Contrail Service Orchestration Release Notes

Release 3.3.0
28 January 2019
Revision 9

These Release Notes accompany Release 3.3.0 of Juniper Networks® Contrail Service Orchestration (CSO). They contain installation and upgrade information, and they describe new and changed features, limitations, and known and resolved issues in the software.

Contents

Introduction

Juniper Networks Contrail Service Orchestration (CSO) transforms traditional branch networks, offering opportunities for high flexibility of the network, rapid introduction of new services, automation of network administration, and cost savings. The solution supports both Juniper Networks and third-party virtualized network functions (VNFs) that network providers use to create network services.

CSO Release 3.3.0 is a secure software-defined WAN (SD-WAN) solution that builds on the capabilities of CSO Release 3.2 and the Cloud CPE solution. The following are the highlights of the features available in Release 3.3.0:

- SD-WAN
  - Dual CPE (spoke redundancy) support
  - Backup link support
  - LTE Modem (USB) support on NFX250 devices
  - High availability for virtual route reflectors (VRRs)
  - Support for traffic type profiles
  - Cloud spoke sites on Amazon Web Services (AWS) virtual private cloud (VPC)
  - SD-WAN reports
- Security management
  - Firewall policy intents based on users or user groups
  - CSO integration with Active Directory (AD) through Juniper Identity Management Service (JIMS)
  - Offline signature download
- Infrastructure
  - Device Return Material Authorization (RMA)
  - Upgrade from CSO Release 3.2.1 to Release 3.3.0
  - Multiregion support (centralized deployment)
  - Health check for infrastructure components
  - Provisioning VMware ESXi VMs using the provisioning tool
- Miscellaneous
  - Personalization of the unified Administration and Customer Portal
  - Single-sign on (SSO) initiated by an identity provider (IdP)

CSO can be implemented by service providers to offer network services to their customers or by Enterprise IT departments in a campus and branch environment. In these release

notes, service providers and Enterprise IT departments are called *service providers*, and the consumers of their services are called *customers*.

The solution offers the following deployment models:

- Cloud CPE Centralized Deployment Model (*centralized deployment*)

  In the centralized deployment, customers access network services in a service provider's cloud. Sites that access network services in this way are called *cloud sites* in these release notes.

- Cloud CPE Distributed Deployment Model (*distributed deployment*)

  In the distributed deployment, customers access network services on a CPE device, located at a customer's site. These sites are called *on-premise sites* in these release notes.

- A combined centralized and distributed deployment

  In this combined deployment, the network contains both cloud sites and on-premise sites. One customer can have both types of sites; however, you cannot use the same network service package for cloud sites and on-premise sites. If you require the same network service for cloud sites and on-premise sites, you must create two identical network service packages with different names.

CSO uses the following components for the NFV environment:

- When end users access network services in the cloud:
  - Network Service Orchestrator provides ETSI-compliant management of the life cycle of network service instances.

    This application includes RESTful APIs that you can use to create and manage network service catalogs.
  - Contrail Cloud Platform provides the underlying software-defined networking (SDN), NFV infrastructure (NFVI), and the virtualized infrastructure manager (VIM).
- When end users access network services on a local CPE device:
  - Network Service Orchestrator, together with Network Service Controller, provides ETSI-compliant management of the life cycle of network service instances.
  - Network Service Controller provides the VIM.
  - The CPE device provides the NFV infrastructure .

The following CSO components connect to Network Service Orchestrator through its RESTful API:

NOTE: **The Administration and Customer Portals are unified into a single portal with role-based access control (RBAC) enforcement.**

- Administration Portal, which you use to set up and manage your virtual network and customers through a GUI.
- Customer Portal, which is an application that you can provide to customers to enable them to manage sites and services for their organizations through a GUI.
- The Designer Tools, which enable design, creation, management, and configuration of network services through a GUI. Network services are stored in the network service catalog.
- Service and Infrastructure Monitor, which works with Icinga, an open-source enterprise monitoring system to provide real-time data about CSO, such as the status of virtualized network functions (VNFs), virtual machines (VMs), and physical servers; information about physical servers' resources; components of a network service (VNFs and VMs hosting a VNF); counters and other information for VNFs.

  You can deploy CSO in a trial or production environment. Table 1 on page 5 shows the number of sites and VNFs supported for each environment.

Table 1: Number of Sites and VNFs Supported

Contrail Service Orchestration Environment Type	Number of VNFs Supported for a Centralized Deployment	Number of Sites and VNFs Supported for a Distributed Solution	Number of Tenants	Number of Sites Per Tenant	Number of Sites Supported for an SD-WAN Deployment
Trial environment without HA	10 VNFs	25 sites, 2 VNFs per site	5	1	Up to 5 full mesh sites
				5	Up to 25 hub and spoke sites
Trial environment with HA	100 VNFs, 20 VNFs per Contrail compute node	200 sites, 2 VNFs per site	10	5	Up to 50 full mesh sites
				20	Up to 200 hub and spoke sites
Production environment without HA	500 VNFs, 20 VNFs per Contrail compute node	200 sites, 2 VNFs per site	10	5	Up to 50 full mesh sites
				20	Up to 200 hub and spoke sites
Production environment with HA	500 VNFs, 20 VNFs per Contrail compute node	3000 sites, 2 VNFs per site	50	10	Up to 500 full mesh sites
			50	60	Up to 3000 hub and spoke sites

Installation and Upgrade

NOTE: During the installation or upgrade process, ensure that you save the passwords for each infrastructure component when they are displayed on the console because these passwords are encrypted and are not displayed again.

In addition, during the installation, ensure that you save the Administration Portal password that is displayed on the console. For the upgrade, you must log in using the password configured for the previously installed version of CSO.

Installation Instructions

A full-version installer is available for CSO Release 3.3.0, which can be used for both trial and production environments.

NOTE: You do not need Internet access from the CSO server to install CSO.

For more information, follow the instructions in the Deployment Guide or the README file that is included with the software installation package.

Upgrade Instructions

NOTE: You can upgrade to CSO Release 3.3.0 only from CSO Release 3.2.1. If your installed version of CSO is not Release 3.2.1, then you must perform a fresh installation of CSO 3.3.0.

If your installed version is CSO Release 3.2.1, you can use a script (**upgrade.sh**) to directly upgrade to CSO Release 3.3.0. You can roll back to CSO Release 3.2.1, if the upgrade is unsuccessful.

For more information, see *Upgrading Contrail Service Orchestration Overview* in the Deployment Guide.

Post-Installation and Post-Upgrade Instructions

- After you successfully install or upgrade CSO, you must do the following:
  - Configure SMTP settings—After you log in for the first time to the CSO GUI, you must configure the SMTP settings for your deployment on the SMTP Settings page (**Administration > SMTP**).
  - Configure name servers on CPE devices—Use custom properties to provide the name server details when you are adding a tenant.
- Accessing GUIs—We recommend that you use Google Chrome version 60 or later to access the CSO GUIs. For more information, see "Accessing the CSO GUIs" on page 32.

Installation and Upgrade Limitations

- For SD-WAN deployments, CPE devices behind NAT are not supported.
- If the Kubernetes minion node in the central or regional microservices virtual machine (VM) goes down, the pods on the minion node are moved to the Kubernetes master node. When you bring the minion node back up, the pods do not automatically rebalance across the nodes.
- In CSO Release 3.3.0, the VM on which the virtual route reflector (VRR) is installed supports only one management interface.
- Before upgrading vSRX by using CSO, execute the **request system storage cleanup** command on the vSRX by using Junos OS CLI.

Software Installation Requirements for NFX250 Network Services Platform

The NFX250 requires Junos OS Release 15.1X53-D472 for the CSO Release 3.3.0.

When you set up a distributed deployment with an NFX250 device, you must use Administration Portal or the CSO API to:

1. Upload the software image to CSO.
2. Specify this image as the boot image when you configure activation data.

For more information, see https://www.juniper.net/documentation/en_US/release-independent/junos/information-products/pathway-pages/nfx-series/product/ .

Software Downloads

Table 2 on page 8 displays the supported versions and download links for CSO Release 3.3.0 and associated software components.

Table 2: CSO and Associated Software Components

Product	Supported Version	Download Link
Contrail Service Orchestration	3.3.0	https://www.juniper.net/support/downloads/?p=cso
Juniper Identity Management Service (JIMS)	1.1.0R1	Pre-bundled with CSO and also available here: https://www.juniper.net/support/downloads/?p=jims#sw
Contrail Analytics	4.0.3.0-162	Pre-bundled with CSO
Contrail Cloud Platform	3.2.5	https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69888.html
NFX250 CPE device	Junos OS 15.1X53-D472	• Installation Package: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74823.html • Installation Media (USB): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74824.html
SRX Series CPE device	Junos OS 15.1X49-D133	• SRX300, SRX320, SRX340, SRX345, and SRX550 High Memory Services Gateway (SRX550M): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74828.html • SRX1500: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74830.html • SRX1500 (USB): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74832.html • SRX1500 (PXE): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74833.html • SRX4100, SRX4200: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74829.html • SRX4100, SRX4200 (PXE): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74834.html • SRX4100, SRX4200 (USB): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74831.html
vSRX	Junos OS 15.1X49-D133	• vSRX (Upgrade compressed tar file (TGZ)): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74836.html • vSRX (KVM Appliance): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74837.html • vSRX (Hyper-V Image): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74838.html • vSRX (VMware Appliance with SCSI virtual disk (.ova)): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74839.html • vSRX (VMware Appliance with IDE virtual disk (.ova)): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/74840.html
MX Series (hub device)		https://www.juniper.net/support/downloads/

New and Changed Features in Contrail Service Orchestration Release 3.3.0

This section describes the new features or enhancements to existing features in CSO Release 3.3.0.

Centralized Deployment

- **Support for multiple regions in a centralized deployment—**From CSO Release 3.3.0 onward, you can configure a maximum of three regions for centralized deployments. The regions are used to group services for various business reasons such as location, proximity, service distribution, and load. During CSO installation, you can add the regions and deploy the infrastructure services on the regions. You can assign the regions while creating a point of presence (POP). If you do not select the regions, then the default region **regional** is selected.

Device Management

- **Support for Return Material Authorization (RMA) for a device—**From CSO Release 3.3.0 onward, you can recall a defective device, and replace it with a new or restored device by using the RMA process. Configurations that are customized using configuration templates are automatically restored during this process.

  

  NOTE: In CSO Release 3.3.0, the RMA process is not completely automated. You must manually push licenses, application signatures, certificates, and policies to complete the RMA process.

  When the new or restored device is in the **PROVISIONED** state:

  - In SD-WAN deployments, you can proceed to configure the device by manually pushing application signatures, certificates, and policies.

- In hybrid WAN deployments, service chains are restored automatically.

SD-WAN

- **Device redundancy support—**CSO Release 3.3.0 supports device redundancy for large enterprise SD-WAN on-premise spoke sites. You can configure an SD-WAN site with two CPE devices to act as primary and secondary devices and protect the site against device and link failures. If the primary device fails, the secondary device takes over the traffic processing.

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

- **Support for configuring the backup link during site addition—**From CSO Release 3.3.0 onward, you can optionally specify a backup link when you add a site. When the primary links are down, the site can use the backup link to route traffic. When a primary link comes back online, CSO monitors the performance on the primary link and when the primary link meets the SLA requirements, the traffic is switched back to the primary link. However, SLA data is not monitored for the backup link.
- **Support for LTE as a backup link on NFX250 devices—**CSO Release 3.3.0 supports LTE as a backup link on NX250 devices. If an LTE access type is configured for a WAN link, then, by default, the WAN link is used only as a backup link. You can configure the LTE access type while creating an on-premise spoke site.
- **High availability for virtual route reflectors—**CSO Release 3.3.0 supports high availability (HA) for virtual route reflectors (VRRs). In an SD-WAN solution, multiple VRRs can be installed on the regional servers. BGP sessions are established between hub-and-spoke devices and VRRs.
- **Support for traffic type profiles—**CSO Release 3.3.0 introduces traffic type profiles that enable MSP administrators and tenant administrators to configure CoS parameters that meet specific business requirements. Traffic type profiles enable you to define a traffic type and to configure parameters such as priority, buffer and bandwidth allocations, probe parameters, and DiffServ code point (DSCP) values for the traffic type.
- **Support for cloud spoke sites on AWS VPC—**From CSO Release 3.3.0 onward, a tenant administrator can create and configure a cloud spoke site for an SD-WAN endpoint in an Amazon Web Services (AWS) virtual private cloud (VPC). To create a cloud spoke site, log in to Customer Portal and select **Sites > Site Management > Add > Cloud Spoke**. You must select the **vSRX_AWS_SDWAN_Endpoint_option_1** device template while creating a cloud spoke site.
- **Support for generating SD-WAN reports—**From CSO Release 3.3.0 onward, you can generate SD-WAN reports to view the SLA performance of all sites in a tenant and specific sites in a tenant. Using SD-WAN report definitions, you can generate the following SD-WAN reports:

  - **SD-WAN Tenant Performance Reports**—Enable you to view the parameters that measure the SLA performance across all sites in a tenant.
  - **SD-WAN Site Performance Reports**—Enable you to view the parameters that measure the SLA performance of specific sites in a tenant. You can generate reports for up to five sites in a tenant.

- **Support for viewing application visibility filtered based on departments**—From CSO Release 3.3.0 onward, you can filter and view the application visibility data for departments within a single tenant.

Security Management

- **Support for offline download of signature database**—From CSO Release 3.3.0 onward, when there is no Internet connectivity, CSO provides the option to download the signature database either from a local webserver hosted on your PC, or from any webserver accessible through the intranet.

  NOTE: You must first download the signature database from the Juniper Networks-hosted webserver to your local webserver, before performing an offline download.

- **Support for user-based firewall policy intents**—From CSO Release 3.3.0 onward, you can define user-based firewall policy intents, which enable you to permit, reject, or deny traffic based on users or user groups, on SRX Series devices and vSRX instances.
- **Support for Juniper Identity Management Service**—CSO Release 3.3.0 supports the Juniper Identity Management Service (JIMS). JIMS collects user identity information from a configured Active Directory and makes it available to SRX Series devices or vSRX instances.

  You can download and install JIMS, configure the CSO client on JIMS to obtain user identity information from the configured Active Directory, and use CSO and JIMS to manage user-based firewall policy intents on SRX Series devices and vSRX instances.

Unified Portal

- **Support for single-sign on (SSO) initiated by an Identity Provider (IdP)**—From CSO Release 3.3.0 onward, the Identity Provider (IdP) initiation method is supported to authenticate MSP and tenant users. In this method, users are authenticated by using the SSO Server and then the CSO application is launched.
- **Personalizing the unified Administration and Customer Portal**—From CSO 3.3.0 onward, you can personalize the unified Administration and Customer Portal. You can personalize the login page, top-left logo, and reports, and apply a font style and color palette to the left navigation bar and menu. You can also create, edit, and delete a custom color palette. You can also upload custom font styles and preview the custom color palette settings before you apply the settings.

Miscellaneous

- **Health check for infrastructure components—**From CSO Release 3.3.0 onward, you can run a script (**components_health.sh**) to perform a health check of all infrastructure components. This script detects whether any infrastructure component has failed and displays the health status of each infrastructure component.
- **Provisioning VMware ESXi VMs by using the provisioning tool—**From CSO Release 3.3.0 onward, if you use VMware ESXi VMs, you can use the provisioning tool—**provision_vm_ESXI.sh**—to create and configure VMs for CSO.
- **Ability to push a license to devices—**From CSO Release 3.3.0 onward, MSP administrators can apply licenses to devices from Administration portal. MSP administrators can select any of the uploaded licenses from the License Files page, click **Push License**, and select the devices to which they want to apply the license.

  If licenses are available for a tenant, the licenses are pushed to the device as part of the ZTP workflow.

Unsupported Features

The CSO Release 3.3.0 documentation describes some features that are present in the application but that have not yet been fully qualified by Juniper Networks. If you use any of these features before they have been fully qualified, it is your responsibility to ensure that the feature operates correctly in your targeted configuration.

The following features are present but unsupported in this release:

- **Support for Application Quality of Experience (AppQoE)—**CSO Release 3.3.0 supports AppQoE (on SRX series devices and vSRX instances) to improve the user experience at the application level. AppQoE is enabled when the SD-WAN mode for the tenant is set to **Real-time Optimized**. In real-time-optimized mode, CSO monitors the end-to-end application traffic for class-of-service (CoS) and SLA compliance.

  NOTE: **Because AppQoE is an unsupported feature, SD-WAN with full mesh topology for dynamic policies is also not supported.**

- **Support for RMA on dual CPE SRX devices**

Servers, Software, and Network Devices Tested

- Node Servers and Servers Tested in Contrail Service Orchestration
- Software Tested for COTS Servers
- Network Devices and Software Tested for the Contrail Cloud Platform (Centralized Deployment)
- Network Devices and Software Tested for Use with CPE Devices (Distributed Deployment)

Node Servers and Servers Tested in Contrail Service Orchestration

CSO uses commercial off-the-shelf (COTS) node servers or servers for both the centralized and distributed deployments for the following functions:

- CSO central and regional servers
- Contrail Analytics servers
- Contrail Cloud Platform in the centralized deployment

Table 3 on page 13 lists the node servers and servers that have been tested for these functions in CSO. You should use these specific node servers or servers for CSO.

Table 3: COTS Node Servers and Servers Tested in the Cloud CPE and SD-WAN Solutions

Option	Vendor	Model	Type
1	QuantaPlex	T41S-2U 4-Node server	Multinode server accepting 4 nodes
2	Supermicro	SuperServer Model SYS-2028TPHC1TR-OTO-4	Multinode server accepting 4 nodes
3	Dell	PowerEdge R420 rack server	1U rack-mounted server

Software Tested for COTS Servers

Table 4 on page 13 shows the software that has been tested for CSO. You must use these specific versions of the software when you implement CSO.

Table 4: Software Tested for the COTS Nodes and Servers

Description	Version
Operating system for all COTS nodes and servers	Ubuntu 14.04.5 LTS NOTE: Ensure that you perform a fresh install of Ubuntu 14.04.5 LTS on the CSO servers in your deployment because upgrading from a previous version to Ubuntu 14.04.5 LTS might cause issues with the installation.
Operating system for VMs on CSO servers	• Ubuntu 14.04.5 LTS for VMs that you configure manually and not with the provisioning tool. • The provisioning tool installs Ubuntu 14.04.5 LTS in all VMs.
Hypervisor on CSO servers	KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0
Additional software for CSO servers	Secure File Transfer Protocol (SFTP)
Software defined networking (SDN) for a centralized deployment	Contrail Cloud Platform Release 3.2.5 with Heat v2 APIs
Contrail Analytics	Contrail Release 4.0.3.0-162

Network Devices and Software Tested for the Contrail Cloud Platform (Centralized Deployment)

The Contrail Cloud Platform has been tested with:

- The network devices described in Table 5 on page 14.
- The software described in Table 6 on page 14.

You must use these specific versions of the software for CSO Release 3.3.0.

Table 5: Network Devices Tested for the Centralized Deployment

Function	Device	Model	Quantity
SDN gateway router	Juniper Networks MX Series 3D Universal Edge Router	MX80-48T router with two 10-Gigabit Ethernet XFP optics	1
Management switch	Juniper Networks EX Series Ethernet Switch	EX3300-48T switch with: • 48 10/100/1000-Gigabit Ethernet interfaces • 4 built-in 10-Gigabit Ethernet SFP transceiver interfaces	1
Data switch	Juniper Networks QFX Series Switch	QFX 5100-48S-AFI switch with: • 48 SFP+ transceiver interfaces • 6 QSFP+ transceiver interfaces	1

Table 6: Software Tested in the Centralized Deployment

Function	Software and Version
Operating system for MX Series router	Junos OS Release 14.2R3
Operating system for QFX Series switch	Junos OS Release 13.2X51-D38
Hypervisor on CSO servers	KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0
Element management system software	EMS microservice Junos Space Network Management Platform Release 15.1R1 (See *VNFs Supported by the Cloud CPE Solution* for VNFs that require this product)
Software defined networking (SDN), including Contrail Analytics, for a centralized deployment	Contrail Release 3.2.5 with OpenStack Mitaka
Virtualized infrastructure manager (VIM) and virtual machine (VM) orchestration	OpenStack Mitaka
Authentication and Authorization	OpenStack Mitaka
Network Functions Virtualization (NFV) Orchestrator	CSO Release 3.3.0

Network Devices and Software Tested for Use with CPE Devices (Distributed Deployment)

The distributed deployment has been tested with:

- The network devices described in Table 7 on page 15.
- The software described in Table 8 on page 16.

  You must use these specific versions of the software when you implement the distributed deployment.

Table 7: Network Devices Tested for the Distributed Deployment and SD-WAN Implementation

Function	Device	Model
PE router and IPsec concentrator (Hybrid WAN distributed deployment only)	Juniper Networks MX Series 3D Universal Edge Router	• MX960, MX480, or MX240 router with a Multiservices MPC line card • MX80 or MX104 router with Multiservices MIC line card • Other MX Series routers with a Multiservices MPC or Multiservices MIC line card See MPCs Supported by MX Series Routers and MICs Supported by MX Series Routers for information about MX Series routers that support Multiservices MPC and MIC line cards.
Cloud hub device (SD-WAN implementation only)	Juniper Networks MX Series 3D Universal Edge Router Juniper Networks SRX Series Services Gateway	• MX104, MX240, MX480, or MX960 router with an Multiservices MIC line card. See MPCs Supported by MX Series Routers and MICs Supported by MX Series Routers for information about MX Series routers that support Multiservices MPC and MIC line cards. • SRX1500 Services Gateway • SRX4100 Services Gateway • SRX4200 Services Gateway
On-premise hub device (SD-WAN implementation only)	Juniper Networks SRX Series Services Gateway	• SRX1500 Services Gateway • SRX4100 Services Gateway • SRX4200 Services Gateway
CPE device (Hybrid WAN deployment) or spoke device (SD-WAN implementation)	• NFX250 Series Network Services Platform • SRX Series Services Gateway • vSRX on an x86 server	• NFX250-LS1 device • NFX250-S1 device • NFX250-S2 device • SRX300 Services Gateway • SRX320 Services Gateway • SRX340 Services Gateway • SRX345 Services Gateway • SRX550 High Memory Services Gateway (SRX550M) • vSRX

Table 8: Software Tested in the Distributed Deployment and SD-WAN Solution

Function	Software and Version
Hypervisor on CSO servers	KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0
Authentication and Authorization	OpenStack Mitaka
Network Functions Virtualization (NFV)	CSO Release 3.3.0
Contrail Analytics	Contrail Release 4.0.3.0-162
NFX Software	Junos OS Release 15.1X53-D472
Routing and Security for NFX250 device	vSRX KVM Appliance 15.1X49-D133
Operating system for vSRX used as a CPE device on an x86 server	vSRX KVM Appliance 15.1X49-D133
Operating system for SRX Series Services Gateway used as a CPE device or spoke device	Junos OS Release 15.1X49-D133
Operating system for MX Series router used as PE router	Junos OS Release 16.1R3.00
Operating system for MX Series Router used as a hub device for an SD-WAN implementation	Junos OS Release 16.1R5.7
Operating system for SRX Series Services Gateway used as a hub device for an SD-WAN implementation	Junos OS Release 15.1X49-D133

Hardware, Software, and Virtual Machine Requirements for Contrail Service Orchestration

Minimum Hardware Requirements for Contrail Service Orchestration

Table 3 on page 13 lists the makes and models of node servers and servers that you can use in CSO. When you obtain node servers and servers for CSO, we recommend that you:

- Select hardware that was manufactured within the last year.
- Ensure that you have active support contracts for servers so that you can upgrade to the latest firmware and BIOS versions.

The number of node servers and servers that you require depends on whether you are installing a trial or a production environment.

Table 9 on page 17 shows the required hardware specifications for node servers and servers in a trial environment.

Table 9: Trial Environments (Without HA and With HA)

Function	Trial Environment (Without HA)	Trial Environment (With HA)
Node or Server Specification		
Storage	Greater than 1 TB of one of the following types: • Serial Advanced Technology Attachment (SATA) • Serial Attached SCSI (SAS) • Solid-state drive (SSD)	Greater than 1 TB of one of the following types: • SATA • SAS • SSD
CPU	One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.5 GHz or higher specification	One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.5 GHz or higher specification
Network interface	One Gigabit Ethernet (GE) or 10 Gigabit Ethernet interface	One Gigabit Ethernet or 10 Gigabit Ethernet interface
CSO Servers (includes Contrail Analytics in a VM)		
Number of nodes or servers	1	3
vCPUs per node or server	48	48
RAM per node or server	256 GB	256 GB
Contrail Cloud Platform for a Centralized Deployment		
Number of nodes or servers	1	4–8 • 3 nodes for Contrail controller, and analytics • 1–4 Contrail compute nodes
vCPUs per node or server	48	48
RAM per node or server	256 GB	256 GB

Table 10 on page 17 shows the required hardware specifications for node servers and servers in a production environment.

Table 10: Production Environment (Without HA and with HA)

Server Function	Values
Node or Server Specification	

Table 10: Production Environment (Without HA and with HA) (continued)

Server Function	Values
Storage	Greater than 1 TB of one of the following types: • SATA • SAS • SSD
CPU	One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.5 GHz or higher specification
Network interface	One Gigabit Ethernet (GE) or 10 Gigabit Ethernet interface
CSO Servers	
Number of nodes or servers for a production environment without HA	3 • 1 central server • 1 regional server NOTE: This specification is for a single region setup; for each additional region that you add, another regional server is needed. • 1 Contrail Analytics server
Number of nodes or servers for a production environment with HA	9 • 3 central servers • 3 regional servers NOTE: This specification is for a single region setup; for each additional region that you add, three more regional servers are needed. • 3 Contrail Analytics servers
vCPUs per node or server	48
RAM per node or server	256 GB
Contrail Analytics Server for a Distributed Deployment	
Number of nodes or servers	1
vCPUs per node or server	48
RAM per node or server	256 GB
Contrail Cloud Platform for a Centralized Deployment	
Number of nodes or servers	4–28 • 3 nodes for Contrail controller, and analytics • 1–25 Contrail compute nodes

Table 10: Production Environment (Without HA and with HA) (continued)

Server Function	Values
vCPUs per node or server	48
RAM per node or server	256 GB

Software and Virtual Machine Requirements

You must use the software versions that were tested in CSO. This section shows the VMs required for each type of environment.

Table 11 on page 19 shows complete details about the VMs required for a trial environment without HA.

Table 11: Details of VMs for a Trial Environment (Without HA)

Name of VM	Components That Installer Places in VM	Resources Required	Ports to Open
csp-installer-vm	—	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm	All microservices, including GUI applications	• 8 vCPUs • 48 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm	Third-party applications used as infrastructure services	• 4 vCPUs • 24 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm	All microservices, including GUI applications	• 4 vCPUs • 24 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-regional-sblb	Load balancer for device to Fault Management Performance Management (FMPM) microservice connectivity	• 4 vCPUs • 8 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy virtualized network functions (VNFs) that use this EMS	• 4 vCPUs • 16 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.

Table 11: Details of VMs for a Trial Environment (Without HA) (continued)

Name of VM	Components That Installer Places in VM	Resources Required	Ports to Open
csp-contrailanalytics-1	Contrail Analytics for centralized and distributed deployments	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-vrr-vm	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-central-k8mastervm	Central K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	-
csp-regional-k8mastervm	Regional K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	-

Table 12 on page 20 shows complete details about VMs and microservice collections required for a production environment without HA.

Table 12: Details of VMs for a Production Environment Without HA

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-installer-vm	—	• 4 vCPUs • 64 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.

Table 12: Details of VMs for a Production Environment Without HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-regional-sblb	Load balancer for device to microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-elkvm	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-elkvm	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-vrr-vm	VRR	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-central-k8mastervm	Central K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	-
csp-regional-k8mastervm	Regional K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	-
csp-contrailanalytics-1	Contrail Analytics for centralized and distributed deployments	• 48 vCPUs • 256 GB RAM • 1 TB hard disk storage	See Table 15 on page 27.

Table 13 on page 21 shows complete details about the VMs for a trial environment with HA.

Table 13: Details of VMs for a Trial Environment with HA

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-installer-vm	—	• 4 vCPUs • 48 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.

Table 13: Details of VMs for a Trial Environment with HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-central-infravm1	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm2	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm3	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-lbvm1	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-lbvm2	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-lbvm3	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm1	All microservices, including GUI applications	• 8 CPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm2	All microservices, including GUI applications	• 8 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm3	All microservices, including GUI applications	• 8 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm1	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm2	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.

Table 13: Details of VMs for a Trial Environment with HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-regional-infravm3	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm1	All microservices, including GUI applications	• 8 CPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm2	All microservices, including GUI applications	• 8 CPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm3	All microservices, including GUI applications	• 8 CPUs • 32 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-lbvm1	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-lbvm2	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-lbvm3	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-contrailanalytics-1	Contrail Analytics for centralized and distributed deployments	• 16 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.

Table 13: Details of VMs for a Trial Environment with HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-vrr-vm1	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.
csp-vrr-vm2	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 15 on page 27.

Table 14 on page 24 shows complete details about VMs and microservice collections required for a production environment with HA.

Table 14: Details of VMs for a Production Environment with HA

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-installer-vm	—	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm1	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm2	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-infravm3	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-lbvm1	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-lbvm2	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-lbvm3	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.

Table 14: Details of VMs for a Production Environment with HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-central-msvm1	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm2	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-central-msvm3	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm1	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm2	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-infravm3	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm1	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm2	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-msvm3	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 15 on page 27.
csp-regional-lbvm1	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-lbvm2	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.

Table 14: Details of VMs for a Production Environment with HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-regional-lbvm3	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-elkvm1	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-elkvm2	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-central-elkvm3	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-elkvm1	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-elkvm2	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-elkvm3	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-regional-sblb3	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.

Table 14: Details of VMs for a Production Environment with HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required	Ports to Open
csp-vrr-vm1	Virtual route reflector (VRR)	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-vrr-vm2	Virtual route reflector (VRR)	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 15 on page 27.
csp-contrailanalytics-1	Contrail Analytics server	• 48 vCPUs • 256 GB • 1 TB hard disk storage	See Table 15 on page 27.
csp-contrailanalytics-2	Contrail Analytics server	• 48 vCPUs • 256 GB • 1 TB hard disk storage	See Table 15 on page 27.
csp-contrailanalytics-3	Contrail Analytics server	• 48 vCPUs • 256 GB • 1 TB hard disk storage	See Table 15 on page 27.

Table 15 on page 27 shows the ports that must be open on all VMs in CSO to enable the following types of CSO communications:

- External—CSO user interface (UI) and CPE connectivity
- Internal—Connectivity between CSO components

The provisioning tool opens these ports on each VM; however, if you provision the VMs manually, you must manually open the ports on each VM.

Table 15: Ports to Open on CSO VMs

Port Number	CSO Communication Type	Port Function
22	External and internal	SSH
80	Internal	HAProxy
83	External	Network Service Designer
179	External	BGP for VRR
443	External and internal	HTTPS, including Administration Portal and Customer Portal
514	Internal	Syslog receiving port

Table 15: Ports to Open on CSO VMs (continued)

Port Number	CSO Communication Type	Port Function
1414	Internal	Cassandra Java Virtual Machine (JVM)
1936	External	HAProxy status page
1947	External	Icinga service
2181	Internal	ZooKeeper client
2216	External	CSO telemetry converter
2379	Internal	etcd client communication
2380	Internal	etcd peer
2888	Internal	ZooKeeper follower
3000	External	Grafana
3306	Internal	MySQL
3514	External	Contrail Analytics Syslog receiving port
3888	Internal	ZooKeeper leader
4001	Internal	SkyDNS etcd discover
4443	Internal	HAProxy
4505, 4506	Internal	Salt communications
5000, 5001	Internal	Keystone public
5044	Internal	Beats
5543	Internal	Logstash UDP
5601	External	Kibana UI
5665	Internal	Icinga API
5666	Internal	icinga nrpe
5671	Internal	RabbitMQ SSL listener
5672	Internal	RabbitMQ client
6000	Internal	Swift Object Server

Table 15: Ports to Open on CSO VMs (continued)

Port Number	CSO Communication Type	Port Function
6001	Internal	Swift Container Server
6002	Internal	Swift Account Server
6379	Internal	Redis
6543	Internal	Virtualized Network Function manager (VNFM)
7000	Internal	Kubernetes API server
7804	External	Device connectivity
8006	Internal	Network Service Orchestrator
8016	Internal	Notification engine
8080	Internal	cAdvisor
8081	Internal	Contrail Analytics
8082	Internal	Device Management Service (DMS) central
8083	Internal	Activation Service (AS) central
8085	Internal	DMS Schema
8086	Internal	Contrail Analytics
8090, 8091	Internal	Generic container
8528	Internal	Arango Cluster
8529	Internal	Arango DB
8530	Internal	Arango Cluster
8531	Internal	Arango Cluster
9042	Internal	Cassandra native transport
9090	Internal	Swift Proxy Server
9091	Internal	xmltec-xmlmail tcp
9101	External and internal	HA proxy exporter
9102	Internal	jetdirect

Table 15: Ports to Open on CSO VMs (continued)

Port Number	CSO Communication Type	Port Function
9160	Internal	Cassandra
9200	Internal	Elasticsearch
10000	Internal	Docker repository from CSP installer
10248	Internal	kubelet healthz
10255	Internal	kubelet
15100	Internal	Logstash TCP
15672	Internal	RabbitMQ management
30000-32767	Internal	Kubernetes service node range
30900	External	Prometheus
30901	Internal	Kubernetes
35357	Internal	Keystone private

VNFs Supported

CSO supports the Juniper Networks and third-party VNFs listed in Table 16 on page 30.

Table 16: VNFs Supported by Contrail Service Orchestration

VNF Name	Version	Network Functions Supported	Deployment Model Support	Element Management System Support
Juniper Networks vSRX	vSRX KVM Appliance 15.1X49-D133	• Network Address Translation (NAT) • Demonstration version of Deep Packet Inspection (DPI) • Firewall • Unified threat management (UTM)	• Centralized deployment • Distributed deployment supports NAT, firewall, and UTM.	Element Management System (EMS) microservice, which is included with CSO
LxCIPtable (a free, third party VNF based on Linux IP tables	14.04	• NAT • Firewall	Centralized deployment	EMS microservice
Cisco Cloud Services Router 1000V Series (CSR-1000V)	3.15.0	Firewall	Centralized deployment	Junos Space Network Management Platform

Table 16: VNFs Supported by Contrail Service Orchestration (continued)

VNF Name	Version	Network Functions Supported	Deployment Model Support	Element Management System Support
Riverbed SteelHead	9.2.0	WAN optimization	Distributed deployment, NFX250 platform only	EMS microservice

NOTE: From CSO Release 3.3.0 onward, service chaining with the Silver Peak VX VNF is not supported.

Licensing

You must have licenses to download and use the Juniper Networks CSO. When you order licenses, you receive the information that you need to download and use CSO. If you did not order the licenses, contact your account team or Juniper Networks Customer Care for assistance.

CSO licensing model depends on whether you use a centralized or distributed deployment:

- For a centralized deployment, you need licenses for Network Service Orchestrator and for Contrail Cloud Platform. You can either purchase both types of licenses in one Cloud CPE MANO package or you can purchase each type of license individually.

  You also need licenses for:

  - Junos OS software for the MX Series router, EX Series switch, and QFX Series switch in the Contrail Cloud Platform.
  - VNFs that you deploy.
  - (Optional) Licenses for Junos Space Network Management Platform, if you deploy VNFs that require this EMS.

- For a distributed deployment, you need licenses for Network Service Orchestrator and for Network Service Controller.

  You also need licenses for the following items, depending on which you use in your deployment.

  - The vSRX application that provides the security gateway for the NFX250 device or the vSRX implementation used as a CPE device.
  - VNFs that you deploy.
  - Junos OS software for the MX Series router, including licenses for subscribers.
  - Junos OS software for the SRX Services Gateways.

- For a combined centralized and distributed deployment, you need licenses for components for both types of deployment.

Accessing the CSO GUIs

NOTE: We recommend that you use Google Chrome Version 60 or later to access the CSO GUIs.

Table 17 on page 32 shows the URLs and login credentials for the GUIs for a non redundant CSO installation.

Table 17: Access Details for the GUIs

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

Table 17: Access Details for the GUIs (continued)

GUI	URL	Login Credentials
Kibana This tool provides a visual representation of log files. You can use it to monitor: • Network services in a central or regional POP • Microservices in the deployment	http://*infra-vm-IP-Address* | *ha-proxy-IP-Address*:5601 Where: *infra-vm-IP-Address*—IP address of the VM that hosts the infrastructure services for a central or regional POP. Use this option to monitor network services. *ha-proxy-IP-Address*—IP address of high availability (HA) proxy. Use this option to monitor the microservices. • For a deployment without HA, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: http://192.0.2.2:5601	Login credentials are not needed.
Grafana and Prometheus These tools provide monitoring and troubleshooting for the infrastructure services in CSO. You use Prometheus to create queries for the infrastructure services and Grafana to view the results of the queries in a visual format.	• Prometheus—*ha-proxy-IP-Address*:30900 • Grafana—*ha-proxy-IP-Address*:3000 Where: *ha-proxy-IP-Address*—IP address of HA proxy • For a deployment without HA, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: http://192.0.2.2:30900	For Grafana, specify the username and password. The default username is **admin** and the default password is **admin**. For Prometheus, the login credentials are not needed. After the upgrade, to log in to Administration Portal, you must specify the cspadmin password of the previously installed version.

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Juniper Networks CSO Release 3.3.0.

AWS Spoke

- When an AWS spoke site is being provisioned and the vSRX instance is coming up, all traffic is stopped for 16–30 minutes. After the device is activated and if intent-based policies are configured, the traffic flows as configured.
- The cloud formation template includes a new route table to forward traffic to the vSRX device. If you have configured manual routing between your subnets and VMs, then the new route table replaces the manual routing with only one route forwarding the traffic to the vSRX device.
- In CSO Release 3.3.0, the supported Junos OS release for the AWS spoke is Junos OS Release 15.1X49.D133. You can use the Amazon Machine Image (AMI) number for the respective region from Table 18 on page 34 or obtain the latest AMI for Release 15.1X49.D133 from the AWS Marketplace.:

  1. In the cloud formation template, paste the AMI ID in the Custom Image ID field.

     NOTE: **You must specify the Custom Image ID field because not doing so results in failure during stack creation or provisioning.**

  2. Proceed with the workflow for the cloud formation template in AWS.

Table 18: AMI IDs for Different Regions

Region	AMI ID
US East (N. Virginia)	ami-4a09d335
US East (Ohio)	ami-416a5924
US West (N. California)	ami-8f9182ef
US West (Oregon)	ami-7dc4a705
Canada (Central)	ami-90c342f4
EU (Frankfurt)	ami-f45c071f
EU (Ireland)	ami-d37155aa
EU (London)	ami-5936d63e
Asia Pacific (Singapore)	ami-dd0928a1
Asia Pacific (Sydney)	ami-a2bf77c0
Asia Pacific (Seoul)	ami-faaa0494
Asia Pacific (Tokyo)	ami-48c0dc34

Table 18: AMI IDs for Different Regions (continued)

Region	AMI ID
Asia Pacific (Mumbai)	ami-6f735400
South America (Sao Paulo)	ami-a6b6e7ca
AWS GovCloud (US)	ami-74c85c15

- When you create a cloud spoke site, the default link fields and backup link fields are not applicable.

Policy Deployment

- The deployment of fiirewall policies with UTM profiles fails on sites (devices) on which UTM licenses are not present. Ensure that you install the required licenses before deploying firewall policies that are associated with UTM profiles.

  In addition, when you add new sites or departments, firewall policies that are automatically deployed to the sites might fail if licenses are not installed. In such cases, install the licenses on the applicable sites and re-deploy the failed policy.
- After ZTP of SD-WAN CPE, you must install APBR licenses and app signatures prior to deploying SD-WAN policies through the administrator portal GUI .
- An SD-WAN policy deployment is successful even if there is no matching WAN link meeting the SLA. This is expected behavior and is done so that when a WAN link matching the SLA becomes available, traffic is routed through that link.
- When you are creating an SLA Profile and want to specify the advanced configuration, then you must specify maximum upstream rate, maximum upstream burst size, maximum downstream rate and maximum downstream burst size.

Security Management

- Intrusion prevention system (IPS) is not supported. Therefore, in the IPS report, the attack name from the IPS signatures is displayed as UNKNOWN.
- ISSL Proxy is not supported on SRX300 and SRX320 series devices.
- Firewall rules are pushed to the device depending on the order in which the firewall policy intents are resolved.

Site and Tenant Workflow

- In the Configure Site workflow, use IP addresses instead of hostnames for the NTP server configuration.
- CSO uses hostname-based certificates for device activation. The regional microservices VM hostname must be resolvable from CPE.
- CSO uses RSA key based authentication when establishing an SSH connection to a managed CPE device. The authentication process requires that the device has a

configured root password, and you can use the Administration Portal to specify the root password in the device template.

To specify a root password for the device:

1. Log in to the Administration Portal.
2. Select **Resources>Device Templates**.
3. Select the device template and click **Edit**.
4. Specify the plain text root password in the **ENC_ROOT_PASSWORD** field.
5. Click **Save**.

- When you try to deploy a LAN segment on an SRX Series spoke device, the CSO GUI allows you to select more than one port for a LAN segment. However, for SRX Series devices, only one port for a LAN segment can be deployed; multiple ports in a LAN segment can be deployed only on NFX Series devices.
- Tenant Administrator users cannot delete sites.
- On a site with an NFX Series device, if you deploy a LAN segment without the VLAN ID specified, CSO uses an internal VLAN ID meant for internal operations and this VLAN ID is displayed in the UI. There is no impact on the functionality.
- CSO does not push the default class-of-service configuration on the hub device. You must configure this configuration manually to ensure that the hub configuration is synchronized with the spoke configuration.
- On a cloud hub shared by multiple tenants, by default, CSO does not add a default route and no security policies are configured for the traffic to reach the Internet. You must add the default route and the required security policies for the site traffic to reach the Internet through the cloud hub.

Topology

- Changing the DHCP IP address on the OAM interface is not supported.
- Hybrid-WAN and SD-WAN deployments using the same MX as a hub is not supported.
- When using MX as a SD-WAN hub, NAT configuration must be done on MX Series routers using Stage-2 configuration templates.
- DHCP configuration on WAN links on a SD-WAN hub is not supported.
- Automatic hub-meshing is not supported. Hub-meshing must be performed manually in order for traffic to flow between the hubs.

- The customization of AUX, LAN, and OAM subnets in device profiles (templates) is not supported. Therefore, do not modify the device profiles to change the existing subnets for AUX, LAN, and OAM. Use the default subnets provided in the device profiles.
- You cannot manage an MX Series cloud hub by using an Internet link. Use MPLS as the OAM link to manage the MX Series cloud hub.

User Interface

- When you use Mozilla Firefox to access the CSO GUIs, a few pages do not work as expected. We recommend that you use Google Chrome version 60 or later to access the CSO GUIs.

General

- When you edit a tenant, changing the deployment plan from Hybrid WAN to SD-WAN or vice versa is not supported, although the field is displayed as editable.
- For a centralized deployment, use the following procedure to check that the JSM Heat resource is available in Contrail OpenStack on the Contrail Controller node.

NOTE: **This procedure must be performed on all the Contrail Controller nodes in your CSO installation.**

1. Log in to the Contrail Controller node as root.
2. To check whether the JSM Heat resource is available, execute the **heat resource-type-list | grep JSM** command.

   If the search returns the text **OS::JSM::Get Flavor**, the file is available in Contrail OpenStack.

3. If the file is missing, do the following:

   a. Use Secure Copy Protocol (SCP) to copy the **jsm_contrail_3.py** file to the following directory:

      - For Heat V1 APIs, the **/usr/lib/python2.7/dist-packages/contrail_heat/resources** directory on the Contrail Controller node.
      - For Heat V2 APIs, the **/usr/lib/python2.7/dist-packages/vnc_api/gen/heat/resources** directory on the Contrail Controller node.

      

      NOTE: The jsm_contrail_3.py file is located in the /root/Contrail_Service_Orchestration_3.3/scripts directory on the VM or server on which you installed CSO.

   b. Rename the file to **jsm.py** in the Heat resource directory to which you copied the file.

   c. Restart the Heat services by executing the **service heat-api restart && service heat-api-cfn restart && service heat-engine restart** command.

   d. After the services restart successfully, verify that the JSM Heat resource is available as explained in Step 2. If it is not available, repeat Step 3.

- In vCPE deployments, when a tenant object is created through Administration Portal or the API for a centralized deployment, ContrailOpenStack adds a default security group for the new tenant. This default security group denies inbound traffic and you must manually update the security group in Contrail OpenStack to allow ingress traffic from different networks. Otherwise, Contrail OpenStack might drop traffic.
- In vCPE deployments, CSO does not provide a remote procedure call (RPC) to get the device identifier for a specific site. You can use multiple API calls or the license installation tool to obtain the device identifier for a specific site.
- From CSO Release 3.3.0 onward, service chaining with the Silver Peak VX VNF is not supported.
- After you successfully upgrade from CSO Release 3.2.1 to Contrail Service Orchestration (CSO) Release 3.3.0, ensure that you download the application signatures before installing signatures on the device. This is a one-time operation after the upgrade.
- NFX Series devices that are used to form a cluster in a dual CPE scenario must be configured with the same time zone.

Known Issues

AWS Spoke

- The AWS device activation process takes up to 30 minutes. If the process does not complete in 30 minutes, a timeout might occur and you must retry the process. You do not need to download the cloud formation template again.

  To retry the process:

  1. Log in to Customer Portal.
  2. Access the Activate Device page, enter the activation code, and click **Next**.
  3. After the **CREATE_COMPLETE** message is displayed on the AWS server, click **Next** on the Activate Device page to proceed with device activation.

- For an AWS spoke, during the activation process, the device status on the Activate Device page is displayed as **Detected** even though the device is down.

  Workaround: None.

  Bug Tracking Number: CXU-19779.

CSO HA

- In a CSO HA environment, two RabbitMQ nodes are clustered together, but the third RabbitMQ node does not join the cluster. This might occur just after the initial installation, if a virtual machine reboots, or if a virtual machine is powered off and then powered on.

  Workaround: Do the following:

  1. Log in to the RabbitMQ dashboard for the central microservices VM (http://*central-microservices-vip*:15672) and the regional microservices VM (http://*regional-microservices-vip*:15672).

  2. Check the RabbitMQ overview in the dashboards to see if all the available infrastructure nodes are present in the cluster.

  3. If an infrastructure node is not present in the cluster, do the following:

     a. Log in to the VM of that infrastructure node.

     b. Open a shell prompt and execute the following commands sequentially:

        rabbitmqctl stop_app

        service rabbitmq-server stop

        rabbitmqctl stop_app command

        rm -rf /var/lib/rabbitmq/mnesia/

        service rabbitmq-server start

        rabbitmqctl start_app

  4. In the RabbitMQ dashboards for the central and regional microservices VMs, confirm that all the available infrastructure nodes are present in the cluster.

  Bug Tracking Number: CXU-12107

- In an HA setup, the time configured for the CAN VMs might not be synchronized with the time configured for the other VMs in the setup. This can cause issues in the throughput graphs.

  Workaround:

  1. Log in to can-vm1 as the root user.

  2. Modify the **/etc/ntp.conf** file to point to the desired NTP server.

  3. Restart the NTP process.

After the NTP process restarts successfully, can-vm2 and can-vm3 automatically resynchronize their times with can-vm1.

Bug Tracking Number: CXU-15681.

- In an HA setup, after the central or regional microservices server goes down, policy deployments are stuck in the **In Progress** state.

  Workaround: Contact Juniper Networks Technical Support.

  Bug Tracking Number: CXU-20099.

- In an HA setup, when a failed infrastructure VM recovers, it might not join the ArangoDB cluster.

  Workaround:

  1. Log in to the centralinfravm3 VM.
  2. Execute the **service arangodb3.cluster stop** command.
  3. Log in to the centralinfravm2 VM.
  4. Execute the **service arangodb3.cluster stop** command.
  5. Log in to the centralinfravm1 VM.
  6. Execute the **service arangodb3.cluster stop** command.
  7. On the centralinfravm1 VM, execute the **service arangodb3.cluster start** command and wait for 20 seconds for the command to finish executing.
  8. On the centralinfravm2 VM, execute the **service arangodb3.cluster start** command and wait for 20 seconds for the command to finish executing.
  9. On the centralinfravm3 VM, execute the **service arangodb3.cluster start** command and wait for 20 seconds for the command to finish executing.

  Bug Tracking Number: CXU-20430.

- In some cases, when the power fails, the ArangoDB cluster does not form.

  Workaround: Use the workaround specified for CXU-20340.

  Bug Tracking Number: CXU-20346.

- When an HA setup comes back up after a power outage, MariaDB instances do not come back up on the VMs.

  Workaround:

You can recover the MariaDB instances by executing the **recovery.sh** script (that is packaged with the CSO installation package:

1. Log in to the installer VM.
2. Navigate to the current deployment directory for CSO; for example, **/root/Contrail_Service_Orchestration_3.3/**,
3. Execute the **./recovery.sh** command and follow the instructions.

Bug Tracking Number: CXU-20260,

SD-WAN

- In CSO Release 3.3.0, the LTE link can be only a backup link. Therefore, the SLA metrics are not applicable and default values of zero might be displayed on the Application SLA Performance page, which can be ignored.

  Workaround: None.

  Bug Tracking Number: CXU-19943

- In a dual CPE spoke, non-cacheable applications do not work when the initial path is on CPE0 and the APBR path selected is on CPE1.

  Workaround: None.

  Bug Tracking Number: PR1340331

- In an SRX Series dual CPE site, when the application traffic takes the Z-mode path, the application throughput reported in the Administration Portal GUI is lower than the actual data throughput.

  Workaround: None.

  Bug Tracking Number: PR1347723.

- If all the active links, including OAM connectivity to CSO, are down and the LTE link is used for traffic, and if the DHCP addresses change to a new subnet, the traffic is dropped because CSO is unable to reconfigure the device.

  Workaround: None.

  Bug Tracking Number: CXU-19080.

- On the Site SLA Performance page, applications with different SLA scores are plotted at the same coordinate on the x-axis.

  Workaround: None.

  Bug Tracking Number: CXU-19768.

- When all local breakout links are down, site to Internet traffic fails even though there is an active overlay to the hub.

  Workaround: None.

Bug Tracking Number: CXU-19807

- When the CPE device is not able to reach CSO, DHCP address changes on WAN interfaces might not be detected and reconfigured.

  Workaround: None.

  Bug Tracking Number: CXU-19856

- When the OAM link is down, the communication between the CPE devices and CSO does not work even though CSO can be reached over other WAN links. There is no impact to the traffic.

  Workaround: None.

  Bug Tracking Number: CXU-19881.

- In a full mesh topology, the GRE IPsec down alarms are not created for some overlays during link failure.

  Workaround: None.

  Bug Tracking Number: CXU-20403.

- If you specify an MPLS link without local breakout capability as the backup link, then Internet breakout traffic is dropped because the overlay link to hub will not be used for Internet traffic if local breakout is enabled for the site.

  Workaround: Configure an Internet or an LTE link as the backup link.

  Bug Tracking Number: CXU-20447.

- If you define an SLA profile for a static SD-WAN policy but do not remove the default values for the SLA parameters and deploy the policy, the policy is deployed as a dyanmic SD-WAN policy.

  Workaround: When you define the SLA profile for a static SD-WAN policy, ensure that you remove the default values for the SLA parameters.

  Bug Tracking Number: CXU-20499.

- If you modify the path preference of an existing SLA profile that has already been deployed and redeploy the SD-WAN policy, the path of the SLA profile is not updated on the CPE device.

  Workaround: Modify the path preference in an SLA profile that is not yet deployed.

  Bug Tracking Number: CXU-20540.

- For non-cacheable applications, in a hub- and- spoke topology, on link switchover, in some cases, the traffic between the hub and spoke might take an incorrect physical path because the existing session flow is not updated. However, there is no traffic loss.

  Workaround: None.

  Bug Tracking Number: PR1341274

- In the bandwidth-optimized SD-WAN mode, when the same SLA is used in the SD-WAN policy for different departments and an SLA violation occurs, two link switch events

that appear identical, because the department name is missing from the event details, are displayed.

Workaround: None.

Bug Tracking Number: CXU-20529.

- When you configure a high delay and loss on the OAM link, the link switch might be delayed or might not occur.

  Workaround: None.

  Bug Tracking Number: CXU-20562.

- For a tenant with bandwidth-optimized SD-WAN mode, the SLA performance for the site is always displayed as 0/100.

  Workaround: None.

  Bug Tracking Number: CXU-20563.

Security Management

- If you create firewall policy with more than 10 firewall policy intents and deploy the firewall policy on a tenant with 45 or more sites, the policy deployment fails.

  Workaround: None.

  Bug Tracking Number: CXU-20292

- If you create a NAT pool, specify the **Translation** as **Port/Range**, configure the port as a range, and enter an incorrect starting port number, then you cannot enter the ending port number and the NAT pool is created with a single port value instead of a range.

  Workaround: When you create a NAT pool with a port range, ensure that the starting port number is between 1024 and 65,335, and then enter the corresponding ending port number between 1024 and 65,335.

  Bug Tracking Number: CXU-20366.

- On the Active Database page in Customer Portal, the wrong installed device count is displayed. The count displayed is for all tenants and not for a specific tenant.

  Workaround: None.

  Bug Tracking Number: CXU-20531.

Site and Tenant Workflow

- The tenant delete operation fails when CSO is installed with an external Keystone.

  Workaround: You must manually delete the tenant from the Contrail OpenStack user interface.

  Bug Tracking Number: CXU-9070

- When both the OAM and data interfaces are untagged, ZTP fails when using a NFX Series platform as CPE.

Workaround: Use tagged interfaces for both OAM and data.

Bug Tracking Number: CXU-15084

- The tenant creation job might fail if connectivity from CSO to the VRR is lost during job execution.

  Workaround: If the tenant creation job fails and the tenant is created in CSO, delete the tenant and retrigger the tenant creation.

  Bug Tracking Number: CXU-16884

- If the tenant name exceeds 16 characters, the activation of the SRX Series hub device fails.

  Workaround: Delete the tenant and re-create a new tenant with name that has less than 16 characters and retry the activation workflow.

  Bug Tracking Number: PR1344369.

- For tenants with a large number of spoke sites, the tenant deletion job fails because of token expiry.

  Workaround: Retry the tenant delete operation.

  Bug Tracking Number: CXU-19990.

- In some cases, on the Monitor Overview page (**Monitoring > Overview**) for a site, the ZTP status is displayed incorrectly when you hover over the site.

  Workaround: None.

  Bug Tracking Number: CXU-20226.

- In some cases, if automatic license installation is enabled in the device profile, after ZTP is complete, the license might not be installed on the CPE device even though license key is configured successfully.

  Workaround: Reinstall the license on the CPE device by using the Licenses page on the Administration Portal.

  Bug Tracking Number: PR1350302.

- In the scenario where the redirect service from Juniper (redirect.juniper.net) is not being used, after you upgrade an NFX device to Junos OS Release 15.1X53-D472, the device is unable to connect to the regional server because the phone home server certificate (**phd-ca.crt**) is reverted to the factory default.

  Workaround: Manually copy the regional certificate to the NFX device.

  Bug Tracking Number: PR1350492.

- LAN segments with overlapping IP prefixes are not supported across tenants or sites.

  Workaround: Create LAN segments with unique IP prefixes across tenants and sites.

  Bug Tracking Number: CXU-20347.

- In a hub and spoke topology with multi-tenancy (network segmentation) enabled, the reverse traffic from the hub to the originating spoke might not take the same path as the traffic in the forward direction. There is no traffic loss.

Workaround: None.

Bug Tracking Number: CXU-20494.

- In the Configure Site workflow for a full mesh topology with multitenancy enabled, the option to connect the CPEs only to the hub is not supported; that is, if you specify **false** for the **used_for_meshing** parameter, this option is ignored.

  Workaround: None.

  Bug Tracking Number: CXU-20495.

- For hybrid WAN tenants, during site creation, all the VIMs in the system are displayed even though a specific VIM is already assigned during the tenant creation.

  Workaround: None.

  Bug Tracking Number: CXU-20371.

- When you use DHCP for the activation of a dual CPE device, ZTP might fail because the device takes longer than expected to to connect to the Device Connectivity Service (DCS).

  Workaround: Retry the failed ZTP job.

  Bug Tracking Number: CXU-20467.

- During site addition, if you create a department but do not assign a LAN segment to that department, during the site activation, the firewall policy deployment fails.

  Workaround: Do one of the following:

  - Go to the *Site-Name* page, and on the LAN tab, add a new LAN segment to the department that did not have any LAN segments assigned during site creation.
  - Alternatively, during site addition, when you create a department, ensure that you assign at least one LAN segment to that department.

  Bug Tracking Number: CXU-20502.

- When the primary and backup interfaces of the CPE device uses the same WAN interface of the hub, the backup underlay might be used for Internet or site-to-site traffic even though the primary links are available.

  Workaround: Ensure that you connect the WAN links of each CPE device to unique WAN links of the hub.

  Bug Tracking Number: CXU-20564.

Topology

- When a spoke is recalled, the configuration remains on the hub. When the spoke is reprovisioned, the activation fails and an error message indicating that the source and destination addresses of the tunnel cannot be the same is displayed in the logs.

  Workaround: Clean up the configuration of the recalled spoke in the hub and reprovision the spoke with a new name.

  Bug Tracking Number: CXU-20441.

User Interface

- When you bring down or bring up an AWS availability zone, there might be a momentary slowdown in the response time of the Administration Portal GUI and some in-progress jobs might be affected.

  Workaround: Wait for five minutes and retry the failed jobs.

  Bug Tracking Number: CXU-20463.

General

- If you create VNF instances in the Contrail cloud by using Heat Version 2.0 APIs, a timeout error occurs after 120 instances are created.

  Workaround: Contact Juniper Networks Technical Support.

  Bug Tracking Number: CXU-15033

- When you upgrade the gateway router by using the CSO GUI, after the upgrade completes and the gateway router reboots, the gateway router configuration reverts to the base configuration and loses the IPsec configuration added during Zero Touch Provisioning (ZTP).

  Workaround: Before you upgrade the gateway router by using the CSO GUI, ensure that you do the following:

  1. Log in to the Juniper Device Manager (JDM) CLI of the NFX Series device.
  2. Execute the **virsh list** command to obtain the name of the gateway router (*GWR_NAME*).
  3. Execute the **request virtual-network-functions *GWR_NAME* restart** command, where *GWR_NAME* is the name of the gateway router obtained in the preceding step.
  4. Wait a few minutes for the gateway router to come back up.
  5. Log out of the JDM CLI.
  6. Proceed with the upgrade of the gateway router by using the CSO GUI.

Bug Tracking Number: CXU-11823.

- CSO may not come up after a power failure.

  Workaround:

  1. Log in to the installer VM.
  2. Navigate to the **/root/Contrail_Service_Orchestration_3.3/** directory.
  3. Run the **reinitialize_pods.py** script as follows:

     ./python.sh recovery/components/reinitialize_pods.py
  4. SSH to the VRR by using the VRR IP address to check if you are able to access the VRR.

     If there is an error in connecting (**port 22: Connection refused**), then you must recover the VRR by following step 5 through 21.
  5. Log in to physical server hosting the VRR.
  6. Execute the **virsh destroy vrr** command to destroy the VRR.

     

     WARNING: **Do not execute the virsh undefine vrr command because doing so will cause the VRR configuration to be lost and the configuration cannot be recovered.**
  7. Delete the VRR image that is located in the **/root/ubuntu_vm/vrr/vrr-15.1R6.7.qcow2** directory.
  8. Copy the fresh VRR image from the **/root/disks/vrr-15.1R6.7.qcow2** directory to the **/root/ubuntu_vm/vrr/vrr-15.1R6.7.qcow2** directory.
  9. Execute the **virsh start vrr** command and wait for approximately 5 minutes for the command to finish executing.
  10. Execute the **virsh list –all** command to check if the VRR is running or not.

      If the VRR is not running, check that the image that was copied was the uncorrupted image and re-try the steps from step 7.
  11. If the VRR is running, navigate to the **/root/ubuntu_vm/vrr/** directory.
  12. Run the **./vrr.exp** command to push the base configuration to the VRR.

13. Check if the VRR is reachable from the regional microservices VM. If the VRR is reachable, proceed to step 14. If the VRR is not reachable:

    a. Log in to the VRR.

    b. Check if the base configuration was pushed properly:

    - If the base configuration was pushed properly, re-check if the VRR is reachable from the regional microservices VM. If the VRR is reachable, proceed to step 14.
    - If the base configuration was not pushed properly:

      i. Add the necessary routes to reach CSO.

      ii. Re-check if the VRR is reachable from the regional microservices VM. If the VRR is reachable, proceed to 14.

14. Import the POP by using the URL https://*central-ms-ip*:443/tssm/import-pop, where *central-ms-ip* is the IP address of the central microservices VM.

15. Use POSTMAN to import the VRR.

NOTE: *Do not* import the VRR until the VRR is reachable from the regional microservices VM.

The following is the JSON format for the VRR. (In the JSON below, <*vrr-ip-address*> is the IP address of the VRR and <*vrr-password*> is the password that was configured for the VRR.

```
{               "input": {                                "job_name_prefix":
 "ImportPop",                                    "pop": [{
                          "dc_name": "regional",
                     "device": [{
                      "name": "vrr-<vrr-ip-address>",
                                             "family": "VRR",
                                                 "device_ip":
"<vrr-ip-address>",
        "assigned_device_profile": "VRR_Advanced_SDWAN_option_1",
                                                          "authentication": {

   "password_based": {
                                             "username": "root",

        "password": "<vrr-password>"
                                          }
                                },
                     "management_state": "managed",
                                          "pnf_package": "null"
                                    }],
          "name": "regional"                                  }]
    }}
```

16. Verify whether the VRR is imported properly:
    a. Log in to the CSO Administration Portal.
    b. Click **Resources > POPs > Import POPs > Import History** and confirm that the **ImportPop** job is running and that it has completed successfully.
17. On the Tenants page, add a tenant named **recovery**.
18. After the tenant is successfully created, log in to the VRR and access the Junos OS CLI.
19. Execute the **show configuration|display set** and verify that the tenant configuration (for the previously-configured tenants) is recovered.
20. Execute the **show bgp summary** and check that the BGP status to the hub and spokes are **Established**.
21. If the status is **Not Established**, add the routes for the OAM traffic of the hub and spokes to the VRR and recheck the status.

Bug Tracking Number: CXU-16530

- If you run the script to revert the upgraded setup to CSO Release 3.2.1, in some cases, the ArangoDB cluster becomes unhealthy.

  Workaround:

  1. Log in to the centralinfravm3 VM.
  2. Execute the **service arangodb3 stop** command and wait for 30 seconds.
     - If the command executes successfully, proceed to Step 3.
     - If there is no progress after 30 seconds:
       a. Press Ctrl+c to abort the command.
       b. Execute the **kill -9 `ps -ef | grep arangod | grep -v grep | awk {'print $2'}`** command.
  3. Log in to the centralinfravm2 VM.
  4. Execute the **service arangodb3 stop** command and wait for 30 seconds.
     - If the command executes successfully, proceed to Step 5.
     - If there is no progress after 30 seconds:
       a. Press Ctrl+c to abort the command.

b. Execute the **kill -9 `ps -ef | grep arangod | grep -v grep | awk {'print $2'}`** command.

5. Log in to the centralinfravm1 VM.

6. Execute the **service arangodb3 stop** command and wait for 30 seconds.

   - If the command executes successfully, proceed to Step 7.
   - If there is no progress after 30 seconds:

     a. Press Ctrl+c to abort the command.

     b. Execute the **kill -9 `ps -ef | grep arangod | grep -v grep | awk {'print $2'}`** command.

7. On the centralinfravm3 VM, execute the **service arangodb3 stop** command and wait for 20 seconds for the command to finish executing.

8. On the centralinfravm2 VM, execute the **service arangodb3 stop** command and wait for 20 seconds for the command to finish executing.

9. On the centralinfravm1 VM, execute the **service arangodb3 stop** command and wait for 20 seconds for the command to finish executing.

10. Execute the **netstat -tuplen | grep arangod** command on all *three* central infrastructure VMs to check the status of the ArangoDB cluster. If the port binding is successful for all the central infrastructure VMs, then the ArangoDB cluster is healthy.

    The following is a sample output.

    ```
    tcp6 0 0 :::8528 :::* LISTEN 0 54213 9220/arangodb
    tcp6 0 0 :::8529 :::* LISTEN 0 44018 9327/arangod
    tcp6 0 0 :::8530 :::* LISTEN 0 91216 9289/arangod
    tcp6 0 0 :::8531 :::* LISTEN 0 42530 9232/arangod
    ```

    Bug Tracking Number: CXU-20397.

- On a CPE device configured with an LTE backup link, LTE link flaps are observed when the CPE device is running for a longer period.

  Workaround: None.

  Bug Tracking Number: PR1349613.

- In an HA setup, when you upgrade from JCS 3.2.1 to JCS 3.3.0, the Kubernetes system pods for the central or regional load balancer VM are in the Terminating state. This causes the load balancer VM to be in the Not Ready state, which causes the health check to fail during the upgrade.

Workaround:

1. On the installer VM:
   - If the central load balancer VM is in the Not Ready state, execute the **salt 'csp-central-lbvm*' cmd.run 'reboot'** command.
   - If the regional load balancer VM is in the Not Ready state, execute the **salt 'csp-regional-lbvm*' cmd.run 'reboot'** command.
2. Wait for some time until the nodes are in the Ready state.
3. Rerun the **upgrade.sh** script to continue with the upgrade.

Bug Tracking Number: CXU-20271.

- The provisioning of CPE devices fails if all VRRs within a redundancy group are unavailable.

  Workaround: Recover the VRR that is down and retry the provisioning job.

  Bug Tracking Number: CXU-19063

- In the centralized deployment, after you import a POP, the CPU, memory, and storage allocation are displayed as zero.

  Workaround: Refresh the UI, and the correct information is displayed.

  Bug Tracking Number: CXU-19105

- The CSO health check displays the following error message: **ERROR: ONE OR MORE KUBE-SYSTEM PODS ARE NOT RUNNING**

  Workaround:

  1. Log in to the central microservices VM.
  2. Execute the **kubectl get pods –namespace=kube-system** command.
  3. If the kube-proxy process is not in the Running state, execute the **kubectl apply –f /etc/kubernetes/manifests/kube-proxy.yaml** command.

     Bug Tracking Number: CXU-20275.

- In a department, if there are two LAN segments with DHCP enabled, only one DHCP server setting is deployed on the device.

  Workaround: Enable DHCP only for one LAN segment in a department.

  Bug Tracking Number: CXU-20519.

- The Grant RMA operation fails for a multihomed hub device.

  Workaround: None.

  Bug Tracking Number: CXU-20457.

- After the upgrade, the health check on the standalone Contrail Analytics Node (CAN) fails.

  Workaround:

  1. Log in to the CAN VM.
  2. Execute the **docker exec analyticsdb service contrail-database-nodemgr restart** command.
  3. Execute the **docker exec analyticsdb service cassandra restart** command.

  Bug Tracking Number: CXU-20470.

- When the LTE modem is disconnected or disabled in the NFX250 CPE device, an alarm is triggered. However, the underlay link status on the Sites page might not display the alarm.

  Workaround: None.

  Bug Tracking Number: CXU-20492.

- For a vSRX CPE, the auto-deployment of license fails with an error message indicating that no license is found even though a license exists on the vSRX instance.

  Workaround: Manually deploy the license by using the Push License workflow from the CSO GUI.

  Bug Tracking Number: CXU-20558.

- The load services data operation or health check of the infrastructure components might fail if the data in the Salt server cache is lost because of an error.

  Workaround: If you encounter a Salt server-related error, do the following:

  1. Log in to the installer VM.
  2. Execute the **salt '*' deployutils.get_role_ips 'cassandra'** command to confirm whether one or more Salt minions have lost the cache.
     - If the output returns the IP address for all the Salt minions, this means that the Salt server cache is fine; proceed to step 7.
     - If the IP address for some minions is not present in the output, this means that the Salt server has lost its cache for those minions and must be rebuilt as explained from step 3.
  3. Navigate to the current deployment directory for CSO; for example, **/root/Contrail_Service_Orchestration_3.3/**.
  4. Redeploy the central infrastructure services (up to the NTP step):
     a. Execute the **DEPLOYMENT_ENV=central ./deploy_infra_services.sh** command.

b. Press Ctrl+c when you see the following message on the console:

```
2018-04-10 17:17:03 INFO utils.core Deploying roles set(['ntp']) to servers
 ['csp-central-msvm', 'csp-contrailanalytics-1', 'csp-central-k8mastervm',
 'csp-central-infravm']
```

5. Redeploy the regional infrastructure services (up to the NTP step):

   a. Execute the **DEPLOYMENT_ENV=regional ./deploy_infra_services.sh** command.

   b. Press Ctrl+c when you see a message similar to the one for the central infrastructure services.

6. Execute the **salt '*' deployutils.get_role_ips 'cassandra'** command and confirm that the output displays the IP addresses of all the Salt minions.

7. Re-run the load services data operation or the health component check that had previously failed.

Bug Tracking Number: CXU-20815.

Resolved Issues

The following issues are resolved in Juniper Networks CSO Release 3.3.0.

- You can use the Administration Portal to upload licenses to CSO; however, you cannot use the Administration Portal to install licenses on physical or virtual devices that CSO manages. You must use the APIs or the license installation tool to install licenses on devices.
- If sites are removed without first undeploying the associated policies, the removal of SLA profiles fails.

  Bug Tracking Number: CXU-13179
- ZTP fails on SRX 3xx Series device CPE because DHCP bindings already exist on CPE.

  Bug Tracking Number: CXU-13446
- If you create a firewall policy and deploy it to the device, and subsequently create one or more firewall policy intents without redeploying the policy, the firewall policy is automatically deployed to the device when there is a change in the topology, such as the addition of a new site, department, or LAN segment.

  Bug Tracking Number: CXU-15794
- In rare cases, site routes are not advertised to the hub, which results in the sites not being reachable.

  Bug Tracking Number: CXU-16145.
- Sorting by **Administrator** on the Tenants page displays an error message.

Bug Tracking Number: CXU-16642

- Automatic Policy deployments on new site addition (for example, auto NAT, firewall, SD-WAN) can sometimes fail because trusted certificates might be installed on the device in parallel.

  Bug Tracking Number: CXU-16652

- When configuring an SRX Series spoke device in a multihoming topology with a cloud hub and enterprise hub, Administration Portal displays a **Primary Hub and Secondary Hub must belong to a same Device Family** error message.

  Bug Tracking Number: CXU-16662

- The preferred link for the underlay is not displayed in the GUI.

  Bug Tracking Number: CXU-16785.

- When the MX Series device is used as a hub, site-to-site traffic in the reverse direction in the hub-and-spoke topology might not take the desired path from the hub to the originating spoke. However, there is no traffic loss.

  Bug Tracking Number: CXU-15970.

- When you perform the microservices VM failure tests and the Kubernetes nodes go to the **not ready** state, some Docker pods might not come up in the **running** state.

  Bug Tracking Number: CXU-16541.

- Dynamic SLA is not supported for full mesh topology; only static policies are supported.

  Bug Tracking Number: CXU-17087.

- Site names across tenants must be unique, that is, you cannot use the same site name across tenants.

  Bug Tracking Number: CXU-17483.

- When you create an antivirus profile, the engine types Kaspersky and Juniper Express are not supported.

  Bug Tracking Number: CXU-17901.

- ZTP for SRX Series devices does not work with a redirect server because a BOOTSTRAP complete message is not received when ZTP is initiated through a redirect server.

  Bug Tracking Number: CXU-14099

Documentation Updates

There are no errata or changes in CSO Release 3.3.0 documentation:

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using 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:

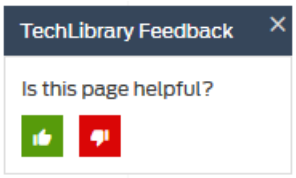

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- 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 J-Care or Partner Support Service 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 http://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/.

Revision History

28 January 2019—Revision 9

24 Oct 2018—Revision 8

30 Jul 2018—Revision 7

23 May 2018—Revision 6

17 May 2018—Revision 5

19 April 2018—Revision 4

13 April 2018—Revision 3

11 April 2018—Revision 2

31 March 2018—Revision 1, CSO Release 3.3.0