

   Multiple licenses can be pushed to a single device.

The SP Administrator adds CSO licenses to the application. You can assign the added licenses to your tenants. The following procedure describes this process.

1. Click **Administration > Licenses > CSO Licenses**

   The CSO Licenses Page is displayed. All assigned licenses and the license counts appear in the list

2. Click the checkbox next to the license you want to assign.

3. Click the **Update Assignment** button.

   The **Assign CSO License** window appears and shows the quantity for this license and the number available for assignment to tenants

4. From the **Tenants** section, click the Add icon (**+**) button to add a new assignment.

   A new row on the list will appear.

5. From the **Tenant** pull-down menu, select the tenant.

6. Enter the number of licenses to assign to this tenant in the **Quantity** field. Alternatively, you can click the up and down arrows on the right of the field until the appropriate number appears in the field.

7. Click **OK**.

   The window will close and the **CSO Licenses** page will update immediately.

For more information about CSO licenses, see *About the Device License Files Page*.

3

CHAPTER

SD-WAN Deployment

SD-WAN Deployment Overview | 42

Contrail SD-WAN Deployment Architectures | 42

Initial SD-WAN Deployment | 64

SD-WAN Deployment Overview

This walkthrough highlights the steps that you need to complete to deploy an SD-WAN solution using the hub-and-spoke topology with the provider hub device located in the service provider's cloud. We use an NFX250 Series device as the CPE and an SRX Series device as the provider hub in the SP cloud. We indicate where in the CSO GUI you need to go to complete each step. The document also provides some explanation of the choices that you need to make. It assumes that this is the first SD-WAN deployment you are attempting.

Additional information about using the Administration Portal GUI for any of the steps below can be found in the *Contrail Service Orchestration Administration Portal User Guide*.

The topology shown in Figure 15 on page 42 is a reference for this SD-WAN deployment.

Figure 15: SD-WAN Example Deployment Topology

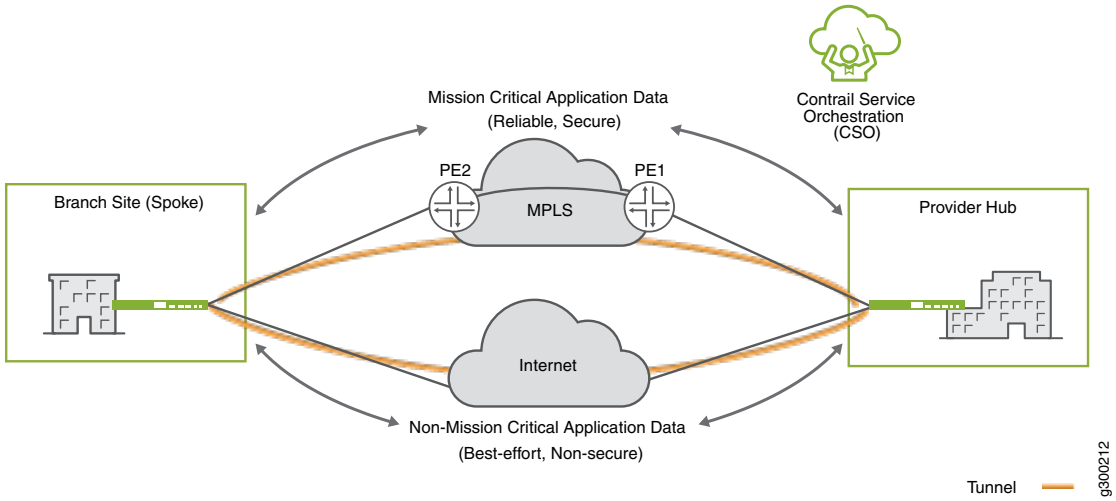

Contrail SD-WAN Deployment Architectures

An SD–WAN implementation offers a flexible and automated way to route traffic from site to site. As shown in Figure 16 on page 43, a basic SD–WAN architecture includes just a few basic elements

- Multiple sites
- Multiple connections between sites that form the underlay network

- Multiple overlay tunnels
- A controller

Figure 16: SD-WAN Architecture

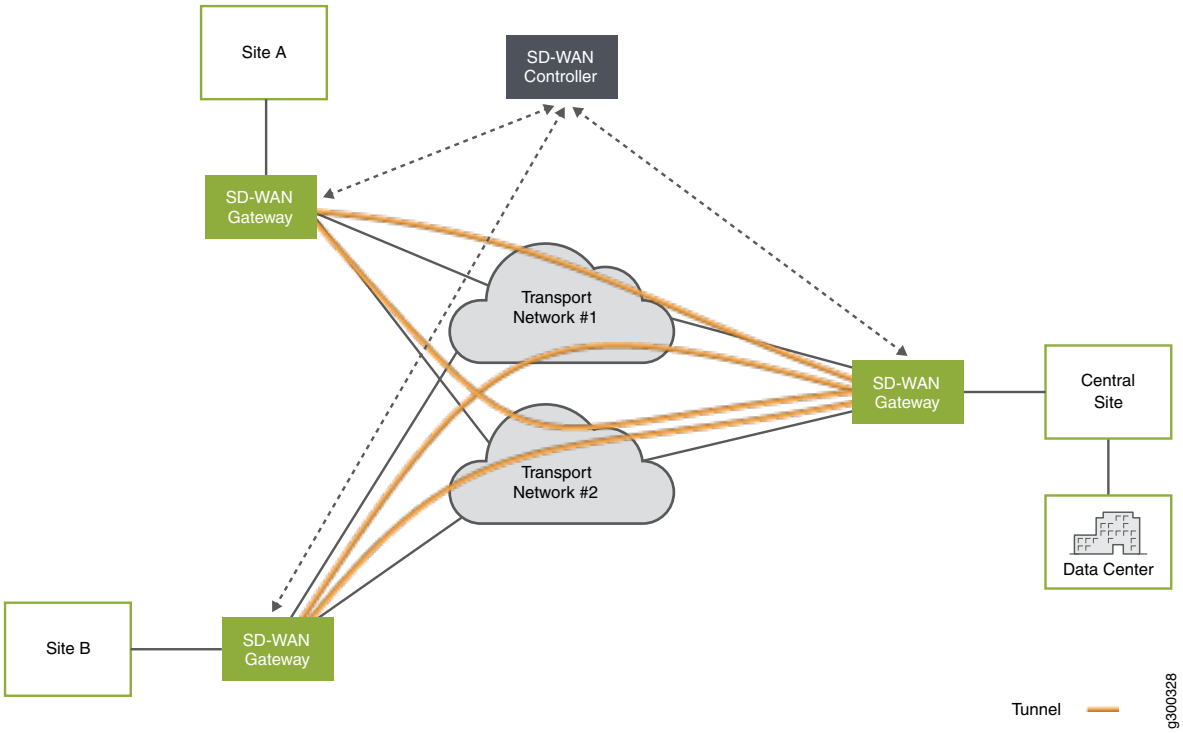

The SD–WAN controller, built in to CSO, acts as an orchestration layer and provides an interface, allowing the operator to setup and manage the devices at the sites.

Contrail SD-WAN Reference Architecture

Juniper Networks Contrail SD–WAN solution architecture, shown in Figure 17 on page 44 uses a hub–and–spoke topology and is based on the Hybrid WAN model, with CPE devices located at customer branch sites. On the local side of the site, the CPE devices connect to LAN segments and participate in dynamic routing protocols with other LAN devices. On the WAN side, the CPE devices connect across two or more links to a provider hub device. Because the SD-WAN model uses a hub–and–spoke topology, traffic travels from site to site through the provider hub. By default, traffic going to the Internet also flows through the provider hub device.

Figure 17: Contrail SD-WAN Reference Architecture

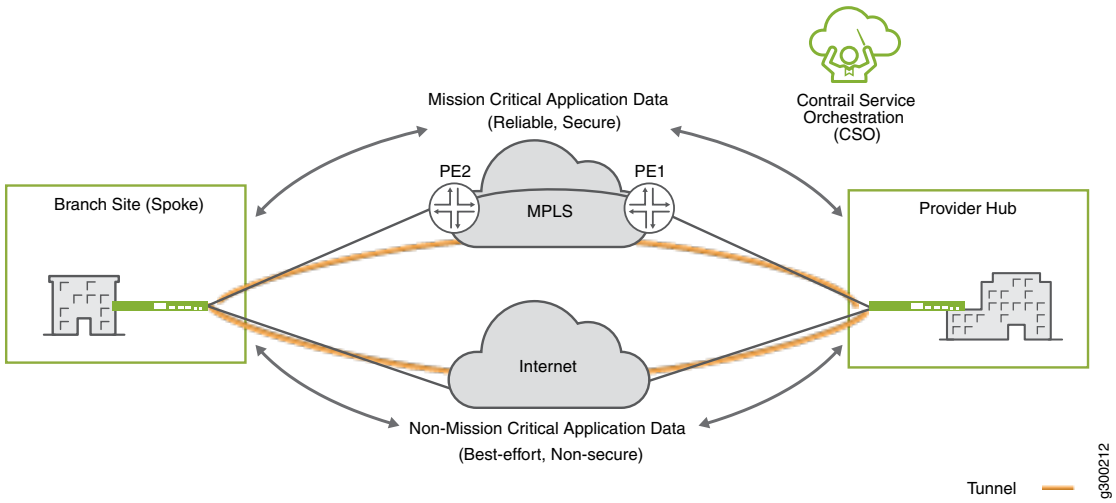

The SD-WAN orchestrator and controller functions are implemented through Juniper Networks Contrail Service Orchestration (CSO) software. The CSO platform uses policies and SLA parameters to differentiate and direct traffic flows across the available paths as desired.

The following sections describe these architectural elements in more detail.

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, as shown in Figure 18 on page 45.

Figure 18: On-Premises Spoke Devices

NFX Network Services Platform

The NFX Network Services Platform 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 Contrail 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 45 shows the supported NFX hardware and required Junos OS software release version for each supported model.

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

Table 6: NFX Hardware and Software Matrix for On-Premises Spoke Devices

Platform	Models Supported	Junos OS Software Release Versions
NFX150 Network Services Platform	• NFX150–S1 • NFX150–S1E • NFX150–C–S1 • NFX150–C–S1–AE/AA • NFX150–C–S1E–AE/AA	18.2X85–D12 19.3R2-S1
NFX250 Network Services Platform	• NFX250–LS1 • NFX250–S1 • NFX250–S2	15.1X53–D497.0 18.4R3

SRX Series Devices and vSRX Virtual Firewall

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 46 shows the supported SRX hardware and required Junos OS software release version.

Table 7: SRX Hardware and Software Matrix for On-Premises Spoke Devices

Platform	Models Supported	Junos OS Software Release Versions
SRX Series	• SRX4200 • SRX4100 • SRX550M • SRX345 • SRX340 • SRX320 • SRX300	15.1X49–D172 19.3R2-S1
	SRX1500	19.3R2-S1
vSRX Virtual Firewall	vSRX	15.1X49-D172 19.3R2-S1

NOTE: For the most up to date information on hardware and software support for CSO, see the Contrail Service Orchestration Release Notes.

Cloud Spoke Devices

A Contrail SD–WAN spoke device, in the form of a vSRX, can be located in an AWS VPC. The vSRX serves as the cloud spoke device; 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 Contrail SD-WAN Design and Architecture Guide.

Provider Hub Devices

The Contrail SD–WAN solution supports two deployment topologies (discussed later in this guide): 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 the cloud-hosted deployment of CSO, 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.

Figure 19 on page 47 and Table 8 on page 47 show the provider hub devices supported in a CSO SD-WAN environment.

Figure 19: SD-WAN Provider Hub Devices

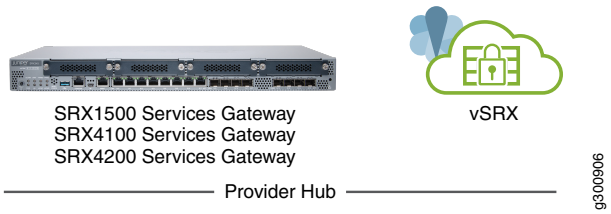

Table 8: Provider Hub Devices and Supported Software

Role	Supported Device Types	Required Junos OS Software Version
Provider Hub	• SRX4200 • SRX4100 • SRX1500	15.1X49–D172
	vSRX	15.1X49–D172

NOTE: For the most up to date information on hardware and software support for CSO, see the Contrail Service Orchestration Release Notes.

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 Contrail 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 spoke site. SRX1500, SRX4100, and SRX4200 devices can serve this function. The spoke site that functions this way, must be configured as an *enterprise hub site* during site creation. Creating 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 49 shows the enterprise hub devices supported in a CSO SD-WAN environment.

Table 9: Enterprise Hub Devices and Supported Software

Role	Supported Device Types	Required Junos OS Software Versions
Enterprise Hub	• SRX4200 • SRX4100 • SRX1500	15.1X49–D172 19.3R2-S1
	vSRX	15.1X49–D172 19.3R2-S1

NOTE: For the most up to date information on hardware and software support for CSO, see the Contrail Service Orchestration Release Notes.

Underlay (Physical) Network

The underlay network includes the physical connectivity between devices in the SD–WAN environment. This layer of the network has no awareness of the customer LAN segments, it simply provides reachability between on-premises devices.

Figure 20 on page 50 shows a sample underlay network for a hub–and–spoke SD–WAN deployment (the details apply equally to a dynamic mesh setup). Each spoke site typically has multiple paths to the hub site; in this case, one through the private MPLS cloud, and one over the Internet.

Figure 20: SD-WAN Underlay Network

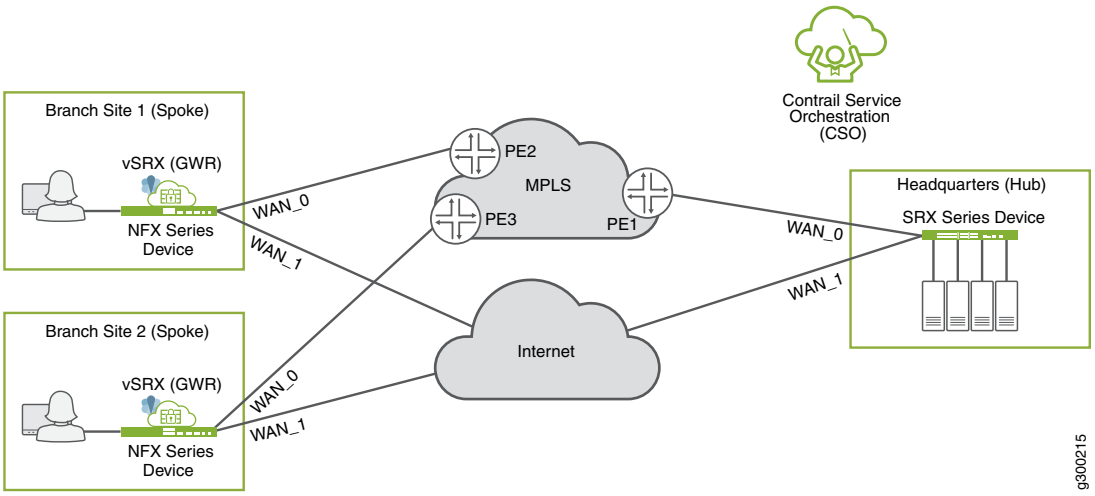

Each on-premises device (or site) can have up to four WAN links, including a satellite link that can be used for OAM. During configuration, CSO identifies the devices' WAN–facing interfaces as WAN_0 through WAN_3.

Note that:

- The WAN interfaces can be VLAN tagged or untagged, as per design requirements.
- The on-premises devices' Internet-facing interfaces can be attached to different service provider networks.

WAN Access Options

Each WAN access type listed below can be used for ZTP, data, or OAM traffic. All the links can be leveraged for data traffic simultaneously.

- MPLS
- Ethernet
- LTE

NOTE: LTE WAN access supported using a dongle on NFX250 Series devices.

LTE WAN access supported using a built-in interface on NFX150 Series devices.

LTE WAN access supported using a mini–PIM in slot 1 of SRX300 Series devices.

All of the previously mentioned LTE interfaces are supported for ZTP.

Only supported for Hub–and–Spoke SD–WAN deployments with single CPE.

Full/Dynamic Mesh deployments are not supported.

Dual CPE configurations are not supported.

LTE APN settings can be localized for the installation region during the ZTP process.

- ADSL/VDSL (ADSL/VDSL support for WAN links and ZTP on NFX Series devices starting in CSO Release 4.0.0 and on SRX300 Series devices starting in CSO Release 5.0.3.)
- Broadband
- MPLS and broadband
- Satellite link

WAN Interface Types - Data and OAM

The WAN interfaces are used primarily to send and receive user traffic (data). At least one of the WAN interfaces must also be used for management (OAM) traffic. The OAM interface is used to communicate with CSO, and allows CSO to manage the on-premises device.

Figure 21 on page 52 illustrates these two interface types.

Figure 21: WAN Interface Types

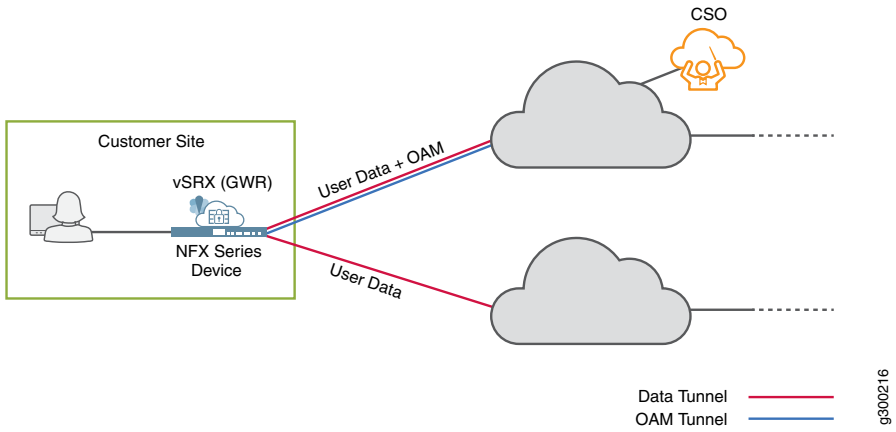

Note that:

- The on-premises device's OAM interface must be able to reach CSO. The connectivity can be supplied strictly using CSO-orchestrated overlay networks. You do not need pre-existing underlay network connectivity for this. Starting in CSO release 5.0.1, CSO automatically selects an IP address for the on-premises device's OAM interface. This ensures that the address is unique within the entire CSO deployment and prevents human error.
- To ensure secure communication over the WAN, the on-premises device initiates the connection to CSO.
- Device-initiated connections can work across intermediate NAT devices.
- The user-and-OAM-data interface can use a single IP address for both functions.

Overlay (Tunnels) Network

The overlay network includes the logical tunnel connectivity between devices in the SD-WAN environment. This layer of the network has awareness of the customer LAN segments, and is responsible for transporting customer traffic between sites.

Figure 22 on page 53 shows an overlay network for a hub-and-spoke environment. Each spoke site has two tunnels to carry traffic to the hub site: one through the private MPLS cloud, and one over the Internet.

Figure 22: SD-WAN Hub-and-Spoke Overlay

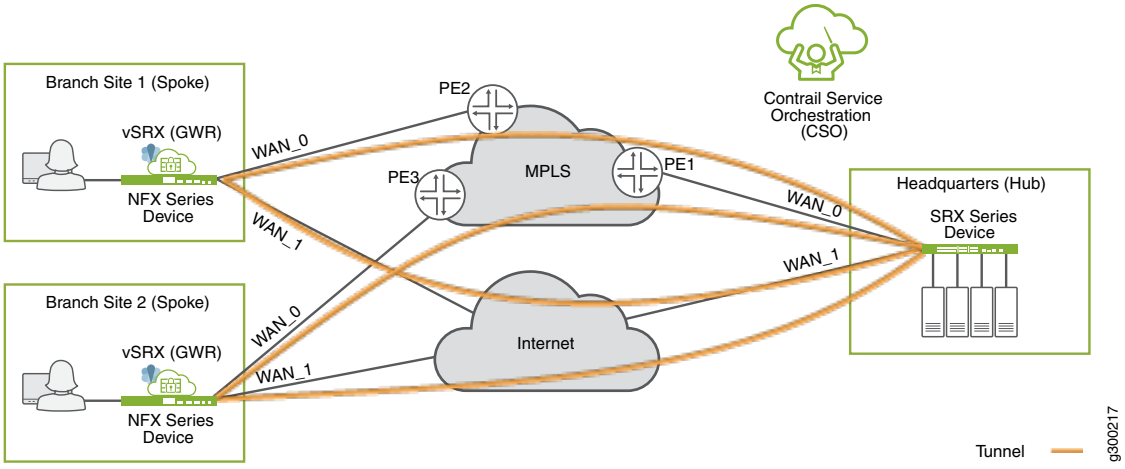

The tunnels have two encapsulation options: MPLSoGRE or MPLSoGREoIPsec. CSO automatically provisions and establishes these tunnels as part of the deployment process.

Overlay Deployment Topologies

The SD–WAN solution supports hub–and–spoke or dynamic mesh deployment topologies. A dynamic mesh topology is similar to a full mesh topology wherein every site is capable of connecting directly to any other site. But with dynamic mesh, the connections (tunnels) are brought up on-demand, thereby reducing the continual load on any one site. A single tenant can support both hub-and-spoke and dynamic mesh topologies.

Hub and Spoke

With the hub–and–spoke topology, all spoke sites are connected to at least one hub site, as shown in Figure 23 on page 54. Spoke sites cannot communicate directly with other spoke sites.

Figure 23: SD-WAN Hub-and-Spoke Topology

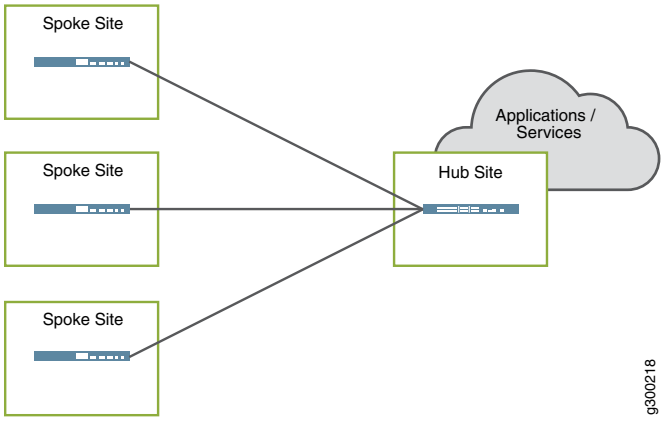

The hub sites used can be either provider hub or enterprise hub sites. When an enterprise hub site is used, the provider hub (if any) is used as backup only. This topology is preferred when applications and services are centralized at the hub site.

Dynamic Mesh

With the dynamic mesh topology, all sites are interconnected using overlay tunnels, as shown in Figure 24 on page 55, and each site can communicate directly with every other site through the tunnels. Although the figure shows the DATA_AND_OAM connection on the MPLS link, WAN_0, this function can be performed on either the MPLS or Internet links.

Figure 24: SD-WAN Dynamic Mesh Topology

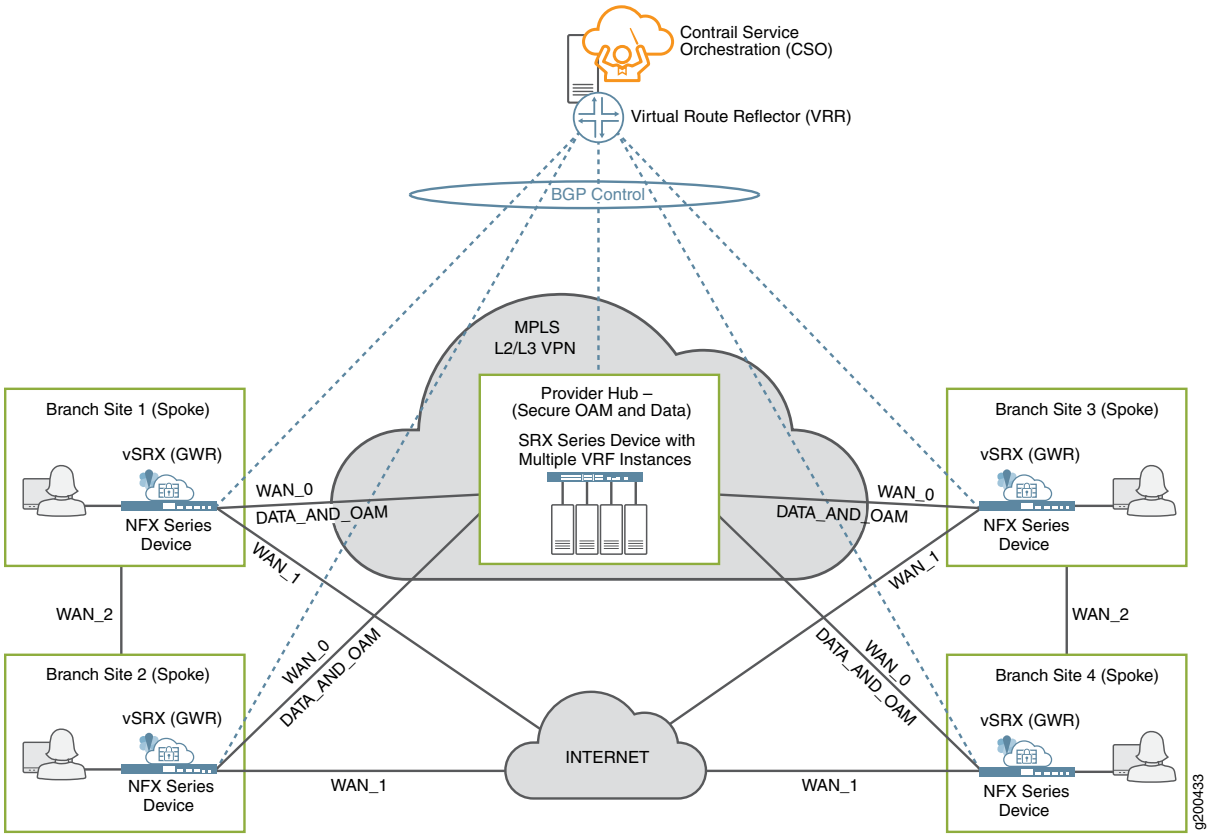

This topology is well suited for deployments where applications and services are not centralized.

> **NOTE:** Both hub–and–spoke and full mesh topologies require adding a secure OAM network overlay, and thus an OAM Hub, to the deployment.

When spoke devices are added to a dynamic mesh topology, the administrator configuring the sites must assign a mesh tag to each WAN interface. Only two devices with matching mesh tags can form the VPN connection to allow communication. Interfaces with mismatched mesh tags can never communicate directly.

Orchestration and Control

Orchestration and controller functions are implemented through Juniper's Contrail Service Orchestration (CSO) software. CSO software offers a Web–based UI to manage the SD-WAN environment, as shown in Figure 25 on page 56.

Figure 25: CSO Login Screen

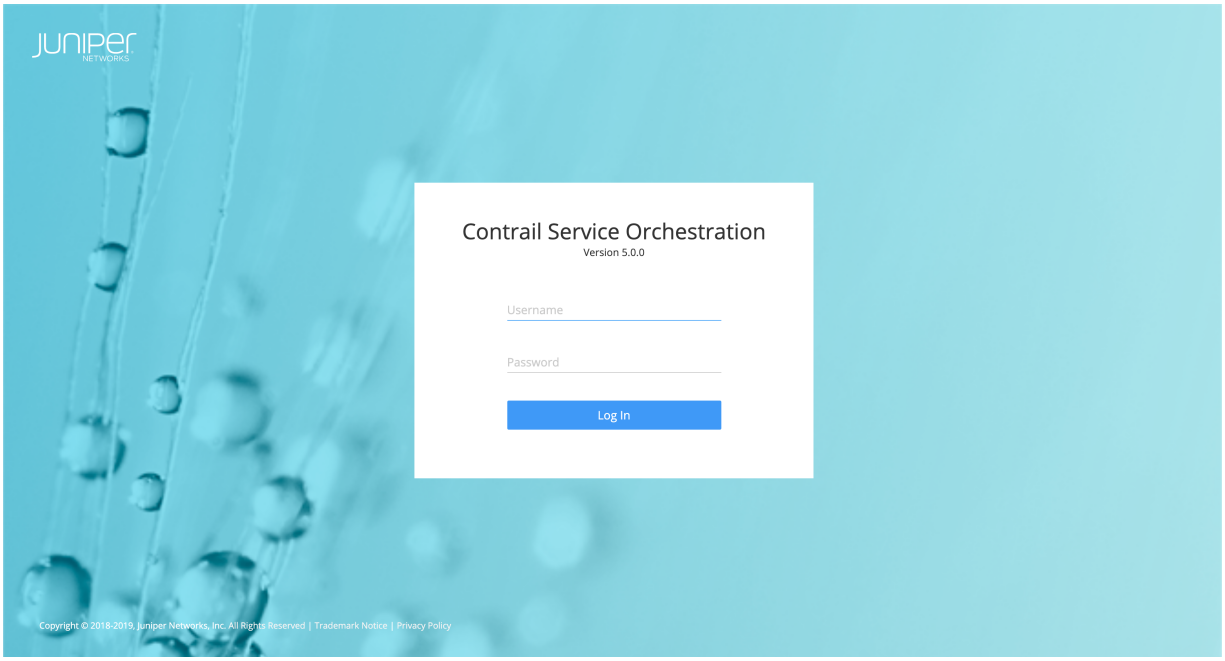

The Service Orchestration Layer contains the Network Service Orchestrator (NSO). The orchestration software has a global view of all resources and enables tenant management, providing end–to–end traffic orchestration, visibility, and monitoring. The Domain Orchestration Layer contains the Network Service Controller (NSC). The orchestration software works together with the controller to manage on-premises (CPE) devices, and provide topology and CPE lifecycle management functionality.

At the user level, CSO provides the interface to deploy, manage, and monitor the devices in the SD–WAN network through the NSC. At the network level, NSC includes a vRR that allows each site to advertise its local routes to remote sites.

For more information regarding SD–WAN architecture, see Contrail SD-WAN Design and Architecture Guide.

Secure OAM Network

SD–WAN deployments include a secure OAM overlay network to provide end-to-end secure communications between on-premises devices and CSO. This is true regardless of whether your CSO software is deployed on-premises or as a cloud-hosted deployment. In a cloud-hosted deployment, the provider hub devices, and thus, one end of the OAM network is owned and managed by the SP. As shown in Figure 26 on page 57, dedicated, IPsec-encrypted OAM tunnels enable on-premises devices to send management, routing, and logging traffic securely over the network to a provider hub. The provider hub then forwards the traffic to CSO.

Figure 26: Secure OAM Tunnels

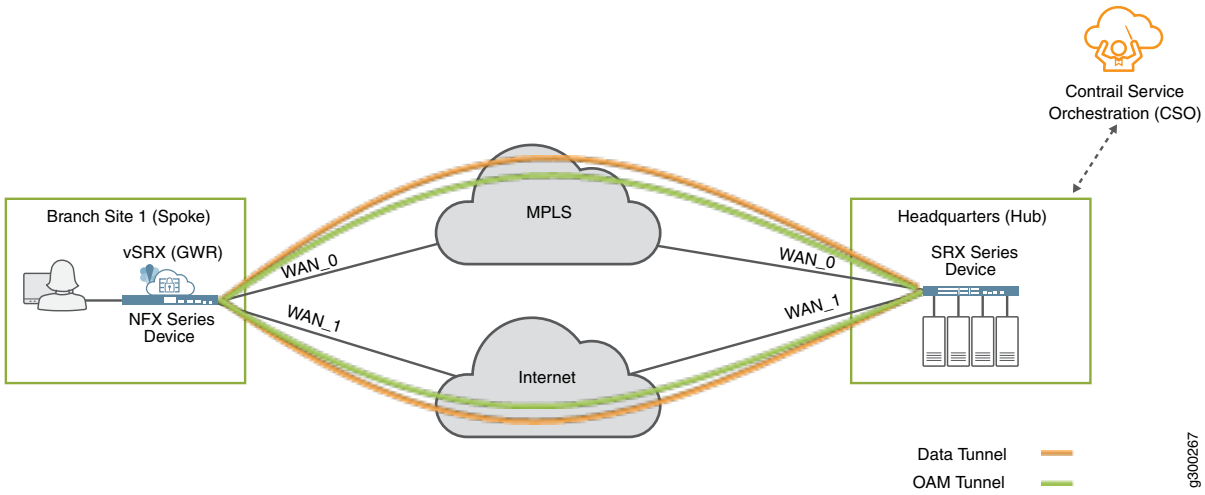

Integration with Deployment Topologies

Both the hub–and–spoke and dynamic mesh deployment topologies must use secure OAM tunnels.

Hub and Spoke

With the hub–and–spoke topology, each spoke site now has two sets of connections to the provider hub site: an overlay tunnel carrying data, and a separate, dedicated IPsec overlay tunnel carrying OAM traffic, as shown in Figure 27 on page 57.

Figure 27: OAM Tunnels in the Hub-and-Spoke Topology

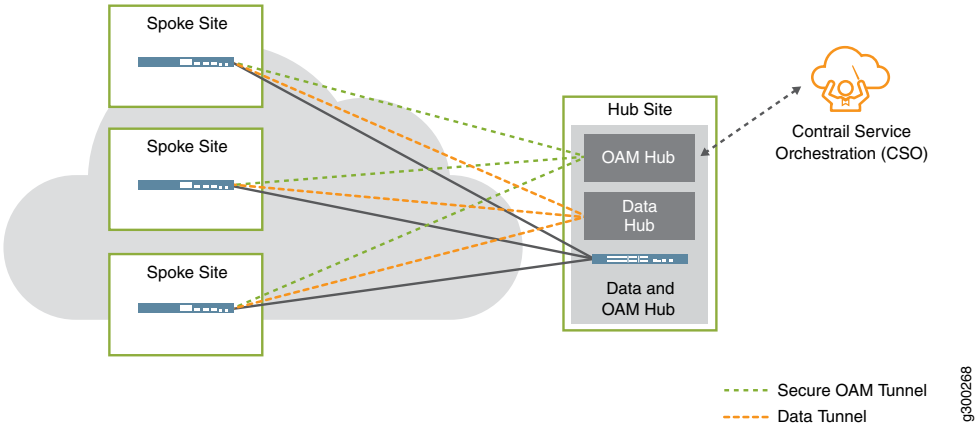

Dynamic Mesh

Since a normal full mesh topology would not include a hub device for data traffic, one must be added. As shown in Figure 28 on page 58, each spoke site has a new connection: a separate, dedicated IPsec overlay tunnel carrying OAM traffic to the provider hub.

Figure 28: OAM Tunnels in the Full Mesh Topology

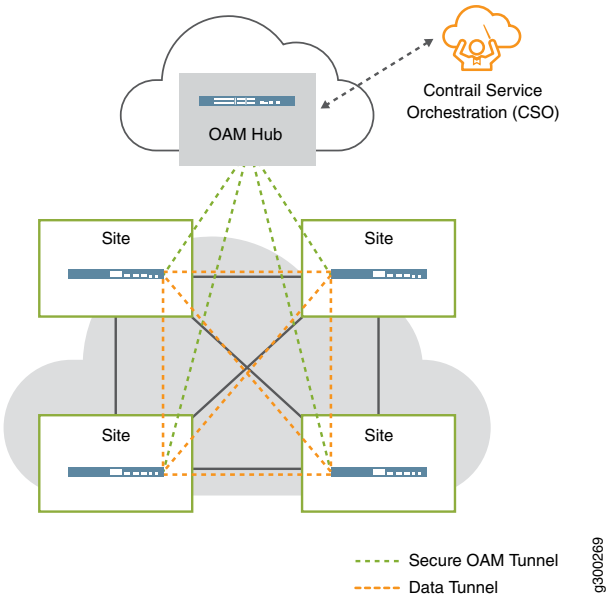

OAM Hub Design Options

There are two ways to implement the OAM hub in an on–premises CSO deployment, depending on design requirements. As shown in Figure 29 on page 59, the options are as follows:

- Data and OAM tunnels terminate on same provider hub device—This is a good option for small deployments, where the single hub device can handle both the data and OAM traffic.
- Data and OAM tunnels terminate on separate provider hub devices—This option can be useful for larger deployments where the main hub device's resources are needed to service the overlay tunnels carrying data traffic; a second hub device can be used to terminate the OAM tunnels.

Figure 29: OAM Tunnels - Provider Hub Design Options

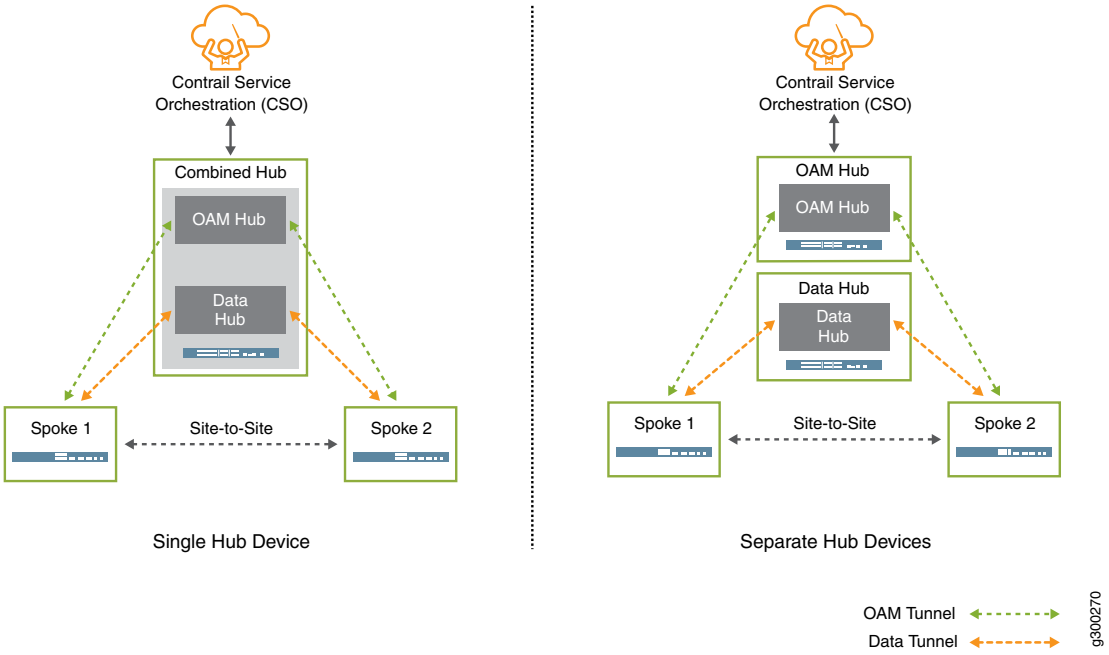

NOTE: In a cloud-hosted CSO deployment the OAM hub is provided as part of the service.

However, an OpCo administrator can deploy a DATA_ONLY or an OAM_AND_DATA hub. In the case of a DATA_ONLY hub, the DATA hub has an IPsec secured tunnel to the OAM_HUB. In the case of an OAM_AND_DATA hub, the OpCo administrator is required to set up the IPsec secured connection between the OAM_AND_DATA HUB and CSO.

Usage Notes on Provider Hub Design Options

- An OAM hub can support multiple tenants, or can be dedicated to a single tenant.
- Connectivity from the provider hub(s) to CSO should be private and secured, as it is not covered by the OAM tunnels.
- We recommended that you implement multiple OAM hubs for redundancy and to ensure no loss of management or monitoring of the on-premises devices.

  For a cloud-hosted CSO deployment, OAM hub redundancy is handled by the SP Administrator so cannot be addressed by an OpCo or tenant administrator.

- When a spoke site is multi-homed to multiple hub devices, one OAM tunnel should terminate on each hub.
- On-premises devices using NAT are supported for hub-and-spoke deployments.

Zero Touch Provisioning

One of the key features of the Contrail SD–WAN solution is the ability to “plug–and–play” new spoke devices using ZTP (autoinstallation). In CSO, the ZTP process is implemented with the help of an Internet–located redirect server. For true ZTP, the use of the redirect server is required. The redirect server itself is discussed in the next section.

A high-level list of steps performed during ZTP looks like:

- Before performing ZTP, add the appropriate CSO SSL certificate to the redirect server.
- When a spoke device first comes online, it uses a local DHCP server to obtain an IP address and name server information.
- The spoke device then contacts the redirect server, which provides the DNS name and certificate for the CSO installation.
- The spoke device then contacts the CSO server to obtain its initial configuration and Junos OS software update (if required).

> **NOTE:** CSO Release 4.1 and later include enhancements that reduce the bandwidth required for ZTP to 2 Mbps.

Usage Notes for ZTP

- At least one of the device’s WAN interfaces must obtain its IP address from a DHCP server in order to also be assigned a DNS name server and a default route.
- Both CSO and the redirect server must be reachable over the same WAN interface.
- The ZTP process can be run from any WAN interface on the spoke device, including a satellite link.
- The download of the initial configuration can require significant amount of time, especially on slow links, due to the size of configuration and Junos OS software.

Redirect Server

The redirect server is an Internet–located, Juniper–owned–and–managed server that is integral to the ZTP process. The function of the server is to hold and distribute the DNS names and SSL certificates of all CSO instances so that newly installed spoke devices can locate and authenticate to their designated CSO instance.

During site configuration, the serial number of the spoke device is entered into CSO. This leaves the specific CSO instance in a state where it is expecting contact from a device with this serial number. When the redirect server is contacted by the spoke device, it uses the serial number to:

- Pair the device with the proper instance of CSO
- Pass the proper SSL certificate to the spoke device
- Pass the DNS name of the proper CSO instance to the spoke device

With the above information, the spoke device can contact CSO and receive its initial configuration, including Junos OS software update (if required).

Design Considerations for CSO and Redirect Server

The redirect server is located on the Internet, and cannot be moved.

In Figure 30 on page 61, both the redirect server and CSO are located on the Internet. In this case, the spoke device obtains and uses IP addressing and other information provided through its Internet–facing interface, and can then reach both the redirect server (first) and CSO (second) through that same interface. This is the only CSO deployment option available for the Juniper–hosted CSO instance in Release Version 5.0.

Figure 30: CSO and Redirect Servers on Internet

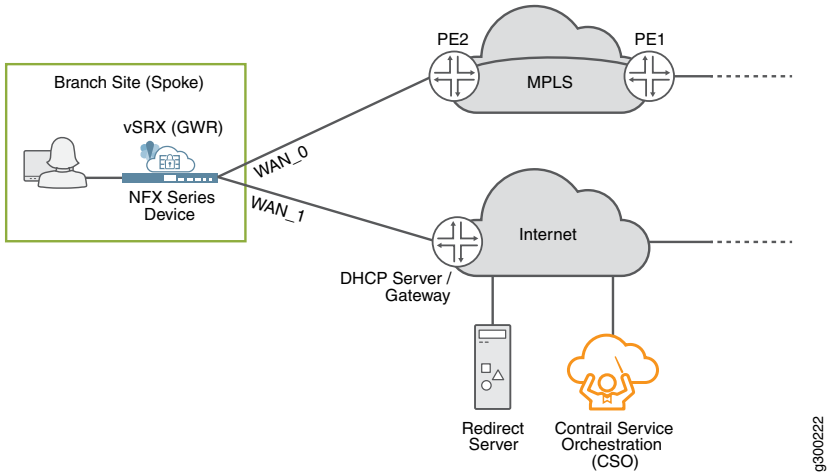

Bypassing the Redirect Server

In some cases, it may not be desirable or practical to use the Juniper redirect server. In these cases, a spoke device can be pre–staged to include reachability information for the CSO server before being shipped to the customer location.

Service Chaining in Contrail SD-WAN

Starting in CSO Release 4.0, service chaining is available for SD-WAN environments. Service chaining is a concept wherein multiple network services instantiated in software and running on x86 hardware are linked, or chained together in an end-to-end fashion. This allows the one physical device to provide the services normally provided by multiple devices. Service chaining can be performed on NFX Series devices, as shown in Figure 31 on page 62.

Figure 31: Service Chaining in an SD-WAN Environment

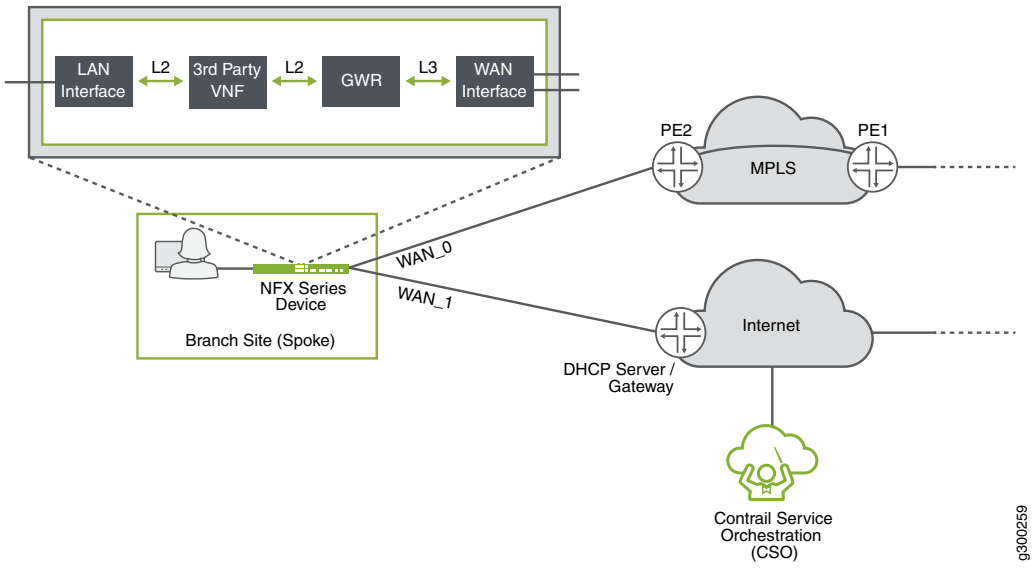

Starting in CSO Release 4.0, the following third-party virtual network functions (VNFs) are supported: *Fortigate-VM* and *Single-legged Ubuntu VM*.

NOTE:

- Currently only Layer 2 VNF mode is supported in SD-WAN service chains.
- For a complete list of supported VNFs in CSO, see the *VNF Overview* in the *Contrail Service Orchestration Designer Tools User Guide*.

Three Planes, Four Layers

To bring all of the above elements together, the Contrail SD–WAN Architecture can be thought of in three planes, comprised of four functional layers:

1. Data Plane:
   - Includes the underlay network; provides physical connectivity
   - Includes the overlay network; provides tunnels for tenant data traffic
2. Control Plane—Includes the routing protocols which flow through the OAM tunnels
3. Management Plane—Includes the overlay tunnels for the secure OAM network

Figure 32 on page 63 illustrates this concept.

Figure 32: Three Planes, Four Layers

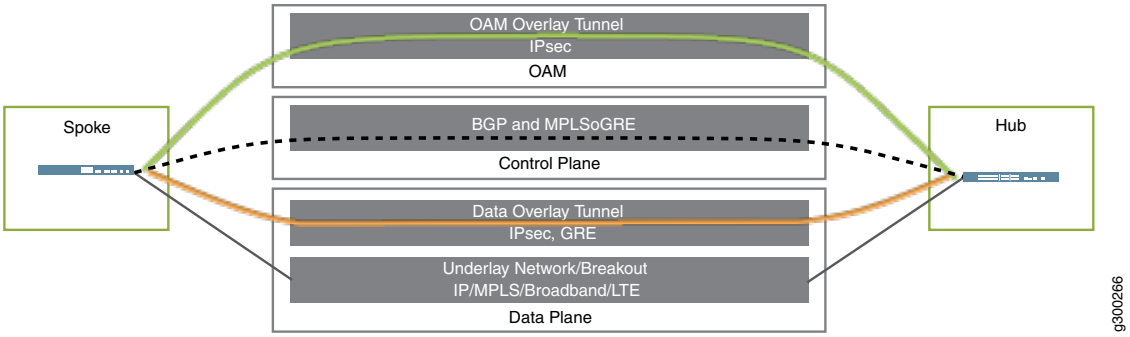

Release History Table

Release	Description
4.0	Starting in CSO Release 4.0, service chaining is available for SD-WAN environments.
4.0	Starting in CSO Release 4.0, the following third-party virtual network functions (VNFs) are supported: *Fortigate-VM* and *Single-legged Ubuntu VM*.

Initial SD-WAN Deployment

IN THIS SECTION

- Before You Begin | 65
- Download Intrusion Protection System (IPS) and Application Signatures | 65
- Upload Licenses | 68
- Add a New Tenant | 68
- Modify Device Templates | 70
- Choose a Point of Presence (POP) for the Hub | 74
- Add a Provider Hub Device to Your Tenant | 74
- Add an Enterprise Hub to Your Tenant | 78
- Add an On-Premises Spoke for the Tenant | 82
- Install a License on a Device | 88
- Install an Application Signature on a Device | 88
- Add Firewall and NAT Policies to the Topology | 89
- Add SD-WAN SLA-Based Steering Profiles and Policy | 91

This document describes the steps required to create a basic SD-WAN deployment. Figure 33 on page 65 shows an overview of the steps that will be covered in this deployment example.

Figure 33: Basic SD-WAN Deployment Workflow

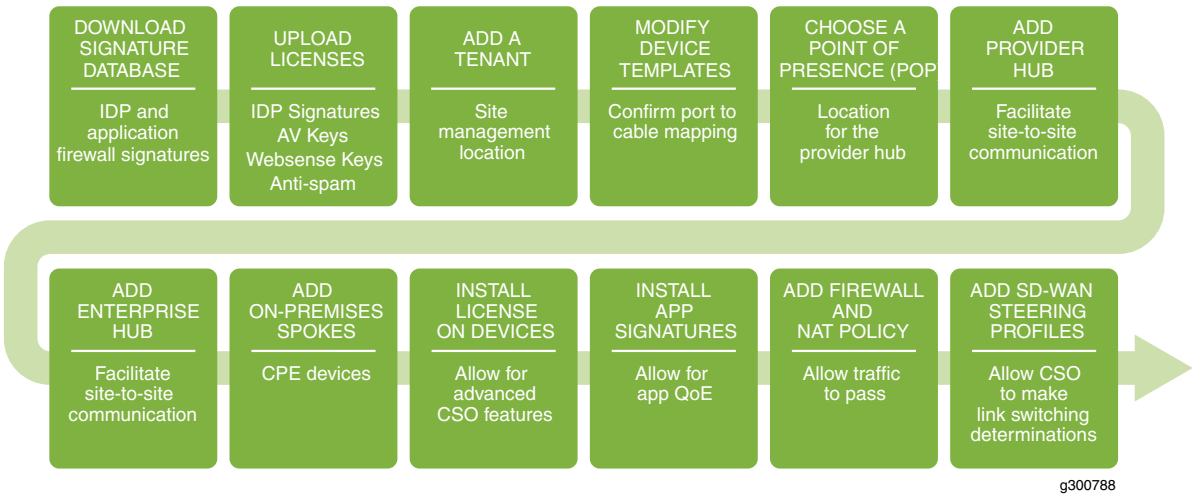

Before You Begin

This example uses hardware-based SRX devices in the roles of provider and enterprise hub and on-premises spoke devices. The vSRX Series of devices could be used in place of either the hub or spoke devices. NFX series devices could also be used for the spoke devices.

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.

Download Intrusion Protection System (IPS) and Application Signatures

This section details how to download the IPS and application firewall signature databases from Juniper Networks onto your CSO installation. Downloading the signature databases makes the IPS and application firewall signatures available to install on your hub and CPE device after it has been activated.

From this point on in this deployment example, we assume that you are logged in to CSO as an OpCo administrator.

The user name part of your credentials is the e-mail address that was used when your CSO account was set up. When an account is initially setup, CSO sends an e-mail to that address with a link that includes a one-time activation code. Clicking the link takes you to the CSO login page which then prompts you to set a password. This is a one-time activity. Subsequent logins to the Administration Portal use the e-mail address and your newly-set password as your login credentials.

> **NOTE:** If you are working with an on-premises installation of CSO, you can login as the cspadmin user (or equivalent) to perform all of these steps.

1. Enter the login credentials for the Administration Portal.

2. Navigate to the **Administration > Signature Database** page.

   On this page, there is a list of available database versions, their publish dates, update summaries, and detector versions. The active database (if there is one) is in its own section at the top of the list. There is a list of available signature databases is in the section below and it is sorted from newest (at the top) to oldest.

3. Click the **Signature Download Settings** button in the upper right corner of the window.

   The **Signature Download Settings** window appears as shown in Figure 34 on page 67.

Figure 34: Signature Download Settings

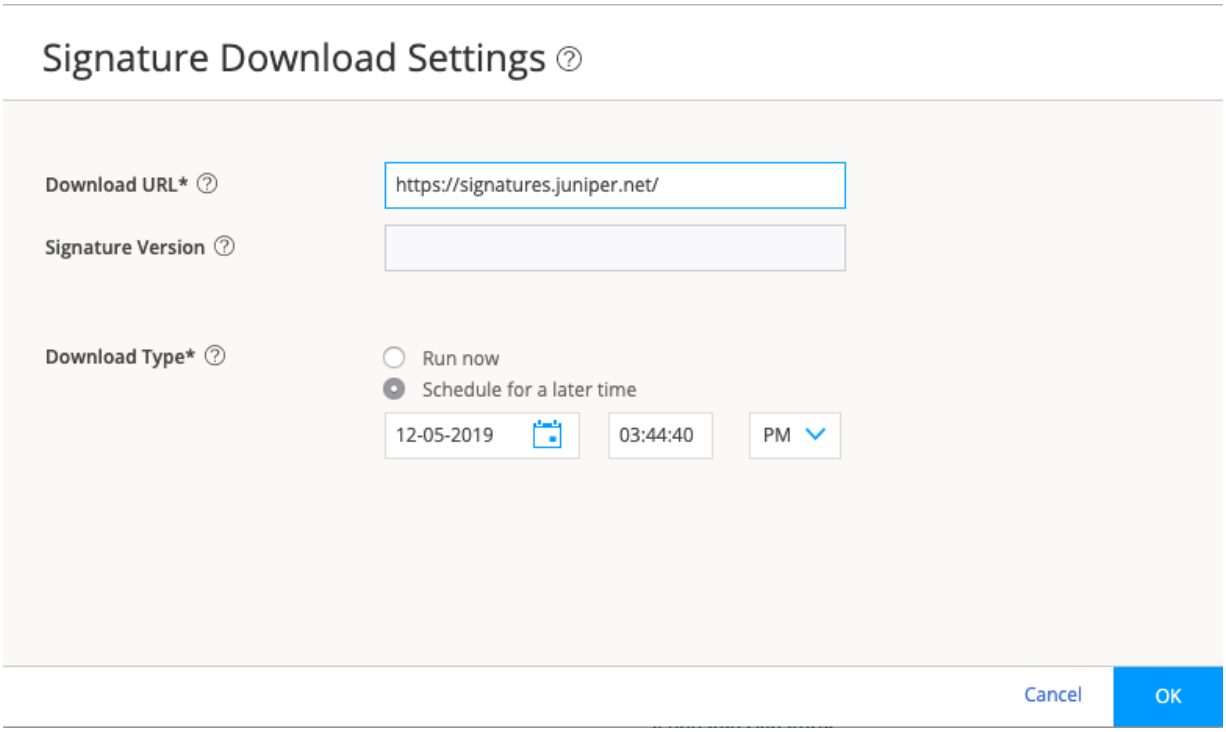

The **Download URL** field is pre-populated with Juniper's signature download URL. If you have previously downloaded the desired signature pack to another location (URL), enter that URL here.

4. If you are downloading the signature database from a location other than https://signatures.juniper.net, then you must enter the signature version that you want to download. If you are downloading from https://signatures.juniper,net, then you can leave this field empty.

5. Select whether to download the signatures now, or at a later time.

6. Click **OK**.

   A notification will appear at the top of the screen indicating whether the job has been scheduled or is running immediately.

Once the download completes successfully, the new database version number appears in the **Active Database** portion of the page. The new signature database is available to all of your tenants and their sites. To see the application signatures included in the database, navigate to **Configuration > Shared Objects > Application Signatures**.

Starting with CSO Release 5.0.2, you can define your own custom application signatures for use in SD-WAN policy. For more information regarding this optional step, see Contrail Service Orchestration Administration Portal User Guide.

Upload Licenses

The licenses that you upload to CSO using this procedure are available to be pushed to your tenant devices during the ZTP process or after they are provisioned. You need to install licenses on the hub and spoke SRX devices (physical and virtual) that you use in SD-WAN solutions. The licenses allow access to virtual network services such as application-based routing, application monitoring, and vSRX security features.

To upload the license for your devices:

1. Navigate to the **Administration > Licenses > Device Licenses** page.

   On this page is a list of all available device licenses. Since you have not installed any licenses yet, the list is empty.

2. Click the Add icon (**+**) button at the top-right part of the list to add a license.

   The **Add License** window appears.

3. Click the **Browse** button.

   This lets you locate the license file on your computer.

4. Select a tenant or *All Tenants* from the Tenant pull-down menu.

   This associates the license file with a particular tenant or all tenants. If the license is associated with a particular tenant, then the licenses can only be applied to devices that belong to that tenant.

5. (Optional) Enter a description of the license file if desired.

You can repeat this procedure to upload as many licenses as you have.

Add a New Tenant

In this section we use the Administration Portal to add a tenant to CSO. OpCo administrators in on-premises or cloud-hosted CSO installations and SP administrators in an on-premises CSO installation can add tenants.

1. Select **Tenants** from the left-navigation panel.

2. If there are no existing tenants, the **Add Tenant** button is displayed on the center of the page.
   Click the **Add Tenant** button, if it is available, to add a new tenant.
   If there are existing tenants, click the Add icon (**+**) at the upper right portion of the window to add a new tenant.

3. In the Add Tenant window that appears:
   - Enter a name for your tenant such as **Tenant1**.
   - Fill in the **Admin User** information.

     The e-mail address is used as the login name (username) for this user.
   - Select the checkbox next to the **Tenant AdminRole Name** in the **Available** list.
   - Click the **Right Arrow** button between the **Available** and **Selected** lists to move the **Tenant Admin** to the **Selected** list.
     If you want to see the access and permissions of the available roles, click the **Tenant Admin** or **Tenant Operator** names.
   - The password expiration default is 180 days.

     You can set any value between 1 and 365.
   - Click **Next**.

     The window advances to the **Deployment Info** tab.
   - In the **Deployment Info** window, select the **SD-WAN** card in the **Services** section.

     Depending on how your tenant was configured, you may see one or more of the following in addition to the SD-WAN card: Hybrid-WAN, Next Gen Firewall, and LAN. For this example, select only SD-WAN.

     This activates the **SD-WAN Mode** section of the window.
   - The **Realtime Optimized** radio button is selected by default with the SD-WAN service.

     You cannot change this selection.
   - Click **Next**.

     The window advances to the **Tenant Properties** tab. For this example, browse the tenant properties but do not make any changes.
     See the CSO Administration Portal User Guide for more information about the settings on the **Tenant Properties** tab.
   - Click **Next**.

     The window advances to the **Summary** section. Review the summary.
   - Click **OK**.

     A pop-up message appears that tells you that the Add Tenant job was started. After some time, your new tenant appears in the list of tenants.

Modify Device Templates

In this section, we examine device templates that we use for this example.

For the SRX in the Enterprise Hub Role

1. Navigate to **Resources > Device Templates**.
2. Find the device template named **SRX as SDWAN Hub**.
3. Select the checkbox next to that template.
4. Click the **Clone** button.
5. Enter a **Display Name** and a **Name** for the cloned template.

   CSO shows the **Display Name** in various workflow locations but uses the **Name** behind the scenes. For this example, we name the template SRX_as_SD-WAN_Hub.
6. Click **OK**.
7. Select the checkbox next to the cloned template and then select **Template Settings** from the **Edit Device Template** pull-down menu.

   A new window titled **Template Settings for *Display Name*** appears as shown in Figure 35 on page 71. You must scroll down to see these particular settings in the template.

Figure 35: Partial Template Settings for SRX as SD-WAN Hub

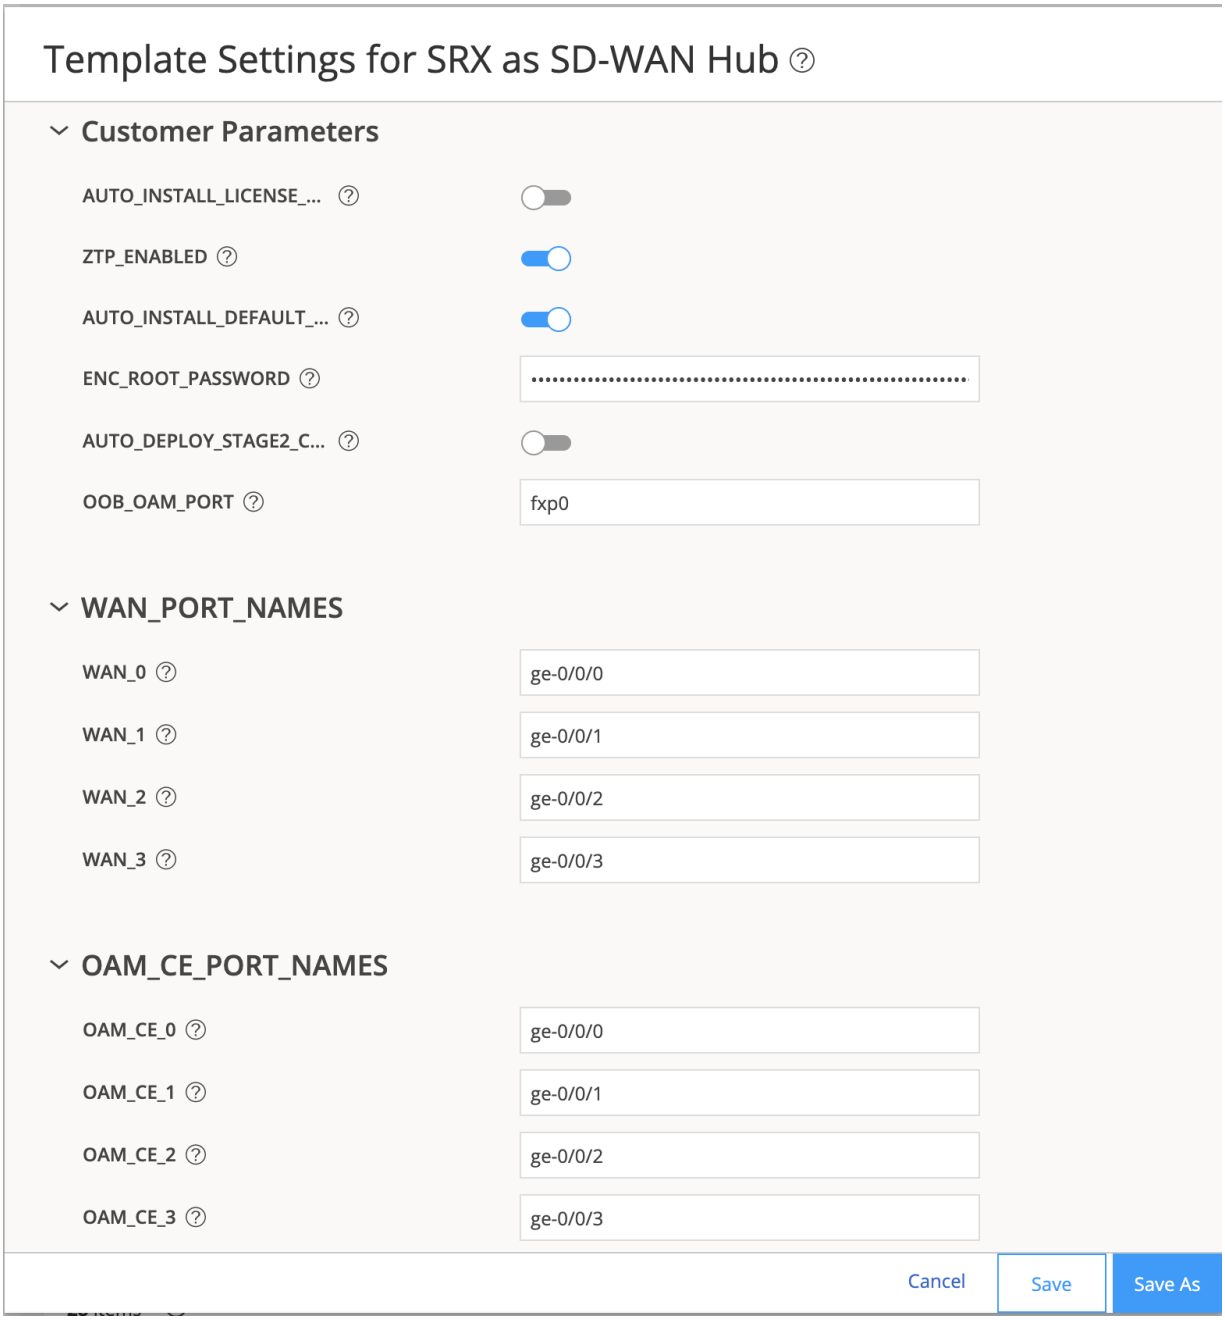

Note the device port names (**ge-0/0/0**, etc) for the WAN and OAM ports. If your hub device is not cabled to match, then adjust the port names in the template as needed.

8. Select **Save** when finished.

For the SRX in the CPE device role

9. Find the device template named **SRX as SD-WAN CPE** and select the checkbox next to its name.

10. Click the **Clone** button.

11. Enter a **Display Name** and a **Name** for the cloned template.

    The **Display Name** is whatCSO uses when selecting the template for use.

12. From the **Edit Device Template** pull-down menu, select **Template Settings**.

    The **Template Settings for <*Display Name*>** window appears as shown in Figure 36 on page 73. You must scroll down to see these particular settings in the template.

Figure 36: Partial Template Settings for SRX as SD-WAN CPE

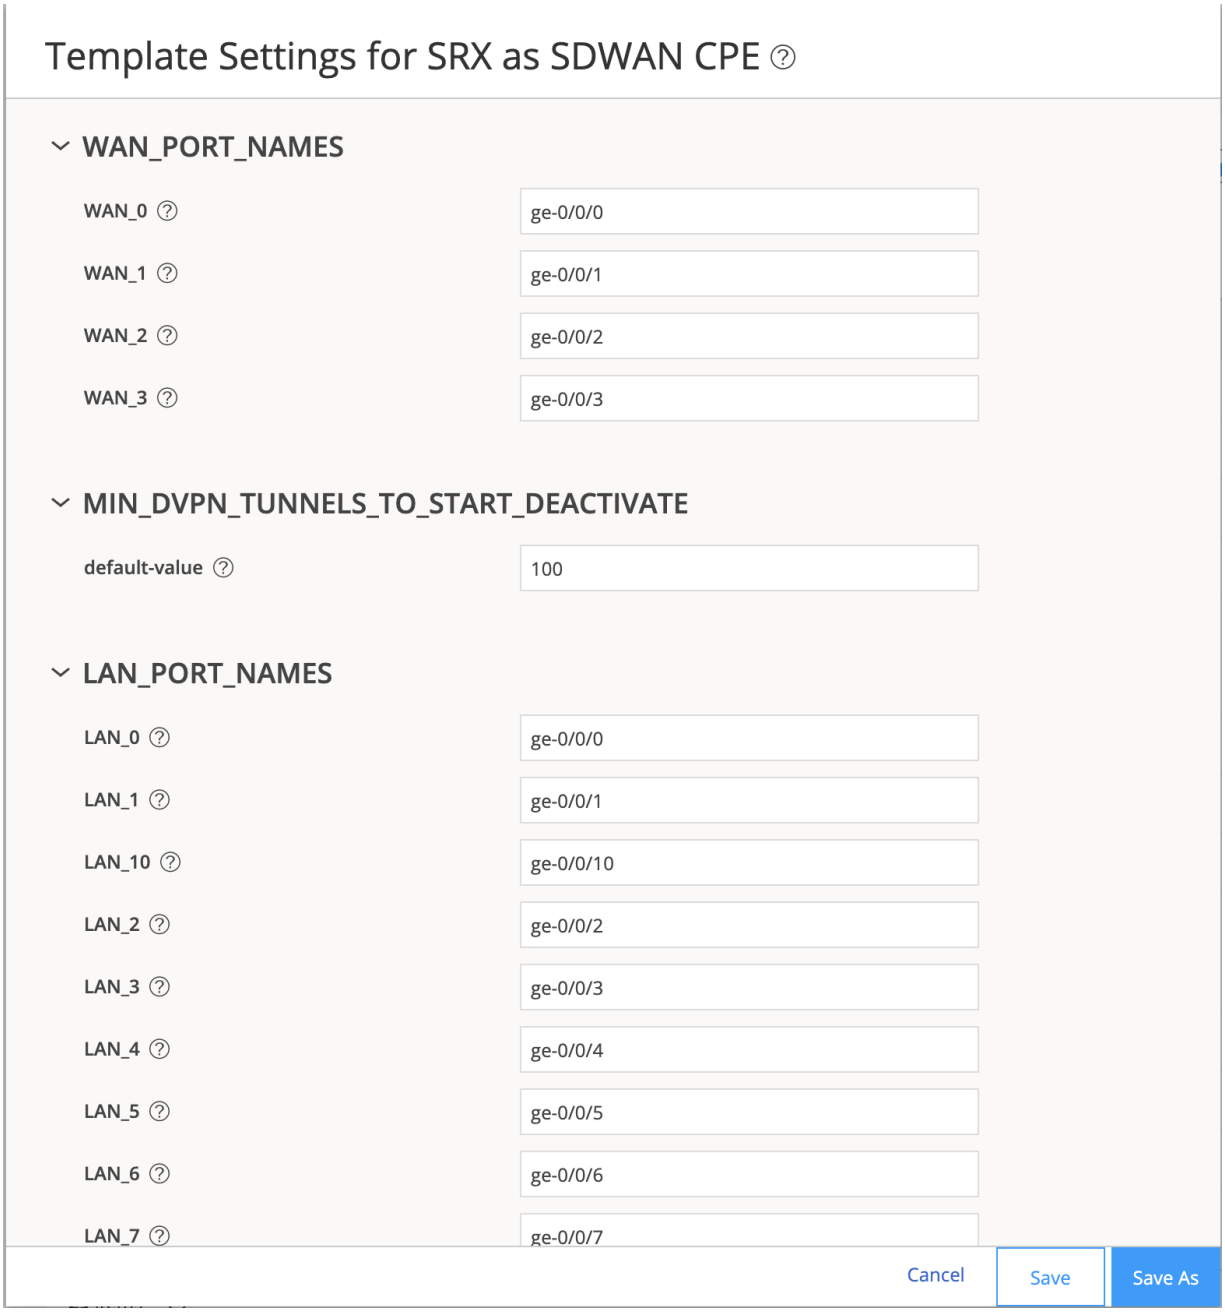

Note the device port names (**ge-0/0/0**, etc) for the WAN and OAM ports. If your CPE device is not cabled to match, then adjust the port names in the template as needed.

13. Click Save when finished.

The templates will be used later when you deploy the enterprise hub and CPE spoke devices.

Choose a Point of Presence (POP) for the Hub

A POP is a location within the service provider's cloud in which PE routers and IPSec Concentrators are located. It is a regionally located access point through which customer sites gain access to provider hub devices that are placed within. The hubs are either DATA_ONLY, OAM_ONLY, or OAM_AND_DATA hubs. SPs often place POPs in their network so that they are geographically close to customer sites.

> **NOTE:** The SP administrator is the only administrator with the privileges to create POPs. In a cloud-hosted CSO deployment, tenants choose the appropriate POP from a list of available POPs created by the SP administrator. In an on-premises CSO deployment, you (as the cspadmin user) create the POP in which the hub device resides.

To choose or add a POP (for cloud-hosted CSO):

1. Navigate to the **Resources > POPs** page.

   Here you can see a list of POPs available to you.

2. Make note of the POP name(s) and location(s) so that you can choose the appropriate one when adding your devices.

To add a POP (for on-premises CSO):

1. Navigate to the **Resources > POPs** page as the SP administrator.

   Here you can see a list of existing POPs.

2. Click on the Add icon (**+**) to add a POP and fill in the information in the **Add POP** window that appears. Currently, all POPs are regional.
   Give the POP a name and, optionally, address information so that its location can be displayed on CSO monitoring maps.

Add a Provider Hub Device to Your Tenant

A provider hub device resides in a regional POP within the service provider 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 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. As their names imply, they have different capabilities. At least one of the provider hubs in each tenant must have OAM capabilities. Adding multiple OAM-capable provider hubs helps to balance OAM traffic loads in large CSO deployments. In cloud-hosted versions of CSO, the OAM-capable hubs are clearly labeled.

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

The following two procedures describe how to add provider hub devices to your CSO installation and tenant. The first procedure describes adding a provider hub device to an on-premises version of CSO. It can only be done by an SP administrator. In cloud-hosted CSO, the addition of the hub devices to the system is carried out by Juniper Networks.
The second procedure is carried out at the OpCo or tenant level in both on-premises and cloud-hosted CSO versions. It makes the provider hubs added in the first procedure available for use by tenants.

Add Provider Hubs for On-Premises CSO

1. Navigate to the **Resources > Provider Hub Devices** page as the SP administrator user (cspadmin or equivalent).

2. At the top-right part of the page, click the Add icon (**+**).

   A new window appears titled **Add Provider Hub Device**.

3. Fill in the **Site Information** section as follows:
   - **Name**: Name the provider hub something that makes sense, like **PH-OAM-DATA-1**.
   - **Management Region**: **Regional**

     There is currently no other option for this.
   - **Site Capability**: **DATA_AND_OAM**

     This allows both operation, administration, and maintenance (OAM) and user data to traverse this device. It ensures that CSO can manage CPE devices through this provider hub device.
   - **POP**: **Select the POP that you just created from the pull-down menu**.
   - **Authentication Type**: **Pre Shared Key**

     You can choose Public Key Infrastructure if you have the proper certificates set up.
     CSO supports single and multi-level PKI certificates.
   - (Optional) **Advanced Configuration**: Change the information in this section as appropriate for your network.
4. Click **Next**.

The window advances to the **WAN** tab.

5. In the **Device Template Section**, select **SRX** as the **Device Series** from the pull-down menu.

6. Select **SRX_as_SDWAN_Hub** from the carousel of available device templates (cards).

7. In the **Device Information** section, enter the device serial number.

8. Leave the **Auto Activate** button active (blue).

9. (Optional) If you want to upgrade the device image for your SRX Series device, select the new boot image from the list. The boot image is the device image that was previously uploaded to the image management system in CSO. The boot image is used to upgrade the device during the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure and uses the image that is present on the device.

10. In the **Management Connectivity** section fill in the form as follows:

    - Leave the **Loopback IP Prefix** blank.
      CSO automatically configures the proper loopback IP during the ZTP process, based on information contained in the device template and CSO databases.
    - OAM Interface: Enter the appropriate interface, such as **ge-0/0/0** as the **OAM Interface** of the provider hub.

      The interface selected must match the device template and your network cabling.
    - OAM VLAN: Leave this field blank.

      > **NOTE:** You can enter a VLAN ID if one is needed in your network. If you specify an OAM VLAN ID, then all in-band OAM traffic reaches the site through the selected OAM interface. The range is 0 through 65535.

    - OAM IP Prefix: Enter an IP address prefix, such as **10.100.100.11/32**.

      This is the IP address prefix for the OAM network. Secure OAM traffic is passed across this network in IPsec tunnels. The OAM IP Prefix must be unique across the entire management network.

      > **NOTE:** For SRX Series services gateways like we are using in this example, always use a **/32** prefix.

    - OAM Gateway: Enter an IP address, such as **10.100.100.1**.

This is the IP address of the next-hop on the management network through which CSO connectivity is established.

- EBGP Peer-AS: Leave this field blank.

  This is the external BGP peer autonomous system number. It is used to peer with a PE router in the SP network (if any).
  Enter a value here if needed in your network.

11. In the **WAN Links** section, fill in the information as follows:

    - Leave the **WAN_0 (ge-0/0/0)** slider button enabled.

      The physical device interface is already chosen from the value in the device template and cannot be altered here.

    - Link Type: Select **MPLS**.
    - Address Assignment: Enter **Static**.
    - Static IP Prefix: Enter an IP address prefix, such as **172.21.22.2/29**.

      This represents the provider hub address of the hub-to-CPE network connection.

    - Gateway IP Address: Enter an IP address, such as **172.21.22.1**.

      This is the IP address of the spoke (SRX or NFX CPE device) at the customer site.

    - Enable the **WAN_1 (ge-0/0/1)** slider button.

      The physical device interface is already chosen from the value in the device template and cannot be changed here.

    - Link Type: Select **Internet**.
    - Address Assignment: Enter **Static**.
    - Static IP Prefix: Enter an IP address prefix, such as **192.0.2.2/29**.

      This represents the provider hub address of the hub-to-CPE network connection.

    - Gateway IP Address: Enter an IP address, such as **192.0.2.1**.

      This is the IP address of the spoke (SRX or NFX CPE device) at the customer site.

    > **NOTE:** Enable the other WAN interfaces for your provider hub device as appropriate.

12. Click **OK** when you're finished.

    The **Activate Device** window pops up.
    The device shows up in the in the **Provisioned** state when this window shows the operation completed successfully.

You can dismiss this window by clicking **OK** before the operation is complete. To track the progress, navigate to **Monitor > Jobs** and click on the job name.

Add a Provider Hub to Your Tenant

1. Navigate to the **Resources > Provider Hub Devices** page.

   Here you can see a list of all cloud hub devices, their assigned POP, site associations, status, model, serial number, and OS version.

2. Make note of the names of the Provider Hub devices available to you.

3. Navigate to the **Resources > Site Management** page.

4. From the **Add** menu, select **Add Provider Hub**.

   The **Add Provider Hub for *Tenant Name*** window appears.

5. Select a **Service POP** from the pull-down menu.

6. Select a **Hub Device Name** from the pull-down menu.

   As mentioned previously, the you must add at least one provider hub with OAM capabilities.
   You can repeat this process to add as many provider hubs as you want from this POP to your tenant.

7. Click **OK**.

   The provider hub device is added to the list.

Add an Enterprise Hub to Your Tenant

Unlike a provider hub which is shared amongst multiple tenants, an enterprise hub acts as the primary hub device for spoke sites belonging to a single tenant. Tenants that have an enterprise hub installed can use it for site-to-site VPNs between spoke sites. In this case, the provider hub becomes a backup hub for the same VPNs. The site-to-site VPNs initially created through the enterprise hub can be dynamically switched to direct site-to-site VPNs based on the (user configurable) dynamic VPN threshold settings for the tenant.

An enterprise hub is added to a tenant by a tenant administrator using the Customer Portal in CSO. To add an enterprise hub to your new tenant:

1. Enter the Customer Portal for your tenant.

   SP and OpCo administrators access the Customer Portal by navigating to **Tenants** and clicking the tenant name from the list. This puts these administrators into the tenant administrator role for that tenant.
   Tenant administrators are automatically placed in their Customer Portal upon successful login.

2. Navigate to **Resources > Site Management**.

   The **Sites** page appears.

3. From the **Add** pull-down menu, select **Enterprise Hub**.

   The **Add Enterprise Hub for *Site Name*** window appears.

4. In the **Site Information** section, enter a name that makes sense for your site.

   Choose site names carefully because they cannot be changed after the site is added.

5. In the **Site Capabilities** section, select the **SD-WAN** card from the **WAN Capabilities** area.

6. In the configuration section, the **Primary Provider Hub** pull-down menu should already be populated with the name of the provider hub added earlier.

   The **On-demand VPN Threshold**, **Address and Contact Information**, and **Advanced Configuration** sections are all optional.

   > **NOTE:** Most settings made while creating sites cannot be changed once the site is provisioned. The exceptions are: **Address and Contact Information** and the NTP settings available in the **Advanced Configuration Settings** section.

7. Click **Next**.

   The page advances to the **WAN** tab.

8. Click the left arrow (**<**) or right arrow (**>**) until you see the **SRX as SD-WAN CPE** card. Click on that card.

9. In the **Device Information** section, enter the device serial number.

10. Leave the **Auto Activate** button active (blue).

11. (Optional) If you want to upgrade the device image for your SRX Series device, select the new boot image from the list. The boot image is the device image that was previously uploaded to the image management system in CSO. The boot image is used to upgrade the device during the ZTP process. If the boot image is not provided, then the device skips the automatic upgrade procedure and uses the image that is present on the device.

12. In the **Management Connectivity** section fill in the form as follows:

    - Leave the **Loopback IP Prefix** blank.
      CSO automatically configures the proper loopback IP during the ZTP process, based on information contained in the device template and CSO databases.
    - OAM Interface: Enter the appropriate interface, such as **ge-0/0/0** as the **OAM Interface** of the provider hub.

      The interface selected must match the device template and your network cabling.
    - OAM VLAN: Leave this field blank.

      > **NOTE:** You can enter a VLAN ID if one is needed in your network. If you specify an OAM VLAN ID, then all in-band OAM traffic reaches the site through the selected OAM interface. The range is 0 through 65535.

    - OAM IP Prefix: Enter an IP address prefix, such as **10.100.100.11/32**.

      This is the IP address prefix for the OAM network. Secure OAM traffic is passed across this network in IPsec tunnels. The OAM IP Prefix must be unique across the entire management network.

      > **NOTE:** For SRX Series services gateways like we are using in this example, always use a **/32** prefix.

    - OAM Gateway: Enter an IP address, such as **10.100.100.1**.

      This is the IP address of the next-hop on the management network through which CSO connectivity is established.
    - EBGP Peer-AS: Leave this field blank.

      This is the external BGP peer autonomous system number. It is used to peer with a PE router in the SP network (if any).
      Enter a value here if needed in your network.

13. In the **WAN Links** section, fill in the information as follows:

- Leave the **WAN_0 (ge-0/0/0)** slider button enabled.

  The physical device interface is already chosen from the value in the device template and cannot be changed here.
- Link Type: Select **MPLS**.
- Address Assignment: Enter **Static**.
- Static IP Prefix: Enter an IP address prefix.

  This represents the hub-side address of the hub-to-CPE network connection.
- Gateway IP Address: Enter an IP address.

  This is the IP address of the spoke (SRX or NFX CPE device) at the customer site.
- Enable the **WAN_1 (ge-0/0/1)** slider button.

  The physical device interface is already chosen from the value in the device template and cannot be changed here.
- Link Type: Select **Internet**.
- Address Assignment: Enter **Static**.
- Static IP Prefix: Enter an IP address prefix.

  This represents the hub-side address of the hub-to-CPE network connection.
- Gateway IP Address: Enter an IP address.

  This is the IP address of the spoke (SRX or NFX CPE device) at the customer site.

> **NOTE:** Enable the other WAN interfaces for your provider hub device as appropriate.

14. Expand the **Advanced Settings** section by clicking on the right arrow **>** icon.

15. Enable the **Use for Full Mesh** slider button (set to blue).

16. Select the **Internet** mesh tag from the **Mesh Tag** pull-down menu.

17. Ensure that the proper **Overlay Peer Interface** is selected.

    This is the interface on the provider hub that this enterprise hub will use as a BGP peer interface.

18. Click **Next**.

    The page advances to the **LAN** tab.

19. Click the **Add LAN Segment** button.

The **Add LAN Segment** window appears.

20. Enter a **Name** for the LAN segment.

21. Ensure that the **Department** pull-down menu has **Default** selected.

22. Enter a valid **Gateway Address/Mask**.

    This value is used as the gateway address for devices deployed on this LAN segment.

23. Select and move the appropriate port(s) from the **Available** list to the **Selected** list.

    Click the blue arrow to the right of a port name to move it from one list to the other.

24. Click **Save**.

    The new LAN segment is listed in the **LAN** tab of the **Add Enterprise Hub for *Site Name*** window.

25. Click **Next**.

    The page advances to the **Summary** tab.

26. Click **OK** when you're finished reviewing the summary tab information.

The **Activate Device** window pops up.
The device shows up in the in the **Provisioned** state when this window shows the operation completed successfully.

You can dismiss this window by clicking **OK** before the operation is complete. To track the progress, navigate to **Monitor > Jobs** and click on the job name.

Add an On-Premises Spoke for the Tenant

In this section, we continue in the Customer Portal for the newly configured tenant to create an on-premises spoke with an SRX CPE device.

This procedure begins in the **Tenants** window of the Administration Portal at the list of tenants.

1. Click on the name of the tenant that you created.

   This will take you to the Customer Portal for that tenant.

2. Navigate to the **Resources > Site Management** page.

3. In the **Site Management** window that appears, select **Add On-Premise Spoke Site (Manual)** from the **Add** pull-down menu.

   The **Add On-Premise Spoke Site for *Tenant*** page appears.

4. In the **Site Information** section, enter a name that makes sense for your site

   Choose site names carefully because they cannot be changed after the site is added.

5. In the **Site Capabilities** section, select the type of WAN and LAN capabilities you want for this site.

   The available site capabilities are based on the tenant capabilities defined during tenant creation. You can choose one WAN capability in addition to one optional LAN capability.

   For this example, choose only **SD-WAN**.

6. In the **Configuration** section, the **Provider Hub** and **Enterprise Hub** pull-down menus are already populated with the previously added hub devices.

   The **On-demand VPN Threshold**, **Address and Contact Information**, and **Advanced Configuration Settings** sections are all optional.

   > **NOTE:** Most settings made while creating sites cannot be changed once the site is provisioned. The exceptions are: **Address and Contact Information** and the NTP settings available in the **Advanced Configuration** section.

7. Click **Next**.

   The page advances to the **WAN** tab.

8. Next to **Device Series**, select **SRX** from the pull-down menu.

   A horizontal list of device template cards applicable to SRX Series devices is shown.

9. Click the left arrow (**<**) or right arrow (**>**) until you see the **SRX as SD-WAN CPE** card. Click on that card.

10. In the **Device Information** section, enter the device serial number.

11. Leave the **Auto Activate** button active (blue).

12. (Optional) If you want to upgrade the device image for your SRX Series device, select the new boot image from the list. The boot image is the device image that was previously uploaded to the image management system in CSO. The boot image is used to upgrade the device during the ZTP process. If

the boot image is not provided, then the device skips the automatic upgrade procedure and uses the image that is present on the device.

13. In the **WAN Links** section, fill in the information as follows:

    - Leave the **WAN_0 (ge-0/0/0)** slider button enabled.

      The physical device interface is already chosen from the value in the device template and cannot be changed here.

    - Link Type: Select **MPLS**.
    - Address Assignment: Enter **Static**.
    - Static IP Prefix: Enter an IP address prefix.

      This represents the hub-side address of the hub-to-CPE network connection.

    - Gateway IP Address: Enter an IP address.

      This is the IP address of the spoke (SRX or NFX CPE device) at the customer site.

    - Enable the **WAN_1 (ge-0/0/1)** slider button.

      The physical device interface is already chosen from the value in the device template and cannot be changed here.

    - Link Type: Select **MPLS**.
    - Address Assignment: Enter **Static**.
    - Static IP Prefix: Enter an IP address prefix.

      This represents the hub-side address of the hub-to-CPE network connection.

    - Gateway IP Address: Enter an IP address.

      This is the IP address of the spoke (SRX or NFX CPE device) at the customer site.

    > **NOTE:** Enable the other WAN interfaces for your provider hub device as appropriate.

14. Expand the **Advanced Settings** section by clicking on the right arrow **>** icon.

15. Enable the **Use for Full Mesh** slider button (set to blue).

16. Select the **MPLS** mesh tag from the **Mesh Tag** pull-down menu.

17. Enable the **Use for OAM Traffic** slider button (set to blue).

    Figure 37 on page 85 below shows an example of the settings described above.

Figure 37: WAN_0 Configuration Example

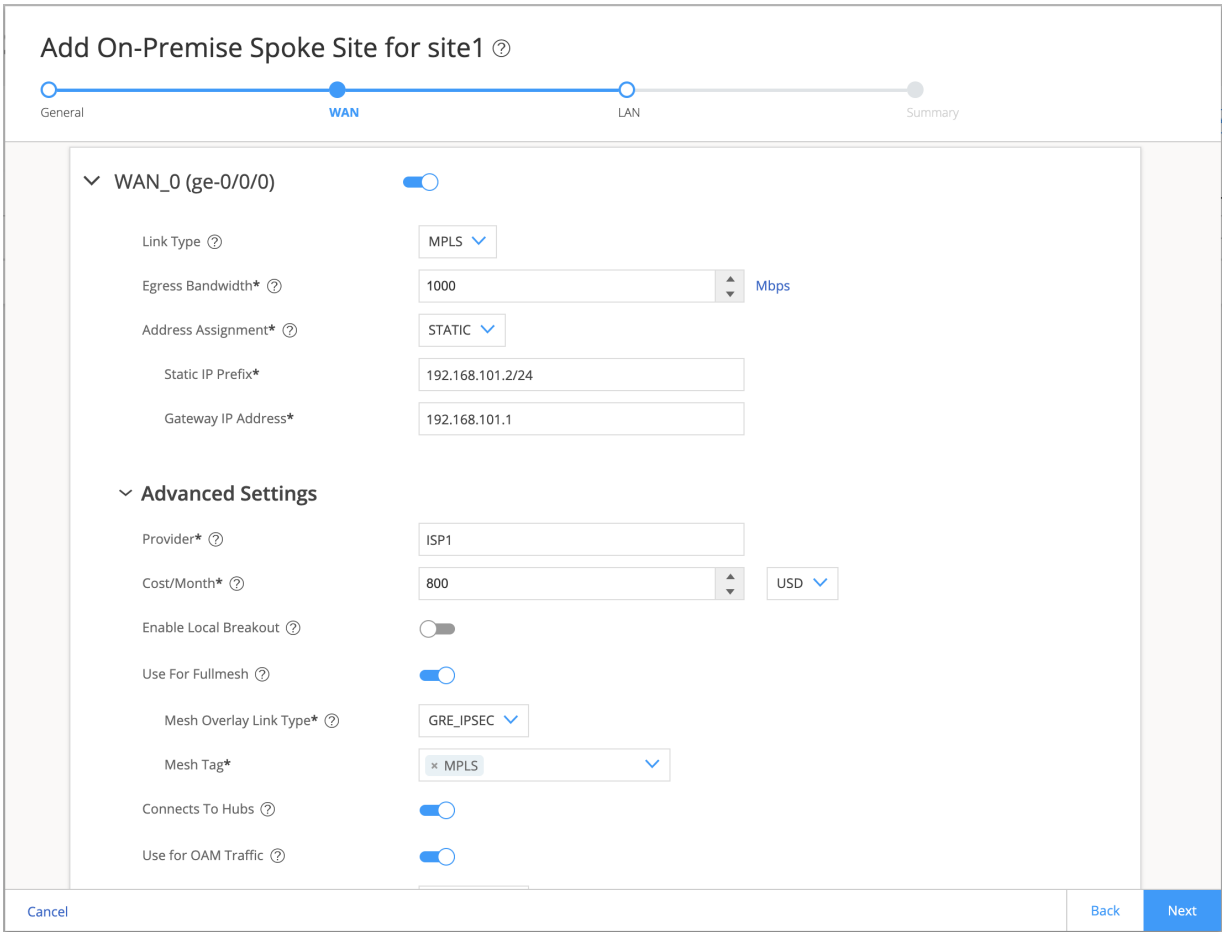

18. Select the **Enable** button next to **Wan_1**.

    The physical device interface is already chosen from the value in the device template and cannot be changed here.

    - Link Type: Select **Internet**.
    - Address Assignment: Enter **Static**.
    - Static IP Prefix: Enter an IP address prefix.

      This represents the hub-side address of the hub-to-CPE network connection.

      NOTE: Enable the other WAN interfaces for your provider hub device as appropriate.

19. Expand the **Advanced Settings** section by clicking on the right arrow **>** icon.

20. Enable the **Enable Local Breakout** button (set to blue).

Leave the **Breakout Options** pull-down menu set to **Use for breakout and WAN traffic**.

21. Enable the **Autocreate Source NAT Rule** button (set to blue).

    Leave the **Translation** pull-down menu set to **Interface**.

22. Enable the **Use for Full Mesh** slider button (set to blue).

23. Select the **Internet** mesh tag from the **Mesh Tag** pull-down menu.

24. Enable the **Use For OAM Traffic** slider button (set to blue).

25. The **Overlay Peer Device** is automatically set to the provider hub device.

    Ensure that the **Overlay Peer Interface** pull-down menu is set to the proper interface.

    Figure 38 on page 86 below shows an example of the WAN_1 configuration as described above.

Figure 38: WAN_1 Configuration Example

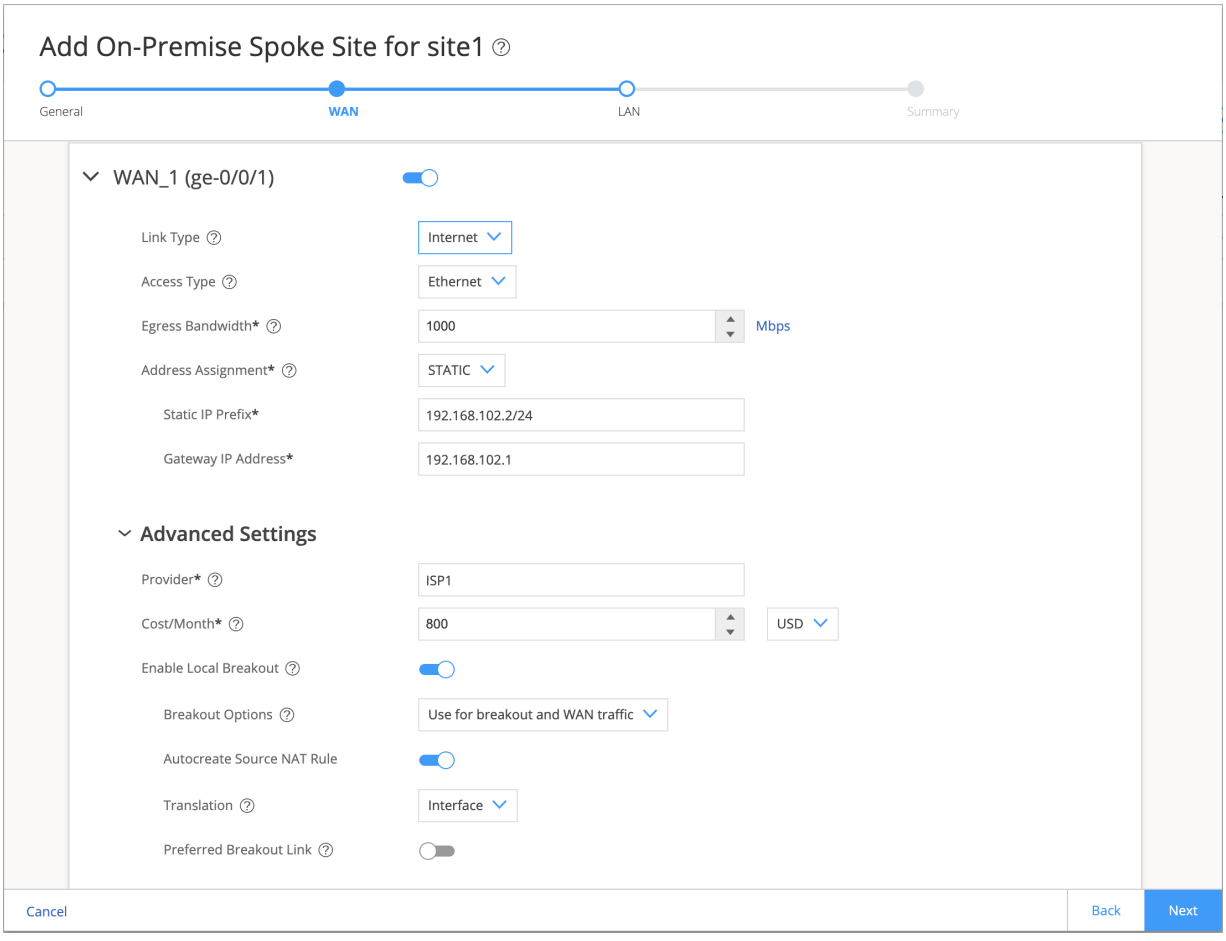

26. Click **Next** when finished.

    The window advances to the **LAN** section.

27. Click the **Add LAN Segment** button.

    A new window appears titled **Add LAN Segment**.

    Fill in the following information in this window:

    - Name: **LAN2**

      > **NOTE:** This can be any name that makes sense in your deployment.

    - VLAN ID: Leave this field blank.

      > **NOTE:** Enter a VLAN ID if required at the remote site.

    - Department: Leave this field as **Default**.

      In CSO, spoke site departments equate to security zones on the CPE device. In this example, the Default security zone will be used later when we create security policies. Creating multiple departments for the spoke site creates multiple security zones with the same names on the CPE device.

      If you have departments set up already and the proper department is not shown, you can create one by clicking on the **Create Department** link.

    - Gateway Address/Mask: Enter an IP address and mask.

      Specify a unique and valid IPv4 address with subnet mask. This address is the default gateway for endpoints in this LAN segment.

    - DHCP: **Off**

      The An SRX Series device can provide DHCP server services for the remote LAN. For this example, leave DHCP set to off.

    - CPE Ports: Select **LAN_2 (ge-0/0/2)** by clicking the checkbox next to it.

    - Click the right arrow **>** icon to move **LAN_2 (ge-0/0/2)** from the available list to the selected list.

28. Click **Save** when finished.

    The **Add LAN Segment** window closes.

29. Click **Next**.

    The window advances to the **Summary** section.

30. Review the **Summary** section.

31. Click **OK** when you're finished reviewing this section.

    A device activation window pops up and displays the progress of your site deployment.

Install a License on a Device

To install a license on a device, use the Administration Portal.

1. Navigate to **Administration > Licenses > Device Licenses**.

   In the pop-up window that appears,

2. Click the checkbox next to the license file that you uploaded in Step 3.

3. Click the **Push License** button at the upper-right part of the list and select **Push**.

   The **Push License** window appears.

4. Select the name of the tenant that you created previously from the **Tenant** pull-down menu.

   Your sites and devices appear under **Sites and Devices**.

5. Select the checkbox next to your tenant site to push the license to the CPE device at that site.

Install an Application Signature on a Device

This step allows the CPE device to obtain the signature database needed for application identification.

To install an application signature:

1. Navigate to **Adminstration > Signature Database**.

   From the signature download you completed previously, you can now see the **Active Database** section has the number of the downloaded database listed.

2. Click the **Install on Device** link under the **Actions** column.

   In the new window that appears, you can elect to push the signatures to any device listed.

3. Select the checkbox next to the NFX250 device.

4. Click **OK**.

Add Firewall and NAT Policies to the Topology

In this section, we use the Customer Portal to add and deploy an intent-based firewall policy that allows the CPE-side LAN segments to pass traffic between each other and to the Internet. This requires adding two intents to the policy. One for department–to–department, and one for department–to–any address.

1. In the Customer Portal for your tenant, navigate to **Configuration > Firewall > Firewall Policy**.

   This brings up the **Firewall Policy** page with a list of existing policies.

2. Click the Add icon (**+**) to add a new policy.

   The **Add Firewall Policy** page appears.

3. Give the policy a name that makes sense, like FirstFirewallPolicy.

4. Click the checkbox to enable the policy for all sites.

   The new policy is added to the list.

5. Click on the policy name in the list to bring up the **Intents** for that policy.

   It shows that there are no **Enterprise Intenets** or **Zone-based Intents** for this policy.

   Allow Site-to-Site Traffic

6. Click the Add icon (**+**) to add a new intent.

   The window changes to reveal the intent editor page.

7. Click the Add icon (**+**) in the **Select Source** field.

   A list of possible sources appears.

8. Select **Default** from the **Departments [DEPT]** section of the list.

   Your choice is added to the source section.

9. Click the Add icon (**+**) labeled **Select Action**.

   A list containing available actions is shown. Select **Allow** from that list.

10. Your action choice is added between the source and destination sections.

11. Click the Add icon (**+**) in the **Destination** section.

    A list of possible destinations appears.

12. Select **Default** from the **Departments [DEPT]** section of the list.

    Your choice is added to the destination section.

13. (Optional) Enable the **Logging** slider switch (turns blue).

    We recommend that you log all deny or drop actions within firewall intents. Turning logging on for an accept action creates a lot of logs.

14. Click the **Save** button.

    The new intent is shown under the policy.

    Allow Outgoing Site Traffic

15. Click the Add icon (**+**) to add a second intent to the policy.

    The window changes to reveal the intent editor page.

16. Click the Add icon (**+**) in the **Select Source** field.

    A list of possible sources appears.

17. Select **Default** from the **Departments [DEPT]** section of the list.

    Your choice is added to the source section.

18. Click the Add icon (**+**) labeled **Select Action**.

    A list containing available actions is shown. Select **Allow** from that list.

19. Your action choice is added between the source and destination sections.

20. Click the Add icon (**+**) in the **Destination** section.

    A list of possible destinations appears.

21. Select **Any** from the **Addresses [ADDR]** section of the list.

    Your choice is added to the destination section.

22. (Optional) Enable the **Logging** slider switch (turns blue).

It is recommended to log all deny or drop actions within firewall intents. Turning logging on for an accept action creates a lot of logs.

23. Click the **Save** button.

24. Click the **Deploy** button.

    This brings up a **Deploy** window. Here you can select to run the policy deployment now or schedule it to run later.

25. Click **Deploy**.

    Deployment progress bars appear as CSO deploys the policy.

Add SD-WAN SLA-Based Steering Profiles and Policy

In this section, we use the Customer Portal to select a path-based steering profile and apply it to the SD-WAN Policy to specify that You Tube traffic should pass over the WAN_1 overlay link rather than the default link, WAN_0.

1. Navigate to **Configuration > SD-WAN > Path-Based Steering Profiles**.

2. Click the Add icon (**+**) to create a new profile.

   This brings up a **Add PathProfile** window.

   In the new window, fill in the following information.

   - Name: Enter a name for the profile, such as **Internet-Path**.
   - Traffic Type Profile: Select **INTERNET**.
   - Path Preference: Enter **Internet**.

     Priority value 1 is the highest priority. Higher priority profiles (lower numbers) take precedence over lower priority ones during SD-WAN events.

3. Click **OK**.

   The window closes and the new policy appears in the list.

4. Navigate to **Configuration > SD-WAN > SD-WAN Policy**.

   This brings up the SD-WAN policy page which includes a list of all SD-WAN policies.

5. Click the Add icon (**+**) at the upper right part of the list to create a new policy.

The policy builder screen appears with the **Source** section activated.

The default value for Source is **All Sites**.

The default value for Application is **Any**.

Use the default values for these fields.

6. Click the **+ Select Destination** field.

7. Type **YouTube** at the text-insertion point.

   This brings up a list of available applications.

8. Select **YouTube** from the list.

9. Click **+ Select Profile**.

   This brings up a list of available profiles.

10. Select **Internet-Path** from the **Path-Based Profiles [SLA]** section of the list.

11. Click **Save**.

    This closes the builder window and shows the list of SD-WAN Policies.

12. Click the **Deploy** button.

    This brings up a **Deploy** window. Here you can select to run the policy deployment now or schedule it to run later.

13. Click **Deploy**.

    Deployment progress bars appear as CSO deploys the policy. When it finishes, the **Total Intents** count increases from 0 to 1.

4

CHAPTER

SD-LAN Deployment

SD-LAN Deployment | **94**

SD-LAN Deployment

IN THIS SECTION

- SD-LAN Deployment Overview | 94
- SD-LAN Deployment | 98

SD-LAN Deployment Overview

The SD-LAN deployment focuses on branch site LAN connectivity using specific EX Series switches and Virtual Chassis. Once deployed, you can manage the connected branch site LANs through the EX switch. You can also manage many aspects of the EX switch or Virtual Chassis itself.

There are several options for deploying an SD-LAN solution:

- Behind an Internet Gateway device as a standalone LAN switch in a new SD-LAN deployment as shown in Figure 39 on page 95.

Figure 39: SD-LAN Behind an Internet Gateway Device

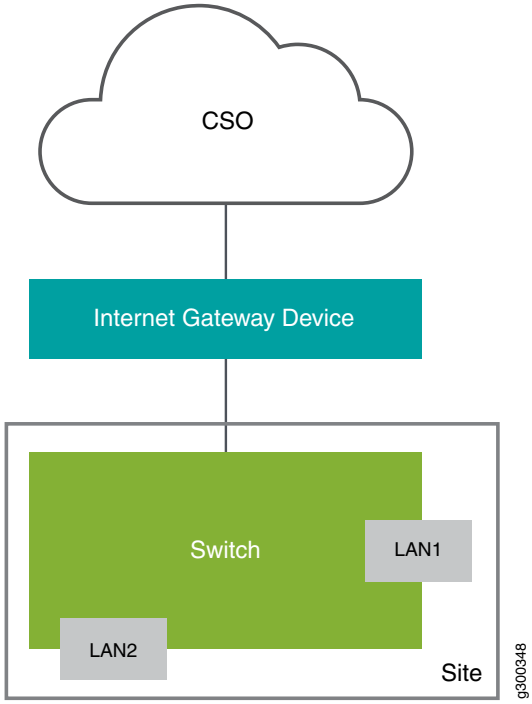

- Behind a CSO-managed CPE device as part of either a new SD-WAN deployment or an extension of an existing SD-WAN deployment, as shown in Figure 40 on page 96.

Figure 40: SD-LAN Behind a CPE Device

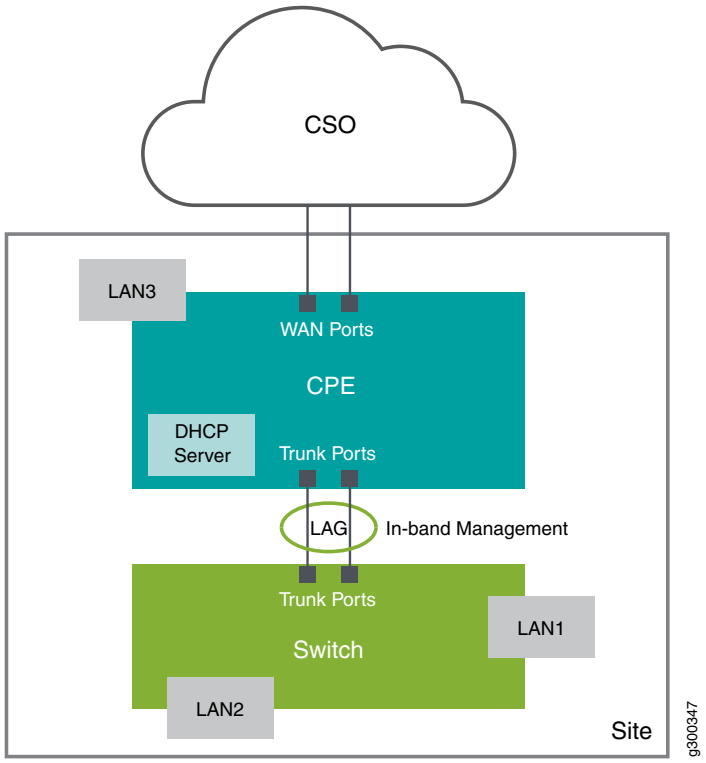

NOTE: You cannot deploy an EX Series LAN switch behind an NFX150 Series CPE device.

- Behind a CSO-managed NGFW device as part of either a new NGFW deployment or as an extension of an existing NGFW deployment, as shown in Figure 41 on page 97.

Figure 41: SD-LAN Behind NGFW

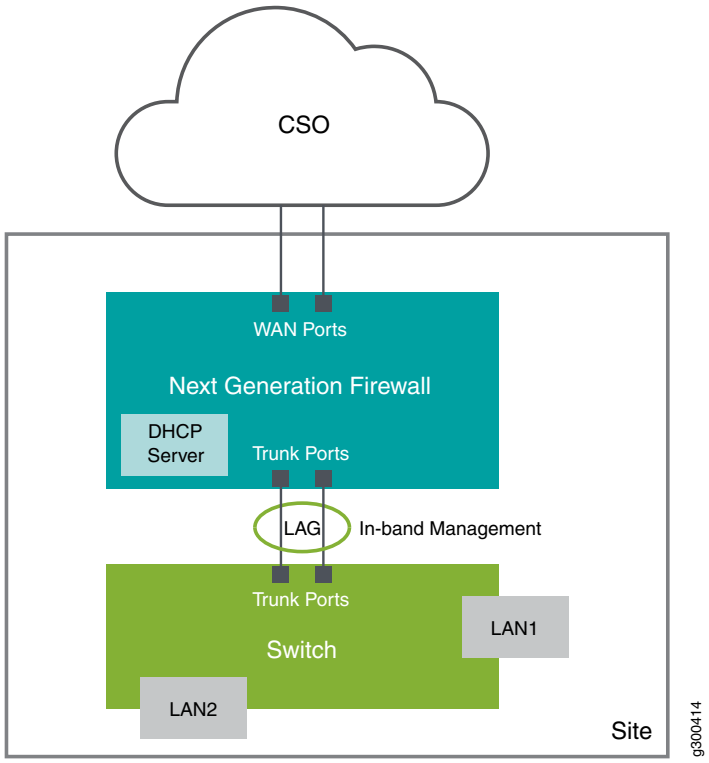

It is important to note the Internet Gateway Device and the CPE device shown in Figure 39 on page 95 and Figure 40 on page 96. The LAN switch deployed at a branch site must be deployed behind an Internet gateway device that is capable of routing traffic to CSO.

An SD-LAN 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 LAN service available. This service is included in the tenant configuration by the tenant administrator during tenant onboarding. The remainder of this document provides the steps that you need to perform in order to complete an SD-LAN deployment in CSO.

Table 10 on page 98 shows the switching and WiFi platforms currently supported for SD-LAN.

Table 10: Hardware and Software Matrix for Devices in an SD-LAN Deployment

Role	Platform	Models Supported	Software Release Versions
SD-LAN Devices	EX Series Switches	• EX2300 • EX3400 • EX4300 • EX4600 • EX4650	Junos OS 18.4R2 Junos OS 18.4R3
	Mist Access Points	• AP41 • AP43 • AP61	Mist AP Firmware 0.3.x and later

NOTE: For the most up to date information on hardware and software support for CSO, see the Contrail Service Orchestration Release Notes.

SD-LAN Deployment

The procedure you follow to complete this task varies slightly depending on whether you are in the role of a CSO tenant administrator or OpCo administrator. A note is used where needed to account for these variances.

This procedure makes the following assumptions:

- You have already established your login credentials for CSO.
- The tenant for which you are creating the LAN site is called **ExampleCo**, and has already been created.
- The **ExampleCo** tenant was added with LAN service capabilities.

If any of these things are not true, see *Accessing Administration Portal*, *Accessing Customer Portal*, or *Adding a Single Tenant* as needed.

This example demonstrates deploying a standalone SD-LAN behind an Internet gateway device.

The steps to deploy an SD-LAN site are as follows:

1. Log in to CSO using your login credentials.

   > **NOTE:** If you are an OpCo administrator, navigate to **Tenants** in the left-nav bar and select **ExampleCo** from the list of tenants on the tenants page. If you are the tenant administrator, you will be placed in the Customer Portal for **ExampleCo** upon successful login.

2. In the Customer Portal for **ExampleCo**, Navigate to **Resources > Site Management**.

   The **Sites** page appears.

3. Click the **Add** button and select **Add On-Premise Spoke (Manual)** from the list of options.

   The **Add On-Premise Spoke Site for ExampleCo** page appears.

4. In the **Site Information** section, give the site a name such as **LAN-Site1**.

5. In the **Site Capabilities** section, click the **LAN** icon.

   Depending on the configuration of the **ExampleCo** tenant, there may be other icons available. Only select **LAN** for this example.

6. Click the right arrow icon **>** next to **Address and Contact Information** to expand this section.

   None of the fields are required, but adding address information for the site allows CSO to place an icon for the site on maps on the monitoring page and show how it is linked to CSO.

7. Click the right arrow icon **>** next to **Advanced Configuration**.

   The two required fields, **Name Server IP List** and **NTP Server** are both pre-populated for you. Make changes as needed for your network to any of the fields.

8. Click **Next**.

   The wizard skips past the **WAN** page to the **LAN** page.

9. In the **Device Profile** section, fill in the **Device Name**.

10. Select the appropriate **Device Type** from the pull-down menu.

11. (Optional) Select the appropriate **Device Model** from the pull-down menu.

12. In the **Switch Details** section, enable the **Virtual Chassis** slider button if you are deploying a Virtual Chassis at the remote site. Otherwise, leave the slider off.

    Enter the **Serial Number** of the switch or the primary member of the Virtual Chassis in the field.

13. The **Auto Activate** button is turned on by default. Turn it off if you want to disable auto-activation and use an activation code instead.

    If you left **Auto Activate** turned on, skip to step 16.

14. (Optional) If you turned off **Auto Activate**, enter an activation code in the field that appears.

    The code can be any combination of letters and numbers.

    Remember this code.

15. The **Zero Touch Provisioning** (ZTP) button is turned on by default. Turn it off if the switch is not upgraded to a Junos OS image version with support for a Phone-Home Client. If ZTP is disabled, you must manually copy (by using CLI), the Stage-1 configuration on to the switch.

    ZTP, if left on, begins immediately after the activation procedure.

16. (Optional) Enter LAN information for the branch site.

    This optional step allows you to define where the remote site LANs are connected to the EX switch. You can define as many LANs as needed by following the next 5 steps.

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

       The **Add LAN Segment** window appears.

    b. Enter a name for the LAN segment, such as **LAN1**, in the field provided.

    c. (Optional) Enter a **VLAN ID** for the LAN segment.

       If no VLAN ID is needed, you can safely remove the pre-populated value from the field.

    d. Click **Save** when finished.

       You can add as many LAN segments as you need by repeating this procedure.

17. Click **Next**.

    The wizard advances to the **Summary** page.

18. Review the configuration on the **Summary** page.

19. Click **OK** when satisfied, or click **Back** as needed to make any changes

If you need to edit anything, you can click the **Edit** links within the summary to go directly to that page of the wizard.

If you left auto-activate turned on, the activation procedure begins at this point. The **Site Activation** page appears. Skip to Step 20.

(Optional) If you turned off auto-activate, your site appears in the list with a status of **Configured** and you are now ready to activate the site:

a. Click the *site name* link.

   This takes you to the site page for this site with the **Overview** tab highlighted.

b. Click the **Devices** tab.

c. Click the **Check-box** next to the device name.

   The **Stage1 Config** button becomes active.

d. Click the **Stage1 Config** button.

   A new window appears containing the stage 1 configuration for this device.

e. Click the **Copy to Clipboard** button.

f. Click **OK**.

   The window closes.

g. Using a console or SSH connection, install the copied configuration on the EX switch and commit it.

   Assuming that the required network connectivity is in place from the EX switch, the switch connects back to CSO using an outbound SSH connection. When this connection is completed, the device will be activated in CSO; its status changes from **Expected** to **Provisioned**.

20. The **Site Activation** window proceeds through **Prestage Device** to **Detect Device** to **Bootstrap Device** and, finally to **Provision Device**.

Each stage will report success as it completes its operation. The window can be closed at any point. While the activation process is running, the **Site Status** column in the site list reports **Activating** and provides a link to **View** the activation wizard's progress. The **Site Status** changes to **Provisioned** once all the steps are successfully completed.

NOTE: In the event of an error or delay, you can open a read-only SSH session to the device from CSO. This will allow you to troubleshoot connection or other issues.

Once deployed, you can monitor and manage the switch through the Customer Portal's Switch Port Operational View.

5

CHAPTER

Standalone Next-Generation Firewall (NGFW) Deployment

Next-Generation Firewall (NGFW) Deployment

IN THIS SECTION

- NGFW Deployment Overview | 104
- NGFW Deployment Architecture | 105
- NGFW Deployment | 106

NGFW Deployment Overview

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 and Hybrid WAN deployments which focus on secure site-to-site connectivity and remote VNF deployment. A high-level view of the spoke site with NGFW is shown in Figure 4 on page 18.

Figure 42: NGFW Spoke Site

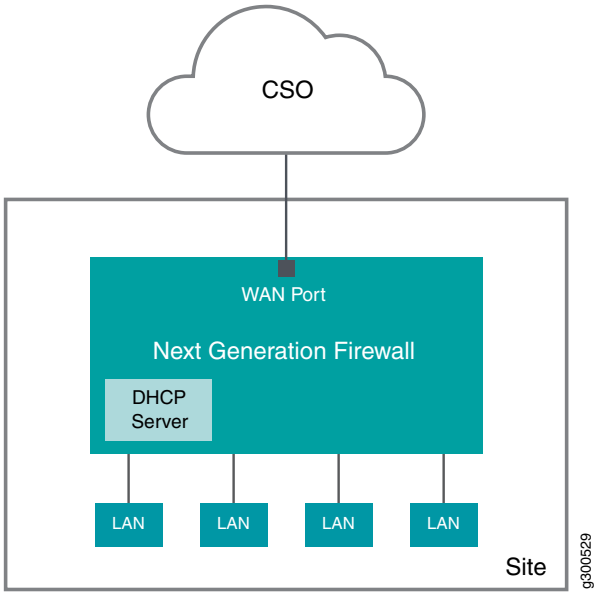

An NGFW deployment is carried out in the Customer Portal of CSO as a site deployment. The tenant under which the site is deployed must have the NGFW service available. This service is included in the tenant configuration by the tenant administrator during tenant onboarding. The remainder of this document

provides a brief discussion of the architecture, and the steps that you need to perform in order to complete a NGFW deployment in CSO.

NGFW Deployment Architecture

The architecture used in this example is described below.

The architecture for a cloud-hosted, CSO-managed NGFW deployment is very similar to any standalone firewall deployment as shown in Figure 4 on page 18. There is only one WAN port needed for communication with CSO. This port must get its IP address and gateway information from an available DHCP server. The gateway must provide a path to the Internet so that the NGFW can communicate with Juniper's redirect server.

CSO provisions the device and adds logging functionality. Optionally, default FW and NAT policies can be added during the initial provisioning process. After provisioning the site administrator can push additional GE, NAT, UTM, or IPS policies to the device.

Device monitoring is supported via the CSO GUI where you can view application and security logging data.

The remaining ports on the NGFW can be used for LAN communication at the site. Additionally, an EX Series access switch can be added after the NGFW deployment. This addition allows for further LAN management within the site, including the ability to add CSO-managed Mist WiFi access points to the site.

Table 11 on page 105shows the security devices supported in an NGFW deployment.

Table 11: Hardware and Software Matrix for Devices in an NGFW Deployment

Role	Platform	Models Supported	Junos OS Software Release Versions
NGFW Devices	SRX Series Security Gateways	• SRX300 • SRX320 • SRX340 • SRX345 • SRX550M	18.4R1 19.3R2-S1

NOTE: For the most up to date information on hardware and software support for CSO, see the Contrail Service Orchestration Release Notes.

NGFW Deployment

The procedure you follow to complete this task varies slightly depending on whether you are in the role of a CSO tenant administrator or OpCo administrator. A note is used where needed to account for these variances.

This procedure makes the following assumptions:

- You have already established your login credentials for CSO.
- The tenant for which you are creating the NGFW site is called **Example_Company**, and has already been created.
- The **Example_Company** tenant was added with NGFW WAN services capabilities.
- There is a working DHCP server available from which the WAN port of the NGFW device will obtain:
  - IP address
  - Address of a gateway router that can route traffic to the Internet

If any of these things are not true, see *Accessing Administration Portal*, *Accessing Customer Portal*, or *Adding a Single Tenant* as needed.

The steps to deploy an NGFW site are as follows:

1. Login to CSO using your login credentials.

   NOTE: If you are an OpCo administrator, navigate to **Tenants** in the left-navigation panel and select **Example_Company** from the list of tenants on the tenants page. If you are the tenant administrator, you will be placed in the Customer Portal for **Example_Company**.

2. In the Customer Portal for **Example_Company**, Navigate to **Resources > Site Management**.

   The **Sites** page appears.

3. Click the **Add** button and select **Add On-Premise Spoke (Manual)** from the list of options.

   The **Add On-Premise Spoke Site for Example_Company** page appears.

4. In the **Site Information** section, give the site a name such as **NGFW-Site1**.

5. In the **Site Capabilities** section, click the **Next Gen Firewall** icon.

   Depending on the configuration of the **Example_Company** tenant, there may be other icons available. Only select **Next Gen Firewall** for this example.

6. Click the right arrow icon > next to **Address and Contact Information** to expand this section.

   None of the fields are required, but adding address information for the site allows CSO to place an icon in the correct location for the site on maps on the monitoring page and show how it is linked to CSO. Without an address, CSO will place an icon at a default site.

7. Click the right arrow icon > next to **Advanced Configuration**.

   The two required fields, **Name Server IP List** and **NTP Server** are both pre-populated for you. Make changes as needed for your network to any of the fields.

8. Click **Next**.

   The wizard advances to the **WAN** page.

9. In the **Device Information** section, fill in the serial number of the SRX device you are onboarding.

10. The **Auto Activate** button is turned on by default. Turn it off if you want to disable auto-activation and use an activation code instead.

    Auto-activation, if left on, begins immediately after this add spoke site procedure is completed.

11. The **Zero Touch Provisioning** (ZTP) button is turned on by default. Turn it off if you want to pre-stage the device.

    ZTP, if left on, begins immediately after the activation procedure, if enabled.

12. Select the appropriate **In-Band Management** port from the pull-down list.

> **NOTE:** In-Band Management refers to management traffic that uses a connection that also carries non-management traffic. In this case, the in-band management port is the WAN port over which the device communicates with both CSO and the Internet.

13. Select a firewall policy from the pull-down list.

    CSO has a built-in firewall policy called **Default_Fw_Policy** that is provided for you. This policy is a zone-based policy intent that allows all traffic from any address in the trust zone to reach any address in the untrust zone.

14. Select a NAT policy from the pull-down list.

    CSO has a built-in NAT policy called **Default_NAT_Policy** that is provided for you. This policy is a Source-NAT policy that translates the source IP address of any traffic originating in the trust zone to the IP address of the trust-zone interface.

15. Click **Next**.

    The wizard advances to the **Summary** page.

16. Review the configuration on the **Summary** page as shown in Figure 43 on page 108.

Figure 43: NGFW Add Site Summary

The summary lists in text everything could be set in the wizard's GUI.

> **NOTE:** At the bottom of the summary page a **Save JSON** link is shown that allows you to download a JSON file of this site configuration. This JSON configuration can be modified for other sites so that they can be quickly imported without using the wizard workflow.

17. Click **OK** when satisfied, or click **Back** as needed to make any changes.

    If you need to edit anything, you can click the **Edit** links within the summary to go directly to that page of the wizard.

    The **Site Activation** wizard appears when you click **OK**.

18. If you left auto-activate turned on, the activation procedure begins at this point with the **Site Activation** page appearing.

If you turned off ZTP, you must copy the set commands from the **Pre-Stage Device** section of the **Site Activation** wizard. If you left ZTP on, it will begin as part of the site activation wizard.

19. The **Site Activation** window proceeds through **Prestage Device** to **Detect Device** to **Bootstrap Device** and, finally to **Provision Device**.

Each stage will report success as it completes its operation. The window can be closed at any point. While the activation process is running, the **Site Status** column in the site list reports **Activating** and provides a link to **View** the activation wizard's progress. The **Site Status** changes to **Provisioned** once all the steps are successfully completed.

6
CHAPTER

Hybrid WAN Deployment (uCPE)

Hybrid WAN Deployment Overview | **111**

Hybrid WAN Deployment Architecture | **112**

Hybrid WAN Deployment | **114**

Hybrid WAN Deployment Overview

This walkthrough highlights the steps you need to complete to deploy a Hybrid WAN solution. We'll use an NFX250 device as the CPE and an SRX Series device as the hub located in the SP cloud. We'll show you in the CSO GUI where, you need to go to complete each step. The document also provides some explanation of the choices you need to make at each step. It assumes that this is your first Hybrid WAN deployment.

In the Hybrid WAN deployment (also known as *distributed*), customers access network services from a CPE device located at the customer's site. These customer sites are called on-premises sites.Figure 44 on page 111 shows a simplified Hybrid WAN deployment.

Figure 44: Simplified Hybrid WAN Deployment

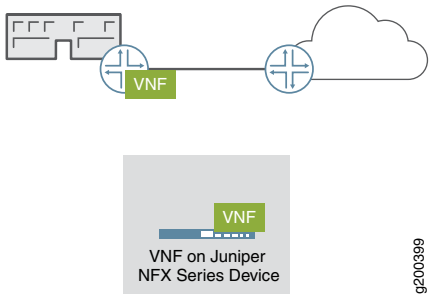

Initial configuration of the CPE device at the site is automated through the use of zero touch provisioning (ZTP) that is orchestrated through CSO. CSO also monitors the CPE device and its services, and can push software and configuration updates to the devices remotely, reducing operating expenses.

This deployment model is useful in environments where service delivery from the service provider's cloud is costly. In fact, CSO has been designed to require only modest bandwidth, needing as little as 30 kbps for probe and OAM traffic over Hybrid WAN connections where there are only a few sessions active. When AppQoE is involved, the bandwidth requirement increases to somewhere between 105 kbps and 2 Mbps, depending on the number of sessions.

During ZTP operations, if new device images are needed, they can be downloaded as part of the ZTP process, or pre-staged on the device. In those circumstances, the bandwidth requirement increases to a maximum of 5 Mbps only when device image download is needed. This makes these solutions applicable even in cases where connection bandwidth is limited or noisy.

The Hybrid WAN deployment uses a CPE device such as an NFX Series Network Services platform or SRX Series Services Gateway at the customer site and thus supports private hosting of network services at a site. The Hybrid WAN deployment can be extended to offer software defined wide area networking (SD-WAN) capabilities.

NOTE: If an SRX Series device is used as the CPE device at the customer site, it can not host VNFs. It can still offer all of the built-in services inherent in an SRX Series device.

In the Hybrid WAN deployment model, there is typically only one path from the on-premises site back to headquarters or the service provider cloud. The following sections describe the high-level architecture of a Hybrid WAN deployment and provide a walkthrough of how to set up CSO for Hybrid WAN.

Hybrid WAN Deployment Architecture

In the Hybrid WAN deployment, the Contrail Services Orchestration (CSO) software resides in the service provider's cloud, and is operated by the service provider to provide network services to the CPE devices at customer sites.

Figure 45 on page 112 shows a simple diagram of the Hybrid WAN solution. The cloud represents the service provider network to which the customer site is connected.

Figure 45: Hybrid WAN Solution

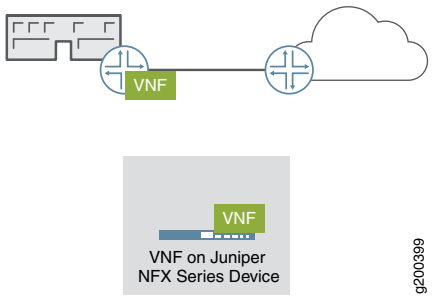

As mentioned previously, the Hybrid WAN deployment makes use of on-premises CPE devices in order to localize the delivery of network services and provide gateway router functionality. In this case, the Juniper Networks NFX Series or SRX Series devices act as the CPE devices.

In the case of an NFX Series device acting as the CPE, the gateway router function is provided by a built-in vSRX VNF and network services are hosted and provided from within an NFX device that is located at the customer site. This makes the network services extremely responsive from the point of view of the customer LAN, while negating the need for customer traffic to traverse the WAN in order to access the services.

In the case of an SRX Series device as the managed CPE device, only services native to the SRX (such as firewall, NAT, and UTM) can be provisioned and managed at the customer site by CSO. Other services(such

as WAN optimization) must be provisioned and managed separately from the SRX and cannot be managed by CSO.

In addition to the CPE devices, the Hybrid WAN solution also makes use of a provider edge (PE) router in the service provider cloud. The PE router terminates IPSec tunnels and provides policy-based access to the service provider's MPLS network. The PE and CPE devices communicate over one or more WAN links and make use of MPLS/GRE or IPSec tunnels for secure transport. Supported device types for a Hybrid WAN deployment and required software versions are shown in Table 12 on page 113.

Table 12: Hardware and Software Matrix for CPE Devices in a Hybrid WAN Deployment

Role	Platform	Models Supported	Junos OS Software Release Versions
CPE device	NFX250 Network Services Platforms	• NFX250-LS1 device • NFX250-S1 device • NFX250-S2 device	15.1X53-D497 18.4R3
	NFX150 Network Services Platforms	• NFX150-S1 device • NFX150-S1E device • NFX150-C-S1 device • NFX150-C-S1-AE/AA device • NFX150-C-S1E-AE/AA device	18.2X85-D12 19.3R2-S1
	SRX Series Services Gateways	• SRX300 • SRX320 • SRX340 • SRX345 • SRX4100 • SRX4200	15.1X49-D172 19.3R2-S1
		SRX1500	19.3R2-S1
	vSRX on an x86 server	vSRX	15.1X49-D172 19.3R2-S1

NOTE: For the most up to date information on hardware and software support for CSO, see the Contrail Service Orchestration Release Notes.

Selection of services, and some service management capabilities can be allocated to the customer by the service provider using the CSO Administration Portal. The customer would then access the allowed services and management capabilities by using the Customer Portal.

CSO manages the lifecycle of the VNFs hosted on the NFX CPE devices from creation in Network Designer, through instantiation, deployment, and finally through replacement or retirement.

> **NOTE:** Designer tools such as Network Designer are only available for on-premises deployments of CSO.

Hybrid WAN Deployment

IN THIS SECTION

- Modify Device Templates | **114**
- Add and Configure a New Tenant | **115**
- Add and Configure a Site for the Tenant | **116**

Modify Device Templates

From this point on in this deployment example, we assume that you are logged in to CSO as an OpCo administrator. The user name part of your credentials is an e-mail address that was used when your CSO account was set up. When an account is initially setup, CSO sends an e-mail to that address with a link that includes a one-time activation code. Clicking the link takes you to the CSO login page which then prompts you to set a password.

In this section, we modify an existing device template so that it works for this example.

1. Enter the login credentials for the Administration Portal.
2. Navigate to **Resources > Device Templates**.
3. Find the device template named **NFX250 as Managed Internet CPE**.

4. Select the check-box next to the template and then select **Template Settings** from the **Edit Device Template** pull-down menu.

   A new window titled Template Settings appears.

5. In the Template Settings window, ensure that the following things are set:

   - ACTIVATION_CODE_ENABLED: **ON**

     By requiring an activation code, a CPE device will not be allowed to communicate with CSO until the tenant has activated a site using the activation code. The value of the activation code will be set later in the process.

   - AUTO_DEPLOY_STAGE2_CONFIG: **OFF**

     Stage 2 configurations are configurations that can be added to a device after the initial, stage 1, provisioning of the device. This setting prevents the automatic deployment of a stage 2 configuration.

   - OOB_MGMT_ENABLED: **OFF**

     This setting ensures that the **jmgmt0** interface is not enabled on the NFX device. Since this is a managed Internet service and the NFX device will be sitting on the customer's premise, this might be a useful setting to prevent unwanted login by the tenant.

   - WAN_0**ge-0/0/11**

   Do not change any other settings.

6. Select **Save** when finished.

Add and Configure a New Tenant

In this section, we use the Administrator Portal to add a tenant to CSO.

To add a tenant:

1. Select **Tenants** from the left-navigation panel.

2. Click the**Add Tenant** button.

   If there are no tenants created yet, **Add Tenant** appears as a button. If there are tenants, click the Add icon (**+**)to create a new tenant.

3. In the **Add Tenant** window:

   - Enter a name for your tenant such as **Tenant1**.
   - Fill in the **Admin User** information.

- Select the check-boxes next to both**Roles** in the **Available** section and click the arrow link to move them to the **Selected** section.
- The password expiration defaults to 180 days.

  You can set any value from 1 to 365 days.
- Click **Next**.
- In the **Deployment Info** window, select the **Hybrid WAN** icon.
- Click **Next**.

  The window advances to the **Tenant Properties** section. For this example, browse the Tenant properties but do not make any changes.
- Click **Next**.

  The window advances to the **Summary** section. Review the summary.
- Click **OK**.

  A pop-up message appears that tells you that the **Add Tenant** job was started. After some time, your new tenant appears in the list of tenants.

Add and Configure a Site for the Tenant

In this section, we move to the Customer Portal for the newly configured tenant in order to create a site.

This procedure begins in the **Tenants** window of the Administration Portal, at the list of tenants.

1. Click on the name of the tenant that you just created.

   This will take you to the Customer Portal for that tenant. The **Dashboard** appears.

2. Select **Resources > Site Management** link from the left-navigation bar.

3. The **Sites** window appears. Click the **Add On-Premise Spoke (Manual)**.

   A new window titled **Add On-Premise Spoke Site for *Tenant*** appears.

4. Fill out the information in the **Site Information** section.

   Enter a site name that makes sense, like: **site1**.

   If you fill in the address information, CSO will use it to display the site on maps on some of the monitoring pages.

5. Click **Next**.

The **WAN** section appears.

6. Under **Device Template**, select the **NFX250** from the **Device Series** pull-down menu.

   This displays the device templates for NFX250 Series devices.

7. Select **NFX250 as Managed Internet CPE**.

8. Fill out the information In the **Device Information** section.

9. Disable **Auto Activate**.

   This enables the **Activation Code** field.

   > **NOTE:** If you leave **Auto Activate** enabled, CSO will attempt to connect with the device as soon as the site configuration is complete.

10. Enter an activation code for the NFX250.

    This code must be entered when the device is powered on for the first time in its final location.

    > **NOTE:** You cannot modify any settings for the WAN_0 interface because there are strict requirements for this device template that the WAN_0 must be an Internet-facing interface.

11. Click **Next** when finished.

    The window advances to the **Summary** section.

12. Review the **Summary** section.

13. Click **OK** when you're finished reviewing your entries.

    You'll see pop-up messages appear for site–creation job start and site–creation job finished.

    > **NOTE:** In the event of an error or delay, you can open a read–only SSH session to the device from CSO. This allows you to troubleshoot connection issues or other problems.

CHAPTER 7

Appendix A - Network Function Virtualization in Contrail Service Orchestration

Network Function Virtualization in the Contrail Service Orchestration Deployments | **119**

VNFs Supported by the Contrail Service Orchestration Solutions | **121**

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 46 on page 119, 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 46: Network Function Virtualization

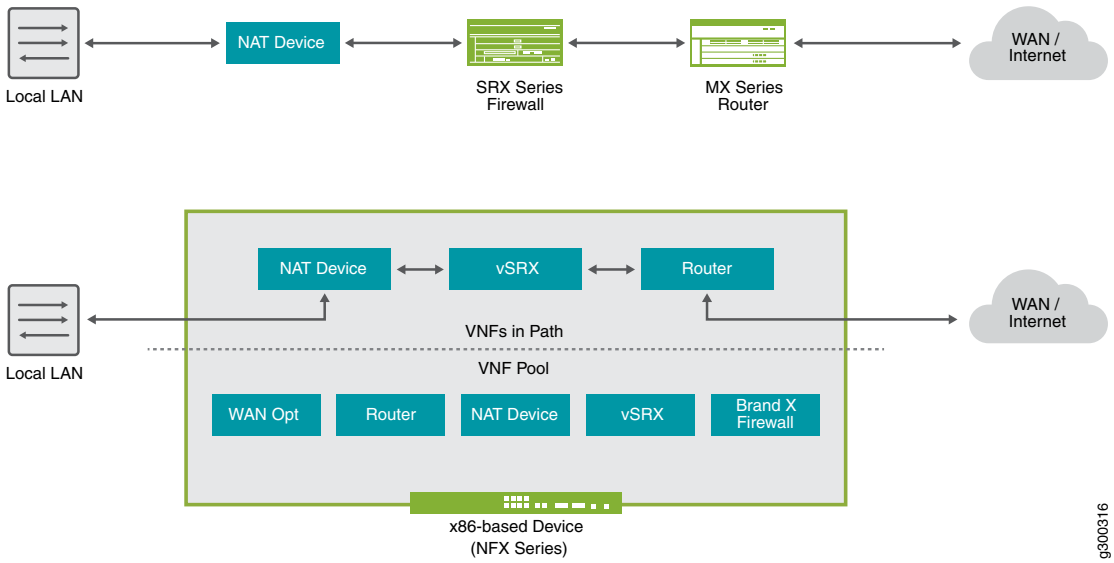

Juniper's CSO solutions comply with European Telecommunications Standards Institute (ETSI) standards for lifecycle management of network service instances.

The Contrail SD-WAN Solution uses the following components for the Network Functions Virtualization (NFV) environment:

- For the Hybrid WAN and SD-WAN deployments:
  - Network Service Orchestrator, together with Network Service Controller, provides ETSI-compliant management of the life cycle of network service instances.
  - Network Service Controller provides service-chaining and the VIM.

  - The CPE device provides the NFV infrastructure (NFVI).

Other CSO components connect to 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 CSO Release 5.1.1 and later, the functionality of these *Designer Tools* (used for on-premises CSO deployments) is available in *Configuration Templates* (used for cloud-delivered CSO deployments).

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 the Contrail Service Orchestration Solutions

Contrail Service Orchestration (CSO) supports Juniper Networks and third-party VNFs listed in Table 13 on page 121.

Table 13: VNFs Supported by Contrail Service Orchestration

VNF Name	Version	Network Functions Supported	Deployment Model Support
Juniper Networks vSRX	vSRX KVM Appliance 15.1X49-D123	• Network Address Translation (NAT) • Demonstration version of Deep Packet Inspection (DPI) • Firewall • Unified threat management (UTM)	Hybrid WAN and SD-WAN deployments supports NAT, firewall, and UTM.
Fortinet	5.6.3	Firewall	Hybrid WAN and SD-WAN deployments–NFX250 and NFX150 platforms.
Single-legged Ubuntu	16.04	Firewall	Hybrid WAN and SD-WAN deployments–NFX250 and NFX150 platforms.

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 platform 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 in a cloud-hosted CSO deployment.

NOTE: Currently, SD-WAN deployments support only layer 2 (L2) service chains while Hybrid WAN deployments can support L2 and L3 service chains.

In a cloud-hosted deployment, CSO only contains those VNFs installed by Juniper Networks' administrators. Requests for additional VNFs must be made through your account manager and Professional Services.

8

CHAPTER

Appendix B - Manual Staging of NFX

Install Junos OS Software onto an NFX Series Device from a USB Drive | **124**

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.

RELATED DOCUMENTATION

- Building Blocks Used for Contrail Service Orchestration Deployments | **20**
- Hybrid WAN Deployment Overview | **111**
- SD-WAN Deployment Overview | **42**