

Contrail Service Orchestration

CSO SD-WAN and SD-LAN – Design and Architecture Guide

Published
2020-09-28

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

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

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

Contrail Service Orchestration CSO SD-WAN and SD-LAN – Design and Architecture Guide
Copyright © 2020 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About the Documentation | vii

Documentation and Release Notes | vii

Documentation Conventions | vii

Documentation Feedback | x

Requesting Technical Support | x

Self-Help Online Tools and Resources | xi

Creating a Service Request with JTAC | xi

1

About This Guide

Inside and Outside the Guide | 2

Coverage and Audience | 2

CSO Terminology | 2

References | 5

2

Introduction

Deployment Models and Building Blocks | 7

Introduction | 7

Juniper Networks Cloud CPE Solution | 10

Contrail SD-WAN Model | 11

Contrail SD-LAN Model | 12

Next-Generation Firewall (NGFW) Model | 16

Building Blocks Used for Contrail Service Orchestration Deployments | 17

Administrators | 18

Portals | 18

Tenants | 19

Topologies | 19

Points of Presence (POPs) | 22

Sites | 23

On-Premises Spoke Devices | 28

Standalone Next-Generation Firewall (NGFW) | 29

SD-LAN Devices | 29

- Virtual Route Reflector (VRR) | 29
- SLA-Based Steering Profiles and Policies | 30
- Path Based Steering Profiles | 31
- Intent-based Firewall Policies | 31
- Software Image Management | 31

Contrail Service Orchestration Solutions

SD-WAN Solution | 33

Overview | 33

SD-WAN Requirements | 34

Contrail SD-WAN Solution Highlights | 35

Target Customers | 35

Contrail SD-WAN Deployment Architectures | 36

Contrail SD-WAN Reference Architecture | 37

Spoke Devices | 38

On-Premises Spoke Devices | 38

Cloud Spoke Devices | 40

Spoke Redundancy | 40

Provider Hub Devices | 40

Provider Hubs | 40

Provider Hub Redundancy | 41

Enterprise Hub Sites and Devices | 41

Underlay (Physical) Network | 43

WAN Access Options | 43

WAN Interface Types - Data and OAM | 44

Overlay (Tunnels) Network | 45

Overlay Deployment Topologies | 45

Orchestration and Control | 47

Secure OAM Network | 48

Integration with Deployment Topologies | 49

OAM Hub Design Options | 50

Usage Notes on Provider Hub Design Options | 51

Zero Touch Provisioning | 52

Usage Notes for ZTP | 52

Redirect Server | 52

Service Chaining in Contrail SD-WAN | 53

Three Planes, Four Layers | 54

Two Basic SD-WAN Use Cases | 55

Managed SD-WAN - Overlay Access | 55

Enterprise SD-WAN - Overlay | 57

Secure and Redundant OAM Network | 58

SD-LAN Solution | 62

SD-LAN Deployment Architectures | 63

SD-LAN Architectures | 64

SD-LAN with EX Series Switches | 66

Standalone SD-LAN Site | 66

Extend SD-WAN Site | 67

SD-LAN Behind NGFW Site | 68

Managed WLAN Using Mist Wireless Access Points | 69

NGFW Solution | 70

NGFW Deployment Architecture | 70

NGFW Architecture | 70

NGFW Deployment Usage Notes | 71

4

Operations and Management

Network Operation | 73

Network Operation | 73

vRR Design | 73

Control Plane Resiliency | 74

Route Distribution and Separation | 75

APBR and SLA Management - Control Plane | 77

Data Plane Operation | 78

Mesh Tags and Dynamic Mesh VPNs | 81

Internet Breakout | 86

Network Security | 92

Orchestration and Management Within CSO | 94

Orchestration and Management Using CSO | 94

Architecture | 94

Orchestration Layers | 95

Infrastructure Services and Microservices | 96

Platform Characteristics | 97

Multitenancy with RBAC | 97

High Availability and Scale | 97

Programmability and Integration | 98

Extensibility and Customization | 99

Telemetry and Analytics Capability | 100

Intent-Based Policies | 100

Upgrade and Backward Compatibility | 102

Element Management | 102

CSO User Interface | 106

Operational Workflows - Overview | 110

Operational Workflows | 110

Resiliency and High Availability | 111

Resiliency and High Availability | 111

Network Control Plane | 111

Headless Forwarding | 111

Data Plane | 112

Spoke Redundancy | 112

Using NFX Series Devices | 112

Using SRX Series Devices | 112

Failover Scenarios | 113

Hub Redundancy (CPE Multihoming) | 114

Design Options | 114

About the Documentation

IN THIS SECTION

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

Use this design and architecture guide as a high-level reference for understanding the components and capabilities of Juniper Networks SD-WAN, SD-LAN, and NGFW solutions within the Contrail Service Orchestration platform.

Documentation and Release Notes

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

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

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

Documentation Conventions

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

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

Convention	Description	Examples
<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@# set system domain-name <i>domain-name</i>
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"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

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

Convention	Description	Examples
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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



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

Requesting Technical Support

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

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

1

PART

About This Guide

[Inside and Outside the Guide](#) | 2

Inside and Outside the Guide

IN THIS CHAPTER

- Coverage and Audience | 2
- CSO Terminology | 2
- References | 5

Coverage and Audience

This guide discusses design and architecture elements for Juniper's Contrail SD-WAN, SD-LAN, and NGFW solutions within Juniper's Contrail Service Orchestration (CSO) software. The guide covers features and functionality up to and including CSO Release 5.3.0.

This guide is appropriate for network designers, planners, and architects, as well as network engineers and administrators who need to understand the solution at a higher level.

CSO Terminology

Some of the terms used throughout this guide might be unfamiliar or specific to Juniper Networks Contrail SD-WAN, SD-LAN, or NGFW solutions. [Table 3 on page 2](#) provides definitions for the terminology used throughout this guide.

Table 3: CSO Terminology

Term	Definition
Branch	A tenant site, connected to other sites in either a full mesh or hub-and-spoke topology. Also known as a <i>spoke site</i> .

Table 3: CSO Terminology (*continued*)

Term	Definition
CPE	Customer-premises equipment—A device placed at a remote customer spoke site that provides services (such as WAN routing or firewall filtering) for the remote site. The CPE allows the remote site to connect with a hub or other spoke sites. Legacy CPE devices provide single services, newer CPE devices (such as the NFX Series and SRX Series devices) provide multiple services to enable the SD-WAN, SD-LAN, and NGFW solutions. See also <i>on-premises spoke device</i> .
CSO	Contrail Service Orchestration—A Juniper Networks software product that facilitates the Contrail SD-WAN, SD-LAN, and NGFW solutions. You access CSO through a graphical user interface (GUI) to harness its built-in automation capabilities, which enable you to provision, manage, and monitor your WAN, campus, and branch networks.
Dynamic Mesh	A resource conserving method for implementing full-mesh topologies. All of the sites in the full mesh are included in the topology; but the site-to-site VPNs are not brought up until traffic crosses a user-defined threshold called the Dynamic VPN threshold.
Enterprise Hub	A single-tenant on-premises spoke device deployed as a hub at an enterprise hub site. The enterprise hub can serve as the hub portion of a hub-and-spoke topology. When deployed like this, the provider hub (if any) serves as a backup hub to the enterprise hub for site-to-site communications.
Enterprise Hub Site	A special type of spoke site with enhanced capabilities that approximate those of a provider hub site.
Hub Site	A site that acts as a hub for traffic from multiple spokes in a hub-and-spoke topology. In the absence of an enterprise hub, all spoke-to-spoke traffic flows through the provider hub. See also <i>Provider Hub</i> and <i>Enterprise Hub</i> .
MANO	Management and Orchestration
Mesh Tag	A text-based label for WAN interfaces on CPE devices. Mesh tags enable SLA-based dynamic VPN creation between customer sites. Only interfaces with matching mesh tags can form a VPN.
Microservices	Lightweight, modular building blocks that implement a specific function and communicate with other functions using well defined interfaces (e.g. RESTful APIs). Can be scaled independently.
MP-BGP	Multiprotocol BGP—A routing protocol used for large-scale, multi-tenancy deployments
NGFW	Next-generation firewall—An SRX Series Services Gateway placed at a remote customer site that acts as a CPE and provides WAN and advanced security services.
NSC	Network Service Controller—The SD-WAN controller layer of CSO, provides topology and CPE lifecycle management functionality, as well as site-to-site routing and reachability.

Table 3: CSO Terminology (*continued*)

Term	Definition
On-premises spoke device	See <i>CPE</i> .
OpCo	<p>Operating Company—Typically a service provider who has multiple large tenants. A single instance of CSO can have multiple OpCos, each with multiple tenants.</p> <p>NOTE: An OpCo administrator is the highest level of administrator available in cloud-hosted CSO deployments.</p>
POP	Point of Presence—Typically a physical location where the provider has assets used to deploy one or more of the available solutions. Assets are network devices such as edge routers, provider hubs, and server resources. The POP can also be a data center where the provider can deploy CSO.
SD-LAN	Software-defined local area network—Uses CSO to provision, manage, and monitor LAN segments located at remote sites. Can include EX Series switches, EX Series Virtual Chassis, and Mist WiFi access points.
SD-WAN	Software-defined wide area network—Uses CSO to provision, manage, and monitor on-premises spoke devices, provider hubs, and enterprise hubs located across a WAN environment. Typically includes the use of NFX Series Network Services Platforms and SRX Series Services Gateways.
Secure OAM	Juniper Networks security-focused implementation of operations, administration, and management (OAM) functions within CSO.
Site	Any customer location, such as an on-premises spoke, an enterprise hub, or cloud spoke.
Spoke	A tenant branch site in a hub-and-spoke topology.
Tenant	Typically an enterprise customer with many branches (sites) who subscribes to the offerings provided by the service provider. Sites are provisioned within a tenant. One tenant cannot see the sites or assets of another.
PNF	Physical Network Function—Network service provided by a physical device, such as firewall services provided by an SRX Series Services Gateway.

Table 3: CSO Terminology (*continued*)

Term	Definition
Provider Hub	<p>A multitenant hub device located in a POP on the service provider's network. A provider hub can terminate IPsec tunnels for both overlay and secure OAM networks. Provider hub devices can also terminate MPLSoGRE and MPLSoGREoIPsec tunnels. Only an SP administrator or OpCo administrator can add, modify, or delete provider hub devices.</p> <p>NOTE: In cloud-hosted versions of CSO, an OpCo administrator can add only DATA_ONLY hubs. OAM capable hubs can only be added by the SP administrator.</p>
VNF	Virtualized Network Function—Network service provided by software running in a virtual environment, such as the vSRX virtual firewall.
ZTP	Zero touch provisioning, also known as autoinstallation.

References

This guide is hosted on the [Contrail Service Orchestration Documentation](#) page along with several other guides, including:

- [CSO Deployment Guide](#)
- [CSO Installation and Upgrade Guide](#)
- [CSO Administration Portal User Guide](#)
- [CSO Customer Portal User Guide](#)
- [Contrail Service Orchestration Monitoring and Troubleshooting Guide](#)
- and more

2

PART

Introduction

Deployment Models and Building Blocks | 7

Deployment Models and Building Blocks

IN THIS CHAPTER

- Introduction | 7
- Juniper Networks Cloud CPE Solution | 10
- Contrail SD-WAN Model | 11
- Contrail SD-LAN Model | 12
- Next-Generation Firewall (NGFW) Model | 16
- Building Blocks Used for Contrail Service Orchestration Deployments | 17

Introduction

Service providers have historically deployed dedicated physical customer premise equipment (CPE) to provide enterprise services. However, enterprises no longer have the patience to wait months, even years, for a new version of a firewall or other service that can easily be offered now as a software-based solution. With the advent of network virtualization, and an increased emphasis on end-to-end provisioning, automation, analytics, and network visibility, a new system is needed to provide these services and meet the increasing demand of end customers.

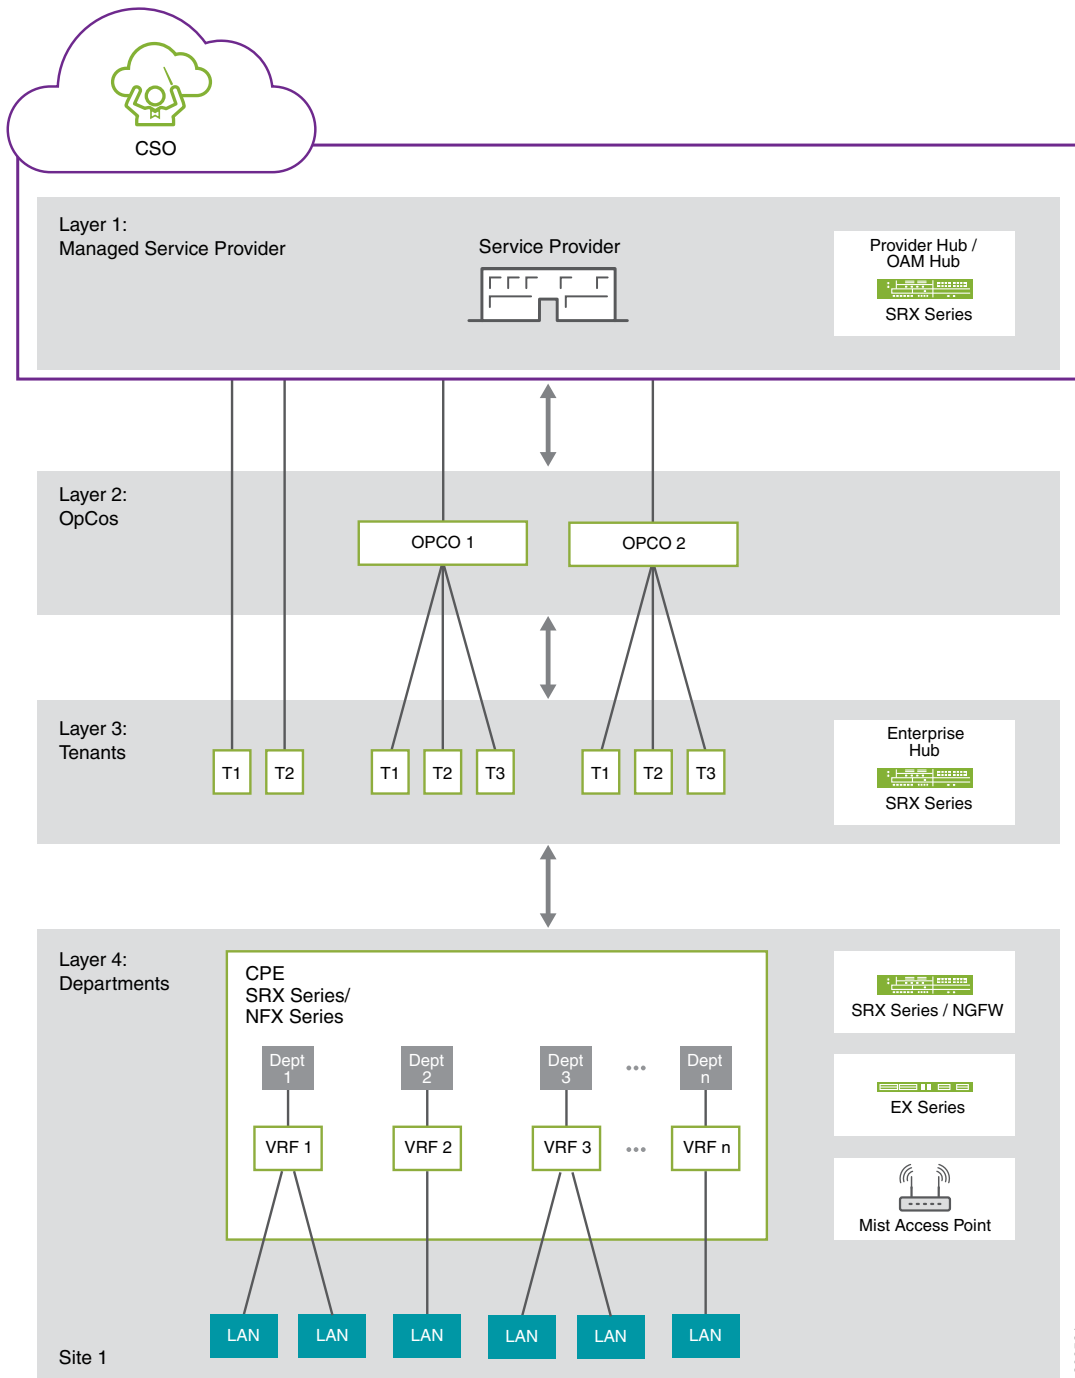
Contrail Service Orchestration is an orchestration and management platform capable of provisioning several types of WAN connectivity solutions for remote sites. CSO is available in two deployment options:

- **Cloud-hosted installation**—A Juniper-owned, software-as-a-service (SaaS) deployment. Smaller service providers called *operating companies* (OpCos) and multiple individual tenants subscribe to this cloud-based service.
- **On-premises installation**—A software package that you download and install on your servers and networks. In this model, you are considered the service provider (SP).

The cloud-hosted model reduces some of the complexity and overhead involved in managing the servers, virtual machines, and orchestration and management infrastructures needed to run CSO in an on-premises deployment. For a cloud-hosted SaaS deployment, [Figure 1 on page 8](#) shows that the service provider environment is in the cloud. Thus, all of the virtual machines and back-end infrastructure needed to run CSO are hidden from the view of the subscribers below. Juniper Networks, acting as the SP Administrator, creates and provides access to all of the top-level elements of CSO such as PoPs, provider hubs, network

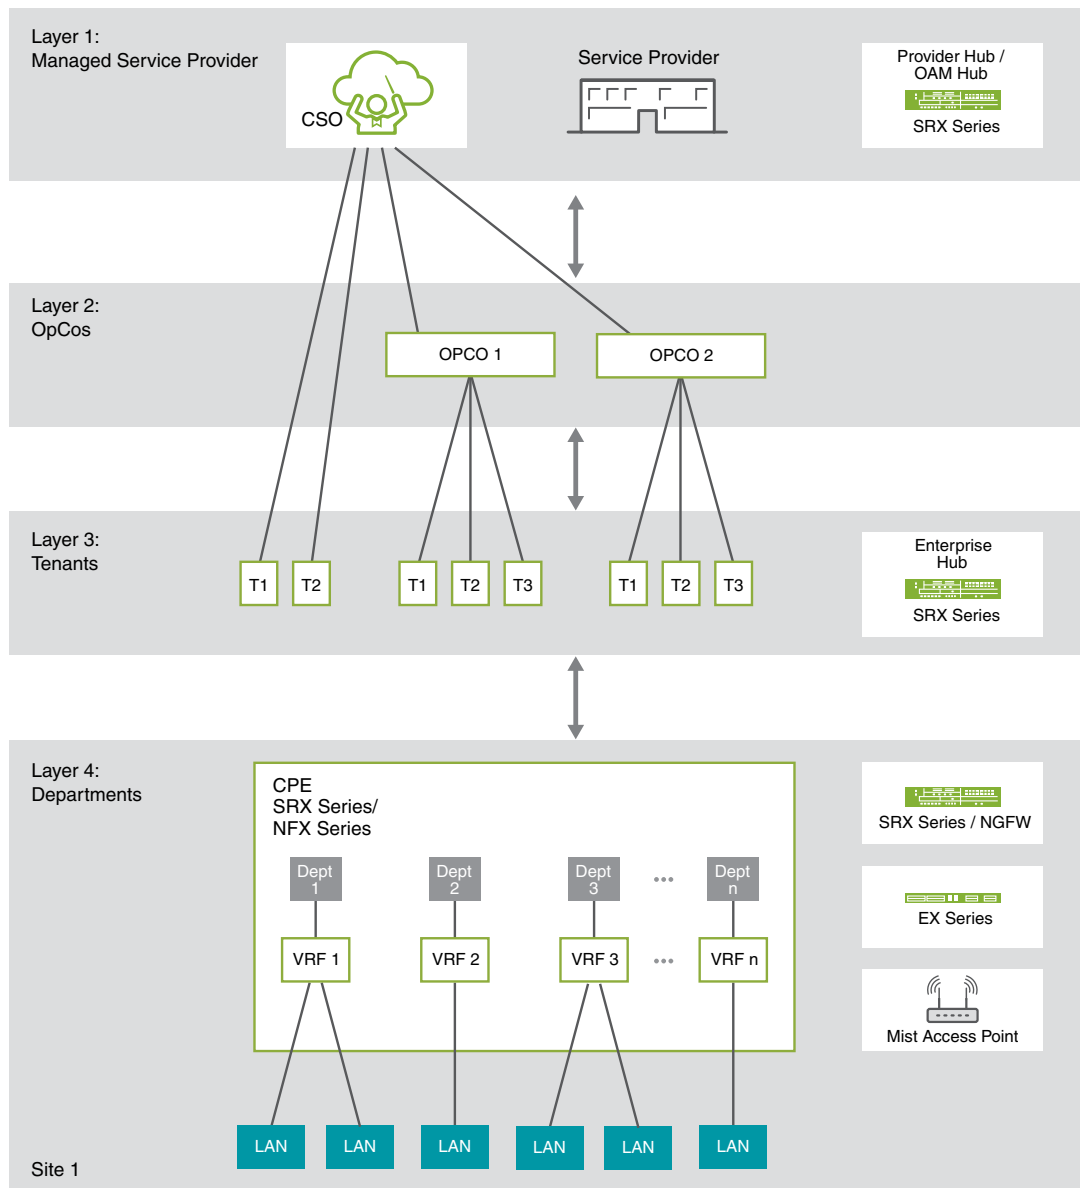
services, and Operating Company (OpCo) administrators. OpCo and tenant subscribers only need to subscribe to CSO as a service so they can manage their own infrastructures to make use of Contrail SD-WAN, SD-LAN, and NGFW solutions.

Figure 1: Cloud-Hosted CSO Deployment



For an on-premises installation, [Figure 2 on page 9](#) shows CSO deployed within a service provider (SP) network. Thus, an SP who deploys CSO using this model has full control and responsibility for the entirety of the CSO deployment. The SP Administrator has full access to create all top-level elements such as PoPs, provider hub devices, network services, OpCos, tenants, and so on.

Figure 2: On-Premises CSO Deployment



g300522

Regardless of deployment type, some of the highlights of Juniper Networks Contrail SD-WAN solution include:

- Integrated Security – Full security suite with NGFW, UTM, etc.
- Single Orchestrator – CPE zero touch provisioning, VNF deployment, managed security, SD-WAN services
- Adherence to open standards – Not a book-ended solution. Easily interoperable with existing SP and enterprise infrastructure and third-party CPEs through open APIs and protocols, with software deployable on public as well as private clouds
- Full routing and MPLS stacks – Support for BGP/OSPF/IS-IS/MPLS/VRRP, etc. on WAN and LAN; scalable architecture with distributed SD-WAN hubs
- Carrier grade appliance – Innovative branch device (NFX Series) with service chaining support for 3rd-party VNFs
- End-to-end management and orchestration – Feature rich, horizontally scalable, easy-to-use orchestration platform

Highlights of Juniper Networks Contrail SD-LAN solution include those mentioned above along with:

- LAN management and visibility – Provision, deploy, configure, and monitor EX Series LAN switches, virtual chassis (VCs) and Mist wireless access points at remote sites
- Managed LAN security – Provision, deploy, configure and monitor standalone SRX Series next-generation firewall (NGFW) devices at remote sites

The remaining sections of this guide describe the components and architectures of SD-WAN, SD-LAN, and NGFW solutions. Some sections are only applicable to cloud-hosted or on-premises deployments. There are notes throughout that explain when a topic applies only to one or the other type of deployment.

Juniper Networks Cloud CPE Solution

To understand the SD-WAN, SD-LAN, and Next Generation Firewall (NGFW) solutions, it helps to know about the Cloud CPE solution. Prior to the Cloud CPE solution, the customer-premises equipment (CPE) purchasing model required enterprises to rip out dedicated hardware devices and replace them with new hardware devices whenever they needed services and capabilities not offered by their existing devices. In contrast, the Cloud CPE solution placed modern CPE devices at remote sites that were capable of running virtualized services (such as a virtual firewall, Network Address Translation, and other services). By switching to this model, customers could add new services and software on these newer CPE devices without removing and replacing the devices themselves.

Originally, the Cloud CPE solution was focused on providing remote sites with access to virtualized network functions (VNFs) that were hosted in a service provider's (SP) cloud. These cloud-hosted network services could be deployed at strategic locations within the SP cloud to provide a better user experience to the SP's remote customers.

Advances in overlay networking have allowed Juniper's Contrail SD-WAN, SD-LAN, and NGFW solutions to further enhance the performance, security, and connectivity options for remote customers. Based on the original Cloud CPE solution, these newest solutions offer end-to-end provisioning of Layer 4 through Layer 7 network services through an open, cloud-based architecture. They support both Juniper Networks and third-party VNFs as well as physical network functions, provided by on-premises devices, that service providers can use to offer network services. They support various connectivity options for the CPE devices based on existing underlay networks. The addition of software-defined overlay networks on top of the underlay connections provides enhanced connectivity, increased performance, and the ability to use VNFs and physical network functions (PNFs) deployed directly at remote customer locations.


Network security is provided at the SP and customer sites as well as when data is in transit.

Contrail SD-WAN Model

Contrail SD-WAN extends the Cloud CPE solution to provide a flexible way to route traffic to its destination, and as a result provides the automation of Layer 3 connectivity as well as distributed Layer 4 to Layer 7 services. Based on the distributed Cloud CPE deployment model, this implementation uses intelligent CPE devices located at customer sites to connect to hub devices, using a variety of deployment models. Traffic can flow between CPE sites, and from a CPE site to the Internet.

The SD-WAN model is shown in [Figure 3 on page 11](#).

Figure 3: SD-WAN Deployment Model

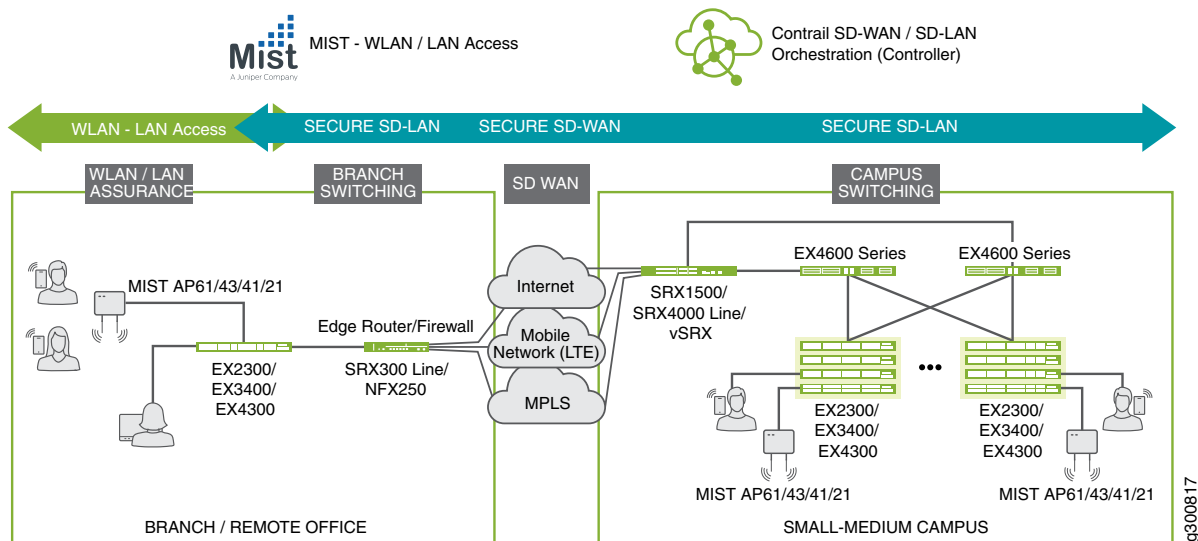


The Contrail SD-WAN solution brings SDN-type capabilities to enterprises, offering agility, automation, and rapid automated recovery from failed WAN links, while containing WAN service costs. You can add connectivity options such as broadband or cellular Internet connections to your existing IP/MPLS VPN services, allowing you to prioritize critical traffic across the connections, as well as move traffic proactively to a backup link if the primary link's quality degrades enough to put a service-level agreement (SLA) at risk.

Contrail SD-LAN Model

Contrail SD-LAN extends the Contrail SD-WAN solution by allowing tenant administrators to deploy, manage, and maintain LAN devices such as EX Series LAN access switches, EX Series Virtual Chassis, and Mist wireless access points. This allows you to monitor remote-site LANs and manage deployed access switches, Virtual Chassis, and WLAN access points as shown in [Figure 4 on page 12](#).

Figure 4: SD-LAN Overview

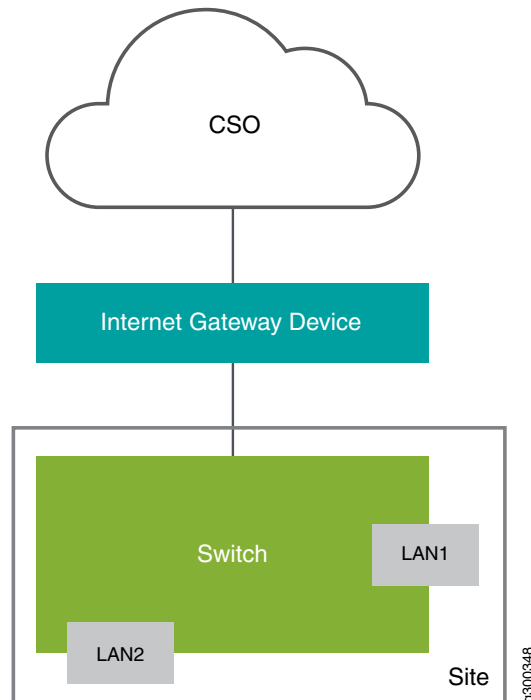


As shown in the diagram, you can deploy EX Series switches or Virtual Chassis at a branch site behind an on-premises spoke device. You can also place them in a campus or data center environment behind a provider hub or enterprise hub. In either case, you can add Mist WiFi access points to the LAN segments. Such flexibility enables you to use CSO to provision, manage, and monitor LAN segments across your WAN infrastructure.

There are several ways that you can provision the EX Series switches:

- By adding a site containing one or more switches and connecting the site to an Internet gateway device, as shown in [Figure 5 on page 13](#).

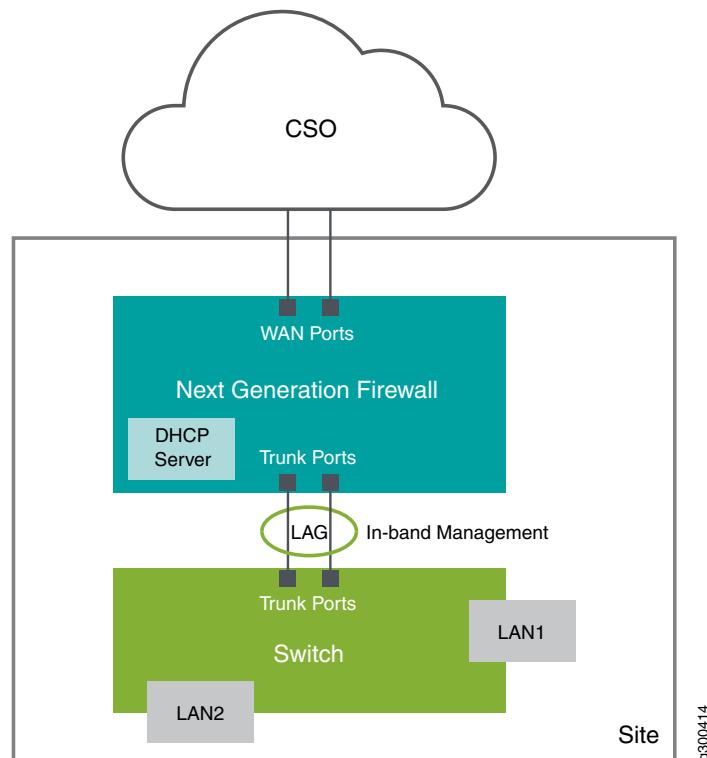
Figure 5: Switch Behind Internet Gateway Device



- By adding a site that contains both an NGFW and a switch.
- By adding a switch to an NGFW site that is already provisioned and managed by CSO.

For either of these NGFW scenarios, [Figure 6 on page 14](#) shows where you place the switch in the network.

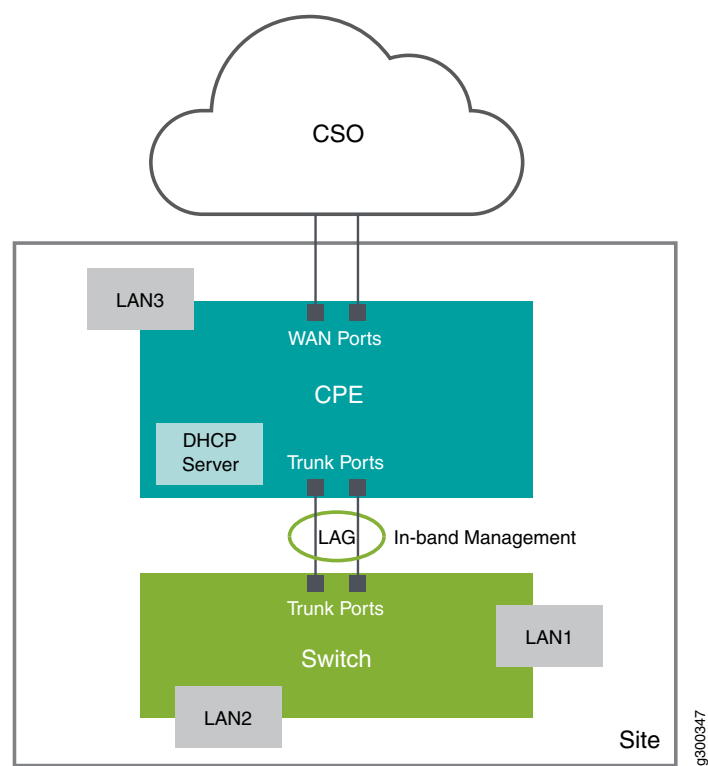
Figure 6: Switch Behind NGFW



- By adding a site with an SD-WAN CPE device and a switch.
- By adding a switch to an SD-WAN CPE device that is already provisioned and managed by CSO.

For either of these SD-WAN CPE scenarios, [Figure 7 on page 15](#) shows a LAN switch behind a CPE device.

Figure 7: Switch Behind CPE Device



The EX Series switches that you can deploy as SD-LAN switches include the EX2300, EX3400, EX4300, EX4600, and EX4650 switches and Virtual Chassis.

In any of these scenarios using SD-LAN, you can monitor your LAN switch or Virtual Chassis for the following:

- Resource utilization
- Status of ports
- Alerts and alarms generated on the switch
- Ports consuming the maximum bandwidth
- Ports with the maximum number of errors
- Ports with the maximum packet loss

Virtual Chassis support on EX Series switches is shown in [Table 4 on page 15](#).

Table 4: EX Series Virtual Chassis Support

EX Series Switches	Number of Virtual Chassis Members
EX2300	4

Table 4: EX Series Virtual Chassis Support (*continued*)

EX Series Switches	Number of Virtual Chassis Members
EX3400	10
EX4300	10
EX4600	10
EX4650	2

NOTE: CSO does not support mixed-mode virtual chassis configurations.

The Virtual Chassis must be pre-built with master, backup, and the minimal number of other member devices before provisioning in CSO. Only the master member's serial number is needed during the initial provisioning of a Virtual Chassis. For details about creating EX Series Virtual Chassis, see:

- *Configuring an EX2300, EX3400, or EX4300 Virtual Chassis*
- *Configuring EX4600 Switches in a Mixed or Non-Mixed Virtual Chassis*
- *Configuring an EX4650 or a QFX Series Virtual Chassis*

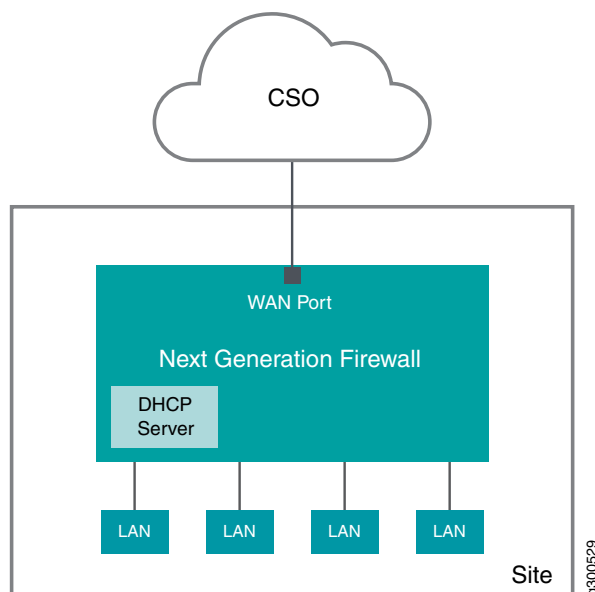
Adding Mist WLAN access points to your managed EX switches at remote sites allows CSO to automatically manage the access points. This gives CSO even more visibility into your remote sites.

See the [CSO Deployment Guide](#) for more details.

Next-Generation Firewall (NGFW) Model

You can deploy a standalone next-generation firewall (NGFW) device at remote branch spoke sites. NGFW deployment provides remote network security through the use of SRX Series devices as customer-premises equipment (CPE) at a spoke site. This solution offers managed security and LAN visibility to a single location without providing CSO-managed site-to-site connectivity or VNFs, like the Contrail SD-WAN solution provides. [Figure 8 on page 17](#) shows a simplified NGFW deployment topology.

Figure 8: Standalone NGFW



Building Blocks Used for Contrail Service Orchestration Deployments

IN THIS SECTION

- Administrators | 18
- Portals | 18
- Tenants | 19
- Topologies | 19
- Points of Presence (POPs) | 22
- Sites | 23
- On-Premises Spoke Devices | 28
- Standalone Next-Generation Firewall (NGFW) | 29
- SD-LAN Devices | 29
- Virtual Route Reflector (VRR) | 29
- SLA-Based Steering Profiles and Policies | 30
- Path Based Steering Profiles | 31
- Intent-based Firewall Policies | 31
- Software Image Management | 31

Contrail Service Orchestration (CSO) uses conceptual and logical elements as building blocks to complete deployments in the GUI. This document provides some discussion about those elements and their use in CSO. For more detailed discussions regarding these elements, see the *Contrail Service Orchestration Administration Portal User Guide* and *Contrail Service Orchestration Customer Portal User Guide* available at https://www.juniper.net/documentation/product/en_US/contrail-service-orchestration..

Administrators

CSO uses a hierarchical, domain-based administration framework. After CSO installation, the first administrator is named *cspadmin* by default. This administrator is also known as the global service provider (SP) administrator. This SP administrator has full read and write access to all of the CSO platform from the global domain. In a cloud-hosted CSO deployment, the *cspadmin* role is reserved for Juniper Networks. The SP administrator can create, edit, and delete other administrators and operators who are subject to role-based access controls (RBAC) that assign them privileges to the rest of the objects in CSO.

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

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

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

Portals

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

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

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

For more information about Administration and Customer Portals, see the *Contrail Service Orchestration Administration Portal User Guide* and *Contrail Service Orchestration Customer Portal User Guide* available at https://www.juniper.net/documentation/product/en_US/contrail-service-orchestration.

Tenants

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

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

Topologies

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


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

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

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

This is the topology of a simple SD-WAN, solution wherein the SP provides network services to its on-premises customers but does not allow them to communicate with one another. [Figure 9 on page 20](#) shows an example where the virtual network functions (VNFs) are located at an on-premises site, but the site has no access to other sites belonging to the tenant.

Figure 9: Distributed CPE / Simple SD-WAN



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

It is also the topology of the NGFW and SD-LAN solutions. The NGFW solution provides for remote site security with SRX Series next-generation firewall devices. The SD-LAN solution provides for remote site LAN management with EX Series LAN access switches and Virtual Chassis. [Figure 10 on page 20](#) and [Figure 11 on page 21](#) below show high-level examples of these two solutions.

Figure 10: Standalone NGFW

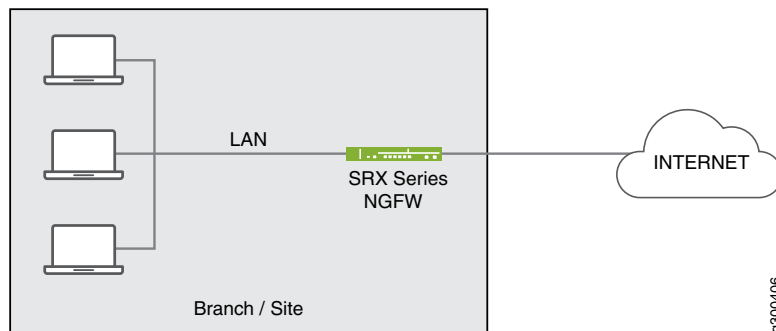




Figure 11: SD-LAN Solution with EX Switch



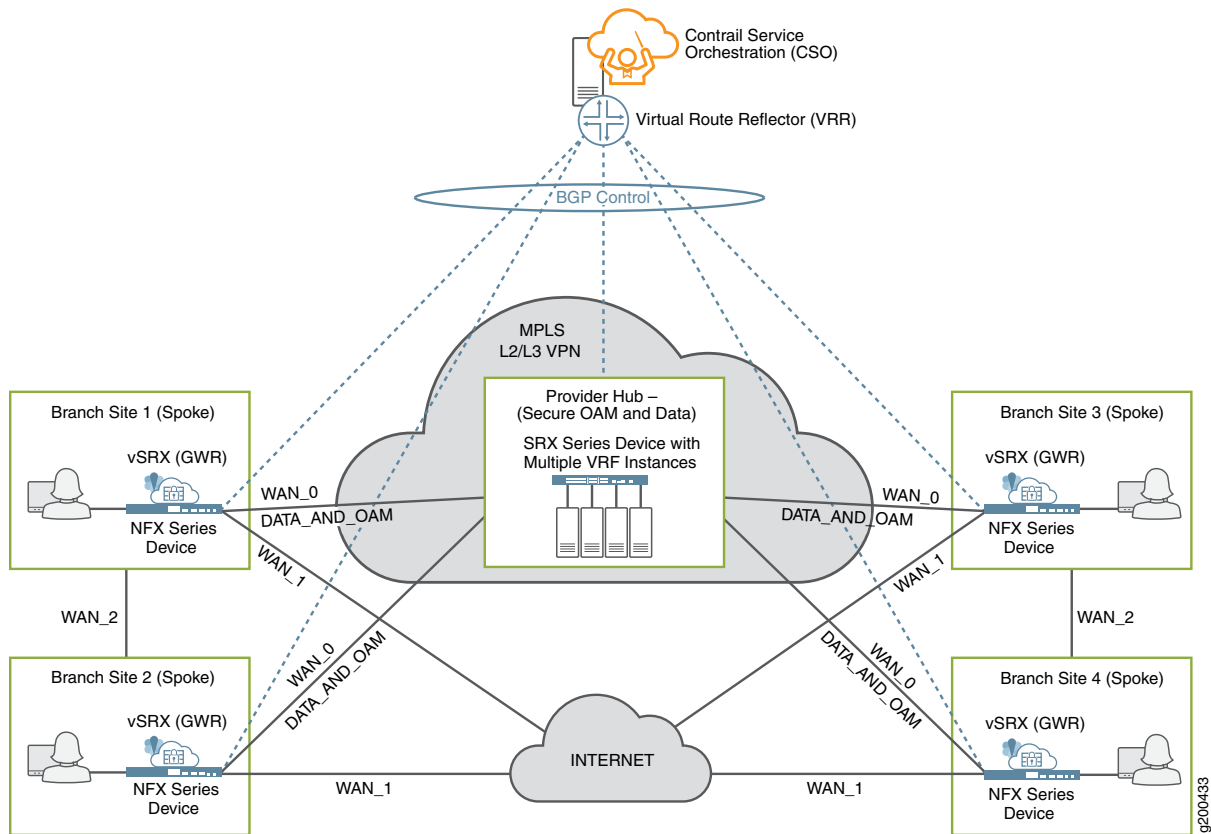
- Hub-and-Spoke Topology**—This topology is available for SD-WAN deployments. Given that SD-WAN is intended specifically to enable and enhance the efficacy of WAN communication using network overlays, this topology does allow for communication from site to site. Specifically, if one site needs to communicate with another site, that communication goes through the hub on its way to the other site. [Figure 12 on page 21](#) shows a very basic example of a hub-and-spoke topology. VNFs can be deployed at any of the locations shown.

Figure 12: Hub-and-Spoke Topology



- Dynamic Mesh Topology**—This topology is also available for SD-WAN deployments. Direct site-to-site communication is allowed. [Figure 13 on page 22](#) shows a very basic example of a full mesh topology. VNFs can be deployed at any of the branch locations shown. This topology requires more overlay networks than the hub-and-spoke topology so CSO allows for the creation of a full mesh topology as a construct, but the tunnels from one site to another are created dynamically, (or on-demand) based on traffic thresholds thereby conserving resources and improving overall performance.

Figure 13: Dynamic Mesh Topology




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

Points of Presence (POPs)

A POP is a placeholder, usually at the telco edge or enterprise datacenter, where network services can be deployed and underlay network connections are made to remote sites, as shown in [Figure 14 on page 23](#). POPs can have PE routers and provider hubs (both data and OAM type).

Figure 14: Points of Presence (POPs)



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

Sites

Sites are the branch offices or remote locations from which customers access the network services provided by the CSO solutions. A site is assigned to a POP and the type of sites available depend on the type of deployment you are creating: SD-WAN, SD-LAN, or NGFW. Sites are created by the Tenant administrator. Sites can be named whatever makes sense for the Tenant. [Table 5 on page 24](#) lists what types of sites can be created within each deployment.

Table 5: Site Types by Deployment

Deployment	Available Site Types	Uses	Service Notes
SD-WAN	On-Premises Spoke	NFX Series or SRX Series devices placed at customer sites in either a hub-and-spoke or full mesh topology.	<p>An on-premises spoke has the following capabilities:</p> <p>SRX Series</p> <ul style="list-style-type: none"> • The SRX300 Line of Services Gateways support ADSL and LTE interfaces. • The SRX1500 and SRX300 Line of Services Gateways support PPPoE on WAN Ethernet interfaces. • SRX Series devices deployed as on-premises spoke devices cannot host VNF-based network services. <p>NFX Series</p> <ul style="list-style-type: none"> • NFX Series devices used as on-premises spoke devices support ADSL, VDSL, and LTE access links, which can also be used for ZTP. The DSL access links allow configuration of PPPoE. Starting with CSO Release 4.0, LTE access links can be used as primary DATA, OAM, or DATA_OAM links. • NFX Series devices support PPPoE on WAN Ethernet interfaces. • Supports local breakout when using a dynamic mesh topology. <p>NOTE: ZTP using an xDSL interface will not work if the link is PPPoE. If the link is bridged and uses DHCP, then ZTP will work on xDSL interfaces.</p>
	Cloud Spoke	vSRX placed in a tenant's Amazon Web Services (AWS) Virtual Private Cloud (VPC)	

Table 5: Site Types by Deployment (*continued*)

Deployment	Available Site Types	Uses	Service Notes
			<p>A cloud spoke has the following capabilities:</p> <ul style="list-style-type: none"> • Firewall and UTM services are available to protect the customer's resources in an AWS VPC. • Connectivity between VPC resources and on-premises sites. • WAN_0, WAN_1, and LAN interfaces need to be predefined in the VPC. • Two elastic IP addresses need to be reserved in the VPC to attach to WAN interfaces later. • VPC should be created and attached to an Internet gateway. • Only a hub-and-spoke topology is supported.

Table 5: Site Types by Deployment (continued)

Deployment	Available Site Types	Uses	Service Notes
			<ul style="list-style-type: none"> • The hub must have public IP addresses on its WAN interfaces. • The hub WAN interface type should be set as Internet during onboarding.
	Provider Hub	<p>SRX Series devices placed in a central role in a service provider cloud. The hub devices establish IPSec tunnels with the spoke sites. Provider hub devices are multi-tenant (shared amongst multiple sites) through the use of VRF instances configured on them.</p> <p>In a cloud-hosted CSO deployment, an OpCo can create Provider Hub sites, but not the hub devices themselves. In this case, available hub devices are created by the SP administrator and made available to the lower-level administrators. Tenant administrators are not allowed to add Provider Hub sites.</p>	<p>A provider hub has the following capabilities:</p> <ul style="list-style-type: none"> • You must specify the capability of the hub devices when setting up the site. Specifying OAM capabilities (OAM Hub) allows the hub to help create secure OAM networks with the CPE devices. This option is only available for on-premises CSO deployments. • For cloud-hosted CSO, data hub sites can be added only by an OpCo. • A hub device is required for the dynamic mesh topology. • Local breakout is not supported on hub sites.
	Enterprise Hub	Provides additional hub-like capabilities to a normal spoke site.	

Table 5: Site Types by Deployment (*continued*)

Deployment	Available Site Types	Uses	Service Notes
			<p>An enterprise hub has the following capabilities:</p> <ul style="list-style-type: none"> • Can behave as a normal spoke. • Acts as an anchor point for spokes for dynamic VPN creation. • Provides an on-premises central breakout option. • Can host a data center department. • Can import BGP and OSPF routes from the LAN-side L3 device to create a data center dynamic LAN segment. • Automatically meshed with other enterprise hubs that belong to the same tenant.

Table 5: Site Types by Deployment (*continued*)

Deployment	Available Site Types	Uses	Service Notes
			<ul style="list-style-type: none"> Regular spoke sites can be assigned to associate with enterprise hubs. Supports local, central, and cloud breakout profiles with intent-based rules for more granular breakout control.
SD-LAN	Switch	Used to access EX Series switches in a branch location.	CSO can manage EX Series switches located behind an SD-WAN spoke site, an NGFW security device, or a third-party Internet gateway.
	Access Point	CSO automatically recognizes any Mist WiFi access points attached to a switch in a branch location.	For SD-WAN and NGFW environments, CSO can detect and manage Mist access points located behind a switch.

On-Premises Spoke Devices

On-premises spoke devices (also known as customer premises equipment [CPE]) are devices placed at remote locations in the site types mentioned previously. You can deploy on-premises spoke devices in SD-WAN environments. [Figure 15 on page 28](#) shows available on-premises spoke device types.

Figure 15: On-Premises Spoke Devices



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

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

Standalone Next-Generation Firewall (NGFW)

SRX Series devices can be used as standalone firewalls, managed by CSO in the customer LAN. CSO supports the use of SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100, and SRX4200 Security Gateways as well as the vSRX for this purpose. In this next-generation firewall (NGFW) scenario, the SRX acts as a CPE device but provides no site-to-site or site-to-hub communications as with an SD-WAN solution.

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

SD-LAN Devices


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

In addition, CSO supports dynamic routing protocols such as BGP and OSPF in the local LAN. Therefore, when SD-LAN is configured using any of the above devices, routes to the site LANs can be updated dynamically with BGP or OSPF.

Virtual Route Reflector (VRR)

The VRR is part of CSO's SD-WAN controller. It is one of the virtual machines that gets 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. You make this selection with the **Configure Site** workflow for site onboarding. The OAM interfaces can be implemented using dedicated IPsec tunnels which allows CPE devices to be behind NAT. [Figure 16 on page 30](#) illustrates the concept of the VRR

Figure 16: VRR Overview



SLA-Based Steering Profiles and Policies

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

You can set:

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

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

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

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

Path Based Steering Profiles

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

Intent-based Firewall Policies

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

See *Firewall Policy Overview* for more information.

Software Image Management

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

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

3

PART

Contrail Service Orchestration Solutions

SD-WAN Solution | 33

SD-LAN Solution | 62

NGFW Solution | 70

SD-WAN Solution

IN THIS CHAPTER


- Overview | 33
- SD-WAN Requirements | 34
- Contrail SD-WAN Solution Highlights | 35
- Target Customers | 35
- Contrail SD-WAN Deployment Architectures | 36
- Two Basic SD-WAN Use Cases | 55
- Secure and Redundant OAM Network | 58

Overview

Typical branch offices can have multiple WAN connection types, including MPLS, Internet (such as LTE and ADSL), and so on. In these traditional networks, the MPLS-based connection typically includes performance guarantees known as *service level agreements* (SLAs) to ensure business-critical applications can function properly. The Internet connection often provides an alternative link for backup and load balancing purposes. However, with Internet access offerings providing ever-increasing bandwidth, many applications can now reasonably be routed over the Internet link.

Software-defined wide area networking (SD-WAN) is primarily thought of as a connectivity solution, implemented as an overlay on top of traditional WAN access. An SD-WAN solution provides the ability to make use of the links in whichever way an enterprise customer wishes, as shown in [Figure 17 on page 34](#).

Figure 17: High-Level View of SD-WAN Interconnectivity



In an SD-WAN environment, low-priority traffic can use the lower-cost Internet link(s), while more important traffic can travel across better quality links (such as those provided by an MPLS network). Link usage can also be assigned per application. With an SD-WAN solution, an enterprise customer can mix and match cost optimization with SLA requirements as they see fit.

SD-WAN Requirements

The key components of an SD-WAN solution center around application awareness, visibility, and performance. An SD-WAN solution must generally provide the following types of functionality:

- Multiple connection types – MPLS, Internet, LTE, ADSL, etc.
- Secure site-to-site connectivity - tunneling and VPNs
- An intuitive interface for managing WAN connections
- Ability to make use of all available uplink paths
- Ability to optimize use of WAN connection for cost savings
- Application-aware performance monitoring over WAN links
- Dynamic spoke/endpoint learning and reachability

In addition, modern SD-WAN solutions have evolved to offer even broader capabilities, including:

- Automation of end-to-end solution provisioning
- Enterprise network modeling, network definition
- Zero touch provisioning (ZTP) of on-premises devices, including establishing connectivity

- Provisioning of multiple node types (spoke, hub, concentrators, etc.)
- Dynamic path selection, and ability to load balance across multiple WAN connections
- End-to-end, application-level SLAs through continuous path measurement
- Dynamic application steering to counteract link degradation
- End-to-end visibility and monitoring of devices, connectivity, and application performance
- Support for 3rd-party services
- Intent-based policy creation to define traffic treatment
- Security through enterprise-wide policies

Juniper Networks Contrail SD-WAN solution, as described in this document, provides a full, end-to-end solution that aims to offer all of the above functionality.

Contrail SD-WAN Solution Highlights

Highlights of Juniper's Contrail SD-WAN solution include:

- Integrated Security – full security suite with NGFW, UTM, etc.
- Single Orchestrator – CPE zero touch provisioning, VNF deployment, managed security, SD-WAN services
- Adherence to open standards – not a lock-ended solution, easily interoperable with existing SP/enterprise infrastructure and third-party CPEs through open APIs and protocols, with software deployable on public as well as private clouds
- Full routing and MPLS stacks - support for BGP/OSPF/IS-IS/MPLS, etc. on WAN/LAN, scalable architecture with distributed SD-WAN gateways
- Support for BGP on the underlay networks
- Carrier grade appliance – innovative branch device (NFX Series) with service chaining support for 3rd-party VNFs
- End-to-end management/orchestration – feature rich, horizontally scalable, easy-to-use orchestration platform

Target Customers

The Contrail SD-WAN on-premises deployment is architected to address the following types of customers:

- Service providers with existing MPLS networks
 - Offers ways to provide SD-WAN as a value-added offering that leverages existing MPLS core networks
- Managed service providers (also known as OpCos or MSPs), building networks on top of service providers and offering end-to-end solutions (first-mile, middle-mile, last-mile)
 - Provides logical separation from underlay provider using overlay networks
- Large enterprises, trying to build their own end-to-end overlay-based network
 - Enables independence from underlay/transport provider
 - More control and agility to meet network requirements

In addition to the benefits above, the cloud-hosted, software-as-a-service (SaaS) deployment of CSO is designed for broadly distributed enterprise and MSP customers who:

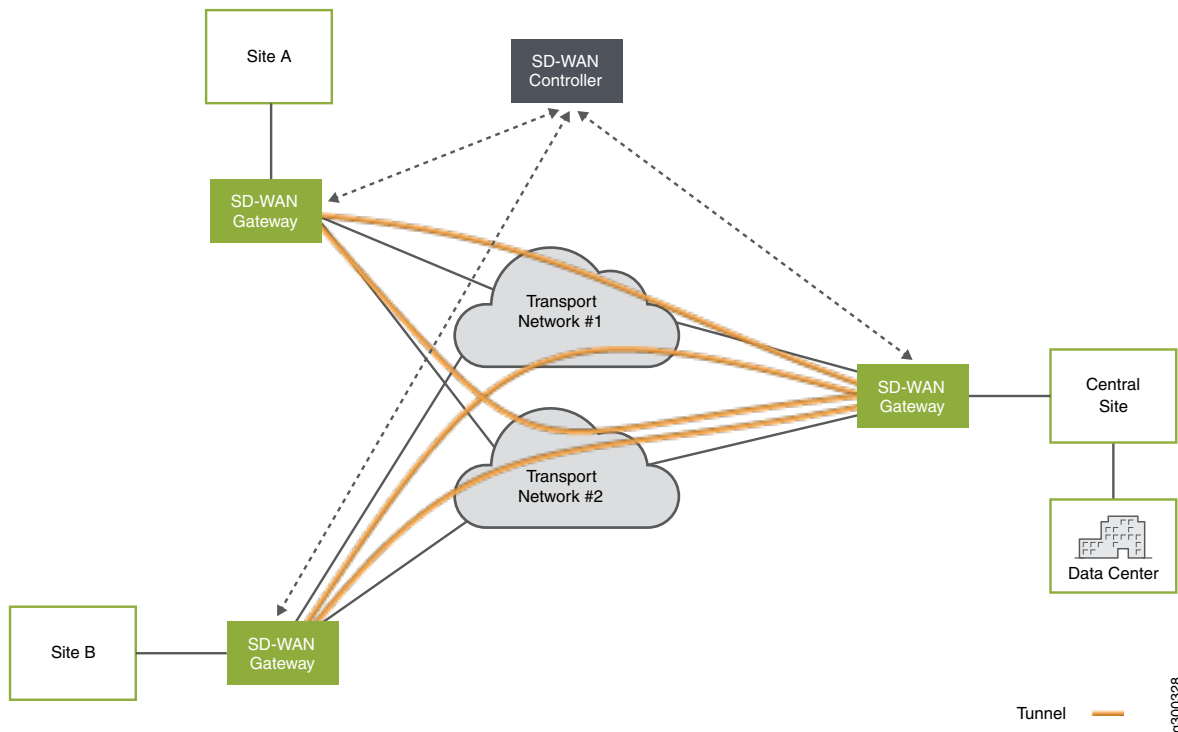
- Prefer a cloud consumption model
- Demand ease of use and quick turn up of cloud enabled branch services
- Don't have the cloud Infrastructure to host and operate an SD-WAN solution
- Have the business need to consume cloud services in a pay as you grow SaaS model
- Have limited or no IT personnel on-site in remote branches

Contrail SD-WAN Deployment Architectures

An SD-WAN implementation offers a flexible and automated way to route traffic from site to site. As shown in [Figure 18 on page 37](#), a basic SD-WAN architecture includes just a few basic elements

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

Figure 18: SD-WAN Architecture

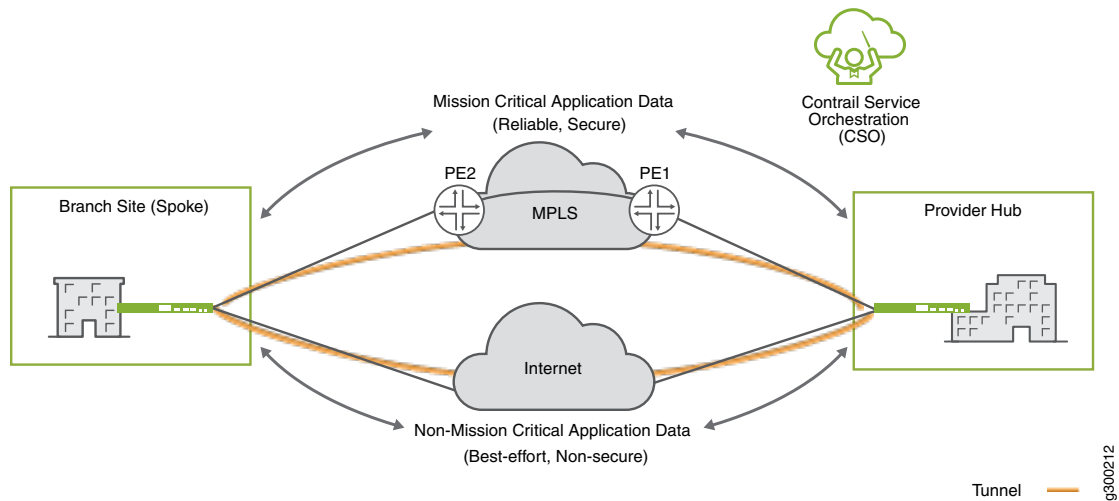


The SD-WAN controller, built in to CSO, acts as an orchestration layer and provides an interface, allowing the operator to setup and manage the devices at the sites.

Contrail SD-WAN Reference Architecture

Juniper Networks Contrail SD-WAN solution architecture, shown in [Figure 19 on page 38](#) uses a hub-and-spoke topology, with CPE devices located at customer branch sites. On the local side of the site, the CPE devices connect to LAN segments and participate in dynamic routing protocols with other LAN devices. On the WAN side, the CPE devices connect across two or more links to a provider hub device. Because the SD-WAN model uses a hub-and-spoke topology, traffic travels from site to site through the provider hub device. By default, traffic going to the Internet also flows through the provider hub device.

Figure 19: Contrail SD-WAN Reference Architecture



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

The following sections describe these architectural elements in more detail.

Spoke Devices

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

On-Premises Spoke Devices

On-premises spoke devices can be either NFX Series devices or specific SRX Series devices, as shown in [Figure 20 on page 38](#).

Figure 20: On-Premises Spoke Devices



NFX Series Network Services Platform

An NFX Series Network Services Platform used as an on-premises spoke device can host a range of multivendor VNFs, support service chaining, and be managed by orchestration software in the cloud. NFX Series devices eliminate the operational complexities of deploying multiple physical network devices at a customer site and offer a substantial improvement over traditional, single function CPE devices.

A key VNF supported on the NFX Series platform is the vSRX Virtual Firewall. In the Contrail SD-WAN solution, the vSRX instance with routing and switching capabilities performs the gateway router function. It also provides the same feature-rich security services found on standard SRX series devices.

[Table 6 on page 39](#) shows the NFX Series hardware that you can implement as an on-premises spoke device.

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

Table 6: NFX Series Hardware – On-Premises Spoke Devices

Platform	Models Supported
NFX150 Network Services Platform	<ul style="list-style-type: none"> • NFX150-S1 • NFX150-S1E • NFX150-C-S1 • NFX150-C-S1-AE/AA • NFX150-C-S1E-AE/AA
NFX250 Network Services Platform	<ul style="list-style-type: none"> • NFX250-LS1 • NFX250-S1 • NFX250-S2

SRX Series Devices and vSRX Virtual Firewalls

A physical SRX Series security device can be used in place of the NFX Series platform to provide the gateway router function, as can a vSRX instance installed on a server. [Table 7 on page 40](#) shows the SRX hardware and vSRX virtual firewalls that you can implement as on-premises spoke devices.

Table 7: SRX Series Hardware and vSRX Firewalls – On-Premises Spoke Devices

Platform	Models Supported
SRX Series	<ul style="list-style-type: none"> • SRX4200 • SRX4100 • SRX1500 • SRX550M • SRX345 • SRX340 • SRX320 • SRX300
vSRX Virtual Firewalls	vSRX vSRX 3.0

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

Cloud Spoke Devices

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

Spoke Redundancy

Two redundant CPE devices can be used at spoke sites to protect against device and link failures. For more detail, see the Resiliency and High Availability section. of the [Contrail SD-WAN Design and Architecture Guide](#).

Provider Hub Devices

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

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

Provider Hubs

In a service provider (SP) environment, the service provider hosts a *provider hub* device in their network. The provider hub device acts as a point of presence (POP) or connection point. It is typically a shared

device, providing hub functionality to multiple customers (tenants) through the use of virtual routing and forwarding instances (VRF). The SP administrator and the OpCo administrator can both manage the provider hub device.

In the cloud-hosted deployment of CSO, the SP administrator role is performed by Juniper Networks as the cspadmin user (or equivalent). The OpCo administrator role can be assigned to a user by the SP administrator, but the OpCo administrator does not have SP administrator privileges.

[Figure 21 on page 41](#) and [Table 8 on page 41](#) show the provider hub devices supported in a CSO SD-WAN environment.

Figure 21: SD-WAN Provider Hub Devices




Table 8: Provider Hub Devices

Role	Supported Device Types
Provider Hub	<ul style="list-style-type: none">• SRX4200• SRX4100• SRX1500• vSRX• vSRX 3.0

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

Provider Hub Redundancy

Two redundant provider hub devices can be used at one POP to protect against device and link failures, and to provide upstream multi-homing for spoke sites. For more detail, see the [“Resiliency and High Availability” on page 111](#) topic in this guide.

Enterprise Hub Sites and Devices

A special type of spoke device, called an *enterprise hub device*, can be deployed as the CPE at an on-premises spoke site. SRX1500, SRX4100, and SRX4200 devices can serve this function. The spoke site that functions

this way, must be configured as an *enterprise hub site* during site creation. Creating an enterprise hub site opens additional functionality for the site:

- Can act as the anchor point for site-to-site communications on the customer’s network.
- Can act as the central breakout node for the customer’s network.
- Offers a specialized department called the *data-center department*.
- Supports data center dynamic LAN segments with BGP and OSPF route imports, including default routes, from the LAN-side Layer 3 device.
- Allows for intent-based breakout profiles to create granular breakout behavior based on department, application, site, and so on.

In an enterprise environment, the enterprise hub is owned by the customer (tenant) and usually resides within an enterprise data center. Only the customer’s spoke sites can connect to the enterprise hub device. OpCo administrators and tenant administrators can manage the enterprise hub. [Figure 22 on page 42](#) and [Table 9 on page 42](#) show the enterprise hub devices supported in a CSO SD-WAN environment.

Figure 22: SD-WAN Enterprise Hub Devices




Table 9: Enterprise Hub Devices

Role	Supported Device Types
Enterprise Hub	<ul style="list-style-type: none">• SRX4200• SRX4100• SRX1500• vSRX• vSRX 3.0


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

Underlay (Physical) Network

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

Figure 23 on page 43 shows a sample underlay network for a hub-and-spoke SD-WAN deployment (the details apply equally to a dynamic mesh setup). Each spoke site typically has multiple paths to the hub site; in this case, one through the private MPLS cloud, and one over the Internet.

Figure 23: SD-WAN Underlay Network



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

Note that:

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

WAN Access Options

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

- MPLS
- Ethernet
- LTE

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

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

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

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

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

Full-cone and restrictive NAT deployments are supported.

Dual CPE configurations are not supported.

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


- ADSL/VDSL (ADSL/VDSL support for WAN links and ZTP on NFX Series devices starting in CSO Release 4.0.0, and ADSL support on the SRX300 Line of Services Gateways starting in CSO Release 5.2.0.)
- Broadband
- MPLS and broadband
- Satellite link

WAN Interface Types - Data and OAM

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

Figure 24 on page 44 illustrates these two interface types.

Figure 24: WAN Interface Types



Note that:

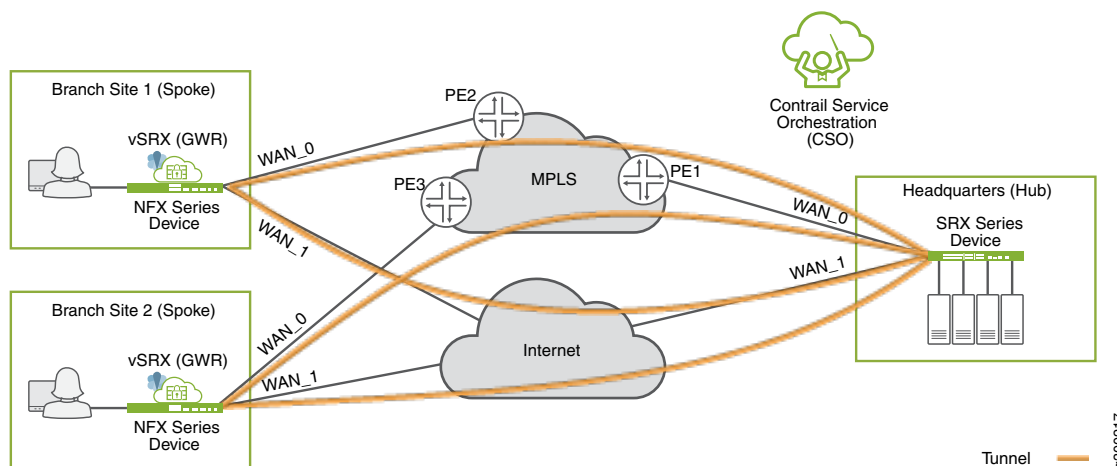
- The on-premises device's OAM interface must be able to reach CSO. The connectivity can be supplied strictly using CSO-orchestrated overlay networks. You do not need pre-existing underlay network connectivity for this. Starting in CSO release 5.0.1, CSO automatically selects an IP address for the on-premises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

Overlay (Tunnels) Network

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

Figure 25 on page 45 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 25: SD-WAN Hub-and-Spoke Overlay



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

Overlay Deployment Topologies


The SD-WAN solution supports hub-and-spoke or dynamic mesh deployment topologies. A dynamic mesh topology is similar to a full mesh topology wherein every site is capable of connecting directly to

any other site. But with dynamic mesh, the connections (tunnels) are brought up on-demand, thereby reducing the continual load on any one site. A single tenant can support both hub-and-spoke and dynamic mesh topologies.

Hub and Spoke

With the hub-and-spoke topology, all spoke sites are connected to at least one hub site, as shown in [Figure 26 on page 46](#). Spoke sites cannot communicate directly with other spoke sites.

Figure 26: SD-WAN Hub-and-Spoke Topology




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

Dynamic Mesh

With the dynamic mesh topology, overlay tunnels between participating sites enable the sites to communicate directly with each other, as shown in [Figure 27 on page 47](#). 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 27: SD-WAN Dynamic Mesh Topology



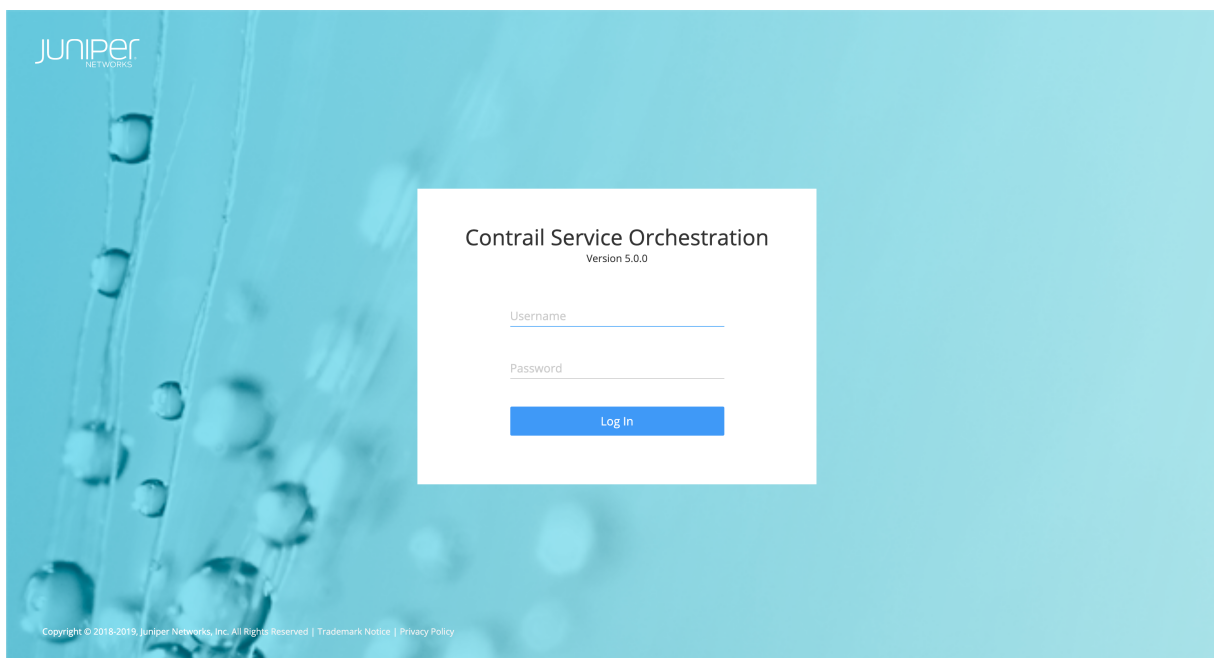
This topology is well suited for deployments where applications and services are not centralized.

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

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

Orchestration and Control

Orchestration and controller functions are implemented through Juniper's Contrail Service Orchestration (CSO) software. CSO software offers a Web-based UI to manage the SD-WAN environment, as shown in [Figure 28 on page 48](#).

Figure 28: CSO Login Screen


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

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

Secure OAM Network

SD-WAN deployments include a secure OAM overlay network to provide end-to-end secure communications between on-premises devices and CSO. This is true regardless of whether your CSO software is deployed on-premises or as a cloud-hosted deployment. In a cloud-hosted deployment, the provider hub devices, and thus, one end of the OAM network is owned and managed by the SP. As shown in [Figure 29 on page 49](#), dedicated, IPsec-encrypted OAM tunnels enable on-premises devices to send management, routing, and logging traffic securely over the network to a provider hub. The provider hub then forwards the traffic to CSO.

Figure 29: Secure OAM Tunnels



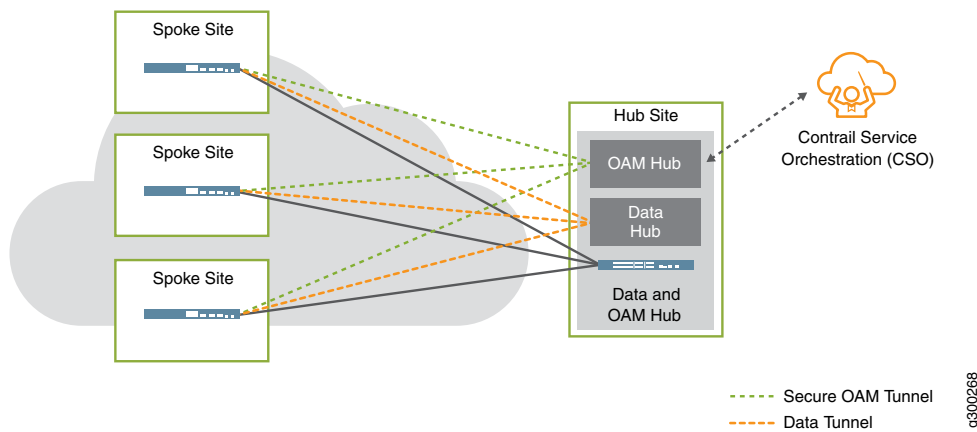
Integration with Deployment Topologies

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

Hub and Spoke

With the hub-and-spoke topology, each spoke site now has two sets of connections to the provider hub site: an overlay tunnel carrying data, and a separate, dedicated IPsec overlay tunnel carrying OAM traffic, as shown in [Figure 30 on page 49](#).


Figure 30: OAM Tunnels in the Hub-and-Spoke Topology



Dynamic Mesh

Since a normal full mesh topology would not include a hub device for data traffic, one must be added. As shown in [Figure 31 on page 50](#), each spoke site has a new connection: a separate, dedicated IPsec overlay tunnel carrying OAM traffic to the provider hub.

Figure 31: OAM Tunnels in the Full Mesh Topology




OAM Hub Design Options

There are two ways to implement the OAM hub in an on-premises CSO deployment, depending on design requirements. As shown in [Figure 32 on page 51](#), the options are as follows:

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

Figure 32: OAM Tunnels - Provider Hub Design Options



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

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

Usage Notes on Provider Hub Design Options

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

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

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

Zero Touch Provisioning

One of the key features of the Contrail SD-WAN solution is the ability to “plug-and-play” new spoke devices using ZTP (autoinstallation). The following is a high-level list of steps performed during ZTP:

- If you implement the on-premises version of CSO, you need to add the appropriate CSO SSL certificate to the redirect server before performing ZTP.

NOTE: If you deploy the cloud-delivered version of CSO, Juniper Networks configures the redirect server for you.

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

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

Usage Notes for ZTP

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

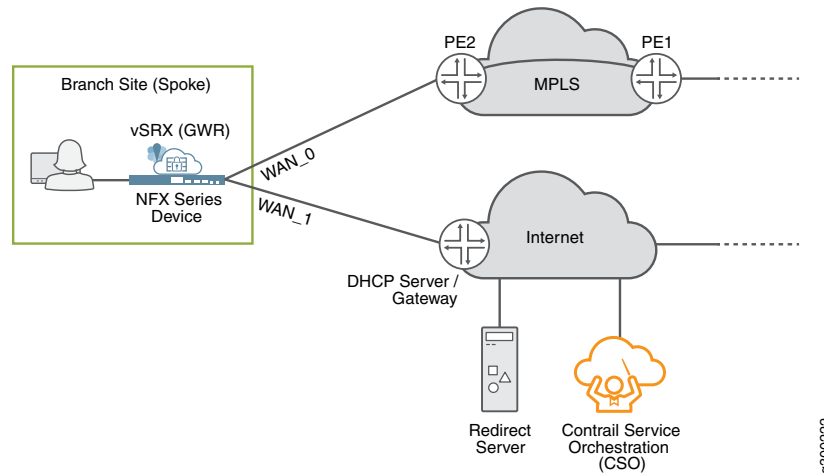
Redirect Server

The redirect server is an Internet-located, Juniper-owned-and-managed server that is integral to the ZTP process. The server enables each spoke device to locate and authenticate with its designated CSO instance. With the assistance of the redirect server, the spoke device can contact CSO and receive its initial configuration, including a Junos OS software update (if required).

For on-premises deployments of CSO, the administrator configures WAN ports on the spoke devices to connect to both the Internet and the redirect server. For cloud-delivered CSO, Juniper Networks handles this configuration for you.

In [Figure 33 on page 53](#), both the redirect server and CSO are located on the Internet. The spoke device obtains and uses IP addressing and other information provided through its Internet-facing interface, and can reach both the redirect server and CSO through that same WAN interface.


Figure 33: CSO and Redirect Server on Internet



Service Chaining in Contrail SD-WAN

Starting in CSO Release 4.0, service chaining is available for SD-WAN environments. Service chaining is a concept wherein multiple network services instantiated in software and running on x86 hardware are linked, or chained together in an end-to-end fashion. This allows the one physical device to provide the services normally provided by multiple devices. Service chaining can be performed on NFX Series devices, as shown in [Figure 34 on page 54](#).

Figure 34: Service Chaining in an SD-WAN Environment



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

NOTE:

- Currently only Layer 2 VNF mode is supported in SD-WAN service chains.

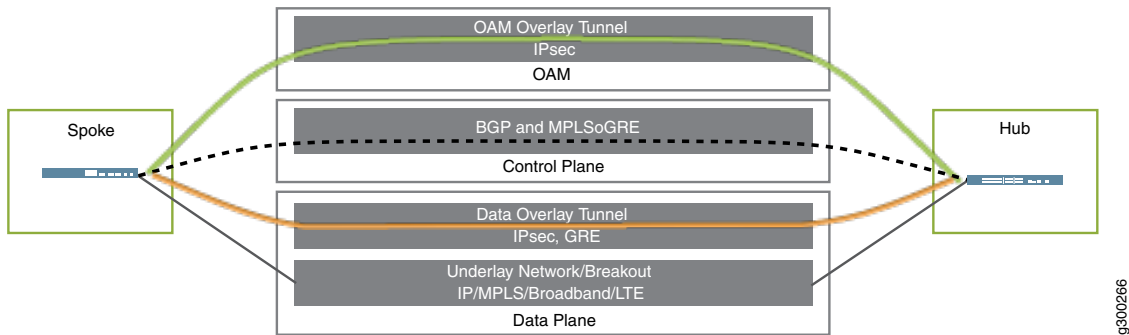
Three Planes, Four Layers

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

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

Figure 35 on page 55 illustrates this concept.

Figure 35: Three Planes, Four Layers



Release History Table

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

Two Basic SD-WAN Use Cases

Two SD-WAN use cases are described below. These use cases illustrate variations around which devices constitute the hubs: a separate SRX Series device (in addition to the MX Series PE device providing underlay connectivity) dedicated to providing SD-WAN overlay connectivity for CPE devices; or a dedicated SRX Series device used for terminating overlay connectivity.

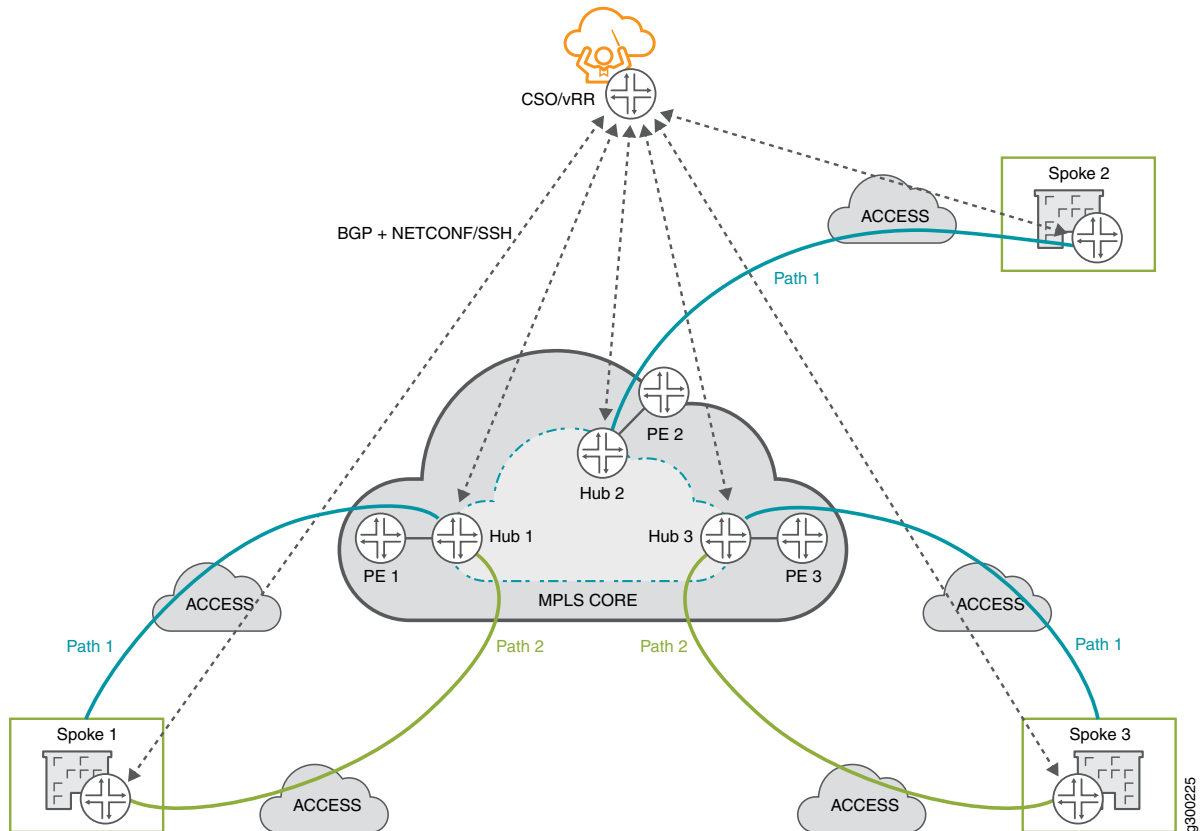
Managed SD-WAN - Overlay Access

This use case is most applicable when the provider wants to take advantage of their existing network, but maintain separation between the existing infrastructure and new SD-WAN infrastructure.

As shown in [Figure 36 on page 56](#), the existing PE devices deployed at POPs remain in place and continue to form that function. In addition, SD-WAN hub devices are deployed at POPs alongside the PE devices to terminate overlay tunnels from the spoke sites.

Again CSO manages the hub and spoke devices. In this use case it also makes use of its vRR to establish BGP sessions with the devices. The vRR advertises reachability information to all devices to provide site-to-site connectivity.

Figure 36: Managed SD-WAN Use Case - Overlay Access



Implementation characteristics:

- The core MPLS infrastructure is managed by the provider.
- The access links can be MPLS or Internet.
- The overlay tunnels extend from the spoke site CPE devices to the dedicated SD-WAN hub devices.
- Multiple overlay encapsulations are supported.
 - MPLSoGRE (CE-PE/MPLS access)
 - MPLSoGREoIPsec (Internet access)
- SRX Series devices are used as provider hubs for IPsec termination.
- The SRX Series devices peer with the PEs for connectivity.
- Provider hubs can be shared across multiple tenants.

Enterprise SD-WAN - Overlay

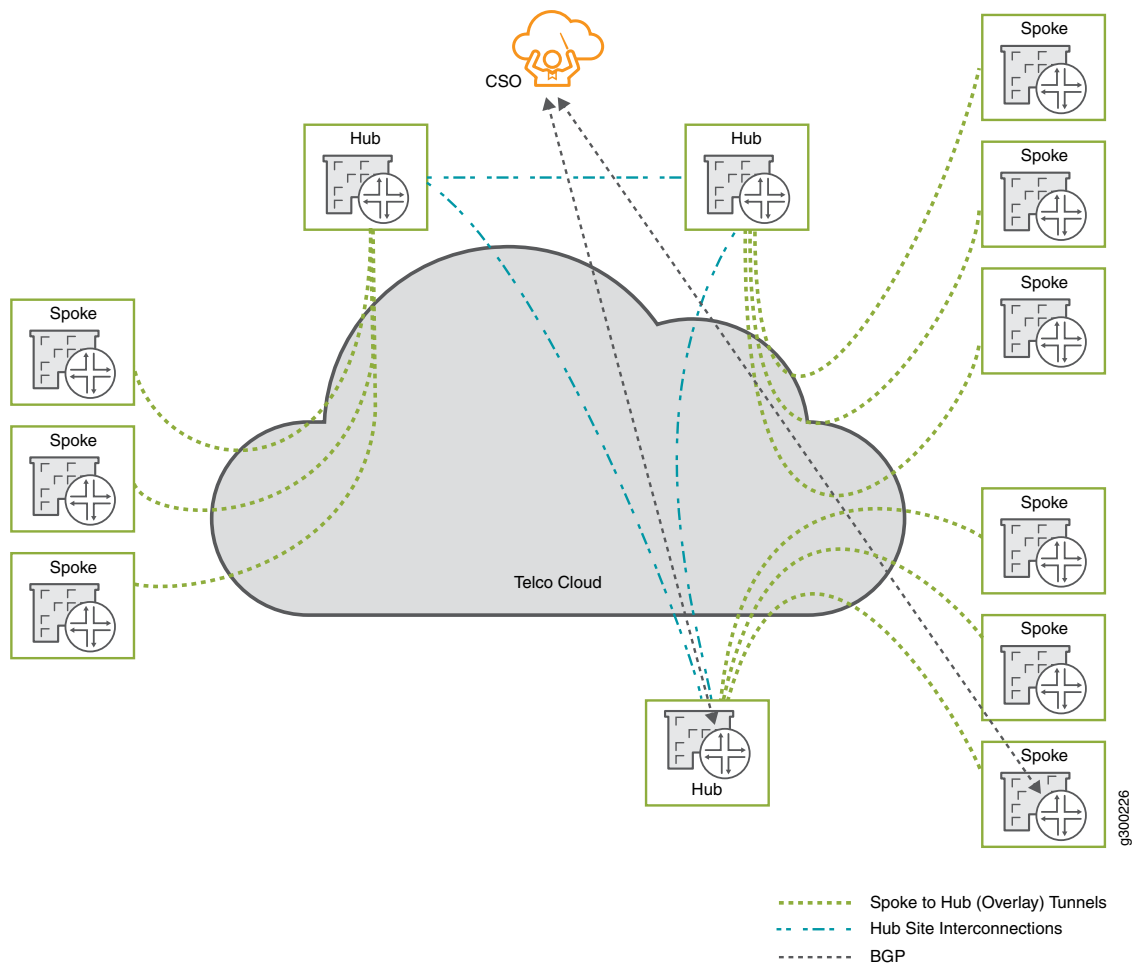
This use case is most applicable to larger enterprises to enable full, end-to-end overlay network connectivity, completely independent of the underlying provider networks.

With this use case, the enterprise customer (tenant or OpCo) owns the hub devices, not the provider. Only spoke sites belonging to this tenant can connect to the enterprise hub devices.

Figure 37 on page 57 illustrates this use case. The enterprise hub devices are located at the customer sites, with overlay tunnels to each of the spoke sites. The hub devices are also interconnected through a provider service such as an MPLS VPN, providing full site-to-site connectivity.

Again, CSO manages all hub and spoke devices, and its vRR advertises reachability information to all devices.

Figure 37: Enterprise SD-WAN Use Case - Overlay



Implementation characteristics:

- The overlay tunnels extend from the spoke site CPE devices to the hub devices.
- The overlay tunnels use MPLSoGRE or MPLSoGREoIPsec encapsulation, as appropriate.
- SRX1500, SRX4100, or SRX4200 Series devices can be used as enterprise hubs for IPsec termination.
- Enterprise hub sites are located at customer sites.
- PE resiliency can be implemented by connecting CPE WAN links to primary and secondary PE nodes.

CSO establishes BGP peering relationships between the CPE and PE nodes. See [Adding an On-Premises Spoke Site with SD-WAN Capability](#) for details.


NOTE: Only supported when local breakout is configured on the CPE WAN link.

- BGP underlay route advertising can be configured to the primary and secondary PE nodes from CPE devices when local breakout is enabled on the WAN interface. See [Adding an On-Premises Spoke Site with SD-WAN Capability](#) for details.

Secure and Redundant OAM Network

Contrail SD-WAN deployments include a secure OAM overlay network to provide end-to-end secure communications between on-premises devices and CSO. As shown in [Figure 29 on page 49](#), dedicated, IPsec-encrypted OAM tunnels enable on-premises devices to send management, routing, and logging traffic securely over the network to a provider hub. The hub then forwards that traffic to CSO.

Figure 38: Secure OAM Tunnels

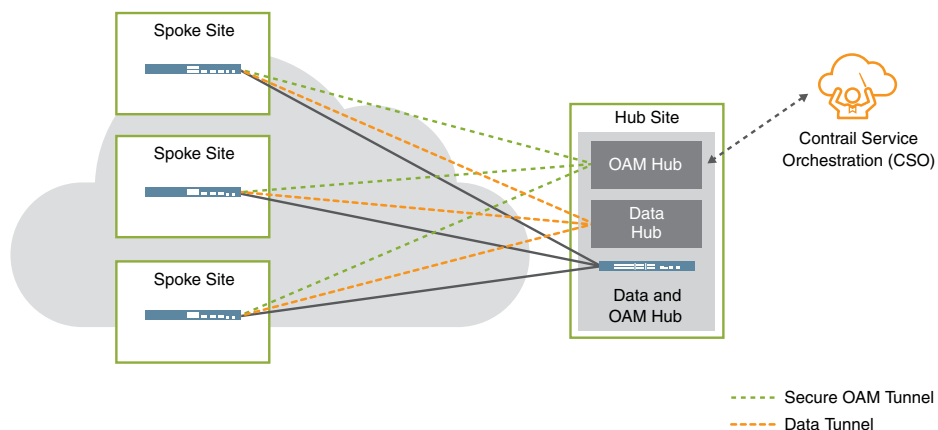


The sites in both the hub-and-spoke and dynamic mesh deployment topologies must use at least one secure OAM tunnel. You accomplish this by setting one of the WAN links for use with OAM during the site onboarding process.

BEST PRACTICE: We recommend having at least two of your WAN links set for use as OAM as shown in [Figure 29 on page 49](#).

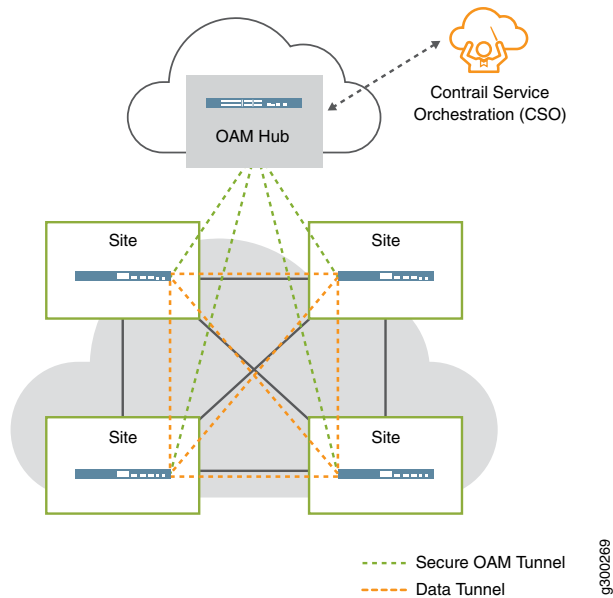
With the hub-and-spoke topology, each spoke site now has two sets of connections to the provider hub site: an overlay tunnel carrying data, and a separate, dedicated IPsec overlay tunnel carrying OAM traffic, as shown in [Figure 30 on page 49](#).

Figure 39: OAM Tunnels in the Hub-and-Spoke Topology



Since a normal dynamic mesh topology would not include a hub device for data traffic, one must be added for the secure OAM traffic. As shown in [Figure 31 on page 50](#), each spoke site has a new connection: a separate, dedicated IPsec overlay tunnel carrying OAM traffic to the provider hub.

Figure 40: OAM Tunnels in the Full Mesh Topology




OAM Provider Hub Design Options

There are two ways to implement the OAM hub, depending on design requirements. As shown in [Figure 32 on page 51](#), the options are as follows:

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

Figure 41: OAM Tunnels - Provider Hub Design Options



Usage Notes on Provider Hub Design Options:

- An OAM provider hub can support multiple tenants, or can be dedicated to a single tenant.
- Connectivity from the provider hub(s) to CSO should be private and secured, as it is not covered by the OAM tunnels.
- We recommended that you implement multiple OAM provider hubs for redundancy and to ensure no loss of management or monitoring of the on-premises devices.
- When a spoke site is multi-homed to multiple hub devices, one OAM tunnel should terminate on each hub. There is no configuration needed in CSO other than configuring multi-homing and specifying the two hubs. CSO automatically terminates one OAM tunnel on each hub device.
- On-premises devices behind NAT are supported for hub-and-spoke and dynamic mesh deployments.

SD-LAN Solution

IN THIS CHAPTER

- SD-LAN Deployment Architectures | 63
- SD-LAN with EX Series Switches | 66
- Managed WLAN Using Mist Wireless Access Points | 69

SD-LAN Deployment Architectures

This topic describes the SD-LAN deployment architectures. There are three basic SD-LAN architectures:

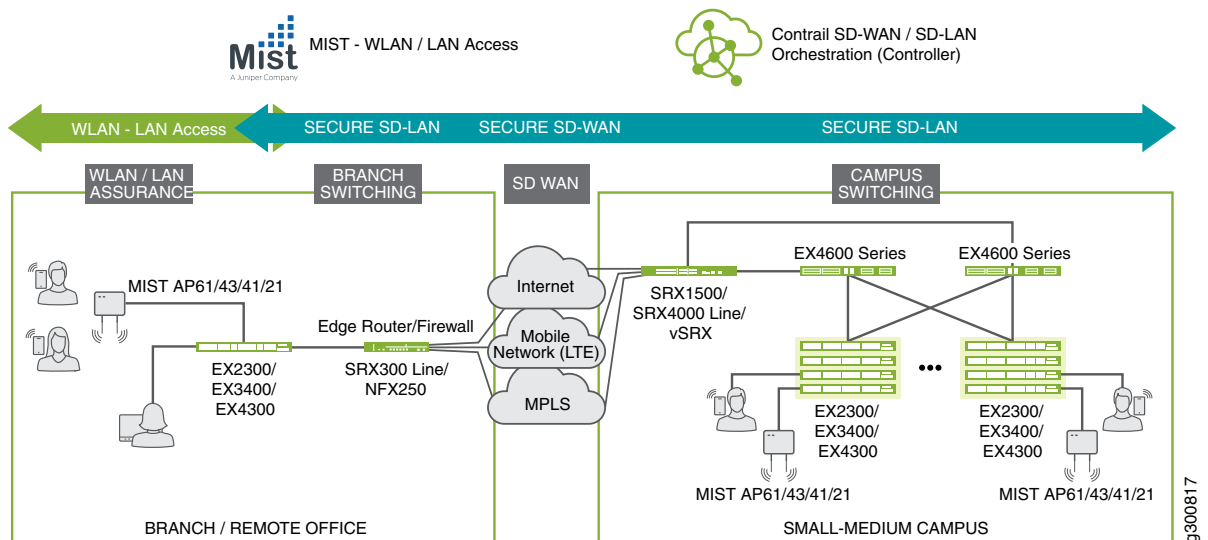
1. **LAN segment connected behind an SD-WAN hub**—For this model, the LAN is in the data center and you need to provision CSO to use dynamic routing protocols for connectivity between the hub and LAN.
2. **LAN segment connected by switches or Virtual Chassis behind an SD-WAN spoke**—For this model, EX Series switches or Virtual Chassis connect to an on-premises spoke device (NFX Series or SRX Series) at a remote site. You can also place Mist access points behind a switch.

NOTE: Juniper Networks recommends this model because it provides flexibility and the opportunity to expand your LAN with additional switches and WiFi access points over time.

3. **LAN segment connected directly to an SD-WAN spoke**—For this model, the LAN connects directly to the on-premises spoke device.

Figure 4 on page 12 shows a sample design of how you can deploy SD-LAN in a remote branch location and a data center, and then connect the sites together across an SD-WAN.

Figure 42: SD-LAN Topology



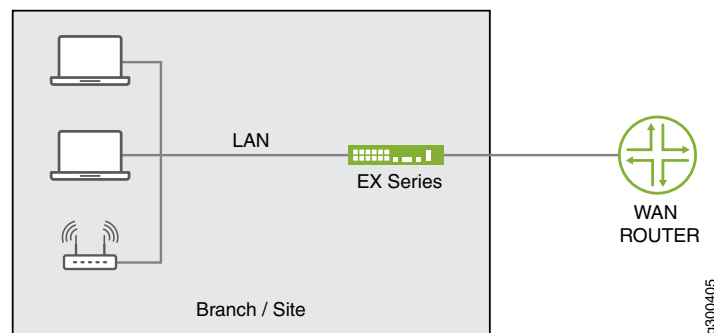
SD-LAN allows you to deploy, provision, manage, and monitor EX Series access switches, EX Series Virtual Chassis, and SRX Series Security Gateways deployed at your spoke sites. CSO can also recognize Mist access points behind an EX Series switch or Virtual Chassis. The LAN devices can be deployed either as standalone devices or behind existing SD-WAN CPE devices. In addition to device visibility, the SD-LAN

solution allows CSO to learn the details of the on-premises spoke site's LANs through the use of dynamic routing protocols such as BGP and OSPF.

SD-LAN Architectures

When you implement SD-LAN with an EX switch or Virtual Chassis as shown in [Figure 43 on page 64](#), the LAN device must be connected to a WAN router device. The WAN router device can be an existing CPE on-premises spoke device as part of an SD-WAN solution, a standalone NGFW device as described later, or a third-party router. The WAN router serves as the gateway to other spoke sites through enterprise hub devices and, ultimately, to CSO.

Figure 43: SD-LAN Using EX Switch



The EX switches can be provisioned using zero-touch provisioning (ZTP). a Virtual Chassis can be provisioned in CSO using ZTP, but the Virtual Chassis itself must be fully formed prior to provisioning. CSO can manage the following aspects of the deployed EX switch or Virtual Chassis:

- Configure system-wide settings around DHCP, SNMP, Radius, and so on
- Configure device specific LAN settings such as VLANs, interfaces, 802.1x, POE, port security, and so on through the use of stage 2 configuration.

The EX switch or Virtual Chassis gives visibility into the local LAN and allows for the inclusion of Mist access points into the CSO orchestration model. When deployed in any of these ways, any existing Mist access points that are connected to the EX switch, at that site, are automatically seen by CSO.

Usage Notes on SD-LAN Architecture

EX Series switches and Virtual Chassis are supported in CSO. Once a switch or Virtual Chassis has been provisioned, you can:

- **Configure and monitor the ports of the EX switch or Virtual Chassis**—You can either configure the ports by accessing each port individually or by using a port profile, from the Ports tab of the Devices page in the Customer Portal UI. You can configure and deploy port authentication profiles to implement network access control (NAC), and firewall filters to enforce security on the switch ports. After you configure the switch ports, you can monitor the ports from the Devices page of the Customer Portal UI.
- **Perform an image upgrade for Virtual Chassis members**—CSO supports the upgrade of images for an EX Series access switch or Virtual Chassis. Images for each member of the Virtual Chassis are upgraded one after the other in the order – Linecard, Backup, and Master.
- **Add firewall configurations for EX Series switches**—You can configure firewall filters for EX Series switches and Virtual Chassis. A firewall filter defines the rules to permit or deny packets that are transiting a switch port. You can assign the firewall filter as an ingress filter or egress filter to a switch port either while manually configuring the port or through port profiles.
- **Set up RMA for defective EX Series switches**—You can initiate the Return Material Authorization (RMA) workflow for a defective EX Series switch (physical standalone switch) when the switch is behind an SRX Series device acting as an SD-WAN CPE, next-generation firewall, or internet gateway.

You can set up RMA for an EX Virtual Chassis member when the Virtual Chassis is deployed as a standalone switch (that is, behind an internet gateway) only.

SD-LAN with EX Series Switches

For those tenants with LAN service capabilities, you can add an SD-LAN site using EX Series access switches and Virtual Chassis. CSO supports EX2300, EX3400, EX4300, EX4600, and EX4650 switches. The number of members supported in a Virtual Chassis are shown in [Table 10 on page 66](#).

Table 10: EX Virtual Chassis Member Count by Switch

Switch Model	Maximum Virtual Chassis Member Count
EX2300	4
EX3400	10
EX4300	10
EX4600	10
EX4650	2

The VC must be fully formed prior to provisioning in CSO. CSO only requires the master member's serial number. Other member devices are automatically discovered during provisioning.

The use cases below show variations in how the EX switches are added to remote sites within CSO and apply equally to individual switches or VC.

Software defined LAN (SD-LAN) using EX Series access switches can be implemented in the following two ways:

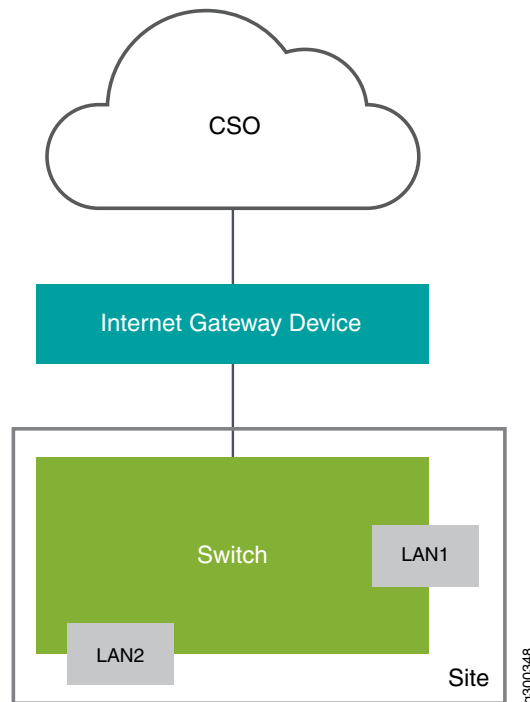
- As an extension of an existing spoke/branch site
- As a standalone spoke/branch site

Standalone SD-LAN Site

For this solution, the EX switch at the remote site must be installed behind an Internet Gateway device. The Internet Gateway device provides the WAN routing necessary for the switch to reach CSO.

[Figure 44 on page 67](#) shows a simple example of this.

Figure 44: Standalone SD-LAN Site

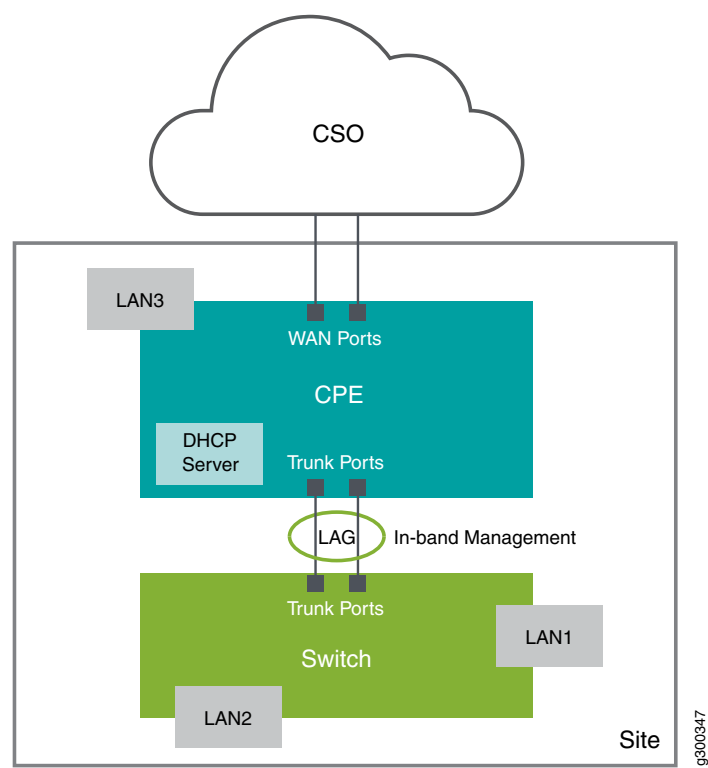


To implement a site like this, you must create a new site and only select the LAN service capability during the site onboarding process.

Extend SD-WAN Site

For existing SD-WAN spoke sites, CSO allows you to add an EX Series LAN switch to a branch/spoke site, behind a CPE device, as shown in [Figure 45 on page 68](#). In this case the CPE device provides the WAN routing needed to reach CSO. If the CPE device is an SRX Series device, the routing is handled natively on the SRX security gateway itself. If the CPE device is an NFX Series device, then the built-in vSRX VNF handles the WAN routing duties.

Figure 45: SD-LAN Behind CPE

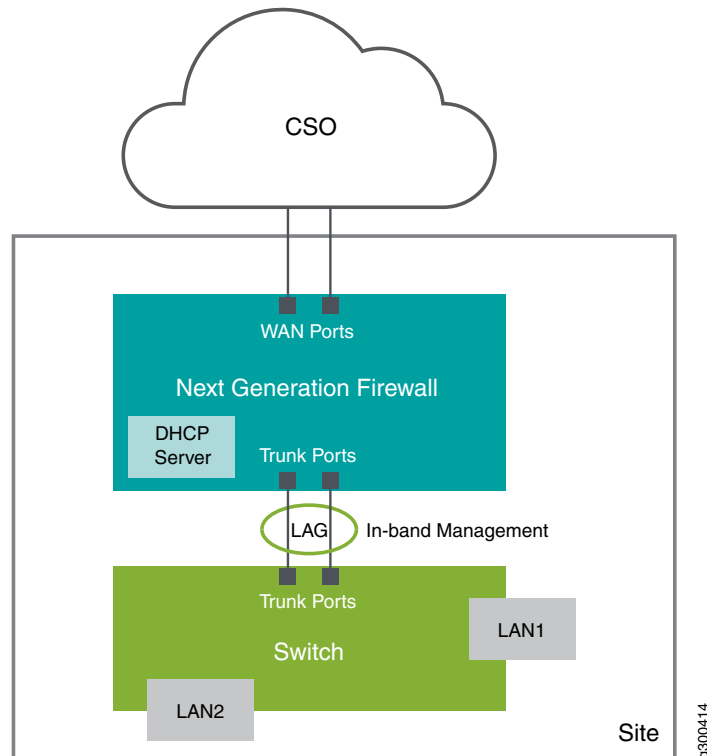


To implement a site like this, you can add the EX switch behind an existing branch/spoke site, or you can create a new site and select both SD-WAN and LAN service capabilities during the site onboarding process.

SD-LAN Behind NGFW Site

Another option for deploying an SD-LAN site is to deploy it behind a NGFW as shown in [Figure 46 on page 69](#). In this case, the NGFW provides the WAN routing needed to get to CSO.

Figure 46: SD-LAN Behind NGFW



Regardless of how the switch is deployed, CSO provides the same ability to configure, monitor, and maintain the device at the remote site.

Managed WLAN Using Mist Wireless Access Points

In addition to EX Series switches and Virtual Chassis, CSO supports the provisioning and management of Mist wireless access points. The use case is described below.

Once visibility into the remote LAN is available with any of the SD-LAN use cases, CSO allows you to use, manage, and maintain Mist wireless access points at the remote branch/spoke sites. Once connected to the EX switch in an SD-LAN solution, CSO will be able to manage multiple Mist access points. If the access point is connected to the EX switch prior to the SD-LAN deployment, then CSO will automatically provision and activate the access point during the site onboarding process. If the access point is added to the switch after SD-LAN deployment, then CSO allows you to add the access point by adding it to the existing site.

NGFW Solution

IN THIS CHAPTER

- [NGFW Deployment Architecture | 70](#)

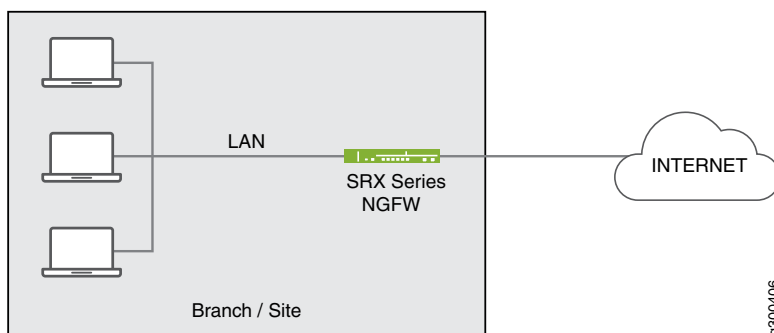
NGFW Deployment Architecture

This topic describes the next-generation firewall (NGFW) deployment architecture for standalone NGFW using SRX Series Security Gateways.

NGFW Architecture

The NGFW architecture offers strong security services for remote sites, along with WAN connectivity. It can also offer LAN services and LAN visibility similar to the SD-LAN architecture. When you use an SRX Series device at an on-premises spoke site as a standalone NGFW, the WAN routing functions are performed on the SRX Series device itself. This architecture allows the SRX Series device to perform all of its built-in security functions (such as firewall and NAT) while providing visibility into the LANs that exist at your spoke sites. [Figure 47 on page 70](#) shows an SRX Series device connected to both the WAN and an onsite LAN.

Figure 47: NGFW and SD-LAN



As mentioned previously, an NGFW site can exist on its own or be extended later with the addition of EX Series LAN switches or Virtual Chassis at any time after provisioning and deployment.

NGFW Deployment Usage Notes

With an NGFW deployment you can:

- **Enable WAN connectivity for sites**—When you provision NGFW service capabilities for a tenant, any site belonging to that tenant can use the NGFW device as its WAN link back to CSO.
- **Enable automatic LAN connectivity**—The NGFW device can provide addressing for a connected LAN by using a built-in DHCP server.
- **Create custom application signatures in firewall policies**—CSO supports custom application signatures in firewall policies, in addition to the existing support in SD-WAN policies.
- **Create customized IPS signatures, static groups, and dynamic groups**—You can create, modify, or delete customized intrusion prevention system (IPS) signatures, IPS signature static groups, and IPS signature dynamic groups. In addition, you can clone predefined or customized IPS signatures, static groups, and dynamic groups. You can then use the IPS signatures, static groups, and dynamic groups in an IPS profile that can contain one or more IPS or exempt rules.
- **Import policy configurations**—CSO supports the import of policy configurations from next-generation firewall devices. The following features are supported:
 - Manage next-generation firewall sites for enterprise customers with brown field deployments.
 - Discover existing policy configuration while onboarding NGFW devices (without enabling ZTP).
 - Import policy configurations from Firewall and NAT policy pages.
 - Deploy policies after import into CSO.

You enable an NGFW deployment in CSO by using the Customer Portal to add an NGFW site. A tenant assigned to the NGFW site must have the NGFW service available. To add the NGFW service, a tenant administrator includes the NGFW service in the tenant configuration during the onboarding process.

4

PART

Operations and Management

Network Operation | **73**

Orchestration and Management Within CSO | **94**

Operational Workflows - Overview | **110**

Resiliency and High Availability | **111**

Network Operation

IN THIS CHAPTER

- [Network Operation | 73](#)

Network Operation

When deploying CSO as an on-premises deployment, it is helpful to know how the network operates and what protocols are in use. When working with a cloud-hosted deployment, the concepts are all the same, but the details and control are invisible to subscribers; they are the responsibility of the team that installs CSO in the cloud.

As with most networks, the Contrail SD-WAN solution generally operates in two planes:

- Control plane – OAM and routing traffic
- Data (forwarding) plane – user traffic

Control Plane Operation

The control plane for the Contrail SD-WAN solution centers around the CSO platform. More specifically:

- CSO's Network Service Controller (NSC) layer implements the control plane using vRRs.
- All sites across all tenants establish MP-IBGP peerings with the vRR.
- CSO uses a single private AS number for all tenants, with route targets for tenant separation.
- Tenant route separation is provided both by the vRR and by multi-tenant hub devices using BGP extended communities.

vRR Design

All CSO deployments include one or more vRR instances, which provide control plane functionality for the SD-WAN environment. [Figure 48 on page 74](#) shows a general example where the on-premises devices at each site peer with the vRR.

Figure 48: Control Plane - Single vRR Design

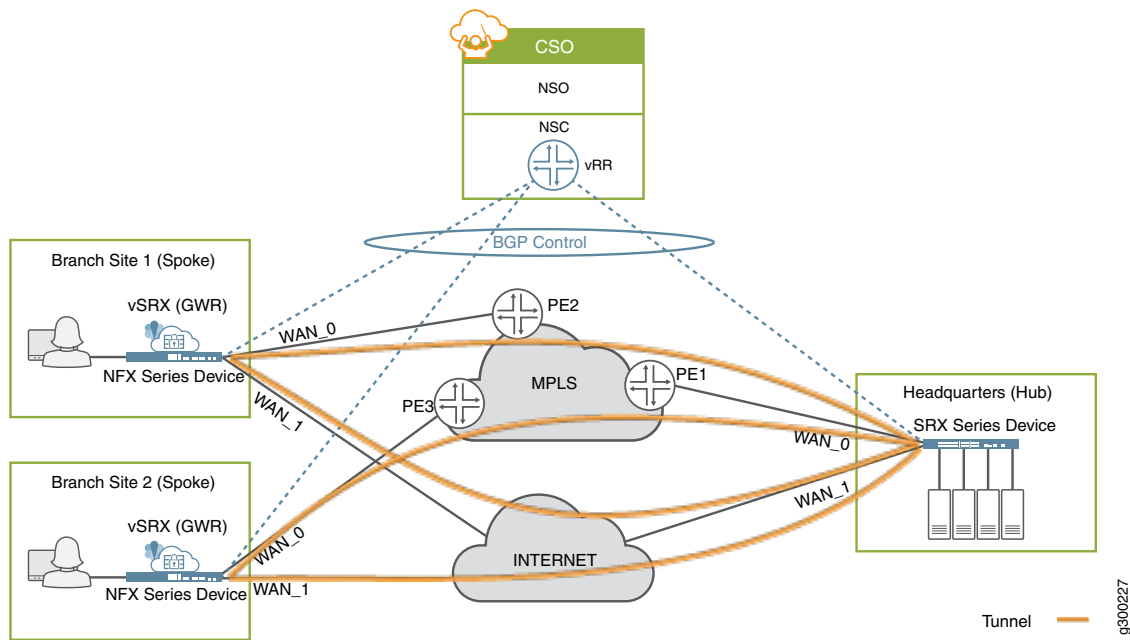


Figure 49 on page 74 shows an example of the CLI output from the vRR.

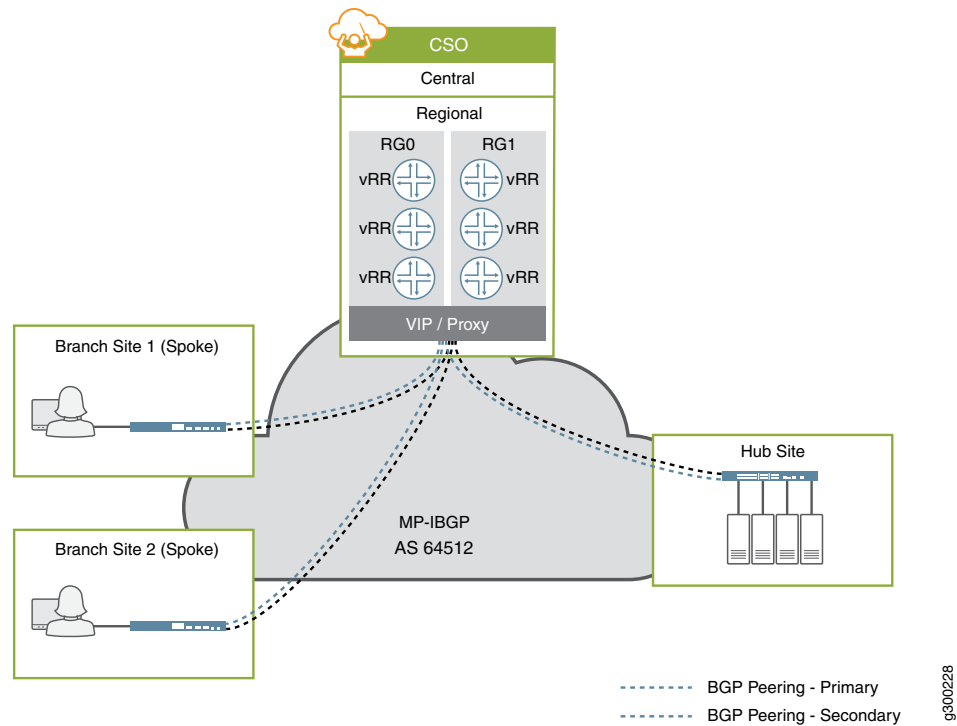
Figure 49: Sample CLI Output from vRR

```
root@vrr-10.209.17.142> show bgp summary
Groups: 3 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
bgp.l3vpn.0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn State!#Active/Received/Accepted/Damped...
192.168.10.2 64512 16 28 0 86 2:02:39 Establ
bgp.l3vpn.0: 13/13/13/0
192.168.30.2 64512 14 57 0 89 2:19:53 Establ
bgp.l3vpn.0: 24/24/24/0
192.168.110.2 64512 16 28 0 38 2:02:33 Establ
bgp.l3vpn.0: 26/26/26/0
root@vrr-10.209.17.142>
```

Control Plane Resiliency

CSO Release 3.3 and later supports the installation of multiple vRRs to provide redundancy and scale. CSO separates the vRRs into two redundancy groups (RGs), and makes a single virtual IP address visible to the network. As part of a site’s configuration, CSO establishes BGP peering sessions between the device and a vRR in each RG. If the primary vRR fails or connectivity is lost, the second vRR continues to receive and advertise LAN routes for the connected sites, thereby providing redundancy. This design is illustrated in Figure 50 on page 75.

Figure 50: Control Plane - Multi-vRR Design




Route Distribution and Separation

The Contrail SD-WAN solution uses Junos OS virtual routing and forwarding (VRF) instances and MP-BGP route targets to provide tenant route separation and enable multi-tenancy.

These concepts can be well illustrated using an MPLS VPN environment as an example. As shown in [Figure 51 on page 76](#), each customer is assigned a unique route target value, and all sites of the customer VPN use that route target value. When a router advertises a customer's routing information it attaches the appropriate route target value based on which customer VRF originated the advertisements. The receiving router uses the attached route target value to identify the customer VRF into which the received routing information should be placed.

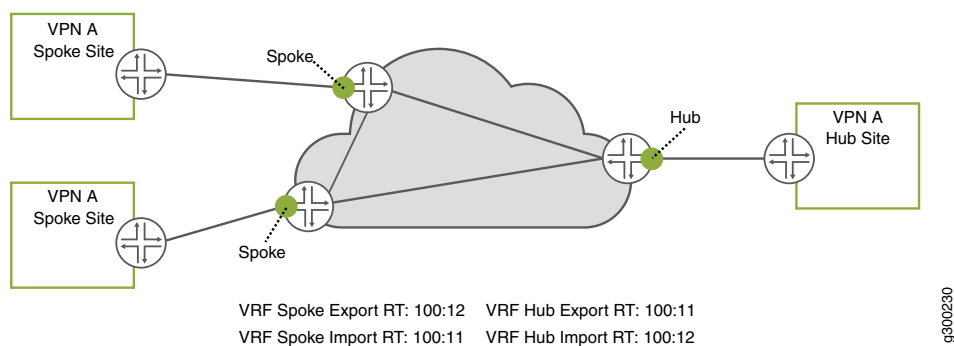
Figure 51: Route Separation Example - MPLS VPNs



An MPLS VPN hub-and-spoke environment uses route targets differently, as shown in [Figure 52 on page 76](#). For each customer, every spoke VRF attaches the same route target value when sending routing information. The receiving router accepts routes with that same route target value and installs them into hub VRF. By contrast, the hub VRF attaches a different route target value when sending routing information, and the receiving routers accept and install routes with that same route target value into spoke VRFs.

With this setup, only the hub VRF accepts routes from the spoke VRFs, and only the spoke VRFs accept routes from the hub VRF. Using this method, the spoke sites need very little routing information (perhaps just a default route) as they only need reachability to the hub site, thereby keeping routing tables small and churn-free.

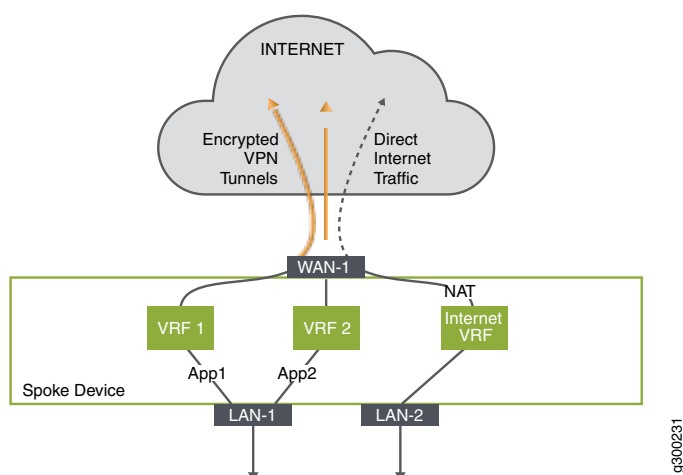
Figure 52: Route Separation Example - Hub-and-Spoke MPLS VPN



The hub and spoke example above serves as a good foundation, as the Contrail SD-WAN solution implements route distribution and separation in the same way when forwarding traffic from one site to another, or when breaking out traffic to the local internet.

[Figure 53 on page 77](#) shows a spoke site example where the spoke device is configured with two overlay tunnels and local breakout, with all traffic flowing out the same interface. Each traffic path has its own VRF, and route targets are assigned appropriately at the spoke and hub sites to ensure proper tenant route separation.

Figure 53: Route Separation - SD-WAN Spoke Site



APBR and SLA Management - Control Plane

Advanced policy-based routing (APBR) enables you to define routing behavior and path selection per application (group). The APBR mechanism classifies sessions based on well-known applications and user-defined application signatures and uses policy intents to identify the best possible route for the application. Dynamic application-based routing makes it possible to define policies that will switch WAN links on the fly based on the application's defined SLA parameters.


Real-Time Optimized - AppQoE

Starting with Release 3.3.1, CSO supports Application Quality of Experience (AppQoE), a data plane-level mechanism that provides better scalability and faster decision making. Working in conjunction with APBR, AppQoE functions at the device level; that is, the devices themselves perform SLA measurements across the available WAN links, and then dynamically map the application traffic to the path that best serves the application's SLA requirement. This is all done without the need for the CSO controller to distribute SLA-specific routes.

With AppQoE, when an SLA violation occurs, only traffic corresponding to the application that reported the SLA violation is moved to an alternate link; any other traffic using the link is unaffected.

With real-time optimized SLA management only the default VRF is required, as shown in [Figure 54 on page 78](#). The default VRF uses ECMP across all the links. The next hop selection per SLA happens in the data path (described in the data-plane section).

Figure 54: Real-Time Optimized (AppQoE) Routing Architecture



In this case, the MPLS label is used only to identify the tenant.

NOTE: AppQoE is enabled when the SD-WAN mode for the tenant is set to *Real-time Optimized*. This is the default mode for SD-WAN deployments.

Note the following about AppQoE:

- Only supported on SRX and vSRX devices.
- Both ends must use the same Junos OS version and the same configuration.
- Multi-homing is supported.


Data Plane Operation

This section discusses how a packet is forwarded in a hub-and-spoke topology.

When a user at a spoke site sends traffic through the on-premises CPE device, and the packet is not locally switched or sent direct to the Internet, it is sent over a tunnel to the hub device. This packet from the customer LAN is first encapsulated inside an MPLSoGRE header with the GRE destination as one of the

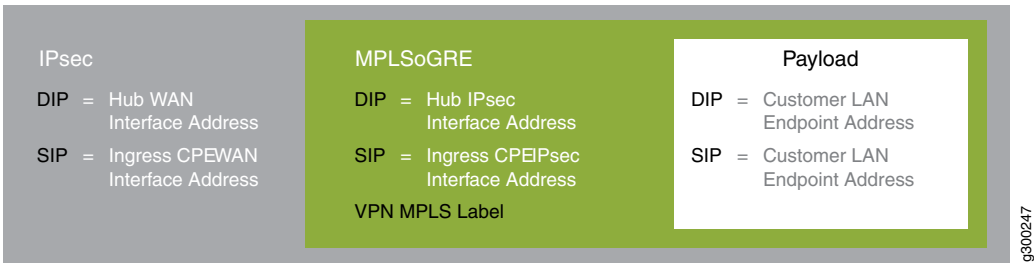
WAN links of the hub device. The MPLS label in the MPLSoGRE header identifies the VRF to be used for forwarding the packet at the hub site. The resulting packet header is shown in [Figure 55 on page 79](#).

Figure 55: Packet Header - MPLSoGRE



If the tunnel between the spoke and hub site is configured to use IPsec, the MPLSoGRE packet is then further encrypted and encapsulated in an IPsec header that uses tunnel mode. The resulting packet header is shown in [Figure 56 on page 79](#).

Figure 56: Packet Header - MPLSoGREoIPsec




At the hub, the IPsec header is first decrypted. The resulting packet's MPLSoGRE header is used to terminate the GRE tunnel and perform a lookup in the appropriate VRF, as identified using the MPLS label. Based on the route lookup in the VRF, the packet is then either forwarded towards another spoke site, or out of the SD-WAN environment. If forwarded to another spoke, the hub device encapsulates the packet as described as above.

Design Options

[Figure 57 on page 80](#) illustrates how the tunnels are typically deployed using the packet headers described above. GREoIPSec tunnels are generally used over the Internet path, given the need for secure packet transport over the public network. GRE tunnels are generally used over MPLS paths, though the GREoIPSec option may also be used as appropriate.

Figure 57: Tunnel Design Options



APBR and SLA Management - Data Plane

As noted earlier, tenants can choose one SD-WAN mode of SLA management for application traffic:

- Real-time optimized – Device-level SLA management, using AppQoE

AppQoE is a data plane-level mechanism that provides better scalability and faster decision making. With AppQoE, link switching occurs at the application level in the data path of the devices; the devices themselves perform SLA measurements across the available WAN links, without the need of the CSO controller.

Link monitoring occurs using two types of inline probes:

- Passive Probes
 - Inline probes that ride along with application traffic
 - Mimic the burstiness of the application flows
 - Enable monitoring of RTT, jitter, packet loss for the application session
 - Used to monitor currently used path for SLA compliance, detect SLA violation
- Active Probes

- Periodic probes (based on configuration) that gather SLA data on all potential paths
- Used to determine the original best path for the traffic
- Used to monitor alternate paths

NOTE: AppQoS is enabled when the SD-WAN mode for the tenant is set to *Real-time Optimized*.

Tunnel Liveliness

To avoid blackholing traffic, appropriate liveness checks are enforced in the overlay network. The Contrail SD-WAN solution uses two mechanisms to ensure liveness:

- IPsec dead peer detection (DPD), where it is used
- GRE keepalives

Mesh Tags and Dynamic Mesh VPNs

As mentioned in the deployment models discussion, dynamic mesh is Juniper's resource-saving implementation of full-mesh VPNs within CSO. This section describes the operation of mesh tags and dynamic mesh VPNs that they enable.

Mesh Tags

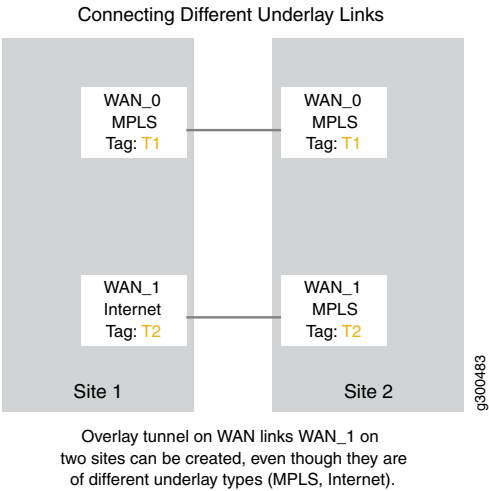
Mesh tags are text-based labels applied to the WAN interfaces of CPE and hub devices during the onboarding process in CSO. CSO is shipped with two default mesh tags: Internet and MPLS. You can create your own mesh tags using the CSO Administration Portal. On-demand, or dynamic, VPNs can only be formed between WAN interfaces that share the same mesh tag.

The following discussion explains how mesh tags work and some of the use cases to which they apply.


As mentioned above, one mesh tag is applied to each WAN interface of the CPE device at each site. On spoke devices such as the NFX150 and NFX250, and most SRX devices, only one mesh tag can be applied to each WAN interface. On provider hub and enterprise hub devices such as the SRX4x00 Series devices, multiple mesh tags can be applied to each interface due to the increased VPN capabilities of the devices.

The following list helps to illustrate the various use cases in which mesh tags and dynamic mesh VPNs come into play.

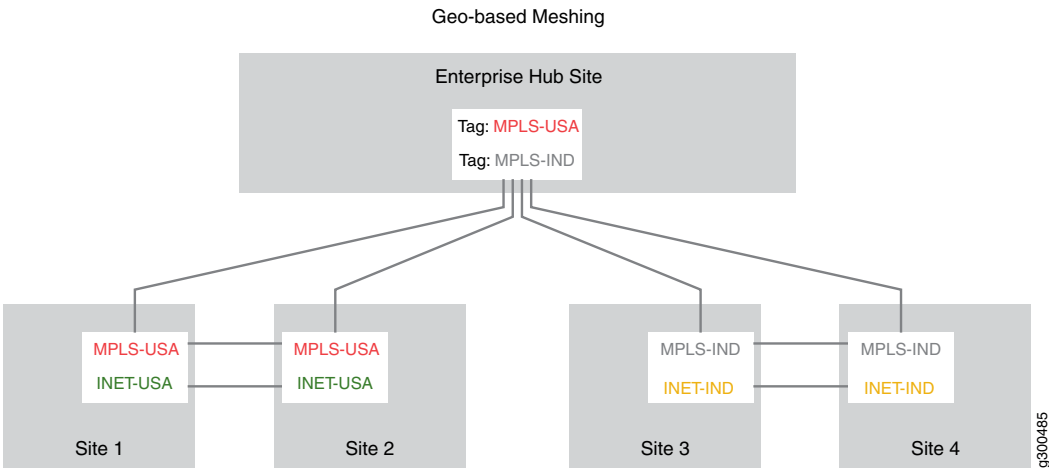
• Connecting Different Underlay Links



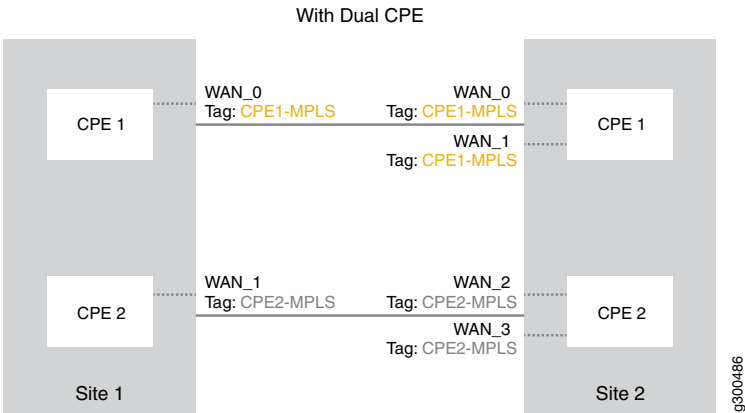
• Site-to-Site Tunnels Based on Capacity



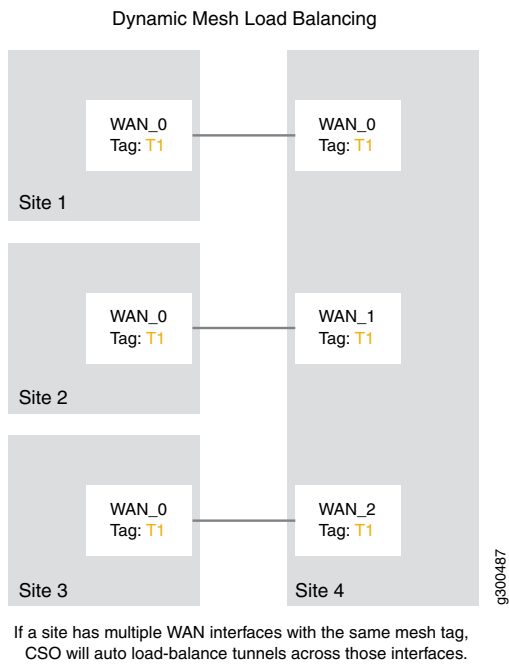
• Geo-Based Meshing



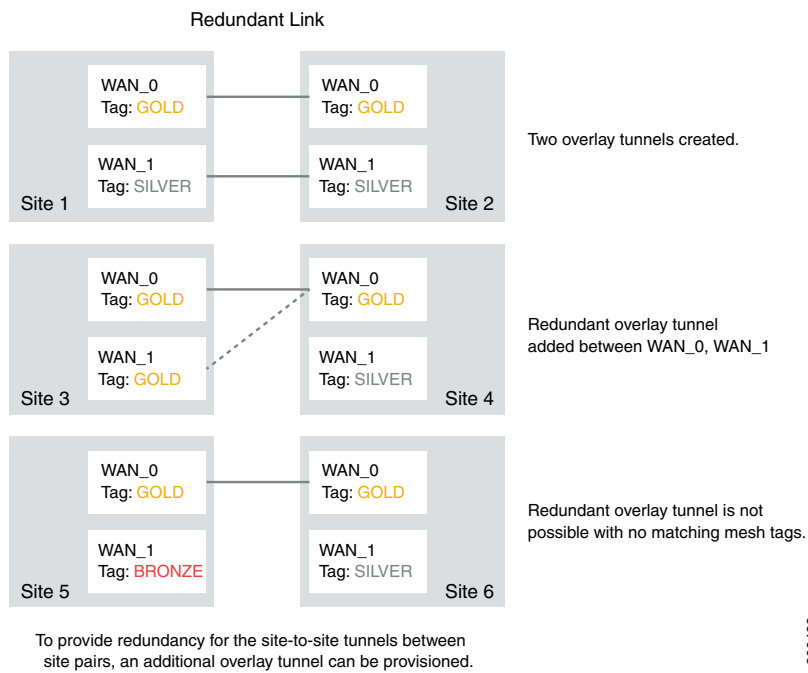
• With Dual CPE



• Dynamic Mesh Load Balancing




• Redundant Link



Dynamic Mesh VPNs

Figure 58 on page 85 shows a dynamic mesh VPN topology between three spoke sites and describes how the site-to-site VPN is brought up.

Figure 58: Dynamic Mesh Operation



1–Sites and tunnels to Hub sites provisioned using ZTP. Site to site traffic goes through the site to hub data tunnels.	4–CSO configures on-demand site-to-site tunnels between the site-pairs.
2–CSO receives syslog messages from the devices containing details about traffic rates.	5–Site-to-site traffic now switches to the newly formed site-to-site tunnels.
3–CSO recognizes that the traffic between Phoenix Site 1 and Houston Site 2 exceeds KPI thresholds.	

NOTE: Tunnel deletion is also controlled and automated by CSO using traffic thresholds and syslog messaging.

Internet Breakout

While traffic destined for the Internet can be sent across the overlay tunnels and through a central site, the tunnels are more typically intended to support site-to-site traffic. For non-SD-WAN destinations, local breakout provides the option to send the traffic out of the local on-premises device directly to the Internet. Local breakout allows the tenant to use its network bandwidth optimally at each site and to avoid incurring the cost of carrying all traffic to the central site.

Local breakout is an important feature in SD-WAN deployments, as many enterprises nowadays use SaaS services that are hosted outside the corporate network. Since most of these SaaS apps use SSL as the transport and also support single sign-on with the enterprise AAA systems, security concerns are addressed despite sending traffic directly over the Internet.

WAN Interface Options

An on-premises device's WAN (MPLS and Internet) interfaces can support tunneled and local breakout traffic in any combination:

- Tunneled traffic only
- Tunneled and local breakout traffic
- Local breakout traffic only


Design Options

There are multiple ways to implement local breakout, depending on design requirements.

Breakout at Spoke

Local breakout at spoke sites allows users to access the Internet directly without having to send traffic over the overlay network towards the hub, thus helping to conserve tunnel bandwidth. This option can be implemented on either the Internet or MPLS WAN links. [Figure 59 on page 86](#) illustrates this concept.

Figure 59: Local Breakout at Spoke

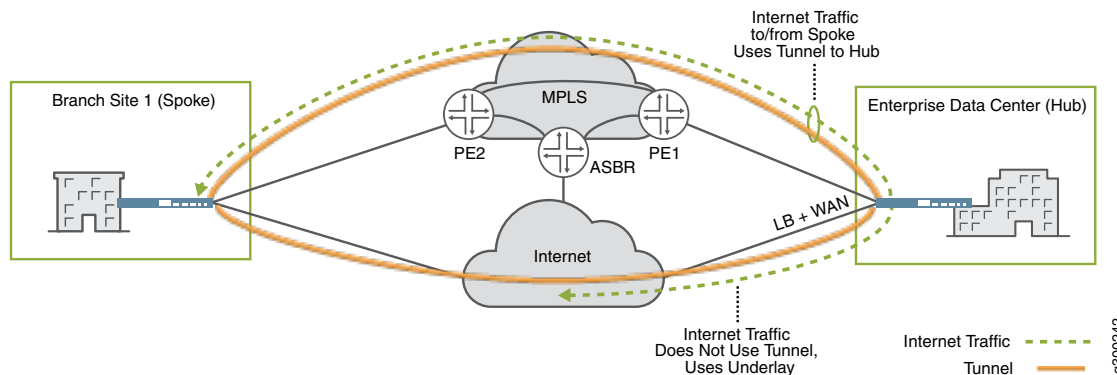


When using local breakout, you can specify either interface-based or pool-based NAT.

Breakout at Provider Hub (Central Breakout)

Central breakout at a provider hub site enables hub-and-spoke deployments where spoke sites forward Internet-destined traffic through the overlay network to the provider hub device, which then forwards the traffic out to the Internet as shown in [Figure 60 on page 87](#).

Figure 60: Local Breakout at Hub



Central breakout at the hub site is enabled differently than at a spoke site. It can be configured manually in CSO through Stage 2 templates.

Central breakout can also be provided to spoke sites through an Enterprise Hub site. In this scenario, the enterprise hub can either perform local breakout using an underlay network for forwarding or it can receive the default route from the Datacenter department and propagate it to the spokes.

When central breakout is offered at both the provider hub and enterprise hub through the default route method, the default route from the enterprise hub is preferred using BGP local preference.

Cloud Breakout

Another breakout option for Internet-destined traffic, Cloud Breakout, is available to spoke and enterprise hub sites. When cloud breakout is enabled, the spoke site or the enterprise hub site forwards Internet-destined traffic to Zscaler for further security-related processing before it is sent to the Internet. The Zscaler account must be active and accessible prior to sending traffic through the breakout.

Usage Notes for Cloud Breakout

- Generic routing encapsulation (GRE) tunnels that use public IP addresses for the WAN links are supported for cloud breakout.
- When using GRE tunnels, the CPE devices cannot be behind NAT.
- When you configure cloud breakout settings, you can specify IPsec phase 1 parameters, phase 2 parameters, and domain name.
- You can specify IP address or hostname validation for cloud breakout nodes.

- CSO auto-populates FQDN, preshared keys, and WAN link information and provides the option to change the auto-populated values.
- CSO supports high-availability between the WAN links of an SD-WAN spoke site and the cloud breakout node.
- WAN link nodes can be configured as active/passive or active/active.
- A maximum of two WAN links can be defined between the SD-WAN spoke site and the cloud breakout node.

Order of Preference for Scenarios with Multiple Breakout Options

If multiple breakout options are available to the CPE at the spoke site and there is no breakout policy specified, then the order of preference for breakout is:

1. Datacenter department/enterprise hub
2. Local breakout/Cloud breakout
3. Provider hub (Central)

If multiple breakout options are available to an enterprise hub site, the order of preference for breakout traffic is:

Without SD-WAN policy:

1. Datacenter department
2. Hub

With SD-WAN policy:

1. Local breakout/Cloud breakout
2. Datacenter department
3. Provider hub (Central)

Use Cases for Local Breakout

Some use cases for local breakout are described below.


Service Provider Data Center

In this use case, the Enterprise customer uses the service provider's SD-WAN service for site-to-site inter-connectivity. The customer also uses value-added services hosted out of the service provider's data center.

At the spoke site, the on-premises device's MPLS-facing WAN interface is configured to support both tunneled and local breakout traffic. As shown in [Figure 61 on page 89](#), traffic flows across the network as follows:


- Inter-site (SD-WAN) traffic travels across the MPLS network using the overlay tunnel.
- DC-destined traffic uses local breakout and travels directly across the underlay MPLS network.

Figure 61: Local Breakout at Spoke to DC Located in Telco Cloud



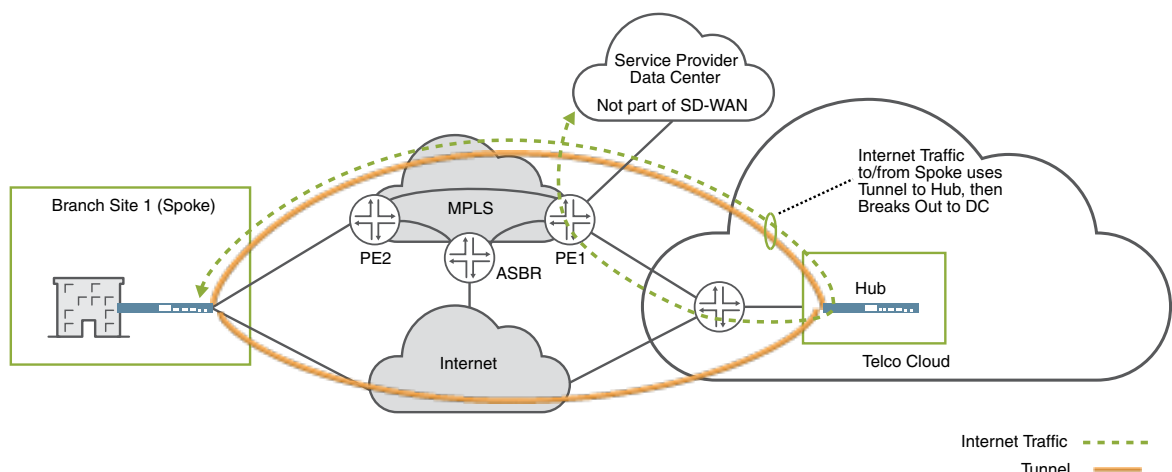
As a variation on this scenario, the data center could be located elsewhere on the MPLS network, perhaps at a POP, as shown in [Figure 63 on page 90](#). In this case, traffic flows remain generally the same as above.

Figure 62: Local Breakout at Spoke to DC Located at POP



As another variation on this scenario, DC-destined traffic could use the overlay tunnel, breakout at the hub device, and double back to the DC, as shown in [Figure 63 on page 90](#).

Figure 63: Local Breakout at Hub to DC Located at POP



This option has some drawbacks:

- It uses more tunnel bandwidth.
- It may increase latency as the on-premises device at the spoke site processes and encapsulates more traffic.
- It increases the load on the hub device.
- It creates a suboptimal path, causing traffic to flow through the tunnels to the hub device, only to have to double back to get to the DC.

However, it also has some advantages:

- Using the overlay tunnels, DC-bound traffic can take advantage of SLA services and choose the best path dynamically, thus improving network performance for those applications.
- Additional security functions can be offered centrally.

Migration to SD-WAN

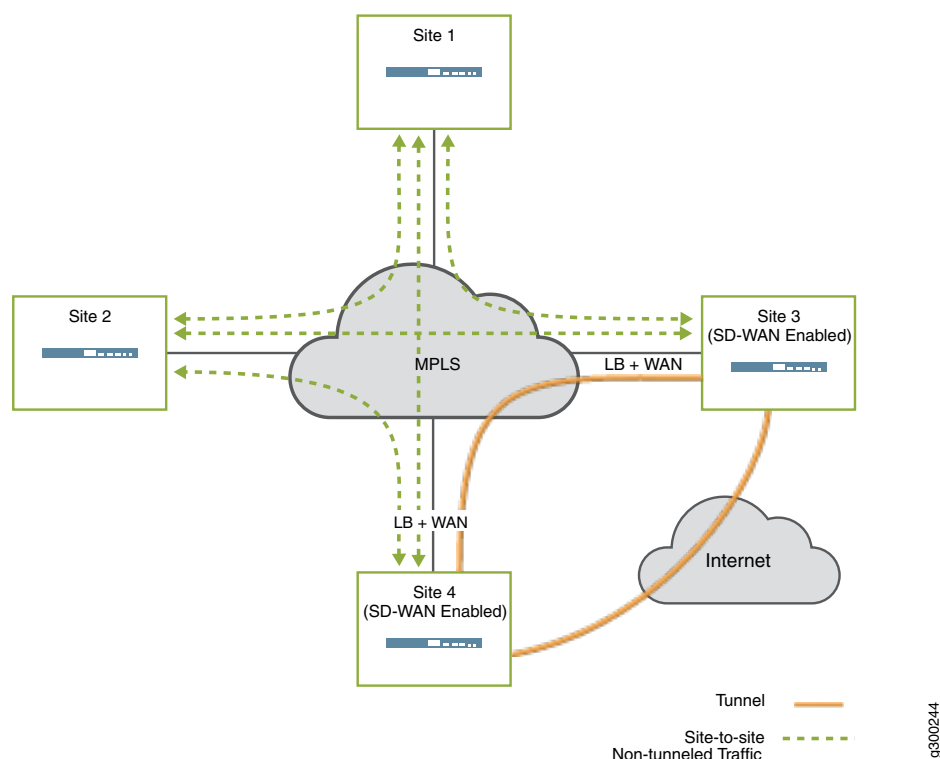
In this use case, the enterprise customer has multiple large locations and uses the service provider's existing MPLS service to provide a full mesh between sites. The customer wants to migrate to SD-WAN, and the implementation is likely to be incremental. Nevertheless, it is critical to maintain connectivity between sites at all times.

Figure 64 on page 91 illustrates a scenario where the migration is underway. SD-WAN functionality has been added to Site 3 and Site 4, while the other sites have not yet been migrated. At each SD-WAN-enabled site, the on-premises device's MPLS-facing WAN interface is configured to support both tunneled and local breakout traffic. Traffic flows across the network as follows:

- Traffic between the SD-WAN-enabled sites can use the overlay tunnel.

- Traffic between an SD-WAN-enabled site and a legacy site uses local breakout and travels directly across the underlay MPLS network.

Figure 64: Local Breakout to Support Migration to SD-WAN



In this case, local breakout is the key to maintaining connectivity between the migrated sites and legacy sites.

Local breakout and NAT

When traffic flows from a tenant VRF to the Internet, NAT must typically be used to translate from the tenant's private network space to the Internet (public) network space.

At spoke sites, the on-premises devices can use Auto-NAT to automatically perform source NAT on all local breakout traffic. At hub sites, Auto-NAT is not available; however, the CSO UI supports manual creation of NAT rules for these on-premises devices.

Local Breakout and DNS

Configuring an on-premises device as a DHCP server for LAN segments allows you to specify DNS server information for end hosts. For a site with local breakout enabled, it is generally recommended to specify more than one name server: an internal server for corporate domain name resolution, and a public or ISP server for Internet-destined local breakout traffic.

Network Security

One of the important security considerations for SD-WAN architectures is providing security for data not only at rest, but also in motion. Data security has been enhanced to allow for the use of multi-level PKI for the data and OAM tunnels. This allows CSO to receive multi-level CA certificates from a CA server, push multiple CA certificates to CPE devices, renew and revoke multiple CA certificates on CPE devices.

CSO supports simple certificate enrollment protocol (SCEP), starting with CSO release 4.1. This allows CSO to:

- Act as SCEP server
- Act as SCEP client
 - Certificate revocation
 - Certificate auto-renew
- Deploy certificates to a CPE/site
- Manage certificates on CPE (site)
- Provide GUI support for CA Server information
- Site/CPE certificate renewals
- Microsoft CA/NDES support
- Broker certificates for each site/CPE

A back-end API is provided for programmatic access to PKI features.

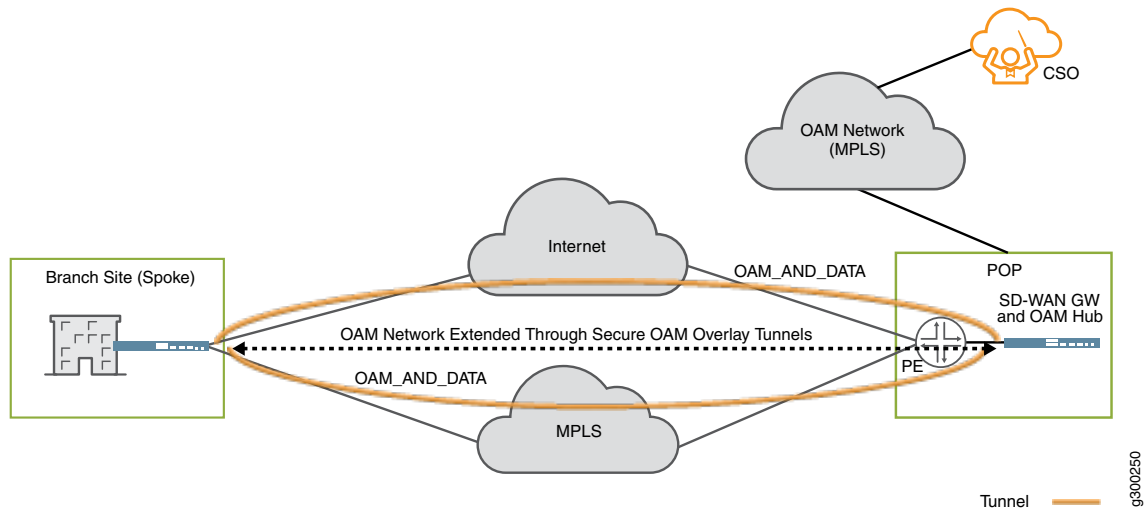
Data Plane

Data plane connections can be configured to use IPsec with PKI-based authentication. When used, the local on-premises device encrypts traffic before transmitting it over the network to the remote site and authentication is handled with public-private key pairs.

Management and Control Plane

CSO connects to and configures on-premises devices using SSH for console and NETCONF connections. Starting with CSO Release 4.0, dedicated OAM overlay tunnels help to enhance secure, end-to-end communications between on-premises devices and CSO. IPsec-encrypted and PKI authenticated OAM tunnels, shown in [Figure 65 on page 93](#), enable on-premises spoke devices to send management, routing, and logging traffic securely over the network to a provider hub. The hub then forwards the traffic to CSO.

Figure 65: Management and Control Plane Security - Secure OAM Network



For more detail, see the [“Secure and Redundant OAM Network”](#) on page 58 section earlier in this guide.

Orchestration and Management Within CSO

IN THIS CHAPTER

- Orchestration and Management Using CSO | 94
- Platform Characteristics | 97

Orchestration and Management Using CSO

The following management and orchestration information pertains to the hardware, software, and services of the CSO platform itself. Therefore, this discussion pertains mostly to those who are implementing an on-premises deployment of CSO. While the same elements exist in cloud-hosted versions of CSO, subscribed users generally have little need to understand the concepts discussed in these sections.


The Contrail SD-WAN, SD-LAN, and NGFW solutions implement management and orchestration with Contrail Service Orchestration (CSO) software. CSO is a scalable and cloud deployable multi-tenant software platform that abstracts the complexity involved in creating and managing network services. Essentially, CSO provides the automation and integration framework for the various components of the solution.

The CSO platform is metadata-driven and uses templates to represent network and resource services. It uses *intent-based policies*, designed to translate the higher-level business rules such as, “send traffic type a, destined for endpoint b, across link c” into repeatable actionable tasks and executes them under the control of a flexible orchestration engine.

Architecture

Some key characteristics of the CSO architecture are shown in [Figure 66 on page 94](#).

Figure 66: CSO Architecture Characteristics



g300232

These characteristics include:

- Container-based, microservices architecture that allows each functional component to be independently deployed and scaled
- Hierarchical central-regional services that can support a large number of network elements (VNFs, PNFs, etc.) across multiple geographical locations
- An orchestration platform to allow full lifecycle management of network devices and virtualized network services, as well as monitoring and visualization
- Open, plugin-based, multi-vendor VNF and EMS support
- Standards-based REST API for OSS/BSS integration
- Flexible deployment options such as on-premises, public cloud, and private cloud.

Orchestration Layers

CSO software is built with multiple layers of abstraction for usability and scalability, as shown in [Figure 67 on page 95](#). The platform implements these layers using orchestration software and controller software.

Figure 67: CSO Orchestration Layers



The Service Orchestration Layer contains the Network Service Orchestrator. The orchestration software has a global view of all resources, including both virtual network functions as well as physical devices.

The orchestration software also enables tenant management, providing end-to-end traffic orchestration, visibility, and monitoring. In addition, Enterprise customers can login to a Customer Portal to enable and manage their own set of services on demand.

The Domain Orchestration Layer contains the Network Service Controller. The orchestration software works together with the controller to manage on-premises (CPE) devices. The controller provides topology and CPE lifecycle management functionality; it also monitors device and link status, and passes this information to the orchestration layer.

The two layers are connected using standard Web-based REST APIs, and both the orchestration layer and the controller layer expose their own sets of APIs, which can be used by any external OSS system to integrate with CSO.

Infrastructure Services and Microservices

CSO uses a fully distributed, docker container-based microservices architecture. The platform consists of several infrastructure services and microservices, which are deployed across the central and regional nodes. Each of these microservices can be independently scaled and deployed, to enable the overall system to scale as needed.

For HA deployments, multiple sets of microservices can be deployed to allow for the failure of orchestrator components. An overlay connection is used between the sets of services to allow for seamless communication of microservices.

Some important microservices include:

- **Tenant site and service management (TSSM):** Provides APIs for tenant, site and service management
- **VNF manager:** Provides APIs to manage virtualized networking services
- **Intent-based Policy and SLA management (PSLAM):** Provides policy and SLA profile object management service to enable SD-WAN functions.
- **Routing manager:** Provides APIs to manage routing operations such as creating VPNs, interfaces to vRRs, enabling routing on CPE devices, etc.
- **Telemetry:** Provides APIs used by fault monitoring and performance monitoring system for collecting service check results from telemetry agents.
- **Activation service:** Provides network activation functions to enable zero touch provisioning of devices.

NOTE: Installation and upgrade of CSO do not apply to cloud-hosted releases. The information on microservices can be found in the [CSO Installation and Upgrade Guide](#) which is specific to the latest on-premises version of CSO which is 5.1.0

Platform Characteristics

Multitenancy with RBAC

The CSO platform has built-in multitenancy support, enabling multiple tenants to coexist on the system. Multitenancy is based on the [OpenStack Keystone](#) model. In this model, each object in the database belongs to a specific tenant and is assigned a tenant ID. When an administrator is granted certain roles to a specific tenant, he or she is authorized with certain rights to all objects belonging to that tenant.

The server enforces multitenancy RBAC at the API level. A user must authenticate with CSO's Identity and Access Management (IAM) microservice to acquire the access token first before any CSO APIs can be invoked. On each API call, the API server enforces the multitenancy RBAC by making sure the object's tenant ID matches the assigned tenant IDs in the access token, and the REST URI is defined in the assigned roles.

RBAC in CSO is object-based. To simplify RBAC application, CSO has pre-defined user roles which provide users assigned to those roles read-only or read-and-write access to specific objects. Custom roles allow administrators to grant users additional access privileges to those or other specific objects.

High Availability and Scale

As noted above, the CSO installation architecture for small environments does not provide HA. The small setup includes one instance each of a central VM and a regional VM; any VM failure renders CSO non-operational. CSO can also scale out for larger environments, using multiple servers with load balancing between them. These servers typically work in active/active HA mode, and services are duplicated across servers. The loss of a server does not impact CSO functionality.

A key design principle is that there is no in-memory state. All the states are transactional and maintained in a database using a Job Manager. CSO ensures that if a node fails, the Job Manager automatically detects the incomplete job and assigns the process to an alternate server for processing.

All CSO infrastructure services, such as database services and message buses, use proven open source software that supports multi-node clustering for HA and scale. These infrastructure service clusters are fine tuned for large scale deployment. The primary CSO configuration and analytics database is built on Cassandra, which is known for its scalability and fault tolerance on commodity hardware and in cloud environments.

All CSO microservices are stateless and do not hold any state between API calls. The application states are kept in the database. Microservices communicate with each other only through RESTful APIs or the highly available message bus. Microservice RESTful APIs are designed to be idempotent (making the same call repeatedly produces the same result) and highly fault tolerant over commodity hardware or cloud environments. CSO microservices are packaged as Docker containers and orchestrated by Kubernetes. Because of the stateless nature and idempotent APIs, each microservice can scale linearly and independently.

Kubernetes allows each microservice to scale up and down automatically based on CPU usage. Kubernetes can also monitor the health of CSO microservice instances, and auto-heal failed instances.

The CSO platform can be deployed on-premises, or in a hybrid or public cloud infrastructure. When deployed across multiple availability zones of the public or private cloud, the platform can survive power and network failures across centers.

Programmability and Integration

All CSO microservices make their functionality accessible via RESTful APIs. Some of these APIs are meant for consumption by other microservices or applications running on CSO, but most are exposed to be consumed by external systems, such as northbound OSS/BSS applications. This allows providers and end customers to automate various tasks, processes, and workflows by invoking these APIs from scripts or backend systems. All microservice APIs are generated from data model descriptions in YANG and can be categorized at a high level as:

- CRUD APIs to create, read, update, and delete resources in the system. These are synchronous APIs that return status and details using HTTP. The caller can define a tenant topology, add or delete sites to this topology, enable no-touch activation of devices at the customer site, setup network connections defined in the topology, enable end user configuration of on-premises devices, monitor device and link status, and more.
- RPC (Remote Procedure Call) APIs to perform operations on these resources. These are typically asynchronous APIs that return completion status and results using Advanced Message Queuing Protocol (AMQP) notifications. The caller can specify an exchange and a routing key for the response message, and the CSO microservice will publish the result notification to that exchange using the specified routing key.

CSO microservices also publish various messages to certain documented exchanges created in the AMQP server, including various resource state change events and alerts. External systems can consume these messages and perform various tasks, thus allowing them to create event-driven automation tasks. One can configure new rules in the FMPM microservice to generate specific alerts and also post alerts on different message buses like Kafka.

The APIs exposed by CSO can be categorized as shown in [Table 11 on page 98](#).

Table 11: CSO APIs

Catalog management	APIs to manage network service descriptors and VNFs
VIM/POP Management	APIs to create define and manage VIM and POP data centers
Topology Management	API to insert and manage end-to-end CPE service topology (logical)
Site/Customer Creation	APIs to manage customer/site objects and association with service topology nodes.

Table 11: CSO APIs (*continued*)

Network Design APIs	APIs to define virtualized services and service chains
Site Activation	APIs to notify vCPE/uCPE device deployment, topology and service placement.
Identity Management	APIs to manage Identity for both enterprise and service provider users
Bootstrap Service	APIs for configuration and management device activation service
Service Placement/Instantiation	APIs to position and manage service chains in customer topology
Device and Service Monitoring	APIs to monitor status of devices, network services, and services topology
Root Cause Analysis/Troubleshooting	APIs trace and correlation engine for events, alarms and logs
Zero touch and Device Management	APIs for activating, provisioning and managing NFX/SRX
Image Management	APIs to manage NFX, SRX, EX, and EX VC software images
SD-WAN	APIs for link provisioning, auto-VPN, discover-VPN, distributed routing
Abstracted Routing	APIs for creating L2/L3 service chains
Public Key Infrastructure (PKI)	APIs for working the PKI security features

For detailed list of APIs, see [Contrail Service Orchestration API Reference](#).

Extensibility and Customization

CSO is architected to allow easy extension and customization of its microservices. These capabilities can be categorized into three main building blocks:

- Plugin-based architecture:** Various microservices, such as EMS, FMPPM, VNFM, Flex, etc., have a plugin-based architecture to allow their behavior to be extended and customized using plugins that can be created and installed without requiring any code changes in the microservice itself. These microservices ship with a certain set of plugins, and new plugins can be created and added in the field.
- Customization of site connectivity topology and activation workflows:** For every site, the WAN-side connection topology, as well as the configuration deployed to the on-premises device(s) during its activation are modeled as device templates. These templates can be modified, or new ones created, in the field to customize the activation workflows and configurations based on each service provider's unique requirements.

- **Designer Tools:** Based on the MANO architecture, new network service packages (VNFs, etc.) can be created on the fly using the designer-tools set of microservices in CSO, and then assigned to tenants. This allows the service provider to create new NFV service offerings for their customers.

Telemetry and Analytics Capability

An important capability of the CSO platform is its ability to collect telemetry data from different devices/VNFs and use it to:

- Store as time series data and make the data query-able from Northbound Applications and the CSO UI to display as charts and graphs.
- Create events for microservices to be able to react to. For example, SLA metrics collected from the devices are published to analyze for link SLA violations, so that the relevant applications can take the appropriate action.
- Publish selected data to Northbound listening applications over Kafka and RabbitMQ.

CSO uses Contrail Analytics nodes to store time series data. Contrail Analytics by itself is a horizontally scalable component that provides high availability as well as the ability to query data through REST APIs. The data from the time series is exposed through CSO APIs to the UI and Northbound applications.

Intent-Based Policies

CSO's user interface puts a strong focus on simplifying and automating many of the main functions an operator needs to perform. This simplification is enabled by modeling enterprise objects and using intent-based policies to configure them.

Intent-based policies allow an operator to configure policies using constructs such as departments, sites, site groups, and application groups. The policy is applied to all relevant devices that match the parameters specified in the matching construct; the operator does not have to worry about configuring the policy explicitly on the devices.

Intents can be expressed as part of various workflows, as described below:

- **Site Onboarding**—During site or hub onboarding, the following intents can be specified:
 - Default link - tenant admin can choose a default link; used as the default overlay path for all traffic which doesn't have a policy saying otherwise.
 - Application breakout - enables site administrators to designate that certain application traffic be routed directly to the Internet from the spoke site.
 - Central breakout - enables Internet-destined traffic to break out directly to the Internet at the enterprise hub.
 - Department breakout - enables site administrators to designate that all Internet-destined traffic from a specific local department be routed directly to the Internet from the spoke site.

- Hub breakout - enables site administrators to designate that all Internet-destined traffic be route directly to the Internet from the provider hub device.
- Site group - allows the same policies to be deployed across a group of sites with similar characteristics.
- Site local Internet breakout - enables site administrators to designate that all Internet-destined traffic be routed directly to the Internet from the spoke site.
- Zscaler breakout - Allow all Internet-destined traffic to be routed to a Zscaler implementation prior to going to the Internet. This breakout can be done locally, centrally, or at the provider hub.

NOTE: Although the above intents can be specified during the site onboarding process, they are not applied until after ZTP.

- **SD-WAN Intent Policy Creation**—Steering and breakout profiles can be created to be used in SD-WAN policies.

Two types of profiles are supported:

- Path-Based Steering Profile – operator explicitly specifies a preferred path for traffic. Traffic matching an SD-WAN policy using this profile will take the preferred path.
- Breakout Profile - operator specifies a breakout type of either, local breakout using underlay networking, backhaul using hub sites for breakout traffic, or local breakout using a cloud-based platform such as Zscaler. The operator also specifies a traffic type profile and preferred path for the breakout traffic. If a WAN link type that matches the preferred path is available at the CPE and enabled for breakout, then the traffic will use that link for breakout traffic. If *any* is selected as the preferred path, then CSO will use all available links that are enabled for breakout in a load-balancing fashion.

An SD-WAN policy can be created by specifying the following elements:

- Source endpoint(s) - site groups, departments
- Destination endpoint(s) - application/application groups
- Action - Steering profile or breakout profile

The operator simply needs to select these high-level elements from the available drop-down menus, and then deploy the policy. CSO takes care of translating these intents into configurations that are pushed to the relevant network devices.

Security Intent-Based Policies

To create firewall policies, the operator does not need to specify the location and connectivity information of the endpoints; instead, CSO uses existing topology information to determine how the relevant endpoints are connected and creates the appropriate security policies to be deployed to the appropriate policy enforcement points.

Firewall policy intents can be defined using the following elements as source and destination identifiers:

- Site
- Department (SRX security zone: Up to 25 departments supported starting in CSO version 4.1)
- Application (L7: signature based)
- Services (protocol based)
- Address objects representing hosts, networks, IP ranges, etc.

Firewall intents are order insensitive, meaning the operator does not have to arrange the intents in the proper order. CSO analyzes all firewall intents and converts them to security policies statements in correct order.

Upgrade and Backward Compatibility

CSO supports seamless upgrades from previous versions, including infrastructure services and microservices upgrade, data migration, device connectivity, and configuration.

The upgrade procedure is an 'offline' activity; all microservices are shut down while the upgrade is in progress. However, network devices (CPEs, hubs, etc.) and the SD-WAN environment overall continue to function normally.

The CSO data model and APIs maintain backward compatibility such that latest version of all CSO microservices support (read/write) data created by previous versions. Migration scripts/additional workflows can also be executed as part of the upgrade process.

Element Management

CSO includes a set of microservices that provide scalable, multi-vendor element management capabilities. These capabilities are used to provide SD-WAN services by managing, orchestrating, and controlling the physical and virtual networking devices that make up the overall solution.

These devices can generally be brought under CSO management in two ways:

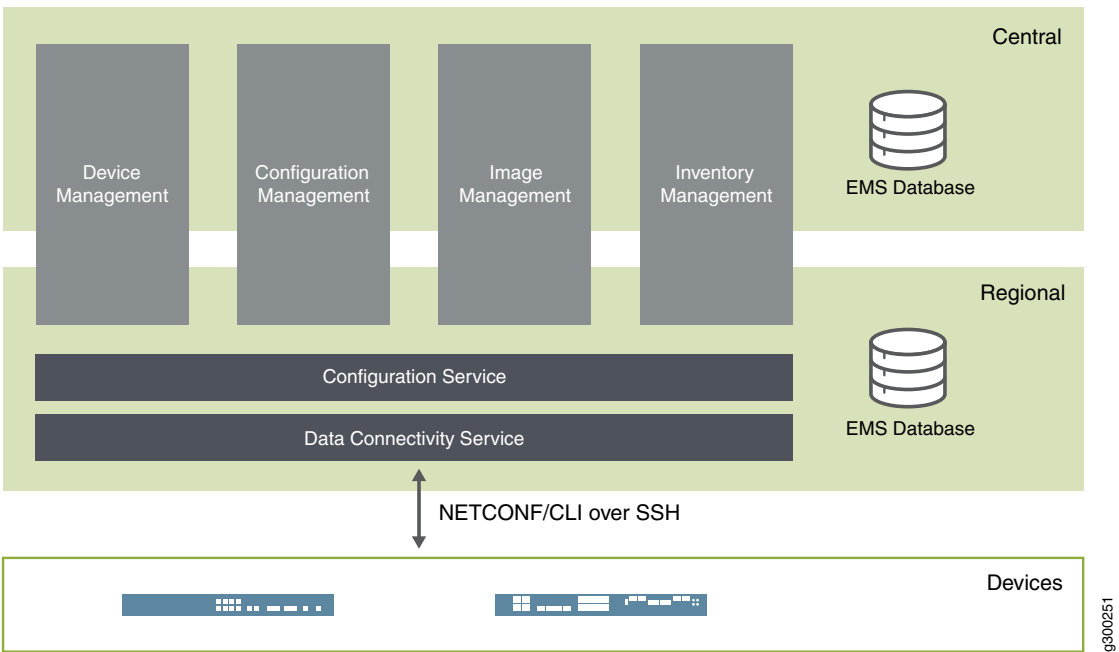
- If the device is already provisioned, it can be discovered by CSO and brought under its management by providing the device's management IP address and admin account credentials. A provider hub device located in a service provider POP is typically discovered using this option.
- For devices that need to be automatically brought online and provisioned, CSO employs a zero-touch mechanism to bring the device under its management. By providing the serial number of the expected device at each site, CSO creates a device object in its database corresponding to each device and prepares the image and configuration that needs to be delivered to it. When the device arrives at the site and is racked up and powered on, it will reach out to the Juniper redirect service (<https://redirect.juniper.net>) to learn how to reach its regional CSO instance. Upon contacting the CSO server, the device receives an assigned software image and initial configuration. Once up and running, CSO performs further actions

on the device, such as bringing up required VMs, provisioning overlay tunnels, installing a telemetry agent, etc.

CSO interacts with network devices using NETCONF or CLI sessions over SSH, thus ensuring that all management communications use a secure, encrypted channel. CSO supports both password-based authentication as well as SSH key-based authentication to the device.

Figure 68 on page 103 illustrates the various microservices that work together to provide CSO's element management capabilities, and how they are distributed across central and regional servers.

Figure 68: CSO Microservices for Element Management



Microservice	Description
Activation Service	Supports secure zero-touch activation of CPE devices through draft-ietf-netconf-zerotouch .
Device Management Service	Manages the lifecycle of devices; devices include VNFs, PNFs, CPEs, PEs, IPsec concentrators, etc.
Config Management Service	Manages the lifecycle of configuration objects, including their versioning as well as their deployment onto devices.
Image Management Service	Maintains a repository of device images and other software packages, and manages the deployment and installation of these onto devices.

Microservice	Description
Inventory Service	Takes care of discovering and managing physical and logical inventory resources on devices.
Template Service	Manages all templates on-boarded into the system, and provides APIs for rendering them using different template engines via plugins; templates can be used to generate configuration or operational commands.
FMPM Provider Service	Centralized service that maintains all FM and PM data, and provides APIs for collecting and querying the data.
FMPM Collector Service	Distributed service that is responsible for collection of FM and PM data from managed entities.
Config Service	Provides APIs to execute commands on managed devices, and acts as the gateway between all microservices and managed devices; has a plugin-based architecture to support multiple management protocols, such as NETCONF/SSH, CLI/SSH, and REST/HTTP.
Device Connectivity Service	Takes care of transport connection establishment and authentication between CSO and the managed devices.

CSO Behind NAT

CSO can be installed behind a NAT gateway. When used, managed devices can reach CSO through a publicly exposed IP address. This option is specified during initial CSO installation, and requires some additional manual configuration of NAT rules once setup is complete.


CSO in the Cloud

While CSO is often installed within the service provider's network, it can also be installed in the cloud, depending on design requirements.

CSO in Public Cloud

[Figure 69 on page 105](#) shows CSO located in an AWS VPC and accessible through a private connection. This is known as a cloud-hosted CSO deployment. CSO Release 5.0.0 is a Juniper-managed, cloud-hosted CSO deployment that follows this model.

Figure 69: CSO in Public Cloud




Implementation characteristics:

- The CSO installation uses private IP addressing.
- The NAT gateway provides a public-facing IP address for CSO.
- The connection between CSO and the hub device uses an MPLS network or a private Internet connection, such as AWS Direct Connect.
- The hub device must use a public IP address for OAM.
- The hub device's IP address must be directly reachable from CSO.
- The spoke device initiates its connection to CSO using the public IP address on the NAT gateway.

CSO on Internet

Figure 70 on page 106 shows CSO located at some other on-Internet location, such as in a private cloud, and accessible directly over the Internet.

Figure 70: CSO on Internet



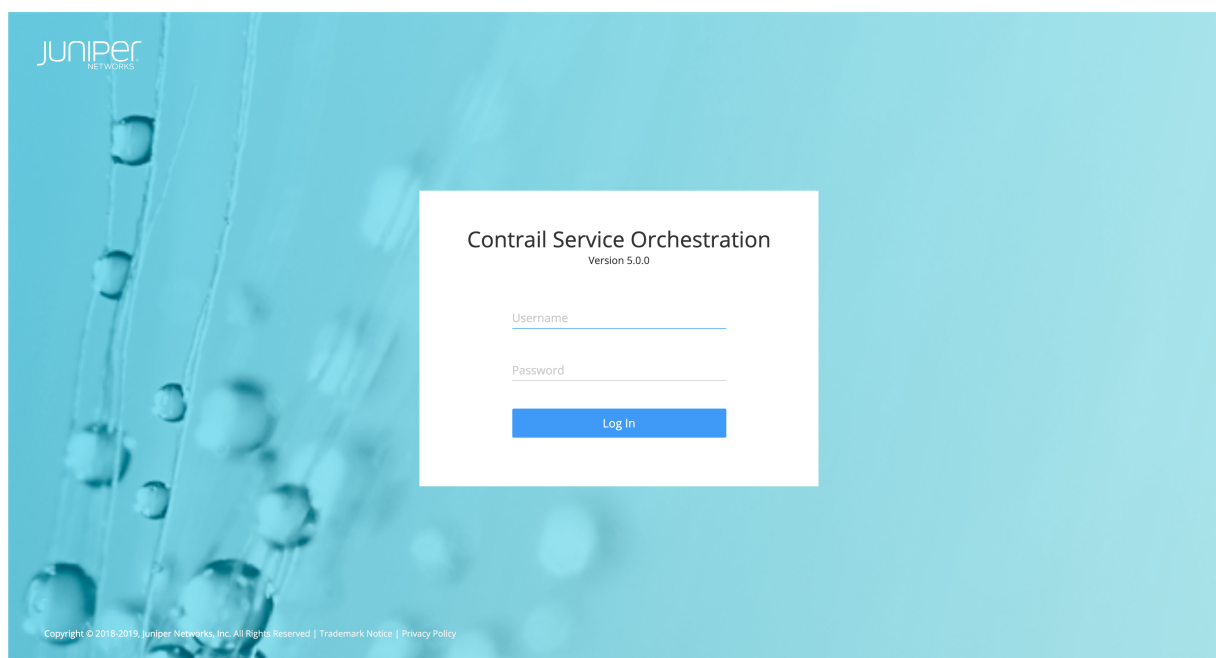
Implementation characteristics:

- The CSO installation uses private IP addressing.
- The NAT gateway provides a public-facing IP address for CSO.
- The connection between CSO and the hub device uses the public Internet.
- The hub device must use a public IP address for OAM.
- The hub device's IP address must be directly reachable from CSO.
- The spoke device initiates its connection to CSO using the public IP address on the NAT gateway.

CSO User Interface

CSO software offers a single Web-based UI to create, configure, and monitor tenants, sites, devices, network topologies, and security and SD-WAN policies. A sample screenshot of the dashboard is shown in [Figure 71 on page 107](#).

Figure 71: CSO User Interface - Dashboard View



Web UI Architecture

The CSO Web UI uses a lightweight framework for building single-pane-of-glass user interfaces in a decoupled way. The UI allows workflows to be dynamically created from independently developed and deployed plugins, which allow the UI to be extended dynamically in a customer environment without any impact on existing functionality.

The UI architecture supports a single, unified dashboard that hosts monitoring widgets. A thumbnail view of the widgets is provided by the framework, and the operator can drag and drop the widgets to compose customized monitoring views. The UI includes a “preferences” API that can be used to read and write UI-related user preferences, such as a preferred sort order or visible subset of columns for a grid instance. These preferences are preserved across user sessions.

Personas

There are two main personas in the Web UI:

- **Service Provider admin**—global access to all operating companies, tenants, and customers; access CSO through the Administration Portal
- **Tenant admin**—customer-specific access; access CSO through the Customer Portal

NOTE: In cloud-hosted deployments of CSO, the Service Provider administrator still logs into the Administration Portal, but as an OpCo administrator. The global SP Admin role is reserved for the hosts of the cloud-hosted deployment.

Operating Companies (OpCos)

CSO Release 4.0 and later supports operating companies (OpCos) in a service provider environment.

In cases where a global service provider is required to have regional business entities to manage customers on a regional basis (for regulatory, billing, or operational reasons), the OpCo construct enables the service provider to extend their CSO platform to enable each regional entity to independently offer SD-WAN services to its own tenants and customers.

When supporting OpCos, the CSO multitenant hierarchy has three levels:

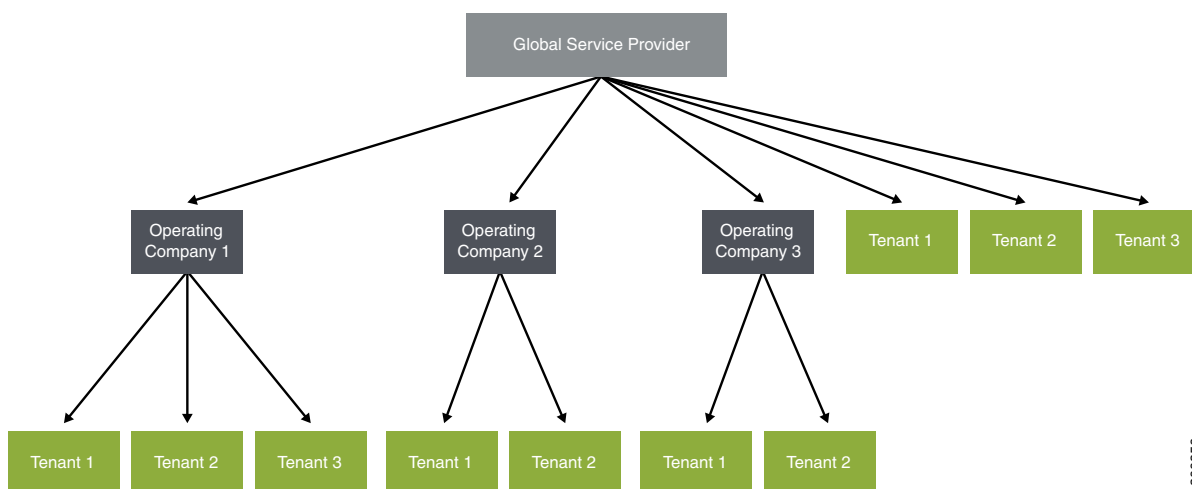
- **Global service provider**—Contains one or more operating companies and its tenants, manages resources at the service provider level, and shares common resources with operating companies and tenants.

NOTE: In a cloud-hosted CSO deployment, there is no user access to the Global Service Provider role/hierarchy.

- **Operating company**—A region-specific service provider that can manage its tenants and provide services to them. Tenants managed by one OpCo are isolated from tenants of another OpCo.
- **Tenant**—Uses the resources provided by the global service provider or OpCo.

Figure 72 on page 108 shows the relationship between the global service provider, operating companies, and tenants.

Figure 72: CSO Multitenant Hierarchy



For more details on CSO portals, user types, and personas, see the [CSO Administration Portal User Guide](#) and [CSO Customer Portal User Guide](#) for Release 5.0.

Release History Table

Release	Description
4.1	Up to 25 departments supported starting in CSO version 4.1
4.0	CSO Release 4.0 and later supports operating companies (OpCos) in a service provider environment.

Operational Workflows - Overview

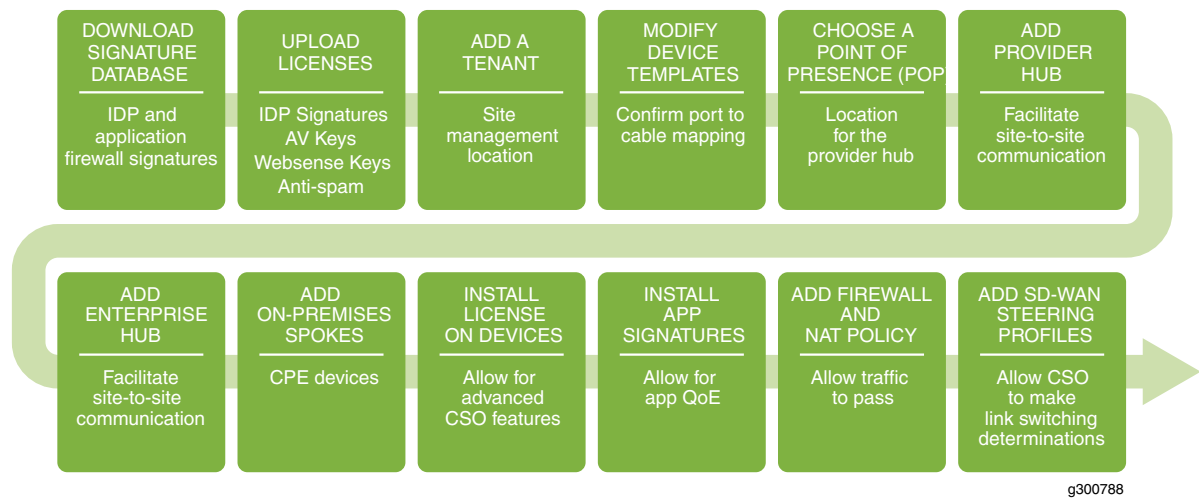
IN THIS CHAPTER

- Operational Workflows | 110

Operational Workflows

While this guide does not cover Contrail SD-WAN or SD-LAN deployment, [Figure 73 on page 110](#) provides a brief description of an end-to-end workflow that the operator might be expected to follow for SD-WAN.

Figure 73: SD-WAN Workflow



For more details on SD-WAN deployment workflows, see the [CSO Deployment Guide](#).

Resiliency and High Availability

IN THIS CHAPTER

- [Resiliency and High Availability | 111](#)

Resiliency and High Availability

The Contrail SD-WAN solution is resilient and highly available at all layers. As a result, the network works seamlessly across failures with as little downtime as possible. The following sections discuss high availability at each layer.

Network Control Plane

The control plane itself is a distributed entity in the Contrail SD-WAN solution. The control plane is enabled using vRRs, which peer with the on-premise devices and set up routing dynamically based on information from the Routing Manager and Policy/SLA Manager microservices.

Route reflectors are deployed in a hierarchical structure. The on-premise devices peer with their closest regional route reflector, which itself peers with the other route reflectors.

Headless Forwarding

If on-premise devices lose connectivity to the route reflector in the SD-WAN controller, the devices are still able to continue forwarding traffic. This is referred to as headless operation. This situation will be sub-optimal as the controller cannot monitor and suggest new routes, but the paths still continue to exist and traffic will be forwarded in a best effort manner.

In headless mode, no new configuration or policy changes are made to the device, and no new data is reported from the device. Once connectivity is restored, the device checks in with the controller to ensure it has the latest routing and configuration information.

Data Plane

CSO Release 3.3 and later support on-premise device redundancy. A site can include a cluster of two nodes, acting as primary and secondary, to protect against device and link failures. If the primary node fails, or the links to it are down, traffic will flow through the secondary node.

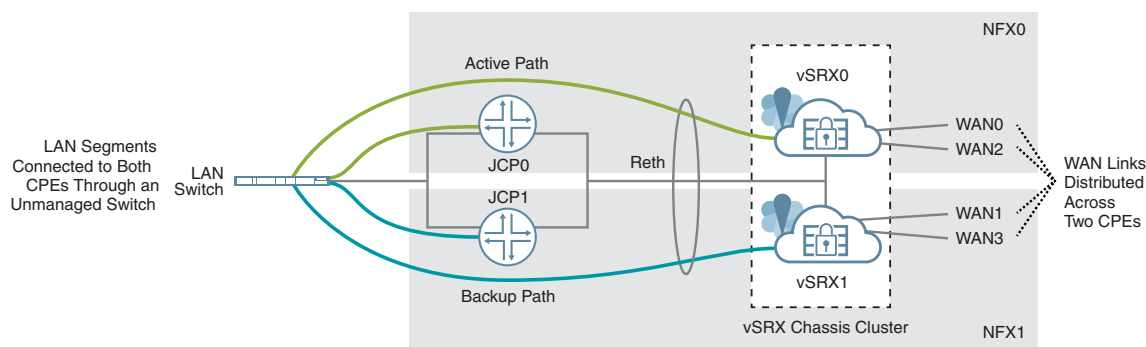
Spoke Redundancy

Spoke sites can include redundancy by interconnecting two CPE devices to create a single, logical, secure router. NFX Series or SRX Series devices can be used.

Using NFX Series Devices

Figure 74 on page 112 shows a spoke redundancy setup using NFX Series devices, each with a vSRX Virtual Firewall installed. The two CPE devices are interconnected by creating an SRX chassis cluster to form a single logical node. The cluster uses a redundant Ethernet (reth) interface to connect to the Junos Control Plane (JCP) component, which acts as a switch to provide connectivity in and out of the devices.

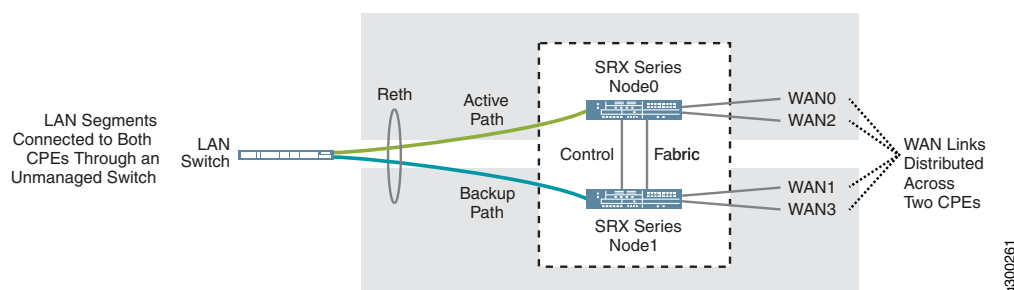
Figure 74: Dual CPE Devices - NFX Series Network Services Platforms



Using SRX Series Devices

Figure 75 on page 113 shows a spoke redundancy setup using SRX Series devices. The two CPE devices are interconnected using chassis clustering to form a single logical node.

Figure 75: Dual CPE Devices - SRX Series Devices



Again, multiple connections on both sides of the devices provide link redundancy. The LAN side uses active/backup links, which are bundled as a reth interface on the SRX cluster. The WAN side uses all four active WAN links, distributed across the two devices.

Failover Scenarios

Table x describes how a spoke redundancy setup will react to various failure scenarios.

Scenario	NFX Behavior	SRX Behavior
Device failure	Node failover of vSRX cluster	Node failover of SRX chassis cluster
GWR vSRX VM failure	Node failover of vSRX cluster	N/A
LAN-side link failure	JCP - LAG based protection for individual link failures vSRX - Reth failover to the other cluster node if all LAN links to a node fail	LAG based protection for individual link failures Reth failover to the other cluster node if all LAN links to a node fail
WAN-side link failure	Same as single-CPE - ECMP across remaining links until SLA enforcement from SD-WAN controller	
Interconnect physical link failure	JCP - LAG based protection	None built in; can add LAG based protection using two interconnected switches between the nodes

Usage Notes

You must use the same device model of NFX Series or SRX Series device and the devices (primary and secondary) must have the same version of Junos OS installed.

The following SD-WAN features are not supported when using spoke redundancy:

- LTE WAN backup link

- Service chain support

For more information on spoke redundancy, see [Device Redundancy Support Overview](#) in the CSO User Guide.

Hub Redundancy (CPE Multihoming)

For hub-and-spoke topologies, redundancy can also be provided on the hub side by deploying two hub devices in an active/backup setup. If the primary hub goes down, or all overlay tunnels to the primary hub fail, traffic switches over to the secondary hub. When the primary hub comes up again and tunnels are established, traffic moves back to the primary hub.

Dual hub mode can also be used in primary/secondary mode. For example, a hub may be primary for half of the spokes, and secondary for the other half. This way the load is distributed in an active/active manner across all pairs of hub devices. Note that this mode requires meshing the hub devices to maintain flow symmetry across the network.


Design Options

There are several ways to implement redundancy between hub and spoke devices, depending on design requirements:

- **Single Spoke Device Multihomed to Dual Hub Devices; Single Access**

[Figure 76 on page 114](#) shows how a single spoke device could be multihomed to dual hub devices with single access.

Figure 76: One Spoke Device, One Tunnel to Each Hub




In this scenario, the hub devices are in an active/backup configuration, and spoke site prefixes are routed to the active hub.

- **Single Spoke Device Multihomed to Dual Hub Devices; Multiple Access**

[Figure 77 on page 115](#) shows how a single spoke device could be multihomed to dual hub devices with multiple access.

Figure 77: One Spoke Device, Two Tunnels to Each Hub




As in the previous scenario, the hub devices are in an active/backup configuration, and spoke site prefixes are routed to the active hub. In addition, the spoke site has overlay links to each hub through each access network. APBR routes traffic from CPE device to active hub over all possible overlays.

- **Clustered Spoke Devices Multihomed to Dual Hub Devices; Multiple Access**

Figure 78 on page 115 shows how dual CPE devices could be multihomed to dual hub devices with multiple access from each CPE device.

Figure 78: Spoke Cluster, One Tunnel to Each Hub



As in the previous scenario, the hub devices are in an active/backup configuration, spoke site prefixes are routed to the active hub, and APBR routes traffic from active CPE device to active hub over all possible overlays. In this scenario, the CPE devices are also in an active/passive configuration.

The spoke site has eight overlay tunnels:

- Active CPE to active hub - two active links
- Active CPE to backup hub - two backup links
- Backup CPE to active hub - two active links
- Backup CPE to backup hub - two backup links

NOTE: Both hubs to which a CPE device is multihomed must be the same type of device.

NOTE: If using NAT, a hub switchover due to a primary hub failure may cause site-to-Internet and site to cloud application sessions to flap, as the NAT behavior adjusts to the change. Site-to-site sessions will continue to work through the switchover.