Cloud CPE Solution Release Notes

Release 3.1
20 December 2017
Revision 10

These Release Notes accompany Release 3.1 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

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

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 is a secure software-defined WAN (SD-WAN) solution that builds on the capabilities of the existing Cloud CPE Solution. 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

Copyright © 2017, Juniper Networks, Inc.

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: **In Cloud CPE Solution Release 3.1, 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.

Copyright © 2017, Juniper Networks, Inc.

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 Instructions

Two CSO installers are available for Release 3.1:

- **A plug-and play-installer for a demo environment**—This installer offers a snapshot version of CSO that you can install in a few minutes by copying the installer tar file to the CSO server, expanding the tar file, and running the installer. The installation settings, such as the VM resources, are fixed and you cannot modify them.
- **A full-version installer**—You must use this installer for a production environment and you can use it for install a demo environment if you want to customize the installation settings.

  After copying the installer tar file to the CSO server and expanding the file, you provision the VMs, and run a script to create a file of settings for the installation. You then run the installer, which takes approximately an 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.

Copyright © 2017, Juniper Networks, Inc.

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.

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 3.1.

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 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
- 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

Copyright © 2017, Juniper Networks, Inc.

New and Changed Features

This section describes the new features or enhancements to existing features in Cloud CPE Solution Release 3.1.

SD-WAN

- **Support for managing and deploying SD-WAN policies**—From Cloud CPE Solution Release 3.1 onward, you can define and deploy SD-WAN policies for applications or application groups based on service-level agreement (SLA) requirements. SD-WAN policies help in optimum utilization of the WAN links and efficient distribution of traffic. Policies are applied at site, site group, or department level. You can also schedule your policy deployment for a later date and time.
- **Support for creating SLA profiles for applications and application groups**—From Cloud CPE Solution Release 3.1 onward, you can create tenant-level SLA profiles and associate the SLA profiles with applications or application groups. (In this context, the term *applications* refers to applications that do not need a Secure Sockets Layer (SSL) inspection.) An SLA profile consists of defined target metrics, which include the following:
  - Throughput, latency, packet loss, jitter, and delay
  - An assured class of service
  - Upstream and downstream rates for its applications
- **Support for creating hub sites**—Cloud CPE Solution Release 3.1 supports the creation of hub sites for tenants. You create a hub site by selecting the site type as an on-premise hub or during the creation of the site.
- **Support for four SD-WAN-enabled links per site**—Starting with Cloud CPE Solution Release 3.1, you can configure up to four WAN links per site that support SD-WAN. You can configure these links as MPLS or broadband links. In releases before Cloud CPE Solution Release 3.1, you can configure only two WAN links.
- **Support for monitoring SLA performance**—Cloud CPE Solution Release 3.1 supports the SLA-performance monitoring of tenants, sites, and applications that have met and those that have not met their defined SLA values in a specified period.
- **Support for monitoring SD-WAN events**—Cloud CPE Solution Release 3.1 supports the monitoring of SD-WAN events. SD-WAN events are triggered when the SLA requirements for a site are not met on its designated WAN link and the site switches WAN links to meet its SLA requirements.
- **Support for SD-WAN alert definitions**—From Cloud CPE Solution Release 3.1 onward, you can create, edit, and delete SD-WAN alert definitions. An alert definition consists

Copyright © 2017, Juniper Networks, Inc.

of data criteria for triggering alerts that warn you about issues in your SD-WAN environment. Alert definitions also define the necessary action required to resolve issues based on the severity of the alert.

- **Support for creating site groups**—Cloud CPE Solution Release 3.1 enables you to create site groups, which are a collection of one or more sites, for policy management. A site group enables you to apply a policy to all the sites in a group simultaneously. You create site groups from the Create Site Group page (**Site** > **Site Groups** > **Create Site Group**).
- **Viewing the bandwidth capacity of a WAN link**—From Cloud CPE Solution Release 3.1 onward, you can view the maximum bandwidth capacity of a WAN link. To view bandwidth capacity of a WAN link, hover over the WAN link connected to a site on the **WAN** tab of the *Site-Name* page (**Sites > Site Management > *Site-Name* > WAN**).
- **Support for grouping LAN segments into departments**—From Cloud CPE Solution Release 3.1 onward, you can group LAN segments within a site into departments. You use departments to apply specific policies to LAN segments that are members of a department. You can create, view, edit, or delete departments from the Departments page (**Configuration > Shared Objects > Departments**).

Security Management

- **Support for NAT policy management**—Cloud CPE Solution Release 3.1 enables you to create, modify, and delete Network Address Translation (NAT) policies and rules. In Cloud CPE Solution Release 3.1, only source-NAT and static-NAT management are supported.
- **Support for intent-based firewall policy**—Cloud CPE Solution Release 3.1 enables you to create, modify, and delete firewall intents associated with a firewall policy. Firewall policies are intent-based, which means that they can incorporate both Transport Layer (Layer 4) and Application Layer (Layer 7) application firewall constructs in a single intent. In addition, policies are automatically assigned to devices based on the endpoints chosen in the definition of the intent, and do not need to be assigned to specific devices manually. Firewall intents consist of source and destination endpoints; the endpoints can be applications (L7), sites, IP addresses, IP address groups, site groups, departments, and services. (In this context, the term *applications* refers to applications that do not need an SSL inspection.)
- **Support for schedules in firewall policy**—From Cloud CPE Solution Release 3.1 onward, you can create, modify, clone, and delete firewall policy schedules. A schedule enables you to run an intent for a specified period either on a one-time or on a recurring basis based on how the schedule is created.
- **Support for services in firewall and NAT policies**—From Cloud CPE Solution Release 3.1 onward, you can create, modify, clone, and delete services or service groups.

  A service refers to an application on a device, such as Domain Name System (DNS). Services are based on protocols and ports used by an application, and when added to a policy, a configured service can be applied across all devices. The protocols available to create a service include TCP, UDP, SUN-RPC, MS-RPC, ICMP, ICMPv6, and so on.

Copyright © 2017, Juniper Networks, Inc.

You can combine services together to form a service group. Service groups are useful when you want to apply the same policy to multiple services because by doing this you can create and work with fewer policies.

- **Support for security dashboards**—From Cloud CPE Solution Release 3.1 onward, a security dashboard page displays information such as top events, top denials, top applications, top source and destination IP addresses, top traffic, and top infected hosts.
- **Support for application visibility**—From Cloud CPE Solution Release 3.1, you can view information on bandwidth consumption, session establishment, and the risks associated with your network applications. Analyzing your network applications provides useful security management information, such as applications that can lead to data loss, bandwidth overconsumption, time-consuming applications, and personal applications that can elevate business risks.
- **Support for security alerts**—From Cloud CPE Solution Release 3.1, you can create, edit, and delete security alert definitions. Alerts are used to notify administrators about significant events within the system and warn them about problems in your monitored environment.

  An alert definition consists of data criteria for triggering an alert. An alert is triggered when the event threshold exceeds the defined data criteria.
- **Support for security events and system log messages**—From Cloud CPE Solution Release 3.1 onward, you can view security events associated with firewall, Web filtering, IPsec VPN, content filtering, antispam, antivirus, and IPS events.

  Security events include the system log messages of the device and critical information such as the number of events, virus instances found, interfaces that are down, attacks, CPU spikes, reboots, and sessions.
- **Ability to collect and view device events**—From Cloud CPE Solution Release 3.1 onward, you can troubleshoot a device by using device events. Device events include the following:
  - Routine operations—for example, user login into the configuration database
  - Failure and error conditions—for example, failure to access a configuration file
  - Emergency or critical conditions—for example, device power failure due to excessive temperature
- **Support for generating reports**—From Cloud CPE Solution Release 3.1 onward, you can generate reports to view the summary of network activity and overall network status of CPE devices. Using reports, you can:
  - Create, edit, delete, and clone reports, preview reports in PDF, and send reports by e-mail
  - Schedule reports based on the defined filters
  - Schedule reports based on the available default reports
  - Generate reports with multiple sections, where each section has its own criteria

Copyright © 2017, Juniper Networks, Inc.

- **Support for application signature management**—Cloud CPE Solution Release 3.1 enables you to create, modify, clone, delete, and view custom application signature groups, and view predefined application signatures.
- **Support for active database**—From Cloud CPE Solution Release 3.1, you can download and install the application firewall signature database on CPE devices. This database includes signature definitions of attacks and applications that can be used to identify applications for tracking firewall policies and quality of service prioritization.
- **Support for address management**—Cloud CPE Solution Release 3.1 enables you to create, edit, and delete addresses and address groups. Addresses and address groups are used in firewall and NAT services.

Unified Portal

- **Support for a unified Administration and Customer Portal**—Cloud CPE Solution Release 3.1 now supports a unified portal for both service provider users and tenant users and for the services managed and consumed by the administrators and tenants.

  The unified portal contains the features of vCPE, uCPE, and SD-WAN for both the Administration and Customer Portals; enforces RBAC, which prevents tenants from accessing administrator data; and supports different backend authentication methods for service provider users and tenant users.

- **Support for SSO authentication**—Cloud CPE Solution Release 3.1 supports single sign-on (SSO) authentication for the unified portal. You can authenticate and authorize users by using one of the following authentication methods:

  - Local—User accounts are maintained locally in Contrail Service Orchestration (CSO), and users are authenticated and authorized by CSO.
  - Authentication by using an SSO server—User accounts are maintained in the service provider's SSO server, but authorization information is stored in CSO. Users are authenticated by using the SSO server.
  - Authentication and authorization by using an SSO server—User accounts and user roles are maintained in the service provider's SSO server. Users are authenticated by using the SSO server and authorized by CSO by using Security Assertion Markup Language (SAML) attributes.

  You can configure one SSO server for a service provider and another for all its tenants. The following SSO servers are supported:

  - Okta
  - OneLogin

- **Support for role-based access control (RBAC)**—Cloud CPE Solution Release 3.1 enables you to add, view, edit, and delete tenant and service provider users. The following roles are available:

Copyright © 2017, Juniper Networks, Inc.

  - **MSP Administrator**—Users with the MSP Administrator role have full access to the Administration Portal UI or API capabilities. They can use the UI or APIs to add one or more users with MSP Administrator or MSP Operator roles, onboard tenants, and add the first tenant administrator during the onboarding process. They can also add tenant administrators or operators by switching the scope to a specific tenant.
  - **MSP Operator**—Users with the MSP Operator role have read-only access to the Administration Portal and APIs.
  - **Tenant Administrator**—Users with the Tenant Administrator role have full access to the Customer Portal and APIs. They can add one or more users with Tenant Administrator or Tenant Operator roles.
  - **Tenant Operator**—Users with the Tenant Operator role have read-only access to the Customer Portal and APIs.

  You can assign MSP roles to service provider users and Tenant roles to tenant users.

- **Security enhancements related to login credentials**—Cloud CPE Solution Release 3.1 includes the following security-related enhancements:
  - You can use the **Forgot Password** link on the Login page to reset your password.
  - After you log in for the first time, you are prompted to change your password.

    Passwords must conform to the password rules specified in the UI.
  - You can specify the duration (in days) after which the password expires and must be changed. Users who do not change their password before the specified duration elapses are automatically logged out.
  - Your account is locked after five consecutive unsuccessful login attempts.

- **Support for switching the tenant scope**— From Cloud CPE Solution Release 3.1 onward, Administration Portal users can change the tenant scope from all tenants to a specific tenant by using the tenant switcher displayed on the banner. When you switch scope from all tenants to a specific tenant, the menu and pages displayed are almost the same as those displayed for Customer Portal users, with some additional actions visible to the Administration Portal users. When you switch back to the **All Tenants** scope, the menu and pages for the Administration Portal are displayed.

Copyright © 2017, Juniper Networks, Inc.

Miscellaneous

- **Full-service chaining support on the NFX Series CPE device**—From Cloud CPE Solutions Release 3.1 onward, Network Service Controller (NSC) offers an European Telecommunications Standards Institute (ETSI) NFV-compliant virtualized infrastructure manager (VIM) to instantiate VNFs on NFX Series CPE devices.

  This instantiation method provides optimal activation of third-party VNFs and full-service chaining on NFX Series CPE devices. Consequently, all traffic from a LAN connected to an NFX Series CPE device first traverses the VNF on the device, and then passes through the vSRX gateway device, and finally exits through the WAN link.

- **Activation process for CPE devices**—From Cloud CPE Solution Release 3.1 onward, you can activate SRX300 Services Gateway and NFX250 Network Services Platform devices in the following ways:

  - By connecting a computer to the LAN port of the device and entering the activation code through your browser
  - By specifying the activation code in Customer Portal

- **Support for upgrading and deploying an image**—From Cloud CPE Solution Release 3.1 onward, you can upgrade and deploy an image on a single device or multiple devices on a per-site basis or across all sites of a tenant. A device can be a physical network function (PNF) or a virtual network function (VNF).

  You can also schedule the deployment of images.

- **Support for configuration deployment**—From Cloud CPE Solution Release 3.1 onward, you can deploy SD-WAN and security policies immediately or schedule the deployment for a later date and time.

- **Support for viewing policies at the device level and site level**—From Cloud CPE Solution Release 3.1 onward, you can view the policies assigned to a CPE device (**Resources** > **Devices** > ***Device-Name*** > **Policies**) and the policies assigned to a tenant site (**Sites** > **Site Management** > ***Site-Name*** > **Policies**). You can view the following information about policies:

  - List of all policies applicable to a tenant or site
  - Details about the tenant user who last updated the policy
  - Time when the policy was last updated
  - Deployment status of the policy
  - Number of intents applicable to the site compared to the total number of intents applicable to the tenant

- **Support for job management**—From Cloud CPE Solution Release 3.1 onward, you can view the list of jobs that are currently running and the jobs that are scheduled to run

Copyright © 2017, Juniper Networks, Inc.

later. You can also specify whether you want to run a job immediately or schedule it for a later date and time.

- **Customer Portal Dashboard**—From Cloud CPE Solution Release 3.1 onward, you can view a customized view of network services by using the widgets on the user-configurable Dashboard page.

  You can drag the widgets from the carousel at the top of your dashboard to your workspace, where you can add, remove, and rearrange them. The dashboard automatically adjusts the placement of the widgets to dynamically fit on your browser window without changing their order. You can manually reorder the widgets by using the drag and drop option. In addition, you can press and hold the top portion of the widget to move it to a new location.

Servers, Software, and Network Devices Tested

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 13 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

Copyright © 2017, Juniper Networks, Inc.

Software Tested for COTS Servers

Table 3 on page 14 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	• Centralized deployment: Contrail Cloud Platform Release 3.0.2, or VMware ESXi Version 5.5.0 • Distributed deployment: 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 Cloud Platform Release 3.0.2 with Heat v2 APIs
Contrail Analytics	Contrail Release 4.0.0.15

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 14.
- The software described in Table 5 on page 15.

  You must use these specific versions of the software for the Cloud CPE Solution Release 3.1.

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: • 48 10/100/1000-Gigabit Ethernet interfaces • 4 built-in 10-Gigabit Ethernet SFP transceiver interfaces	1

Copyright © 2017, Juniper Networks, Inc.

Table 4: Network Devices Tested for the Contrail Cloud Platform *(continued)*

Function	Device	Model	Quantity
Data switch	Juniper Networks QFX Series Switch	QFX 5100-48S-AFI switch with: • 48 SFP+ transceiver interfaces • 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 EX Series switch	Junos OS Release 12.3R10
Operating system for QFX Series switch	Junos OS Release 13.2X51-D38
Hypervisor on Contrail Service Orchestration servers	Contrail Release 3.0.2 with OpenStack (Liberty or Kilo), or VMware ESXi Version 5.5.0
Element management system software	EMS microservice Junos Space Network Management Platform Release 15.1R1 (See *VNFs Supported by the Cloud CPE Solution* for VNFs that require this product)
Software defined networking (SDN) for a centralized deployment	Contrail Cloud Platform Release 3.0.2
Contrail Analytics	Contrail Release 4.0.0.15
Virtualized infrastructure manager (VIM) and virtual machine (VM) orchestration	OpenStack Liberty or Kilo
Authentication and Authorization	OpenStack Liberty or Kilo
Network Functions Virtualization (NFV)	Contrail Service Orchestration Release 3.1

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 16.
- The software described in Table 7 on page 16.

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

Copyright © 2017, Juniper Networks, Inc.

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 • MX80 or MX104 router with MX-MIC line card • 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 • SRX4100 Services Gateway • SRX4200 Services Gateway	—
CPE device	• NFX250 Series Network Services Platform • SRX Series Services Gateway • vSRX on an x86 server	• NFX250-LS1 device • NFX250-S1 device • NFX250-S2 device • SRX300 Services Gateway • SRX320 Services Gateway • SRX340 Services Gateway • SRX345 Services Gateway • 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
Contrail Analytics	Contrail Release 4.0.0.15
NFX Software	Junos OS Release 15.1X53-D47
Routing and Security for NFX250 device	vSRX KVM Appliance 15.1X49-D101
Operating system for vSRX used as a CPE device on an x86 server	vSRX KVM Appliance 15.1X49-D101
Operating system for SRX Series Services Gateway used as a CPE device	Junos OS Release 15.1X49-D101
Operating system for MX Series router used as PE router	Junos OS Release 16.1R3.00
Operating system for SRX Series Services Gateway used as a hub device for an SD-WAN implementation	Junos OS Release 15.1X49-D101

Copyright © 2017, Juniper Networks, Inc.

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

Minimum Hardware Requirements for the Cloud CPE Solution

Table 2 on page 13 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 17 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		
Storage	Greater than 1 TB of one of the following types: • Serial Advanced Technology Attachment (SATA) • Serial Attached SCSI (SAS) • Solid-state drive (SSD)	Greater than 1 TB of one of the following types: • SATA • SAS • SSD
CPU	One 64-bit dual processor, type Intel Sandybridge, such as Intel Xeon E5-2670v3 @ 2.4 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)		

Copyright © 2017, Juniper Networks, Inc.

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

Function	Demo Environment (no HA)	Trial HA Environment
Number of nodes or servers	1 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 • 3 nodes for Contrail controller and analytics • 1–4 Contrail compute nodes
vCPUs per node or server	16	48
RAM per node or server	64 GB	256 GB

Table 9 on page 18 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: • SATA • SAS • 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	

Copyright © 2017, Juniper Networks, Inc.

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

Server Function	Values
Number of nodes or servers for a non-HA environment	3 • 1 central server • 1 regional server • 1 Contrail Analytics server
Number of nodes or servers for an HA environment	9 • 3 central servers • 3 regional servers • 3 Contrail Analytics servers
vCPUs per node or server	48
RAM per node or server	256 GB
Contrail Analytics Server for a Distributed Deployment	
Number of nodes or servers	1
vCPUs per node or server	48
RAM per node or server	256 GB
Contrail Cloud Platform for a Centralized Deployment	
Number of nodes or servers	4–28 • 3 nodes for Contrail controller and Contrail Analytics • 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.

Table 10 on page 20 shows complete details about the VMs required for a demo environment. HA is not included with the demo environment.

Copyright © 2017, Juniper Networks, Inc.

Table 10: Details of VMs for a Demo Environment

Name of VM	Components That Installer Places in VM	Resources Required	Ports to Open
csp-installer-vm	—	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm	Third-party applications used as infrastructure services	• 6 CPU • 32 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm	All microservices, including GUI applications	• 6 vCPUs • 40 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm	Third-party applications used as infrastructure services	• 6 vCPUs • 32 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.
csp-regional-msvm	All microservices, including GUI applications	• 6 vCPUs • 32 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb	Load balancer for device to Fault Management Performance Management (FMPM) microservice connectivity	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy virtualized network functions (VNFs) that use this EMS	• 4 vCPUs • 32 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.
csp-contrailanalytics-1	Contrail Analytics for a distributed deployment For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform.	• 8 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-vrr-vm	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.

Table 11 on page 21 shows complete details about VMs and microservice collections required for a production environment without HA.

Copyright © 2017, Juniper Networks, Inc.

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 • 64 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm	Third -party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm	Third -party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-msvm	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb	Load balancer for device to microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-elkvm	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-elkvm	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-vrr-vm	VRR	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.

Table 12 on page 22 shows complete details about the VMs for a trial HA environment.

Copyright © 2017, Juniper Networks, Inc.

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 • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm1	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm2	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm3	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-lbvm1	Load-balancing applications	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-lbvm2	Load-balancing applications	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm1	All microservices, including GUI applications	• 8 CPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm2	All microservices, including GUI applications	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm1	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm2	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm3	Third-party applications used as infrastructure services	• 8 vCPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.

Copyright © 2017, Juniper Networks, Inc.

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-msvm1	All microservices, including GUI applications	• 8 CPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-msvm2	All microservices, including GUI applications	• 8 CPUs • 48 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-lbvm1	Load-balancing applications	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-lbvm2	Load-balancing applications	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-contrailanalytics-1	Contrail Analytics for a distributed deployment For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform.	• 16 vCPUs • 64 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-vrr-vm	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage	See Table 14 on page 27.

Table 13 on page 24 shows complete details about VMs and microservice collections required for a production environment with HA.

Copyright © 2017, Juniper Networks, Inc.

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 • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm1	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm2	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-infravm3	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-lbvm1	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-lbvm2	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-lbvm3	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm1	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm2	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-central-msvm3	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm1	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.

Copyright © 2017, Juniper Networks, Inc.

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-infravm2	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-infravm3	Third-party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-msvm1	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-msvm2	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-msvm3	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage	See Table 14 on page 27.
csp-regional-lbvm1	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-lbvm2	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-lbvm3	Load-balancing applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-elkvm1	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-central-elkvm2	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.

Copyright © 2017, Juniper Networks, Inc.

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-elkvm3	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-elkvm1	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-elkvm2	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-elkvm3	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-regional-sblb3	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.
csp-vrr-vm	Virtual route reflector (VRR)	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage	See Table 14 on page 27.

Table 14 on page 27 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.

Copyright © 2017, Juniper Networks, Inc.

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
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

Copyright © 2017, Juniper Networks, Inc.

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

Port Number	CSO Communication Type	Port Function
5671	Internal	RabbitMQ SSL listener
5672	Internal	RabbitMQ client
6000	Internal	Swift Object Server
6001	Internal	Swift Container Server
6002	Internal	Swift Account Server
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

Copyright © 2017, Juniper Networks, Inc.

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

Port Number	CSO Communication Type	Port Function
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 29.

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-D101	• Network Address Translation (NAT) • Demonstration version of Deep Packet Inspection (DPI) • Firewall	• Centralized deployment • 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 • 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:

Copyright © 2017, Juniper Networks, Inc.

- 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 to access the Contrail Service Orchestration (CSO) GUIs.**

Table 16 on page 30 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	https://*central-IP-Address* where: *central-IP-Address*—IP address of the VM that hosts the microservices for the central POP For example: *https://192.0.2.1*	Specify the OpenStack Keystone username and password. The default username is **cspadmin** and the default password is **passw0rd**.

Copyright © 2017, Juniper Networks, Inc.

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

GUI	URL	Login Credentials
Customer Portal	Same as the URL used to access the Administration Portal	Specify the credentials when you create the Customer either In Administration Portal or with API calls.
Designer Tools—Log into Network Service Designer and click the menu in the top left of the page to access the other designer tools.	https://*central-IP-Address*:83 where: *central-IP-Address*—IP address of the VM that hosts the microservices for the central POP For example: *https://192.0.2.1:83*	Specify the OpenStack Keystone username and password. The default username is **cspadmin** and the default password is **passw0rd**.
Service and Infrastructure Monitor This tool provides monitoring and troubleshooting for network services in a hybrid WAN deployment.	http://*ha-proxy-IP-Address*:1947/icingaweb2 where: *ha-proxy-IP-Address*—IP address of HA proxy. • For a non-high availability (non-HA) deployment, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: *http://192.0.2.1:1947/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 This tool provides a visual representation of log files. You can use it to monitor: • Network services in a central or regional POP • Microservices in the deployment	http://*infra-vm-IP-Address* | *ha-proxy-IP-Address*:5601 where: *infra-vm-IP-Address*—IP address of the VM that hosts the infrastructure services for a central or regional POP. Use this option to monitor network services. *ha-proxy-IP-Address*—IP address of high availability (HA) proxy. Use this option to monitor the microservices. • For a non-HA deployment, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: *http://192.0.2.2:5601*	Login credentials are not needed.

Copyright © 2017, Juniper Networks, Inc.

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

GUI	URL	Login Credentials
Grafana and Prometheus 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 • Grafana—*ha-proxy-IP-Address*:3000 where: *ha-proxy-IP-Address*—IP address of HA proxy • For a non-HA deployment, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: *http://192.0.2.2:30900*	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.

Copyright © 2017, Juniper Networks, Inc.

- 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 as follows:
        - 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/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 both heat resources directories.
     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 you use Customer Portal to activate a network service on a network link, you must configure the following settings for each VNF in the network service:
  - Hostname
  - NTP server
  - DNS name server

Copyright © 2017, Juniper Networks, Inc.

- When you use Mozilla Firefox to access the Contrail Service Orchestration (CSO) GUIs, a few pages do not work as expected. Therefore, we recommend that you use Google Chrome version 60 or later to access the CSO GUIs.
- 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.

  [CXU-1242]

- 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.

  [CXU-3630]

- 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.

  [CXU-3631]

- 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**.

  [CXU-5990]

- You can use the logs on an NFX250 device to review the status of the device's activation through Administration Portal.

  [CXU-6188]

- In Cloud CPE Solution Release 3.1, intrusion prevention system (IPS) is not supported. Therefore, in the IPS report, the attack name from the IPS signatures is displayed as **UNKNOWN**.

Copyright © 2017, Juniper Networks, Inc.

- In Cloud CPE Solution Release 3.1, 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, 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.
- High availability for the Kubernetes master node is not supported. If the VM goes down, bring it up immediately for CSO to be operational.
- 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** , 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.

- In Cloud CPE Solution Release 3.1, 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

Copyright © 2017, Juniper Networks, Inc.

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, 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, 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.

Copyright © 2017, Juniper Networks, Inc.

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 login 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:

Copyright © 2017, Juniper Networks, Inc.

```
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.

- 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]

- You cannot edit the settings for a customer in Administration Portal.

Copyright © 2017, Juniper Networks, Inc.

Workaround: Use Administration Portal to import a JSON file that contains the correct data for the customer. **[CXU-4538]**

- You can only view 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.

  Workaround: Before you activate the NFX250 device, specify the HugePages count on the device and reboot it.

  [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.

  Workaround: Complete the following actions:

  1. Log in to the CSO central infrastructure VM as root.

  2. View the username that CSO uses to access Contrail Analytics.

     ```
     root@host:~/# etcdctl get /csp/infra/contrailanalytics/queryuser
     <username>
     ```

     <username> is the actual value that the query returns.

  3. View the password that CSO uses to access Contrail Analytics.

     ```
     root@host:~/# etcdctl get /csp/infra/contrailanalytics/querypassword
     <password>
     ```

     <password> is the actual value that the query returns.

  4. Log in to the CSO regional infrastructure VM as root.

  5. Repeat Steps 4 and 5 to verify the username and password on the regional host.

     If the username and password on both the central and regional infrastructure VMs match the values configured in Contrail Analytics, you do not need to take further action. The default username configured in Contrail Analytics is admin and the default password is contrail123.

  If the username and password on the central or regional infrastructure VMs do not match the values in Contrail Analytics, update them as follows:

  1. Log in to the appropriate CSO infrastructure VM as root.

  2. Update the username and password with the values configured in Contrail Analytics.

Copyright © 2017, Juniper Networks, Inc.

```
root@host:~/# etcdctl set /csp/infra/contrailanalytics/queryuser
contrail-analytics-username
root@host:~/# etcdtl set /csp/infra/contrailanalytics/querypassword
<contrail-analytics-password>
```

contrail-analytics-username is the actual username, and *<contrail-analytics-password>* is the actual password.

[CXU-5873]

- The configuration for a CPE device might not be removed when you deactivate the device in Administration Portal.

  Workaround: To deactivate the CPE device, first delete the configuration from the CPE device with Customer Portal, and then deactivate the device with 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.

  Workaround: Create a new request. **[CXU-6474]**

- 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]

- In some cases, when multiple system log messages (syslogs) or queries are being processed, the Contrail Analytics Node (CAN) crashes (Docker restarts).

  Workaround: Do the following:

  1. Log in to the CAN as **root**.
  2. Restart the **analytics** and **controller** Docker containers.
  3. Log in to the **analytics** Docker container by using the **docker exec -it** *analytics-docker-id* **bash** command, where *analytics-docker-id* is the ID of the **analytics** Docker container.
  4. Verify that the following entries are present in the **/etc/contrail/contrail-collector.conf** file:

     [STRUCTURED_SYSLOG_COLLECTOR]

     # TCP & UDP port to listen on for receiving structured syslog messages

     port=3514
  5. Verify that the following entries are available in the **/etc/contrail/contrail-analytics-api.conf** file:

     aaa_mode=no-auth

Copyright © 2017, Juniper Networks, Inc.

6. If either of the entries is not present, restart the services by executing the following commands:

   service contrail-collector restart

   service contrail-analytics-api restart

7. Log in to the **controller** Docker container by using the **docker exec -it *controller_docker_id* bash** command, where *controller_docker_id* is the ID of the controller Docker container.

8. Verify that the following entries are available in the **/usr/lib/python2.7/dist-packages/vnc_cfg_api_server/vnc_cfg_api_server.py** file:

   (under) def __init__(self, args_str=None):

   self.aaa_mode = 'no-auth'

[CXU-10838]

- 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]

- Traffic not classified by SD-WAN policies follows the default routing behavior.

  Workaround: Configure SD-WAN policies for all the traffic originating from the LAN. [CXU-10716]

- 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]

- In some cases, an operation might fail with the **Authentication required** error message because of an expired token.

  Workaround: Retry the operation. [CXU-10809]

- 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.

Copyright © 2017, Juniper Networks, Inc.

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 remove a cloud hub from a tenant, the corresponding router is removed from the All Tenants scope.

  Workaround: None. [CXU-11796]

- 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

 Copyright © 2017, Juniper Networks, Inc.

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 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 log out fails.

  Workaround:

  1. Reload the CSO login page.
  2. Enter the username and press the Tab button on your keyboard or click the mouse outside the username field.

     You are automatically logged in to the CSO GUI.
  3. Click the **Logout** link in the banner to log out.

Copyright © 2017, Juniper Networks, Inc.

You are logged out of the CSO GUI and the SAML 2.0 session.

[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**".

  Workaround: None. [CXU-11874]

- 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.

Copyright © 2017, Juniper Networks, Inc.

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, 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]

- In some cases, traffic fails to flow over an overlay tunnel.

  Workaround: Reboot the vSRX or SRX Series device to ensure that the traffic flows normally. [CXU-11921]

- When you log in to CSO as a Tenant Administrator user, the **Configure Site** workflow is not available.

  Workaround: Log in as an MSP Administrator user and switch to the tenant for which you want to configure the site. The **Configure Site** workflow becomes available. [CXU-11922]

Copyright © 2017, Juniper Networks, Inc.

- ZTP activation of an SRX Series device by using the phone home client (PHC) fails.

  Workaround: If the activation fails with the error message **Ztp-activation finished incomplete for ems-device** and no other error messages are present in the activation logs, then the MSP Administrator (in the All Tenants scope) can retry the activation job by navigating to the Jobs page (**Monitor > Jobs**), selecting the failed job, and clicking the Retry Job button. [CXU-11926]

- 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
```


Copyright © 2017, Juniper Networks, Inc.

```
                                        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
      ```

Copyright © 2017, Juniper Networks, Inc.

```
 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]

- You see an import error message when you use the license tool, because the **tssmclient** package is missing from the license tool files.

  Workaround: complete the following procedure:

  1. Copy the **tssmclient.tar.gz** file from http://www.juniper.net/support/downloads/?p=cso#sw to the

Copyright © 2017, Juniper Networks, Inc.

***csoVersion*/licenseutil/csoclients/** directory on the installer VM, where *csoVersion* is the name of the installer directory created when you extract the TAR file for the installation package.

2. Access the ***csoVersion*/licenseutil/csoclients/** directory and extract the tssmclient directory from the TAR file.

   For example, if the installer directory is called csoVersion:

   ```
   root@host:~/# cd csoVersion/licenseutil/csoclients/
   root@host:~/csoVersion/licenseutil/csoclients# tar xvzf tssmclient.tar.gz
   ```

   This command replaces the **tssmclient** folder and its contents.

3. Run the license tool using the instructions in the Deployment Guide.

[CXU-12054]

- 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.

Copyright © 2017, Juniper Networks, Inc.

[CXU-12107]

- Some variables in the CSO and NSC installer packages do not have the correct values.

  Workaround: After you extract the TAR file for the installation package, access the installer directory, which has the same name as the installer package, and execute the following sed command:

  For example, if the name of the installer package is **csoVersion.tar.gz**:

  ```
  root@host:~/# cd csoVersion
  root@host:~/csoVersion# sed -i s@gcr.io/google_containers@csp-installer:10000@g
  salt/file_root/kubeminion/files/manifests/kube-proxy.yaml;sed -i s@"timedatectl
  set-ntp"@ntpdate@g salt/file_root/ntp/init.sls;sed -i
  s@"add_admin_portal_documentation(server_dict\['ip'\]"@"#add_admin_portal_documentation(server_dict\['ip'\])"@g
  micro_services/core.py;
  ```

  You can then use the files and tools in the installer directory to perform operations such as provisioning VMs, creating configuration files, and installing the solution.

  [CXU-12113]

- When you try to install the **Distributed_Cloud_CPE_Network_Service_Controller_3.1** package, the load services data module fails with an import error in the **publish_data_to_design_tools** function.

  Workaround:

  1. Navigate to the ***untar-dir*/micro_services/** directory, where *untar-dir* is the directory where you unzipped the installation tar file.

  2. Open the **load_services_data.py** file in an editor and comment out the **publish_data_to_design_tools** function in the file as follows:

     ```
     '''
     # publish data to design-tools
     publish_data_to_design_tools(
         token,
         regions['central']['vip_ip'],
         regions['central']['vip_port'],
         data_str,
         http_protocol
     )
     '''
     ```

  3. Save the **load_services_data.py** file.

  4. Run the **load_service_data.sh** script to continue with the installation.

  [CXU-12137]

- When you deploy a firewall policy, the deployment fails with the message **Fail to invoke mapper to create snapshot with reason null**.

  Workaround: Do the following:

Copyright © 2017, Juniper Networks, Inc.

1. Log in to the central infrastructure virtual machine (VM) as root.
2. Start the ZooKeeper CLI by executing the **/usr/share/zookeeper/bin/zkCli.sh** command.
3. Execute the **delete /secmgt/sm_initialized** command.
4. Exit the ZooKeeper CLI by executing the **quit** command.
5. Log out of the central infrastructure VM.
6. Log in to the central microservice VM as root.
7. At the shell prompt, execute the **kubectl get pods | grep secmgt-sm** command to find out the name of the security management pod.
8. Restart the pod by executing the **kubectl delete pod** *pod-name* command, where *pod-name* is the name of the security management pod.
9. Wait until the security management pod is in the **1/1 running** state.
10. Log out of the central microservice VM.
11. Re-deploy the firewall policy.

[CXU-12151]

Resolved Issues

The following issues are resolved in Juniper Networks Cloud CPE Solution Release 3.1.

- Deactivating an SRX Series Services Gateway acting as a CPE device might not remove all configuration settings from the device. [CXU-3754]

Documentation Updates

This section lists the errata and changes in the Cloud CPE Solution Release 3.1 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.

Copyright © 2017, Juniper Networks, Inc.

- **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

Copyright © 2017, Juniper Networks, Inc.

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

20 December 2017—Revision 10

Copyright © 2017, Juniper Networks, Inc.

30 October 2017—Revision 9

5 October 2017—Revision 8

3 October 2017—Revision 7

26 September 2017—Revision 6

13 September 2017—Revision 5

8 September 2017—Revision 4

1 September 2017—Revision 3

22 August 2017—Revision 2

21 August 2017—Revision 1, Cloud CPE Solution Release 3.1

Copyright © 2017 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

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