



---

# Contrail Service Orchestration

## Contrail Service Orchestration (CSO) Deployment Guide

Release  
4.1.0



---

Modified: 2019-04-01

Juniper Networks, Inc.  
1133 Innovation Way  
Sunnyvale, California 94089  
USA  
408-745-2000  
[www.juniper.net](http://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*  
4.1.0

Copyright © 2019 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 . . . . .	xi
	Documentation and Release Notes . . . . .	xi
	Documentation Conventions . . . . .	xi
	Documentation Feedback . . . . .	xiii
	Requesting Technical Support . . . . .	xiv
	Self-Help Online Tools and Resources . . . . .	xiv
	Creating a Service Request with JTAC . . . . .	xv
<b>Chapter 1</b>	<b>Solutions Overview . . . . .</b>	<b>17</b>
	About this Deployment Guide . . . . .	17
	Contrail Service Orchestration Solutions Overview . . . . .	17
	Building Blocks Used for Contrail Service Orchestration Deployments . . . . .	22
	Administrators . . . . .	22
	Portals . . . . .	22
	Tenants . . . . .	23
	Topologies . . . . .	23
	Points of Presence (POPs) . . . . .	25
	Sites . . . . .	26
	Customer Premises Equipment (CPE) . . . . .	29
	Virtual Route Reflector (VRR) . . . . .	29
	Service-Level Agreement (SLA) Profiles and Policies . . . . .	30
	Firewall Policies . . . . .	31
	Network Function Virtualization in the Contrail Service Orchestration Deployments . . . . .	32
	Number of Sites and VNFs Supported in Contrail Service Orchestration . . . . .	36
	VNFs Supported by the Contrail Service Orchestration Solutions . . . . .	37
<b>Chapter 2</b>	<b>Deployment Tools . . . . .</b>	<b>41</b>
	Contrail Service Orchestration (CSO) Deployment Tools . . . . .	41
	Contrail Services Orchestration (CSO) GUIs . . . . .	41
	Designing and Publishing Network Services . . . . .	44
	Contrail Service Orchestration License Tool . . . . .	45
	Overview of the License Page . . . . .	45
<b>Chapter 3</b>	<b>SD-WAN Deployment . . . . .</b>	<b>47</b>
	SD-WAN Deployment Overview . . . . .	47
	SD-WAN Deployment Architectures . . . . .	48
	SD-WAN Reference Architecture . . . . .	49
	Spoke Devices . . . . .	50
	Hub Devices (SD-WAN Gateway) . . . . .	52
	Underlay (Physical) Network . . . . .	53
	Overlay (Tunnels) Network . . . . .	55

	SD-WAN Orchestrator/Controller . . . . .	57
	Your First SD-WAN Deployment . . . . .	59
	Before You Begin . . . . .	61
	Download Application Signatures . . . . .	61
	Upload Licenses . . . . .	62
	Create and Configure a New Tenant . . . . .	62
	Enable Application Traffic Type Profile . . . . .	63
	Modify Device Templates . . . . .	64
	Upload Software Image for vSRX . . . . .	65
	Create a Point of Presence (POP) for the Hub Site . . . . .	66
	Create Cloud Hub Device . . . . .	66
	Create and Configure the Tenant's Hub Site . . . . .	69
	Create and Configure a Spoke Site for the Tenant . . . . .	69
	Install License on Device . . . . .	72
	Install Application Signature . . . . .	72
	Add Firewall and NAT Policies to the Topology . . . . .	73
	Create SD-WAN SLA Profiles and Policies . . . . .	74
<b>Chapter 4</b>	<b>Distributed CPE Deployment (uCPE) . . . . .</b>	<b>77</b>
	Hybrid WAN Deployment Overview . . . . .	77
	Hybrid WAN (Distributed) Deployment Architecture . . . . .	78
	Your First Hybrid WAN (Distributed) Deployment . . . . .	80
	Install Junos Software onto NFX from USB Port . . . . .	80
	Modify Device Templates . . . . .	83
	Create and Configure a New Tenant . . . . .	84
	Create and Configure a Site for the Tenant . . . . .	85
<b>Chapter 5</b>	<b>Centralized CPE Deployment (vCPE) . . . . .</b>	<b>87</b>
	Centralized Deployment Overview . . . . .	87
	Centralized Deployment Architecture Overview . . . . .	87
	Setting Up a Centralized Deployment . . . . .	88
	Create Network Service . . . . .	90
	Create POP . . . . .	92
	Add Tenant . . . . .	94
	Allocate Network Service . . . . .	95
	Create Cloud Site . . . . .	96
<b>Chapter 6</b>	<b>Appendix A - Contrail Cloud Reference Architecture for Centralized Deployment . . . . .</b>	<b>99</b>
	About this Reference Architecture . . . . .	99
	Architecture of the Contrail Cloud Implementation in the Centralized Deployment . . . . .	99
	Architecture of the Contrail Cloud Implementation . . . . .	99
	Architecture of the Servers . . . . .	101
	Architecture of the Contrail Nodes . . . . .	102
	Cabling the Hardware for the Centralized Deployment . . . . .	103
	Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment . . . . .	106
	Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment . . . . .	107

	Configuring the MX Series Router for the Contrail Cloud Implementation in a Centralized Deployment . . . . .	109
	Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment . . . . .	111
<b>Chapter 7</b>	<b>Appendix B - Uploading VNF Images for Centralized Deployment . . . . .</b>	<b>113</b>
	Uploading the vSRX VNF Image for a Centralized Deployment . . . . .	113
	Uploading the LxCIPtable VNF Image for a Centralized Deployment . . . . .	115
	Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment . . . .	117
<b>Chapter 8</b>	<b>Appendix C - Manual Staging of NFX . . . . .</b>	<b>119</b>
	Install Junos Software onto NFX from USB Port . . . . .	119



# List of Figures

<b>Chapter 1</b>	<b>Solutions Overview</b> .....	<b>17</b>
	Figure 1: Basic SD-WAN Concept .....	19
	Figure 2: Hybrid WAN Deployment .....	20
	Figure 3: Centralized Deployment .....	21
	Figure 4: Combined Deployment .....	21
	Figure 5: Centralized CPE .....	23
	Figure 6: Distributed CPE (or Hybrid WAN) .....	24
	Figure 7: Hub-and-Spoke Topology .....	24
	Figure 8: Dynamic Mesh Topology .....	25
	Figure 9: Points of Presence (POPs) .....	25
	Figure 10: CPE Devices .....	29
	Figure 11: VRR Overview .....	30
	Figure 12: Network Function Virtualization .....	32
	Figure 13: NFV Components of the Cloud CPE Solution .....	34
<b>Chapter 2</b>	<b>Deployment Tools</b> .....	<b>41</b>
	Figure 14: Designer Tools Overview .....	44
<b>Chapter 3</b>	<b>SD-WAN Deployment</b> .....	<b>47</b>
	Figure 15: SD-WAN Example Deployment Topology .....	48
	Figure 16: SD-WAN Architecture .....	49
	Figure 17: SD-WAN Reference Architecture .....	50
	Figure 18: On-premise Spoke Devices .....	50
	Figure 19: SD-WAN Hub Devices .....	53
	Figure 20: SD-WAN Underlay Network .....	54
	Figure 21: WAN Interface Types .....	55
	Figure 22: SD-WAN Hub-and-Spoke Overlay .....	56
	Figure 23: SD-WAN Hub-and-Spoke Topology .....	56
	Figure 24: SD-WAN Full Mesh Topology .....	57
	Figure 25: CSO Dashboard .....	58
	Figure 26: CSO Abstraction Layers .....	58
	Figure 27: SD-WAN Deployment Workflow .....	60
<b>Chapter 4</b>	<b>Distributed CPE Deployment (uCPE)</b> .....	<b>77</b>
	Figure 28: Simplified Hybrid WAN Deployment .....	77
	Figure 29: Distributed CPE .....	78
<b>Chapter 5</b>	<b>Centralized CPE Deployment (vCPE)</b> .....	<b>87</b>
	Figure 30: Simplified Centralized Deployment .....	87
<b>Chapter 6</b>	<b>Appendix A - Contrail Cloud Reference Architecture for Centralized Deployment</b> .....	<b>99</b>

Figure 31: Architecture of Contrail Cloud Implementation . . . . .	100
Figure 32: Architecture of Servers in the Central POP for a Non-Redundant Installation . . . . .	101
Figure 33: Architecture of Servers in the Central POP for a Redundant Installation . . . . .	102
Figure 34: Logical Representation of Contrail Controller Nodes . . . . .	102
Figure 35: Logical Representation of Contrail Compute Nodes . . . . .	103



# List of Tables

	<b>About the Documentation</b> . . . . .	<b>xi</b>
	Table 1: Notice Icons . . . . .	xii
	Table 2: Text and Syntax Conventions . . . . .	xii
<b>Chapter 1</b>	<b>Solutions Overview</b> . . . . .	<b>17</b>
	Table 3: Site Types by Deployment . . . . .	27
	Table 4: Number of Sites and VNFs Supported . . . . .	36
	Table 5: Number of Sites, Tenants, and Tunnels Supported for a Full-Mesh SD-WAN Deployment . . . . .	36
	Table 6: VNFs Supported by Contrail Service Orchestration . . . . .	38
<b>Chapter 2</b>	<b>Deployment Tools</b> . . . . .	<b>41</b>
	Table 7: Access Details for the GUIs . . . . .	42
<b>Chapter 3</b>	<b>SD-WAN Deployment</b> . . . . .	<b>47</b>
	Table 8: NFX Hardware and Software Matrix for On-premise Spoke Devices . . . . .	51
	Table 9: SRX Hardware and Software Matrix for On-premise Spoke Devices . . . . .	51
	Table 10: Hub Devices . . . . .	53
<b>Chapter 4</b>	<b>Distributed CPE Deployment (uCPE)</b> . . . . .	<b>77</b>
	Table 11: Hardware and Software Matrix for CPE Devices in a Hybrid WAN Deployment . . . . .	79
<b>Chapter 5</b>	<b>Centralized CPE Deployment (vCPE)</b> . . . . .	<b>87</b>
	Table 12: Hardware and Software Matrix for the PE Router in the Centralized Deployment Model . . . . .	88
<b>Chapter 6</b>	<b>Appendix A - Contrail Cloud Reference Architecture for Centralized Deployment</b> . . . . .	<b>99</b>
	Table 13: Connections for EX Series Switch . . . . .	103
	Table 14: Connections for QFX Series Switch . . . . .	104
	Table 15: Connections for MX Series Router . . . . .	105



# About the Documentation

- Documentation and Release Notes on page xi
- Documentation Conventions on page xi
- Documentation Feedback on page xiii
- Requesting Technical Support on page xiv

## 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 xii 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 xii defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
<b>Bold text like this</b>	Represents text that you type.	To enter configuration mode, type the <b>configure</b> command:  user@host> <b>configure</b>
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> <b>show chassis alarms</b>  No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> <li>Introduces or emphasizes important new terms.</li> <li>Identifies guide names.</li> <li>Identifies RFC and Internet draft titles.</li> </ul>	<ul style="list-style-type: none"> <li>A policy <i>term</i> is a named structure that defines match conditions and actions.</li> <li><i>Junos OS CLI User Guide</i></li> <li>RFC 1997, <i>BGP Communities Attribute</i></li> </ul>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name:  [edit] root@# <b>set system domain-name</b> <i>domain-name</i>


Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"><li>To configure a stub area, include the <b>stub</b> statement at the <b>[edit protocols ospf area area-id]</b> hierarchy level.</li><li>The console port is labeled <b>CONSOLE</b>.</li></ul>
< > (angle brackets)	Encloses optional keywords or variables.	<b>stub &lt;default-metric <i>metric</i>&gt;;</b>
(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.	<b>broadcast   multicast</b>  <b>(<i>string1</i>   <i>string2</i>   <i>string3</i>)</b>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<b>rsvp { # Required for dynamic MPLS only</b>
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	<b>community name members [ <i>community-ids</i> ]</b>
Indentation and braces ( { } )	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options {   static {     route default {       nexthop <i>address</i>;       retain;     }   } }</pre>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"><li>In the Logical Interfaces box, select <b>All Interfaces</b>.</li><li>To cancel the configuration, click <b>Cancel</b>.</li></ul>
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select <b>Protocols&gt;Ospf</b> .

## 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](mailto: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 on page 17](#)
- [Contrail Service Orchestration Solutions Overview on page 17](#)
- [Building Blocks Used for Contrail Service Orchestration Deployments on page 22](#)
- [Network Function Virtualization in the Contrail Service Orchestration Deployments on page 32](#)
- [Number of Sites and VNFs Supported in Contrail Service Orchestration on page 36](#)
- [VNFs Supported by the Contrail Service Orchestration Solutions on page 37](#)

### About this Deployment Guide

---

The intent of this deployment guide is to provide a comprehensive understanding of the available Contrail Service Orchestration (CSO) solutions. In order to do that, we will:

- Briefly discuss each of the available solutions
- Discuss the building blocks used in every deployment
- Discuss the tools used to put the blocks together
- Provide an end-to-end walkthrough of each of the solutions that covers the specifics involved in deploying them

This guide is hosted on the Contrail Service Orchestration Documentation page, alongside several other guides, including:

- [CSO Installation and Upgrade Guide](#)
- [CSO User Guide](#)
- [CSO Monitoring and Troubleshooting Guide](#)
- [CSO Design and Architecture Guide](#)
- And more

### Contrail Service Orchestration Solutions Overview

---

Juniper Networks Cloud Customer Premises Equipment (CPE) and SD-WAN solutions offer automated service delivery to branch network environments, leading to cost savings

over traditional branch networks, while improving network agility and reducing configuration errors.

Traditional branch networks use many dedicated network devices with proprietary software to provide services and require extensive equipment refreshes every 3-5 years to accommodate advances in technology. Both configuration of standard services for multiple sites and customization of services for specific sites are labor-intensive activities. As branch offices rarely employ experienced IT staff on site, companies must carefully plan network modifications and analyze the return on investment of changes to network services.

In contrast, the Cloud CPE solutions enable a branch site to access network services based on Juniper Networks and third-party virtualized network functions (VNFs) that run on commercial off-the-shelf (COTS) servers located in a central office (CO) or on a CPE device located at the customer site. This approach maximizes the flexibility of the network, enabling use of standard services and policies across sites and enabling dynamic updates to existing services. Customization of network services is fast and easy, offering opportunities for new revenue and quick time to market.

CSO provides a flexible and scalable micro-service architecture platform for deploying new service offerings. 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 and their customers.

CSO offers multiple deployment solutions that benefit both the service providers and their customers. The solutions are split into two overall groups, Cloud CPE solutions and SD-WAN solutions. The Juniper Networks Cloud Customer Premises Equipment (CPE) and the SD-WAN solutions both address the difficulties in traditional CPE deployments like:

- the need for multiple hardware and software platforms to deploy multiple network services
- long wait times for service instantiation
- network disruption for service instantiation
- fixed service offerings


CSO uses these deployment solutions to transform traditional branch networks, offering opportunities for highly flexible networks, rapid introduction of new services, automation of network administration, and cost savings. The solutions can be implemented by service providers for their customers or by Enterprise IT departments in a campus and branch environment. In this documentation, service providers and Enterprise IT departments are called service providers, the users of their network services are called customers, and solution and deployment are used interchangeably.

The following list briefly describes each of the available CSO deployment models.

## **SD-WAN Deployment Model**

The SD-WAN solution offers a flexible and automated way to route traffic through the cloud using overlay networks. This solution uses CPE devices located at on-premises sites. At its most basic, an SD-WAN solution needs multiple sites, multiple connections between sites, and a controller as shown in [Figure 1 on page 19](#).

**Figure 1: Basic SD-WAN Concept**



The CPE devices, or spokes, have a WAN side and a LAN side. On the WAN side, hub-and-spoke and full mesh topologies are supported. The CPE devices will use at least two and up to four interfaces as connection paths to cloud-based hubs, cloud-based spokes, other on-premises sites, or to the Internet. CSO allows you to give preference to one path over another for any given traffic. Thus, business-critical traffic could be routed through the service provider's cloud-based hub using MPLS/GRE while non-critical traffic could be routed over the Internet connection through an IPsec tunnel. Each path can have a service level agreement (SLA) profile applied which monitors the path for latency, congestion, and jitter and accounts for path preference. Should the path fail to meet one or more of the required parameters, traffic will be 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. Starting with CSO Release 4.0.0, spoke devices can also provide service chains of network services in addition to the routing flexibility already available.

One CSO installation can support a combined centralized and distributed deployment and an SD-WAN solution simultaneously.

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.


### Cloud CPE Distributed Deployment Model

The Cloud CPE Distributed Deployment Model is also known as distributed deployment, Hybrid WAN or uCPE. Since the CSO GUI uses the term Hybrid WAN, we also use it in this deployment guide.

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* in this documentation. In the deployment workflows used in the CSO GUI, this deployment

is known as Hybrid WAN. [Figure 2 on page 20](#) illustrates a simplified Hybrid WAN deployment.

**Figure 2: 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 30kbps 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 105kbps and 2Mbps, 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 5Mbps only when device image download is needed. This makes these solutions applicable even in cases where connection bandwidth is limited or noisy.

The distributed CPE 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 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.

## Cloud CPE Centralized Deployment Model

The Cloud CPE Centralized Deployment Model is also known as centralized deployment or vCPE.

In the centralized deployment, customers access network services remotely from a service provider's cloud. Sites that access network services in this way are called *service edge sites* in this documentation. [Figure 3 on page 21](#) illustrates a simplified centralized deployment.


*Figure 3: Centralized Deployment*

The only equipment that needs to be configured in this deployment resides at the service provider's cloud. This deployment model is useful when few remote sites are accessing services and cost of traffic back to the CO for service delivery is not an issue. The centralized deployment offers a fast migration route and this deployment is the recommended model for sites that can accommodate network services, particularly security services, in the cloud. There are no CPE devices deployed at customer sites in a centralized deployment. All network services are deployed in the service provider's cloud.

#### **A Combined Centralized and Distributed Deployment**

In this deployment, the network contains both service edge sites and on-premises sites. A customer can access network services from both service edge sites and on-premises sites. However, you cannot use the same network service at both locations. If you require the same network service at both the service edge and on-premises, you must create two identical network services with different names and deploy one at the service edge site and the other at the on-premise site.

[Figure 4 on page 21](#) illustrates a simplified combined deployment.

*Figure 4: Combined Deployment*

Implementing a combination deployment in which some sites use the centralized deployment and some sites use the distributed deployment provides flexible access based on customer site capabilities and cost factors.

Since the combined deployment is simply a combination of the centralized and distributed deployments, this guide does not provide an end-to-end walkthrough of this deployment option.

## Building Blocks Used for Contrail Service Orchestration Deployments

---

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 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 administrator or global admin. This administrator has full read and write access to the entirety of the CSO platform from the global domain. He or she 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. Successful login as cspadmin places the user in the Administration Portal of the global domain; the user can switch into the Customer Portal of any OpCo or Tenant.

The next level of administrator is the Operating Company or OpCo 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. 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 last 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 all three levels, Global, OpCo, and Tenant. While operator users are not, strictly speaking, 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 both Administration and Customer Portals available. Access to any given portal is controlled by a user's login. If your login does not grant access to an administration portal, then you cannot see or access any of the elements of an administration portal.

Administration portals allow tenant creation, OpCo 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 global administrators can see the **Tenants** page in the Administration Portal

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

## Tenants

CSO uses the tenant element to logically separate one customer from another. A global service provider (SP) administrator creates one tenant to represent each customer for which they will provide network services. If the global SP is logically split up into multiple OpCos, then the individual OpCo administrators can create tenants that represent their customers.

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 SP and 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.


## Topologies

There are, essentially, four network topologies supported in CSO. When defining a tenant, the global administrator must decide if that tenant will be able to use:

- **The 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.


It is within this topology that a Centralized CPE solution and its network services are deployed. [Figure 5 on page 23](#) shows an example where the VNFs used in the Centralized CPE are deployed in the SP cloud.

*Figure 5: Centralized CPE*

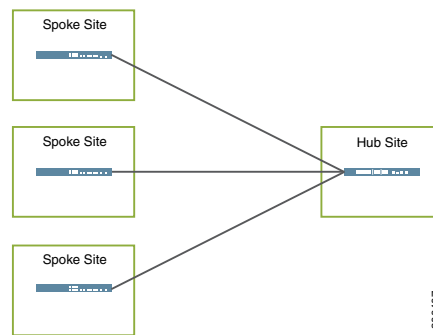


- **Standalone Topology**—This topology is one in which the customers, or users of network services remain separate from each other with no means of communication amongst themselves.

This is the topology of Distributed CPE, or Hybrid WAN solutions 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 24](#) shows an example where the VNF functions are located on-premises, but the on-premises site has no access to other sites.

*Figure 6: Distributed CPE (or Hybrid WAN)*

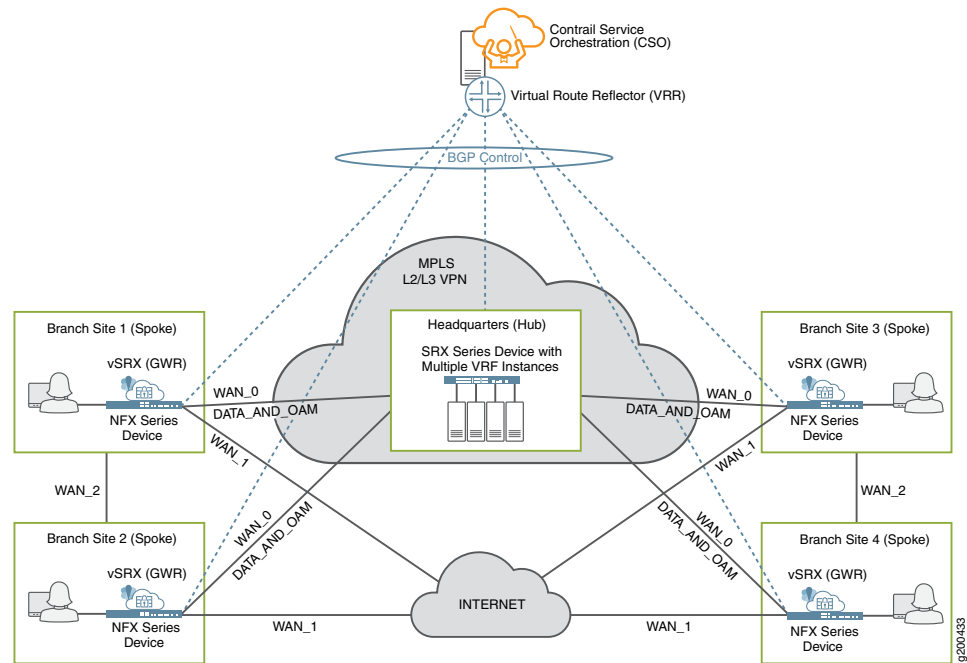
- Hub-and-Spoke Topology**—This topology is available for SD-WAN deployments. Given that SD-WAN is intended specifically to enable and enhance the efficacy of WAN communication using network overlays, this topology does allow for communication from site to site. Specifically, if one site needs to communicate with another site, that communication goes through the hub on its way to the other site. [Figure 7 on page 24](#) shows a very basic example of hub-and-spoke topology. VNFs can be deployed at any of the locations shown.

*Figure 7: 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 8 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 consideration must be given where resources are constrained. Tunnels from one site to another are created on-demand thereby conserving resources and improving overall performance.



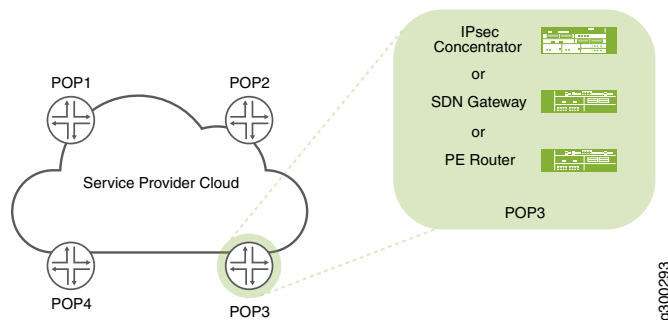
Figure 8: Dynamic Mesh Topology



## Points of Presence (POPs)

A POP is a place, usually at the SP Cloud edge, where network services can be deployed and underlay network connections are made to remote sites. POPs can have PE router, IPsec concentrator, and SDN Gateway devices assigned to them.

Figure 9: 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.

## 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 for creation depend on the type of deployment you are creating: SD-WAN, Hybrid WAN, or Centralized CPE. Sites can be created by the Global administrator, the OpCo administrator, or the Tenant administrator. Sites can be named whatever makes sense for the SP or Tenant. [Table 3 on page 27](#) lists what types of sites can be created within each deployment.

**Table 3: Site Types by Deployment**

Deployment	Available Site Types	Uses	Service Notes
SD-WAN	On-premise Spoke	Use this site type for locating NFX Series or SRX Series devices at customer sites in either a hub-and-spoke or full mesh topology.	<p>SRX Series devices deployed as on-premises spoke devices can not host VNF-based network services.</p> <p>NFX devices used as on-premise spoke devices can support ADSL, VDSL, and LTE access links, but cannot 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.</p> <p>Local breakout is supported on this type of site when using the full mesh topology.</p>
	Cloud Hub	Use this type of site for locating MX Series or SRX series in a SP cloud. The cloud hub devices are used for establishment of IPSec tunnels. Cloud hub devices are multi-tenant (shared amongst multiple sites) through the use of VRF instances configured on them.	<p>You must specify the capability of the cloud 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.</p> <p>A cloud hub device is required for the dynamic mesh topology.</p> <p>Local breakout is not supported on Cloud Hub sites.</p>
	Cloud Spoke	This type of site is specifically for deploying a vSRX in a tenant's Amazon Web Services (AWS) Virtual Private Cloud (VPC)	<p>Firewall and UTM services are available to protect the customer's resources in AWS VPC.</p> <p>Connectivity between VPC resources and on-premise sites.</p> <p>WAN_0, WAN_1, and LAN interfaces need to be predefined in VPC.</p> <p>Two elastic IP addresses need to be reserved in VPC to attach to WAN interfaces later.</p> <p>VPC should be created and attached to an Internet gateway.</p> <p>Only hub-and-spoke topology supported.</p> <p>Hub needs to have public IPs on in its WAN interfaces.</p> <p>Hub WAN interface type should be set as Internet during onboarding.</p>
	Gateway Site	Use this type of site, along with SRX4x00 Series Services Gateway devices, to provide additional capabilities to those of a normal spoke site.	

Table 3: Site Types by Deployment (continued)

Deployment	Available Site Types	Uses	Service Notes
			<p>This type of site has the following capabilities:</p> <ul style="list-style-type: none"> <li>• Can behave as a normal spoke</li> <li>• Anchor point for spokes for dynamic VPN creation</li> <li>• Provides on-premise central breakout option</li> <li>• Can host a data center department. Can import BGP and OSPF routes from the LAN-side L3 device. Thus creating a dynamic LAN segment.</li> <li>• Automatically meshed with other gateway sites</li> <li>• Regular spoke sites can be assigned to associate with a gateway site.</li> <li>• Supports local, central and cloud breakout profiles with intent-based rules for more granular breakout control.</li> </ul>
Hybrid WAN/Distributed CPE	Local Service Edge	<p>Automation point in the network closest to the customer CPE location, for example: the PE-Router, where centralized service VNFs can be attached.</p> <p>Use this site type for customers who access the Internet through a VPN in the SP cloud.</p>	The Local Service Edge site acts as a service attachment point. During site creation, an optional VRF can be created on the gateway router to handle routing and forwarding specifically for the centralized service attachment point.
	Regional Service Edge	<p>Automation point deeper in the service provider network that performs centralized services for many branches, for example: a hub router deep inside the enterprise network.</p> <p>Use this site type for each branch location in the customer network. For use with customers who access the Internet from their local site.</p>	The end-point is identified only by route-target. The centralized VNF (network-service) connectivity is orchestrated only by peering using BGP routing protocols. No configuration changes are made to the hub router.

Table 3: Site Types by Deployment (continued)

Deployment	Available Site Types	Uses	Service Notes
Centralized CPE	Local Service Edge	<p>The local service edge is the automation point in the network closest to the customer location (SDN gateway) where centralized service VNFs can be attached.</p> <p>Use this site type for customers who access the Internet through a VPN in the SP cloud.</p>	<p>Site acts as a service attachment point. During site creation, an optional VRF can be created on the gateway router to handle routing and forwarding specifically for the centralized service attachment point.</p> <p>If a PNE is configured, then it must be associated with a site and a VNF must be created on the PNE for each associated site.</p>
	Regional Service Edge	<p>Automation point deeper in the service provider network that performs centralized services for many branches, for example: a hub router deep inside the enterprise network.</p> <p>Use this site type for each branch location in the customer network. For use with customers who access the Internet from their local site.</p>	<p>The end-point is identified only by route-target. The centralized VNF (network-service) connectivity is orchestrated only by peering using BGP routing protocols. No configuration changes are made to the hub router.</p>

## 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 in Hybrid WAN deployments or as on-premise spoke devices in SD-WAN deployments. [Figure 10 on page 29](#) shows available CPE device types.

Figure 10: CPE Devices



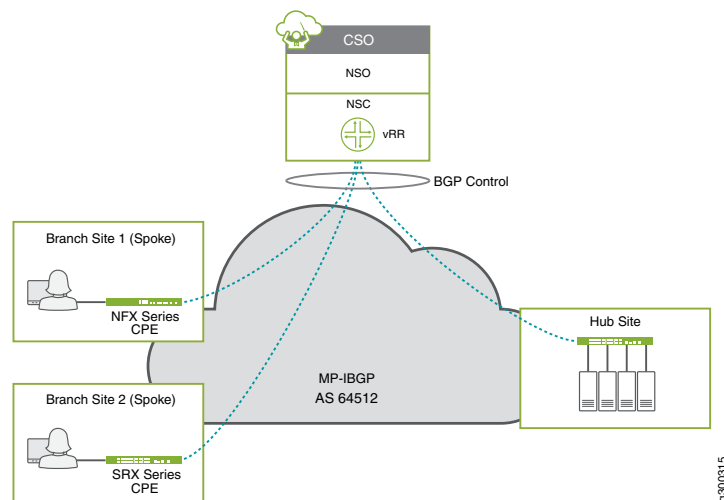
NFX250 and NFX150 Series Network Services Platforms, SRX300, SRX 550M, SRX4100, SRX4200, and vSRX Series Services Gateways can all be deployed as CPE devices. The NFX series devices provide the ability to host VNFs that can be deployed within the Hybrid WAN and SD-WAN solutions. The SRX Series devices 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.

## 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. Starting in CSO Release 4.0.0, the OAM interfaces can be implemented using dedicated IPsec tunnels which allows CPE and hub devices to be behind NAT. [Figure 11 on page 30](#) illustrates the concept of the VRR

**Figure 11: VRR Overview**



The number of VRRs available varies by CSO installation size (small, medium, or large). More VRRs can be added to your installation if needed. See [CSO Installation and Upgrade Guide](#) for details.

## Service-Level Agreement (SLA) Profiles and Policies

CSO allows for the creation of SLA profiles that can be mapped to SD-WAN policies for traffic management in an SD-WAN deployment. SLA profiles are created for applications or groups of applications for all tenants. An SLA profile consists of a set of configurable constraints that can be defined in the unified portal for both the Administration and Customer Portals.

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



**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 [Configuring Application SLA Profiles](#) in the *Contrail Service Orchestration User Guide* for more details.

## Firewall Policies

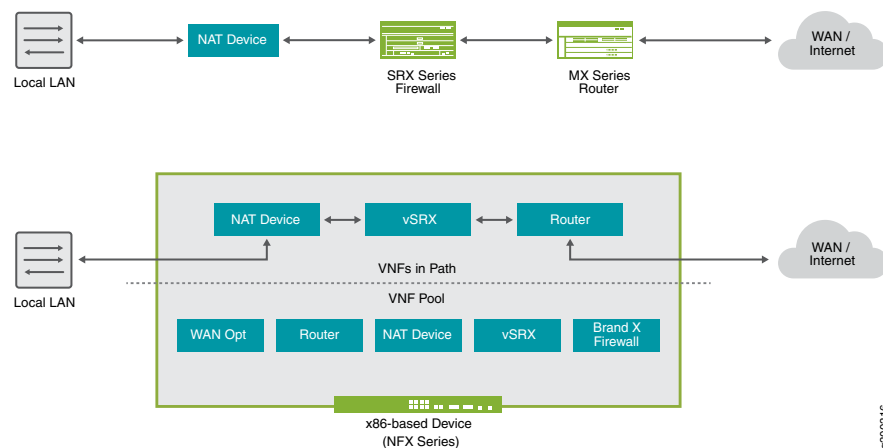
Accessed through the Customer Portal, CSO presents firewall policies as *intent-based* policies. Firewall policies provide security functionality by enforcing intents on traffic that passes through a device. Traffic is permitted or denied based on the action defined as the firewall policy intent. If 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.

## 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 12 on page 32](#) 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 12: Network Function Virtualization*



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

The Cloud CPE and SD-WAN Solutions use the following components for the Network Functions Virtualization (NFV) environment:

- For the centralized deployment:
  - Network Service Orchestrator provides ETSI-compliant management of the life cycle of network service instances.
  - This application includes RESTful APIs that you can use to create and manage network service catalogs.
  - Contrail OpenStack provides the following functionality:
    - Underlying software-defined networking (SDN) to dynamically create logical service chains that form the network services
    - NFV infrastructure (NFVI).
    - Virtualized infrastructure manager (VIM)
- For the Distributed CPE 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.
- Service and Infrastructure Monitor, which works with Icinga, an open source enterprise monitoring system to provide real-time data about the Cloud CPE solution, such as the status of virtualized network functions (VNFs), virtual machines (VMs), and physical servers; information about physical servers' resources; components of a network service (VNFs and VMs hosting a VNF); counters and other information for VNFs.

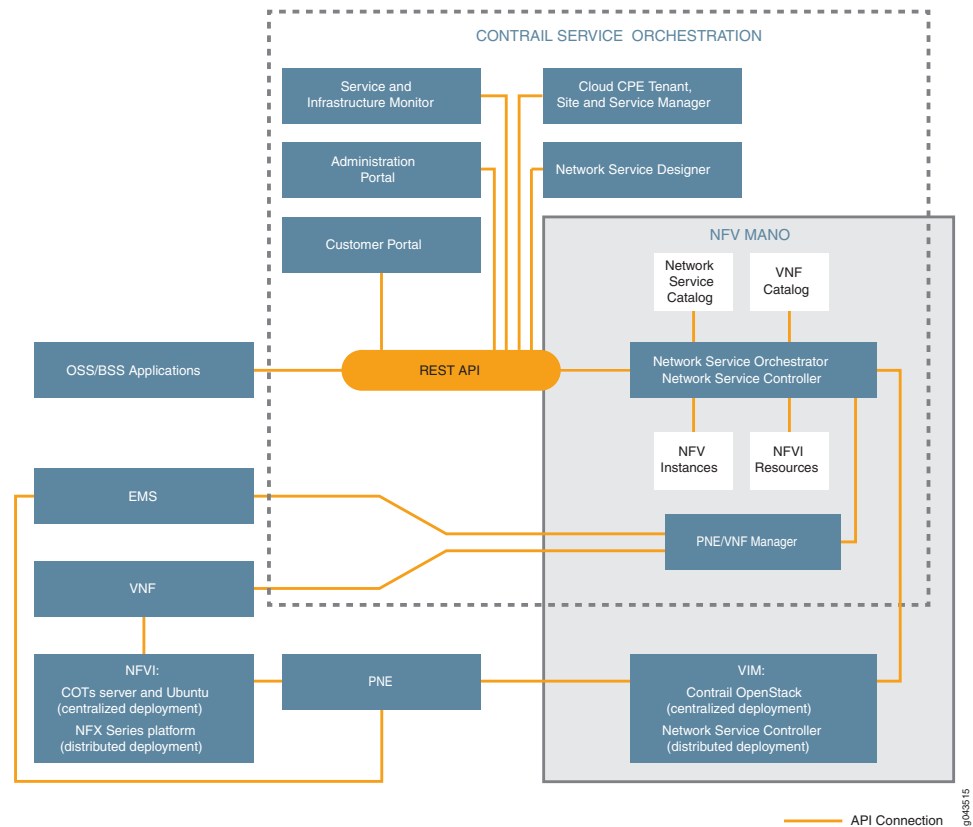
The Cloud CPE solution extends the NFV model through the support of physical network elements (PNEs). A PNE is a networking device in the deployment that you can configure through CSO, but not use in a service chain. Configuration of the PNE through CSO as opposed to other software, such as Contrail 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 the MX Series router that acts as an SDN gateway in a centralized deployment.

In the distributed deployment, VNFs reside on a CPE device located at a customer site. The NFX250 and NFX150 are switches that host the vSRX application as a VNF to enable routing and IPSec VPN access with the service provider's POP. MX Series routers, configured as provider edge (PE) routers, provide managed Layer 1 and Layer 2 access

and managed MPLS Layer 3 access to the POP. Network Service Controller provides the VIM, NFVI, and device management for the NFX. Network Service Controller includes Network Activator, which enables remote activation of the NFX Series device when the site administrator connects the device and switches it on.

Figure 13 on page 34 illustrates how the components in the Cloud CPE solution interact and how they comply with the ETSI NFV MANO model.

Figure 13: NFV Components of the Cloud CPE Solution



OSS/BSS applications and Contrail Service Orchestration (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.

The following process describes the interactions of the components when a customer requests the activation of a network service:

1. Customers send requests for activations of network services through Customer Portal or OSS/BSS applications.
2. Service and Infrastructure Monitor is continuously tracking the software components, hardware components, and processes in the network.
3. Network Service Orchestrator receives requests through its northbound REST API and:
  - For the centralized deployment:
    - a. Accesses information about the network service and associated VNFs from their respective catalogs, and communicates this information to the VIM, which is provided by Contrail OpenStack.
    - b. Sends information about the VNF to VNF Manager.
  - For the distributed deployment, accesses information about the network service and associated VNFs from their respective catalogs, and communicates this information to the Network Service Controller.
4. The VIM receives information from Network Service Orchestrator and:
  - For the centralized deployment:
    - The VIM creates the service chains and associated VMs in the NFVI, which is provided by the servers and Ubuntu. Contrail OpenStack creates one VM for each VNF in the service chain.
    - VNF Manager starts managing the VNF instances while the element management system (EMS) performs element management for the VNFs.
  - For the distributed deployment, Network Service Controller creates the service chains and associated VMs in the NFVI, which is provided by the CPE device.
5. The network service is activated for the customer.

The PNE fits into the NFV model in a similar, though not identical, way to the VNFs.

- For the centralized deployment:
  1. Network Service Orchestrator receives the request through its northbound REST API and sends information about the PNE to PNE/VNF Manager.
  2. PNE/VNF Manager receives information from Network Service Orchestrator and sends information about the PNE to the EMS.
  3. VNF Manager starts managing the VNF instances and the EMS starts element management for the VNFs.
  4. The PNE becomes operational.

- For the distributed deployment:
  1. Network Service Orchestrator receives the request through its northbound REST API.
  2. Network Service Controller receives information from Network Service Orchestrator and starts managing the PNE.
  3. The PNE becomes operational.

#### Related Documentation

- [Contrail Service Orchestration Solutions Overview on page 17](#)
- [Number of Sites and VNFs Supported in Contrail Service Orchestration on page 36](#)
- [VNFs Supported by the Contrail Service Orchestration Solutions on page 37](#)

## Number of Sites and VNFs Supported in Contrail Service Orchestration

CSO supports three environment types: small, medium, and large. The small environment does not include any high availability features. [Table 4 on page 36](#) shows the number of sites and VNFs supported for each environment.

**Table 4: Number of Sites and VNFs Supported**

Deployment Type	Number of VNFs Supported for a Centralized Deployment	Number of Sites and VNFs Supported for a Distributed Deployment	Number of Sites Supported for a Hub and Spoke SD-WAN Deployment
Small	10 VNFs	Up to 500, 2 VNFs per site	Up to 500
Medium	100 VNFs, 20 VNFs per Contrail compute node	Up to 3500, 2 VNFs per site	Up to 3500
Large	500 VNFs, 20 VNFs per Contrail compute node	Up to 6000, 2 VNFs per site	Up to 6000

The following table provides the number of sites and tunnels supported by full-mesh deployments:

**Table 5: Number of Sites, Tenants, and Tunnels Supported for a Full-Mesh SD-WAN Deployment**

Description	Scale
Number of full-mesh DVPN tunnels supported per tenant	50000
Number of full-mesh DVPN tunnels supported across a CSO installation	125000

**Table 5: Number of Sites, Tenants, and Tunnels Supported for a Full-Mesh SD-WAN Deployment (continued)**

Description	Scale
Number of full-mesh tenants qualified across a CSO installation	200 tenants with 10 sites per tenant
Number of full-mesh sites qualified for a given tenant	250 sites
Maximum number of events per second that can be processed by SD-WAN log processing	90000
Number of tunnels supported on NFX250	600 tunnels
Number of tunnels supported on SRX4100 and SRX 4200	1500 tunnels

Each environment has different requirements for:

- The number and specification of node servers and servers. See *Minimum Requirements for Servers and VMs*
- The number and specification of virtual machines (VMs). *Provisioning VMs on Contrail Service Orchestration Nodes or Servers*

**Related Documentation**

- *Minimum Requirements for Servers and VMs*
- *Provisioning VMs on Contrail Service Orchestration Nodes or Servers*

## VNFs Supported by the Contrail Service Orchestration Solutions

Contrail Service Orchestration (CSO) supports Juniper Networks and third-party VNFs listed in [Table 6 on page 38](#).

**Table 6: VNFs Supported by Contrail Service Orchestration**

VNF Name	Version	Network Functions Supported	Deployment Model Support	Element Management System Support
Juniper Networks vSRX	vSRX KVM Appliance 15.1X49-D123	<ul style="list-style-type: none"> <li>Network Address Translation (NAT)</li> <li>Demonstration version of Deep Packet Inspection (DPI)</li> <li>Firewall</li> <li>Unified threat management (UTM)</li> </ul>	<ul style="list-style-type: none"> <li>Centralized deployment</li> <li>Hybrid WAN and SD-WAN deployments supports NAT, firewall, and UTM.</li> </ul>	Element Management System (EMS) microservice, which is included with CSO
LxCIPtable (a free, third party VNF based on Linux IP tables)	14.04	<ul style="list-style-type: none"> <li>NAT</li> <li>Firewall</li> </ul>	Centralized deployment	EMS microservice
Cisco Cloud Services Router 1000V Series (CSR-1000V)	3.15.0	Firewall	Centralized deployment	Junos Space Network Management Platform
Riverbed SteelHead	9.2.0	WAN optimization	Hybrid WAN deployment—NFX250 and NFX150 platforms.	EMS microservice
Fortinet	5.6.3	Firewall	Hybrid WAN and SD-WAN deployments—NFX250 and NFX150 platforms.	EMS microservice
Single-legged Ubuntu	16.04	Firewall	Hybrid WAN and SD-WAN deployments—NFX250 and NFX150 platforms.	EMS microservice

Immediately after installation, CSO does not contain any VNFs. You have to upload the VNFs to the CSO platform through the Administration Portal or through API calls.

You can use these 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.



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

#### Related Documentation

- [Uploading the vSRX VNF Image for a Centralized Deployment on page 113](#)
- [Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 115](#)

- [Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 117](#)





## CHAPTER 2

# Deployment Tools

- [Contrail Service Orchestration \(CSO\) Deployment Tools on page 41](#)
- [Contrail Services Orchestration \(CSO\) GUIs on page 41](#)
- [Designing and Publishing Network Services on page 44](#)
- [Contrail Service Orchestration License Tool on page 45](#)

## 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).

- **CSO License Tool**

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

## 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 7 on page 42](#) for information about logging into the Contrail Service Orchestration GUIs.

**Table 7: Access Details for the GUIs**

GUI	URL	Login Credentials
Administration Portal	<p><code>https://central-IP-Address</code></p> <p>Where:</p> <p><i>central-IP-Address</i>—IP address of the VM that hosts the microservices for the central POP</p> <p>For example:</p> <p><code>https://192.0.2.1</code></p>	<p>Specify the OpenStack Keystone username and password.</p> <p>The default username is <b>cspadmin</b>.</p> <p>After initial install, Specify the autogenerated cspadmin password that is displayed on the console after the installation is complete. You will be prompted to change the random password when you first login.</p> <p>After an upgrade, you must specify the cspadmin password of the previously installed version.</p>
Customer Portal	Same as the URL used to access the Administration Portal	Specify the credentials when you create the Customer either In Administration Portal or with API calls.
Designer Tools—Log into Network Service Designer and click the menu in the top left of the page to access the other designer tools.	<p><code>https://central-IP-Address:83</code></p> <p>Where:</p> <p><i>central-IP-Address</i>—IP address of the VM that hosts the microservices for the central POP</p> <p>For example:</p> <p><code>https://192.0.2.1:83</code></p>	<p>Specify the OpenStack Keystone username and password.</p> <p>The default username is <b>cspadmin</b>.</p> <p>Specify the autogenerated cspadmin password that is displayed on the console after the installation is complete.</p> <p>After the upgrade, you must specify the cspadmin password of the previously installed version.</p>

Table 7: Access Details for the GUIs (continued)

GUI	URL	Login Credentials
<p>Kibana</p> <p>This tool provides a visual representation of log files. You can use it to monitor:</p> <ul style="list-style-type: none"> <li>• Network services in a central or regional POP</li> <li>• Microservices in the deployment</li> </ul>	<p><code>http://infra-vm-IP-Address   ha-proxy-IP-Address:5601</code></p> <p>Where:</p> <p><i>infra-vm-IP-Address</i>—IP address of the VM that hosts the infrastructure services for a central or regional POP. Use this option to monitor network services.</p> <p><i>ha-proxy-IP-Address</i>—IP address of high availability (HA) proxy. Use this option to monitor the microservices.</p> <ul style="list-style-type: none"> <li>• For a deployment without HA, use the IP address of the VM that hosts the microservices for the central POP.</li> <li>• For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO.</li> </ul> <p>For example:</p> <p><code>http://192.0.2.2:5601</code></p>	<p>Starting with CSO Release 4.1, you must login to access the Kibana logs. The login credentials are:</p> <ul style="list-style-type: none"> <li>• Username: <b>admin</b></li> <li>• Password: <b>&lt;Random password generated for Elasticsearch during CSO installation&gt; See the <code>setup_assist</code> file for the password.</b></li> </ul>

## Release History Table

Release	Description
4.1	Starting with CSO Release 4.1, you must login to access the Kibana logs. The login credentials are:

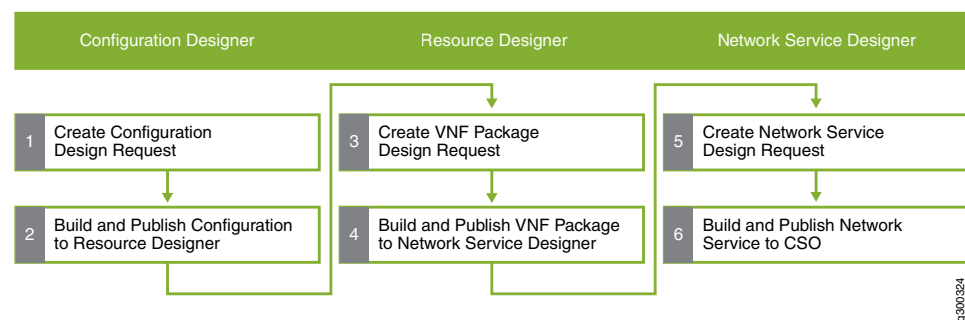
## Related Documentation

- [Designing and Publishing Network Services on page 44](#)
- [Contrail Service Orchestration Solutions Overview on page 17](#)

## Designing and Publishing Network Services

The Contrail Service Orchestration (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 all of the Cloud CPE and SD-WAN solutions. 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 44](#) shows an overview of the workflow used within the Designer Tools application.

Figure 14: Designer Tools Overview



- First, you use *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.
- Next, you use *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 *Network Service Designer* to:
  - Design service chains for network services using the VNF packages that you created with Resource Designer.
  - Configure the network services.
  - Publish network services to the network service catalog.

You use the same process to create network services for centralized CPE, Hybrid WAN, and SD-WAN deployments. You cannot, however, share network services between a centralized deployment and a Hybrid WAN deployment that are managed by one Contrail Service Orchestration installation. In this case, you must create two identical services, one for the centralized deployment and one for the Hybrid WAN deployment. The same is true for SD-WAN deployments, the same network service can not be shared between an on-premise 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.

You can also use *Configuration Designer* to create workflows for device templates.

For detailed information about using the Designer Tools, see the [Contrail Service Orchestration User Guide](#).

#### Related Documentation

- [Contrail Services Orchestration \(CSO\) GUIs on page 41](#)
- [Setting Up a Centralized Deployment on page 88](#)
- [Hybrid WAN Deployment Overview on page 77](#)

## Contrail Service Orchestration License Tool

- [Overview of the License Page on page 45](#)

### Overview of the License Page

SRX and vSRX Series devices can be used in both the distributed cloud CPE and SD-WAN solutions as CPE devices or as cloud or site 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 licensing page is available in the Administration Portal or the Customer Portal by navigating to **Administration > 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.

- vSRX VNFs in a centralized deployment
- The following items in a distributed deployment:
  - vSRX gateway router on an NFX Series device
  - vSRX or SRX Series CPE devices
- vSRX or SRX Series CPE devices in an SD-WAN deployment

#### To upload a license to CSO for later push to an SRX device:

1. Login to CSO as an authorized user—License management is available to both tenant administrators and the global administrator. Operators can not upload licenses to CSO or push them to devices.
2. Navigate to the **Administration > 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 **+** at the top-right part of the list.

This brings up a pop-up window in which you locate and describe the new license file.

4. Click the **Browse** button to locate the license file that was emailed 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 **check box** 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.  
  
If you are logged in as the global administrator, you will see a pull-down list of tenants. Below that will be the sites and devices to which you can push the license files.
3. Select the appropriate device, and click **Push Licenses**. Multiple licenses can be pushed to a single device.

See *Contrail Service Orchestration User Guide* for additional details about the CSO license page.

- See Also**
- [Contrail Services Orchestration \(CSO\) GUIs on page 41](#)
  - [Designing and Publishing Network Services on page 44](#)

## CHAPTER 3

# SD-WAN Deployment

- [SD-WAN Deployment Overview on page 47](#)
- [SD-WAN Deployment Architectures on page 48](#)
- [Your First SD-WAN Deployment on page 59](#)

### SD-WAN Deployment Overview

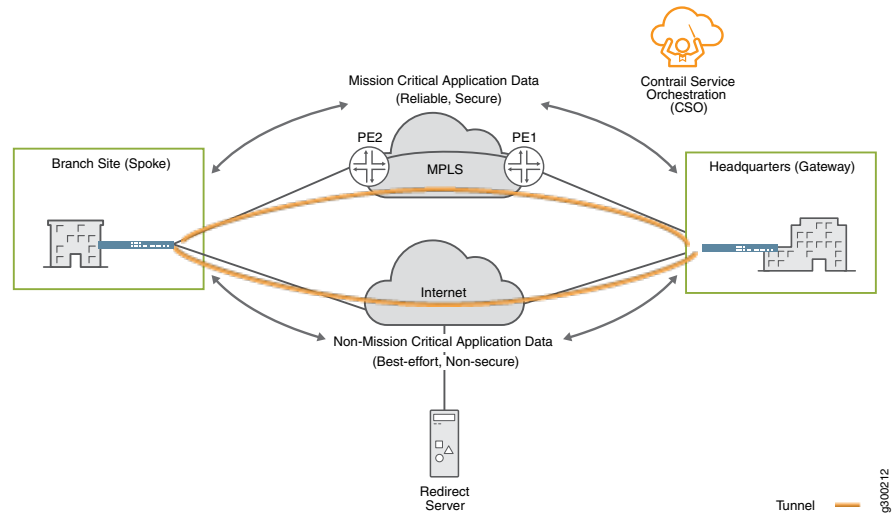
---

This walkthrough highlights the steps, or workflows, that you need to complete in order to deploy an SD-WAN solution using the hub-and-spoke topology with the 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 Hub which is located 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 at each step. It assumes that this is the first deployment you are attempting.

Additional information about using the GUI for any of the steps below can be found in the [Contrail Service Orchestration User Guide](#).

We use the topology shown in [Figure 15 on page 48](#) as a reference for this SD-WAN deployment.

Figure 15: SD-WAN Example Deployment Topology



Before discussing the details of the deployment, we talk about the architecture of SD-WAN deployments in CSO. The architecture discussion helps to identify the specific pieces used in an SD-WAN deployment so that the flow of operations discussed later in [“Your First SD-WAN Deployment” on page 59](#) can be easily understood.

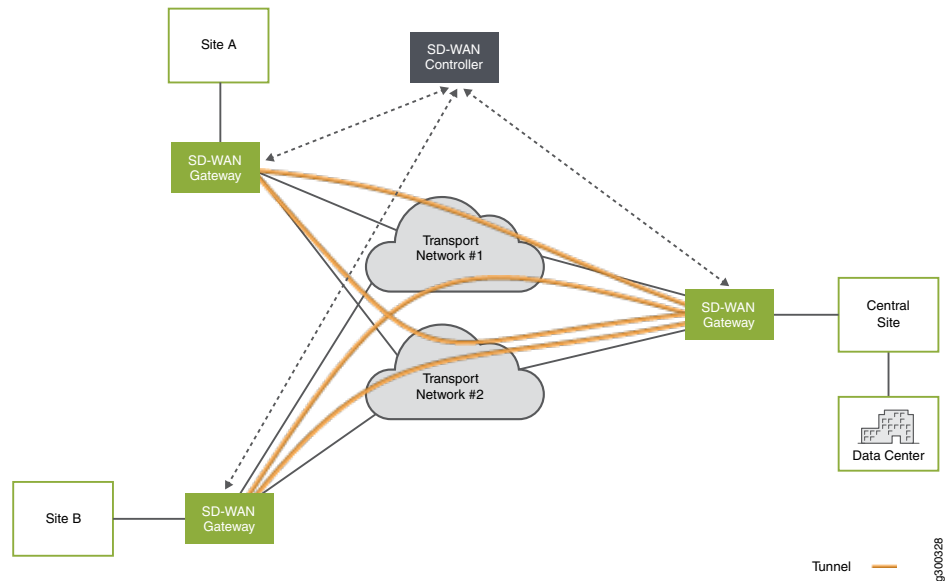
## SD-WAN Deployment Architectures

An SD-WAN implementation offers a flexible and automated way to route traffic from site to site. As shown in [Figure 16 on page 49](#), a basic SD-WAN architecture includes just a few basic elements

- Multiple sites
- Multiple connections between sites that form the underlay network
- A controller
- Multiple overlay tunnels



Figure 16: SD-WAN Architecture

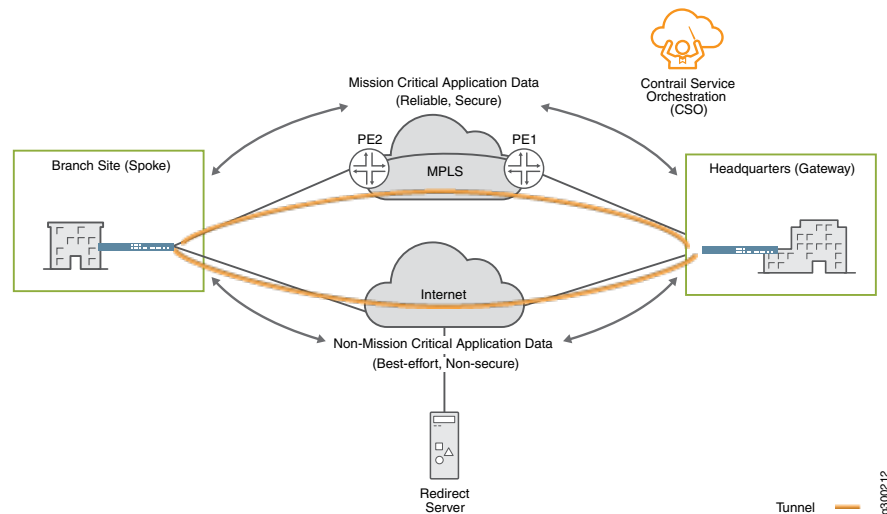


The SD-WAN controller, built-in to CSO, acts as an orchestration layer and provides an interface, allowing the operator to setup and manage the devices at the sites.

### SD-WAN Reference Architecture

Juniper's Contrail SD-WAN solution architecture, shown in [Figure 17 on page 50](#) using a hub-and-spoke topology, is based on the distributed cloud CPE model, with CPE devices located at customer branch sites. On the local side of the site, the CPE devices connect to LAN segments; on the WAN side, the CPE devices connect across two or more links to a hub device. With the SD-WAN model (in a hub-and-spoke topology), traffic travels from site to site through the hub. By default, traffic going to the Internet also flows through the hub device.

Figure 17: SD-WAN Reference Architecture



The SD-WAN orchestrator and controller functions are implemented through Juniper's 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" (GWR), 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-premise spoke and cloud spoke.

### On-premise Spoke Devices

Figure 18: On-premise Spoke Devices



On-premise spoke devices can be either NFX Network Services devices or specific SRX Series Services Gateways. [Table 8 on page 51](#) shows the supported NFX hardware and required Junos OS software release version for each supported model.

### 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 devices eliminate the operational complexities of deploying multiple physical network devices at a customer site.

A key VNF supported on the NFX platform is the vSRX Virtual Firewall. In the Contrail SD-WAN solution, the vSRX instance performs the GWR function, given its routing and switching capabilities. It also provides the same feature-rich security services found on a standard SRX Series Services Gateway.



**NOTE:** The NFX150 includes SRX functionality natively built in.

**Table 8: NFX Hardware and Software Matrix for On-premise Spoke Devices**

Platform	Models Supported	Junos OS Software Release Version
NFX150 Network Services Platform	<ul style="list-style-type: none"> <li>NFX150-S1</li> <li>NFX150-S1E</li> <li>NFX150-C-S1</li> <li>NFX150-C-S1-AE/AA</li> <li>NFX150-C-S1E-AE/AA</li> </ul>	18.2X85-D11
NFX250 Network Services Platform	<ul style="list-style-type: none"> <li>NFX250-LS1</li> <li>NFX250-S1</li> <li>NFX250-S2</li> </ul>	15.1X53-D496.0

#### *SRX Series Services Gateways and vSRX Virtual Firewall*

A physical SRX device can be used in place of the NFX platform to provide the GWR function, as can a vSRX instance installed on a server. [Table 9 on page 51](#) shows the supported SRX hardware and required Junos OS software release version.

**Table 9: SRX Hardware and Software Matrix for On-premise Spoke Devices**

Platform	Models Supported	Junos OS Software Release Version
SRX Series Services Gateway	<ul style="list-style-type: none"> <li>SRX4100 Services Gateway</li> <li>SRX4200 Services Gateway</li> <li>SRX300 Services Gateway</li> <li>SRX320 Services Gateway</li> <li>SRX340 Services Gateway</li> <li>SRX345 Services Gateway</li> <li>SRX550M Services Gateway</li> </ul>	151_X49_D161
vSRX Series Virtual Services Gateway	vSRX	151_X49_D170

*Gateway Sites and Devices*—Starting with CSO Release 4.1, a special type of spoke device, called a Gateway Device, can be deployed as the CPE at an on-premise spoke. Only

SRX4100 and SRX4200 Services Gateway devices can serve this function. The spoke site that functions this way, must be configured as a gateway site during site creation. Creating a gateway site with an SRX4x00 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
- Provides for a new, 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, etc.

#### *Cloud Spoke Devices*

Starting in CSO Release 3.3, an SD-WAN spoke device can be located in an AWS VPC. A vSRX instance in the VPC serves as the cloud spoke device; once the endpoint comes online it acts like any other spoke device.

#### *Spoke Redundancy*

Starting with CSO Release 3.3, dual 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](#).

## Hub Devices (SD-WAN Gateway)

The 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 gateway device. In a hub-and-spoke deployment, there is at least one hub device and one or more spoke devices.

A hub device acts as the SD-WAN gateway, terminating MPLS/GRE and IPsec tunnels from spoke devices. The hub is sometimes referred to as an SD-VPN gateway, given its role as an IPsec concentrator.



**NOTE:** Starting with CSO Release 4.0, the on-premise hub is no longer supported.


---

#### *Cloud Hub*

In a service provider environment, a cloud hub is owned by the service provider and the hub device resides within the service provider's network (POP). It is typically a "shared" device, providing hub functionality to multiple customers (tenants).

In an enterprise environment, the cloud hub is owned by the customer (tenant) and the hub device resides within the enterprise data center. Only the customer's spoke sites can connect to its hub device.

Figure 19: SD-WAN Hub Devices



As of CSO Release 4.1, the supported hub devices are shown in [Table 10 on page 53](#)

Table 10: Hub Devices

Role	Supported Device Types	Required Junos OS Software Version	Usage Notes
Cloud Hub	<ul style="list-style-type: none"> <li>MX Series Devices with Services Line Cards</li> <li>SRX1500 Services Gateway</li> <li>SRX4100 Services Gateway</li> <li>SRX4200 Services Gateway</li> <li>vSRX</li> </ul>	<p>For MX Series Devices: 16.1R5.7</p> <p>For SRX Series Devices: 15.1X49-D161</p> <p>For vSRX: 15.1X49-D170</p>	<p>The requirement for the Services Line Card (MIC or MPC) is there so that the MX can terminate IPsec tunnels.</p> <p>See <a href="#">MPCs Supported by MX Series Routers</a> and <a href="#">MICs Supported by MX Series Routers</a> for information about MX Series routers that support Multiservices MPC and MIC line cards</p>

### Hub Redundancy

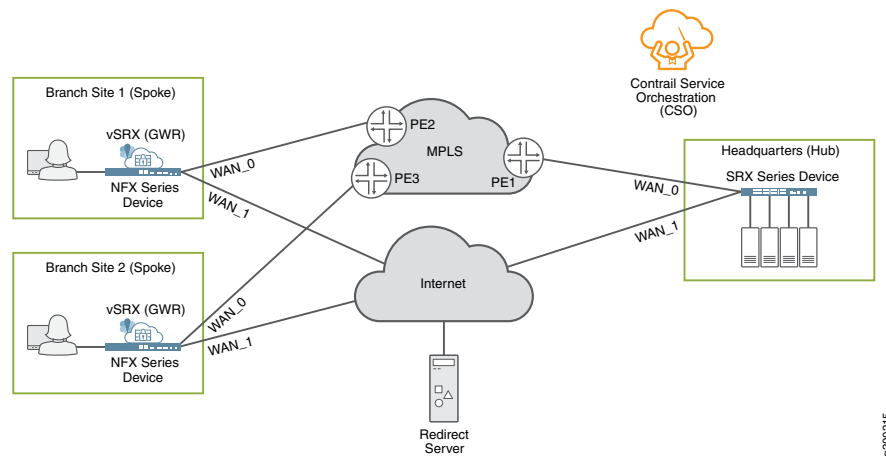
Starting with CSO Release 3.3, dual hub devices can be used to protect against device and link failures, and to provide upstream multihoming for spoke sites. For more detail, see the *Resiliency and High Availability* section of the [Contrail SD-WAN Design and Architecture Guide](#).

## 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-premise devices.

[Figure 20 on page 54](#) shows a sample underlay network for a hub-and-spoke SD-WAN deployment (the details apply equally to a full 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



Starting in CSO release 4.1, each on-premise 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.
- All on-premise devices' MPLS network-facing interfaces must be attached to a single Layer 3 VPN provider.
- The on-premise 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 (LTE WAN access supported using a dongle on NFX250 devices starting with CSO Release 3.3. LTE WAN access supported on NFX150 devices starting with CSO Release 4.0.0.)
- ADSL/VDSL (ADSL/VDSL WAN access supported on NFX devices starting with CSO Release 4.0.0)
- Broadband
- MPLS and broadband
- Satellite link


The NFX150 Series devices include integrated LTE capabilities. They are shipped from the factory with LTE APN settings configured for the entire world. Starting with CSO release 4.1, the LTE APN settings can be localized for the installation region during the ZTP process.

#### *WAN Interface Types - Data and OAM*

The WAN interfaces are used primarily to forward 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-premise device.

Figure 21 on page 55 illustrates these two interface types.

*Figure 21: WAN Interface Types*



Note that:

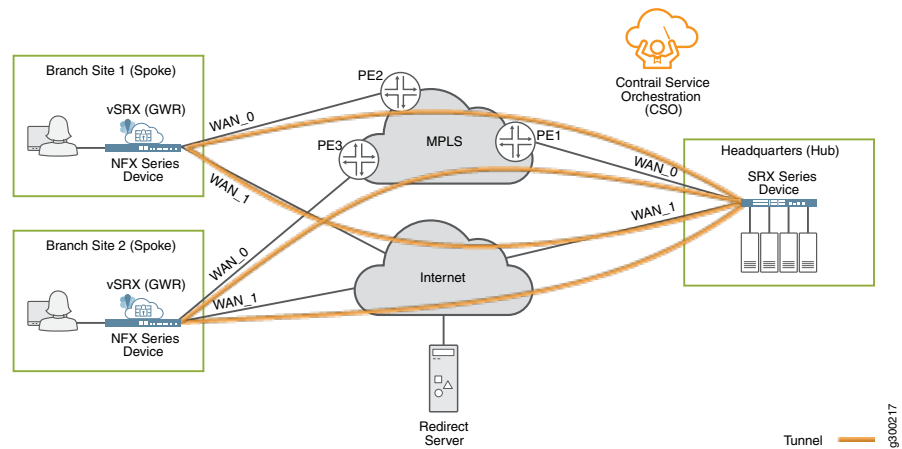
- The on-premise device's OAM interface must be able to reach CSO.
- To ensure secure communication over the WAN, the on-premise device initiates the connection to CSO.
- Device-initiated connections can work across intermediate NAT devices.
- The user-and-OAM-data interface can use a single IP address 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 56 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



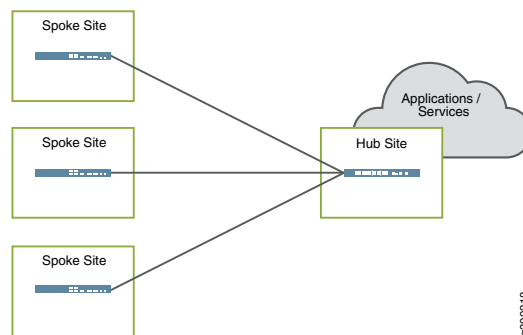
The tunnels have two encapsulation options: MPLSoGRE or MPLSoGREoIPsec. CSO automatically provisions and establishes these tunnels as part of the deployment process.

#### 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.

With the hub-and-spoke topology, all spoke sites are connected to at least one hub site, as shown in [Figure 23 on page 56](#). Spoke sites cannot communicate directly with other spoke sites.

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

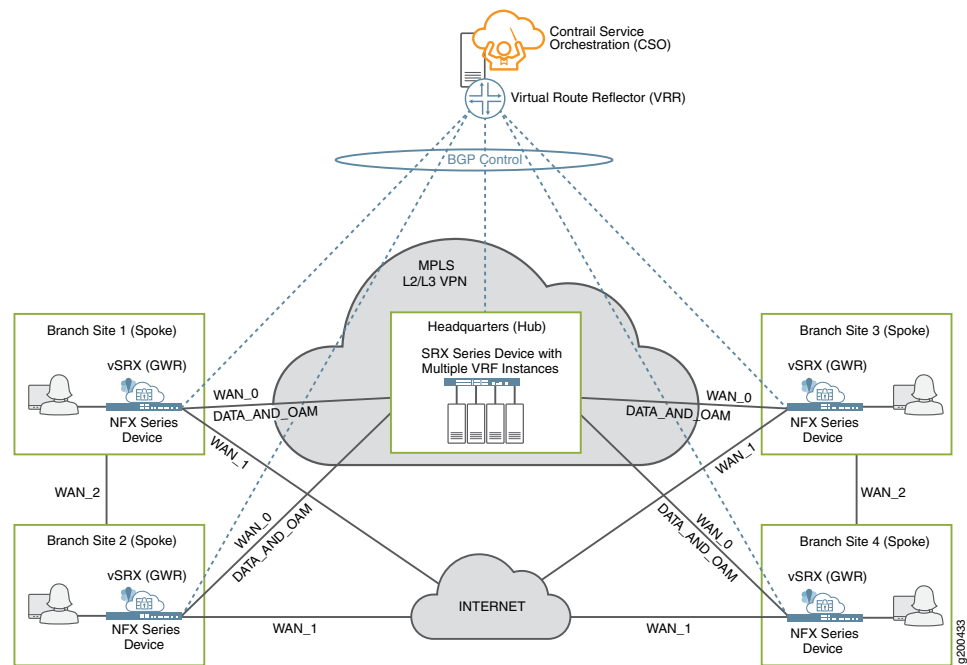


This topology is preferred when applications and services are centralized at the hub site.

With the dynamic mesh topology, all sites are interconnected, as shown in [Figure 24 on page 57](#), and each site can communicate directly with every other site. 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 Full Mesh Topology



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

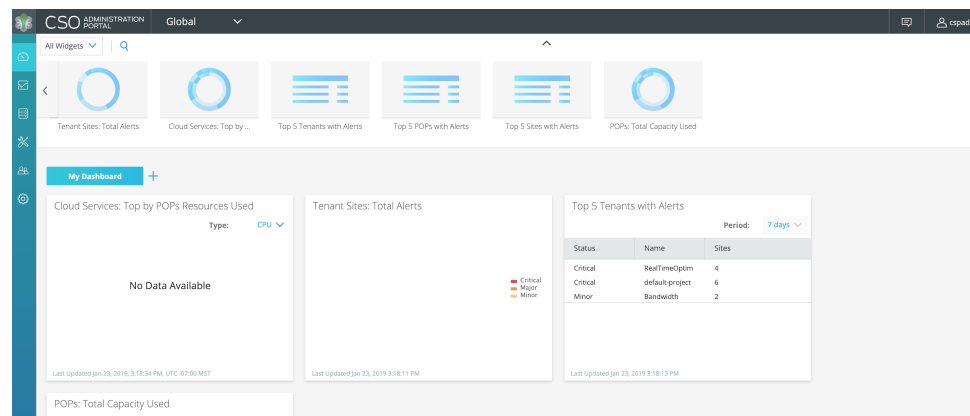


**NOTE:** Starting with CSO Release 4.0, both hub-and-spoke and full mesh topologies require adding a secure OAM network overlay, and thus an OAM Hub, to the deployment.

## SD-WAN Orchestrator/Controller

The SD-WAN 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 58](#).

Figure 25: CSO Dashboard



CSO is built with multiple layers of abstraction for usability and scalability, as shown in [Figure 26 on page 58](#). The platform implements these layers using orchestration software and controller software.

Figure 26: CSO Abstraction Layers



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-premise (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](#).

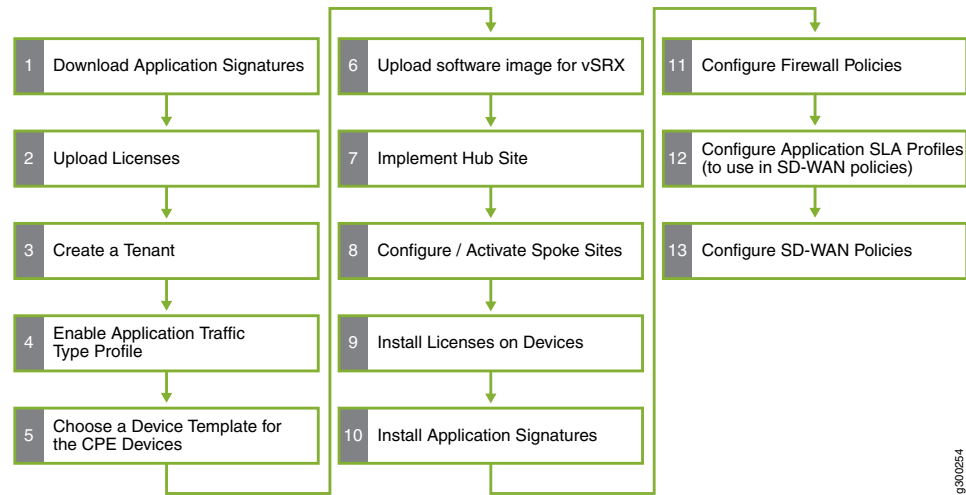
Release History Table

Release	Description
4.1	Starting with CSO Release 4.1, a special type of spoke device, called a Gateway Device, can be deployed as the CPE at an on-premise spoke. Only SRX4100 and SRX4200 Services Gateway devices can serve this function.
4.1	As of CSO Release 4.1, the supported hub devices are shown in <a href="#">Table 10 on page 53</a>
4.1	Starting in CSO release 4.1, each on-premise device (or site) can have up to four WAN links
4.0	Starting with CSO Release 4.0, the on-premise hub is no longer supported.
4.0	LTE WAN access supported on NFX150 devices starting with CSO Release 4.0.0
4.0	ADSL/VDSL WAN access supported on NFX devices starting with CSO Release 4.0.0
4.0	Starting with CSO Release 4.0, both hub-and-spoke and full mesh topologies require adding a secure OAM network overlay, and thus an OAM Hub, to the deployment.
3.3.0	Starting in CSO Release 3.3, an SD-WAN spoke device can be located in an AWS VPC. A vSRX instance in the VPC serves as the cloud spoke device; once the endpoint comes online it acts like any other spoke device.
3.3.0	Starting with CSO Release 3.3, dual CPE devices can be used at spoke sites to protect against device and link failures.
3.3.0	Starting with CSO Release 3.3, dual hub devices can be used to protect against device and link failures, and to provide upstream multihoming for spoke sites.
3.3	LTE WAN access supported using a dongle on NFX250 devices starting with CSO Release 3.3

## Your First SD-WAN Deployment

This document describes the steps required in order to create your first SD-WAN deployment. [Figure 27 on page 60](#) shows an overview of the steps that will be covered in this deployment example.

Figure 27: SD-WAN Deployment Workflow



93010254

- [Before You Begin on page 61](#)
- [Download Application Signatures on page 61](#)
- [Upload Licenses on page 62](#)
- [Create and Configure a New Tenant on page 62](#)
- [Enable Application Traffic Type Profile on page 63](#)
- [Modify Device Templates on page 64](#)
- [Upload Software Image for vSRX on page 65](#)
- [Create a Point of Presence \(POP\) for the Hub Site on page 66](#)
- [Create Cloud Hub Device on page 66](#)
- [Create and Configure the Tenant's Hub Site on page 69](#)
- [Create and Configure a Spoke Site for the Tenant on page 69](#)
- [Install License on Device on page 72](#)
- [Install Application Signature on page 72](#)
- [Add Firewall and NAT Policies to the Topology on page 73](#)
- [Create SD-WAN SLA Profiles and Policies on page 74](#)

## Before You Begin

- Provision your VMs and complete CSO installation according to the steps discussed in [Contrail Service Orchestration Install and Upgrade Guide](#) and



**NOTE:** If you are provisioning your VMs on a KVM-based hypervisor, you must complete the steps in [Creating a Data Interface for a Distributed Deployment](#) prior to provisioning. This step creates a required bridge interface for the VMs to communicate with the CPE devices.

- Purchase an Advanced Policy-based Routing license for a vSRX. You must purchase a license that includes the **appid-sig** feature.
- Download the required vSRX KVM appliance software image from the [Juniper Networks Software Download](#) site. For CSO Release 4.1.0, the required version is **15.1X49-D160**.
- Download the vSRX 15.1X49-D161 qcow2 software image from [www.juniper.net/support/downloads](http://www.juniper.net/support/downloads) to your computer. This software image is needed in order to instantiate the vSRX VNF on the NFX device later in the deployment.



**NOTE:** Make note of the physical interfaces that you select for use throughout this deployment example. These interfaces need to be connected to form the underlay networks over which the data and management traffic will travel.

## Download Application Signatures

This section details how to download application signatures from Juniper onto your CSO installation. Downloading the signature database makes the application signatures available to install on your CPE device after it has been activated in a later step. These signatures are used for application identification within CSO.

From this point on in this deployment example, we assume that your CSO software is installed at 192.168.101.12 and that you know the login credentials for the **cspadmin** user of the Administration Portal.

1. Open your web browser and in the URL field, enter **https://192.168.101.12**
2. Enter the login credentials for the Administration Portal.

By default, the username is **cspadmin** and the password is randomly generated during installation. If this is the first time logging into the Administration Portal, you must set a new password for the **cspadmin** user.

3. 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 newest database is at the top of the list.

4. Click the **Full Download** link under the **Actions** column.

A pop-up window appears that shows the progress of the download. You can watch the progress here or dismiss the window by clicking OK. If you dismiss the progress window before the job completes, you can still access the job information by looking in **Monitor > Jobs**. The download job appears at the top of the list.

Once the download completes successfully, the new database version number appears in the **Active Database** portion of the page.

## Upload Licenses

The licenses that you upload using this procedure are available to be pushed to your tenant devices during the ZTP process.

To upload the license for your vSRX gateway router (GWR) device:

1. Navigate to the **Administration > 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. This brings up a window in which you click the **Browse** button to locate the license file that you purchased for the vSRX.

2. Click the **+** 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 it 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.

## Create and Configure a New Tenant

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

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

2. Click the **Add Tenant** button

If there are no tenants created yet, Add Tenant will be a button. If there are tenants, click the “+” to create 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
- Select the check-boxes next to all three **Roles** in the **Available** section and click the arrow link to move them to the **Selected** section

- Set the User Password to never expire

- Click Next

- In the **Deployment Type** window, select the check-box next to **SD-WAN Sites**

This activates the **SD-WAN Mode** section of the window.

- Select the **Bandwidth Optimized** radio button

- 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.

The preceding steps show only one of many possible settings that can be used to create an SD-WAN tenant

## Enable Application Traffic Type Profile

You can customize class-of-service and probe parameters with traffic type profiles. All traffic type profiles are disabled by default. A maximum of six traffic type profiles can be enabled at one time.

To enable application traffic type profiles:

1. Navigate to the **Configuration > Application Traffic Type Profiles** page.

Here you can see the built-in application traffic type profiles.

Click **OK**.

2. Click the **check box** next to Internet.

3. Click the **Pencil** icon at the upper right part of the list.

In the new window that appears, you can see the parameters that make up this profile.

4. Click the **Enable Toggle Switch** next to **Status**.

5. Click **OK**

This enables the profile for use in an Application SLA Profile that you create later.

## Modify Device Templates

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

1. Navigate to **Resources > Device Templates**

2. Find the device template named **NFX250 as SD-WAN CPE**.

3. 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

4. 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.

- **USE\_SINGLE\_SSH\_TO\_NFX: ON**

Do not change any other settings.

5. Select **Save** when finished.

6. Find the device template named **SRX as SDWAN Hub** and select the check-box next to its name.



7. From the **Edit Device Template** pull-down menu, select **Template Settings**
8. In the **Template Settings** window that appears, make sure the following options are set:
  - **ACTIVATION\_CODE\_ENABLED: Off**
  - **ZTP\_ENABLED: Off**
  - **WAN\_0: ge-0/0/3**
  - **WAN\_1: ge-0/0/1**
  - **WAN\_2: ge-0/0/0**
  - **WAN\_3: ge-0/0/2**Leave all the other settings at their default.
9. Click Save when finished.

## Upload Software Image for vSRX

The NFX appliance that you are using as a CPE will be in factory-default state. Therefore it will not have any vSRX images to instantiate. During the zero touch provisioning (ZTP) process, the NFX downloads the GWR (vSRX) image from CSO.

To upload a software image:

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

Here you can see the software images that have been uploaded to CSO.
2. Click the **+** button to create a new image.

The **Upload Image** page that pops up requires that you fill in all of the fields except Description and Supported Platform.
3. Name the image **vsrx-vmdisk-15.1.qcow2**
4. Select **VNF Image** as the image type.
5. Click **Browse** and select the **.qcow2** software image that you downloaded previously.
6. Select **Juniper** as the Vendor.
7. Select **juniper-vsrx** as the Family.

8. Fill in the Major Version Number, Minor Version Number, and Build Number as **15, 1,** and **X49-D161**, respectively.
9. Click **Upload**. CSO displays a progress window as the file is uploaded.

## Create a Point of Presence (POP) for the Hub Site

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 customers gain access to the CSO Portals and cloud hub devices that are placed within. SPs often place POPs in their network so that they are geographically close to customer sites.

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

Here you can see a list of POPs. If you have not created any POPs, the list is empty.

2. At the top-right part of the list, click the **+** icon to create a new POP.

A pop-up window appears that requires you to enter basic information about the POP such as POP name and Address Information.

3. Give the POP a name that makes sense, like **bay-area-pop**, and enter the appropriate address information. CSO uses this information to place the POP on a map in certain monitoring screens.

4. Click **Next** 4 times to advance through the next 4 screens.

Since we will not be adding devices, virtual infrastructure management (VIMs), or element management systems (EMS) to this POP, we can just advance through these pages until we arrive at the summary page

## Create Cloud Hub Device

A cloud hub device resides in a regional POP within the service provider's network or cloud. Cloud hub devices can be shared amongst multiple tenants through the use of virtual routing and forwarding (VRF) instances configured on the hub itself.

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

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

2. At the top-right part of the list, click the **+** icon to add a cloud hub device.

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

3. Fill in the following information in this window:

- Name the hub something that makes sense, like **Cloud-Hub-1**
- Management Region: **Regional**

There is currently no other option for this.

- POP: **<Select the POP that you just created from the pull-down menu>**
- 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 on-premises CPE devices through this hub device.

- Device Template: **SRX\_as\_SDWAN\_Hub**

Other options for the hub device template also populate the list. The list is built from the Device Templates on the **Resources > Device Templates** page. Multiple tenants can share this hub. There is usually one hub per POP.

4. In the **Configuration** section, select the **Connectivity** tab, and fill in the form as follows:

- Authentication: **Pre-shared Key**

You can choose Public Key Infrastructure if you have the proper certificates set up.

- Loopback IP Prefix: Enter a 32-bit IP address prefix such as **10.10.10.123/32** as the for the CPE device.

Be sure to use an address that works within your network. This address is used for BGP peering. The IP address prefix must be a /32 IP address prefix and must be unique across the entire management network.

- OAM Interface: Enter **ge-0/0/3** as the OAM Interface of the Cloud Hub device.



**NOTE:** The device template that was modified earlier contained interface assignments for WAN\_0 and WAN\_1 interfaces. You must choose an unused interface.

- OAM VLAN ID: **<Leave 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**

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 must be established.

- Click the check box under the WAN\_0 section to enable the WAN\_0 interface of the Hub device.

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

- Link Type: <Leave as MPLS>



**NOTE:** Internet is the other available link type. Since there is usually only one MPLS connection to any given service provider, any other WAN connections that you set up will likely have the link type set to Internet.

- Address Assignment: **Static**
- Static IP Prefix: <Enter an IP address prefix, such as 172.21.22.1/29>

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

- Gateway IP Address: Enter an IP address, such as 172.21.22.2

This is the IP address of the GWR on the NFX250 at the customer site.



**NOTE:** Enable the other WAN interfaces for your Cloud Hub device as appropriate.

- Select the Devices tab
- Under **Device Details**, Enter the serial number of the hub device.

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

A pop-up message tells you that the device is being added. When the add completes, the list refreshes and shows the new cloud hub device in **EXPECTED** state under **Management Status**.

6. Select the **check-box** next to the new cloud hub device

7. Click the **Activate Cloud Hub Device** button at the top right of the list

A new window appears that shows the stages of activation. The stages should flow from EXPECTED to DEVICE\_DETECTED to Stage-one configuration applied to Bootstrap successful to Device Active.

8. Click **Ok**.

The **Activate Device** window closes and your device is listed as PROVISIONED in the **Management Status** column. Once your cloud hub device is in the PROVISIONED state, you can proceed to the next step.

## Create and Configure the Tenant's Hub Site

In this section, we continue in the Customer Portal for your new tenant to create a cloud hub site that will connect with the spoke site that we created in the previous section.

A cloud hub site is the site on the SP's network at which the cloud hub device resides. The cloud hub site is associated with a POP.

Ensure that you are on the **Sites** page in the Customer Portal of your new tenant.

1. From the **Add** pull-down menu on the **Sites** page, select Cloud Hub Site

A new window, titled Add Cloud Site, appears.

2. Fill in the information requested in this window as follows:

- Site Name: <Enter a site name that makes sense for your deployment, such as **cloud-hub-site1**>.



**NOTE:** Site name must match hub device name.

- Cloud Hub Type: **Cloud Hub**
- In the **Address Section**, fill in appropriate address information.  
The fields in this section are optional.
- In the **Contact Information**, fill in appropriate contact information.  
The fields in this section are optional.
- In the **Configuration** section, select the **POP** and **Hub Device Name**

The POP must exist and the hub device must be activated for it to show up in the list.

3. Click **OK** when finished

## Create and Configure a Spoke 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** is displayed

2. Select **Sites** link from the left-nav bar

3. In the **Sites** window that appears, click the **Add On-premise Spoke Site**

A new window titled **Add Site for <Tenant>** appears.

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

The only required information in this window is the site name.

5. Click Next

This brings up the **Connectivity Requirements** section.

6. Under **Connection Plan**, click the left (<) or right (>) arrow until you see the **NFX250 as SD-WAN CPE** box. Click on that box.

This activates the **Connectivity Requirements for the Selected Plan** section.

7. Select the **Enable** check-box next to **Wan\_0**

Fill in the following connectivity requirements

- Type: **MPLS**
- Access Type: **Ethernet**
- Subscribed Bandwidth: **2Mbps**
- Provider: **MPLS-Service-Provider**
- Cost/Month: **1000**

This number is used in SD-WAN link-switch calculations.

- Local Breakout: Off

8. Select the **Enable** check-box next to **Wan\_1**

Fill in the following connectivity requirements

- Type: **MPLS**
- Access Type: **Ethernet**
- Subscribed Bandwidth: **25Mbps**
- Provider: **Internet-Service-ProviderA**
- Cost/Month: **100**

This number is used in SD-WAN link-switch calculations.

- Local Breakout: Off

9. Click Next when finished

The window advances to the **Additional Requirements Section**

10. Select **Wan\_1** in the Default Links field

This setting sets the default forwarding path of all traffic. All traffic leaving the spoke will traverse the WAN\_1 link. WAN\_0 should go mostly unused until SD-WAN policies are created that will cause mission critical data to be sent over the WAN\_0 link.



**NOTE:** If you accidentally select the wrong link as the default, you can remove it from the list by clicking the small 'x' on the left of the link name in the field. You can add the proper link before or after removing the improper one.

11. Click Next when finished

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

A notification appears in this section indicating that you must create at least one LAN segment.

12. Click the Add LAN Segment button

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

Fill in the following information in this window:

- Name: **LAN2**
- Type: Directly Connected
- Ports: **LAN\_2 (ge-0/0/2)**
- VLAN ID: <Leave blank>
- IP Address Prefix: **172.40.1.2/24**

The address shown is just an example. When deploying, use an address prefix that makes sense for the site you are creating.

- Department: <Leave as Default>

In CSO, spoke site departments equate to security zones on the GWR. 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 GWR.

- DHCP: Off

13. Click Save when finished

The **Add LAN Segment** window closes

14. Click Next

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

15. Review the **Summary** section

16. Click **OK** when you're finished reviewing

You will see pop-up messages appear for site creation job start and site creation job finished.

We create the spoke site first so that we can establish the departments (security zones) that will be used by the tenant. We cannot create a hub site until this is determined. If you attempt to create a hub site before creating a spoke site, CSO displays an error.

Likewise, one of the steps in configuring the spoke site is to associate it with a hub. Therefore, we cannot configure the site until after the hub has been created.

## Install License on Device

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

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

In the pop-up window that appears,

2. Click the check box 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 **check box** next to your tenant site to push the license to the CPE device at that site.

## Install Application Signature

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

To install an application signature:

1. Navigate to **Administration > 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 **check box** next to the NFX250 device
4. Click **OK**

## Add Firewall and NAT Policies to the Topology

In this section, we use the Customer Portal for your new tenant and create an intent-based firewall policy that blocks **icmp-ping** traffic.

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

This brings up the **Firewall Policy** page. Here you can see a list of policies. If this is the first time looking at the page, the list is empty.

2. Click the **+** to create a new firewall policy.

This brings up a policy builder window with the **Select Source** area active.

3. Select **Any** from the **ADDR** section of the list.

4. Click the **Action** circle

This brings up a list of actions: Allow, Deny, and Reject.

5. Select **Deny** from the list

This changes the action circle to match the icon for the Deny action.

6. Click the **+** in the Select Destination section

This brings up a list of destinations.

7. Click the **View More Results** link at the bottom of the list

This causes the a slide-out panel to appear on the right side of the screen.

8. Enter **icmp-ping** in the field at the top of the panel

This filters all of the list sections below the field. The numbers to the right of each collapsed section adjust to show the count of **icmp-ping** contained in each section.

9. Click the **>** next to the **Services [SVCS]** section

This shows the **icmp-ping** item within the services section.

10. Click the **check-box** next to **icmp-ping**

This activates the **Check Mark** button at the top of the panel.

11. Click the **Check Mark**

This adds the **icmp-ping** service to the destination list.

12. Click **Save**

This closes the policy builder and shows the new policy in the list.

13. 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.

14. Click **Deploy**

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

The policy can be implemented at any time for any site within this tenant.

## Create SD-WAN SLA Profiles and Policies

In this section, we use the Customer Portal to configure an SD-WAN Application SLA Profile and SD-WAN Policy to specify that Microsoft Outlook traffic should pass over the WAN\_0 overlay link rather than the default link, WAN\_1.

1. Navigate to **Configuration > SD-WAN > Application SLA Profile**

2. Click the **+** to create a new profile

This brings up a **Create SLA Profile** window.

In the new window, fill in the following information

- Name: <Enter a name for the profile, such as: **Internet-SLA**
- Priority: **5**

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

- Traffic Type Profile: **INTERNET**
- Path Preference: **Internet**
- Packet Loss: **Drag the slider to 20%**
- RTT: **20**
- Jitter: **10**
- Throughput: **5**

3. Click **OK**

This causes the window to close. The new policy shows 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 **+** at the upper right part of the list to create a new policy

This brings up a policy builder with the **Source** section activated.

The Source defaults to **All Sites**. The Application section defaults to **Any**.

Leave these at their default.

6. Click **+ Select SLA Profile**

This brings up a list of available profiles.

7. Select **Internet-SLA** from the list.

8. Click **Save**

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

9. 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.

10. Click **Deploy**

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



## CHAPTER 4

# Distributed CPE Deployment (uCPE)

- [Hybrid WAN Deployment Overview on page 77](#)
- [Hybrid WAN \(Distributed\) Deployment Architecture on page 78](#)
- [Your First Hybrid WAN \(Distributed\) Deployment on page 80](#)


### Hybrid WAN Deployment Overview

---

This walkthrough highlights the steps, or workflows, that you need to complete in order to deploy a Hybrid WAN solution. We use an NFX250 Series device as the CPE and an SRX Series device as the Hub which is located 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 at each step. It assumes that this is the first deployment you are attempting.

In the distributed deployment, customers access network services from a CPE device located at the customer's site. These sites are called on-premise sites in this documentation. In the deployment workflows used in the CSO GUI, this deployment is known as Hybrid WAN. [Figure 28 on page 77](#) illustrates a simplified distributed deployment.

*Figure 28: 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 30kbps 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 105kbps and 2Mbps, 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 5Mbps only when device image download is needed. This makes these solutions applicable even in cases where connection bandwidth is limited or noisy.

The distributed CPE 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 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-premise 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 (Distributed) Deployment Architecture

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

Figure 29 on page 78 shows a simple diagram of the distributed CPE solution. The cloud represents the service provider network to which the customer site is connected.

*Figure 29: Distributed CPE*



As mentioned previously, the distributed Cloud CPE deployment makes use of on-premises CPE devices in order to localize the delivery of network services and provide gateway router (GWR) functionality. In this case, the Juniper Networks NFX Series or SRX Series devices act as the CPE devices. In the case of NFX as CPE, the GWR function is provided by a built-in vSRX VNF and network services are hosted and provided from within the NFX 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, 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.

The distributed Cloud CPE deployment also makes use of a provider edge (PE) router in the service provider cloud. The PE router acts as a IPsec concentrator, terminating IPsec tunnels, and a PE router, providing 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.

**Table 11: Hardware and Software Matrix for CPE Devices in a Hybrid WAN Deployment**

Role	Platform	Models Supported	Junos OS Software Release Version
PE Router and IPsec Concentrator (Hybrid WAN deployment only)	MX Series 3D Universal Edge Router	<ul style="list-style-type: none"> <li>MX960, MX480, or MX240 router with a Multiservices MPC line card</li> <li>MX80 or MX104 router with Multiservices MIC line card</li> <li>Other MX Series routers with a Multiservices MPC or Multiservices MIC line card</li> </ul> <p>See <a href="#">MPCs Supported by MX Series Routers</a> and <a href="#">MICs Supported by MX Series Routers</a> for information about MX Series routers that support Multiservices MPC and MIC line cards.</p>	Junos OS Release 16.1R3.00
CPE device (Hybrid WAN deployment) or spoke device (SD-WAN implementation)	<ul style="list-style-type: none"> <li>NFX Series Network Services Platforms</li> <li>SRX Series Services Gateways</li> <li>vSRX on an x86 server</li> </ul>	<ul style="list-style-type: none"> <li>NFX250-LS1 device</li> <li>NFX250-S1 device</li> <li>NFX250-S2 device</li> <li>NFX150-S1 device</li> <li>NFX150-S1E device</li> <li>NFX150-C-S1 device</li> <li>NFX150-C-S1-AE/AA device</li> <li>NFX150-C-S1E-AE/AA device</li> <li>SRX300 Services Gateway</li> <li>SRX320 Services Gateway</li> <li>SRX340 Services Gateway</li> <li>SRX345 Services Gateway</li> <li>SRX4100 Services Gateway</li> <li>SRX4200 Services Gateway</li> <li>vSRX</li> </ul>	<p><i>For NFX250:</i> Junos OS Release 15.1X53-D496</p> <p><i>For NFX150:</i> Junos OS Release 18.2X85-D11</p> <p><i>For SRX Series:</i> Junos OS Release 15.1X49-D161</p>

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

## Your First Hybrid WAN (Distributed) Deployment

---

- [Install Junos Software onto NFX from USB Port on page 80](#)
- [Modify Device Templates on page 83](#)
- [Create and Configure a New Tenant on page 84](#)
- [Create and Configure a Site for the Tenant on page 85](#)

### Install Junos Software onto NFX from USB Port

This section details how to install Junos OS software version 15.1X53-D496.0 onto an NFX250 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. This procedure is for an NFX250 device.

#### Before You Begin

---

In order for this procedure to succeed, you must have the following

- Physical access to the USB port of the NFX device
- A USB drive of at least 4GB containing the Junos OS Software image, 15.1X53-D496.0, inserted into the USB port of the NFX
- Access to the console port of the NFX 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 NFX250. This DHCP server must be able to provide IP address, name server, and default gateway to the NFX upon request.

The following procedures contain comments that are added to clarify the steps that are discussed.

1. Ensure that the USB drive containing the Junos OS software image is inserted in the USB port of the NFX device.

This allows you to boot the NFX from the USB drive.

2. Access the NFX 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 device.
4. Power on the NFX 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 will install 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 is made up of multiple components that load and boot in a specific order. See [NFX 250 Overview](#) for details. The PFE of the NFX may take a few minutes to complete the boot and allow the jsxe0 interface to obtain its address from DHCP.

You can login to the console of the NFX as **root** 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 for different purposes. You are looking 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) in order 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
Pseudo CB 0
Routing Engine 0
FPC 0         REV 04    650-066113   DXXXXXXXXXX3   RE-NFX250-S2
CPU
PIC 0         REV 04    BUILTIN     BUILTIN        FPC CPU
Base-T-2x1G SFP-
```

```

Power Supply 0
Fan Tray 0                fan-ctrl-0 0, Front
to Back Airflow - AF0
Fan Tray 1                fan-ctrl-0 1, Front
to Back Airflow - AF0

```

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.

## Modify Device Templates

From this point on in this deployment example, we assume that your CSO software is installed at 192.168.101.12 and that you know the login credentials for the **cspadmin** user of the Administration Portal.

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

1. Open your web browser and in the URL field, enter **https://192.168.101.12**
2. Enter the login credentials for the Administration Portal.

By default, the username is **cspadmin** and the password is randomly generated during installation. If this is the first time logging into the Administration Portal, you must set a new password for the **cspadmin** user.

3. Navigate to **Resources > Device Templates**
4. Find the device template named **NFX250 as Managed Internet CPE**.
5. 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

6. 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\_Oge-0/0/11

Do not change any other settings.

7. Select Save when finished.

## Create and Configure a New Tenant

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

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

2. Click the **Add Tenant** button

If there are no tenants created yet, Add Tenant will be a button. If there are tenants, click the "+" to create 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
- Select the check-boxes next to all three **Roles** in the **Available** section and click the arrow link to move them to the **Selected** section
- Set the User Password to never expire

If needed, you can configure password expiry rules here.

- Click Next
- In the **Deployment Type** window, select the check-box next to **Hybrid WAN Sites**
- 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.

## Create 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** is displayed

2. Select **Sites** link from the left-nav bar

3. In the **Sites** window that appears, click the **Add Spoke Site - Hybrid**

A new window titled **Add Site for <Tenant>** appears.

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

The only required information in this window is the site name. 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 in some of the monitoring windows.

5. Click Next

This brings up the **Connectivity Requirements** section.

6. Under **Connection Plan**, click the left (<) or right (>) arrow until you see the **NFX250 as MAnaged Internet CPE** box. Click on that box.

This activates the **Connectivity Requirements for the Selected Plan** section.



**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.

7. Click Next when finished

The window advances to the **Summary**

8. Review the **Summary** section

9. Click **OK** when you're finished reviewing

You will see pop-up messages appear for site-creation job start and site-creation job finished.

10. Click the **check-box** next to the site you just created

11. Click the **Configure Site** button

This brings up a new window titled Configure Site <site-name>.

12. In the Configuration Section, click the **Advanced Config** tab.

On this tab, fill in the following information:

- **Name Server IP List:** <Click the pull-down menu, if no results are found, **enter the IP address of a DNS name server**>.

This is a required field.



**NOTE:** You must press enter when you have completed the IP address entry. If you don't the entry will be lost.

- **Ntp server IP List:** <This is an optional field. However, it is a good idea to enable NTP whenever possible. **Enter the IP address of an NTP server**>.
- **Select timezone:** <This is an optional field. However, it is a good idea to set this to the appropriate time zone. **Select the appropriate time zone for this site**>.

13. Click the **Devices** tab

On this tab, fill in the following information:

- **Serial Number:** <Enter the serial number of your NFX250 device>

This is a required field.

- **Activation Code:**

We create the spoke site first so that we can establish the departments (security zones) that will be used by the tenant. We cannot create a hub site until this is determined. If you attempt to create a hub site before creating a spoke site, CSO displays an error.

One of the steps in configuring the spoke site is to associate it with a hub. Therefore, we cannot configure the site until after the hub has been created.

## CHAPTER 5

# Centralized CPE Deployment (vCPE)

- [Centralized Deployment Overview on page 87](#)
- [Setting Up a Centralized Deployment on page 88](#)


### Centralized Deployment Overview

---

The Cloud CPE Centralized Deployment Model (centralized deployment or vCPE) differs from the other deployments discussed in this guide in that the centralized deployment requires that the service provider install and maintain a Contrail Cloud instance in their network. This Contrail Cloud instance is needed in order to host the VNFs and the CSO installation.

In the centralized deployment, customers access network services remotely from a service provider's cloud. Sites that access network services in this way are called service edge sites in this documentation. [Figure 30 on page 87](#) illustrates a simplified centralized deployment.

*Figure 30: Simplified Centralized Deployment*



The following sections provide a high-level look at the Centralized Deployment Architecture and a walkthrough of one possible way to set up CSO for Centralized Deployment. *Appendix A* presents the details of a Centralized Deployment Reference Architecture.

### Centralized Deployment Architecture Overview

CSO's Centralized Deployment Model uses some of the architectural elements used by the other deployment models, but not all of them. Due to its centralized nature, this model doesn't support:

- CPE devices at remote sites as used in Hybrid WAN and SD-WAN deployments

- Overlay networks as used in SD-WAN deployments
- VNFs hosted anywhere outside of the Contrail implementation inside the SP cloud

Now that we have mentioned what is not supported, we'll tell you what architectural elements are supported.

In the Centralized deployment, you need to have:

- A Contrail implementation
- A provider edge (PE) router that provides access to the SP cloud for the remote sites
- A POP, either central or regional in which the PE router resides
- A specific Virtual Infrastructure Manager (VIM)
- (Optional) An Element Management System (EMS)

Generally, the connection from the remote site to the Service Provider cloud can be over any transport so long as a connection can be made. Any L2 or L3 connectivity works. We recommend some sort of VPN connection in order to secure the connection from bad actors. Juniper supports the use of an MX Series router as the provider edge (PE) device. Use of a Services line card is required if using an MX Series router to terminate IPsec VPN traffic. [Table 12 on page 88](#) shows the hardware and software that can be used as the PE router in a centralized deployment

**Table 12: Hardware and Software Matrix for the PE Router in the Centralized Deployment Model**

Role	Platform	Models Supported	Junos OS Software Release Version
PE Router	MX Series 3D Universal Edge Router	<ul style="list-style-type: none"> <li>• MX960, MX480, or MX240 router with a Multiservices MPC line card</li> <li>• MX80 or MX104 router with Multiservices MIC line card</li> <li>• Other MX Series routers with a Multiservices MPC or Multiservices MIC line card</li> </ul> <p>See <a href="#">MPCs Supported by MX Series Routers</a> and <a href="#">MICs Supported by MX Series Routers</a> for information about MX Series routers that support Multiservices MPC and MIC line cards.</p>	Junos OS Release 16.1R3.00

The SP network is usually an MPLS network. CSO is deployed as part of the Contrail implementation either in the SP Data Center or elsewhere in the cloud.

## Setting Up a Centralized Deployment

Before you set up a centralized deployment, complete the following tasks:

- Configure a Contrail Cloud installation. See [Appendix A](#) for details.
- Install Contrail Service Orchestration. See [CSO Installation and Upgrade Guide](#) for details.
- Upload VNF images. See the following topics for details:
  - [Uploading the vSRX VNF Image for a Centralized Deployment on page 113](#)



- [Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 115](#)
- [Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 117](#)

- Install VNF licenses.

You can use the **Licenses** page to install vSRX licenses. See [“Contrail Service Orchestration License Tool” on page 45](#) for details.

- Publish network services with Network Service Designer. See [“Designing and Publishing Network Services” on page 44](#) for details.
- Access Contrail Cloud and add the following rule to the default security group in the Contrail project.

```
Ingress IPv4 network 0.0.0.0/0 protocol any ports any
```

In this Centralized deployment example, you will use the Contrail Service Orchestration (CSO) platform to instantiate a centralized CPE solution to provide the customer with Internet access. As part of this solution, you will provide basic firewall functionality and NAT the outgoing customer traffic to the interface address of the VNF.

To do that, you will:

- Create a Network Service using Network Service Designer
- Create a POP
- Create a Tenant
- Create a Tenant Site
- Deploy a VNF to the tenant site

In this deployment example, we assume that your CSO software is installed at 192.168.101.12 and that you know the login credentials for the **cspadmin** user of the Administration Portal. This ensures that you have full access to both the Administration Portal and the Customer Portal of your CSO installation. If policy doesn't allow you to have the cspadmin user credentials, then you will need to have Administration Portal credentials with Tenant Admin, Tenant Operator, and Configure Sites roles enabled. You will also need a Customer Portal login so that you can create sites.

To set up a centralized deployment.

## Create Network Service

In this part of the deployment, you create a network service using CSO Designer Tools. For an overview about using Designer Tools, see [“Designing and Publishing Network Services” on page 44](#). You access the CSO Designer Tools at the same URL as the CSO Administration Portal, but on port 83. For example, if the IP address of the Administration Portal is 192.168.101.12, then the URL for Designer Tools would be: **https://192.169.101.12:83**

1. Login using the **cspadmin** or equivalent credentials.

2. Click on the **New Request** button

The window is replaced with a multi-tabbed form that you fill out to complete the request.

3. On the **Request Information** tab, fill in the following information:

- Name: **nat-vnf**
- Deployment Type: **vCPE-Only**

The rest of the fields on the page can be left blank. However, as you create more and more services, it will become useful to fill in at least some of the information for each request.

4. Click **Next**

The page advances to the **Service Chain and Design Goals** tab.

5. Drag the **NAT** building block from the **Function Palette** to the **Functional Service Design** area

The NAT building block will stick in the **Functional Service Design** area.

6. Click on the **+ Add Goal** link on the left side of the window

A pop-up window, titled **New Goal** appears.

7. Select **Session** from the pull-down type menu of the **New Goal** window.

This expands the **New Goal** window to show multiple fields including Goal Value, Acceptable Value, and Must Value.

8. Enter **100** in the **Goal Value** field.

Leave the other fields blank for this example.

9. Click **Save**

10. Click **Next**

This advances the window to the **Summary** page.

11. Review the summary and click **Create**

This resets the window and selects the requests tab on the left. Your new request appears in the main part of the window.

12. Click the **Begin** button once it becomes available.

This advances the window to the **Build** tab with the **Functional Service Design** (NAT) section is at the top and **Network Service Design** is at the bottom of the window.

13. Click on the **NAT** icon in the **Functional Service Design** section.

This changes the icon from grey to colored.

14. From the right portion of the **Network Service Design** section, find the **vSRX** box, drag it to the left, and drop it on the **Network Service Design** pane.

This allows the vSRX service to be placed in a service chain.

15. Click on the **I** icon in the **Network Service Design** banner that splits the page.

This pulls down a menu with **I** and **E** icons on it. The **I** stands for ingress and the **E** stands for egress. You will connect these icons to the service chain in the next step.

16. Click the **I** icon and then click the small circle on the left side of the **vSRX** service box in the **Network Service Design** pane.

This attaches an ingress point to the vSRX VNF.

17. Click the **E** icon and then click the small circle on the right side of the **vSRX** service box in the **Network Service Design** pane.

This attaches an egress point to the vSRX VNF.

18. Click the **Management** tab at the bottom of the window.

Ensure that there is an arrow pointing to an **M** icon. This indicates that the management interface is added.

19. Click on the **Functional Configuration** button.

This is a grey bar at the right side of the window. This pops up a new window with a title that ends in **vnf**.

20. The **Basic Configuration** tab is selected. Configure the following:

- DNS Servers: <Enter the IP address of a known DNS server such as **8.8.8.8**>
- NTP Servers: <Enter the IP address of a known NTP server on your network>  
**10.210.8.72**
- Select the **NAT** tab
- NAT Source Address: <Enter an address prefix for source NAT traffic, like  
192.0.2.0/24>
- NAT Destination Address: <Enter an address prefix for destination NAT traffic, like  
172.16.24.0/24>
- Click **OK** to close this window.

21. Back in the **Build** window, click the **Save** button

The save button is shaped like an old-fashioned floppy disk.

22. Click the **Publish** icon

The publish icon looks like a cloud with an up-arrow inside of it. Clicking it brings us a **Publish NSD** window.

23. Click the **Publish** button in the **Publish NSD** window.

24. Click the **Cspadmin** logo at the top right of the window and select **Logout** from the pull-down menu.

You can now close the browser or the tab in the browser for the Network Service Designer.

## Create POP

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 customers gain access to the network services that are deployed within. SPs often place POPs in their network so that they are geographically close to customer sites.

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

Here you can see a list of POPs. If you have not created any POPs, the list is empty.

2. At the top-right part of the list, click the **+** icon to create a new POP.

A pop-up window appears that requires you to enter basic information about the POP such as POP name and Address Information.

3. Give the POP a name that makes sense, like **east-region-pop**, and enter the appropriate address information. CSO uses this information to place the POP on a map in certain monitoring screens.

4. Click **Next** twice

This advances the window past the **Devices** page to the **VIM** page. VIMs are virtual infrastructure managers. Since the Centralized Deployment requires the use of Contrail Cloud, you must create a VIM for each POP that you create.

5. In the **Add VIM** area, click the **+** to add a new VIM.

This brings up a new window titled **Add Cloud VIM**.

6. In the **Add Cloud VIM** window, fill in the following information:

- Name: <Enter a name that makes sense for this VIM, like **contrail-cloud**.

#### Connection Information

- IP address: <Enter the IP address of the Contrail Controller Node in the Contrail Cloud Platform, such as: **192.168.10.225**>
- Auth URL: <Enter the authentication URL for the Contrail OpenStack Keystone, such as: **http://192.168.10.225:35357/v3/**>
- User Name: <Enter the user name for logging into CSO, such as: **tenantadmin**>
- Password: <Enter the password for the user above>
- Domain: <Specify the name of the Contrail OpenStack domain that you configured for the Contrail Cloud Platform, such as: **Default**>
- Tenant: <Specify the name of the Contrail OpenStack tenant that you configured for the Contrail Cloud Platform, such as: **admin**>

#### Network Information (Resource Pools)

- Click the **+** above the resource pools list to add a resource pool

Fill in the following information in the resource pool:

- Pool Name: **internal**
- Compute Zone: **nova**
- Click the **check mark** to save the entry
- Does Management Network Exist: <**Yes** or **No**>

Select Yes to use an existing network in Contrail or select No to create a new network in Contrail.

- Management Network Name: <Enter the name of the management network in Contrail, such as: **mgmt-network**>

Specify the name of the existing management network in Contrail or the new management network that you want to create in Contrail.

- Click the **check-mark** to save the resource pool.

### Internet Network Information

- Click the **+** above the list of Internet Networks  
Fill in the following information in the Internet Networks
- Network Name: **public**
- Exists: **No**
- Route Target: <Click Edit and then enter an appropriate route target like: **64512:10000**>
- Subnet: <Click Edit and then enter an appropriate subnet like: **172.40.5.0/24**>
- Click the **check-mark** to save the Internet Network Information.

### Service Profile Information

- Click the **+** above the list of service profiles  
Fill in the following information in the Service Profile
  - Profile Name: <Enter a profile name that makes sense for this deployment, such as: **first-profile**>
  - Tenant Name: <Enter a tenant name that makes sense for this deployment, such as: **Tenant1**>
  - Domain Name: <Enter a Contrail domain name that makes sense for this deployment, such as: **Default**>
  - User Name: <Enter the CSO Tenant Admin user name, or equivalent, such as: **cspadmin**>
  - Password: <Enter the password for the user above>
  - Click the **check-mark** to save the resource pool.
7. Click **Save** to complete the VIM configuration.  
This clears the VIM create window and returns to the **Add POP** window.
  8. Click **Next** twice to advance to the **Summary** tab of the **Add POP** window  
Review the summary information to confirm that it is what you intended.
  9. Click **OK** to finish creating the POP  
You will see notification pop-ups that inform you when the job is started and when it's finished.

## Add Tenant

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

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

2. Click the **Add Tenant** button

If there are no tenants created yet, Add Tenant will be a button. If there are tenants, click the “+” to create 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
- Select the check-boxes next to all three Roles in the Available section and click the arrow link to move them to the Selected section
- Set the User Password to never expire
- Click Next
- In the **Deployment Type** window, select the check-box next to **Hybrid WAN**
- Click Next

The window advances to the **Tenant Properties** section.

- Expand the **Service Profiles (Optional)** section by clicking on the > icon or the gear icon.
- Click the + to add a service profile

From the **VIM Name** pull-down, select the VIM that you created in the previous section.

From the **Service Profile Name** pull-down, select the service profile that you created in the previous section.

- Click **Save**
- 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.

## Allocate Network Service

Return to the list of tenants in the **Tenants** window. In this section, we will allocate network services to your new tenant.

1. In the list of tenants, there is an **Assigned Services** column. Click the link **Allocate Network Services**

This brings up a new window in which you can see all the available network services. The network service **nat-vnf** should be listed there based on the work you did in the Create Network Service section.

2. Click the **check-box** next to the **nat-vnf** network service and then click **OK**.

The number of assigned services will change from 0 to 1.

## Create Cloud Site

There are 3 types of cloud sites available: Local Service Edge, Regional Service Edge, Hybrid Spoke Site. For this deployment, we will create a Regional Service Edge site. See , , and in the *CSO User Guide* for more information about these site types.

1. Click on the link that is named for the tenant that you just created.

This takes you to the CSO Customer Portal.

2. Navigate to the **Sites** page by clicking **Sites** tab on the left-nav bar.

3. Click the **Add Regional Service Edge Site** button

This brings up a new window titled **Add Regional Service Edge Site**. Fill in the following information:

- Site Name: <Enter a name for the site that makes sense to you, like **New-York1**>
- Fill out the **Address** and **Contact Information** sections as appropriate.

None of this information is required but is used in monitoring and alerting functions in CSO.

- In the **Configuration** section
  - Service POP: <Select the recently created POP from the pull-down menu, like **east-region-pop**>
  - VIM: <Select the recently created VIM from the pull-down menu, like **contrail-cloud**>
  - Resource Pool: <Select the recently created resource pool from the pull-down menu, like **internal**>
  - Route Target: <Enter a route target for the virtual network, like **64512:1**>
  - Virtual Network Name: <Enter a unique string of alphanumerics and special characters, such as **CustomerA-VNet**>

This network name is a representation of your network in the cloud

- Left Subnet Prefix: <Select one or more IP prefixes from the list>
- In the **Service Attachment Points** section, set Local Internet Breakout: **ON**

4. Click **OK**

You will see messages pop up indicating the start and stop of the Site Creation Job. Wait for the job to complete successfully.

5. Click on the Site Name Link (New-York1)



This brings up the site-specific window for New-York1.

6. Click the **Services** tab

7. Click on the attachment point between the Site and Local Breakout icons.

This brings up the **Deploy Network Services** menu.

8. Drag the **nat-vnf** network service onto the attachment point and drop it there.

This brings up the **Service** window.

9. In the **Service** window, click the **Basic Configuration** tab.

Fill in the following information on this page:

- Host Name: **vsrx-nat-vnf**
- Loopback Address: <Enter an IP address for the loopback interface, such as: **192.168.100.100**>
- DNS Server: <Enter the IP address of a DNS server such as: **8.8.8.8**>
- NTP Server: <Enter the IP address of an NTP server such as: **10.210.8.72**>
- Enable Default Screens: **Disable**
- Ping Prefix List: **0.0.0.0/0**

10. Click on the **NAT** tab

Enter the following information on this page:

- Nat Source Address: **0.0.0.0/0**
- Policy Name: <Enter a policy name that makes sense, like: **nat-outgoing**>
- Source Zone: **left**
- Destination zone: **right**
- Source Address: **any**
- Source Address: **any**
- Action: **permit**
- Application: **any**

11. Click **OK**

This takes you back to the Site specific window.

12. Click the **Start Service** button to deploy the VNF

This causes a pop-up confirmation window to appear.

13. Click **OK** in the confirmation window

The deployment status shows the percentage of completion for the deployment job. Wait for the status to reach 100%. This can take 10 minutes or more.

You can verify the effects of your changes by logging in to the CLI of the various devices and confirming:

- Learned BGP routes in each table for the local network that is attached through the GRE tunnel to the vSRX VNF.
- MPLS label for traffic from **CustomerA-VNet** instance to the left interface of the VNF is using one label. and
- Traffic in **mgmt** instance is using another label
- Traffic in the **public** instance to the right interface is using a third label.
- If you login to the host-hub device and leave a ping running to the DNS name server at 8.8.8.8, you can then login to the vSRX VNF and see security flows by using the **show security flow session** command.

**Related  
Documentation**

- [Contrail Services Orchestration \(CSO\) GUIs on page 41](#)
- [Designing and Publishing Network Services on page 44](#)

## CHAPTER 6

# Appendix A – Contrail Cloud Reference Architecture for Centralized Deployment

- [About this Reference Architecture on page 99](#)
- [Architecture of the Contrail Cloud Implementation in the Centralized Deployment on page 99](#)
- [Cabling the Hardware for the Centralized Deployment on page 103](#)
- [Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 106](#)
- [Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 107](#)
- [Configuring the MX Series Router for the Contrail Cloud Implementation in a Centralized Deployment on page 109](#)
- [Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment on page 111](#)

### About this Reference Architecture

---

The following sections discuss a reference architecture that could be used to deploy Contrail Cloud and CSO in the SP cloud. It is only meant as an example. Specific device ports have been called out in the examples. However, there are no specific requirements for using the exact devices, software versions, or ports in your Centralized Deployment.

### Architecture of the Contrail Cloud Implementation in the Centralized Deployment

---

This section describes one possible architecture of the components in the Contrail Cloud implementation used in the centralized deployment.

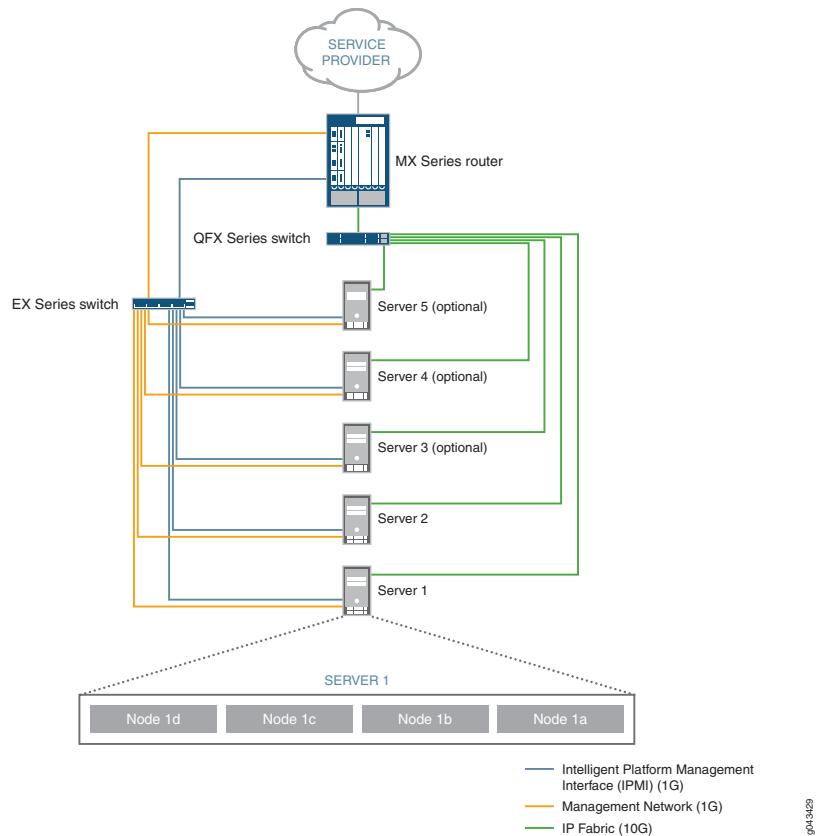
- [Architecture of the Contrail Cloud Implementation on page 99](#)
- [Architecture of the Servers on page 101](#)
- [Architecture of the Contrail Nodes on page 102](#)

### Architecture of the Contrail Cloud Implementation

The centralized deployment uses the Contrail Cloud implementation to support the service provider's cloud in a network point of presence (POP). The Contrail Cloud

implementation consists of the hardware platforms, including the servers, and Contrail OpenStack software. [Figure 31 on page 100](#) illustrates the Contrail Cloud implementation. The Contrail Service Orchestration (CSO) software is installed on one or more servers in the Contrail Cloud implementation to complete the deployment.

**Figure 31: Architecture of Contrail Cloud Implementation**



In the Cloud CPE Centralized Deployment Model:

- The MX Series router provides the gateway to the service provider's cloud.
- The EX Series switch provides Ethernet management and Intelligent Platform Management Interface (IPMI) access for all components of the Cloud CPE Centralized Deployment Model. Two interfaces on each server connect to this switch.
- The QFX Series switch provides data access to all servers.
- The number of servers depends on the scale of the deployment and the high availability configuration. You must use at least two servers and you can use up to five servers.
- Each server supports four nodes. The function of the nodes depends on the high availability configuration and the type of POP.

## Architecture of the Servers


The configuration of the nodes depends on whether the Contrail Cloud implementation is in a regional POP or central POP and on the high availability configuration. Each node is one of the following types:

- Contrail Service Orchestration node, which hosts the Contrail Service Orchestration software.
- Contrail controller node, which hosts the Contrail controller and Contrail Analytics.
- Contrail compute node, which hosts the Contrail Openstack software and the virtualized network functions (VNFs).

The Contrail Cloud implementation in a central POP contains all three types of node. [Figure 32 on page 101](#) shows the configuration of the nodes in the Contrail Cloud implementation in the central POP for a deployment that offers neither Contrail nor Contrail Service Orchestration high availability:

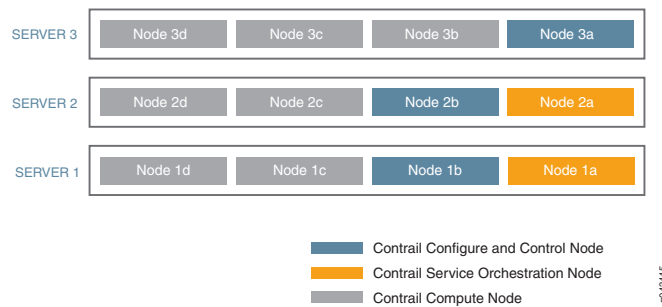
- Server 1 supports one Contrail controller node, two Contrail compute nodes, and one Contrail Service Orchestration node.
- Server 2 and optional servers 3 through 5 each support four Contrail compute nodes.

*Figure 32: Architecture of Servers in the Central POP for a Non-Redundant Installation*



[Figure 33 on page 102](#) shows the configuration of the nodes in the Contrail Cloud implementation in the central POP for a deployment that offers both Contrail and Contrail Service Orchestration high availability:

- Servers 1, 2, and 3 each support one Contrail controller node for Contrail redundancy.
- Servers 1 and 2 each support one Contrail Service Orchestration node for Contrail Service Orchestration redundancy.
- Other nodes on servers 1, 2, and 3 are Contrail compute nodes. Optional servers 4 through 7 also support Contrail compute nodes.

**Figure 33: Architecture of Servers in the Central POP for a Redundant Installation**

The Contrail Cloud implementation in a regional POP contains only Contrail nodes and not Contrail Service Orchestration nodes. In a deployment that does not offer Contrail high availability, the regional Contrail Cloud implementations support:


- One Contrail controller node and three Contrail compute nodes on server 1.
- Four Contrail compute nodes on server 2 and on optional servers 3 through 5.

In a deployment that offers Contrail high availability, the regional Contrail Cloud implementations support:

- One Contrail controller node for Contrail redundancy on servers 1, 2, and 3.
- Three Contrail compute nodes on servers 1, 2, and 3.
- Four Contrail compute nodes on optional servers 4 through 7.

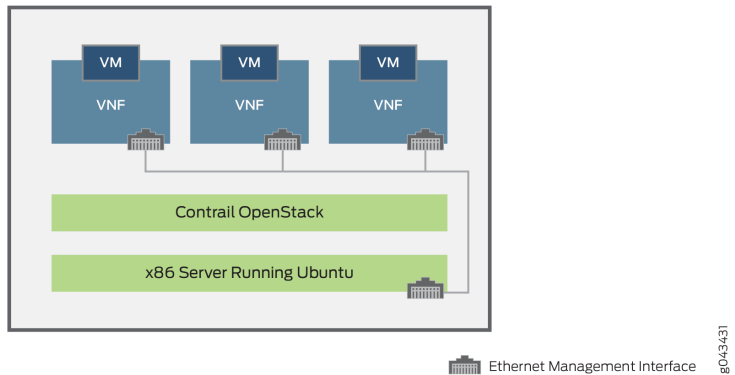
## Architecture of the Contrail Nodes

Each Contrail controller node uses Contrail vRouter over Ubuntu and kernel-based virtual machine (KVM) as a forwarding plane in the Linux kernel. Use of vRouter on the compute node separates the deployment's forwarding plane from the control plane, which is the SDN Controller in Contrail OpenStack on the controller node. This separation leads to uninterrupted performance and enables scaling of the deployment. [Figure 34 on page 102](#) shows the logical representation of the Contrail controller nodes.

**Figure 34: Logical Representation of Contrail Controller Nodes**

A Contrail compute node hosts Contrail OpenStack, and the VNFs. Contrail OpenStack resides on the physical server and cannot be deployed in a VM. Each VNF resides in its own VM. [Figure 35 on page 103](#) shows the logical representation of the Contrail compute nodes.

Figure 35: Logical Representation of Contrail Compute Nodes



- Related Documentation**
- [Network Function Virtualization in the Contrail Service Orchestration Deployments on page 32](#)

## Cabling the Hardware for the Centralized Deployment

This section describes how to connect cables among the network devices and servers in the Contrail Cloud implementation. See [Architecture of the Contrail Cloud Implementation in the Centralized Deployment](#) for more information.

To cable the hardware:

1. Connect cables from the EX Series switch to the other devices in the network.  
See [Table 13 on page 103](#) for information about the connections for the EX Series switch.
2. Connect cables from the QFX Series switch to the other devices in the network.  
See [Table 14 on page 104](#) for information about the connections for the QFX Series switch.
3. Connect cables from the MX Series router to the other devices in the network.  
See [Table 15 on page 105](#) for information about the connections for the MX Series router.

Table 13: Connections for EX Series Switch

Interface on EX Series Switch	Destination Device	Interface on Destination Device
eth0 (management interface)	EX Series switch	ge-0/0/41
ge-0/0/0	Server 1	IPMI
ge-0/0/1	Server 2	IPMI

**Table 13: Connections for EX Series Switch (continued)**

Interface on EX Series Switch	Destination Device	Interface on Destination Device
ge-0/0/2	Server 3	IPMI
ge-0/0/3	Server 4	IPMI
ge-0/0/4	Server 5	IPMI
ge-0/0/5	Server 6	IPMI
ge-0/0/6	Server 7	IPMI
ge-0/0/20	Server 1	eth0
ge-0/0/21	Server 2	eth0
ge-0/0/22	Server 3	eth0
ge-0/0/23	Server 4	eth0
ge-0/0/24	Server 5	eth0
ge-0/0/25	Server 6	eth0
ge-0/0/26	Server 7	eth0
ge-0/0/41	EX Series switch	eth0 (management interface)
ge-0/0/42	QFX Series switch	eth0 (management interface)
ge-0/0/44	MX Series router	fxp0
ge-0/0/46	MX Series router	ge-1/3/11
ge-0/0/47	Server 1	eth1

**Table 14: Connections for QFX Series Switch**

Interface on QFX Series Switch	Destination Device	Interface on Destination Device
eth0 (management interface)	EX Series switch	ge-0/0/42
xe-0/0/0	Server 1	eth2
xe-0/0/1	Server 2	eth2
xe-0/0/2	Server 3	eth2



**Table 14: Connections for QFX Series Switch (continued)**

Interface on QFX Series Switch	Destination Device	Interface on Destination Device
xe-0/0/3	Server 4	eth2
xe-0/0/4	Server 5	eth2
xe-0/0/5	Server 6	eth2
xe-0/0/6	Server 7	eth2
xe-0/0/20	Server 1	eth3
xe-0/0/21	Server 2	eth3
xe-0/0/22	Server 3	eth3
xe-0/0/23	Server 4	eth3
xe-0/0/24	Server 5	eth3
xe-0/0/24	Server 6	eth3
xe-0/0/25	Server 7	eth3
xe-0/0/46	MX Series router	xe-0/0/0
xe-0/0/47	MX Series router	xe-0/0/1

**Table 15: Connections for MX Series Router**

Interface on MX Series Router	Destination Device	Interface on Destination Device
fxp0 (management interface)	EX Series switch	ge-0/0/44
ge-1/3/11	EX Series switch	ge-0/0/46
xe-0/0/0	QFX Series switch	xe-0/0/46
xe-0/0/1	QFX Series switch	xe-0/0/47
ge-1/0/0 and ge-1/0/1 or xe-0/0/2 and xe-0/0/3, depending on the network	Service provider's device at the cloud	—

**Related Documentation** • [Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 106](#)

- [Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 107](#)
- [Configuring the MX Series Router for the Contrail Cloud Implementation in a Centralized Deployment on page 109](#)
- [Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment on page 111](#)

## Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment

Before you configure the EX Series switch, complete any basic setup procedures and install the correct Junos OS software release on the switch.

To configure the EX Series switch:

1. Define VLANs for the IPMI ports. For example:

```
user@switch# set interfaces interface-range ipmi member-range ge-0/0/0 to
ge-0/0/19
user@switch# set interfaces interface-range ipmi unit 0 family ethernet-switching
port-mode access
user@switch# set interfaces interface-range ipmi unit 0 family ethernet-switching
vlan members ipmi
user@switch# set interfaces vlan unit 60 family inet address 172.16.60.254/24
user@switch# set vlans ipmi vlan-id 60
user@switch# set vlans ipmi l3-interface vlan.60
```

2. Define a VLAN for the management ports. For example:

```
user@switch# set interfaces interface-range mgmt member-range ge-0/0/20 to
ge-0/0/46
user@switch# set interfaces interface-range mgmt unit 0 family ethernet-switching
port-mode access
user@switch# set interfaces interface-range mgmt unit 0 family ethernet-switching
vlan members mgmt
user@switch# set interfaces vlan unit 70 family inet address 172.16.70.254/24
user@switch# set vlans mgmt vlan-id 70
user@switch# set vlans mgmt l3-interface vlan.70
```

3. Define a static route for external network access. For example:

```
user@switch# set routing-options static route 0.0.0.0/0 next-hop 172.16.70.253
```

### Related Documentation

- [Building Blocks Used for Contrail Service Orchestration Deployments on page 22](#)
- [Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 107](#)

- [Configuring the MX Series Router for the Contrail Cloud Implementation in a Centralized Deployment on page 109](#)

## Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment

Before you configure the QFX Series switch, complete any basic setup procedures and install the correct Junos OS software release on the switch.

To configure the QFX Series switch:

1. Configure the IP address of the Ethernet management port. For example:

```
user@switch# set interfaces vme unit 0 family inet address 172.16.70.251/24
```

2. Configure integrated routing and bridging (IRB). For example:

```
user@switch# set interfaces irb unit 80 family inet address 172.16.80.254/24
```

3. Configure a link aggregation group (LAG) for each pair of server ports. For example:

```
user@switch# set interfaces xe-0/0/0 ether-options 802.3ad ae0
user@switch# set interfaces xe-0/0/20 ether-options 802.3ad ae0
user@switch# set interfaces ae0 mtu 9192
user@switch# set interfaces ae0 aggregated-ether-options lacp active
user@switch# set interfaces ae0 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae0 unit 0 family ethernet-switching interface-mode
access
user@switch# set interfaces ae0 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/1 ether-options 802.3ad ae1
user@switch# set interfaces xe-0/0/21 ether-options 802.3ad ae1
user@switch# set interfaces ae1 mtu 9192
user@switch# set interfaces ae1 aggregated-ether-options lacp active
user@switch# set interfaces ae1 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae1 unit 0 family ethernet-switching interface-mode
access
user@switch# set interfaces ae1 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/2 ether-options 802.3ad ae2
user@switch# set interfaces xe-0/0/22 ether-options 802.3ad ae2
user@switch# set interfaces ae2 mtu 9192
user@switch# set interfaces ae2 aggregated-ether-options lacp active
user@switch# set interfaces ae2 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae2 unit 0 family ethernet-switching interface-mode
access
user@switch# set interfaces ae2 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/3 ether-options 802.3ad ae3
```

```
user@switch# set interfaces xe-0/0/23 ether-options 802.3ad ae3
user@switch# set interfaces ae3 mtu 9192
user@switch# set interfaces ae3 aggregated-ether-options lacp active
user@switch# set interfaces ae3 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae3 unit 0 family ethernet-switching interface-mode
    access
user@switch# set interfaces ae3 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/4 ether-options 802.3ad ae4
user@switch# set interfaces xe-0/0/24 ether-options 802.3ad ae4
user@switch# set interfaces ae4 mtu 9192
user@switch# set interfaces ae4 aggregated-ether-options lacp active
user@switch# set interfaces ae4 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae4 unit 0 family ethernet-switching interface-mode
    access
user@switch# set interfaces ae4 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/5 ether-options 802.3ad ae5
user@switch# set interfaces xe-0/0/25 ether-options 802.3ad ae5
user@switch# set interfaces ae5 mtu 9192
user@switch# set interfaces ae5 aggregated-ether-options lacp active
user@switch# set interfaces ae5 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae5 unit 0 family ethernet-switching interface-mode
    access
user@switch# set interfaces ae5 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/6 ether-options 802.3ad ae6
user@switch# set interfaces xe-0/0/26 ether-options 802.3ad ae6
user@switch# set interfaces ae6 mtu 9192
user@switch# set interfaces ae6 aggregated-ether-options lacp active
user@switch# set interfaces ae6 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae6 unit 0 family ethernet-switching interface-mode
    access
user@switch# set interfaces ae6 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/7 ether-options 802.3ad ae7
user@switch# set interfaces xe-0/0/27 ether-options 802.3ad ae7
user@switch# set interfaces ae7 mtu 9192
user@switch# set interfaces ae7 aggregated-ether-options lacp active
user@switch# set interfaces ae7 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae7 unit 0 family ethernet-switching interface-mode
    access
user@switch# set interfaces ae7 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/8 ether-options 802.3ad ae8
user@switch# set interfaces xe-0/0/28 ether-options 802.3ad ae8
user@switch# set interfaces ae8 mtu 9192
user@switch# set interfaces ae8 aggregated-ether-options lacp active
user@switch# set interfaces ae8 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae8 unit 0 family ethernet-switching interface-mode
    access
user@switch# set interfaces ae8 unit 0 family ethernet-switching vlan members data
```

4. Configure a VLAN for data transmission. For example:

```
user@switch# set vlans data vlan-id 80
user@switch# set vlans data l3-interface irb.80
```

5. Configure OSPF routing. For example:

```
user@switch# set interfaces irb unit 80 family inet address 172.16.80.254/24
user@switch# set protocols ospf area 0.0.0.0 interface irb.80 passive
```

6. Configure the interface that connects to the MX Series router. For example:

```
user@switch# set interfaces xe-0/0/46 ether-options 802.3ad ae9
user@switch# set interfaces xe-0/0/47 ether-options 802.3ad ae9
```

```
user@switch# set interfaces ae9 aggregated-ether-options lACP active
user@switch# set interfaces ae9 aggregated-ether-options lACP periodic fast
user@switch# set interfaces ae9 unit 0 family inet address 172.16.10.253/24
```

```
user@switch# set protocols ospf area 0.0.0.0 interface ae9.0
```

#### Related Documentation

- [Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 106](#)
- [Configuring the MX Series Router for the Contrail Cloud Implementation in a Centralized Deployment on page 109](#)

## Configuring the MX Series Router for the Contrail Cloud Implementation in a Centralized Deployment

Before you configure the MX Series router, complete any basic setup procedures and install the correct Junos OS software release on the switch.

To configure the MX Series router:

1. Configure interfaces, IP addresses, and basic routing settings. For example:

```
user@router# set interfaces ge-1/0/0 unit 0 family inet address 10.87.24.77/28
user@router# set interfaces lo0 unit 0 family inet address 172.16.100.1/32
user@router# set routing-options route-distinguisher-id 172.16.100.1
user@router# set routing-options autonomous-system 64512
user@router# set protocols ospf area 0.0.0.0 interface lo0.0
```

```
user@router# set interfaces ge-1/0/0 unit 0 family inet service input service-set s1
service-filter ingress-1
user@router# set interfaces ge-1/0/0 unit 0 family inet service output service-set s1
service-filter ingress-1
```

2. Configure the interfaces that connect to the QFX Series switch. For example:

```
user@router# set chassis aggregated-devices ethernet device-count 2
user@router# set interfaces xe-0/0/0 gigether-options 802.3ad ae0
user@router# set interfaces xe-0/0/1 gigether-options 802.3ad ae0
user@router# set interfaces ae0 aggregated-ether-options lacp periodic fast
user@router# set interfaces ae0 unit 0 family inet service input service-set s1
service-filter ingress-1
user@router# set interfaces ae0 unit 0 family inet service output service-set s1
service-filter ingress-1
user@router# set interfaces ae0 unit 0 family inet address 172.16.10.254/24
user@router# set protocols ospf area 0.0.0.0 interface ae0.0
```

3. Configure BGP and tunneling for the service provider's cloud. For example:

```
user@router# set chassis fpc 0 pic 0 tunnel-services
user@router# set chassis fpc 0 pic 0 inline-services bandwidth 1g
user@router# set routing-options dynamic-tunnels dynamic_overlay_tunnels
source-address 172.16.100.1
user@router# set routing-options dynamic-tunnels dynamic_overlay_tunnels gre
user@router# set routing-options dynamic-tunnels dynamic_overlay_tunnels
destination-networks 172.16.80.0/24
user@router# set protocols mpls interface all
user@router# set protocols bgp group Contrail_Controller type internal
user@router# set protocols bgp group Contrail_Controller local-address 172.16.100.1
user@router# set protocols bgp group Contrail_Controller keep all
user@router# set protocols bgp group Contrail_Controller family inet-vpn unicast
user@router# set protocols bgp group Contrail_Controller neighbor 172.16.80.2
user@router# set protocols bgp group Contrail_Controller neighbor 172.16.80.3
user@router# set protocols ospf export leak-default-only
```

4. Set up routing. For example:

```
user@router# set routing-options static rib-group inet-to-public
user@router# set routing-options static route 0.0.0.0/0 next-hop 10.87.24.78
user@router# set routing-options static route 0.0.0.0/0 retain
user@router# set routing-options static route 10.87.24.64/26 next-table public.inet.0
user@router# set routing-options rib-groups inet-to-public import-rib inet.0
user@router# set routing-options rib-groups inet-to-public import-rib public.inet.0
user@router# set routing-options rib-groups inet-to-public import-policy
leak-default-only
user@router# set policy-options policy-statement leak-default-only term default
from route-filter 0.0.0.0/0 exact
user@router# set policy-options policy-statement leak-default-only term default then
accept
user@router# set policy-options policy-statement leak-default-only then reject
user@router# set routing-instances public instance-type vrf
user@router# set routing-instances public interface lo0.10
user@router# set routing-instances public vrf-target target:64512:10000
user@router# set routing-instances public vrf-table-label
user@router# set routing-instances public routing-options static route 10.87.24.64/26
discard
```

## 5. Configure NAT. For example:

```

user@router# set services service-set s1 nat-rules rule-napt-zone
user@router# set services service-set s1 interface-service service-interface si-0/0/0.0
user@router# set services nat pool contrailui address 10.87.24.81/32
user@router# set services nat pool openstack address 10.87.24.82/32
user@router# set services nat pool jumphost address 10.87.24.83/32
user@router# set services nat rule rule-napt-zone term t1 from source-address
172.16.80.2/32
user@router# set services nat rule rule-napt-zone term t1 then translated source-pool
openstack
user@router# set services nat rule rule-napt-zone term t1 then translated
translation-type basic-nat44
user@router# set services nat rule rule-napt-zone term t2 from source-address
172.16.80.4/32
user@router# set services nat rule rule-napt-zone term t2 then translated source-pool
contrailui
user@router# set services nat rule rule-napt-zone term t2 then translated
translation-type basic-nat44
user@router# set services nat rule rule-napt-zone term t3 from source-address
172.16.70.1/32
user@router# set services nat rule rule-napt-zone term t3 then translated source-pool
jumphost
user@router# set services nat rule rule-napt-zone term t3 then translated
translation-type basic-nat44
user@router# set firewall family inet service-filter ingress-1 term t1 from source-address
172.16.80.2/32
user@router# set firewall family inet service-filter ingress-1 term t1 from protocol tcp
user@router# set firewall family inet service-filter ingress-1 term t1 from
destination-port-except 179
user@router# set firewall family inet service-filter ingress-1 term t1 then service
user@router# set firewall family inet service-filter ingress-1 term t2 from source-address
172.16.80.4/32
user@router# set firewall family inet service-filter ingress-1 term t2 then service
user@router# set firewall family inet service-filter ingress-1 term t3 from source-address
172.16.70.1/32
user@router# set firewall family inet service-filter ingress-1 term t3 then service
user@router# set firewall family inet service-filter ingress-1 term end then skip

```

**Related  
Documentation**

- [Building Blocks Used for Contrail Service Orchestration Deployments on page 22](#)
- [Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 106](#)
- [Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 107](#)

## Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment

For a centralized deployment, you must configure the physical servers and nodes in the Contrail Cloud implementation and install Contrail OpenStack on the server cluster before you run the installer.

To install Contrail OpenStack:

1. Configure hostnames for the physical servers and nodes.
2. Configure IP addresses for the Ethernet management ports of the physical servers and nodes.
3. Configure DNS on the physical servers and nodes, and ensure that DNS is working correctly.
4. Configure Internet access for the physical servers and nodes.
5. From each server and node, verify that you can ping the IP addresses and hostnames of all the other servers and nodes in the Contrail Cloud implementation.
6. Using Contrail Server Manager, install Contrail OpenStack on the server cluster and set up the roles of the Contrail nodes in the cluster.

You configure an OpenStack Keystone on the primary Contrail controller node in the central Contrail Cloud implementation, and also use this Keystone for:

- Regional Contrail configure and control nodes
- Redundant configure and control nodes in the central Contrail Cloud implementation

Refer to the Contrail documentation for information about installing Contrail OpenStack and configuring the nodes.

7. For each node, use the ETCD keys to specify the same username and password for Contrail.

CSO uses the BASIC authentication mechanism to establish a connection to Contrail.

#### **Related Documentation**



## CHAPTER 7

# Appendix B – Uploading VNF Images for Centralized Deployment

- [Uploading the vSRX VNF Image for a Centralized Deployment on page 113](#)
- [Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 115](#)
- [Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 117](#)

## Uploading the vSRX VNF Image for a Centralized Deployment

---

The Contrail Service Orchestration (CSO) installer places the vSRX image in the `/var/www/html/csp_components` directory on the installer virtual machine (VM) during the installation process. You must copy this image from the installer VM to the Contrail controller node and upload it to make the vSRX virtualized network function (VNF) available in a centralized deployment.

To upload the vSRX VNF image for a centralized deployment:

1. Log in to the installer VM as root.
2. Set up an SSH session as root to the Contrail controller node.
3. Copy the **vSRX-img** file from the installer VM to any directory on the Contrail controller node.

For example, if the IP address of the Contrail controller node is 192.0.2.1, and you want to copy the file to the **root** directory:

```
root@host: /# scp /var/www/html/csp_components/vSRX-img root@192.0.2.1:root
```

4. Check whether you have an OpenStack flavor with the following specification on the Contrail controller node.
  - 2 vCPUs
  - 4 GB RAM
  - 40 GB hard disk storage

For example:

```
root@host: /# openstack flavor list
```

ID	Name	Memory_MB	Disk	Ephemeral	Swap	VCPUs	Is_Public
1	m1.tiny	512	0	0		1	True
2	m1.small	2048	20	0		1	True
3	m1.medium	4096	40	0		2	True
4	m1.large	8192	80	0		4	True
42	m1.nano	64	0	0		1	True
5	m1.xlarge	16384	160	0		8	True
84	m1.micro	128	0	0		1	True

If you do not have a flavor with the required specification, create one.

For example:

```
root@host: /# openstack flavor create m1.vsrx_flavor --ram 4096 --disk 40 --vcpus 2
```

5. Access the directory where you copied the image on the Contrail controller node, and upload it into the Glance software.

For example:

```
root@host: /# cd root
root@host: /root# glance image-create --name vSRX-img --is-public True --container-format bare --disk-format qcow2 < vSRX-img
```



**NOTE:** You must name the image vSRX-img to ensure that the virtual infrastructure manager (VIM) can instantiate the VNF.

To verify that you can manually instantiate the vSRX VNF:

1. Access the OpenStack dashboard.
2. Create an instance of the vSRX image.
3. Select Projects > Instances.

The status of the instance should be spawning or running. You can click the instance to see its console.

If you need to investigate the image further, the default username for the vSRX-img package is root and the password is passwOrd.

#### Related Documentation

- [VNFs Supported by the Contrail Service Orchestration Solutions on page 37](#)
- [Uploading the LxCiPTable VNF Image for a Centralized Deployment on page 115](#)
- [Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 117](#)

## Uploading the LxCIPtable VNF Image for a Centralized Deployment

You use this process to make the LxCIPtable VNF available in a centralized deployment.

To create an LxCIPtable Image:

1. At <http://cloud-images.ubuntu.com/releases/14.04/release/>, determine the appropriate Ubuntu cloud image for your Contrail controller node.
2. Download the appropriate Ubuntu cloud image to the Contrail controller node.

For example:

```
root@host:/# cd /tmp
root@host:/tmp# wget
http://cloud-images.ubuntu.com/releases/14.04/release/ubuntu-14.04-server-cloudimg-amd64-disk1.img
```

3. On the Contrail controller node, upload the Ubuntu image into the Glance software.

```
root@host:/# glance image-create --name IPtables --is-public True --container-format bare
--disk-format qcow2 < ubuntu-14.04-server-cloudimg-amd64-disk1.img
```

4. In a local directory on the Contrail OpenStack node, create a metadata file for the image. For example:

```
root@host:~/images# cat user-data.txt
#cloud-config
password: <PASSWORD>
chpasswd: { expire: False }
ssh_pwauth: True
```

5. Create an instance of the image called **IPtable-temp** in this directory.

```
root@host:~/images# nova boot --flavor m1.medium --user-data=./user-data.txt --image
IPtables IPtable-temp --nic net-id=<management network id>
```

6. From the OpenStack GUI, log in to the instance with the username **ubuntu** and the password specified in the user-data file.

7. Customize the instance.

- a. Set the root password to the value **passwOrd**. For example:



**CAUTION:** You must use the value **passwOrd** for the LxCIPtable VNF to operate correctly.

```
ubuntu@iptables-temp:~$sudo passwd root
Enter new UNIX password:
Retype new UNIX password:
passwd: password updated successfully
ubuntu@iptables-temp:~$
```

- b. In the file **/etc/ssh/sshd\_config**, specify the following setting:

```
PermitRootLogin = yes
```

- c. Restart the service.

```
service ssh restart
```

- d. In the file **/etc/network/interfaces**, modify the eth0, eth1, and eth2 settings as follows:

```
auto eth0
iface eth0 inet dhcp
metric 1
```

```
auto eth1
iface eth1 inet dhcp
metric 100
```

```
auto eth2
iface eth2 inet dhcp
metric 100
```

- e. Verify that IPtables is active.

```
service ufw status
```

8. Take a snapshot of the OpenStack Instance.

- a. Close the instance.

```
sudo shutdown -h now
```

- b. From the OpenStack Instances page, select **Create Snapshot** for this instance, and specify the Name as **Lxclmg**.

- c. Delete the temporary instance that you created in Step 5.

#### Related Documentation

- [VNFs Supported by the Contrail Service Orchestration Solutions on page 37](#)
- [Uploading the vSRX VNF Image for a Centralized Deployment on page 113](#)
- [Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 117](#)

## Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment

You use this process to make the Cisco CSR-1000V VNF available in a centralized deployment.

To create a Cisco CSR-1000V VNF image:

1. Log into the Contrail controller node.
2. Create a new flavor with 3 vCPUs in Contrail OpenStack.

For example:

```
root@host:~# openstack flavor create p1.csr_flavor --ram 4096 --disk 0 --vcpus 3
```

3. Upload the Cisco CSR-1000V image into the Glance software.

For example:

```
root@host:~# glance image-create --name csr1000v-img --is-public True --container-format bare --disk-format qcow2 < cisco-csr1000v-img
```

4. Create an instance of the image called `csr1000v-img` in this directory.

For example:

```
root@host:~/images# nova boot --flavor p1.csr_flavor --image csr1000v-img --nic net-id=MGMT_NET_ID --nic net-id=LEFT_NET_ID --nic net-id=RIGHT_NET_ID --security-group default
```

5. From the OpenStack GUI, log in to the instance using the management IP address as the username and without a password.
6. Configure the following settings for the instance:

```
vrf definition Mgmt-intf
 address-family ipv4
  exit-address-family
 enable password passw0rd
 ip vrf mgmt
```

```
username root privilege 15 password 0 passw0rd
ip ssh version 2
interface GigabitEthernet1
  ip vrf forwarding mgmt
  ip address dhcp
  negotiation auto
line vty 0
  exec-timeout 60 0
  privilege level 15
  password passw0rd
  login local
  transport input telnet ssh
```

7. Take a snapshot of the instance.

- a. Close the instance.

For example:

```
root@host:~/images# sudo shutdown -h now
```

- b. From the OpenStack Instances page, select **Create Snapshot** for this instance, and specify the name of the image as **csr1000v-img**.

#### Related Documentation

- [VNFS Supported by the Contrail Service Orchestration Solutions on page 37](#)
- [Uploading the vSRX VNF Image for a Centralized Deployment on page 113](#)
- [Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 115](#)

## CHAPTER 8

# Appendix C – Manual Staging of NFX

- [Install Junos Software onto NFX from USB Port on page 119](#)

## Install Junos Software onto NFX from USB Port

---

This section details how to install Junos OS software version 15.1X53-D496.0 onto an NFX250 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. This procedure is for an NFX250 device.

### Before You Begin

In order for this procedure to succeed, you must have the following

- Physical access to the USB port of the NFX device
- A USB drive of at least 4GB containing the Junos OS Software image, 15.1X53-D496.0, inserted into the USB port of the NFX
- Access to the console port of the NFX 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 NFX250. This DHCP server must be able to provide IP address, name server, and default gateway to the NFX upon request.

The following procedures contain comments that are added to clarify the steps that are discussed.

1. Ensure that the USB drive containing the Junos OS software image is inserted in the USB port of the NFX device.

This allows you to boot the NFX from the USB drive.

2. Access the NFX 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 device.

4. Power on the NFX 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 will install 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 is made up of multiple components that load and boot in a specific order. See [NFX 250 Overview](#) for details. The PFE of the NFX may take a few minutes to complete the boot and allow the jsxe0 interface to obtain its address from DHCP.

You can login to the console of the NFX as **root** 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 for different purposes. You are looking 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) in order 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
Pseudo CB 0
Routing Engine 0
FPC 0            REV 04    650-066113   DXXXXXXXXXX3  RE-NFX250-S2
CPU
PIC 0           REV 04    BUILTIN     BUILTIN       FPC CPU
Base-T-2x1G SFP-
```

```
Power Supply 0
Fan Tray 0      fan-ctrl-0 0, Front
  to Back Airflow - AF0
Fan Tray 1      fan-ctrl-0 1, Front
  to Back Airflow - AF0
```

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 on page 22](#)
- [Hybrid WAN Deployment Overview on page 77](#)
- [SD-WAN Deployment Overview on page 47](#)