Contrail Service Orchestration Deployment Guide

Release
3.3

Modified: 2018-11-19

Copyright © 2018, Juniper Networks, Inc.

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

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

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

Contrail Service Orchestration Deployment Guide
3.3
Copyright © 2018 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

List of Figures

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

List of Tables

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

About the Documentation

Documentation and Release Notes

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

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

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

Documentation Conventions

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

Copyright © 2018, Juniper Networks, Inc.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*

Copyright © 2018, Juniper Networks, Inc.

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

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

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

Copyright © 2018, Juniper Networks, Inc.

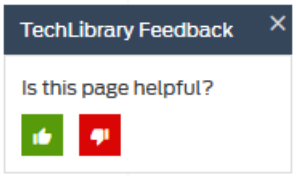

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to **techpubs-comments@juniper.net**. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: https://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: https://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: https://kb.juniper.net/
- Download the latest versions of software and review release notes: https://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/

Copyright © 2018, Juniper Networks, Inc.

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

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 1

Overview of Contrail Service Orchestration

Cloud CPE and SD-WAN Solutions Overview

The Juniper Networks Cloud Customer premises equipment (CPE) and SD-WAN solutions use the Contrail Service Orchestration (CSO) to transform traditional branch networks, offering opportunities for high flexibility of the network, rapid introduction of new services, automation of network administration, and cost savings. The solutions can be implemented by service providers for their customers or by Enterprise IT departments in a campus and branch environment. In this documentation, service providers and Enterprise IT departments are called *service providers*, and the consumers of their services are called *customers*.

The Cloud CPE solution supports both Juniper Networks and third-party virtualized network functions (VNFs) that network providers use to create the network services. The following deployment models are available:

- 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 service edge sites in this documentation.

- Cloud CPE Distributed Deployment Model (distributed deployment), also known as a hybrid WAN 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 this documentation.

Copyright © 2018, Juniper Networks, Inc.

- A combined centralized and distributed deployment

  In this deployment, the network contains both service edge sites and on-premise sites. A customer can have both cloud sites and tenant sites; however, you cannot share a network service between the centralized and distributed deployments. If you require the same network service for the centralized deployment and the distributed deployment, you must create two identical network services with different names.

You must consider several issues when choosing whether to employ one or both types of deployment. The centralized deployment offers a fast migration route and this deployment is the recommended model for sites that can accommodate network services—particularly security services—in the cloud. In contrast, the distributed deployment supports private hosting of network services on a CPE device at a customer's site, and can be extended to offer software defined wide area networking (SD-WAN) capabilities. Implementing a combination network in which some sites use the centralized deployment and some sites use the distributed deployment provides appropriate access for different sites.

The SD-WAN solution offers a flexible and automated way to route traffic through the cloud. Similar to a distributed deployment, this implementation uses CPE devices located at on-premise sites to connect to the LAN segments. Hub-and-spoke and full mesh topologies are supported. The CSO software uses SD-WAN policies and service-level agreement measurements to differentiate and route traffic for different applications.

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

You can either use the solutions as turnkey implementations or connect to other operational support and business support systems (OSS/BSS) through northbound Representational State Transfer (REST) APIs.

Related Documentation

- NFV in the Cloud CPE Solution on page 18
- Topology of the Cloud CPE and SD-WAN Solutions on page 22
- Benefits of the Cloud CPE Solution on page 32

NFV in the Cloud CPE Solution

The Cloud CPE Solution uses the following components for the NFV environment:

- For the centralized deployment:
  - Network Service Orchestrator provides ETSI-compliant management of the life cycle of network service instances.

    This application includes RESTful APIs that you can use to create and manage network service catalogs.
  - Contrail OpenStack provides the following functionality:
    - Underlying software-defined networking (SDN) to dynamically create logical service chains that form the network services

Copyright © 2018, Juniper Networks, Inc.

    - NFV infrastructure (NFVI).
    - Virtualized infrastructure manager (VIM)

- For the distributed deployment:
    - Network Service Orchestrator, together with Network Service Controller, provides ETSI-compliant management of the life cycle of network service instances.
    - Network Service Controller provides service-chaining and the VIM.
    - The CPE device provides the NFV infrastructure (NFVI).

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

- Administration Portal, which you use to set up and manage your virtual network and customers through a graphical user interface (GUI).

  Administration Portal offers role-based access control for administrators and operators. Administrators can make changes; however, operators can only view the portal.

- Customer Portal, a GUI that your customers use to manage sites, CPE devices, and network services for their organizations.

  Customer Portal offers role-based access control for administrators and operators. Administrators can make changes; however, operators can only view the portal.

- Designer Tools:
    - Configuration Designer, which you use to create configuration templates for virtualized network functions (VNFs). When you publish a configuration template, it is available for use in Resource Designer.
    - Resource Designer, which you use to create VNF packages. A VNF package consists of a configuration template and specifications for resources. You use configuration templates that you create with Configuration Designer to design VNF packages. When you publish a VNF package, it is available for use in Network Service Designer.
    - Network Service Designer, which you use to create a network service package. The package offers a specified performance and provides one or more specific network functions, such as a firewall or NAT, through one or more specific VNFs.

- Service and Infrastructure Monitor, which works with Icinga, an open source enterprise monitoring system to provide real-time data about the Cloud CPE solution, such as the status of virtualized network functions (VNFs), virtual machines (VMs), and physical servers; information about physical servers' resources; components of a network service (VNFs and VMs hosting a VNF); counters and other information for VNFs.

The Cloud CPE solution extends the NFV model through the support of physical network elements (PNEs). A PNE is a networking device in the deployment that you can configure through CSO, but not use in a service chain. Configuration of the PNE through CSO as opposed to other software, such as Contrail or Junos OS, simplifies provisioning of the physical device through automation. Combining provisioning and configuration for PNEs and VNFs provides end-to-end automation in network configuration workflows. An

Copyright © 2018, Juniper Networks, Inc.

example of a PNE is the MX Series router that acts as an SDN gateway in a centralized deployment.

In the distributed deployment, VNFs reside on a CPE device located at a customer site. The NFX250 is a switch that hosts the vSRX application to enable routing and IPSec VPN access with the service provider's POP. MX Series routers, configured as provider edge (PE) routers, provide managed Layer 1 and Layer 2 access and managed MPLS Layer 3 access to the POP. Network Service Controller provides the VIM, NFVI, and device management for the NFX250. Network Service Controller includes Network Activator, which enables remote activation of the NFX Series device when the site administrator connects the device and switches it on.

Figure 1 on page 20 illustrates how the components in the Cloud CPE solution interact and how they comply with the ETSI NFV MANO model.

Figure 1: NFV Components of the Cloud CPE Solution

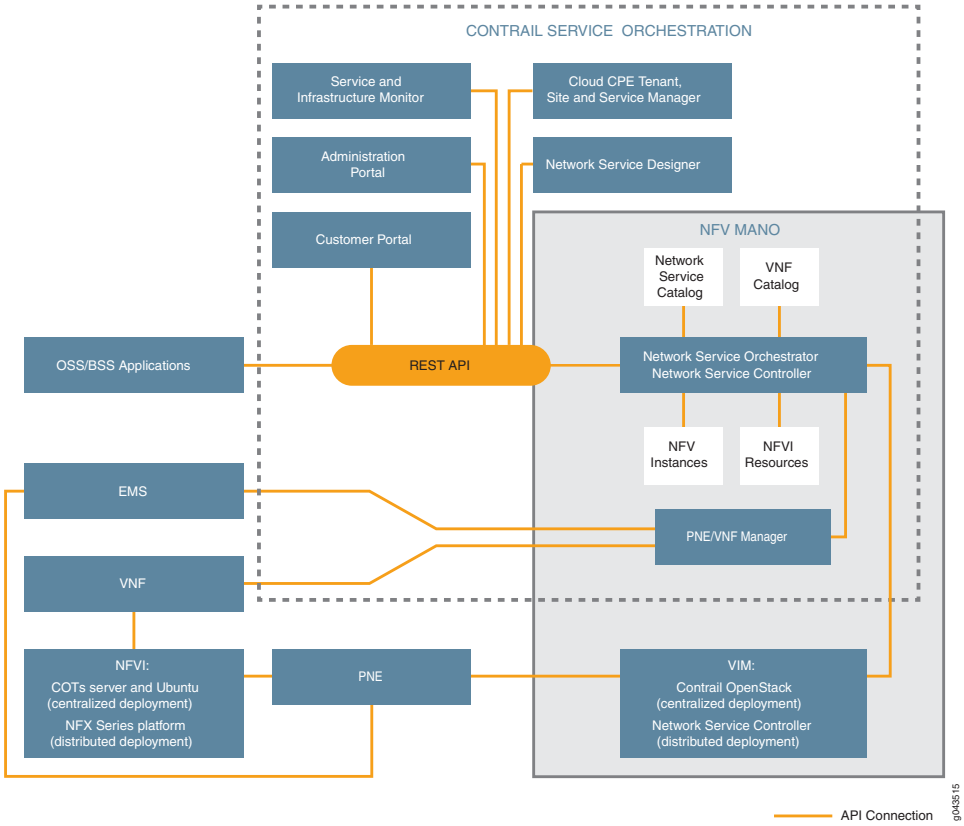

OSS/BSS applications and Contrail Service Orchestration (CSO) components with OSS/BSS capabilities send requests to Network Service Orchestrator through its northbound REST API. Network Service Orchestrator then communicates through its southbound API to the northbound API of the appropriate, directly connected, component. Subsequently, each component in the deployment communicates through its southbound API to the northbound API of the next component in the hierarchy. Components send responses in the reverse direction.

Copyright © 2018, Juniper Networks, Inc.

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

1. Customers send requests for activations of network services through Customer Portal or OSS/BSS applications.

2. Service and Infrastructure Monitor is continuously tracking the software components, hardware components, and processes in the network.

3. Network Service Orchestrator receives requests through its northbound RESTful API and:

   - For the centralized deployment:

     a. Accesses information about the network service and associated VNFs from their respective catalogs, and communicates this information to the VIM, which is provided by Contrail OpenStack.

     b. Sends information about the VNF to VNF Manager.

   - For the distributed deployment, accesses information about the network service and associated VNFs from their respective catalogs, and communicates this information to the Network Service Controller.

4. The VIM receives information from Network Service Orchestrator and:

   - For the centralized deployment:

     - The VIM creates the service chains and associated VMs in the NFVI, which is provided by the servers and Ubuntu. Contrail OpenStack creates one VM for each VNF in the service chain.

     - VNF Manager starts managing the VNF instances while the element management system (EMS) performs element management for the VNFs.

   - For the distributed deployment, Network Service Controller creates the service chains and associated VMs in the NFVI, which is provided by the CPE device.

5. The network service is activated for the customer.

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

- For the centralized deployment:

  1. Network Service Orchestrator receives the request through its northbound RESTful API and sends information about the PNE to PNE/VNF Manager.

  2. PNE/VNF Manager receives information from Network Service Orchestrator and sends information about the PNE to the EMS.

  3. VNF Manager starts managing the VNF instances and the EMS starts element management for the VNFs.

  4. The PNE becomes operational.

Copyright © 2018, Juniper Networks, Inc.

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

Related Documentation

- Cloud CPE and SD-WAN Solutions Overview on page 17
- Topology of the Cloud CPE and SD-WAN Solutions on page 22
- Benefits of the Cloud CPE Solution on page 32

Topology of the Cloud CPE and SD-WAN Solutions

Figure 2 on page 22 shows the topology of the Cloud Customer Premises equipment (CPE) and SD-WAN solutions. You can use one Contrail Service Orchestration (CSO) installation for all or any of the supported solutions and deployments:

- Cloud CPE solution
  - Centralized deployment
  - Distributed (also known as hybrid WAN) deployment
  - Combined centralized and distributed deployment
- SD-WAN solution

Figure 2: Cloud CPE and SD-WAN Solutions Topology

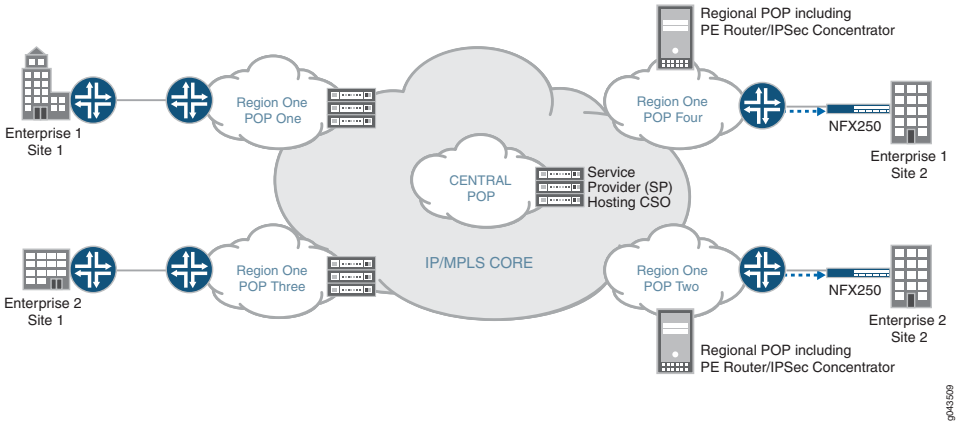

Different sites for an enterprise might connect to different regional POPs, depending on the geographical location of the sites. Within an enterprise, traffic from a site that connects to one regional POP travels to a site that connects to another regional POP through the

Copyright © 2018, Juniper Networks, Inc.

central POP. A site can connect to the Internet and other external links through either the regional POP or the central POP.

Service providers use the central server to set up the Cloud CPE solution through Administration Portal. Similarly, customers activate and manage network services through their own dedicated view of Customer Portal on the central server.

Topologies of the Implementations and Deployments

Centralized Deployment

Figure 3 on page 23 illustrates the topology of a centralized deployment. Customers access network services in a regional cloud through a Layer 3 VPN.

Figure 3: Centralized Deployment Topology

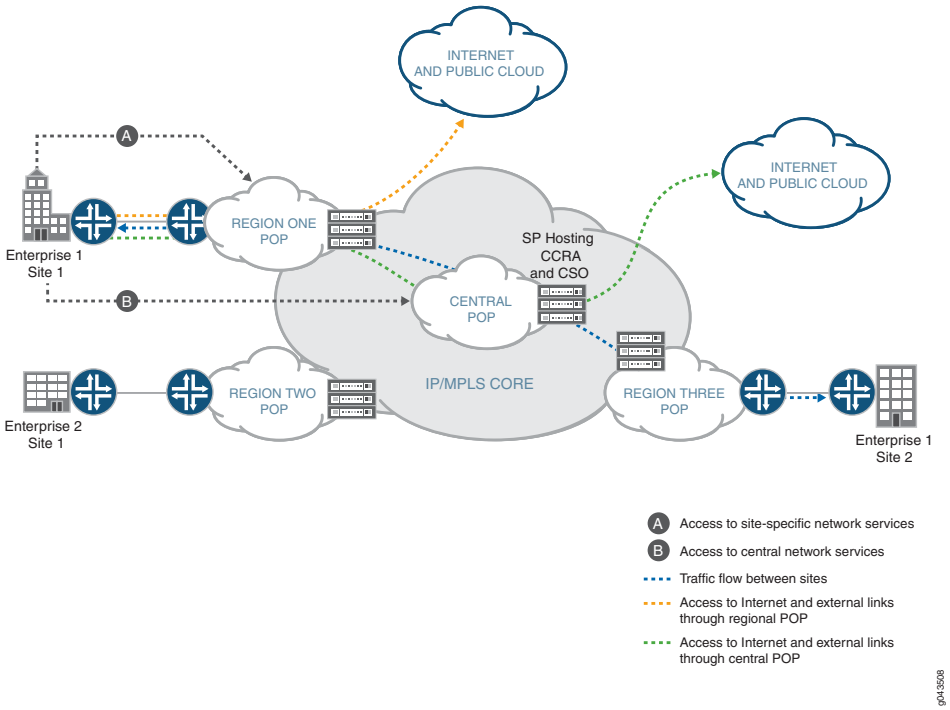

The central and regional POPs contain one or more Contrail Cloud implementations. VNFs reside on Contrail compute nodes and service chains are created in Contrail. You can choose whether to use the CSO OpenStack Keystone on the central infrastructure server or the OpenStack Keystone on the Contrail controller node in the central POP to authenticate CSO operations. The Contrail Cloud implementation provides Contrail Analytics for this deployment.

The MX Series router in the Contrail Cloud implementation is an SDN gateway and provides a Layer 3 routing service to customer sites through use of virtual routing and forwarding (VRF) instances, known in Junos OS as Layer 3 VPN routing instances. A unique routing table for each VRF instance separates each customer's traffic from other customers' traffic. The MX Series router is a PNE.

Copyright © 2018, Juniper Networks, Inc.

Sites can access the Internet directly, through the central POP, or both. Data traveling from one site to another passes through the central POP.

Distributed Deployment

Figure 4 on page 24 illustrates the topology of a distributed deployment.

Figure 4: Distributed Deployment Topology

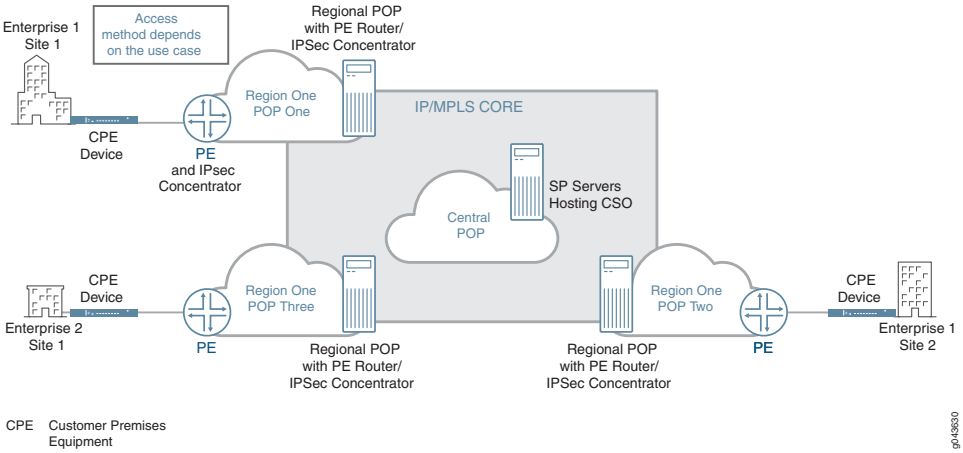

Each site in a distributed deployment hosts a CPE device on which the vSRX application is installed to provide security and routing services. The Cloud CPE solution supports the following CPE devices:

- NFX250 Network Services Platform
- SRX Series Services Gateway
- vSRX

  The vSRX CPE device can reside at a customer site or in the service provider cloud. In both cases, you configure the site in CSO as an on-premise site. Authentication of the vSRX as a CPE device takes place through SSH.

An MX Series router in each regional POP acts as an IPsec concentrator and provider edge (PE) router for the CPE device. An IPsec tunnel, with endpoints on the CPE device and MX Series router, enables Internet access from the CPE device. Data flows from one site to another through a GRE tunnel with endpoints on the PE routers for the sites. The distributed deployment also supports SD-WAN functionality for traffic steering, based on 5-tuple (source IP address, source TCP/UDP port, destination IP address, destination TCP/UDP port and IP protocol) criteria.

Network administrators can configure the MX Series router, the GRE tunnel, and the IPsec tunnel through Administration Portal. Similar to the centralized deployment, the MX Series router in the distributed deployment is a PNE.

The CPE device provides the NFVI, which supports the VNFs and service chains. Customers can configure sites, CPE devices, and network services with Customer Portal.

Copyright © 2018, Juniper Networks, Inc.

The OpenStack Keystone resides on the central infrastructure server and Contrail Analytics resides on a dedicated VM or server.

SD-WAN Solution

The SD-WAN solution supports hub-and-spoke and full mesh VPN topologies.

Figure 5 on page 25 shows the topology of the SD-WAN Solution with a hub and spoke implementation.

Figure 5: SD-WAN Topology

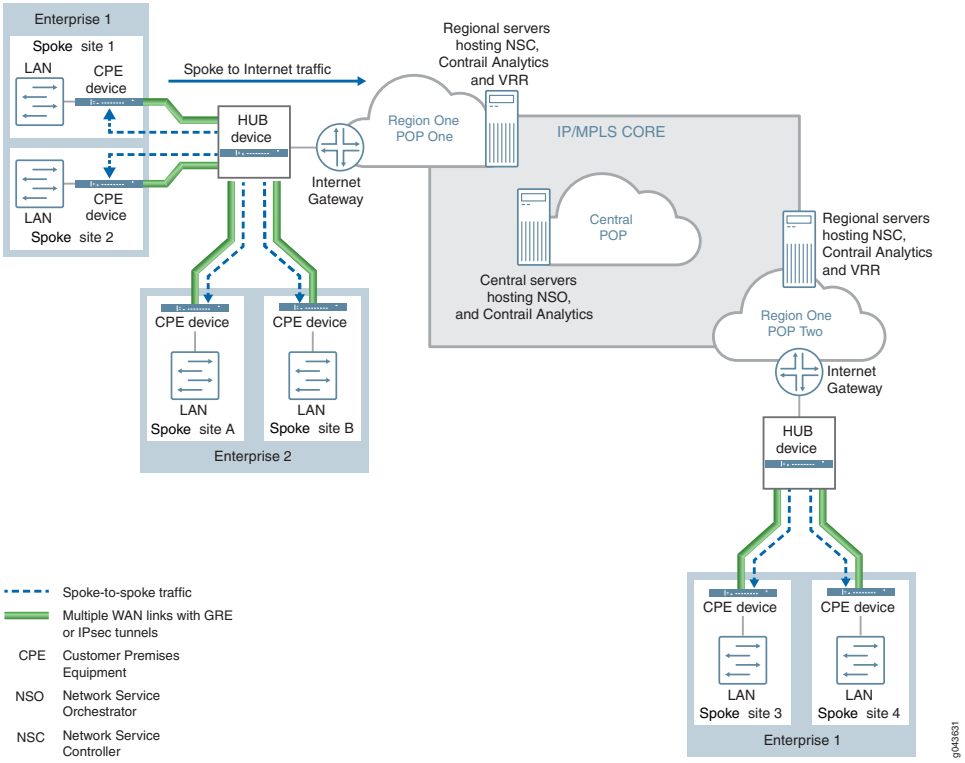

The SD-WAN implementation supports a hub-and-spoke VPN topology, in which CPE devices reside at the spoke sites. The CPE devices are the same as those used in a distributed deployment. The hub device, which is an SRX Series gateway, typically serves all the spoke sites for all the customers in a POP. You can, however, dedicate a hub device to a specific tenant. In the hub-and-spoke topology, all traffic from a LAN segment passes through the hub, whether it is traveling to another of the customer's sites in the same POP or to the Internet.

A virtual route reflector (VRR) resides on a VM on each regional microservices server. During the CSO installation, a VRR is installed on the regional servers. The VRR has a fixed configuration that you cannot modify. Use of a VRR enhances scaling of the BGP network with low cost and removes the need for hardware-based route reflectors that require space in a data center and ongoing maintenance.

Copyright © 2018, Juniper Networks, Inc.

For VRR redundancy, you need create at least two VRRs for a region. We recommend that you create VRRs in even numbers and assign these VRRs equally in different redundancy groups. Each hub or spoke device establishes a BGP peering session with two VRRs that are in different redundancy groups. If the primary VRR fails or connectivity is lost, the BGP peering session remains active because the secondary VRR continues to receive and advertise LAN routes to a device, thereby providing redundancy.

Redundancy groups are formed by logically separating VRRs based on following parameters:

- Physical server affinity—VRRs that reside on a same physical server should not belong to different redundancy group.
- Network affinity—VRRs that reside on a same network should not belong to different redundancy group.

There can be only two redundancy groups—group 0 and group 1. If you do not specify the redundancy group for VRRs, all VRRs are placed in the default redundancy group—group 0—and hub or spoke devices establish a BGP session with only one VRR.

Related Documentation

- Cloud CPE and SD-WAN Solutions Overview on page 17
- NFV in the Cloud CPE Solution on page 18
- Benefits of the Cloud CPE Solution on page 32

Resiliency of the Cloud CPE and SD-WAN Solutions

The Cloud CPE and SD-WAN solutions offer robust implementations with resiliency for the following features:

- High availability of Contrail Service Orchestration (CSO) infrastructure services and microservices in a production environment.

  Each infrastructure service or microservice resides on multiple hosts and if an application on the primary host fails, a corresponding application on another host takes over. Current operations for an application do not recover if a failure occurs; however, any new operations proceed as normal.

- Support for a centralized Cloud CPE deployment on a Contrail OpenStack instance that you configure for high availability.

  The Contrail OpenStack instance includes three Contrail controller nodes in the Contrail Cloud Platform, and provides resiliency for virtualized infrastructure managers (VIMs), virtualized network functions (VNFs), and network services.

- CSO provides additional resiliency for virtualized network functions (VNFs) and network services in the Cloud CPE solution. You can enable or disable automatic recovery of a network service in a centralized deployment. If a network service becomes unavailable due to a connectivity issue with a VNF, Network Service Orchestrator maintains existing instances of the network service in end users' networks and initiates recreation of the VNFs. During this recovery process, the end user cannot activate the network service

Copyright © 2018, Juniper Networks, Inc.

on additional network links. When the problem is resolved, normal operation resumes and end users can activate the network service on additional network links.

Enabling automatic recovery improves reliability of the implementation. Conversely, disabling automatic recovery for a network service allows you to quickly investigate a problem with the underlying VNF. By default, automatic recovery of a network service is enabled.

Related Documentation

- Architecture of the Contrail Cloud Implementation in the Centralized Deployment on page 28
- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76
- Installing and Configuring Contrail Service Orchestration on page 111
- Cloud CPE and SD-WAN Solutions Overview on page 17
- NFV in the Cloud CPE Solution on page 18

Authentication and Authorization in the Cloud CPE and SD-WAN Solutions

The Cloud CPE and SD-WAN solutions use OpenStack Keystone to authenticate and authorize Contrail Service Orchestration (CSO) operations. You can implement the Keystone in several different ways, and you specify which method you use when you install CSO:

- A CSO Keystone, which is integrated with CSO and resides on the central CSO server.

  This option offers enhanced security because the Keystone is dedicated to CSO and is not shared with any other applications. Consequently, this option is generally recommended.

- An external Keystone, which resides on a different server to the CSO server:
  - The Contrail OpenStack Keystone in the Contrail Cloud Implementation for a centralized deployment is an example of an external Keystone.

    In this case, customers and Cloud CPE infrastructure components use the same Keystone token.
  - You can also use an external Keystone that is specific to your network.

See Table 3 on page 28 for guidelines about using the Keystone options with different types of deployments.

Copyright © 2018, Juniper Networks, Inc.

Table 3: Guidelines for Keystone Options for Different Deployments

	Centralized Deployment	Distributed Deployment and SD-WAN Implementation	Combined deployment
The CSO Keystone (recommended)	• Installation of the Keystone occurs with the CSO installation. • After installation, you must use Administration Portal or the API to configure a service profile for each virtualized infrastructure monitor (VIM).	• Installation occurs with the CSO installation. • You do not need to perform any configuration after installation.	• Installation occurs with the CSO installation. • You do not need to perform any configuration after installation for the distributed portion of the deployment. • After installation, you must configure service profiles for VIMs in the centralized portion of the deployment.
The Contrail OpenStack Keystone on the Contrail Cloud Platform (external Keystone)	• Installation occurs with Contrail OpenStack • You specify the IP address and access details for the Contrail OpenStack Keystone when you install CSO.	Not available	• Available for the centralized portion of the deployment. • Installation occurs with Contrail OpenStack. • You specify the IP address and access details for the Contrail OpenStack Keystone when you install CSO.
An external Keystone that is specific to your network.	You specify the IP address and access details for your Keystone when you install CSO.		

Related Documentation

- Cloud CPE and SD-WAN Solutions Overview on page 17

Architecture of the Contrail Cloud Implementation in the Centralized Deployment

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

- Architecture of the Contrail Cloud Implementation on page 28
- Architecture of the Servers on page 29
- Architecture of the Contrail Nodes on page 31

Architecture of the Contrail Cloud Implementation

The centralized deployment uses the Contrail Cloud implementation to support the service provider's cloud in a network point of presence (POP). The Contrail Cloud implementation consists of the hardware platforms, including the servers, and Contrail OpenStack software. Figure 6 on page 29 illustrates the Contrail Cloud implementation. The Contrail Service Orchestration (CSO) software is installed on one or more servers in the Contrail Cloud implementation to complete the deployment.

Copyright © 2018, Juniper Networks, Inc.

Figure 6: Architecture of Contrail Cloud Implementation

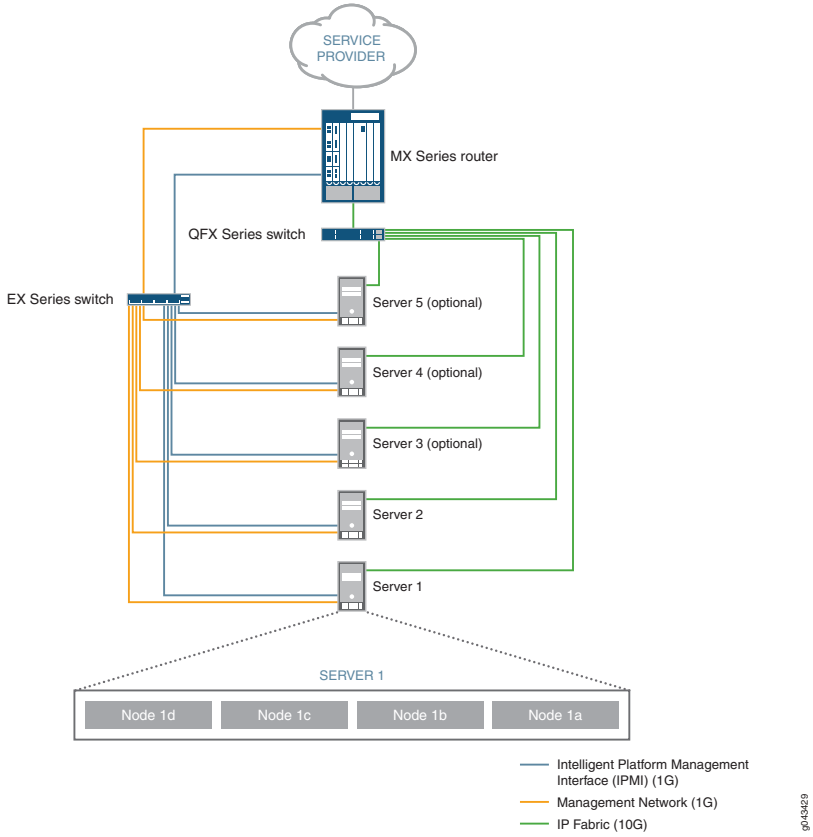

In the Cloud CPE Centralized Deployment Model:

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

Architecture of the Servers

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

- Contrail Service Orchestration node, which hosts the Contrail Service Orchestration software.

Copyright © 2018, Juniper Networks, Inc.

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

The Contrail Cloud implementation in a central POP contains all three types of node. Figure 7 on page 30 shows the configuration of the nodes in the Contrail Cloud implementation in the central POP for a deployment that offers neither Contrail nor Contrail Service Orchestration high availability:

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

Figure 7: Architecture of Servers in the Central POP for a Non-Redundant Installation

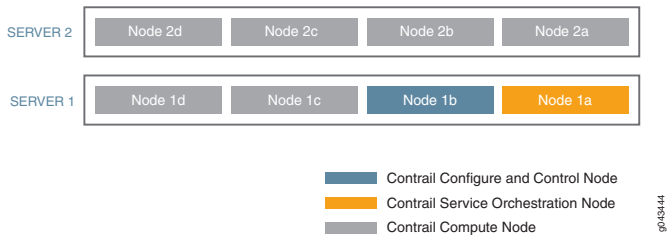

Figure 8 on page 30 shows the configuration of the nodes in the Contrail Cloud implementation in the central POP for a deployment that offers both Contrail and Contrail Service Orchestration high availability:

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

Figure 8: Architecture of Servers in the Central POP for a Redundant Installation

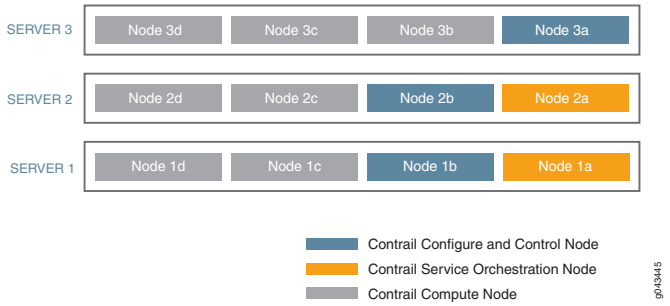

The Contrail Cloud implementation in a regional POP contains only Contrail nodes and not Contrail Service Orchestration nodes. In a deployment that does not offer Contrail high availability, the regional Contrail Cloud implementations support:

Copyright © 2018, Juniper Networks, Inc.

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

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

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

Architecture of the Contrail Nodes

Each Contrail controller node uses Contrail vRouter over Ubuntu and kernel-based virtual machine (KVM) as a forwarding plane in the Linux kernel. Use of vRouter on the compute node separates the deployment's forwarding plane from the control plane, which is the SDN Controller in Contrail OpenStack on the controller node. This separation leads to uninterrupted performance and enables scaling of the deployment. Figure 9 on page 31 shows the logical representation of the Contrail controller nodes.

Figure 9: Logical Representation of Contrail Controller Nodes

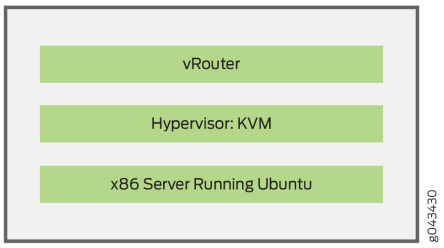

A Contrail compute node hosts Contrail OpenStack, and the VNFs. Contrail OpenStack resides on the physical server and cannot be deployed in a VM. Each VNF resides in its own VM. Figure 10 on page 31 shows the logical representation of the Contrail compute nodes.

Figure 10: Logical Representation of Contrail Compute Nodes

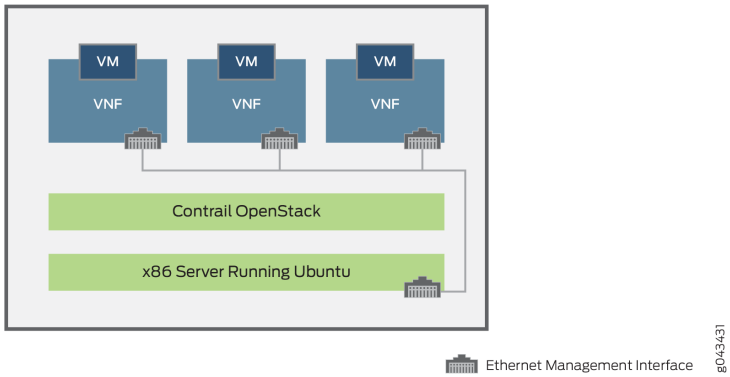

Related Documentation

- Topology of the Cloud CPE and SD-WAN Solutions on page 22
- Resiliency of the Cloud CPE and SD-WAN Solutions on page 26
- Cloud CPE and SD-WAN Solutions Overview on page 17
- NFV in the Cloud CPE Solution on page 18

Benefits of the Cloud CPE Solution

Juniper Networks Cloud Customer Premises Equipment (CPE) solution offers an automated branch network environment, leading to cost savings over traditional branch networks, while improving network agility and reducing configuration errors. The centralized deployment offers a fast migration route through either existing equipment or a Layer 3 network interface device (NID), and is the recommended model for sites that can accommodate network services—particularly security services—in the cloud. Use of a CPE device such as a NFX Series Network Services platform or SRX Series Services Gateway in the distributed deployment supports private hosting of network services at a site and offers software defined wide area networking (SD-WAN) capabilities. Implementing a combination network in which some sites use the centralized deployment and some sites use the distributed deployment provides appropriate access for different sites.

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

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

Use of generic servers and CPE devices with VNFs leads to capital expenditure (CAPEX) savings compared to purchasing dedicated network devices. Set up and ongoing support of the equipment requires minimal work at the branch site: for the centralized deployment, the equipment resides in a central office, and for the distributed deployment, the CPE device uses remote activation to initialize, become operational, and obtain configuration updates. The reduced setup and maintenance requirements, in addition to automated configuration, orchestration, monitoring, and recovery of network services, result in lower operating expenses (OPEX).

Related Documentation

- Cloud CPE and SD-WAN Solutions Overview on page 17
- NFV in the Cloud CPE Solution on page 18

Copyright © 2018, Juniper Networks, Inc.

- Topology of the Cloud CPE and SD-WAN Solutions on page 22

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 2

Specifications

Number of Sites and VNFs Supported in Contrail Service Orchestration

The Cloud CPE solution supports two environment types: a trial environment and a production environment. You can deploy the environments with or without high availability (HA). Table 4 on page 35 shows the number of sites and VNFs supported for each environment.

Table 4: Number of Sites and VNFs Supported

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

Each environment has different requirements for:

Copyright © 2018, Juniper Networks, Inc.

- The number and specification of node servers and servers. See "Minimum Requirements for Servers and VMs" on page 40
- The number and specification of virtual machines (VMs). "Provisioning VMs on Contrail Service Orchestration Nodes or Servers" on page 76

Related Documentation

- Minimum Requirements for Servers and VMs on page 40
- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76
- Hardware and Software Required for Contrail Service Orchestration on page 36

Hardware and Software Required for Contrail Service Orchestration

Contrail Service Orchestration requires commercial off-the-shelf (COTS) node servers or servers, specific network devices, and specific software versions. These sections list the hardware and software that are required and have been tested for the Cloud CPE and SD-WAN solutions.

- Node Servers and Servers Tested in Contrail Service Orchestration on page 36
- Network Devices and Software Tested in the Centralized Deployment on page 37
- Network Devices and Software Tested in the Hybrid WAN Distributed Deployment and the SD-WAN Implementation on page 38

Node Servers and Servers Tested in Contrail Service Orchestration

You use COTS node servers or servers for the following functions:

- Contrail Service Orchestration (CSO) central and regional servers
- Contrail Analytics servers
- Contrail controller and compute nodes in the centralized deployment

Table 5 on page 36 lists the node servers and servers that have been tested for these functions.

Table 5: COTS Node Servers and Servers Tested in the Cloud CPE and SD-WAN Solutions

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

Table 6 on page 37 shows the software that has been tested for COTS servers used in the Cloud CPE solution. You must use these specific versions of the software when you implement the Cloud CPE and SD-WAN solutions.

Copyright © 2018, Juniper Networks, Inc.

Table 6: Software Tested for the COTS Nodes and Servers

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

Network Devices and Software Tested in the Centralized Deployment

Table 7 on page 37 shows the network devices that have been tested for the centralized deployment.

Table 7: Network Devices Tested for the Centralized Deployment

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

Table 8 on page 38 shows the software tested for the centralized deployment. You must use these specific versions of the software when you implement a centralized deployment.

Copyright © 2018, Juniper Networks, Inc.

Table 8: Software Tested in the Centralized Deployment

Function	Software and Version
Operating system for MX Series router	Junos OS Release 14.2R3
Operating system for QFX Series switch	Junos OS Release 13.2X51-D38
Hypervisor on CSO servers	KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0
Element management system software	EMS microservice Junos Space Network Management Platform Release 15.1R1 (See "VNFs Supported by the Cloud CPE Solution" on page 52 for VNFs that require this product)
Software defined networking (SDN), including Contrail Analytics, for a centralized deployment	Contrail Release 3.2.5 with OpenStack Mitaka
Virtualized infrastructure manager (VIM) and virtual machine (VM) orchestration	OpenStack Mitaka
Authentication and Authorization	OpenStack Mitaka
Network Functions Virtualization (NFV)	CSO Release 3.3

Network Devices and Software Tested in the Hybrid WAN Distributed Deployment and the SD-WAN Implementation

Table 9 on page 38 shows the network devices that have been tested for the distributed deployment and the SD-WAN implementation.

Table 9: Network Devices Tested for the Distributed Deployment and SD-WAN Implementation

Function	Device	Model
PE router and IPsec concentrator (Hybrid WAN distributed deployment only)	Juniper Networks MX Series 3D Universal Edge Router	• MX960, MX480, or MX240 router with a Multiservices MPC line card • MX80 or MX104 router with Multiservices MIC line card • Other MX Series routers with a Multiservices MPC or Multiservices MIC line card See MPCs Supported by MX Series Routers and MICs Supported by MX Series Routers for information about MX Series routers that support Multiservices MPC and MIC line cards.

Copyright © 2018, Juniper Networks, Inc.

Table 9: Network Devices Tested for the Distributed Deployment and SD-WAN Implementation (continued)

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

Table 10 on page 39 shows the software tested for the distributed deployment. You must use these specific versions of the software when you implement a distributed deployment.

Table 10: Software Tested in the Distributed Deployment and SD-WAN Solution

Function	Software and Version
Hypervisor on CSO servers	KVM provided by the Ubuntu operating system on the server or VMware ESXi Version 5.5.0
Authentication and Authorization	OpenStack Mitaka
Network Functions Virtualization (NFV)	CSO Release 3.3.0
Contrail Analytics	Contrail Release 4.0.3.0-162
NFX Software	Junos OS Release 15.1X53-D472
Routing and Security for NFX250 device	vSRX KVM Appliance 15.1X49-D133
Operating system for vSRX used as a CPE device on an x86 server	vSRX KVM Appliance 15.1X49-D133

Copyright © 2018, Juniper Networks, Inc.

Table 10: Software Tested in the Distributed Deployment and SD-WAN Solution (continued)

Function	Software and Version
Operating system for SRX Series Services Gateway used as a CPE device or spoke device	Junos OS Release 15.1X49-D133
Operating system for MX Series router used as PE router	Junos OS Release 16.1R3.00
Operating system for MX Series Router used as a hub device for an SD-WAN implementation	Junos OS Release 16.1R5.7
Operating system for SRX Series Services Gateway used as a hub device for an SD-WAN implementation	Junos OS Release 15.1X49-D133

Related Documentation

- Minimum Requirements for Servers and VMs on page 40

Minimum Requirements for Servers and VMs

Minimum Hardware Requirements for Node Servers and Servers

For information about the makes and models of node servers and servers that you can use in the Cloud CPE solution, see Table 5 on page 36. 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.

Table 11 on page 40 shows the specification for the nodes and servers for the Cloud CPE or SD-WAN solution.

Table 11: Specification for Nodes and Servers

Item	Requirement
Storage	Greater than 1 TB of one of the following types: • Serial Advanced Technology Attachment (SATA) • Serial Attached SCSI (SAS) • Solid-state drive (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

Copyright © 2018, Juniper Networks, Inc.

The number of node servers and servers that you require depends on whether you are installing a trial or a production environment, and whether you require high availability (HA).

Table 12 on page 41 shows the required hardware specifications for node servers and servers in the supported environments. The server specifications are slightly higher than the sum of the virtual machine (VM) specifications listed in "Minimum Requirements for VMs on CSO Node Servers or Servers" on page 42, because some additional resources are required for the system software.

Table 12: Server Requirements

Function	Trial Environment without HA	Trial Environment with HA	Production Environment without HA	Production Environment with HA
Contrail Service Orchestration (CSO) Servers				
NOTE: If you use a trial environment without HA and with virtualized network functions (VNFs) that require Junos Space as the Element Management System (EMS), you must install Junos Space on a VM on another server. This server specification for a trial environment without HA does not accommodate Junos Space. For information on Junos Space VM requirements, see Table 13 on page 42.				
Number of nodes or servers	1	3	2 • 1 central server • 1 regional server	6 • 3 central servers • 3 regional servers
vCPUs per node or server	48	48	48	48
RAM per node or server	256 GB	256 GB	256 GB	256 GB
Contrail Analytics Servers for a Hybrid WAN or SD-WAN Deployment				
Number of servers	None—Contrail Analytics is in a VM	None—Contrail Analytics is in a VM	1	3
vCPUs per node or server	—	—	48	48
RAM per node or server	—	—	256 GB	256 GB
Contrail Cloud Platform for a Centralized Deployment				
Number of nodes or servers	1	4–8 • 3 nodes for Contrail controller and analytics • 1–4 Contrail compute nodes	4–28 • 3 nodes for Contrail controller and analytics • 1–25 Contrail compute nodes	4–28 • 3 nodes for Contrail controller and analytics • 1–25 Contrail compute nodes
vCPUs per node or server	16	48	48	48
RAM per node or server	64 GB	256 GB	256 GB	256 GB

Copyright © 2018, Juniper Networks, Inc.

Table 12: Server Requirements (continued)

Function	Trial Environment without HA	Trial Environment with HA	Production Environment without HA	Production Environment with HA
Total Numbers of Servers				
Centralized deployment	2	7–11	6–30	10–34
Hybrid WAN or SD-WAN	1	3	3	9

Minimum Requirements for VMs on CSO Node Servers or Servers

The number and minimum requirements for CSO VMs depends on the deployment environment and whether or not you use HA:

- For a trial environment without HA, see Table 13 on page 42.
- For a trial environment with HA, see Table 14 on page 43.
- For a production environment without HA, see Table 15 on page 46.
- For a production environment with HA, see Table 16 on page 47.

For information about the ports that must be open on all VMs for all deployment environments, see Table 17 on page 50.

Table 13 on page 42 shows complete details about the VMs for a trial environment without HA.

Table 13: Details of VMs for a Trial Environment

Name of VM	Components That Installer Places in VM	Resources Required
csp-installer-vm	—	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage
csp-central-infravm	Third-party applications used as infrastructure services	• 4 CPU • 32 GB RAM • 200 GB hard disk storage
csp-central-msvm	All microservices, including GUI applications	• 4 vCPUs • 32 GB RAM • 200 GB hard disk storage
csp-regional-infravm	Third-party applications used as infrastructure services	• 4 vCPUs • 24 GB RAM • 200 GB hard disk storage

Copyright © 2018, Juniper Networks, Inc.

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

Name of VM	Components That Installer Places in VM	Resources Required
csp-regional-msvm	All microservices, including GUI applications	• 4 vCPUs • 24 GB RAM • 200 GB hard disk storage
csp-regional-sblb	Load balancer for device to Fault Management Performance Management (FMPM) microservice connectivity	• 4 vCPUs • 8 GB RAM • 300 GB hard disk storage
csp-space-vm	Junos Space Virtual Appliance and database—required only if you deploy virtualized network functions (VNFs) that use this EMS	• 4 vCPUs • 16 GB RAM • 200 GB hard disk storage
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 • 48 GB RAM • 500 GB hard disk storage
csp-vrr-vm	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage
csp-central-k8mastervm	Central K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage
csp-regional-k8mastervm	Regional K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage

NOTE: For non-HA trial configurations, we recommend one server with 48 vCPUs and 256 GB RAM. Non-HA trial configurations have been validated with a server with 24 vCPUs and 256GB RAM, but performance issues may occur over longer periods of time.

Table 14 on page 43 shows complete details about the VMs for a trial environment with HA.

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

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-installer-vm	—	• 4 vCPUs • 48 GB RAM • 300 GB hard disk storage

Copyright © 2018, Juniper Networks, Inc.

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

csp-central-infravm1	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage
csp-central-infravm2	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage
csp-central-infravm3	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage
csp-central-lbvm1	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-central-lbvm2	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-central-lbvm3	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-central-msvm1	All microservices, including GUI applications	• 8 CPUs • 64 GB RAM • 500 GB hard disk storage
csp-central-msvm2	All microservices, including GUI applications	• 8 vCPUs • 64 GB RAM • 500 GB hard disk storage
csp-central-msvm3	All microservices, including GUI applications	• 8 vCPUs • 64 GB RAM • 500 GB hard disk storage
csp-regional-infravm1	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage
csp-regional-infravm2	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage
csp-regional-infravm3	Third-party applications used as infrastructure services	• 4 vCPUs • 32 GB RAM • 500 GB hard disk storage

Copyright © 2018, Juniper Networks, Inc.

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

csp-regional-msvm1	All microservices, including GUI applications	• 8 CPUs • 32 GB RAM • 500 GB hard disk storage
csp-regional-msvm2	All microservices, including GUI applications	• 8 CPUs • 32 GB RAM • 500 GB hard disk storage
csp-regional-msvm3	All microservices, including GUI applications	• 8 CPUs • 32 GB RAM • 500 GB hard disk storage
csp-regional-lbvm1	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-regional-lbvm2	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-regional-lbvm3	Load-balancing applications	• 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
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
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 • 48 GB RAM • 300 GB hard disk storage
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	• 4 vCPUs • 24 GB RAM • 300 GB hard disk storage
csp-vrr-vm	Virtual route reflector (VRR)	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage

Table 15 on page 46 shows complete details about the VMs required for a production environment without HA.

Copyright © 2018, Juniper Networks, Inc.

Table 15: Details of VMs for a Production Environment Without HA

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-installer-vm	—	• 4 vCPUs • 64 GB RAM • 300 GB hard disk storage
csp-central-infravm	Third -party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage
csp-central-msvm	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage
csp-regional-infravm	Third -party applications used as infrastructure services	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage
csp-regional-msvm	All microservices, including GUI applications	• 16 vCPUs • 64 GB RAM • 500 GB hard disk storage
csp-regional-sblb	Load balancer for device to microservice connectivity	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage
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
csp-central-elkvm	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage
csp-regional-elkvm	Logging applications	• 4 vCPUs • 32 GB RAM • 300 GB hard disk storage
csp-vrr-vm	VRR	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage
csp-central-k8mastervm	Central K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage

Copyright © 2018, Juniper Networks, Inc.

Table 15: Details of VMs for a Production Environment Without HA (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-regional-k8mastervm	Regional K8 Master VM	• 4 vCPUs • 8 GB RAM • 200 GB hard disk storage
csp-contrailanalytics-1	Contrail Analytics for a distributed deployment. For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform.	• 4 vCPUs • 256 GB RAM • 300 GB hard disk storage

Table 16 on page 47 shows complete details about the VMs for a production environment with HA.

Table 16: Details of VMs for a Production Environment with HA

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

Copyright © 2018, Juniper Networks, Inc.

Table 16: Details of VMs for a Production Environment with HA (continued)

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

Copyright © 2018, Juniper Networks, Inc.

Table 16: Details of VMs for a Production Environment with HA (continued)

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

Table 17 on page 50 shows the ports that must be open on all CSO VMs to enable the following types of CSO communications:

- External—CSO user interface (UI) and CPE connectivity

Copyright © 2018, Juniper Networks, Inc.

- Internal—Between CSO components

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

Table 17: Ports to Open on CSO VMs

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

Copyright © 2018, Juniper Networks, Inc.

Table 17: Ports to Open on CSO VMs (continued)

5543	Internal	Logstash UDP
5601	External	Kibana UI
5665	Internal	Icinga API
5666	Internal	icinga nrpe
5671	Internal	RabbitMQ SSL listener
5672	Internal	RabbitMQ client
6000	Internal	Swift Object Server
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
8529	Internal	ArangoDB
9042	Internal	Cassandra native transport
9090	Internal	Swift Proxy Server
9091	Internal	xmltec-xmlmail tcp
9101	External and internal	HA proxy exporter

Copyright © 2018, Juniper Networks, Inc.

Table 17: Ports to Open on CSO VMs (continued)

9102	Internal	jetdirect
9160	Internal	Cassandra
9200	Internal	Elasticsearch
10248	Internal	kubelet healthz
15100	Internal	Logstash TCP
15672	Internal	RabbitMQ management
30000-32767	Internal	Kubernetes service node range
30900	External	Prometheus
35357	Internal	Keystone private

Related Documentation

- Hardware and Software Required for Contrail Service Orchestration on page 36
- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76

VNFs Supported by the Cloud CPE Solution

The Cloud CPE solution supports the Juniper Networks and third-party VNFs listed in Table 18 on page 52.

Table 18: VNFs Supported by Contrail Service Orchestration

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

Copyright © 2018, Juniper Networks, Inc.

Table 18: VNFs Supported by Contrail Service Orchestration (continued)

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

You must upload VNFs to the Contrail Cloud Platform for the centralized deployment after you install the Cloud CPE solution. You upload the VNF images for the distributed deployment through Administration Portal or API calls.

You can use these VNFs in service chains and configure some settings for VNFs for a service chain in Network Service Designer. You can then view those configuration settings for a network service in Administration Portal. Customers can also configure some settings for the VNFs in their network services through Customer Portal. VNF configurations that customers specify in Customer Portal override VNF configurations that the person who designs network services specifies in Network Service Designer.

Related Documentation

- Uploading the vSRX VNF Image for a Centralized Deployment on page 132
- Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 133
- Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 135
- Installing Licenses with the License Tool on page 154

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 3

Installing and Configuring the Network Devices and Servers for a Centralized Deployment

Cabling the Hardware for the Centralized Deployment

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

To cable the hardware:

1. Connect cables from the EX Series switch to the other devices in the network.

   See Table 19 on page 56 for information about the connections for the EX Series switch.

2. Connect cables from the QFX Series switch to the other devices in the network.

   See Table 20 on page 56 for information about the connections for the QFX Series switch.

3. Connect cables from the MX Series router to the other devices in the network.

   See Table 21 on page 57 for information about the connections for the MX Series router.

Copyright © 2018, Juniper Networks, Inc.

Table 19: Connections for EX Series Switch

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

Table 20: Connections for QFX Series Switch

Interface on QFX Series Switch	Destination Device	Interface on Destination Device
eth0 (management interface)	EX Series switch	ge-0/0/42

Copyright © 2018, Juniper Networks, Inc.

Table 20: Connections for QFX Series Switch (continued)

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

Table 21: Connections for MX Series Router

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- NFV in the Cloud CPE Solution on page 18
- Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 58
- Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 59
- Configuring the MX Series Router in the Contrail Cloud Implementation for a Centralized Deployment on page 61
- Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment on page 64

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

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

To configure the EX Series switch:

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

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

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

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

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

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


Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Hardware and Software Required for Contrail Service Orchestration on page 36
- Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 59
- Configuring the MX Series Router in the Contrail Cloud Implementation for a Centralized Deployment on page 61

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

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

To configure the QFX Series switch:

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

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

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

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

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

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

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

```
user@switch# set interfaces xe-0/0/2 ether-options 802.3ad ae2
user@switch# set interfaces xe-0/0/22 ether-options 802.3ad ae2
user@switch# set interfaces ae2 mtu 9192
user@switch# set interfaces ae2 aggregated-ether-options lacp active
user@switch# set interfaces ae2 aggregated-ether-options lacp periodic fast
```

Copyright © 2018, Juniper Networks, Inc.

```
user@switch# set interfaces ae2 unit 0 family ethernet-switching interface-mode
  access
user@switch# set interfaces ae2 unit 0 family ethernet-switching vlan members data
```

```
user@switch# set interfaces xe-0/0/3 ether-options 802.3ad ae3
user@switch# set interfaces xe-0/0/23 ether-options 802.3ad ae3
user@switch# set interfaces ae3 mtu 9192
user@switch# set interfaces ae3 aggregated-ether-options lacp active
user@switch# set interfaces ae3 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae3 unit 0 family ethernet-switching interface-mode
  access
user@switch# set interfaces ae3 unit 0 family ethernet-switching vlan members data
```

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

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

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

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

```
user@switch# set interfaces xe-0/0/8 ether-options 802.3ad ae8
user@switch# set interfaces xe-0/0/28 ether-options 802.3ad ae8
user@switch# set interfaces ae8 mtu 9192
```


Copyright © 2018, Juniper Networks, Inc.

```
user@switch# set interfaces ae8 aggregated-ether-options lacp active
user@switch# set interfaces ae8 aggregated-ether-options lacp periodic fast
user@switch# set interfaces ae8 unit 0 family ethernet-switching interface-mode
  access
user@switch# set interfaces ae8 unit 0 family ethernet-switching vlan members data
```

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

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

5. Configure OSPF routing. For example:

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

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

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

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

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

Related Documentation

- Hardware and Software Required for Contrail Service Orchestration on page 36
- Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 58
- Configuring the MX Series Router in the Contrail Cloud Implementation for a Centralized Deployment on page 61

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

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

To configure the MX Series router:

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

```
user@router# set interfaces ge-1/0/0 unit 0 family inet address 10.87.24.77/28
user@router# set interfaces lo0 unit 0 family inet address 172.16.100.1/32
```

Copyright © 2018, Juniper Networks, Inc.

```
user@router# set routing-options route-distinguisher-id 172.16.100.1
user@router# set routing-options autonomous-system 64512
user@router# set protocols ospf area 0.0.0.0 interface lo0.0
```

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

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

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

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

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

4. Set up routing. For example:

```
user@router# set routing-options static rib-group inet-to-public
user@router# set routing-options static route 0.0.0.0/0 next-hop 10.87.24.78
user@router# set routing-options static route 0.0.0.0/0 retain
user@router# set routing-options static route 10.87.24.64/26 next-table public.inet.0
user@router# set routing-options rib-groups inet-to-public import-rib inet.0
user@router# set routing-options rib-groups inet-to-public import-rib public.inet.0
user@router# set routing-options rib-groups inet-to-public import-policy
  leak-default-only
```


Copyright © 2018, Juniper Networks, Inc.

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

5. Configure NAT. For example:

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Hardware and Software Required for Contrail Service Orchestration on page 36
- Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 58
- Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 59

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

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

To install Contrail OpenStack:

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

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

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

   Refer to the Contrail documentation for information about installing Contrail OpenStack and configuring the nodes.
7. For each node, use the ETCD keys to specify the same username and password for Contrail.

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 4

Installing and Configuring the Network Devices and Servers for a Distributed Deployment or SD-WAN Solution

Configuring the Physical Servers in a Distributed Deployment

For a distributed deployment, you must configure the Contrail Service Orchestration (CSO) and Contrail Analytics servers (or nodes, if you are using a node server) before you run the installer.

To configure the servers:

1. Configure hostnames for the physical servers.
2. Configure IP addresses for the Ethernet management ports of the physical servers.
3. Configure DNS on the physical servers, and ensure that DNS is working correctly.
4. Configure Internet access for the physical servers and nodes.
5. From each server and node, verify that you can ping the IP addresses and hostnames of all the other servers and nodes in the distributed deployment.
6. For a production environment, install Contrail OpenStack on the Contrail Analytics server.

   Refer to the Contrail documentation for information about installing Contrail OpenStack.

Related Documentation

- Hardware and Software Required for Contrail Service Orchestration on page 36

Copyright © 2018, Juniper Networks, Inc.

- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76

Configuring the MX Series Router in a Distributed Deployment

You need to configure interfaces, virtual routing and forwarding instances (VRFs), and DHCP on the MX Series router with Junos OS. You can, however, use Administration Portal to specify configuration settings for both endpoints of the required IPSec tunnel between the MX Series router and the NFX250 with Administration Portal. When the NFX250 becomes operational, Contrail Service Orchestration (CSO) components set up the tunnel.

Copyright © 2018, Juniper Networks, Inc.

To configure the MX Series router in Junos OS:

1. Configure the interfaces on the MX Series router.

   For example:

```
ge-0/3/7 {
        description "to nfx wan0 i.e. ge-0/0/10";
        vlan-tagging;
        unit 10 {
            description "NFX WAN_0 data";
            vlan-id 10;
            family inet {
                address 195.195.195.1/24;
            }
        }
        unit 20 {
            description "NFX WAN_0 OAM";
            vlan-id 20;
            family inet {
                address 196.196.196.254/24;
            }
        }
    }
    ge-0/3/8 {
        description "to nfx wan1 i.e. ge-0/0/11 FOR IPSEC";
        unit 0 {
            family inet {
                address 198.198.198.1/24;
            }
        }
    }
```

2. Configure a VRF for Operation, Administration, and Maintenance (OAM) traffic between

Copyright © 2018, Juniper Networks, Inc.

Contrail Service Orchestration and the NFX250.

For example:

```
nfx-oam {
        instance-type vrf;
        interface ge-0/0/0.220;
        vrf-target target:64512:10000;
        vrf-table-label;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop 192.168.220.2;
            }
        }
    }
```

3. Configure a VRF for data traffic that travels over the wide area network (WAN).

   Data that travels through the IPSec tunnel also uses this VRF. When you configure the MX endpoint of the IPSec tunnel in Administration Portal, you specify these VRF settings.

   For example:

```
nfx-data {
        instance-type vrf;
        interface ge-0/3/7.10;
        vrf-target target:64512:10001;
        vrf-table-label;
        protocols {
            bgp {
                group nfx-gwr-bgp-grp {
                    type external;
                    family inet {
                        unicast;
                    }
```


Copyright © 2018, Juniper Networks, Inc.

```
                    export send-direct;
                    peer-as 65000;
                    neighbor 195.195.195.2;
                }
            }
        }
    }
```

4. Configure DHCP on the MX Series router.

```
System{
    Services {
        dhcp-local-server {
            group 8-csp-gpr {
                interface ge-0/3/8.0;
            }
        }
    }


access {
    address-assignment {
        pool 8-csp-gpr-pool {
            family inet {
                network 198.198.198.0/24;
                range valid {
                    low 198.198.198.5;
                    high 198.198.198.250;
                }
                dhcp-attributes {
                    domain-name juniper.net;
                    name-server {
```

Copyright © 2018, Juniper Networks, Inc.

```
                        8.8.8.8;
                    }
                }
            }
        }
    }
}
```

Related Documentation

- Hardware and Software Required for Contrail Service Orchestration on page 36
- Topology of the Cloud CPE and SD-WAN Solutions on page 22
- Configuring the Physical Servers in a Distributed Deployment on page 67

Installing and Setting Up CPE Devices

- Preparing for CPE Device Activation on page 72
- Installing and Configuring an NFX250 Device on page 72
- Installing and Configuring an SRX Series Services Gateway or vSRX Instance as a CPE Device on page 73

Preparing for CPE Device Activation

Before customers can activate a CPE device, you must complete the following tasks:

- Specify activation data with Administration Portal or the API for each CPE device, such as:
  - The name of the site for the device
  - The serial number
  - The activation code (NFX250 devices only)

Installing and Configuring an NFX250 Device

An administrator at the customer's site installs the NFX250 and performs the initial software configuration for the NFX250. These are straightforward tasks that involve a limited amount of hardware installation, cabling, and software configuration. See the NFX Series documentation for more information.

When the administrator completes the initial configuration process, the NFX250 device obtains a boot image and configuration image from its regional server and becomes operational.

Copyright © 2018, Juniper Networks, Inc.

Installing and Configuring an SRX Series Services Gateway or vSRX Instance as a CPE Device

An administrator at the customer's site installs and configures an SRX Series Services Gateway or a vSRX instances as a CPE device using the following workflow:

1. Install the hardware and cable the device.
2. Power on the device and access the device console.
3. Log in to Customer Portal and perform the following tasks:
   - Add the site to the network.
   - Apply the initial configuration to the device.
   - Activate the CPE device.

Related Documentation

- Setting Up a Distributed Deployment on page 165
- NFX Series documentation
- SRX Series documentation
- vSRX documentation

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 5

Installing and Configuring Contrail Service Orchestration

Removing a Previous Deployment

You can remove the existing virtual machines (VMs) and perform a completely new installation. This approach makes sense if the architecture of the VMs on the Contrail Service Orchestration node or server has changed significantly between releases.

To remove a previous installation:

1. Remove VMs on the physical server.
   a. Log in to the CSO node or server as root.
   b. View the list of VMs.

      For example:

      ```
      root@host:~/# virsh list --all
      ```

      This command lists the existing VMs.

      ```
      Id  Name   State
      2  csp-ui-vm  running
      ```

Copyright © 2018, Juniper Networks, Inc.

c. Remove each VM and its contents.

For example:

```
root@host:~/# virsh destroy csp-ui-vm
root@host:~/# virsh undefine csp-ui-vm
```

Where, *csp-ui-vm* is the name of VM you want to delete.

d. Delete the Ubuntu source directories and VM.

For example:

```
root@host:~/# rm -rf /root/disks
root@host:~/# rm -rf /root/disks_can
root@host:~/# cd /root/ubuntu_vm
root@host:~/# rm -rf
```

2. Delete the Salt server keys.

For example:

```
root@host:~/# salt-key -D
```

Related Documentation

- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76

Provisioning VMs on Contrail Service Orchestration Nodes or Servers

Virtual Machines (VMs) on the central and regional Contrail Service Orchestration (CSO) nodes or servers host the infrastructure services and some other components. All servers and VMs for the solution should be in the same subnet. To set up the VMs, you can:

- Use the provisioning tool to create and configure the VMs if you use the KVM hypervisor or ESXi VMware on a CSO node or server.

  The tool also installs Ubuntu in the VMs.

- Manually configure Virtual Route Reflector (VRR) VMs on a CSO node or server, if you use the ESXi VMware VM.

The VMs required on a CSO node or server depend on whether you configure:

- A trial environment without high availability (HA). .
- A production environment without HA.
- A trial environment with HA.
- A production environment with HA.

See "Minimum Requirements for Servers and VMs" on page 40 for details of the VMs and associated resources required for each environment.

Copyright © 2018, Juniper Networks, Inc.

The following sections describe the procedures for provisioning the VMs:

Before You Begin

Before you begin you must:

- Configure the physical servers or node servers and nodes.
- The operating system for physical servers must be Ubuntu 14.04.5 LTS.
- For a centralized deployment, configure the Contrail Cloud Platform and install Contrail OpenStack.

Downloading the Installer

To download the installer package:

1. Log in as root to the central CSO node or server.

   The current directory is the home directory.

2. Download the appropriate installer package from https://www.juniper.net/support/downloads/?p=cso#sw.

   - Use the Contrail Service Orchestration installer if you purchased licenses for a centralized deployment or both Network Service Orchestrator and Network Service Controller licenses for a distributed deployment.

     This option includes all the Contrail Service Orchestration graphical user interfaces (GUIs).

   - Use the Network Service Controller installer if you purchased only Network Service Controller licenses for a distributed deployment or SD-WAN implementation.

     This option includes Administration Portal and Service and Infrastructure Monitor, but not the Designer Tools.

3. Expand the installer package, which has a name specific to its contents and the release. For example, if the name of the installer package is **csoVersion.tar.gz**:

   ```
   root@host:~/# tar –xvzf csoVersion.tar.gz
   ```

Copyright © 2018, Juniper Networks, Inc.

The expanded package is a directory that has the same name as the installer package and contains the installation files.

Creating a Bridge Interface for KVM

If you use the KVM hypervisor, before you create VMs, you must create a bridge interface on the physical server that maps the primary network interface (Ethernet management interface) on each CSO node or server to a virtual interface. This action enables the VMs to communicate with the network.

To create the bridge interface:

1. Log in as root on the central CSO node or server.

2. Update the index files of the software packages installed on the server to reference the latest versions.

   ```
   root@host:~/# apt-get update
   ```

3. View the network interfaces configured on the server to obtain the name of the primary interface on the server.

   ```
   root@host:~/# ifconfig
   ```

4. Install the libvirt software.

   ```
   root@host:~/# apt-get install libvirt-bin
   ```

5. View the list of network interfaces, which now includes the virtual interface virbr0.

   ```
   root@host:~/# ifconfig
   ```

6. Open the file **/etc/network/interfaces** and modify it to map the primary network interface to the virtual interface virbr0.

   For example, use the following configuration to map the primary interface eth0 to the virtual interface virbr0:

   ```
   # This file describes the network interfaces available on your system
   # and how to activate them. For more information, see interfaces (5).
   # The loopback network interface
   auto lo
   iface lo inet loopback

   # The primary network interface
   auto eth0
   iface eth0 inet manual
       up ifconfig eth0 0.0.0.0 up

   auto virbr0
   iface virbr0 inet static
   ```


Copyright © 2018, Juniper Networks, Inc.

```
        bridge_ports eth0
        address 192.168.1.2
        netmask 255.255.255.0
        network 192.168.1.0
        broadcast 192.168.1.255
        gateway 192.168.1.1
        dns-nameservers 8.8.8.8
        dns-search example.net
```

7. Modify the default virtual network by customizing the file **default.xml**:

   a. Customize the IP address and subnet mask to match the values for the virbr0 interface in the file **/etc/network/interfaces**

   b. Turn off the Spanning Tree Protocol (STP) option.

   c. Remove the NAT and DHCP configurations.

   For example:

   ```
   root@host:~/# virsh net-edit default
   ```

   Before modification:

   ```
   <network>
        <name>default</name>
        <uuid>0f04ffd0-a27c-4120-8873-854bbfb02074</uuid>
         <forward mode='nat'/>
         <bridge name='virbr0' stp='on' delay='0'/>
         <ip address='192.168.1.2' netmask='255.255.255.0'>
           <dhcp>
             <range start='192.168.1.1' end='192.168.1.254'/>
           </dhcp>
         </ip>
   </network>
   ```

   After modification:

   ```
   <network>
        <name>default</name>
        <uuid>0f04ffd0-a27c-4120-8873-854bbfb02074</uuid>
        <bridge name='virbr0' stp='off' delay='0'/>
        <ip address='192.168.1.2' netmask='255.255.255.0'>
        </ip>
   </network>
   ```

8. Reboot the physical machine and log in as root again.

9. Verify that the primary network interface is mapped to the virbr0 interface.

   ```
   root@host:~/# brctl show
   ```

Copyright © 2018, Juniper Networks, Inc.

```
bridge name   bridge id       STP enabled interfaces
  virbr0      8000.0cc47a010808   no      em1
                                          vnet1
                                          vnet2
```

See Also •

Creating a Data Interface for a Distributed Deployment

For a distributed deployment, you create a second bridge interface that the VMs use to send data communications to the CPE device.

To create a data interface:

1. Log into the central CSO server as root.

2. Configure the new virtual interface and map it to a physical interface.

   For example:

   ```
   root@host:~/# virsh brctl addbr ex: virbr1
   root@host:~/# virsh brctl addif virbr1 eth1
   ```

3. Create an xml file with the name **virbr1.xml** in the directory **/var/lib/libvirt/network**.

4. Paste the following content into the **virbr1.xml** file, and edit the file to match the actual settings for your interface.

   For example:

   ```
   <network>
                 <name>default</name>
                 <uuid>0f04ffd0-a27c-4120-8873-854bbfb02074</uuid>
                 <bridge name='virbr1' stp='off' delay='0'/>
                 <ip address='192.0.2.1' netmask='255.255.255.0'>
                 </ip>
   </network>
   ```

5. Open the **/etc/network/interfaces** file and add the details for the second interface.

   For example:

   ```
   # This file describes the network interfaces available on your system
   # and how to activate them. For more information, see interfaces (5).
   # The loopback network interface
   auto lo
   iface lo inet loopback

   # The primary network interface
   ```


Copyright © 2018, Juniper Networks, Inc.

```
auto eth0
iface eth0 inet manual
  up ifconfig eth0 0.0.0.0 up

auto eth1
iface eth1 inet manual
  up ifconfig eth1 0.0.0.0 up

auto virbr0
iface virbr0 inet static
         bridge_ports eth0
         address 192.168.1.2
         netmask 255.255.255.0
         network 192.168.1.0
         broadcast 192.168.1.255
         gateway 192.168.1.1
         dns-nameservers 8.8.8.8
         dns-search example.net
auto virbr1
iface virbr1 inet static
    bridge_ports eth1
    address 192.168.1.2
    netmask 255.255.255.0
```

6. Reboot the server.

7. Verify that the secondary network interface, eth1, is mapped to the second interface.

```
root@host:~/# brctl show

bridge name   bridge id       STP enabled interfaces
  virbr0      8000.0cc47a010808   no      em1
                                          vnet1
                                          vnet2
  virbr1      8000.0cc47a010809   no      em2
                                          vnet0
```

8. Configure the IP address for the interface.

   You do not specify an IP address for the data interface when you create it.

Customizing the Configuration File for the Provisioning Tool

The provisioning tool uses a configuration file, which you must customize for your network. The configuration file is in YAML format.

To customize the configuration file:

1. Log in as root to the central CSO node or server.

2. Access the **confs** directory that contains the example configuration files. For example, if the name of the installer directory is **csoVersion**

```
root@host:~/# cd csoVersion/confs
```

Copyright © 2018, Juniper Networks, Inc.

3. Access the directory for the environment that you want to configure.

   Table 22 on page 82 shows the directories that contain the example configuration file.

Table 22: Location of Configuration Files for Provisioning VMs

Environment	Directory for Example Configuration File
Trial environment without HA	**cso3.3/trial/nonha/provisionvm**
Production environment without HA	**cso3.3/production/nonha/provisionvm**
Trial environment with HA	**cso3.3/trial/ha/provisionvm**
Production environment with HA	**cso3.3/production/ha/provisionvm**

4. Make a copy of the example configuration file in the **/confs** directory and name it **provision_vm.conf**.

   For example:

   ```
   root@host:~/cspVersion/confs# cp
   /cso3.3/trial/nonha/provisionvm/provision_vm_example.conf provision_vm.conf
   ```

5. Open the file **provision_vm.conf** with a text editor.

6. In the [TARGETS] section, specify the following values for the network on which CSO resides.

   - **installer_ip**—IP address of the management interface of the host on which you deployed the installer.
   - **ntp_servers**—Comma-separated list of fully qualified domain names (FQDN) of Network Time Protocol (NTP) servers. For networks within firewalls, specify NTP servers specific to your network.
   - **physical**—Comma-separated list of hostnames of the CSO nodes or servers.
   - **virtual**—Comma-separated list of names of the virtual machines (VMs) on the CSO servers.

7. Specify the following configuration values for each CSO node or server that you specified in Step 6.

   - **[hostname]**—Hostname of the CSO node or server
   - **management_address**—IP address of the Ethernet management (primary) interface in classless Internet domain routing (CIDR) notation
   - **management_interface**—Name of the Ethernet management interface, virbr0
   - **gateway**—IP address of the gateway for the host

Copyright © 2018, Juniper Networks, Inc.

- **dns_search**—Domain for DNS operations
- **dns_servers**—Comma-separated list of DNS name servers, including DNS servers specific to your network
- **hostname**—Hostname of the node
- **username**—Username for logging in to the node
- **password**—Password for logging in to the node
- **data_interface**—Name of the data interface. Leave blank for a centralized deployment and specify the name of the data interface, such as virbr1, that you configured for a distributed deployment.

8. Except for the Junos Space Virtual Appliance and VRR VMs, specify configuration values for each VM that you specified in Step 6.
   - **[VM name]**—Name of the VM
   - **management_address**—IP address of the Ethernet management interface in CIDR notation
   - **hostname**—Fully qualified domain name (FQDN) of the VM
   - **username**—Login name of user who can manage all VMs
   - **password**—Password for user who can manage all VMs
   - **local_user**—Login name of user who can manage this VM
   - **local_password**—Password for user who can manage this VM
   - **guest_os**—Name of the operating system
   - **host_server**—Hostname of the CSO node or server
   - **memory**—Required amount of RAM in GB
   - **vCPU**—Required number of virtual central processing units (vCPUs)
   - **enable_data_interface**—True enables the VM to transmit data and false prevents the VM from transmitting data. The default is false.
9. For the Junos Space VM, specify configuration values for each VM that you specified in Step 6.
   - **[VM name]**—Name of the VM.
   - **management_address**—IP address of the Ethernet management interface in CIDR notation.
   - **web_address**—Virtual IP (VIP) address of the primary Junos Space Virtual Appliance. (Setting only required for the VM on which the primary Junos Space Virtual Space appliance resides.)
   - **gateway**—IP address of the gateway for the host. If you do not specify a value, the value defaults to the gateway defined for the CSO node or server that hosts the VM.
   - **nameserver_address**—IP address of the DNS nameserver.

Copyright © 2018, Juniper Networks, Inc.

- **hostname**—FQDN of the VM.
- **username**—Username for logging in to Junos Space.
- **password**—Default password for logging in to Junos Space.
- **newpassword**—Password that you provide when you configure the Junos Space appliance.
- **guest_os**—Name of the operating system.
- **host_server**—Hostname of the CSO node or server.
- **memory**—Required amount of RAM in GB.
- **vCPU**—Required number of virtual central processing units (vCPUs).

10. Save the file.

11. Run the following command to start virtual machines.

```
root@host:~/# ./provision_vm.sh
```

The following examples show customized configuration files for the different deployments:

- Trial environment without HA (see Sample Configuration File for Provisioning VMs in a Trial Environment without HA on page 84).
- Production environment without HA (see Sample Configuration File for Provisioning VMs in a Production Environment Without HA on page 87).
- Trial environment with HA (see Sample Configuration File for Provisioning VMs in a Trial Environment with HA on page 90).
- Production environment with HA (see Table 16 on page 47).

Sample Configuration File for Provisioning VMs in a Trial Environment without HA

```
# This config file is used to provision KVM-based virtual machines using lib virt
 manager.

[TARGETS]
# Mention primary host (installer host) management_ip

installer_ip =

ntp_servers = ntp.juniper.net

# The physical server where the Virtual Machines should be provisioned
# There can be one or more physical servers on
# which virtual machines can be provisioned
physical = cso-host

# The list of virtual servers to be provisioned.
server = csp-central-infravm, csp-central-msvm, csp-central-k8mastervm,
csp-regional-infravm, csp-regional-msvm, csp-regional-k8mastervm, csp-installer-vm,
 csp-contrailanalytics-1, csp-vrr-vm, csp-regional-sblb
```


Copyright © 2018, Juniper Networks, Inc.

```
# Physical Server Details
[cso-host]
management_address = 192.168.1.2/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-host
username = root
password = passw0rd
data_interface =

# VM Details

[csp-central-infravm]
management_address = 192.168.1.4/24
hostname = centralinfravm.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 49152
vcpu = 8
enable_data_interface = false

[csp-central-msvm]
management_address = 192.168.1.5/24
hostname = centralmsvm.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 49152
vcpu = 8
enable_data_interface = false

[csp-central-k8mastervm]
management_address = 192.168.1.14/24
hostname = centralk8mastervm.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 8192
vcpu = 4
enable_data_interface = false

[csp-regional-infravm]
management_address = 192.168.1.6/24
hostname = regionalinfravm.example.net
username = root
password = passw0rd
```

Copyright © 2018, Juniper Networks, Inc.

```
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 24576
vcpu = 4
enable_data_interface = false

[csp-regional-msvm]
management_address = 192.168.1.7/24
hostname = regionalmsvm.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 24576
vcpu = 4
enable_data_interface = false

[csp-regional-k8mastervm]
management_address = 192.168.1.15/24
hostname = regionalk8mastervm.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 8192
vcpu = 4
enable_data_interface = false

[csp-installer-vm]
management_address = 192.168.1.10/24
hostname = installervm.example.net
username = root
password = passw0rd
local_user = installervm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 24576
vcpu = 4
enable_data_interface = false

[csp-contrailanalytics-1]
management_address = 192.168.1.11/24
hostname = canvm.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 49152
vcpu = 8
enable_data_interface = false
```


Copyright © 2018, Juniper Networks, Inc.

```
[csp-regional-sblb]
management_address = 192.168.1.12/24
hostname = regional-sblb.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host
memory = 8192
vcpu = 4
enable_data_interface = true


[csp-vrr-vm]
management_address = 192.168.1.13/24
hostname = vrr.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-host
memory = 8192
vcpu = 4


[csp-space-vm]
management_address = 192.168.1.14/24
web_address = 192.168.1.15/24
gateway = 192.168.1.1
nameserver_address = 192.168.1.254
hostname = spacevm.example.net
username = admin
password = abc123
newpassword = jnpr123!
guest_os = space
host_server = cso-host
memory = 16384
vcpu = 4
```

Sample Configuration File for Provisioning VMs in a Production Environment Without HA

```
# This config file is used to provision KVM-based virtual machines using lib virt
 manager.

[TARGETS]
# Mention primary host (installer host) management_ip

installer_ip =

ntp_servers = ntp.juniper.net

# The physical server where the Virtual Machines should be provisioned
# There can be one or more physical servers on
# which virtual machines can be provisioned
physical = cso-central-host, cso-regional-host

# Note: Central and Regional physical servers are used as "csp-central-ms" and
"csp-regional-ms" servers.
```

Copyright © 2018, Juniper Networks, Inc.

```
# The list of servers to be provisioned and mention the contrail analytics servers
 also in "server" list.
server = csp-central-infravm, csp-regional-infravm, csp-installer-vm, csp-space-vm,
 csp-contrailanalytics-1, csp-central-elkvm, csp-regional-elkvm, csp-central-msvm,
 csp-regional-msvm, csp-vrr-vm, csp-regional-sblb


# Physical Server Details
[cso-central-host]
management_address = 192.168.1.2/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-central-host
username = root
password = passw0rd
data_interface =

[cso-regional-host]
management_address = 192.168.1.3/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-regional-host
username = root
password = passw0rd
data_interface =


[csp-contrailanalytics-1]
management_address = 192.168.1.9/24
management_interface =
hostname = canvm.example.net
username = root
password = passw0rd
vm = false


# VM Details

[csp-central-infravm]
management_address = 192.168.1.4/24
hostname = centralinfravm.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-infravm]
management_address = 192.168.1.5/24
hostname = regionalinfravm.example.net
username = root
```


Copyright © 2018, Juniper Networks, Inc.

```
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-space-vm]
management_address = 192.168.1.6/24
web_address = 192.168.1.7/24
gateway = 192.168.1.1
nameserver_address = 192.168.1.254
hostname = spacevm.example.net
username = admin
password = abc123
newpassword = jnpr123!
guest_os = space
host_server = cso-regional-host
memory = 32768
vcpu = 4

[csp-installer-vm]
management_address = 192.168.1.8/24
hostname = installer.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host
memory = 65536
vcpu = 4
enable_data_interface = false

[csp-central-elkvm]
management_address = 192.168.1.10/24
hostname = centralelkvm.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host
memory = 32768
vcpu = 4
enable_data_interface = false


[csp-regional-elkvm]
management_address = 192.168.1.11/24
hostname = regionalelkvm.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host
memory = 32768
```

Copyright © 2018, Juniper Networks, Inc.

```
vcpu = 4
enable_data_interface = false

[csp-central-msvm]
management_address = 192.168.1.12/24
hostname = centralmsvm.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-msvm]
management_address = 192.168.1.13/24
hostname = regionalmsvm.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-sblb]
management_address = 192.168.1.14/24
hostname = regional-sblb.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host
memory = 32768
vcpu = 4
enable_data_interface = true


[csp-vrr-vm]
management_address = 192.168.1.15/24
hostname = vrr.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-regional-host
memory = 8192
vcpu = 4
```

Sample Configuration File for Provisioning VMs in a Trial Environment with HA

```
# This config file is used to provision KVM-based virtual machines using lib virt
 manager.

[TARGETS]
```


Copyright © 2018, Juniper Networks, Inc.

```
# Mention primary host (installer host) management_ip

installer_ip =

ntp_servers = ntp.juniper.net

# The physical server where the Virtual Machines should be provisioned
# There can be one or more physical servers on
# which virtual machines can be provisioned
physical = cso-host1, cso-host2, cso-host3


# The list of virtual servers to be provisioned.
server = csp-central-infravm1, csp-central-infravm2, csp-central-infravm3,
csp-central-msvm1, csp-central-msvm2, csp-central-msvm3, csp-regional-infravm1,
csp-regional-infravm2, csp-regional-infravm3, csp-regional-msvm1,
csp-regional-msvm2, csp-regional-msvm3, csp-contrailanalytics-1, csp-central-lbvm1,
 csp-central-lbvm2, csp-central-lbvm3, csp-regional-lbvm1, csp-regional-lbvm2,
csp-regional-lbvm3, csp-space-vm, csp-installer-vm, csp-vrr-vm1, csp-vrr-vm2,
csp-regional-sblb1, csp-regional-sblb2


# Physical Server Details
[cso-host1]
management_address = 192.168.1.2/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-host1
username = root
password = passw0rd
data_interface =

[cso-host2]
management_address = 192.168.1.3/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-host2
username = root
password = passw0rd
data_interface =


[cso-host3]
management_address = 192.168.1.4/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-host3
username = root
password = passw0rd
data_interface =

# VM Details

[csp-central-infravm1]
```

Copyright © 2018, Juniper Networks, Inc.

```
management_address = 192.168.1.5/24
hostname = centralinfravm1.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-infravm2]
management_address = 192.168.1.6/24
hostname = centralinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-infravm3]
management_address = 192.168.1.7/24
hostname = centralinfravm3.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-msvm1]
management_address = 192.168.1.8/24
hostname = centralmsvm1.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 65536
vcpu = 8
enable_data_interface = false

[csp-central-msvm2]
management_address = 192.168.1.9/24
hostname = centralmsvm2.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
```


Copyright © 2018, Juniper Networks, Inc.

```
memory = 65536
vcpu = 8
enable_data_interface = false

[csp-central-msvm3]
management_address = 192.168.1.9/24
hostname = centralmsvm3.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 65536
vcpu = 8
enable_data_interface = false

[csp-regional-infravm1]
management_address = 192.168.1.10/24
hostname = regionalinfravm1.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-infravm2]
management_address = 192.168.1.11/24
hostname = regionalinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-infravm3]
management_address = 192.168.1.12/24
hostname = regionalinfravm3.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-msvm1]
management_address = 192.168.1.13/24
hostname = regionalmsvm1.example.net
username = root
```

Copyright © 2018, Juniper Networks, Inc.

```
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 32768
vcpu = 8
enable_data_interface = false

[csp-regional-msvm2]
management_address = 192.168.1.14/24
hostname = regionalmsvm2.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 32768
vcpu = 8
enable_data_interface = false

[csp-regional-msvm3]
management_address = 192.168.1.14/24
hostname = regionalmsvm3.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 32768
vcpu = 8
enable_data_interface = false

[csp-space-vm]
management_address = 192.168.1.15/24
web_address = 192.168.1.16/24
gateway = 192.168.1.1
nameserver_address = 192.168.1.254
hostname = spacevm.example.net
username = admin
password = abc123
newpassword = jnpr123!
guest_os = space
host_server = cso-host3
memory = 16384
vcpu = 4


[csp-installer-vm]
management_address = 192.168.1.17/24
hostname = installervm.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 49152
```


Copyright © 2018, Juniper Networks, Inc.

```
vcpu = 4
enable_data_interface = false


[csp-contrailanalytics-1]
management_address = 192.168.1.18/24
hostname = can1.example.net
username = root
password = passw0rd
local_user = installervm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 49152
vcpu = 16
enable_data_interface = false

[csp-central-lbvm1]
management_address = 192.168.1.19/24
hostname = centrallbvm1.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 16384
vcpu = 4
enable_data_interface = false


[csp-central-lbvm2]
management_address = 192.168.1.20/24
hostname = centrallbvm2.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 16384
vcpu = 4
enable_data_interface = false


[csp-central-lbvm3]
management_address = 192.168.1.20/24
hostname = centrallbvm3.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 16384
vcpu = 4
enable_data_interface = false

[csp-regional-lbvm1]
management_address = 192.168.1.21/24
```

Copyright © 2018, Juniper Networks, Inc.

```
hostname = regionallbvm1.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 16384
vcpu = 4
enable_data_interface = false

[csp-regional-lbvm2]
management_address = 192.168.1.22/24
hostname = regionallbvm2.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 16384
vcpu = 4
enable_data_interface = false

[csp-regional-lbvm3]
management_address = 192.168.1.22/24
hostname = regionallbvm3.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 16384
vcpu = 4
enable_data_interface = false

[csp-vrr-vm1]
management_address = 192.168.1.23/24
hostname = vrr1.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-host3
memory = 8192
vcpu = 4

[csp-vrr-vm2]
management_address = 192.168.1.24/24
hostname = vrr2.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-host3
memory = 8192
vcpu = 4

[csp-regional-sblb1]
management_address = 192.168.1.25/24
hostname = regional-sblb1.example.net
```


Copyright © 2018, Juniper Networks, Inc.

```
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 24576
vcpu = 4
enable_data_interface = true

[csp-regional-sblb2]
management_address = 192.168.1.26/24
hostname = regional-sblb2.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 24576
vcpu = 4
enable_data_interface = true
```

Sample Configuration File for Provisioning VMs in a Production Environment with HA

```
# This config file is used to provision KVM-based virtual machines using lib virt
 manager.

[TARGETS]
# Mention primary host (installer host) management_ip

installer_ip =

ntp_servers = ntp.juniper.net

# The physical server where the Virtual Machines should be provisioned
# There can be one or more physical servers on
# which virtual machines can be provisioned
physical = cso-central-host1, cso-central-host2, cso-central-host3,
cso-regional-host1, cso-regional-host2, cso-regional-host3


# The list of servers to be provisioned and mention the contrail analytics servers
 also in "server" list.
server = csp-central-infravm1, csp-central-infravm2, csp-central-infravm3,
csp-regional-infravm1, csp-regional-infravm2, csp-regional-infravm3,
csp-central-lbvm1, csp-central-lbvm2, csp-central-lbvm3, csp-regional-lbvm1,
csp-regional-lbvm2, csp-regional-lbvm3, csp-space-vm, csp-installer-vm,
csp-contrailanalytics-1, csp-contrailanalytics-2, csp-contrailanalytics-3,
csp-central-elkvm1, csp-central-elkvm2, csp-central-elkvm3, csp-regional-elkvm1,
 csp-regional-elkvm2, csp-regional-elkvm3, csp-central-msvm1, csp-central-msvm2,
 csp-central-msvm3, csp-regional-msvm1, csp-regional-msvm2, csp-regional-msvm3,
csp-vrr-vm1, csp-vrr-vm2, csp-regional-sblb1, csp-regional-sblb2,
csp-regional-sblb3

# Physical Server Details
[cso-central-host1]
management_address = 192.168.1.2/24
management_interface = virbr0
```

Copyright © 2018, Juniper Networks, Inc.

```
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-central-host1
username = root
password = passw0rd
data_interface =

[cso-central-host2]
management_address = 192.168.1.3/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-central-host2
username = root
password = passw0rd
data_interface =

[cso-central-host3]
management_address = 192.168.1.4/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-central-host3
username = root
password = passw0rd
data_interface =


[cso-regional-host1]
management_address = 192.168.1.5/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-regional-host1
username = root
password = passw0rd
data_interface =

[cso-regional-host2]
management_address = 192.168.1.6/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-regional-host2
username = root
password = passw0rd
data_interface =

[cso-regional-host3]
management_address = 192.168.1.7/24
management_interface = virbr0
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-regional-host3
```


Copyright © 2018, Juniper Networks, Inc.

```
username = root
password = passw0rd
data_interface =


[csp-contrailanalytics-1]
management_address = 192.168.1.17/24
management_interface =
hostname = can1.example.net
username = root
password = passw0rd
vm = false

[csp-contrailanalytics-2]
management_address = 192.168.1.18/24
management_interface =
hostname = can2.example.net
username = root
password = passw0rd
vm = false

[csp-contrailanalytics-3]
management_address = 192.168.1.19/24
management_interface =
hostname = can3.example.net
username = root
password = passw0rd
vm = false

# VM Details

[csp-central-infravm1]
management_address = 192.168.1.8/24
hostname = centralinfravm1.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-central-infravm2]
management_address = 192.168.1.9/24
hostname = centralinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-central-infravm3]
management_address = 192.168.1.10/24
hostname = centralinfravm3.example.net
```

Copyright © 2018, Juniper Networks, Inc.

```
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host3
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-infravm1]
management_address = 192.168.1.11/24
hostname = regionalinfravm1.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host1
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-infravm2]
management_address = 192.168.1.12/24
hostname = regionalinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host2
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-infravm3]
management_address = 192.168.1.13/24
hostname = regionalinfravm3.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host3
memory = 65536
vcpu = 16
enable_data_interface = false


[csp-space-vm]
management_address = 192.168.1.14/24
web_address = 192.168.1.15/24
gateway = 192.168.1.1
nameserver_address = 192.168.1.254
hostname = spacevm.example.net
username = admin
password = abc123
newpassword = jnpr123!
guest_os = space
```


Copyright © 2018, Juniper Networks, Inc.

```
host_server = cso-central-host2
memory = 32768
vcpu = 4

[csp-installer-vm]
management_address = 192.168.1.16/24
hostname = installervm.example.net
username = root
password = passw0rd
local_user = installervm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-lbvm1]
management_address = 192.168.1.20/24
hostname = centrallbvm1.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 32768
vcpu = 4
enable_data_interface = false


[csp-central-lbvm2]
management_address = 192.168.1.21/24
hostname = centrallbvm2.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-lbvm3]
management_address = 192.168.1.22/24
hostname = centrallbvm3.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host3
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-lbvm1]
management_address = 192.168.1.23/24
hostname = regionallbvm1.example.net
```

Copyright © 2018, Juniper Networks, Inc.

```
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host1
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-lbvm2]
management_address = 192.168.1.24/24
hostname = regionallbvm2.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host2
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-lbvm3]
management_address = 192.168.1.25/24
hostname = regionallbvm3.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host3
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-elkvm1]
management_address = 192.168.1.26/24
hostname = centralelkvm1.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-elkvm2]
management_address = 192.168.1.27/24
hostname = centralelkvm2.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 32768
vcpu = 4
```


Copyright © 2018, Juniper Networks, Inc.

```
enable_data_interface = false

[csp-central-elkvm3]
management_address = 192.168.1.28/24
hostname = centralelkvm3.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host3
memory = 32768
vcpu = 4
enable_data_interface = false


[csp-regional-elkvm1]
management_address = 192.168.1.29/24
hostname = regionalelkvm1.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host1
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-elkvm2]
management_address = 192.168.1.30/24
hostname = regionalelkvm2.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host2
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-regional-elkvm3]
management_address = 192.168.1.31/24
hostname = regionalelkvm3.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host3
memory = 32768
vcpu = 4
enable_data_interface = false

[csp-central-msvm1]
management_address = 192.168.1.32/24
hostname = centralmsvm1.example.net
username = root
password = passw0rd
```

Copyright © 2018, Juniper Networks, Inc.

```
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-central-msvm2]
management_address = 192.168.1.33/24
hostname = centralmsvm2.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-central-msvm3]
management_address = 192.168.1.34/24
hostname = centralmsvm3.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host3
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-msvm1]
management_address = 192.168.1.35/24
hostname = regionalmsvm1.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host1
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-msvm2]
management_address = 192.168.1.36/24
hostname = regionalmsvm2.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host2
memory = 65536
vcpu = 16
enable_data_interface = false
```


Copyright © 2018, Juniper Networks, Inc.

```
[csp-regional-msvm3]
management_address = 192.168.1.37/24
hostname = regionalmsvm3.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host3
memory = 65536
vcpu = 16
enable_data_interface = false

[csp-regional-sblb1]
management_address = 192.168.1.38/24
hostname = regional-sblb1.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host1
memory = 32768
vcpu = 4
enable_data_interface = true

[csp-regional-sblb2]
management_address = 192.168.1.39/24
hostname = regional-sblb2.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host2
memory = 32768
vcpu = 4
enable_data_interface = true

[csp-regional-sblb3]
management_address = 192.168.1.40/24
hostname = regional-sblb3.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host3
memory = 32768
vcpu = 4
enable_data_interface = true


[csp-vrr-vm1]
management_address = 192.168.1.41/24
hostname = vrr1.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-regional-host3
```

Copyright © 2018, Juniper Networks, Inc.

```
memory = 32768
vcpu = 4

[csp-vrr-vm2]
management_address = 192.168.1.42/24
hostname = vrr2.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-regional-host2
memory = 32768
vcpu = 4
```

Provisioning VMs with the Provisioning Tool for the KVM Hypervisor

If you use the KVM hypervisor on the CSO node or server, you can use the provisioning tool to:

- Create and configure the VMs for the CSO and Junos Space components.
- Install the operating system in the VMs:
  - Ubuntu in the CSO VMs
  - Junos Space Network Management Platform software in the Junos Space VM

To provision VMs with the provisioning tool:

1. Log in as root to the central CSO node or server.

2. Access the directory for the installer. For example, if the name of the installer directory is **csoVersion**:

   ```
   root@host:~/# cd /~/csoVersion/
   ```

3. Run the provisioning tool.

   ```
   root@host:~/cspVersion/# ./provision_vm.sh
   ```

   The provisioning begins.

4. During installation, observe detailed messages in the log files about the provisioning of the VMs.

   - **provision_vm.log**—Contains details about the provisioning process
   - **provision_vm_console.log**—Contains details about the VMs
   - **provision_vm_error.log**—Contains details about errors that occur during provisioning

   For example:

   ```
   root@host:~/cspVersion/# cd logs
   ```


Copyright © 2018, Juniper Networks, Inc.

```
root@host:/cspVersion/logs/# tail -f LOGNAME
```

Provisioning VMware ESXi VMs Using the Provisioning Tool

If you use the VMware ESXi (Version 6.0) VMs on the CSO node or server, you can use the provisioning tool—that is, provision_vm_ESXI.sh—to create and configure VMs for CSO.

NOTE: You cannot provision a Virtual Route Reflector (VRR) VM using the provisioning tool. You must provision the VRR VM manually.

Before you begin, ensure that the maximum supported file size for datastore in a VMware ESXi is greater than 512 MB. To view the maximum supported file size in datastore, you can establish an SSH session to the ESXi host and run the **vmfkstools -P *datastorePath*** command.

To provision VMware ESXi VMs using the provisioning tool:

1. Download the CSO Release 3.3 installer package from the Software Downloads page to the local drive.

2. Log in as root to the Ubuntu VM with the kernel version 4.4.0-31-generic, and has access to the internet. The VM must have the following specifications:
   - 8 GB RAM
   - 2 vCPUs

3. Copy the installer package from your local drive to the VM.

   ```
   root@host:~/# scp Contrail_Service_Orchestration_3.3.tar.gz root@VM  :/root
   ```

4. On the VM, extract the installer package.

   For example, if the name of the installer package is Contrail_Service_Orchestration_3.3.tar.gz,

   ```
   root@host:~/# tar –xvzf Contrail_Service_Orchestration_3.3.tar.gz
   ```

   The contents of the installer package are extracted in a directory with the same name as the installer package.

5. Navigate to the **confs** directory in the VM.

   For example:

   ```
   root@host:~/# cd Contrail_Service_Orchestration_3.3/confs
   root@host:~/Contrail_Service_Orchestration_3.3/confs#
   ```

Copyright © 2018, Juniper Networks, Inc.

6. Make a copy of the example configuration file, provision_vm_example_ESXI.conf, that is available in the **confs** directory and rename it provision_vm_ESXI.conf.

   For example:

   ```
   root@host:~/Contrail_Service_Orchestration_3.3/confs# cp
   /cso3.3/trial/nonha/provisionvm/provision_vm_example_ESXI.conf provision_vm.conf
   ```

7. Open the **provision_vm.conf** file with a text editor.

8. In the [TARGETS] section, specify the following values for the network on which CSO resides.

   - **installer_ip**—IP address of the management interface of the VM on which you are running the provisioning script.
   - **ntp_servers**—Comma-separated list of fully qualified domain names (FQDN) of Network Time Protocol (NTP) servers. For networks within firewalls, specify NTP servers specific to your network.

   You need not edit the following values:

   - **physical**—Comma-separated list of hostnames of the CSO nodes or servers are displayed.
   - **virtual**—Comma-separated list of names of the virtual machines (VMs) on the CSO servers are displayed.

9. Specify the following configuration values for each ESXI host on the CSO node or server.

   - **management_address**—IP address of the Ethernet management (primary) interface in classless Internet domain routing (CIDR) notation of the VM network. For example, 192.168.1.2/24.
   - **gateway**—Gateway IP address of the VM network
   - **dns_search**—Domain for DNS operations
   - **dns_servers**—Comma-separated list of DNS name servers, including DNS servers specific to your network
   - **hostname**—Hostname of the VMware ESXi host
   - **username**—Username for logging in to the VMware ESXi host
   - **password**—Password for logging in to the VMware ESXi host
   - **vmnetwork**—Labels for each virtual network adapter. This label is used to identify the physical network that is associated to a virtual network adapter.

     The **vmnetwork** data for each VM is available in the Summary tab of a VM in the vSphere Client. You must not specify **vmnetwork** data within double quotes.

   - **datastore**—Datastore value to save all VMs files.

     The **datastore** data for each VM is available in the Summary tab of a VM in the vSphere Client. You must not specify **datastore** data within double quotes.

Copyright © 2018, Juniper Networks, Inc.

10. Save the **provision_vm.conf** file.

11. Run the **provision_vm_ESXI.sh** script to create the VMs.

```
root@host:~/Contrail_Service_Orchestration_3.3/# ./provision_vm_ESXI.sh
```

12. Copy **provision_vm.conf** file in the installer VM.

    For example:

```
root@host:~/Contrail_Service_Orchestration_3.3/#scp confs/provision_vm.conf
root@installer_VM_IP:/root/Contrail_Service_Orchestration_3.3/confs
```

This action brings up VMware ESXi VMs with the configuration provided in the files.

Manually Provisioning VRR VMs on the Contrail Service Orchestration Node or Server

You cannot use the provision tool—provision_vm_ESXI.sh—to provision the Virtual Route Reflector (VRR) VM. You must manually provision the VRR VM.

To manually provision the VRR VM:

1. Download the VRR Release 15.1F6-S7 software package (.ova format) for VMware from the Virtual Route Reflector page, to a location accessible to the server.

2. Launch the VRR using vSphere or vCenter Client for your ESXi server and log in to the server with your credentials.

3. Set up an SSH session to the VRR VM.

4. Execute the following commands:

```
root@host:~/# configure
root@host:~/# delete groups global system services ssh root-login deny-password
root@host:~/# set system root-authentication plain-text-password
root@host:~/# set system services ssh
root@host:~/# set system services netconf ssh
root@host:~/# set routing-options rib inet.3 static route 0.0.0.0/0 discard
root@host:~/# commit
root@host:~/# exit
```

Verifying Connectivity of the VMs

From each VM, verify that you can ping the IP addresses and hostnames of all the other servers, nodes, and VMs in the CSO.

CAUTION: If the VMs cannot communicate with all the other hosts in the deployment, the installation can fail.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Installing and Configuring Contrail Service Orchestration on page 111

Setting up the Installation Package and Library Access

Copying the Installer Package to the Installer VM

After you have provisioned the VMs, move the installer package from the central server to the installer VM.

1. Copy the installer package file from the central CSO server to the installer VM.
2. Log in to the installer VM as root.
3. Expand the installer package.

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

   ```
   root@host:~/# tar –xvzf csoVersion.tar.gz
   ```

   The contents of the installer package are placed in a directory with the same name as the installer package. In this example, the name of the directory is **csoVersion**.
4. If you have created an installer VM using the provisioning tool, you must copy the **/csoVersion/confs/provision_vm.conf** file from the Ubuntu VM to the **/csoversion/confs/provision_vm.conf** directory of the installer VM.
5. Open the file provision_vm.conf with a text editor.
6. For **installer_ip** in the [TARGETS] section, specify the IP address of the installer VM.
7. Save the file.

Creating a Private Repository on an External Server

You use a private repository to download the libraries required for Contrail Service Orchestration. Use of a private repository for the libraries means that you do not require Internet access during the installation.

You can use a private repository either on the installer VM (the default choice) or on an external server.

- If you use the installer VM for the private repository, it is created when you install the solution, and you can skip this procedure.

Copyright © 2018, Juniper Networks, Inc.

- If you use an external server for the private repository, use the following procedure to create it.

To create the private repository on an external server:

1. Install the required Ubuntu release on the server that you use for the private repository.
2. Copy the installer package to the server.
3. Uncompress the installer package.

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

   ```
   root@host:~/# tar –xvzf csoVersion.tar.gz
   ```

   The contents of the installer package are placed in a directory with the same name as the installer package. In this example, the name of the directory is **csoVersion**.
4. Access the installer directory:

   For example:

   ```
   root@host:~/# cd csoVersion
   ```
5. Execute the **create_private_repo.sh** script to create the private repository.

   ```
   root@host:~/csoVersion#./create_private_repo.sh
   ```

   The script creates the private repository.
6. When you run the **setup_assist** script to create configuration files, specify that you use an external private repository. See "Installing and Configuring Contrail Service Orchestration" on page 111

Related Documentation

- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76
- Installing and Configuring Contrail Service Orchestration on page 111

Installing and Configuring Contrail Service Orchestration

You use the same installation process for both Contrail Service Orchestration (CSO) and Network Service Controller and for both KVM and ESXi environments.

Copyright © 2018, Juniper Networks, Inc.

- Loading Data on page 121
- Performing a Health Check of Infrastructure Components on page 122

Before You Begin

Before you begin:

- Provision the virtual machines (VMs) for the CSO node or server. (See "Provisioning VMs on Contrail Service Orchestration Nodes or Servers" on page 76).
- Copy the installer package to the installer VM and expand it. (See "Setting up the Installation Package and Library Access" on page 110)
- If you have created an installer VM using the provisioning tool, you must copy the **/Contrail_Service_Orchestration_3.3/confs/provision_vm.conf** file from the Ubuntu VM to the **/csoversion/confs/provision_vm.conf** directory of the installer VM.
- If you use an external server rather than the installer VM for the private repository that contains the libraries for the installation, create the repository on the server. (See "Setting up the Installation Package and Library Access" on page 110).

  The installation process uses a private repository so that you do not need Internet access during the installation.
- Determine the following information:
  - The type of deployment environment: Trial or production
  - Whether you use HA.
  - The names of each *regional* region if you use more than one region. The default specifies one *regional* region, called regional.
  - The IP address of the VM that hosts the installer.
  - The timezone for the servers in the deployment, based on the Ubuntu timezone guidelines.

    The default value for this setting is the current timezone of the installer host.
  - The fully qualified domain name (FQDN) of each Network Time Protocol (NTP) server that the solution uses. For networks within firewalls, use NTP servers specific to your network.

    For example: ntp.example.net
  - If you want to access Administration Portal with the single sign-on method, the name of the public domain in which the CSO servers reside. Alternatively if you want to access Administration Portal with local authentication, you need a dummy domain name.
  - For a distributed deployment, whether you use transport layer security (TLS) to encrypt data that passes between the CPE device and CSO.

    You should use TLS unless you have an explicit reason for not encrypting data between the CPE device and CSO.

Copyright © 2018, Juniper Networks, Inc.

- Whether you use the CSO Keystone or an external Keystone for authentication of CSO operations.
  - A CSO Keystone is installed with CSO and resides on the central CSO server.

    This default option is recommended for all deployments, and is required for a distributed deployment. Use of a CSO Keystone offers enhanced security because the Keystone is dedicated to CSO and is not shared with any other applications.
  - An external Keystone resides on a different server to the CSO server and is not installed with CSO.

    You specify the IP address and access details for the Keystone during the installation.
    - The Contrail OpenStack Keystone in the Contrail Cloud Platform for a centralized deployment is an example of an external Keystone.

      In this case, customers and Cloud CPE infrastructure components use the same Keystone token.
    - You can also use your own external Keystone that is not part of the CSO or Contrail OpenStack installation.
- If you use an external Keystone, the username and service token.
- The IP address of the Contrail controller node for a centralized deployment. For a centralized deployment, you specify this external server for Contrail Analytics.
- Whether you use a common password for all VMs or a different password for each VM, and the value of each password.
- The CIDR address of the subnet on which the CSO VMs reside.
- If you use NAT with your CSO installation, the public IP addresses used for NAT for the central and regional regions.
- The primary interface for all VMs.

  The default is eth0.
- The following information for each server and VM in the deployment:
  - Management IP address in CIDR notation

    For example: 192.0.2.1/24
  - FQDN of each host

    For example: central-infravm.example.net
  - Password for the root user

    If you use the same password for all the VMs, you can enter the password once. Otherwise, you must provide the password for each VM.
- For the microservices in the central and each regional region:

Copyright © 2018, Juniper Networks, Inc.

- The IP address of the Kubernetes overlay network address in Classless Interdomain Routing (CIDR) notation.

  The default value is 172.16.0.0/16. If this value is close to your network range, use a similar address with a /16 subnet.

- The range of the Kubernetes service overlay network addresses, in CIDR notation.

  The default value is 192.168.3.0/24.

- The IP address of the Kubernetes service API server, which is on the service overlay network.

  This IP address must be in the range you specify for the Kubernetes Service overlay network. The default value is 192.168.3.1.

- The IP address of the Kubernetes Cluster Domain Name System (DNS)

  This IP address must be in the range you specify for the Kubernetes Service overlay network. The default value is 192.168.3.1.

- The tunnel interface unit range that CSO uses for an SD-WAN implementation with an MX Series hub device.

  You must choose values that are different to those that you configured for the MX Series router. The possible range of values is 0–16385, and the default range is 4000–6000.

- The FQDN that the load balancer uses to access the installation.
  - For a non-HA deployment, the IP address and the FQDN of the VM that hosts the HAproxy.
  - For an HA deployment, the virtual IP address and the associated hostname that you configure for the HAproxy.
- The required number of copies of each microservice.
  - For a deployment without HA —1
  - For a trial deployment with HA—2
  - For a production deployment with HA—3

Creating the Configuration Files

You use an interactive script to create configuration files for the environment topology. The installer uses these configuration files to customize the topology when you deploy the solution.

To run the installation tool:

1. Log in as root to the host on which you deployed the installer.

2. Access the directory for the installer. For example, if the name of the installer directory is **csoVersion**:

Copyright © 2018, Juniper Networks, Inc.

```
root@host:~/# cd /~/csoVersion/
```

3. Run the setup tool:

   ```
   root@host:~/cspVersion/# ./setup_assist.sh
   ```

   The script starts, sets up the installer, and requests that you enter information about the installation.

4. Specify the management IP address of VM that hosts the installer file.

5. Specify the deployment environment:
   - trial—Trial environment
   - production—Production environment
6. Specify whether CSO is behind Network Address Translation (NAT).
   - y—CSO is behind NAT. After you deploy CSO, you must apply NAT rules. For information about NAT rules, see "Applying NAT Rules if CSO is Deployed Behind NAT" on page 137.
   - n—CSO is not behind NAT (default)
7. Accept the default timezone or specify the Ubuntu timezone for the servers in the topology.

8. Specify a comma-separated list of FQDN names of NTP servers.

   For example: ntp.example.net

9. Specify whether the deployment uses high availability (HA).
   - y—Deployment uses HA
   - n—Deployment does not use HA
10. Press enter if you use only one region or specify a comma-separated list or regions if you use multiple regions. You can configure a maximum of three regions. The default region is **regional**.

11. Specify the CSO certificate validity in days.

    The default value is 365 days.

12. Specify whether you need a separate regional southbound load balancer.

    The default value is yes.

13. For a distributed deployment, specify whether you use TLS to enable secure communication between the CPE device and CSO.

Copyright © 2018, Juniper Networks, Inc.

Accept the default unless you have an explicit reason for not using encryption for communications between the CPE device and CSO.

- n—Specifies that TLS is not used.
- y—Specifies use of TLS. This is the default setting.

14. Specify whether you want separate VMs for kubernetes master.

    The default value is no.

15. Specify the e-mail address of Admin User.

16. Specify a domain name to determine how you access Administration Portal, the main CSO GUI:

    - If you want to access Administration Portal with the single sign-on method, specify the name of the public domain in which the CSO servers reside.

      For example: *organization*.com, where *organization* is the name of your organization.

    - If you want to use local authentication for Administration portal, you specify a dummy name.

      For example: example.net

17. Specify whether you use an external Keystone to authenticate CSO operations, and if so, specify the OpenStack Keystone service token.

    - n—Specifies use of the CSO Keystone which is installed with and dedicated to CSO. This default option is recommended unless you have a specific requirement for an external Keystone.
    - y—Specifies use of an external OpenStack Keystone, such as a Keystone specific to your network. Select the IP address and access details for the Keystone.

18. Specify whether you use an external Contrail Analytics server:

    - y—Specifies use of Contrail Analytics in Contrail OpenStack for a centralized or combined deployment.

      You must provide the IP address of the Contrail controller node.

    - n—Specifies use of the Contrail Analytics VM for a distributed deployment.

19. Specify whether you use a common password for all CSO VMs, and if so, specify the password.

20. Specify the following information for the virtual route reflector (VRR) that you create:

    a. Specify whether VRR is behind NAT.

       - Specify the number of VRR instances.
         - For non HA deployments, you must create at least one VRR.
         - For HA deployments, it is recommended that you create VRRs in even numbers and there must be at least two VRRs. Each VRR must be in a different

Copyright © 2018, Juniper Networks, Inc.

redundancy group. If the primary VRR fails or connectivity is lost, the session remains active as the secondary VRR continues to receive and advertise LAN routes to a site, thereby providing redundancy.

- y—VRR is behind NAT. If you are deploying a VRR in a private network, the NAT instance translates all requests (BGP traffic) to a VRR from a public IP address to a private IP address.
- n—VRR is not behind NAT (default).

b. Specify whether you use a common password for all VRRs.

- y—Specify the common password for all VRRs.
- n—Specify the password for each VRR.

c. Specify the public IP address for each VRR that you create. For example, 192.0.20.118/24.

d. Specify the redundancy group for each VRR that you have created.

- For non HA deployments, specify the redundant group of the VRR as zero.
- For HA deployments, the VRRs must be distributed among the redundancy groups. There can be two groups—group 0 and group 1. For example, if you have two VRRs, specify the redundancy group for VRR1 as 0 and the VRR2 as 1.

21. Starting with the central region, specify the following information for each server in the deployment of each region.

The script prompts you for each set of information that you must enter.

- Management IP address with CIDR

  For example: 192.0.2.1/24

- Password for the root user (only required if you use different passwords for each VM)
- The IP address of the Kubernetes overlay network address, in CIDR notation, that the microservices use.

  The default value is 172.16.0.0/16. If this value is close to your network range, use a similar address with a /16 subnet.

- The range of the Kubernetes service overlay network addresses, in CIDR notation.

  The default value is 192.168.3.0/24. It is unlikely that there will be a conflict between this default and your network, so you can usually accept the default. If, however, there is a conflict with your network, use a similar address with a /24 subnet.

- The IP address of the Kubernetes service API server, which is on the service overlay network.

  This IP address must be in the range you specify for the Kubernetes Service overlay network. The default value is 192.168.3.1.

- The IP address of the Kubernetes Cluster DNS server.

Copyright © 2018, Juniper Networks, Inc.

This IP address must be in the range you specify for the Kubernetes Service overlay network. The default value is 192.168.3.1.

- Specify the range of tunnel interface units that CSO uses for an SD-WAN implementation with an MX Series hub device

  The default setting is 4000–6000. You specify values in the range 0–16385 that are different to those that you configured on the MX Series router.

- The IP address and FQDN of the host for the load balancer:
  - For non-HA deployments, the IP address and FQDN of the VM that hosts the HAproxy.
  - For HA deployments, the virtual IP address and associated FQDN that you configure for the HAproxy.
- The number of instances of microservices:
  - For deployments without HA, specify 1.
  - For a trial HA deployment with HA, specify 2.
  - For a production HA deployment with HA, specify 3.

The tool uses the input data to configure each region and indicates when the configuration stage is complete.

22. Configure settings for each region in the deployment:
    - Specify the IP address and prefix of the Kubernetes overlay network that the microservices use.
    - Specify the fully-qualified domain names of the host for the load balancer.
      - For a non-HA deployment, the IP address or FQDN of the VM that hosts the HAproxy.
      - For an HA deployment, the virtual IP address that you configure for the HAproxy.
    - Specify a unique virtual router identifier in the range 0–255 for the HA Proxy VM in each region.

      

      NOTE: **Use a different number for this setting in each region.**

    - Specify the number of instances of microservices:
      - For non-HA installations, specify 1.
      - For HA installations, specify 2.

    The tool uses the input data to configure each region and indicates when the configuration stage is complete.

23. Specify the subnet in CIDR notation on which the CSO VMs reside.

Copyright © 2018, Juniper Networks, Inc.

The script requires this input, but uses the value only for distributed deployments and not for centralized deployments.

24. Specify the range for tunnel interface unit.

25. Accept or specify the primary interface for all VMs.

    The default is eth0. Accept this value unless you have explicitly changed the primary interface on your host of VMs.

26. When all regions are configured, the tool starts displaying the deployment commands.

```
root@host:~/# DEPLOYMENT_ENV=central ./deploy_infra_services.sh
root@host:~/# DEPLOYMENT_ENV=regional ./deploy_infra_services.sh
root@host:~/# DEPLOYMENT_ENV=central ./deploy_micro_services.sh
root@host:~/# DEPLOYMENT_ENV=regional ./deploy_micro_services.sh
```

NOTE: The password for each infrastructure component and the Administration Portal password are displayed on the console after you complete answering the Setup Assistance questions. You must note the password that is displayed on the console as they are not saved in the system. To enhance the password security, the length and pattern for each password is different and the password is encrypted, and passwords in the log file are masked.

Deploying Infrastructure Services

To deploy infrastructure services:

1. Log in as root to the host on which you deployed the installer.

2. Deploy the central infrastructure services and wait at least ten minutes before you execute the next command.

```
root@host:~/# run "DEPLOYMENT_ENV=central ./deploy_infra_services.sh"
```

CAUTION: Wait at least ten minutes before executing the next command. Otherwise, the microservices may not be deployed correctly.

3. Deploy the regional infrastructure services and wait for the process to complete.

```
root@host:~/# run "DEPLOYMENT_ENV=regional ./deploy_infra_services.sh"
```

   If you have configured multiple regions, then you can deploy the infrastructure services on the regions in any order after deploying the central infrastructure.

Copyright © 2018, Juniper Networks, Inc.

NOTE: The **deploy_infra_services.sh** script performs a health check of infrastructure services. If you encounter an error you must rerun the **deploy_infra_services.sh** script again.

Deploying Microservices

To deploy the microservices:

1. Log in as root to the host on which you deployed the installer.

2. Deploy the central microservices and wait at least ten minutes before you execute the next command.

```
root@host:~/# -run "DEPLOYMENT_ENV=central ./deploy_micro_services.sh"
```

CAUTION: Wait at least ten minutes before executing the next command. Otherwise, the microservices may not be deployed correctly.

3. Deploy the regional microservices and wait for the process to complete:

```
root@host:~/# -run "DEPLOYMENT_ENV=regional ./deploy_micro_services.sh"
```

Checking the Status of the Microservices

To check the status of the microservices:

1. Log in as root into the VM or server that hosts the central microservices.

2. Run the following command.

```
root@host:~/# kubectl get pods | grep -v Running
```

If the result is an empty display, as shown below, the microservices are running and you can proceed to the next section.

```
root@host:~/# kubectl get pods | grep -v Running
NAME                                    READY   STATUS             RESTARTS  AGE
```

If the display contains an item with the status *CrashLoopBackOff* or *Terminating*, a microservice is not running.

3. Delete and restart the pod.

```
root@host:~/# kubectl get pods
```


Copyright © 2018, Juniper Networks, Inc.

```
NAME                                READY   STATUS             RESTARTS  AGE
csp.ams-3909406435-4yb0l            1/1     CrashLoopBackOff   0         8m
csp.nso-core-3445362165-s55x8       0/1     Running            0         8m
```

The first item in the display shows the microservice and the second item shows its pod.

```
root@host:~/# kubectl delete pods -l microservice=csp.nso-core
```

4. Wait a couple of minutes, then check the status of the microservice and its pod.

```
root@host:~/#kubectl get pods
```

```
NAME                                READY    STATUS      RESTARTS    AGE

csp.ams-4890899323-3dfd02           1/1      Running     0           1m
csp.nso-core-09009278633-fr234f     0/1      Running     0           1m
```

5. Repeat Steps 1 through 4 for the regional microservices.

Loading Data

After you check that the microservices are running, you must load data to import plug-ins and data design tools.

To load data:

1. Ensure that all the microservices are up and running on the central and each regional microservices host.

2. (Optional) Specify the value of the regional subnet in the file **/micro_services/data/inputs.yaml** on the installer VM. By default, the subnet address is the management address of the regional microservices host that you specify in the **topology.conf** file.

3. Access the home directory of the installer VM.

4. Execute the **./load_services_data.sh** command.

```
root@host:~/#./load_services_data.sh
```

NOTE: You must not execute load_services_data.sh more than once after a new deployment.

Copyright © 2018, Juniper Networks, Inc.

Performing a Health Check of Infrastructure Components

After you install or upgrade CSO, you can run the **components_health.sh** script to perform a health check of all infrastructure components. This script detects whether any infrastructure component has failed and displays the health status of the following infrastructure components:

- Cassandra
- Elasticsearch
- Etcd
- MariaDB
- RabbitMQ
- ZooKeeper
- Redis
- ArangoDb
- SimCluster
- ELK Logstash
- ELK Kibana
- Contrail Analytics
- Keystone
- Swift
- Kubernetes

To check the status of infrastructure components:

1. Login to the installer VM as root.

2. Navigate to the CSO directory in the installer VM.

   For example:

   ```
   root@host:~/# cd Contrail_Service_Orchestration_3.3
   root@host:~/Contrail_Service_Orchestration_3.3#
   ```

3. Run the **components_health.sh** script.

   To check the status of infrastructure components of the central environment, run the following command:

   ```
   root@host:~/Contrail_Service_Orchestration_3.3#./components_health.sh central
   ```

   To check health component of the regional environment, run the following command:

   ```
   root@host:~/Contrail_Service_Orchestration_3.3#./components_health.sh regional
   ```


Copyright © 2018, Juniper Networks, Inc.

To check health component of central and regional environments, run the following command:

```
root@host:~/Contrail_Service_Orchestration_3.3# ./components_health.sh
```

After a couple of minutes, the status of each infrastructure component for central and regional environments are displayed.

For example:

```
**************************************************************************

HEALTH CHECK FOR INFRASTRUCTURE COMPONENTS STARTED IN CENTRAL ENVIRONMENT

**************************************************************************

INFO     Health Check for Infrastructure Component Cassandra Started
INFO     The Infrastructure Component Cassandra is Healthy

INFO     Health Check for Infrastructure Component ElasticSearch Started
INFO     The Infrastructure Component ElasticSearch is Healthy

INFO     Health Check for Infrastructure Component Etcd Started
INFO     The Infrastructure Component Etcd is Healthy

INFO     Health Check for Infrastructure Component MariaDb Started
INFO     The Infrastructure Component MariaDb is Healthy

INFO     Health Check for Infrastructure Component RabbitMQ Started
INFO     The Infrastructure Component RabbitMQ is Healthy

INFO     Health Check for Infrastructure Component ZooKeeper Started
INFO     The Infrastructure Component ZooKeeper is Healthy

INFO     Health Check for Infrastructure Component Redis Started
INFO     The Infrastructure Component Redis is Healthy

INFO     Health Check for Infrastructure Component ArangoDb Started
INFO     The Infrastructure Component ArangoDb is Healthy

INFO     Health Check for Infrastructure Component Sim_Cluster Started
INFO     The Infrastructure Component Sim_Cluster is Healthy

INFO     Health Check for Infrastructure Component Elk_Logstash Started
INFO     The Infrastructure Component Elk_Logstash is Healthy

INFO     Health Check for Infrastructure Component Elk_Kibana Started
INFO     The Infrastructure Component Elk_Kibana is Healthy

INFO     Health Check for Infrastructure Component Keystone Started
INFO     The Infrastructure Component Keystone is Healthy

INFO     Health Check for Infrastructure Component Swift Started
INFO     The Infrastructure Component Swift is Healthy

INFO     Health Check for Infrastructure Component Kubernetes Started
INFO     The Infrastructure Component Kubernetes is Healthy
```

Copyright © 2018, Juniper Networks, Inc.

```
INFO     Health Check for Infrastructure Component Contrail_Analytics Started
INFO     The Infrastructure Component Contrail_Analytics is Healthy


Overall result:

        The following Infrastructure Components are Healthy:

               ['Cassandra', 'ElasticSearch', 'Etcd', 'MariaDb', 'RabbitMQ',
 'ZooKeeper', 'Redis', 'ArangoDb', 'Sim_Cluster', 'Elk_Logstash', 'Elk_Kibana',
 'Keystone', 'Swift', 'Kubernetes', 'Contrail_Analytics']
```

Related Documentation

- Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76
- Generating and Encrypting Passwords for Infrastructure Components on page 125
- Uploading the vSRX VNF Image for a Centralized Deployment on page 132
- Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 133
- Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 135
- Configuring Contrail OpenStack for a Centralized Deployment on page 125

Copyright © 2018, Juniper Networks, Inc.

Generating and Encrypting Passwords for Infrastructure Components

From Contrail Service Orchestration (CSO) Release 3.3 onwards, CSO uses an algorithm to automatically generate a dynamic password for the following infrastructure components:

- Cassandra
- Keystone
- MariaDB
- RabbitMQ
- Icinga
- Prometheus
- ArangoDB

The auto-generated passwords for each infrastructure component and the cspadmin password for Administration Portal are displayed on the console after you complete answering the Setup Assistance questions.

NOTE: **You must note the auto-generated password that is displayed on the console as they are not saved in the system.**

To enhance the password security, the length and pattern for each password is different and the password is encrypted. The passwords in the log file are masked.

Related Documentation

- Installing and Configuring Contrail Service Orchestration on page 111

Configuring Contrail OpenStack for a Centralized Deployment

After you have installed Contrail Service Orchestration (CSO) and uploaded virtualized network functions (VNFs) for a centralized deployment, you must complete the following tasks in Contrail OpenStack.

- Updating the VNF Image Properties on page 126
- Updating the Public Endpoints' IP Addresses on page 126
- Updating the OpenStack Heat Resources on page 127
- Specifying Attributes for Virtual Networks Created in Contrail on page 128
- Configuring the Contrail OpenStack Keystone as the CSO External Keystone on page 128
- Configuring Contrail OpenStack to Communicate with a CSO Keystone on page 131

Copyright © 2018, Juniper Networks, Inc.

Updating the VNF Image Properties

After you have uploaded the VNF images for your centralized deployment, you must update the image properties. To do so:

1. Obtain the identifiers for your VNF images.

   ```
   root@host:/# glance image-list
   ```

2. Execute the following command for each VNF image that you uploaded.

   glance image-update --property hw_cdrom_bus=ide --property hw_disk_bus=ide --property hw_vif_model=e1000 *vnf-image-id*

   Where:

   vnf-image-id—Identifier of the VNF image

   For example:

   ```
   root@host:/# glance image-update --property hw_cdrom_bus=ide --property hw_disk_bus=ide
   --property hw_vif_model=e1000 c79c1ade4f5eed8760fe
   ```

Updating the Public Endpoints' IP Addresses

You must update the deployment's public endpoints' IP addresses to match the management IP address of the Contrail controller node. This action enables Contrail to communicate with CSO. To do so:

1. Copy the endpoint_replace.py script from the CSO installer VM to the Contrail controller node.

   The endpoint_replace.py script is located at the **/root/Contrail_Service_Orchestration_3.3/scripts** directory.

2. Log in to the Contrail controller node as root.

3. Obtain the Keystone service token from the **/etc/contrail/keystone** file.

4. Execute the following command:

   ```
   root@host:/# python endpoint_replace.py --admin-token service-token --management-ip
   contrail-controller-ip-address
   ```

   Where:

   - *service-token*—Service token for the Contrail OpenStack Keystone
   - *contrail-controller-ip-address*—Management IP address of the Contrail controller node

   For example:

Copyright © 2018, Juniper Networks, Inc.

```
root@host:/# python endpoint_replace.py --admin-token 9390f3df14812451541f
--management-ip 192.0.2.1
```

Updating the OpenStack Heat Resources

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 as root.

2. To check whether the JSM Heat resource is available, execute the **heat resource-type-list | grep JSM** command.

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

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

   a. Use Secure Copy Protocol (SCP) to copy the **jsm_contrail_3.py** file 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.py file is located in the **/root/Contrail_Service_Orchestration_3.3/scripts** directory on the VM or server on which you installed CSO.

   b. Rename the file to **jsm.py** in 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.

Copyright © 2018, Juniper Networks, Inc.

Specifying Attributes for Virtual Networks Created in Contrail

A centralized deployment uses Contrail virtual networks for management and Internet traffic. You can create these virtual networks when you set up a centralized deployment in Administration Portal. Alternatively, you can create the networks in Contrail or use existing networks that you created in Contrail. For more information about this subject, see: https://www.juniper.net/documentation/en_US/contrail3.0/topics/task/configuration/creating-virtual-network-juniper-vnc-consolidate.html.

If you create the virtual networks in Administration Portal, CSO automatically sets up the required routing and sharing attributes for the networks. If, however, you create the virtual networks in Contrail, you must:

- Configure routing from the Contrail Service Orchestration (CSO) regional server to both virtual networks.
- Specify that the management virtual network is shared (public).

  This action ensures that the multiple tenants (customers) can access the network.

Configuring the Contrail OpenStack Keystone as the CSO External Keystone

When you install CSO, you can specify that the deployment should use the Contrail OpenStack Keystone as an external Keystone for authentication of CSO operations. If you do so, you must use this procedure to configure the Contrail OpenStack Keystone to authenticate CSO operations. To do so:

1. Log in to the Contrail controller node as root.

2. If you want to execute Keystone commands, set the source path, using the path that you configured during the installation.

   For example:

   ```
   root@host:~/# source /etc/contrail/keystonerc
   ```

3. Set the OpenStack source path.

   For example:

   ```
   root@host:~/# source /etc/contrail/openstackrc
   ```

4. Create a user called cspadmin.

   ```
   root@host:/# openstack user create --domain default --password-prompt cspadmin
   ```

5. Obtain the identifiers (IDs) of the following users:

   - admin
   - cspadmin
   - neutron

Copyright © 2018, Juniper Networks, Inc.

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

ID                      Name
0a3615846a4d689bedf8    admin
20a61f33a15453f21682    cspadmin
41a71e35a152a7c39e69    neutron
```

6. Obtain the ID of the default domain.

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

7. Create a project called default-project.

```
root@host:/# openstack project create --domain default --description "Default Project"
default-project
```

8. Assign the admin role to the admin and cspadmin users.

```
root@host:/# openstack role add admin --user admin --project default-project
root@host:/# openstack role add admin --user cspadmin --project default-project
```

9. Create the roles operator and tenant-operator.

```
root@host:~/# openstack role create operator
root@host:~/# openstack role create tenant-operator
```

10. Obtain the Keystone service token from the **/etc/contrail/keystone** file.

11. If the following groups do not already exist, create them:

    - admin
    - member
    - operator

    ```
    root@host:~/# curl –H “x-auth-token:service-token“ –H “content-type:application/json” –d
    ‘{“group”: {“name”: “group-name”, “domain_id”: “default”}}’ –XPOST
    http://contrail-controller-ip-address:5000/v3/groups
    ```

    where

    - *service-token*—Service token for the Contrail OpenStack Keystone
    - *group-name*—Name of the group
    - *domain_id*—ID of the domain
    - *contrail-controller-ip-address*—Management IP address of the Contrail controller node

    For example:

    ```
    root@host:~/# curl –H “x-auth-token:9390f3df14812451541f“ –H
    “content-type:application/json” –d ‘{“group”: {“name”: “operator”, “2738ef02df227c34ec49”:
    “default”}}’ –XPOST http://192.0.2.1:5000/v3/groups
    ```

Copyright © 2018, Juniper Networks, Inc.

```
root@host:~/# curl –H “x-auth-token:9390f3df14812451541f” –H
“content-type:application/json” –d ‘{“group”: {“name”: “admin”, “2738ef02df227c34ec49”:
“default”}}’ –XPOST http://192.0.2.1:5000/v3/groups
root@host:~/# curl –H “x-auth-token:9390f3df14812451541f” –H
“content-type:application/json” –d ‘{“group”: {“name”: “_member_”, “2738ef02df227c34ec49”:
“default”}}’ –XPOST http:// 192.0.2.1:5000/v3/groups
```

12. Obtain the IDs for the groups:

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

ID                       Name
7df60593f801df3cad04     _member_
5be423fdf76a5d4f8964     admin
3bc8235fd643ae814c3d     operator
```

13. Use the following command to add the admin and cspadmin users to the admin and _member_ groups.

```
root@host:/# curl –g –I –X PUT
http://contrail-controller-ip-address:5000/v3/groups/group-id/users/user-id -H “Accept:
application/json” –H “X-Auth-Token:service-token”
```

where

- *contrail-controller-ip-address*—Management IP address of the Contrail controller node
- *group-id*—ID of the group
- *user-id*—ID of the user
- *service-token*—Service token that you use to access Contrail OpenStack

For example:

```
root@host:/# curl –g –I –X PUT
http://192.0.2.1:5000/v3/groups/5be423fdf76a5d4f8964/users/0a3615846a4d689bedf8
-H “Accept: application/json” –H “X-Auth-Token:9390f3df14812451541f”
root@host:/# curl –g –I –X PUT
http://192.0.2.1:5000/v3/groups/5be423fdf76a5d4f8964/users/20a61f33a15453f21682 -H
“Accept: application/json” –H “X-Auth-Token:9390f3df14812451541f”
root@host:/# curl –g –I –X PUT
http://192.0.2.1:5000/v3/groups/7df60593f801df3cad04/users/0a3615846a4d689bedf8
-H “Accept: application/json” –H “X-Auth-Token:9390f3df14812451541f”
root@host:/# curl –g –I –X PUT
http://192.0.2.1:5000/v3/groups/7df60593f801df3cad04/users/20a61f33a15453f21682 -H
“Accept: application/json” –H “X-Auth-Token:9390f3df14812451541f”
```

14. Use the following command to assign the system_user property to the admin, cspadmin, and neutron users.

```
root@host:/# curl –X PATCH –H “X-Auth-Token:service-token”
http://contrail-controller-ip-address:35357/v3/users/user-id -d ‘{“user”: {“system_user”: 1 }}’
```

where

Copyright © 2018, Juniper Networks, Inc.

- *service-token*—Service token for the Contrail OpenStack Keystone
- *contrail-controller-ip-address*—Management IP address of the Contrail controller node
- *user-id*—ID of the user

For example:

```
root@host:/# curl –X PATCH –H “X-Auth-Token:9390f3df14812451541f”
http://192.0.2.1:35357/v3/users/0a3615846a4d689bedf8 -d ‘{“user”: {“system_user”: 1 }}’
root@host:/# curl –X PATCH –H “X-Auth-Token:9390f3df14812451541f”
http://192.0.2.1:35357/v3/users/20a61f33a15453f21682 -d ‘{“user”: {“system_user”: 1 }}’
root@host:/# curl –X PATCH –H “X-Auth-Token:9390f3df14812451541f”
http://192.0.2.1:35357/v3/users/00d3b0113ae21f270d11 -d ‘{“user”: {“system_user”: 1 }}’
```

Configuring Contrail OpenStack to Communicate with a CSO Keystone

If you use the CSO Keystone with a centralized deployment, you must configure Contrail OpenStack to communicate with that Keystone. To do so:

1. Log in to the Contrail controller node as root.
2. Create a project for the CSO Keystone.

   ```
   root@host:/# openstack project create --domain default --description "CSO Keystone
   project" cso-project1
   ```

3. Assign the admin role to user admin for the project that you created.

   ```
   root@host:~/# openstack role add admin --project cso-project1 –user admin
   ```

4. Create a user, and assign the user to the project that you created.

   For example:

   ```
   root@host:~/# openstack user create --project cso-project1 –password prompt user 1
   ```

5. Assign the admin role to the user that you created.

   ```
   root@host:/# openstack role add admin --project cso-project 1 –user user1
   ```

Related Documentation

- Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment on page 64
- Authentication and Authorization in the Cloud CPE and SD-WAN Solutions on page 27
- Installing and Configuring Contrail Service Orchestration on page 111
- Uploading the vSRX VNF Image for a Centralized Deployment on page 132
- Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 133
- Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 135

Copyright © 2018, Juniper Networks, Inc.

Uploading the vSRX VNF Image for a Centralized Deployment

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

To upload the vSRX VNF image for a centralized deployment:

1. Log in to the installer VM as root.

2. Set up an SSH session as root to the Contrail controller node.

3. Copy the **vSRX-img** file from the installer VM to any directory on the Contrail controller node.

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

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

4. Check whether you have an OpenStack flavor with the following specification on the Contrail controller node.

   - 2 vCPUs
   - 4 GB RAM
   - 40 GB hard disk storage

   For example:

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

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

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

   For example:

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

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

Copyright © 2018, Juniper Networks, Inc.

For example:

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

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

To verify that you can manually instantiate the vSRX VNF:

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

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

   If you need to investigate the image further, the default username for the vSRX-img package is root and the password is passw0rd.

Related Documentation

- VNFs Supported by the Cloud CPE Solution on page 52
- Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 133
- Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 135

Uploading the LxCIPtable VNF Image for a Centralized Deployment

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

To create an LxCIPtable Image:

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

   For example:

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

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

Copyright © 2018, Juniper Networks, Inc.

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

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

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

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

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

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

7. Customize the instance.

   a. Set the root password to the value **passw0rd**. For example:

   

   CAUTION: You must use the value passw0rd for the LxCIPtable VNF to operate correctly.

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

   b. In the file **/etc/ssh/sshd_config**, specify the following setting:

```
PermitRootLogin = yes
```

   c. Restart the service.

```
service ssh restart
```

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

Copyright © 2018, Juniper Networks, Inc.

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

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

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

e. Verify that IPtables is active.

service ufw status

8. Take a snapshot of the OpenStack Instance.

a. Close the instance.

sudo shutdown -h now

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

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

Related Documentation

- VNFs Supported by the Cloud CPE Solution on page 52
- Uploading the vSRX VNF Image for a Centralized Deployment on page 132
- Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 135

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

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

To create a Cisco CSR-1000V VNF image:

1. Log into the Contrail controller node.

2. Create a new flavor with 3 vCPUs in Contrail OpenStack.

For example:

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

Copyright © 2018, Juniper Networks, Inc.

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

   For example:

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

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

   For example:

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

5. From the OpenStack GUI, log in to the instance using the management IP address as the username and without a password.

6. Configure the following settings for the instance:

   ```
   vrf definition Mgmt-intf
    address-family ipv4
    exit-address-family
   enable password passw0rd
   ip vrf mgmt
   username root privilege 15 password 0 passw0rd
   ip ssh version 2
   interface GigabitEthernet1
     ip vrf forwarding mgmt
     ip address dhcp
     negotiation auto
   line vty 0
     exec-timeout 60 0
     privilege level 15
     password passw0rd
     login local
     transport input telnet ssh
   ```

7. Take a snapshot of the instance.

   a. Close the instance.

      For example:

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

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

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- VNFs Supported by the Cloud CPE Solution on page 52
- Uploading the vSRX VNF Image for a Centralized Deployment on page 132
- Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 133

Applying NAT Rules if CSO is Deployed Behind NAT

If you have deployed Contrail Service Orchestration (CSO) behind NAT, you must apply NAT rules after you run the **setup_assit.sh** script on central and regional hosts. The NAT rule set determines the direction of the traffic to be processed.

NOTE: If you do not apply NAT rules after you install or upgrade CSO, you cannot access Administration Portal, Kibana UI, and Rabbit MQ console.

To apply NAT rules:

1. Log in to the installer VM as root.

2. Apply NAT rules for central and regional NAT servers.
   - To quickly apply the NAT rules for central NAT servers:
     a. Copy the following commands and paste them into a text file.

```
iptables -t nat -A PREROUTING -p tcp -d
central-nat-server-public-ip-address--dport 443 -j DNAT --to-destination
northbound-virtual-private-ip-address:443
iptables -t nat -A PREROUTING -p tcp -d
central-nat-server-public-ip-address--dport 35357 -j DNAT --to-destination
northbound-virtual-private-ip-address:35357
iptables -t nat -A PREROUTING -p tcp -d
central-nat-server-public-ip-address--dport 5601 -j DNAT --to-destination
northbound-virtual-private-ip-address:5601
iptables -t nat -A PREROUTING -p tcp -d
central-nat-server-public-ip-address --dport 9200 -j DNAT --to-destination
northbound-virtual-private-ip-address:9200
iptables -t nat -A PREROUTING -p tcp -d
central-nat-server-public-ip-address --dport 1947 -j DNAT --to-destination
northbound-virtual-private-ip-address:1947
```

     b. You must specify IP addresses in the command to match your network configuration.

     c. Copy and paste the updated commands into the CLI.

   - To quickly apply the NAT rules for regional NAT servers:
     a. Copy the following commands and paste them into a text file.

```
iptables -t nat -A POSTROUTING -o virbr0 -p tcp --dport 5601 -d
northbound-virtual-private-ip-address -j SNAT --to-source
regional-management-interface-ip-address
```

Copyright © 2018, Juniper Networks, Inc.

```
iptables -t nat -A PREROUTING -p tcp -d
regional-nat-server-public-ip-address --dport 5601 -j DNAT --to-destination
northbound-virtual-private-ip-address:5601
iptables -t nat -A POSTROUTING -o virbr0 -p tcp --dport 7804 -d
northbound-virtual-private-ip-address -j SNAT --to-source
regional-management-interface-ip-address
iptables -t nat -A PREROUTING -p tcp -d
regional-nat-server-public-ip-address --dport 7804 -j DNAT --to-destination
northbound-virtual-private-ip-address:7804
iptables -t nat -A POSTROUTING -o virbr0 -p tcp --dport 3514 -d
southbound-virtual-private-ip-address -j SNAT --to-source
regional-management-interface-ip-address
iptables -t nat -A PREROUTING -p tcp -d
regional-nat-server-public-ip-address--dport 3514 -j DNAT --to-destination
southbound-virtual-private-ip-address:3514
iptables -t nat -A POSTROUTING -o virbr0 -p tcp --dport 514 -d
southbound-virtual-private-ip-address-j SNAT --to-source
regional-management-interface-ip-address
iptables -t nat -A PREROUTING -p tcp -d
regional-nat-server-public-ip-address --dport 514 -j DNAT --to-destination
southbound-virtual-private-ip-address:514
iptables -t nat -A POSTROUTING -o virbr0 -p tcp --dport 443 -d
southbound-virtual-private-ip-address -j SNAT --to-source
regional-management-interface-ip-address
iptables -t nat -A PREROUTING -p tcp -d
regional-nat-server-public-ip-address --dport 443 -j DNAT --to-destination
northbound-virtual-private-ip-address:443
iptables -t nat -A PREROUTING -d regional-nat-server-public-ip-address/32
-p tcp -m tcp --dport 2216 -j DNAT --to-destination
southbound-virtual-private-ip-address:2216
iptables -t nat -A POSTROUTING -d southbound-virtual-private-ip-address/32
-o virbr0 -p tcp -m tcp --dport 2216 -j SNAT --to-source
regional-management-interface-ip-address
```

b. You must specify IP addresses in the commands to match your network configuration.

c. Copy and paste the updated commands into the CLI.

The NAT rules are applied for central and regional NAT servers, and you can access Administration Portal, Kibana UI, and Rabbit MQ console.

Related Documentation

- Installing and Configuring Contrail Service Orchestration on page 111

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 6

Upgrading to Contrail Service Orchestration Release 3.3

Upgrading Contrail Service Orchestration Overview

If your installed version is Contrail Service Orchestration (CSO) Release 3.2.1, you can use a script to directly upgrade to CSO Release 3.3.

NOTE: **You can upgrade to CSO Release 3.3 only from CSO Release 3.2.1. If your installed version of CSO is not Release 3.2.1, then you must perform a fresh installation of CSO 3.3.**

You can roll back to CSO Release 3.2.1, if the upgrade is unsuccessful.

To upgrade to CSO Release 3.3, you must run the scripts that are available in the Contrail_Service_Orchestration_3.3.tar.gz file in the following order:

1. upgrade.sh—This script upgrades CSO software from Release 3.2.1 to Release 3.3. The upgrade.sh script, puts CSO in maintenance mode, takes a snapshot of all VMs so that you can roll back to the previous release if the upgrade fails (optional), upgrades all microservices and infrastructure components if required, performs health checks at various levels, validates if all VMs, infracomponents, and microservices are up and running, and puts the CSO in live mode.

   

   NOTE: **Before you upgrade ensure that all ongoing jobs are stopped; otherwise during the upgrade the ongoing jobs are stopped. During the upgrade, you experience a downtime as CSO goes into maintenance mode.**

2. revert.sh—Run this script only if the upgrade fails and if you have taken a snapshot of all VMs. This script reverts to the previously installed version.

Copyright © 2018, Juniper Networks, Inc.

Figure 11: High-level Overview of Upgrading to CSO Release 3.3

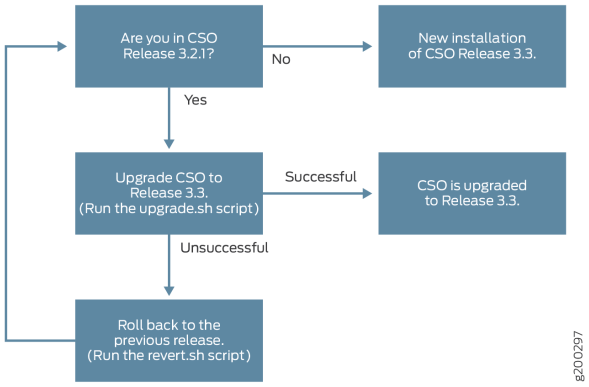

Upgrade to CSO Release 3.3 is independent of the deployment type (HA and non- HA), environment type (trial or production), infrastructure components and microservices used, and the hypervisor type (KVM or VMware ESXi).

To ensure a smooth upgrade, the scripts perform a number of health checks before and after the upgrade. Health checks are performed to determine the operational condition of all components, the host, and VMs. If there is an error during the health check, the upgrade process is paused. You can rerun the script to rectify the error that you encounter.

Following are the types of health checks that are performed:

- Component Health Checks—Checks the operational condition of the infrastructure components and microservices.
- System Health Checks—Checks the following parameters of VMs and the host machine.
  - Available space on the host machine and VMs
  - Operating System (OS) version of the host machine and VMs
  - Kernel version of the host machine and VMs
  - Disk space on the host machine and VMs

Limitations

Upgrade to CSO Release 3.3 has the following limitations:

- The upgrade is applicable only to CSO software and is not applicable to the existing devices and sites in CSO. After a successful upgrade, the existing sites and devices continue to have the same functionality of the previously installed version, that is, CSO Release 3.2.1.
- There are no changes to the device image or device configurations.

Impact of the CSO Upgrade

Table 23 on page 141 describes the impact of the CSO upgrade to Release 3.3.

Copyright © 2018, Juniper Networks, Inc.

Table 23: Impact of the CSO Upgrade

Feature	After the Upgrade
Security Management	• Release 3.3 security management-related features is supported on devices that are onboarded in Release 3.2.1.
SD-WAN	• For the Application Visibility feature, the trend data is reset after the upgrade. You can access the Release 3.2.1 trend data through the REST APIs. • The Application Quality of Experience (AppQoE) feature works only for the tenants that you create in Release 3.3. For more information on AppQoE, see *Application Quality of Experience (AppQoE) Overview* in the *Contrail Service Orchestration User Guide*. • Device Management functions work for Release 3.2.1 sites.
Cloud CPE	• All functionalities of centralized and distributed deployments continues to work on Release 3.2.1 sites or devices that are onboarded in Release 3.2.1. • Multi-region support for centralized deployment is not supported on Release 3.2.1 sites or devices that are onboarded in Release 3.2.1. • Device Management functions work for Release 3.2.1 sites. • High availability (HA) for VRRs are not supported for sites that are created in Release 3.2.1.

Related Documentation

• Upgrading to Contrail Service Orchestration Release 3.3 on page 141

Upgrading to Contrail Service Orchestration Release 3.3

From Contrail Service Orchestration (CSO) Release 3.3, you can directly upgrade the CSO software from Release 3.2.1 by running scripts.

This upgrade procedure is independent of the deployment type (trial and production), environment type (non-HA and HA), infrastructure components and microservices used, and the hypervisor type (KVM or VMware ESXi).

Before you begin the upgrade:

- Ensure that you are in Contrail Service Orchestration (CSO) Release 3.2.1.
- Ensure the installer Virtual Machine (VM) is up and running.
- If you are using VMware ESXi VMs, you must create the **provision_vm.conf** file in the **Contrail_Service_Orchestration_3.3/confs/** directory.

  For example, for a trial environment with HA, you can refer to the **provision_vm.conf** that is available in the **Contrail_Service_Orchestration_3.3/confs/trial/ha/provisionvm/**.
- If you have created a device template by cloning an existing device template, you must specify the new device template name in the **customer_migration.json** file, and place the file in the **/path/Contrail_Service_Orchestration_3.3/ micro_services/data/** location.

  For example, in the following **customer_migration.json** file, you must specify the new device template name in *new-device-template-name*.

Copyright © 2018, Juniper Networks, Inc.

```
{
      "device-profile": [{
        "cloned_profile": "new-device-template-name",   #Specify the new device
 template name.
        "from_profile": "SRX_Advanced_SDWAN_CPE_option_1"  #Name of the device
 template before cloning.
      },{
   "cloned_profile": "new-device-template-nameP",
        "from_profile": "SRX_Advanced_SDWAN_HUB_option_1"
   },{
   "cloned_profile": "new-device-template-name",
        "from_profile": "NFX_Advanced_SDWAN_CPE_option_1"
   }]
}
```

To upgrade to CSO Release 3.3:

1. Download the CSO Release 3.3 installer package from the Software Downloads page to the local drive.

2. Login to the installer VM as root.

3. Copy the installer package from your local folder to the installer VM.

   ```
   root@host:~/# scp Contrail_Service_Orchestration_3.3.tar.gz root@installer VM
   :/root
   ```

4. On the installer VM, extract the installer package.

   For example, if the name of the installer package is Contrail_Service_Orchestration_3.3.tar.gz,

   ```
   root@host:~/# tar –xvzf Contrail_Service_Orchestration_3.3.tar.gz
   ```

   The contents of the installer package are extracted in a directory with the same name as the installer package.

5. Navigate to the CSO Release 3.3 directory in the installer VM.

   ```
   root@host:~/# cd Contrail_Service_Orchestration_3.3
   root@host:~/Contrail_Service_Orchestration_3.3#
   ```

   (Optional) You can view the list of files in the Contrail_Service_Orchestration_3.3.

   ```
   root@host:~/Contrail_Service_Orchestration_3.3# ls
   ```

   The Contrail_Service_Orchestration_3.3.tar.tz file includes the following scripts:

   - upgrade.sh
   - revert.sh

Copyright © 2018, Juniper Networks, Inc.

6. Run the upgrade.sh script.

WARNING: Before you upgrade ensure that all ongoing jobs in Administration Portal and Customer Portal are stopped; otherwise during the upgrade the ongoing jobs are stopped. During the upgrade, you experience a downtime as CSO goes into maintenance mode.

This script upgrades CSO software from Release 3.2.1 to Release 3.3. The upgrade.sh script puts CSO in maintenance mode, takes a snapshot of running status of all VMs (optional), upgrades all microservices and infrastructure components if required, performs health checks at various levels, validates if all VMs, infrastructure components, and microservices are up and running, and puts the CSO in live mode.

If the environment type is production, the upgrade.sh script takes a snapshot of all VMs by default. For trial environment you are prompted to confirm whether you want to take a snapshot.

NOTE: The script does not take a snapshot of installer VM and Virtual Route Reflector (VRR) VM.

```
root@host:~/Contrail_Service_Orchestration_3.3# ./upgrade.sh

INFO     ============================================
INFO                Overall Upgrade Summary
INFO     ============================================
INFO      Configuration Upgrade : success
INFO      System Health Check : success
INFO      CSO Health-Check before Upgrade : success
INFO      CSO Maintenance Mode Enabled : success
INFO      VM Snapshot : success
INFO      Central Infra Upgrade : success
INFO      Regional Infra Upgrade : success
INFO      Microservices pre-deploy scripts execution : success
INFO      Central Microservices Upgrade : success
INFO      Regional Microservices upgrade : success
INFO      Microservices post-deploy scripts execution : success
INFO      CSO Health-Check after Upgrade : success
INFO      Enable CSO Services : success
INFO      Load Microservices Data : success
INFO      Overall Upgrade Status : success
INFO     ============================================
INFO     System got upgraded to 3.3 Successfully.
INFO     ============================================
```

Copyright © 2018, Juniper Networks, Inc.

NOTE: The password for each infrastructure component and the Administration Portal password are displayed on the console after the upgrade is successful. You must note the password that is displayed on the console as they are not saved in the system. To enhance the password security, the length and pattern for each password is different and the password is encrypted, and passwords in the log file are masked.

The time taken to complete the upgrade process depends on the hypervisor type and the environment type. If you are using KVM as the hypervisor, while taking a snapshot all VMs are shut down. If you are using VMware ESXi as the hypervisor, while taking a snapshot all VMs are up and running.

If an error occurs, you must fix the error and rerun the upgrade.sh script. When you rerun the upgrade.sh script, the script continues to execute from the previously failed step.

You can view the following log files that are available at **root/Contrail_Service_Orchestration_3.3/logs**:

- **upgrade_console.log**
- **upgrade_error.log**
- **upgrade.log**

7. (Optional) If you are unable to troubleshoot the error you can roll back to your previous release. Run the revert.sh script.

```
root@host:~/Contrail_Service_Orchestration_3.3# ./revert.sh
```

```
INFO     ================================================
INFO                Overall Revert Summary
INFO     ================================================
INFO     Pre-revert processes: success
INFO     Revert VM Snapshot: success
INFO     Revert Salt Master Configuration: success
INFO     Post-revert processes: success
INFO     CSO Health-Check after Revert: success
INFO     Start Kubernet Pods: success
INFO     Enable CSO Services: success
INFO     ================================================
INFO     CSO successfully reverted to the previously installed version
INFO     ================================================
```

After a successful upgrade, CSO is functional and you can login to Administrator Portal and Customer Portal.

NOTE: After you successfully upgrade from CSO Release 3.2.1 to Contrail Service Orchestration (CSO) Release 3.3, ensure that you download the application signatures before installing signatures on the device. This is a one-time operation after the upgrade.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Upgrading Contrail Service Orchestration Overview on page 139

Adding Virtual Route Reflectors (VRRs) After Upgrading to CSO Release 3.3

To support high availability (HA) for Virtual Route Reflectors (VRRs), you must add VRRs and create redundancy groups after you upgrade to Contrail Service Orchestrator (CSO) Release 3.3.

To add VRRs:

1. Login to the installer VM as root.

2. Navigate to the CSO Release 3.3 directory in the installer VM.

```
root@host:~/# cd Contrail_Service_Orchestration_3.3
root@host:~/Contrail_Service_Orchestration_3.3#
```

3. Run the **add_vrr.sh** script.

```
root@host:~/Contrail_Service_Orchestration_3.3# ./add_vrr.sh
```

The existing VRR details are displayed.

```
=========== Existing VRR Details ========

     host-name               |    redundancy-group

     vrr-192.204.243.28                0

=========================================
```

NOTE: By default, VRRs that are created during Release 3.2.1 belong to the redundancy group, *group 0*.

4. To add VRRs, you are prompted to answer the following questions:

   a. Specify whether VRR is behind NAT.

      - y—VRR is behind NAT. If you are deploying a VRR in a private network, the NAT instance translates all requests (BGP traffic) to a VRR from a public IP address to a private IP address.
      - n—VRR is not behind NAT (default).

   b. Specify whether you use a common password for all VRRs.

      - If you want to use a common password for all VRRs, enter **y** and specify the common password.

- If you want to use a different password for each VRR, enter **n** and specify the password for each VRR.

c. Specify the number of VRR instances.

- For non-HA deployments, you must create at least one VRR.
- For HA deployments, it is recommended that you create VRRs in even numbers and there must be at least two VRRs. Each VRR must be in a different redundancy group. If the primary VRR fails or connectivity is lost, the session remains active as the secondary VRR continues to receive and advertise LAN routes to a site, thereby providing redundancy.

d. For each VRR instance, specify the following:

- Specify the public IP address for each VRR that you create. For example, 192.110.20.118/24.
- Specify the redundancy group for each VRR that you have created.
  - For non-HA deployments, specify the redundancy group of the VRR as zero.
  - For HA deployments, the VRRs must be distributed among the redundancy groups. There can be two groups—group 0 and group 1. For example, if you have two VRRs, specify the redundancy group for VRR1 as 0 and the VRR2 as 1.
- Specify the user name for each VRR.
- Specify the password for each VRR.

  If you have chosen a common password for all VRRs, you are prompted to specify the common password only for the first VRR instance.

You can view the newly added VRRs through the APIs: routing-manager (GET: https://*IP Address of Administration Portal*/routing-manager/vrr-instance) or ems-central (GET: https://*IP Address of Administration Portal*/ems-central/device).

Each hub or spoke device establishes a BGP peering session with VRRs that you have created and assigned to different redundancy groups, thereby providing redundancy.

Related Documentation

- Upgrading Contrail Service Orchestration Overview on page 139

Troubleshooting Upgrade-Related Errors

This topic describes the possible errors that you might encounter while you are upgrading Contrail Service Orchestrator (CSO).

- Salt Synchronization Error on page 147
- Cache Clearance Error on page 148
- Kube-system Pod Error on page 148

Copyright © 2018, Juniper Networks, Inc.

- Kubernetes Node Error on page 149
- Snapshot Error on page 150

Salt Synchronization Error

Problem **Description:** While you are upgrading CSO to Release 3.3 or reverting to the previously-installed release, the upgrade or revert status is displayed as **Going to sync salt...** for a considerable time.

The Salt Master on the installer VM might be unable to reach all Salt Minions on the other VMs, and the salt timeout exception might occur.

Solution Based on the output of the **salt '*' test.ping** command, you must either restart the Salt Master or the Salt Minion.

To resolve the error:

1. Open another instance of installer VM.

2. Run the **salt '*' test.ping** command, to check if the Salt Master on the installer VM is able to reach other VMs.

   ```
   root@host:~/# salt '*' test.ping
   ```

   - If the following error occurs, you must restart the Salt Master.

     Salt request timed out. The master is not responding. If this error persists after verifying the master is up,worker_threads may need to be increased

     ```
     root@host:~/# service salt-master restart
     ```

   - If there are no errors, view the output.

     ```
     root@host:~/# salt '*' test.ping

     csp-regional-sblb.DB7RFF.regional:
     True
     csp-contrailanalytics-1.8V1O2D.central:
         True
     csp-central-msvm.8V1O2D.central:
         True
     csp-regional-k8mastervm.DB7RFF.regional:
         True
     csp-central-infravm.8V1O2D.central:
         False
     csp-regional-msvm.DB7RFF.regional:
         False
     csp-regional-infravm.DB7RFF.regional:
         True
     csp-central-k8mastervm.8V1O2D.central:
         True
     ```

     If the status of a VM is False, you must login to the VM, and restart the Salt Minion.

Copyright © 2018, Juniper Networks, Inc.

```
root@host:~/csp-central-infravm.8V1O2D.central# service salt-minion restart
```

3. Rerun the **salt '*' test.ping** command to verify if the status for all VMs is True.

Cache Clearance Error

Problem **Description:** While you are upgrading CSO to Release 3.3, the following error might occur:

Could not free cache on host server *ServerName*

Solution You must clear the cache on the host server.

To resolve the error:

1. Log in to the host server through SSH.

2. To clear the cache, run the following command:

```
root@host:~/Contrail_Service_Orchestration_3.3# free && sync && echo 3 >
/proc/sys/vm/drop_caches && free
```

The following output is displayed:

```
             total       used       free     shared    buffers     cached
Mem:     264036628  214945716   49090912      15092     198092   71878992
-/+ buffers/cache:  142868632  121167996
Swap:    390233084     473808  389759276
             total       used       free     shared    buffers     cached
Mem:     264036628  142165996  121870632      15092       3256      75792
-/+ buffers/cache:  142086948  121949680
Swap:    390233084     473808  389759276
```

The cache is cleared on the host server.

Kube-system Pod Error

Problem **Description:** While you are upgrading CSO to Release 3.3, the following error might occur:

One or more kube-system pods are not running

Solution Check the status of the *kube-system* pod, and restart *kube-proxy* if required.

To resolve the error:

1. Log in to the central or regional microservices VM through SSH.

2. To view the status of the *kube-system* pod, run the following command:

```
root@host:~/# kubectl get pods –namespace=kube-system
```

The following output is displayed:

Copyright © 2018, Juniper Networks, Inc.

```
NAME                                     READY     STATUS    RESTARTS   AGE
etcd-empty-dir-cleanup-10.213.20.182     1/1       Running   2          3d
kube-addon-manager-10.213.20.182         1/1       Running   2          3d
kube-apiserver-10.213.20.182             1/1       Running   2          3d
kube-controller-manager-10.213.20.182    1/1       Running   6          3d
kube-dns-v11-4cmhl                       4/4       Running    0          3h
kube-proxy-10.213.20.181                 0/1       Error    2          3d
kube-scheduler-10.213.20.182             1/1       Running   6          3d
```

Check the status of *kube-proxy*. You must restart *kube-proxy* if the status is *Error*, *Crashloopback*, or *MatchNodeSelector*.

3. To restart *kube-proxy*, run the following command.

```
root@host:~/# Kubectl apply –f /etc/kubernetes/manifests/kube-proxy.yaml
```

The *kube-system* pod-related error is resolved.

Kubernetes Node Error

Problem **Description:** While you are upgrading CSO to Release 3.3, the following error might occur:

One or more nodes down

Solution Check the status of *kube-master* or *kube-minion* and restart the nodes, if required.

To resolve the issue:

1. Log in to the central or regional microservices VM through SSH.

2. Run the following command to check the status of each node:

```
root@host:~/# kubectl get nodes
```

```
NAME            STATUS     AGE       VERSION
10.213.20.181   Not Ready    3d        v1.6.0
10.213.20.182   Ready        3d        v1.6.0
```

Identify the node that is in the *Not Ready* status. You must restart the node if the status is *Not Ready*.

3. To restart the node that is in the *Not Ready* status, log in to the node through SSH and run the following command:

```
root@host:~/# service kubelet restart
```

4. Rerun the following command to check the status of the node that you restarted.

Copyright © 2018, Juniper Networks, Inc.

```
root@host:~/# kubectl get nodes
```

The kubernetes node-related error is resolved.

Snapshot Error

Problem **Description:** The **upgrade.sh** script, sets CSO to maintenance mode, takes a snapshot of all VMs so that you can roll back to the previous release if the upgrade fails. While you are upgrading to CSO Release 3.3, the snapshot process might fail because of the following reasons:

- Unable to shutdown one or more VMs—You must manually shutdown the VM.
- Unable to take a snapshot for one ore more VMs—You must manually restart the VMs, start kubernetes pods, and set CSO to active mode.

Solution

- To manually shutdown the VMs:

  1. Log in to the CSO node or server as root.

  2. Execute the following command to view the list of VMs.

     ```
     root@host:~/# virsh list --all
     ```

     The VMs are listed as follows:

     ```
     Id    Name                           State
     ----------------------------------------------------
     10    vrr1                           running
     11    vrr2                           running
     40    canvm                          shut off
     41    centralinfravm                 shut off
     43    centralk8mastervm              running
     44    centralmsvm                    shut off
     45    installervm                    running
     46    regional-sblb                  shut off
     47    regionalinfravm                running
     48    regionalk8mastervm             shut off
     49    regionalmsvm                   shut off
     ```

     Identify the VMs that are in *running* state.

  3. Execute the following command to shutdown the VMs that are in *running* state:

     ```
     root@host:~/# virsh shutdown VMName
     ```

If you want to proceed with the upgrade process, you can rerun the **upgrade.sh** script.

Copyright © 2018, Juniper Networks, Inc.

- If you are unable to take the snapshot for one or more VMs, you must:

  1. Log in to the CSO node or server as root.

  2. Execute the following command to view the list of VMs.

     ```
     root@host:~/# virsh list --all
     ```

     The VMs are listed as follows:

     ```
     Id    Name                           State
     ----------------------------------------------------
     10    vrr1                           running
     11    vrr2                           running
     40    canvm                          running
     41    centralinfravm                 running
     43    centralk8mastervm              running
     44    centralmsvm                    running
     45    installervm                    running
     46    regional-sblb                  shut off
     47    regionalinfravm                running
     48    regionalk8mastervm             shut off
     49    regionalmsvm                   running
     ```

     Identify the VMs that are in *shut off* state.

  3. Execute the following command to restart the VMs that are in *shut off* state.

     ```
     root@host:~/# virsh start VMName
     ```

     You must restart the VMs in the following order:

     a. Infrastructure VM

     b. Load balancer VM

     c. Southbound Load balancer VM

     d. Contrail Analytics VM

     e. K8 Master VM

     f. Microservices VM

  4. On the installer VM, run the following commands to start the kubernetes pod:

     a. Execute the following command to check the status of *clear_cache_pods*.

        ```
        root@host:~/# cat upgrade/upgrade.conf | grep clear_cache_pods
        ```

     b. If the status of *clear_cache_pods* is *successful*, execute the following command to start the kubernetes pod.

        ```
        root@host:~/# ./python.sh vm_snapshot/scale_pods.py revert
        ```

  5. Log in to the central infrastructure VM through SSH , and run the following command to set CSO to active mode.

```
root@host:~/# etcdctl set /lb/maintenance false
```

If you want to proceed with the upgrade process, you can rerun the **upgrade.sh** script.

Related Documentation

- Upgrading Contrail Service Orchestration Overview on page 139
- Upgrading to Contrail Service Orchestration Release 3.3 on page 141

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 7

Installing Software Licenses for vSRX and SRX Series Devices

Overview of the License Tool

You can use the license tool to upload and install licenses for the following products:

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

Using this license tool is a quick and convenient way to upload and install licenses simultaneously. You can also use the API to install and upload licenses or to incorporate this functionality into your custom interface.

Contrail Service Orchestration uses the following workflow for uploading and installing licenses:

1. You run the license tool on the installer VM, which communicates with the central microservices host.
2. The central microservices host sends installation instructions to the regional microservices server that manages the CPE device or Contrail Controller node.
3. The regional microservices host executes installation instructions on the CPE device or the Contrail Controller node.

Related Documentation

- Installing Licenses with the License Tool on page 154

Copyright © 2018, Juniper Networks, Inc.

Installing Licenses with the License Tool

The license tool enables you to install and retrieve license information through a command line interface (CLI).

The license installation tool uses the following syntax:

./license_install_util.sh -i *license-id* -p *license-path* -t *tenant-name* - -sitefile *site-list-path* - -install_license - -get_license_info - -service firewall | utm | nat

Table 24 on page 154 describes the arguments and variables for the tool.

Table 24: Keywords and Variables for the License Tool

Arguments and Variables	Function	Requirement
-i *license-id*	Specifies the identifier of the license.	Mandatory for license installation
-p *license-path*	Specifies the path to the license file.	Mandatory for license installation
-t *tenant-name*	Specifies the name of the customer in Contrail Service Orchestration.	• Use for operations concerning all sites for a single customer. • Do not use for operations concerning multiple customers.
- -sitefile *site-list-path*	Specifies the path to a text file that contains a list of comma- or newline-separated sites in Contrail Service Orchestration.	Use for operations concerning multiple customers or a subset of sites for a single customer.
- -install_license	Installs licenses.	Requires either the **-t** or the **sitefile** option
- -get_license_info	Extracts licenses information	Requires either the **-t** or the **sitefile** option
- -service firewall | utm | nat	Specifies the network function for the license.	Mandatory if the site hosts multiple VNFs

- Accessing and Setting Up the License Tool on page 155
- Installing a License on All Sites for One Customer on page 155
- Installing a License for a Specific Service on All Sites for One Customer on page 156
- Installing a License on One or More Sites for Multiple Tenants on page 157
- Installing a License for a Specific Service on One or More Sites for Multiple Tenants on page 157
- Viewing License Information for One Customer's Sites on page 158
- Viewing License Information for One or More Sites on page 158

Copyright © 2018, Juniper Networks, Inc.

Accessing and Setting Up the License Tool

You run the license tool on the installer VM.

To access and set up the license tool:

1. Log in to the installer VM as root.

2. Access the directory that contains the installer. For example, if the name of the installer directory is **csoVersion**

   ```
   root@host:~/# cd csoVersion
   ```

3. Specify the following environment variables:

   - OS_AUTH_URL—URL of the OpenStack Keystone that authorizes Contrail Service Orchestration, including the IP address of the OpenStack Keystone host, port 35357 and the OpenStack version
   - OS_USERNAME—Username for Contrail Service Orchestration
   - OS_PASSWORD—Password for Contrail Service Orchestration
   - OS_TENANT_NAME—OpenStack tenant name, admin
   - TSSM_IP—IP address of the central microservices host
   - REGION_IP—IP address of the regional microservices host

   For example:

   ```
   root@host:~/#export OS_AUTH_URL=http://192.0.2.0:35357/v2.0
   root@host:~/#export OS_USERNAME=cspadmin
   root@host:~/#export OS_PASSWORD=passw0rd
   root@host:~/#export OS_TENANT_NAME=admin
   root@host:~/#export TSSM_IP=192.2.0.1
   root@host:~/#export REGION_IP=192.0.2.2
   ```

Installing a License on All Sites for One Customer

To install a license on all sites for one customer:

1. Run the tool with the following options (see Table 24 on page 154).

   ```
   ./license_install_util.sh -i license-id -p license-path -t tenant-name
   - -install_license
   ```

   For example:

   ```
   root@host:~/#./license_install_util.sh -i JUNOS000001 -p licenses/vsrx-utm-license.txt
   -t test-customer - -install_license
   Total Sites: 2
   Site count for successful license install: 2
   Site count for failed license install: 0
   ```

Copyright © 2018, Juniper Networks, Inc.

2. (Optional) Review the **license_install_results.log** for detailed results.

```
***License Install Status ***

Response:SUCCESS
Site: jd8-site-1
vSRX IP: 10.102.82.36
License Info: license": [
{
"license_id": "JUNOS000001",
"install_status": success
}
]
Response:SUCCESS
Site: jd8-site-2
vSRX IP: 10.102.82.2
License Info: license": [
{
"license_id": "JUNOS000001",
"install_status": success
}
]
```

3. If there is a problem with the license installation, review the **license_install.log** file for troubleshooting information.

Installing a License for a Specific Service on All Sites for One Customer

If you use more than one VNF at a site, you must specify the service when you install the license.

To install a license on all sites for a specific customer:

1. Run the tool with the following options (see Table 24 on page 154).

   ```
   ./license_install_util.sh -i license-id -p license-path -t tenant-name
   -s service-name - -install_license
   ```

   For example:

   ```
   root@host:~/#./license_install_util.sh -i JUNOS000001 -p licenses/vsrx-utm-license.txt
   -t test-customer -s firewall - -install_license
   ```

   ```
   Total Sites: 2
   Site count for successful license install: 2
    Site count for failed license install: 0
   ```

2. (Optional) Review the **license_install_results.log** for detailed results.

3. If there is a problem with the license installation, review the **license_install.log** file for debugging information.

Copyright © 2018, Juniper Networks, Inc.

Installing a License on One or More Sites for Multiple Tenants

To install a license on one or more sites:

1. Create a text file of site names, separated by commas or newline characters.

2. Run the tool with the following options (see Table 24 on page 154).

   ```
   ./license_install_util.sh -i license-id -p license-path
   - -sitefile site-file-name - -install_license
   ```

   For example:

   ```
   root@host:~/#./license_install_util.sh -i JUNOS000001 -p licenses/vsrx-utm-license.txt
   - - sitefile sites.txt
   Total Sites: 2
   Site count for successful license install: 2
   Site count for failed license install: 0
   ```

3. (Optional) Review the **license_install_results.log** for detailed results.

4. If there is a problem with the license installation, review the **license_install.log** file for debugging information.

Installing a License for a Specific Service on One or More Sites for Multiple Tenants

To install a license on one or more sites:

1. Create a text file of site names, separated by commas or newline characters.

2. Run the tool with the following options (see Table 24 on page 154).

   ```
   ./license_install_util.sh -i license-id -p license-path
   - -sitefile site-list-path - -install_license
   ```

   For example:

   ```
   root@host:~/#./license_install_util.sh -i JUNOS000001 -p licenses/vsrx-utm-license.txt
   - -sitefile sites.txt - -service utm
   Total Sites: 2
   Site count for successful license install: 2
   Site count for failed license install: 0
   ```

3. (Optional) Review the **license_install_results.log** for detailed results.

4. If there is a problem with the license installation, review the **license_install.log** file for debugging information.

Copyright © 2018, Juniper Networks, Inc.

Viewing License Information for One Customer's Sites

To view license information for one customer's sites:

1. Run the tool with the following options (see Table 24 on page 154).

   ```
   ./license_install_util.sh -i license-id -p license-path -t tenant-name
   - -get_license_info
   ```

   For example:

   ```
   root@host:~/#./license_install_util.sh -i JUNOS000001 -p licenses/vsrx-utm-license.txt
   -t test-tenant - -get_license_info
   Total Sites: 2
   Site count for successful license info retrieval: 2
   Site count for failed license info retrieval: 0
   Refer license_install_results.log for detailed results, license_install.log
   for debug logs.
   ```

2. (Optional) Review the **license_install_results.log** for detailed results.

   ```
   ***License Information ***
   Site: jd8-site-1
   vSRX IP: 10.102.82.36
   License Info: license": [
   {
   "license_id": "
   JUNOS000001",
   "install_status": success
   }
   ]
   Site: jd8-site-2
   vSRX IP: 10.102.82.2
   License Info: license": [
   {
   "license_id": "
   JUNOS000001",
   "install_status": success
   }
   ]
   ```

3. If there is a problem with operation, review the **license_install.log** file for debugging information.

Viewing License Information for One or More Sites

To view license information for one or more sites:

1. Create a text file of site names, separated by commas or newline characters.

2. Run the tool with the following options (see Table 24 on page 154).

   ```
   ./license_install_util.sh -i license-id -p license-path
   - -sitefile site-list-path - -get_license_info
   ```


Copyright © 2018, Juniper Networks, Inc.

For example:

```
root@host:~/#./license_install_util.sh -i JUNOS000001 -p licenses/vsrx-utm-license.txt
- -sitefile sites.txt - -get_license_info
Total Sites: 2
Site count for successful license info retrieval: 2
Site count for failed license info retrieval: 0
Refer license_install_results.log for detailed results, license_install.log
for debug logs.
```

3. (Optional) Review the **license_install_results.log** for detailed results.

```
***License Information ***
Site: jd8-site-1
vSRX IP: 10.102.82.36
License Info: license": [
{
"license_id": "
JUNOS000001",
"install_status": success
}
]
Site: jd8-site-2
vSRX IP: 10.102.82.2
License Info: license": [
{
"license_id": "
JUNOS000001",
"install_status": success
}
]
```

4. If there is a problem with operation, review the **license_install.log** file for debugging information.

Related Documentation

- Overview of the License Tool on page 153

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 8

Setting Up and Using Contrail Service Orchestration with the GUIs

Accessing the Contrail Services Orchestration GUIs

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

See Table 25 on page 161 for information about logging into the Contrail Service Orchestration GUIs.

Table 25: Access Details for the GUIs

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

Copyright © 2018, Juniper Networks, Inc.

Table 25: Access Details for the GUIs (continued)

GUI	URL	Login Credentials
Designer Tools—Log into Network Service Designer and click the menu in the top left of the page to access the other designer tools.	https://*central-IP-Address*:83 Where: *central-IP-Address*—IP address of the VM that hosts the microservices for the central POP For example: https://192.0.2.1:83	Specify the OpenStack Keystone username and password. The default username is **cspadmin**. Specify the autogenerated cspadmin password that is displayed on the console after the installation is complete. After the upgrade, you must specify the cspadmin password of the previously installed version.
Kibana This tool provides a visual representation of log files. You can use it to monitor: • Network services in a central or regional POP • Microservices in the deployment	http://*infra-vm-IP-Address* | *ha-proxy-IP-Address*:5601 Where: *infra-vm-IP-Address*—IP address of the VM that hosts the infrastructure services for a central or regional POP. Use this option to monitor network services. *ha-proxy-IP-Address*—IP address of high availability (HA) proxy. Use this option to monitor the microservices. • For a deployment without HA, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: http://192.0.2.2:5601	Login credentials are not needed.
Grafana and Prometheus These tools provide monitoring and troubleshooting for the infrastructure services in CSO. You use Prometheus to create queries for the infrastructure services and Grafana to view the results of the queries in a visual format.	• Prometheus—*ha-proxy-IP-Address*:30900 • Grafana—*ha-proxy-IP-Address*:3000 Where: *ha-proxy-IP-Address*—IP address of HA proxy • For a deployment without HA, use the IP address of the VM that hosts the microservices for the central POP. • For an HA deployment, use the virtual IP address that you provide for the HA proxy when you install CSO. For example: http://192.0.2.2:30900	For Grafana, specify the username and password. The default username is **admin** and the default password is **admin**. For Prometheus, the login credentials are not needed. After the upgrade, to log in to Administration Portal, you must specify the cspadmin password of the previously installed version.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Setting Up a Centralized Deployment on page 164
- Setting Up a Distributed Deployment on page 165
- Designing and Publishing Network Services on page 163
- Setting Up Customers' Networks on page 168
- Setting Up the Visual Presentation of Microservice Log Files on page 173
- Cloud CPE and SD-WAN Solutions Overview on page 17

Designing and Publishing Network Services

There are three tools that you use together to design and publish network services for centralized and distributed deployments in a hybrid WAN deployment:

- Firstly, you use Configuration Designer to create configuration templates for virtualized network functions (VNFs). The configuration templates specify the parameters that the customer can configure for a network service.
- Next, you use Resource Designer to create VNF packages. A VNF package specifies the network functions, function chains, performance, and a configuration template that you created in Configuration Designer.
- Finally, you use Network Service Designer to:
  - Design service chains for network services using the VNF packages that you created with Resource Designer.
  - Configure network services.
  - Publish network services to the network service catalog.

You use the same process to create network services for centralized and distributed deployments. You cannot, however, share network services between a centralized deployment and a distributed deployment that are managed by one Contrail Service Orchestration installation. In this case, you must create two identical services, one for the centralized deployment and one for the distributed deployment.

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

For detailed information about using the Designer Tools, see the *Contrail Service Orchestration User Guide*.

Related Documentation

- Accessing the Contrail Services Orchestration GUIs on page 161
- Setting Up a Centralized Deployment on page 164
- Setting Up a Distributed Deployment on page 165
- Setting Up Customers' Networks on page 168
- Cloud CPE and SD-WAN Solutions Overview on page 17

Copyright © 2018, Juniper Networks, Inc.

Setting Up a Centralized Deployment

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

- Configure network devices and servers for the deployment. See the following topics:
  - Cabling the Hardware for the Centralized Deployment on page 55
  - Configuring the EX Series Ethernet Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 58
  - Configuring the QFX Series Switch for the Contrail Cloud Implementation in a Centralized Deployment on page 59
  - Configuring the MX Series Router in the Contrail Cloud Implementation for a Centralized Deployment on page 61
  - Configuring the Physical Servers and Nodes for the Contrail Cloud Implementation in a Centralized Deployment on page 64
- Install Contrail Service Orchestration. See the following topics:
  - Removing a Previous Deployment on page 75
  - Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 76
  - Setting up the Installation Package and Library Access on page 110
  - Installing and Configuring Contrail Service Orchestration on page 111
  - Configuring Contrail OpenStack for a Centralized Deployment on page 125
- Upload VNF images. See the following topics:
  - Uploading the vSRX VNF Image for a Centralized Deployment on page 132
  - Uploading the LxCIPtable VNF Image for a Centralized Deployment on page 133
  - Uploading the Cisco CSR-1000V VNF Image for a Centralized Deployment on page 135
- Install VNF licenses.

  You can use the license tool to install vSRX licenses. See "Installing Licenses with the License Tool" on page 154.
- Publish network services with Network Service Designer.

To set up a centralized deployment.

1. Log in to Administration Portal as a service provider operator.
2. Create the POPs and associated resources.
   - You must create a (Virtualized Infrastructure Manager) VIM for each POP.
   - You can add an MX Series router as a physical network element (PNE) to provide a Layer 3 routing service to customer sites through use of virtual routing and forwarding (VRF) instances.

Copyright © 2018, Juniper Networks, Inc.

- You add the Junos Space element management system (EMS) if you use a VNF that requires this EMS.

3. Add or import customers (tenants) in Administration Portal.
4. Access Contrail and add the following rule to the default security group in the Contrail project.

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

5. Allocate network services to each customer.
6. Upload licenses for other VNFs.
7. Access the view for a specific customer.
8. Create cloud sites for the customer.
   a. Create a regional service edge site for each branch site in the customer's network.
   b. Create a local service edge site if customers access the Internet through the corporate VPN
9. If you configured a PNE, then associate the PNE with the site and configure a VRF for each customer site.

For detailed information about using Administration Portal, see the *Contrail Service Orchestration User Guide*.

Related Documentation

- Accessing the Contrail Services Orchestration GUIs on page 161
- Designing and Publishing Network Services on page 163
- Setting Up Customers' Networks on page 168
- Installing Licenses with the License Tool on page 154

Setting Up a Distributed Deployment

NOTE: **You must send an activation code to the customer for each NFX250 device. The customer's administrative user must provide this code during the NFX250 installation and configuration process. The Juniper Networks Redirect Service uses this code to authenticate the device.**

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

- Publish network services with Network Service Designer.
- Add or import customers (tenants) in Administration Portal.
- Allocate networks services to each customer.

Copyright © 2018, Juniper Networks, Inc.

After you have installed Contrail Service Orchestration and published network services with Network Service Designer, you use Administration Portal to set up the distributed deployment. The following workflow describes the process:

1. Log in to Administration Portal.
2. Access the tenant view for the first customer.
3. Add an on-premise spoke site for each site in the customer's network.

   NOTE: **Alternatively customers can add the spoke sites themselves.**

4. Repeat Step 3 for each customer in the network.
5. Access the All Tenants view for the customers.
6. Add data for the POPs and provider edge (PE) router.
7. Upload images for devices used in the deployment, such as the vSRX gateway and the NFX250 device, to the central activation server.
8. Configure activation data for CPE devices.
9. Upload VNF images.
10. Upload and install licenses:
    1. Upload licenses for vSRX and SRX devices and VNFs with the installer tool (see "Installing Licenses with the License Tool" on page 154).
    2. Upload licenses for other VNFS with Administration Portal.
    3. Manually install licenses for other VNFs.
11. Allocate network services to customers.
12. Activate CPE devices at customer sites.

    NOTE: **Alternatively customers can add the spoke sites themselves.**

    When an administrator installs and configures the NFX250 devices at a customer site, the device automatically interacts with the Redirect Service. The Redirect Service authenticates the device and sends information about its assigned regional server. The device then obtains a boot image and configuration image from the regional server and uses the images to become operational.

    Customers activate SRX Series Services Gateways and vSRX instances acting as CPE devices through Customer Portal.

For detailed information about using Administration Portal, see the *Contrail Service Orchestration User Guide.*

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Accessing the Contrail Services Orchestration GUIs on page 161
- Installing and Setting Up CPE Devices on page 72
- Designing and Publishing Network Services on page 163
- Setting Up Customers' Networks on page 168
- Installing Licenses with the License Tool on page 154

Setting Up an SD-WAN Deployment

To set up an SD-WAN implementation:

1. Access Administration Portal by logging into Contrail Service Orchestration (CSO) with the MSP Administrator login.

2. Create a point of presence (POP) for the SD-WAN deployment.

   

   BEST PRACTICE: **Create different POPs for Hybrid WAN and SD-WAN deployments so that it's clear which physical device (in this case the hub device) to select when you configure the spoke sites.**

3. Access the POP that contains the hub device for the SD-WAN deployment.

4. Add the hub device as a router to the POP.
   - You must supply the serial number of the hub device.
   - Select the SRX_Advanced_SDWAN device template.

   Multiple tenants can share the hub. You typically use one hub for each POP.

   The device should have the status Provisioned and an Activate Device link in the Management Status column on the POPs page in Administration Portal.

5. Activate the hub device in one of the following ways:
   - Use the remote activation utility on the SRX Series device.
   - Manually activate the device.
     a. Copy the Stage 1 configuration from the Routers page in Administration Portal to the SRX Series device console.
     b. Click Activate next to the hub device in the Routers page of Administration Portal.

6. Access the tenant view for the customer.

Copyright © 2018, Juniper Networks, Inc.

7. Add a cloud site to specify which hub site the tenant uses.
8. Create an on-premise spoke site for the customer and specify the LAN segments that connect to the CPE device.
9. Configure network connectivity and device images for the site.
10. Activate the device at the site.
11. Access the All Tenants view.
12. Install Application Signatures on the CPE device.
13. Access the tenant view for the customer.
14. Create and deploy the SD-WAN policy.
15. Create one or more SLA profiles for the customer.
16. Monitor application visibility and SD-WAN events.

    If an SLA violation occurs, CSO automatically switches the traffic from one WAN link to another on the CPE device. You can track these occurrences and view associated alarms in the Monitor Pages in both the All Tenants and specific tenant views.

Related Documentation

-

Setting Up Customers' Networks

After you have set up the network for a customer with Administration Portal, that customer can view, configure, and manage their network through Customer Portal. Customer Portal is actually customer-specific view of Administration Portal. Customers have their own login credentials, which provide role-based access control to the information for their networks. Customers see only their own networks, and cannot view other customers' networks. You can also view and manage each customer's network from Administration Portal, by accessing the view for a specific customer.

With Customer Portal, customers can:

- Add, activate and deactivate sites in the network.

BEST PRACTICE: **Service providers often add sites for customers. Customers typically activate and deactivate sites in their networks.**

Copyright © 2018, Juniper Networks, Inc.

- Configure CPE devices.
- Deploy and manage available network services for a hybrid WAN deployment.
  - Add and configure network services.
  - Disable and remove network services.
  - Monitor network services.

For detailed information about using Customer Portal, see the *Contrail Service Orchestration User Guide.*

Related Documentation

- Accessing the Contrail Services Orchestration GUIs on page 161
- Designing and Publishing Network Services on page 163
- Cloud CPE and SD-WAN Solutions Overview on page 17

Copyright © 2018, Juniper Networks, Inc.

Copyright © 2018, Juniper Networks, Inc.

CHAPTER 9

Monitoring and Troubleshooting

Monitoring and Troubleshooting Overview

You use open-source applications for monitoring and troubleshooting infrastructure services and microservices in Contrail Service Orchestration. These applications offer a visual representation of the metrics in Contrail Service Orchestration with extensive capabilities for analyzing data and monitoring alerts.

Monitoring Infrastructure Services

You use a combination of Prometheus and Grafana to monitor infrastructure services in Contrail Service Orchestration.

- Prometheus is a toolkit for monitoring systems and defining alerts.
- Grafana enables metric analysis and visualization.

You create queries in Prometheus to develop dashboards for infrastructures services, and visualize the dashboards in Grafana. Predefined dashboards for the following applications are included with Contrail Service Orchestration:

- Cassandra
- Kubernetes
- RabbitMQ
- Host metrics
  - Node and server metrics
  - VM metrics

Refer to the documentation for Prometheus and Grafana for information about using these products.

Copyright © 2018, Juniper Networks, Inc.

Monitoring Microservices

Service and Infrastructure Monitor (SIM) provides a continuous and comprehensive monitoring of Contrail Service Orchestration. The application provides both a visual display of the state of the deployment and the ability to view detailed event messages.

Service and Infrastructure Monitor tracks the status of:

- Network services
- Virtualized network functions
- Microservices
- Virtual machines
- Physical servers

For detailed information about using Service and Infrastructure Monitor, see the *Contrail Service Orchestration User Guide*.

You can also use Kibana to view log files and analyze log files in a visual format. See "Setting Up the Visual Presentation of Microservice Log Files" on page 173

Related Documentation

- Viewing and Creating Dashboards for Infrastructure Services on page 172
- Setting Up the Visual Presentation of Microservice Log Files on page 173
- Viewing Information About Microservices on page 174
- Cloud CPE and SD-WAN Solutions Overview on page 17

Viewing and Creating Dashboards for Infrastructure Services

To access and create dashboards for monitoring infrastructure services:

1. Access Grafana at the following URL:

   http://*ha-proxy-IP-Address*:3000

   Where:

   ha-proxy-IP-Address—IP address of high availability (HA) proxy for the infrastructure VMs

2. Select a predefined dashboard.

   The dashboard appears, displaying metrics for the infrastructure service.

3. Access Prometheus at the following URL to create additional dashboards:

   http://*ha-proxy-IP-Address*:30900

Copyright © 2018, Juniper Networks, Inc.

Refer to the documentation for Prometheus and Grafana for more information about using these products. You can also refer to the documentation for the different infrastructure services to determine what type of information to include in your custom dashboards.

Related Documentation

- Monitoring and Troubleshooting Overview on page 171
- Setting Up the Visual Presentation of Microservice Log Files on page 173
- Viewing Information About Microservices on page 174

Setting Up the Visual Presentation of Microservice Log Files

Contrail Service Orchestration includes Kibana and Logstash to enable viewing of logged data for microservices in a visual format.

To set up logging in Kibana:

1. Access Kibana using the URL for the server that you require (see "Accessing the Contrail Services Orchestration GUIs" on page 161).
2. Select **Settings** > **Indices**.
3. Click **Create**.

   This action creates the **csplogs** index file.
4. Log in as root to the installer host and access the installer directory.
5. Copy the **deploy_manager/export.json** file to a location from which you can import it to the Kibana GUI.

   

   NOTE: **Do not change the format of the JSON file. The file must have the correct format to enable visualization of the logs.**
6. In the Kibana GUI, select **Settings** > **Objects**.
7. Click **Import**.
8. Navigate to the location of the **export.json** file that you made available in Step 5.
9. Click **Open**.
10. Confirm overwriting of any existing data.

Copyright © 2018, Juniper Networks, Inc.

11. Refresh the Kibana page.

12. Access the dashboard to view the logs in a visual format.

    Logs appear after an end user activates a network service.

Refer to the Kibana documentation for information about viewing files in a visual format.

Related Documentation

- Accessing the Contrail Services Orchestration GUIs on page 161
- Monitoring and Troubleshooting Overview on page 171
- Viewing and Creating Dashboards for Infrastructure Services on page 172
- Viewing Information About Microservices on page 174

Viewing Information About Microservices

When you log into Kibana, you see the Discover page, which displays a chart of the number of logs for a specific time period and a list of events for the deployment. You can filter this data to view subsets of logs and add fields to the table to find the specific information that you need. You can also change the time period for which you view events.

- Filtering Data in Kibana on page 174
- Troubleshooting Microservices on page 174
- Analyzing Performance on page 175

Filtering Data in Kibana

To filter data in Kibana:

1. Specify a high-level query in the search field to view a subset of the logs.

   You can use keywords from the list of fields in the navigation bar, and specific values for parameters that you configure in Contrail Service Orchestration (CSO), such as a specific customer name or a specific network service.

   For example, specify the following query to view logs concerning requests made for the customer test-customer.

   exists: request_id AND test-customer

2. Select one or more fields from the left navigation bar.

   For example, select request to show details about the request made for this customer.

Troubleshooting Microservices

You can use the troubleshooting dashboard to investigate issues for the microservices.

Copyright © 2018, Juniper Networks, Inc.

To use the troubleshooting dashboard:

1. From the Kibana GUI, select **Dashboard** > **Troubleshooting**.

   The troubleshooting dashboard appears, displaying the following predefined monitoring applications:

   - Log Level Vs Count

     This widget shows the number of logs for each alert level.

   - Status Code Vs Count

     This widget shows the number of logs for each HTTP status code.

   - Service App Name Vs Status Code

     This widget shows a visual representation of the number of logs for each microservice analyzed by HTTP status code.

2. Click on an option, such as an alert level, in a widget to filter the data and drill down to a specific issue.

Analyzing Performance

You can use the troubleshooting dashboard to investigate issues for the microservices.

To use the troubleshooting dashboard:

1. From the Kibana GUI, select **Dashboard** > **Performance Analysis**.

   The performance Analysis dashboard appears, displaying the following predefined monitoring applications:

   - API Vs Min/Average/Max Elapsed time

     This widget shows how long an API associated with a microservice has been in use. You can view minimum, maximum, or average durations.

   - Request ID Vs Timestamp

     This widget shows when an API was called.

   - API Vs Count

     This widget shows the number of times an API has been called.

   - Application Vs API

     This widget shows the level of microservice use analyzed by the type of API call.

   - Request ID Vs Application Vs API

     This widget provides an analysis of requests by API or microservice.

2. Click on an option, such as a request identifier, in a widget to filter the data and drill down to a specific issue.

Copyright © 2018, Juniper Networks, Inc.

Related Documentation

- Monitoring and Troubleshooting Overview on page 171
- Viewing and Creating Dashboards for Infrastructure Services on page 172
- Setting Up the Visual Presentation of Microservice Log Files on page 173

Managing the Microservice Containers

After you deploy the microservices, you can manage the containers with the **deploy_micro_services.sh** script.

- Deleting and Restarting New Pods on page 176
- Clearing the Databases on page 176
- Clearing the Kubernetes Cluster on page 176

Deleting and Restarting New Pods

To restart all pods:

1. Log in to the installer VM as root.

2. Execute the following command to delete and recreate the containers.

```
root@host:~/# run "DEPLOYMENT_ENV=central ./deploy_micro_services.sh --restart_containers"
root@host:~/# run "DEPLOYMENT_ENV=regional ./deploy_micro_services.sh --restart_containers"
```

Clearing the Databases

To clear the Kubernetes databases:

1. Log in to the installer VM as root.

2. Execute the following command to clear the contents of the databases:

```
root@host:~/# run "DEPLOYMENT_ENV=central ./deploy_micro_services.sh --reset_databases"
root@host:~/# run "DEPLOYMENT_ENV=regional ./deploy_micro_services.sh --reset_databases"
```

Clearing the Kubernetes Cluster

To clear the entire Kubernetes cluster:

1. Log in to the installer VM as root.

2. Execute the following command to reset

```
root@host:~/# run "DEPLOYMENT_ENV=central ./deploy_micro_services.sh --reset_cluster"
```


Copyright © 2018, Juniper Networks, Inc.

```
root@host:~/# run "DEPLOYMENT_ENV=regional ./deploy_micro_services.sh --reset_cluster"
```

Related Documentation

- Installing and Configuring Contrail Service Orchestration on page 111

Copyright © 2018, Juniper Networks, Inc.