# Cloud CPE Solution Release Notes

**Release 3.1.2**
**04 January 2018**
**Revision 4**

These Release Notes accompany Release 3.1.2 of the Juniper Networks® Cloud CPE Solution. They contain installation information, and they describe new and changed features, limitations, and known and resolved issues in the software.

**Contents**

## Introduction

The Juniper Networks Cloud Customer Premises Equipment (CPE) solution 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 the network services.

Cloud CPE Solution Release 3.1.2 is a secure software-defined WAN (SD-WAN) solution that builds on the capabilities Cloud CPE Solution Release 3.1. The following are the highlights of the features available in Release 3.1:

- SD-WAN
  - Centralized application, service-level agreement (SLA), and performance management
  - Intent-based advanced policy-based routing (APBR)
  - Traffic visualization and monitoring at a per-application level across branch sites
- Security management
  - Intent-based firewall policies
  - Network Address Translation (NAT) policy management
  - Application visibility and signature management
  - Security reports

The solution 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.

The Cloud CPE solution 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 NFVI.

The following Contrail Service Orchestration (CSO) components connect to Network Service Orchestrator through its RESTful API:

NOTE: **From Cloud CPE Solution Release 3.1 onward, 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 graphical user interface (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 the Cloud CPE solution, such as the status of virtualized network functions (VNFs), virtual machines (VMs), and physical servers; information about physical servers' resources; components of a network service (VNFs and VMs hosting a VNF); counters and other information for VNFs.

  You can deploy the Cloud CPE solution in a demonstration (demo) 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 Sites and VNFs Supported for a Distributed Solution | Number of VNFs Supported for a Centralized Deployment |
|---|---|---|
| Demo non-HA Configuration | 25 sites, 2 VNFs per site | Up to 10 VNFs |
| Production non-HA Configuration | Up to 200 sites, 2 VNFs per site | Up to 500 VNFs, 20 VNFs per Contrail compute node |
| Trial HA Configuration | Up to 200 sites, 2 VNFs per site | Up to 100 VNFs, 20 VNFs per Contrail compute node |
| Production HA Configuration | Up to 2200 sites, 2 VNFs per site | Up to 500 VNFs, 20 VNFs per Contrail compute node |

## Installation



### Installation Notes

You must use the CSO installer for a production environment or a demo environment. After copying the installer TAR file to the CSO server and expanding the file, you provision the virtual machines (VMs), and run a script to create a file of settings for the installation. You then run the installer, which takes approximately one hour to install CSO. Finally, you start the CSO infrastructure services and microservices.

NOTE: If you use the **provision_vm.sh** script to spawn the VMs on the physical servers (host OS), then the physical servers must be connected to the Internet to download certain software packages. After the VMs are spawned, you can proceed with the installation without Internet access.

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

NOTE: If the information in the README file differs from the information in the technical documentation (Deployment Guide or Release Notes), follow the information in the technical documentation.

## Installation Instructions for the NSC Package

If you are installing the Network Service Controller (NSC) package on the installer VM or on the physical server (acting as an installer), after you expand (untar) the installation package, do the following:

1. Access the installer directory by executing the **cd *installer-dir*** command, where *installer-dir* is the name of the installer directory.

   The installer directory has the same name as the installer package. For example, if the installation package name is **nscVersion.tar.gz**, the installer directory is **nscVersion**.

2. Execute the following command to specify certain text variables in the installation files:

```
sed -i '''s@download_url: http.*/@download_url:
http://csp-installer:90/csp_components/@g''' confs/roles_csp.conf;sed
-i '''s@build_server_url_pattern: http.*/@build_server_url_pattern:
http://csp-installer:90/csp_components/@g''' confs/roles_csp.conf;sed
-i '''s@download_url: "http.*/@download_url:
"http://csp-installer:90/csp_components/@g''' confs/roles_csp.conf;sed
-i '''s@build_server_url_pattern: "http.*/@build_server_url_pattern:
"http://csp-installer:90/csp_components/@g''' confs/roles_csp.conf;
```

NOTE: When you copy the command, ensure that you copy it as a single command and remove all line spaces.

After executing the command, you can proceed with the rest of the installation.

## Configuring Name Servers on CPE Devices

To configure the name server on a CPE device, you must use the custom properties to provide the name server details when you are adding a tenant.

## Software Installation Requirements for NFX250 Network Services Platform

The NFX250 requires the Junos OS Release 15.1X53-D47 for the Cloud CPE Solution Release 3.1.2.

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

1. Upload the image to Contrail Service Orchestration.

2. Specify this image as the boot image when you configure activation data.

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

### Software Downloads for SRX Series Devices

The Contrail Service Orchestration (CSO) software package does not contain the images for the SRX300 Series, SRX1500, SRX4100, and SRX4200 devices. You can download these images by using the following links:

- SRX1500: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69369.html
- SRX1500 USB: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69372.html
- SRX1500 Preboot Execution Environment (PXE): https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69374.html
- SRX300 Series: https://www.juniper.net/support/downloads/?p=srx300#sw
- SRX4100 and SRX4200: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69370.html
- SRX4100 and SRX4200 USB: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69373.html
- SRX4100 and SRX4200 PXE: https://webdownload.juniper.net/swdl/dl/secure/site/1/record/69375.html

### Supported Software

For information about the software versions qualified for CSO Release 3.1.2:

- For the centralized deployment, see Table 5 on page 9.
- For the distributed deployment, see Table 7 on page 10.

## New and Changed Features

This section describes the new features or enhancements to existing features in Cloud CPE Solution Release 3.1.2. For new and changed features in Cloud CPE Solution Release 3.1, refer to the *Cloud CPE Solution 3.1 Release Notes* available at https://www.juniper.net/documentation/en_US/nfv3.1/information-products/pathway-pages/3.1/index.html .

- **High availability (HA) support for Kubernetes master node**—From Cloud CPE Solution Release 3.1.2 onward, HA for the Kubernetes master node is supported.

## Servers, Software, and Network Devices Tested

- Node Servers and Servers Tested in the Cloud CPE Solution
- 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 the Cloud CPE Solution

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

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

Table 2 on page 8 lists the node servers and servers that have been tested for these functions in the Cloud CPE solution. You should use these specific node servers or servers for the Cloud CPE solution.

### Table 2: COTS Node Servers and Servers Tested in the Cloud CPE Solution

| 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 3 on page 8 shows the software that has been tested for the Cloud CPE solution. You must use these specific versions of the software when you implement the Cloud CPE solution.

### Table 3: Software Tested for the COTS Nodes and Servers

| Description | Version |
|---|---|
| Operating system for all COTS nodes and servers | Ubuntu 14.04.5 LTS |
| Operating system for VMs on Contrail Service Orchestration servers | Ubuntu 14.04.5 LTS |
| Hypervisor on Contrail Service Orchestration servers | KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0 |
| Additional software for Contrail Service Orchestration servers | Secure File Transfer Protocol (SFTP) |
| Software defined networking (SDN) for a centralized deployment | Contrail Release 3.2.5 with OpenStack Mitaka and Heat v2 APIs |
| Contrail Analytics | Contrail Release 4.0.2.0-35 |

## 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 4 on page 9.
- The software described in Table 5 on page 9.

  You must use these specific versions of the software for the Cloud CPE Solution Release 3.1.2.

### Table 4: Network Devices Tested for the Contrail Cloud Platform

| 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:<br>• 48 10/100/1000-Gigabit Ethernet interfaces<br>• 4 built-in 10-Gigabit Ethernet SFP transceiver interfaces | 1 |
| Data switch | Juniper Networks QFX Series Switch | QFX 5100-48S-AFI switch with:<br>• 48 SFP+ transceiver interfaces<br>• 6 QSFP+ transceiver interfaces | 1 |

### Table 5: 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 17.4R1 |
| Hypervisor on Contrail Service Orchestration servers | KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0 |
| Element management system software | EMS microservice<br>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) for a centralized deployment | Contrail Release 3.2.5 with OpenStack Mitaka |
| Contrail Analytics | Contrail Release 4.0.2.0-35 |
| Virtualized infrastructure manager (VIM) and virtual machine (VM) orchestration | OpenStack Mitaka |

Table 5: Software Tested in the Centralized Deployment *(continued)*

| Function | Software and Version |
|---|---|
| Authentication and authorization | OpenStack Mitaka |
| Network Functions Virtualization (NFV) | Contrail Service Orchestration Release 3.1.2 |

## 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 6 on page 10.
- The software described in Table 7 on page 10.

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

Table 6: Network Devices Tested for the Distributed Deployment and SD-WAN Implementation

| Function | Device | Model | Quantity |
|---|---|---|---|
| PE router and IPsec concentrator (Hybrid WAN distributed deployment only) | Juniper Networks MX Series 3D Universal Edge Router | • MX960, MX480, or MX240 router with MS-MPC line card<br>• MX80 or MX104 router with MX-MIC line card<br>• Other MX Series routers with an MS-MPC or MX-MIC are supported | — |
| Hub device (SD-WAN implementation only) | Juniper Networks SRX Series Services Gateway | • SRX1500 Services Gateway<br>• SRX4100 Services Gateway<br>• SRX4200 Services Gateway | — |
| CPE device | • NFX250 Series Network Services Platform<br>• SRX Series Services Gateway<br>• vSRX on an x86 server | • NFX250-LS1 device<br>• NFX250-S1 device<br>• NFX250-S2 device<br>• SRX300 Services Gateway<br>• SRX320 Services Gateway<br>• SRX340 Services Gateway<br>• SRX345 Services Gateway<br>• vSRX | 1 per customer site |

Table 7: Software Tested in the Distributed Deployment

| Function | Software and Version |
|---|---|
| Hypervisor on Contrail Service Orchestration 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) | Contrail Service Orchestration Release 3.1.2 |

Table 7: Software Tested in the Distributed Deployment *(continued)*

| Function | Software and Version |
|---|---|
| Contrail Analytics | Contrail Release 4.0.2.0-35 |
| NFX software | Junos OS Release 15.1X53-D47 |
| Routing and security for NFX250 device | vSRX KVM Appliance 15.1X49-D123 |
| Operating system for vSRX used as a CPE device on an x86 server | vSRX KVM Appliance 15.1X49-D123 |
| Operating system for an SRX Series services gateway used as a CPE device | Junos OS Release 15.1X49-D123 |
| Operating system for an MX Series router used as a provider edge (PE) router | Junos OS Release 16.1R3.00 |
| Operating system for an SRX Series services gateway used as a hub device for an SD-WAN implementation | Junos OS Release 15.1X49-D123 |

## Hardware, Software, and Virtual Machine Requirements for the Cloud CPE Solution



### Minimum Hardware Requirements for the Cloud CPE Solution

Table 2 on page 8 lists the makes and models of node servers and servers that you can use in the Cloud CPE solution. When you obtain node servers and servers for the Cloud CPE Solution, 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 demo or a production environment.

Table 8 on page 11 shows the required hardware specifications for node servers and servers in a demo environment.

Table 8: Server Requirements for a Demo Environment and a Trial HA Environment

| Function | Demo Environment (no HA) | Trial HA Environment |
|---|---|---|
| *Node or Server Specification* | | |

Table 8: Server Requirements for a Demo Environment and a Trial HA Environment *(continued)*

| Function | Demo Environment (no HA) | Trial HA Environment |
|---|---|---|
| Storage | Greater than 1 TB of one of the following types:<br>• Serial Advanced Technology Attachment (SATA)<br>• Serial Attached SCSI (SAS)<br>• Solid-state drive (SSD) | Greater than 1 TB of one of the following types:<br>• SATA<br>• SAS<br>• SSD |
| CPU | One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.4 Ghz or higher specification | One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.4 Ghz or higher specification |
| Network interface | One 1-Gigabit Ethernet or 10-Gigabit Ethernet interface | One 1-Gigabit Ethernet or 10-Gigabit Ethernet interface |
| *Contrail Service Orchestration Servers (includes Contrail Analytics in a VM )* | | |
| Number of nodes or servers | 1<br>NOTE: If you want to use Junos Space to support virtualized network functions (VNFs) that require this element management system (EMS) in your demo environment, you must install Junos Space in a VM on another server. This server specification for a demo environment does not accommodate Junos Space. See *Details of VMs for a Demo Environment* for information on Junos Space VM requirements. | 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<br>• 3 nodes for Contrail Controller and Analytics<br>• 1–4 Contrail compute nodes |
| vCPUs per node or server | 48 | 48 |
| RAM per node or server | 256 GB | 256 GB |

Table 9 on page 13 shows the required hardware specifications for node servers and servers in a production environment.

## Table 9: Server Requirements for a Production Environment (HA and non-HA)

| Server Function | Values |
|---|---|
| *Node or Server Specification* | |
| Storage | Greater than 1 TB of one of the following types:<br>• SATA<br>• SAS<br>• SSD |
| CPU | One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.4 Ghz or higher specification |
| Network interface | One 1-Gigabit Ethernet or 10-Gigabit Ethernet interface |
| *Contrail Service Orchestration Servers* | |
| Number of nodes or servers for a non-HA environment | 3<br>• 1 central server<br>• 1 regional server<br>• 1 Contrail Analytics server |
| Number of nodes or servers for an HA environment | 9<br>• 3 central servers<br>• 3 regional servers<br>• 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<br>• 3 nodes for Contrail Controller and Contrail Analytics<br>• 1–25 Contrail compute nodes |
| 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 the Cloud CPE solution. This section shows the VMs required for each type of environment.

For non-HA demo deployments, two new VMs are added, one each for central and regional servers. The VMs separate out kubemaster from the kubeminions. The two VMs require 4vCPUs and 8GB RAM each. This is done by oversubscribing the vCPUs. The hypervisor oversubscribes the vCPUs on the servers as 52 vCPUs are utilized now against the available 48 vCPUs on the CCRA specification servers.

Table 10 on page 14 shows complete details about the VMs that need to be deployed for a demo environment. HA is not included with the demo environment.

### Table 10: Details of VMs for a Non-HA Demo Environment

| Components | Name of VM | Components That Installer Places in VM | Resources Required | Ports to Open |
|---|---|---|---|---|
| Installer VM | csp-installer-vm | — | • 4 vCPUs<br>• 16 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| Central Infrastructure services | csp-central-infravm | Third-party applications used as infrastructure services | • 6 CPU<br>• 32 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |
| Central Microservices | csp-central-msvm | All microservices, including GUI applications | • 6 vCPUs<br>• 40 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |
| Regional Infrastructure services | csp-regional-infravm | Third-party applications used as infrastructure services | • 6 vCPUs<br>• 32 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |
| Regional Microservices | csp-regional-msvm | All microservices, including GUI applications | • 6 vCPUs<br>• 32 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |
| Regional SBLB VM | csp-regional-sblb | Load balancer for device to Fault Management Performance Management (FMPM) microservice connectivity | • 4 vCPUs<br>• 16 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |

Table 10: Details of VMs for a Non-HA Demo Environment *(continued)*

| Components | Name of VM | Components That Installer Places in VM | Resources Required | Ports to Open |
|---|---|---|---|---|
| Contrail Analytics VM | csp-contrailanalytics-1 | Contrail Analytics for a distributed deployment<br><br>For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform. | • 8 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| Virtual Route Reflector VM | csp-vrr-vm | Virtual route reflector (VRR) | • 4 vCPUs<br>• 8 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |
| Central K8 Master VM | csp-central-k8mastervm | Kubernetes Master VM for Central | • 4 vCPUs<br>• 8 GB RAM<br>• 200 GB hard disk storage | - |
| Regional K8 Master VM | csp-regional-k8mastervm | Kubernetes Master VM for Regional | • 4 vCPUs<br>• 8 GB RAM<br>• 200 GB hard disk storage | - |

Table 11 on page 15 shows complete details about VMs and microservice collections required for a production environment without HA.

Table 11: 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<br>• 64 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-infravm | Third -party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-msvm | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm | Third -party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |

Table 11: 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-msvm | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-sblb | Load balancer for device to microservice connectivity | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-space-vm | Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-elkvm | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-elkvm | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-vrr-vm | VRR | • 4 vCPUs<br>• 8 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |

Table 12 on page 16 shows complete details about the VMs for a trial HA environment.

Table 12: Details of VMs for a Trial HA Environment

| Name of VM or Microservice Collection | Components That Installer Places in VM | Resources Required | Ports to Open |
|---|---|---|---|
| csp-installer-vm | — | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-infravm1 | Third-party applications used as infrastructure services | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-infravm2 | Third-party applications used as infrastructure services | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |

Table 12: Details of VMs for a Trial HA Environment *(continued)*

| Name of VM or Microservice Collection | Components That Installer Places in VM | Resources Required | Ports to Open |
|---|---|---|---|
| csp-central-infravm3 | Third-party applications used as infrastructure services | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-lbvm1 | Load-balancing applications | • 4 vCPUs<br>• 24 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-lbvm2 | Load-balancing applications | • 4 vCPUs<br>• 24 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-msvm1 | All microservices, including GUI applications | • 8 CPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-msvm2 | All microservices, including GUI applications | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm1 | Third-party applications used as infrastructure services | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm2 | Third-party applications used as infrastructure services | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm3 | Third-party applications used as infrastructure services | • 8 vCPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-msvm1 | All microservices, including GUI applications | • 8 CPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-msvm2 | All microservices, including GUI applications | • 8 CPUs<br>• 48 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-lbvm1 | Load-balancing applications | • 4 vCPUs<br>• 24 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |

Table 12: Details of VMs for a Trial HA Environment *(continued)*

| Name of VM or Microservice Collection | Components That Installer Places in VM | Resources Required | Ports to Open |
|---|---|---|---|
| csp-regional-lbvm2 | Load-balancing applications | • 4 vCPUs<br>• 24 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-space-vm | Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS | • 4 vCPUs<br>• 16 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-contrailanalytics-1 | Contrail Analytics for a distributed deployment<br><br>For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform. | • 16 vCPUs<br>• 64 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-sblb1 | Load balancer for device to FMPM microservice connectivity | • 4 vCPUs<br>• 24 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-sblb2 | Load balancer for device to FMPM microservice connectivity | • 4 vCPUs<br>• 24 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-vrr-vm | Virtual route reflector (VRR) | • 4 vCPUs<br>• 8 GB RAM<br>• 200 GB hard disk storage | See Table 14 on page 21. |

Table 13 on page 18 shows complete details about VMs and microservice collections required for a production environment with HA.

Table 13: 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<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-infravm1 | Third-party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |

Table 13: 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-infravm2 | Third-party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-infravm3 | Third-party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-lbvm1 | Load-balancing applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-lbvm2 | Load-balancing applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-lbvm3 | Load-balancing applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-msvm1 | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-msvm2 | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-central-msvm3 | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm1 | Third-party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm2 | Third-party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-infravm3 | Third-party applications used as infrastructure services | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |

Table 13: 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-msvm1 | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-msvm2 | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-msvm3 | All microservices, including GUI applications | • 16 vCPUs<br>• 64 GB RAM<br>• 500 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-lbvm1 | Load-balancing applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-lbvm2 | Load-balancing applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-lbvm3 | Load-balancing applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-space-vm | Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-elkvm1 | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-elkvm2 | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-central-elkvm3 | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-elkvm1 | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |

Table 13: 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-elkvm2 | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-elkvm3 | Logging applications | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-sblb1 | Load balancer for device to FMPM microservice connectivity | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-sblb2 | Load balancer for device to FMPM microservice connectivity | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-regional-sblb3 | Load balancer for device to FMPM microservice connectivity | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |
| csp-vrr-vm | Virtual route reflector (VRR) | • 4 vCPUs<br>• 32 GB RAM<br>• 300 GB hard disk storage | See Table 14 on page 21. |

Table 14 on page 21 shows the ports that must be open on all VMs in the Cloud CPE Solution 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 14: 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 |
| 443 | External and internal | HTTPS, including Administration Portal and Customer Portal |

Table 14: Ports to Open on CSO VMs *(continued)*

| Port Number | CSO Communication Type | Port Function |
|---|---|---|
| 514 | Internal | Syslog receiving port |
| 1414 | Internal | Cassandra Java Virtual Machine (JVM) |
| 1936 | External | HAProxy status page |
| 1947 | External | Icinga service |
| 2181 | Internal | ZooKeeper client |
| 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 |
| 4505, 4506 | Internal | Salt communications |
| 5000 | External | Keystone public |
| 5044 | Internal | Beats |
| 5543 | Internal | Logstash UDP |
| 5601 | External | Kibana UI |
| 5665 | Internal | Icinga API |
| 5671 | Internal | RabbitMQ SSL listener |
| 5672 | Internal | RabbitMQ client |
| 6000 | Internal | Swift Object Server |
| 6001 | Internal | Swift Container Server |
| 6002 | Internal | Swift Account Server |

Table 14: Ports to Open on CSO VMs *(continued)*

| Port Number | CSO Communication Type | Port Function |
|---|---|---|
| 6379 | Internal | Redis |
| 6543 | Internal | Virtualized Network Function manager (VNFM) |
| 7804 | External | Device connectivity |
| 8006 | Internal | Network Service Orchestrator |
| 8016 | Internal | Notification engine |
| 8080 | Internal | cAdvisor |
| 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 |
| 9042 | Internal | Cassandra native transport |
| 9090 | Internal | Swift Proxy Server |
| 9160 | Internal | Cassandra |
| 9200 | Internal | Elasticsearch |
| 10248 | Internal | kubelet healthz |
| 15100 | Internal | Logstash TCP |
| 15672 | Internal | RabbitMQ management |
| 30000-32767 | Internal | Kubernetes service node range |
| 30900 | External | Prometheus |
| 35357 | Internal | Keystone private |

## VNFs Supported

The Cloud CPE solution supports the Juniper Networks and third-party VNFs listed in Table 15 on page 24.

Table 15: VNFs Supported by the Cloud CPE Solution

| VNF Name | Version | Network Functions Supported | Deployment Model Support | Element Management System Support |
|---|---|---|---|---|
| Juniper Networks vSRX | vSRX KVM Appliance 15.1X49-D123 | • Network Address Translation (NAT)<br>• Demonstration version of Deep Packet Inspection (DPI)<br>• Firewall | • Centralized deployment<br>• Distributed deployment supports NAT and firewall | Element Management System (EMS) microservice, which is included with Contrail Service Orchestration (CSO) |
| LxCIPtable (a free, third-party VNF based on Linux IP tables | 14.04 | • NAT<br>• Firewall | Centralized deployment | EMS microservice |
| Cisco Cloud Services Router 1000V Series (CSR-1000V) | 3.15.0 | Firewall | Centralized deployment | Junos Space Network Management Platform |
| Riverbed SteelHead | 9.2.0 | WAN optimization | Distributed deployment, NFX250 platform only | EMS microservice |
| Silver Peak VX | VXOA 8.0.5.0_61631 | WAN optimization | Distributed deployment, NFX250 platform only | EMS microservice |

## Licensing

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

The Cloud CPE solution 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 GUIs

NOTE: We recommend that you use Google Chrome Version 60 or later or Mozilla Firefox Version 57 or later to access the Contrail Service Orchestration (CSO) GUIs.

Table 16 on page 25 shows the URLs and login credentials for the GUIs for a non-redundant CSO installation.

### Table 16: Access Details for the GUIs

| GUI | URL | Login Credentials |
|---|---|---|
| Administration Portal | http://*infra-vm-IP-Address* \| *ha-proxy-IP-Address*<br><br>where:<br><br>• *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.<br>• *ha-proxy-IP-Address*—IP address of HA proxy. Use this option to monitor microservices.<br><br>NOTE:<br>• For a non-HA deployment, use the IP address of the VM that hosts the microservices for the central POP.<br>• For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO.<br><br>For example, https://192.0.2.1 | Specify the OpenStack Keystone username and password.<br><br>The default username is **cspadmin** and the default password is **passw0rd**. |
| 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. |

Table 16: Access Details for the GUIs *(continued)*

| GUI | URL | Login Credentials |
|---|---|---|
| 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<br><br>where:<br><br>*central-IP-Address*—IP address of the VM that hosts the microservices for the central POP<br><br>For example:<br><br>*https://192.0.2.1:83* | Specify the OpenStack Keystone username and password.<br><br>The default username is **cspadmin** and the default password is **passw0rd**. |
| Service and Infrastructure Monitor<br><br>This tool provides monitoring and troubleshooting for network services in a hybrid WAN deployment. | http://*ha-proxy-IP-Address*:1947/icingaweb2<br><br>where:<br><br>*ha-proxy-IP-Address*—IP address of HA proxy.<br><br>• For a non-high availability (non-HA) deployment, use the IP address of the VM that hosts the microservices for the central POP.<br>• For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO.<br><br>For example:<br><br>*http://192.0.2.1:1947/icingaweb2* | The default username is **icinga** and the password is the common password that you configure for the infrastructure services when you install CSO. The default infrastructure services password is **passw0rd**. |
| Kibana<br><br>This tool provides a visual representation of log files. You can use it to monitor:<br><br>• Network services in a central or regional POP<br>• Microservices in the deployment | http://*infra-vm-IP-Address* \| *ha-proxy-IP-Address*:5601<br><br>where:<br><br>*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.<br><br>*ha-proxy-IP-Address*—IP address of high availability (HA) proxy. Use this option to monitor the microservices.<br><br>• For a non-HA deployment, use the IP address of the VM that hosts the microservices for the central POP.<br>• For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO.<br><br>For example:<br><br>*http://192.0.2.2:5601* | Login credentials are not needed. |

Table 16: Access Details for the GUIs *(continued)*

| GUI | URL | Login Credentials |
|---|---|---|
| Grafana and Prometheus<br><br>These tools provide monitoring and troubleshooting for the infrastructure services in the Cloud CPE solution. 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<br>• Grafana—*ha-proxy-IP-Address*:3000<br><br>where:<br><br>*ha-proxy-IP-Address*—IP address of HA proxy<br><br>• For a non-HA deployment, use the IP address of the VM that hosts the microservices for the central POP.<br>• For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO.<br><br>For example:<br><br>*http://192.0.2.2:30900* | Login credentials are not needed. |

## Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Juniper Networks Cloud CPE Solution Release 3.1.2.

- If the Kubernetes minion node in the central or regional microservices 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.

  To rebalance the pods back to the Kubernetes minion node that was down, do the following:

  1. Check the status of the **kube-proxy** process on the minion node by executing the **kubectl get pods --namespace=kube-system** command.

     A sample output is shown below.

     ```
     root@host:~# kubectl get pods --namespace=kube-system
     NAME                                    READY     STATUS    RESTARTS   AGE
     etcd-empty-dir-cleanup-192.0.2.1    1/1       Running   1          1d
     kube-addon-manager-192.0.2.1        1/1       Running   1          1d
     kube-apiserver-192.0.2.1            1/1       Running   1          1d
     kube-controller-manager-192.0.2.1   1/1       Running   1          1d
     kube-dns-v11-lcs1x                      4/4       Running   4          1d
     kube-proxy-192.0.2.1                1/1       Running   0          1d
     kube-proxy-192.0.2.2                1/1       Unknown   0          1d
     kube-scheduler-192.0.2.1            1/1       Running   1          1d
     kubernetes-dashboard-1579006691-1fvmk   1/1       Running   1          1d
     ```

  2. If the status of the **kube-proxy** process on the Kubernetes minion node is **Unknown**, execute the **kubectl delete pod kube-proxy-*minion-IP-address*--namespace=kube-system --grace-period=0 –force** command, where *minion-IP-address* is the IP address of the minion node that was down.

  3. Verify that the status of the **kube-proxy** process is **Running**.

  4. Execute the command to rebalance the nodes:

     - If you are running a trial HA setup, execute the **kubectl delete pods --all --grace-period=0** command on the Kubernetes master node.
     - If you are running a production HA setup, execute the **kubectl delete pods --all --grace-period=0** command on the Kubernetes master node and the Kubernetes minion node that did not go down.

- We recommend that you do not delete existing LAN segments from a site because this might impact firewall and SD-WAN policy deployments. [CXU-13683]
- 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.pyc** 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.pyc file is located in the /root/Contrail_Service_Orchestration_3.1.2/deployments/central/file_root/contrail_openstack/ directory on the VM or server on which you installed CSO.

   b. Rename the file to **jsm.pyc** 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.

- When a tenant object is created through Administration Portal or the API for a centralized deployment, Contrail OpenStack 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.
- Contrail Service Orchestration does not offer a single 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.
- You can use Administration Portal to upload licenses to Contrail Service Orchestration; however, you cannot use Administration Portal to install licenses on physical or virtual devices that Contrail Service Orchestration manages. You must use the APIs or the license installation tool to install licenses on devices.

- Contrail Service Orchestration 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 Administration Portal to specify the root password in the device template.

  To specify a root password for the device:

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

- You can use the logs on an NFX250 device to review the status of the device's activation.
- In Cloud CPE Solution Release 3.1.2, intrusion prevention system (IPS) is not supported. Therefore, in the IPS report, the attack name from the IPS signatures is displayed as **UNKNOWN**.
- In Cloud CPE Solution Release 3.1.2, the virtual machine (VM) on which the virtual route reflector (VRR) is installed supports only one management interface.
- In Cloud CPE Solution Release 3.1.2, high availability for ArangoDB is not supported. Therefore, ensure that the central infrastructure VM, where ArangoDB is running, is not brought down or does not fail. If the VM is down, bring it up immediately for CSO to be operational.
- In Cloud CPE Solution Release 3.1.2, 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.
- On the SRX Series device, the traffic for any of the following combinations does not flow because of an inability to identify the reverse route for traffic terminating through MPLS:
  - Site to site
  - Site to department or vice versa
  - Department to department

  To enable traffic, you must add a firewall rule permitting traffic from the corresponding department's zone (where the traffic was supposed to terminate) to the Trust zone on the destination site. However, we do not recommend doing this because the rule can conflict with existing firewall intents.
- In Cloud CPE Solution Release 3.1.2, SSL proxy is not supported.

- 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.
- 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.
- When you activate a CPE device with WAN interfaces configured for DHCP:
  - Ensure that all the WAN interfaces configured for DHCP have the IP address allocated from the DHCP server.
  - When multiple WAN links are configured for DHCP, in some cases, all DHCP servers will advertise a default route to the CPE device, which can lead to traffic being routed through an undesired WAN interface, which could then stop the GRE and IPsec tunnels from being operational.

    To avoid this scenario, configure a static route through each WAN interface to reach the tunnel endpoint through the desired WAN interface.
- When you create LAN segments, the LAN segment table does not display the DHCP settings even though the changes are saved successfully.
- When you trigger the ZTP workflow on an NFX Series device, we recommend that you use the activation code for the device and initiate the activation by using the **Activate Device** link in the CSO GUI.
- In the SLA Performance page (**Monitor > Applications > SLA Performance**, the scatter plot displays the SLA name as **UNKNOWN** for the applications for which the SLA is not violated.
- On an NFX Series device, if you try to install the signature database before installing the application identification license, the signature database installation fails.

  Ensure that you first install the application identification license and then install the signature database.

  To install the application identification license:

  NOTE: **Ensure that you have the license ready before you begin this procedure.**

  1. SSH to the vSRX gateway router running on NFX Series device and log in as **root**.
  2. Access the Junos OS CLI and enter the operational mode.
  3. Execute the **show system license** command to view the existing license so that you can verify (in a subsequent step) that the license is added.

     A sample output is as follows:

```
root@host> show system license
License usage:
                                Licenses     Licenses    Licenses    Expiry

  Feature name                        used    installed      needed
  Virtual Appliance                      1            1           0    55
days
  remote-access-ipsec-vpn-client         0            2           0
permanent

Licenses installed:
  License identifier: XXXXXXXXXX
  License version: 4
  Software Serial Number: XXXXXXXX
  Customer ID: XXXXXXXXXXXXXXXX
  Features:
    Virtual Appliance - Virtual Appliance
      count-down, Original validity: 60 days
```

4. Execute the **request system license add terminal** command.

5. Copy the license, paste it into the terminal, and press Ctrl+D.

   If the license is added successfully, a confirmation message is displayed as shown in the following sample output:

```
root@host> request system license add terminal
[Type ^D at a new line to end input,
enter blank line between each license key]
add license complete (no errors)
```

6. Execute the **show system license** command and compare the output with the one obtained in step 3 to verify that the license is added.

   A sample output is as follows:

```
root@host> show system license
License usage:
                                Licenses     Licenses    Licenses    Expiry

  Feature name                        used    installed      needed
  Virtual Appliance                      1            1           0    55
days
  remote-access-ipsec-vpn-client         0            2           0
permanent

Licenses installed:
  License identifier: XXXXXXXXXX
  License version: 4
  Software Serial Number: XXXXXXXX
  Customer ID: XXXXXXXXXXXXXXXX
  Features:
    Virtual Appliance - Virtual Appliance
      count-down, Original validity: 60 days

  License identifier: YYYYYYYYYY
  License version: 4
  Software Serial Number: YYYYYYYYYYYYYY
  Customer ID: YYYYYYYYYY
```

```
Features:
  appid-sig        - APPID Signature
    date-based, 2016-04-05 00:00:00 UTC - 2017-04-06 00:00:00 UTC
  idp-sig          - IDP Signature
    date-based, 2016-04-05 00:00:00 UTC - 2017-04-06 00:00:00 UTC
```

7. Exit the Junos OS CLI and log out of vSRX.

## Known Issues

This section lists known issues for the Juniper Networks Cloud CPE Solution Release 3.1.2.

- On an NFX Series device, application tracking is enabled for department security zones only on pushing an SD-WAN APBR policy. When there is only a firewall policy deployed without SD-WAN, no application visibility is displayed for the NFX Series device.

  Workaround: Deploy an SD-WAN APBR policy to enable application visibility on the NFX Series device. [CXU-12154]

- If the oam-and-data interface goes down, the IBGP session is lost and traffic stops flowing. This causes communication with CSO to be lost and Syslogs are not sent even though there are other WAN interfaces up and running.

  Workaround: None. [CXU-12346]

- For a site, when DHCP is configured on the WAN interface and the LAN segment, the device activation fails.

  Workaround: None. [CXU-13432]

- On the Site Management page, the information displayed on the Overlay Links tab for the cloud hub is incorrect.

  Workaround: The correct information is displayed in the WAN tab of the cloud hub. [CXU-13579]

- If you configure a static SD-WAN policy and a link goes down, it might take approximately three minutes for the gr-0/0/0 interface to be removed from the MPLS/Internet routing table.

  Workaround: None. [CXU-13528]

- If the central infrastructure VM central-infra-vm2 goes down, users might not be able log in to the CSO GUI for two hours (the configured keystone timeout). This is because of an issue with the synchronization of the MySQL token.

  Workaround: Restart the central infrastructure VM central-infra-vm1. Restarting the VM resynchronizes the MySQL token. Users will then be able to log in to the CSO GUI. [CXU-13541]

- If you create an SD-WAN policy and a firewall policy with the source as a department and the department is not associated with a site or a LAN segment, the job created to apply the SD-WAN and firewall policies after ZTP fails.

  Workaround: Associate the department with a site or a LAN segment before performing ZTP. [CXU-13542]

- For sites with device-initiated connections, by default, all site traffic is source NATted at the hub. You cannot apply a different source NAT rule to the hub because the default rule overrides any user-configured source NAT rule.

  Workaround: None. [CXU-13558]

- If you have a tenant with more than one site and deploy a firewall policy to a single site, the policy is deployed only to that site. However, jobs are created to push a dummy firewall policy to other sites, which causes a performance issue on setups with a large number of devices.

  Workaround: None. [CXU-13562]

- On SRX300 Series devices, phone home activation might fail because of a conflict with default configuration on the device.

  Workaround: Copy the stage-1 configuration manually to the device and then trigger the phone home activation workflow. [PR 1312703]

- If you deploy a NAT policy with one or more rules and then delete the policy without first deleting the rules, the configuration on the device is not cleared.

  Workaround: To delete a NAT policy, first delete all the rules associated with the policy and deploy the policy. After the deployment is successful, delete the NAT policy by using the CSO GUI. [CXU-13879]

- Link switching does not occur even though the throughput threshold configured in the SLA profile is crossed because incorrect interface stats are reported for the GRE interface.

  Workaround: None. [CXU-14127]

- The device profile srx-deployment-option-1 assigns OAM traffic to the fxp0 interface, which is not available on the SRX Series Services Gateway.

  Workaround: Edit the stage 1 configuration for the CPE device in Customer Portal to use an interface other than fxp0 for OAM traffic. [CXU-3779]

- The traffic rate value does not change when you monitor a service on an SRX Series Services Gateway in Administration Portal.

  Workaround: None

  [CXU-3822]

- The NFX250 device does not receive communications to block unicast reverse path forwarding (uRPF) because two interfaces on the NFX250 device communicate separately with one interface on the regional microservices server.

  Workaround: Disable the uRPF check in JDM on all interfaces for each NFX Series device.

  [CXU-3854]

- Performance metrics for the NFX Series device are collected through the HTTP interface.

  Workaround: None. [CXU-8710]

- In the detailed view of a site on the Sites page (**Sites > Site Management**), the **Overlay Links** tab displays only GRE links and not GRE over IPsec links.

  Workaround: None. [CXU-10170]

- On the NAT Rules page, if you try to search or use the column filters for departments named **Internet** or **Corporate Network**, the search does not work.

  Workaround: None. [CXU-10406]

- The **Activate Device** link is enabled even though activation is in progress.

  Workaround: After clicking the **Activate Device** link, wait for the activation process to complete; the **Activate Device** link is disabled after activation completes. [CXU-10760]

- If one of the CSO infrastructure nodes (virtual machines) fails (shuts down or restarts), the topology service repeatedly sends out the error message "**Failed to consume message from queue: 'NoneType' object has no attribute 'itervalues'**" to the logs.

  Workaround: After the infrastructure node fails over to the backup node, wait for ten minutes before using any workflows. [CXU-11004]

- If a user who belongs to more than one tenant is deleted, the user is deleted from all tenants.

  Workaround: None. [CXU-11201]

- On the *Site-Name* page (**Sites > *Site-Name***), when you click the **Device** link (in the *Connectivity & Devices* box on the **Overview** tab), you are navigated to the Devices page (**Resources > Devices**) where all available devices are displayed instead of only the device that you selected.

  Workaround: None. [CXU-11339]

- If you modify a site group that is not used in any policy, a GUI notification incorrectly indicates that policies need to be deployed.

  Workaround: Deploy the indicated policies on the device. However, because there are no changes to be deployed, the configuration is not deployed. [CXU-11395]

- If an infrastructure node (virtual machine) goes down, the backup node takes over the tasks handled by the primary node. However, after the node that was down recovers, it does not join the cluster and is stuck in slave mode.

  Workaround: Ensure that both nodes are up before you perform the following tasks:

  1. Log in to the infrastructure node as **root**.
  2. Open a shell prompt and access the redis CLI by executing the **redis-cli -h *node-IP* -p 6379 -c** command, where *node-IP* is the IP address of the infrastructure node that went down.
  3. At the redis CLI prompt, execute the **CLUSTER FAILOVER** command.
  4. Execute the **INFO** command.

5. In the output, under **replication**, ensure that the node is displayed as **Master**.
6. Exit the redis CLI by executing the **QUIT** command.
7. Log out of the infrastructure node.

[CXU-11711]

- If you create a LAN segment with the name LAN0, LAN1, or LAN2, the deployment of the LAN segment fails.

  Workaround: Do not use the names LAN0, LAN1, or LAN2 when you create a LAN segment. [CXU-11743]

- In the device template, if the **ZTP_ENABLED** and **ACTIVATION_CODE_ENABLED** flags are set to true, you cannot proceed with device activation.

  Workaround: Set the **ZTP_ENABLED** flag to true and the **ACTIVATION_CODE_ENABLED** flag to false before proceeding with device activation. [CXU-11794]

- When you upgrade the gateway router (GWR) 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.

  [CXU-11823]

- On rare occasions, the Logspout microservice causes the docker daemon to hog the CPU on the microservices virtual machine.

  Workaround: Restart the Logspout microservice by doing the following:

  1. Log in to the central microservices virtual machine as root.

2. At the shell prompt, run the **kubectl get pod** command to find out the name of the Logspout pod.

3. Restart the pod by executing the **kubectl delete pod *pod-name*** command, where *pod-name* is the name of the Logspout pod.

[CXU-11863]

- If you trigger a device activation on the Activate Device page, the status of the activation is displayed based on the progress of the activation steps completed. However, the device activation process takes between 20 and 30 minutes and if you click **OK** to close the Activate Device page, you cannot go back to the Activate Device page to find out the status.

  Workaround:

  1. Wait for 20 to 30 minutes after you trigger the activation.

  2. On the Sites page, hover over the device icon to see the status:

     - If the device status is **PROVISIONED**, it means that the activation was successful.
     - If the device status is **EXPECTED** or **PROVISION_FAILED**, then the device activation has failed.

       Contact your service provider for further assistance.

  [CXU-11878]

- When a tunnel goes down, the event generated displays different information for the NFX Series and SRX Series devices:

  - When the GRE over IPsec tunnel goes down:
    - The event generated for the vSRX device (running on the NFX Series device) has the description **['Tunnel-id *ID* is inactive']**.
    - The event generated for the SRX Series device has the description **GRE over IPSEC is Down**.
  - When the GRE-only tunnel goes down:
    - The event generated for the vSRX device (running on the NFX Series device) has the description **tunnel-oam-down**.
    - The event generated for the SRX device has the description **GRE tunnel down**.

  Workaround: None. [CXU-11895]

- If you try to delete one or more LAN segments, the confirmation dialog box does not display the list of LAN segments selected for deletion. However, when you click **OK** to confirm the deletion, the LAN segments are deleted successfully.

  Workaround: None. [CXU-11896]

- If the role of an existing user is changed from MSP Operator to MSP Administrator and that user tries to switch the tenant by using the scope switcher in the banner, the tenant switching fails.

  Workaround: Delete the existing user and add an MSP user with the MSP Administrator role. The new user will be able to perform the tenant switch. [CXU-11898]

- In Cloud CPE Solution Release 3.1.2, editing a site is not supported. When you try to edit a site, the message "**unable to retrieve the router info**" is displayed.

  Workaround: Delete the site and add the site again with the modified settings. [CXU-11912]

- If you edit an existing LAN segment that was previously added during site creation, the **Department** field is changed to **Default**.

  Workaround: When you edit a LAN segment, ensure that you select the correct department before saving your changes. [CXU-11914]

- If you apply an APBR policy on a vSRX device or an SRX Series device, in some cases, the APBR rule is not active on the device.

  Workaround:

  1. Log in to the vSRX or SRX Series device in configuration mode.
  2. Deactivate the APBR policy by executing the **delete apply-groups srx-gwr-apbr-policy-config** command.
  3. Commit the configuration by executing the **commit** command.
  4. Activate the APBR policy by executing the **set apply-groups srx-gwr-apbr-policy-config** command.
  5. Commit the configuration by executing the **commit** command.
  6. Log out of the device.

  [CXU-11920]

- On the **Monitor > Overview** page:
  - The number of hub devices is reported as zero even though a cloud hub exists.

    Workaround: The expanded view displays the correct data.
  - When you collapse and expand the map view, the number of links reported is incorrect.

    Workaround: Refresh the page to display the correct data.

  [CXU-11931]

- For GRE-over-IPsec overlays, in some cases, the **event-options** configuration fails to re-enable the gr-0/0/0 interface. As a result, the traffic between the spoke and hub overlay stops flowing even though the overlay is up.

  Workaround: Do the following:

  1. Find out which links are affected by checking the status of the real-time performance monitoring (RPM) probe in the UI:

     a. On the ***Site-Name*** page (**Site > Site Management > Site-Name**), select the **WAN** tab.

     b. On the link between the spoke and the hub, expand the WAN links by clicking the number of WAN links displayed.

     c. For each WAN link, select the link and check the packet loss value displayed on the right side.

        The links for which the packet loss is 100% indicates that the links are down.

  2. Find out which gr-0/0/0 interfaces are down on the spoke device:

     a. Log in to the spoke device as **root**.

     b. Execute the **show interfaces gr-0/0/0.* terse** command.

        An example of the command and output follows:

```
user@host> show interfaces gr-0/0/0.* terse
  Interface Admin Link Proto Local Remote
gr-0/0/0.1 up up inet 192.0.2.1/31
                                    mpls
gr-0/0/0.2 down up inet 192.0.2.2/31
                                    mpls
gr-0/0/0.3 up up inet 192.0.2.3/31
                                    mpls
gr-0/0/0.4 up up inet 192.0.2.4/31
                                    mpls
gr-0/0/0.5 down up inet 192.0.2.5/31
                                    mpls
gr-0/0/0.6 up up inet 192.0.2.6/31
                                    mpls
```

  3. For each disabled interface on the spoke, execute the **show configuration | display set | match "disable" | match "*gr-interface-name* unit *unit-number*"** command to find out whether any disabled statements are present in the configuration, where *gr-interface-name* is the name of the interface and *unit-number* is the unit number that was obtained in Step 2.

     An example of the command and output follows:

```
user@host>show configuration | display set | match "disable" | match "gr-0/0/0
unit 5"
set groups NFX-5-SPOKE_CPE1_WAN_2_GRE_IPSEC_0 interfaces gr-0/0/0 unit 5
disable
```

4. Find out whether there is a problem by doing the following:

   a. Execute the **show interfaces st0* terse** operational command to find out the status of the IPsec tunnels and the IP addresses:

      An example of the command and output follows:

      ```
      user@host> show interfaces st0* terse
      Interface Admin Link Proto Local Remote
      st0 up up
      st0.1 up down inet 192.0.2.7/31
      st0.2 up up inet 192.0.2.8/31  # [IP address match]
      st0.3 up up inet 192.0.2.9/31
      ```

   b. Execute the **show interfaces *gr-down-interface-name*** command, where *gr-down-interface-name* is the name of the interface that was down (obtained in Step 2).

      NOTE: You must execute this command for all the interfaces that were down.

      An example of the command and output follows:

      ```
      user@host> show interfaces gr-0/0/0.5
        Logical interface gr-0/0/0.5 (Index 139) (SNMP ifIndex 579)
          Flags: Down
      Down Point-To-Point SNMP-Traps 0x4000
      IP-Header 192.0.2.10/31:192.0.2.8/31:47:df:64:0000000000000000  # [IP
      address match]
      Encapsulation: GRE-NULL
          Gre keepalives configured: Off, Gre keepalives adjacency state: down
          Input packets : 0
          Output packets: 0
          Security: Zone: trust
          Allowed host-inbound traffic : bootp bfd bgp dns dvmrp igmp ldp msdp
       nhrp
         ospf ospf3 pgm pim rip ripng router-discovery rsvp sap vrrp dhcp finger
       ftp
         tftp ident-reset http https ike netconf ping reverse-telnet reverse-ssh

          rlogin rpm rsh snmp snmp-trap ssh telnet traceroute xnm-clear-text
      xnm-ssl
          lsping ntp sip dhcpv6 r2cp webapi-clear-text webapi-ssl tcp-encap
          Protocol inet, MTU: 9168
            Flags: Sendbcast-pkt-to-re
            Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
              Destination: 192.0.2.11/31, Local: 192.0.2.5/31
          Protocol mpls, MTU: 9156, Maximum labels: 3
      ```

c. For each IP address obtained in Step a, search the output from Step b for a match. (In the sample outputs, the IP address of the st0.2 interface (192.0.2.8/31) matches the IP address in the IP-Header parameter.)

d. If the st0 interface corresponding to the matched IP address is up and the corresponding gr-0/0/0 interface is down, then there is a problem with the configuration.

5. Modify the configuration on the spoke as follows:

a. Log in to the spoke device as **root**.

b. Delete the disabled statements found in Step 3 by executing the **delete** command. An example is provided below:

```
user@host# delete groups NFX-5-SPOKE_CPE1_WAN_2_GRE_IPSEC_0 interfaces
gr-0/0/0 unit 5 disable
```

c. Commit the configuration by executing the **commit** command.

6. Repeat Step 2 through Step 5 for the hub device.

7. Verify that the links are up by following the procedure in Step 1.

[CXU-11996]

- If an SLA profile is defined with only the throughput metric specified, in some cases, the SLA profile is assigned to a link that is down.

  Workaround: In addition to the throughput metric, ensure that you specify at least one more metric (for example, packet loss or latency) for the SLA profile. [CXU-11997]

- In a three-node setup, two nodes are clustered together, but the third node is not part of the cluster. In addition, in some cases, the RabbitMQ nodes are also not part of the cluster. This is a rare scenario, which can 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.

[CXU-12107]

- If you create a site group and add a hub site that references that site group, the creation of the hub site fails.

  Workaround: Add the hub site without using a site group. [CXU-13753]

- If you delete an existing LAN segment, the configuration related to the LAN segment is not deleted from the cloud hub. An attempt to redeploy the same configuration can cause problems. A sample LAN segment configuration is shown below:

  **set groups srx-hub-static-Clouldhub_DefaultVPN-1.4.20.14 routing-options static route 192.0.2.0/24 next-table Default-Clouldhub_DefaultVPN-Clouldhub.inet.0**

  Workaround: To redeploy the same configuration, first delete the LAN segment configuration from the cloud hub by using the Junos OS CLI, and then redeploy the configuration. A sample delete command is shown below:

  **delete groups srx-hub-static-Clouldhub_DefaultVPN-1.4.20.14 routing-options static route192.0.2.0/24**

  [CXU-13812]

- In some cases, when the activation of a spoke device fails, the UI logs do not display the relevant information.

  Workaround: Access the Kibana logs to view the relevant information. [CXU-13941]

- If you try to activate multiple spoke devices, which share a common hub device, at the same time, then in some cases, the activation fails.

  Workaround: Retry the activation on the spoke devices for which the activation fails. [CXU-14066]

- After you restart the physical server that hosts the VRR VM, the VRR VM fails to come back online.

  Workaround: Do the following:

 

1. Destroy the corrupted VRR by executing the **virsh destroy vrr** command from the KVM console.

2. Delete the image file (**/var/lib/libvirt/images/vrr.img**) of the corrupted VRR from the physical server.

3. Copy the VRR image file (**vrr.img**) packaged with your CSO installation from the server or VM on which you installed CSO to the **/var/lib/libvirt/images/**directory on the physical server.

4. Start the VRR by executing the **virsh start vrr** command from the KVM console.

5. Load the base configuration of the VRR and modify the IP address of the VRR and the regional microservices VM. A sample configuration is shown below.

```
set system root-authentication encrypted-password "$ABC123"
set system services ssh
set system services telnet
set system services netconf ssh
set system syslog user * any emergency
set system syslog file messages any notice
set system syslog file messages authorization info
set system syslog file interactive-commands interactive-commands any
set interfaces em0 unit 0 family inet address 192.0.2.6/24
set interfaces em0 unit 0 family mpls
set routing-options rib inet.3 static route 0.0.0.0/0 discard
set routing-options static route 0.0.0.0/0 next-hop 192.0.2.1
set groups vrr-base-config system host-name vrr-192.0.2.6
set groups vrr-base-config system services ssh
set groups vrr-base-config system services telnet
set groups vrr-base-config system syslog user * any emergency
set groups vrr-base-config system syslog file messages any notice
set groups vrr-base-config system syslog file messages authorization info
set groups vrr-base-config system syslog file interactive-commands
interactive-commands any
set groups vrr-base-config routing-options autonomous-system 64512
set groups vrr-base-config protocols bgp hold-time 65535
set groups vrr-base-config protocols bgp group ibgp type internal
set groups vrr-base-config protocols bgp group ibgp local-address 192.0.2.6
set groups vrr-base-config protocols bgp group ibgp family inet-vpn unicast
set groups vrr-base-config protocols bgp group ibgp cluster 192.0.2.6
set apply-groups vrr-base-config
```

6. Create a dummy tenant by using the CSO GUI.

   CSO pushes the configuration to the new VRR including the previously configured tenants and sites, and the network is restored.

[CXU-14922]

- If you configure a site with Internet breakout, the CSO GUI displays a core attachment point on the *Site-Name* page. If you try to start a service on the site, an error message is displayed.

 

Workaround: None. [CXU-14924]

- If you configure Cloud CPE Solution Release 3.1.2 to use an external keystone, tenants created in the external keystone (without using the CSO GUI) are visible in the scope switcher in Administration Portal.

  Workaround: Before using an external keystone for CSO, ensure that it does not contain any data. [CXU-14931]

- In some scenarios, the configuration database is locked, which results in a failure in pushing the configuration to the device through CSO. This issue is caused by the **event-options** configuration.

  Workaround: Log in to the spoke device, and use the Junos OS CLI to remove the **event-options** configuration in the spoke by executing the following commands:

  **set groups VF-SPOKE1_CPE1_WAN_0_GRE_IPSEC_0 event-options policy VF-SPOKE1_CPE1_WAN_0_GRE_IPSEC_0_down then change-configuration commands "set groups VF-SPOKE1_CPE1_WAN_0_GRE_IPSEC_0 interfaces gr-0/0/0.1 disable"**
  **set groups VF-SPOKE1_CPE1_WAN_0_GRE_IPSEC_0 event-options policy VF-SPOKE1_CPE1_WAN_0_GRE_IPSEC_0_up then change-configuration commands "delete groups VF-SPOKE1_CPE1_WAN_0_GRE_IPSEC_0 interfaces gr-0/0/0.1 disable"**
  **set groups VF-SPOKE1_CPE1_WAN_1_GRE_IPSEC_0 event-options policy VF-SPOKE1_CPE1_WAN_1_GRE_IPSEC_0_down then change-configuration commands "set groups VF-SPOKE1_CPE1_WAN_1_GRE_IPSEC_0 interfaces gr-0/0/0.2 disable"**
  **set groups VF-SPOKE1_CPE1_WAN_1_GRE_IPSEC_0 event-options policy VF-SPOKE1_CPE1_WAN_1_GRE_IPSEC_0_up then change-configuration commands "delete groups VF-SPOKE1_CPE1_WAN_1_GRE_IPSEC_0 interfaces gr-0/0/0.2 disable"**

  [CXU-15025]

- 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: If you want to create more than 120 instances (up to a maximum of 500 instances), use the Heat Version 1.0 APIs instead. [CXU-15033]

- For the centralized CPE use case, during the **ns-terminate** operation (to terminate the VNF), the message **4ca27229-4f0e-3a6e-907f-4ea7706a04e9 exception: es index not present** might be present in some logs. However, this does not cause any problems.

  Workaround: None. [CXU-15047]

- For the centralized CPE use case, during the **ns-terminate** operation (to terminate the VNF), the message **EMS device with id *id-number*** might be present in some logs. However, this does not cause any problems.

  Workaround: None. [CXU-15082]

- If you create a site, activate the spoke device using ZTP, and create an SD-WAN policy and deploy it on the site, the job is created and runs successfully. However, even though the policy is deployed on the site, the policy is still listed as awaiting deployment in the UI.

  Workaround: None. [CXU-15101]

- If you deploy a firewall policy with departments (to which LAN segments are linked) and, after the policy is deployed successfully, delete a LAN segment that is the last

LAN segment for the department (referenced in the policy) and redeploy the firewall policy, then a deployment error occurs. If you modify the firewall policy intent to remove the department and then deploy the policy, all the firewall rules are deleted from the device.

Workaround: Create a LAN segment that is associated with the department referenced in the firewall policy and deploy the LAN segment. After the deployment is successful, redeploy the firewall policy; after the deployment is successful, the firewall rules are pushed to the device. [CXU-15138]

- If you add a tenant but do not add any sites to the tenant and then delete the tenant, error messages are displayed in the Administration Portal GUI. However, the tenant is deleted successfully.

  Workaround: None [CXU-15253]

- If you activate an NFX Series device from the CSO GUI, the device status does not change from Expected to Activate for approximately 20 minutes.

  Workaround: None [CXU-15282]

- In an HA setup, after a central microservices VM fails over, the Tenant, Site, and Service Manager (TSSM) API server is sluggish in responding to requests.

  Workaround: Wait for two to three minutes and the TSSM API server responds normally. [CXU-15283]

- In an HA setup, if the central load-balancer VM (as the Virtual Router Redundancy Protocol [VRRP] master) goes down, the switchover takes place in less than a minute and the services run properly with the new VRRP master. However, no jobs are created for approximately 10 minutes.

  Workaround: None. [CXU-15388]

- The CPU utilization of the TSSM core microservice might be high in a long-running setup with a large number (greater than 100) of service instances.

  Workaround: Restart the TSSM core Docker containers by deleting the respective pods, and the CPU utilization returns to normal. [CXU-15434]

- In an HA setup, with three load-balancer VMs, if the master load balancer goes down, one of the remaining load-balancer VMs is switched over as the master. However, after the original load-balancer VM comes up, it is switched over as the master again.

  Workaround: None. [CXU-15441]

- In an HA setup, the time configured for the Contrail Analytics Node (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 re-synchronize their times with can-vm1. [CXU-15681]

- If you activate a spoke device by using ZTP and trigger a recall of the device, the APBR, class-of-service (CoS), and tenant configurations are not deleted from the device.

  Workaround: Before initiating the recall operation, do the following:

  1. Delete the policy intents from the SD-WAN policy.

  2. Deploy the policy to ensure the removal of all APBR-related and CoS-related configurations from the device.

     After the deployment is successful, initiate the recall operation.

  [CXU-15815]

- If you try to create a tenant and the tenant creation fails, retrying the creation of the same tenant fails and the following error message is displayed: **The resource specified already exists**.

  Workaround: None. [CXU-15824]

- If the central server (hosting the csp-central-msvm3, csp-central-infravm3, and csp-central-lbvm3 VMs) goes down, jobs are displayed twice or thrice.

  Workaround: None. However, this issue does not occur after the job services re-establish the connection to the RabbitMQ server. [CXU-15844]

- If you provision an SRX300 CPE device by using ZTP, deploy a firewall policy, and then perform the steps to configure and deploy an SD-WAN policy, the SD-WAN policy deployment fails even though the policy is pushed to the device.

  Workaround: None. [CXU-15918]

- If you create an SD-WAN policy applicable to all sites and deploy the policy on a spoke device, the policy is deployed successfully. When you provision a new spoke device by using ZTP, the deployment of the SD-WAN policy on the new device is triggered automatically. However, the SD-WAN policy deployment fails because of an error in the central microservices VM.

  Workaround: None. [CXU-15949]

- In an HA setup, if one of the central servers is restarted, HAProxy reports that the microservices are flapping (disconnecting and reconnecting).

  Workaround: Depending on the location of the restarted server (central or regional), restart the kubemaster and kubeminions (msvms). [CXU-16161]

- Deletion of an SLA profile fails. However, the corresponding SD-WAN policies are deleted successfully.

  Workaround: None. [CXU-16162]

- In an HA setup, if one of the infrastructure microservices VMs goes down and you trigger report generation, reports are not generated.

  Workaround: Ensure that the VM that was down comes back up and then trigger report generation. [CXU-16222]

- After creating a new spoke site, when you configure and activate the site, site provisioning sometimes fails.

  Workaround: Retry the failed job. [CXU-16349]

- The throughput graphs in the WAN page of the spoke site might not reflect the actual throughput seen in the CPE interfaces running the vSRX image 15.1X49-D123.2

  Workaround: None. [CXU-16471]

- When a redirect server is used for ZTP, the **PHS bootstrap complete** message is not sent to the CSO regional server, because of which ZTP fails.

  Workaround: None. [PR 1315524]

## Resolved Issues

The following issues are resolved in Juniper Networks Cloud CPE Solution Release 3.1.2.

- You cannot edit the settings for a customer in Administration Portal. [CXU-4538]
- You can view only one network service at a time on a CPE device in Customer Portal. [CXU-4551]
- During activation, the NFX250 device reboots and requests that you enter the activation code twice, because there is no default value for the HugesPages count in the Linux Kernel on the device. [CXU-5601, PR1254219]
- After you install CSO, the username and password that CSO uses to access Contrail Analytics might not match the corresponding username and password in Contrail Analytics. [CXU-5873]
- The configuration for a CPE device might not be removed when you deactivate the device in Administration Portal. [CXU-6059]
- You cannot edit the deployment type (vCPE-Only or uCPE-Only) in a request that you create with Network Service Designer. [CXU-6474]
- Traffic not classified by SD-WAN policies follows the default routing behavior. [CXU-10716]
- In some cases, an operation might fail, with the **Authentication required** error message, because of an expired token. [CXU-10809]
- In some cases, when multiple system log messages (syslogs) or queries are being processed, the Contrail Analytics Node (CAN) crashes (Docker restarts). [CXU-10838]
- When a session is transferring more than 2 GB of data, the unit of throughput is incorrectly reported as terabits per second (Tbps) in the CSO GUI. [CXU-11556]
- When you remove a cloud hub from a tenant, the corresponding router is removed from the All Tenants scope. [CXU-11796]

- If the central infrastructure node that created the SAML 2.0 session goes down and you log out of the CSO GUI, the SAML 2.0 session logout fails. [CXU-11867]
- When the deployment of a LAN segment on a device fails, the device status is changed from **PROVISIONED** to **PROVISION_FAILED**. However, when you redeploy the LAN segment and the deployment is successful, the device status is not changed to **PROVISIONED**. Therefore, when you attempt to deploy an SD-WAN or a firewall policy on the device, the deployment fails with the error message **[get_policy_info_list: method execution failed]Site/ Device information not found**. [CXU-11874]
- In some cases, traffic fails to flow over an overlay tunnel. [CXU-11921]
- When you log in to CSO as a Tenant Administrator user, the **Configure Site** workflow is not available. [CXU-11922]
- ZTP activation of an SRX Series device by using the phone home client (PHC) fails. [CXU-11926]
- You see an import error message when you use the license tool, because the **tssmclient** package is missing from the license tool files. [CXU-12054]
- When you deploy a firewall policy, the deployment fails with the message **Fail to invoke mapper to create snapshot with reason null**. [CXU-12151]
- When you onboard an NFX series device by using the default configuration, in some cases there is a connectivity issue between Juniper Device Manager (JDM) and Junos Control Plane (JCP). [CXU-12396]
- By default, CSO uses Heat V2 APIs to bring up network services on Contrail. However, because a bug in CSO, the policy configured does not exchange routes and therefore traffic does not flow through the service chain. [CXU-12863]
- If you modify the default configuration of an SRX340 device to get the IP address by using DHCP, the device activation fails. [CXU-13446]
- During the failover of a link between an NFX Series spoke device and a vSRX hub device, the BGP session goes down even though the virtual route reflector (VRR) is reachable. [CXU-13517]
- If you configure DHCP on an NFX Series or an SRX Series spoke device, in some cases, the spoke might fail to establish a connection to CSO and might fail to send syslog messages to the CAN. [CXU-13567]
- In some cases, one or more security management microservices take more than 20 minutes to come up and are stuck in the same state. [CXU-13726]
- Reports are not generated in HA deployment scenarios. [CXU-14039]
- If you try to activate an SRX300 CPE device by using the redirect server, the phone home activation process does not start. [CXU-14162]
- The reverse path taken by traffic on the hub is different from the forward path. [CXU-14330]
- The Contrail Collector service (Contrail Analytics Release 4.0.2) crashes several times intermittently. However, the service recovers after some time. [CXU-15802]

- BGP routes are withdrawn 10–15 minutes after a BGP session goes down, even though the **hold-timer** value is set to 65,535. When BGP routes are withdrawn, traffic flow is impacted. [PR 1312702]
- Sometimes the IBGP session with the VRR is broken temporarily and comes back up automatically. During that period there might be issues related to link switch. [PR 1312863]
- Some variables in the NSC installer package do not have the correct values. [CXU-14981]

## Documentation Updates

This section lists the errata and changes in the Cloud CPE Solution Release 3.1.2 documentation:

- On the **Connectivity Requirements** tab of the Add On-Premise Site page (**Sites > Site Management > Add > On-Premise Site** ), the help for the **Type** field indicates that you can select the type of WAN link. However, this field only displays the link type of the WAN link (MPLS or Internet), which is specified in the device profile, and cannot be modified.
- **Configuring devices from the POPs landing page**—From Cloud CPE Solution Release 3.1 onward, you can configure devices from the POPs page as follows:

  1. Select **Resources > POPs > *Pop-Name***.

     The *Pop-Name* page appears.

  2. Click the **Routers** tab.

  3. Select the device that you want to configure and click the **Configure Device** button.

     The Stage 2 Config page appears. This page is dynamically rendered based on stage-2 configuration specified in the device profile.

  4. Enter the configuration data on the page.

  5. Click **Save** to save the configuration.

     A confirmation message is displayed and the deployment status changes to **pending deployment**.

  6. Click **Deploy** to save and deploy the configuration.

     A confirmation message is displayed indicating that the job is created and subsequently that the job was successful. You can click **Deploy History** to view the job logs.

  7. Click **Cancel** to go back to the *Pop-Name* page.

## 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 rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/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 http://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: http://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: http://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/

- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

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

## Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- 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 http://www.juniper.net/support/requesting-support.html.

## Revision History

04 January 2018—Revision 4

20 December 2017—Revision 3

15 December 2017—Revision 2

13 December 2017—Revision 1, Cloud CPE Solution Release 3.1.2