



Contrail Service Orchestration

Contrail Service Orchestration Installation and Upgrade Guide

Release

4.1.0



Modified: 2019-02-28

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

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

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

Contrail Service Orchestration Contrail Service Orchestration Installation and Upgrade Guide
4.1.0

Copyright © 2019 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

	About the Documentation	ix
	Documentation and Release Notes	ix
	Documentation Conventions	ix
	Documentation Feedback	xi
	Requesting Technical Support	xii
	Self-Help Online Tools and Resources	xii
	Creating a Service Request with JTAC	xiii
Chapter 1	Introduction	15
	Contrail Service Orchestration Overview	15
Chapter 2	Hardware and Software Requirements	17
	Hardware and Software Tested for Contrail Service Orchestration	17
	Node Servers and Servers Tested in Contrail Service Orchestration	17
	Network Devices and Software Tested in the Hybrid WAN (Distributed CPE) and SD-WAN Deployments	18
	Minimum Requirements for Servers and VMs	20
	Minimum Hardware Requirements for Node Servers and Servers	20
	Minimum Requirements for VMs on CSO Node Servers and Servers	22
Chapter 3	Installing Contrail Service Orchestration with GUI	33
	CSO GUI Installer Overview	33
	About the Downloader Component	34
	About the Installer Component	35
	Installing Contrail Service Orchestration with the GUI Installer	36
	Troubleshooting the CSO GUI Installer-Related Errors	41
	Downloader Component	41
	Installer Component	42
Chapter 4	Installing Contrail Service Orchestration with CLI	43
	Removing a Previous Deployment	43
	Provisioning VMs on Contrail Service Orchestration Nodes or Servers	44
	Before You Begin	45
	Downloading the Installer	45
	Creating a Bridge Interface for KVM	46
	Creating a Data Interface for a Distributed Deployment	47
	Customizing the Configuration File for the Provisioning Tool	49
	Provisioning VMs with the Provisioning Tool for the KVM Hypervisor	72
	Provisioning VMware ESXi VMs Using the Provisioning Tool	73
	Manually Provisioning VRR VMs on the Contrail Service Orchestration Node or Server	75

	Verifying Connectivity of the VMs	76
	Setting Up the Installation Package and Library Access	76
	Copying the Installer Package to the Installer VM	76
	Creating a Private Repository on an External Server	77
	Installing and Configuring Contrail Service Orchestration	78
	Before You Begin	78
	Creating the Configuration Files	81
	Deploying Infrastructure Services	85
	Deploying Microservices	86
	Checking the Status of the Microservices	86
	Loading Data	87
	Taking Snapshot of the VMs	88
	Reverting Snapshots	95
	Performing a Health Check of Infrastructure Components	99
Chapter 5	Post Installation Tasks	103
	Generating and Encrypting Passwords for Infrastructure Components	104
	Applying NAT Rules if CSO Is Deployed Behind NAT	105
	Applying Security Patches	107
	Viewing Information About Microservices	107
	110
Chapter 6	Upgrading Contrail Service Orchestration	111
	Upgrading Contrail Service Orchestration Overview	111
	Limitations	112
	Impact of the CSO Upgrade	112
	Upgrading Contrail Service Orchestration	113
	Adding Virtual Route Reflectors (VRRs) After Upgrading	116
	Troubleshooting Upgrade Related Errors	118
	Salt Synchronization Error	118
	Cache Clearance Error	119
	Kube-system Pod Error	120
	Kubernetes Node Error	121

List of Figures

Chapter 3	Installing Contrail Service Orchestration with GUI	33
	Figure 1: Example Express Install Window	38
	Figure 2: Custom Install Virtual IP Address and Hostname	40

List of Tables

	About the Documentation	ix
	Table 1: Notice Icons	x
	Table 2: Text and Syntax Conventions	x
Chapter 2	Hardware and Software Requirements	17
	Table 3: COTS Node Servers and Servers Tested in the Cloud CPE and SD-WAN Solutions	17
	Table 4: Software Tested for the COTS Nodes and Servers	18
	Table 5: Network Devices Tested for the Distributed Deployment and SD-WAN Implementation	18
	Table 6: Software Tested in the Distributed Deployment and SD-WAN Solution	19
	Table 7: Specification for Node Servers and Servers	20
	Table 8: Server Requirements	21
	Table 9: Details of VMs for a Small Deployment	22
	Table 10: Details of VMs for a Medium Deployment	23
	Table 11: Details of VMs for a Large Deployment	25
	Table 12: Ports to Open on CSO VMs	29
Chapter 3	Installing Contrail Service Orchestration with GUI	33
	Table 13: Resources per Size	35
Chapter 4	Installing Contrail Service Orchestration with CLI	43
	Table 14: Location of Configuration Files for Provisioning VMs	49
Chapter 5	Post Installation Tasks	103
	Table 15: Functions of Microservices	107
Chapter 6	Upgrading Contrail Service Orchestration	111
	Table 16: Impact of the CSO Upgrade	112

About the Documentation

- Documentation and Release Notes on page ix
- Documentation Conventions on page ix
- Documentation Feedback on page xi
- Requesting Technical Support on page xii

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 x defines notice icons used in this guide.

Table 1: Notice Icons







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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>


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.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<pre>[edit] routing-options { static { route default { nexthop <i>address</i>; retain; } } }</pre>
;(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.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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



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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active 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/>

- Join and participate in the Juniper Networks Community Forum:
<https://www.juniper.net/company/communities/>
- Create a service request online: <https://myjuniper.juniper.net>

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

Creating a Service Request with JTAC

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

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

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

CHAPTER 1

Introduction

- [Contrail Service Orchestration Overview on page 15](#)

Contrail Service Orchestration Overview

Contrail Service Orchestration (CSO) is a software platform that designs, secures, automates and runs the entire service life cycle across NFX Series Network Services Platforms, MX Series Routers, and SRX Series Services Gateways, along with the vSRX Virtual Firewall available in public cloud marketplaces.

The solution supports both Juniper Networks and third-party virtualized network functions (VNFs). CSO provides a RESTful APIs to connect with service providers' operational support systems (OSS) and business support systems (BSS) applications. It is responsible for many management and orchestration (MANO) activities in the deployment.

CSO features:

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

- Administration Portal: GUI to manage resources, customers, and availability of network services. It uses the RESTful APIs of other Contrail Service Orchestration components.
- Customer Portal: GUI to manage sites, customer premises equipment (CPE) devices, and network services for organizations.

The portals offer role-based access control (RBAC) for administrators and operators.

This guide provides information about installing and upgrading the Contrail Service Orchestration (CSO) Release 4.1.0 components.

CSO solution software:

- Contrail Service Orchestrator (CSO): Includes both the Network Service Orchestrator (NSO) and Network Service Controller (NSC) to deploy, manage, and monitor CPE devices, and provide topology and CPE lifecycle management functionality.
- CSO Network Service Controller: Includes only the controller layer of CSO to provide topology and CPE lifecycle management functionality, as well as site-to-site routing and reachability, but without the user interface layer. Typically used when integrating with another management system.

CSO Release 4.0.0 and later, you can install CSO by using the new installation GUI as well as through the CLI.



NOTE: The GUI Installation option is available by default for a fresh installation from CSO Release 4.1.0. You can, however, use the CLI, if you prefer.



CAUTION: Currently, the upgrade process is not supported through the GUI.

Deployment types:

CSO 4.1 supports the following types of deployments:

- Small Deployment—Manages approximately 500 devices. You cannot configure high availability with small deployments.
- Medium Deployment—Manages approximately 3500 devices. High availability is supported on medium deployments.
- Large Deployment—Manages approximately 6000 devices. High availability is support on large deployments.

For detailed information on configuring Contrail Service Orchestration, see [Contrail Service Orchestration Deployment Guide](#).

Related Documentation

- *Contrail Service Orchestration (CSO) Deployment Guide*

CHAPTER 2

Hardware and Software Requirements

- [Hardware and Software Tested for Contrail Service Orchestration on page 17](#)
- [Minimum Requirements for Servers and VMs on page 20](#)

Hardware and Software Tested for Contrail Service Orchestration

Contrail Service Orchestration requires commercial off-the-shelf (COTS) node servers or servers, specific network devices, and specific software versions. The following 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 17](#)
- [Network Devices and Software Tested in the Hybrid WAN \(Distributed CPE\) and SD-WAN Deployments on page 18](#)

Node Servers and Servers Tested in Contrail Service Orchestration

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 3 on page 17](#) lists the node servers and servers that have been tested for these functions.

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

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

Table 4 on page 18 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.

Table 4: Software Tested for the COTS Nodes and Servers

Description	Version
Operating system for all COTS nodes and servers	Ubuntu 14.04.5 LTS NOTE: Ensure that you perform a fresh install of Ubuntu 14.04.5 LTS on the CSO servers in your deployment because upgrading from a previous version to Ubuntu 14.04.5 LTS might cause issues with the installation.
Operating system for VMs on CSO servers	<ul style="list-style-type: none"> • 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 onwards NOTE: A mix of different hypervisors across machines is not supported.
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 4.1.2.0-171

Network Devices and Software Tested in the Hybrid WAN (Distributed CPE) and SD-WAN Deployments

Table 5 on page 18 shows the network devices that have been tested for the distributed deployment and the SD-WAN implementation.

Table 5: 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	<ul style="list-style-type: none"> • 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.

Table 5: 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	<ul style="list-style-type: none"> MX104, MX240, MX480, or MX960 router with an Multiservices MIC line card.
	Juniper Networks SRX Series Services Gateway	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.
	vSRX on an x86 server	<ul style="list-style-type: none"> SRX1500 Services Gateway SRX4100 Services Gateway SRX4200 Services Gateway vSRX
CPE device (Hybrid WAN deployment) or spoke device (SD-WAN implementation)	NFX Series Network Services Platforms	<ul style="list-style-type: none"> NFX250-LS1 device
	SRX Series Services Gateways	<ul style="list-style-type: none"> NFX250-S1 device NFX250-S2 device
	vSRX on an x86 server	<ul style="list-style-type: none"> NFX150-S1 NFX150-S1E NFX150-C-S1 NFX150-C-S1-AE/AA NFX150-C-S1E-AE/AA SRX300 Services Gateway SRX320 Services Gateway SRX340 Services Gateway SRX345 Services Gateway vSRX

[Table 6 on page 19](#) shows the software tested for the distributed deployment. You must use these specific versions of the software when you implement a distributed deployment and SD-WAN solution.

Table 6: 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 4.1.0
Contrail Analytics	Contrail Release 4.1.2.0-171
NFX Software	Junos OS Release 15.1X53-D496

Table 6: Software Tested in the Distributed Deployment and SD-WAN Solution (continued)

Function	Software and Version
Routing and Security for NFX250 device	vSRX KVM Appliance 15.1X49-D170
Operating system for vSRX used as a CPE device on an x86 server	vSRX KVM Appliance 15.1X49-D170
Operating system for SRX Series Services Gateway used as a CPE device or spoke device	Junos OS Release 15.1X49-D170
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-D170

Related Documentation

- [Minimum Requirements for Servers and VMs on page 20](#)

Minimum Requirements for Servers and VMs

- [Minimum Hardware Requirements for Node Servers and Servers on page 20](#)
- [Minimum Requirements for VMs on CSO Node Servers and Servers on page 22](#)

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 4 on page 18](#). When you obtain node servers and servers for the Cloud CPE and SD-WAN 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 7 on page 20](#) shows the specification for the node servers and servers for the Cloud CPE or SD-WAN solution.

Table 7: Specification for Node Servers and Servers

Item	Requirement
Storage	<p>Greater than the summation of VMs which are deployed on the node. Storage drive can be one of the following types:</p> <ul style="list-style-type: none"> • Serial Advanced Technology Attachment (SATA) • Serial Attached SCSI (SAS) • Solid-state drive (SSD) <p>NOTE: Solid-state drive (SSD) is preferred storage for better performance.</p>

Table 7: Specification for Node Servers and Servers (continued)

Item	Requirement
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

The number of node servers and servers that you require depends on whether you are installing a small, medium, or large deployment.

[Table 8 on page 21](#) shows the required hardware specifications for node servers and servers in the supported deployments. 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 and Servers” on page 22](#), because some additional resources are required for the system software.

Table 8: Server Requirements

Function	Small Deployment	Medium Deployment	Large Deployment
Contrail Service Orchestration (CSO) Servers			
NOTE: If you install a small deployment 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 small deployment does not include Junos Space. For information about Junos Space VM requirements, see Table 9 on page 22 .			
Number of nodes or servers	1	3	7 <ul style="list-style-type: none"> • 4 central servers • 3 regional servers
vCPUs per node or server	48	48	48
RAM per node or server	256 GB	256 GB	256 GB
Contrail Cloud Platform for a Centralized Deployment			
NOTE: These servers are not needed for Hybrid WAN or SD-WAN solutions			
Number of nodes or servers	1	4–8 <ul style="list-style-type: none"> • 3 nodes for Contrail controller and analytics • 1–4 Contrail compute nodes 	4–28 <ul style="list-style-type: none"> • 3 nodes for Contrail controller and analytics • 1–25 Contrail compute nodes
vCPUs per node or server	4	48	48
RAM per node or server	16 GB	256 GB	256 GB
Total Numbers of Servers			

Table 8: Server Requirements (continued)

Function	Small Deployment	Medium Deployment	Large Deployment
Centralized deployment	2	7–11	10–34
Hybrid WAN or SD-WAN	1	3	7

Minimum Requirements for VMs on CSO Node Servers and Servers

The number of VMs needed and minimum requirements for CSO VMs depend on the deployment environment and whether or not you use high availability (HA):

- For a small deployment, see [Table 9 on page 22](#).
- For a medium deployment, see [Table 10 on page 23](#).
- For a large deployment, see [Table 11 on page 25](#).

For information about the ports that must be open on VMs for all deployments, see [Table 12 on page 29](#).

Use small deployments for managing approximately 500 sites. You cannot configure high availability with small deployments.

[Table 9 on page 22](#) shows details about the VMs for a small deployment.

Table 9: Details of VMs for a Small Deployment

Name of VM	Components That Installer Places in VM	Resources Required
csp-installer-vm	—	<ul style="list-style-type: none"> • 4 vCPUs • 8 GB RAM • 550 GB hard disk storage
csp-central-infravm	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 8 CPU • 48 GB RAM • 500 GB hard disk storage
csp-central-msvm	All microservices, including GUI applications	<ul style="list-style-type: none"> • 8 CPU • 48 GB RAM • 250 GB hard disk storage
csp-contrailanalytics-1	<p>Contrail Analytics for a distributed deployment</p> <p>For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform</p>	<ul style="list-style-type: none"> • 16 vCPUs • 48 GB RAM • 300 GB hard disk storage
csp-regional-sblb	Load balancer for device to Fault Management Performance Management (FMPM) microservice connectivity	<ul style="list-style-type: none"> • 2 vCPUs • 4 GB RAM • 50 GB hard disk storage

Table 9: Details of VMs for a Small Deployment (continued)

Name of VM	Components That Installer Places in VM	Resources Required
csp-space-vm (optional)	Junos Space Virtual Appliance and database—required only if you deploy virtualized network functions (VNFs) that use this EMS	<ul style="list-style-type: none"> • 8 vCPUs • 32 GB RAM • 100 GB hard disk storage
csp-vrr-vm	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 8 GB RAM • 10 GB hard disk storage
csp-central-k8mastervm	Central K8 Master VM	<ul style="list-style-type: none"> • 2 vCPUs • 4 GB RAM • 100 GB hard disk storage

Use medium deployments for approximately 3500 sites. High availability is supported for medium deployments.

[Table 10 on page 23](#) shows details about the VMs for a medium deployment.

Table 10: Details of VMs for a Medium Deployment

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-installer-vm	—	<ul style="list-style-type: none"> • 4 vCPUs • 32 GB RAM • 400 GB hard disk storage
csp-central-infravm1	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-central-infravm2	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-central-infravm3	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-central-lbvm1	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage
csp-central-lbvm2	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage

Table 10: Details of VMs for a Medium Deployment (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-central-lbvm3	Load-balancing applications running only k8master	<ul style="list-style-type: none"> • 2 vCPUs • 4 GB RAM • 100 GB hard disk storage
csp-central-msvm1	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 CPUs • 38 GB RAM • 250 GB hard disk storage
csp-central-msvm2	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-central-msvm3	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-space-vm(optional)	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	<ul style="list-style-type: none"> • 8 vCPUs • 32 GB RAM • 100 GB hard disk storage
csp-central-elkvm1	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 8 GB RAM • 250 GB hard disk storage
csp-central-elkvm2	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 8 GB RAM • 250 GB hard disk storage
csp-central-elkvm3	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 8 GB RAM • 250 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.	<ul style="list-style-type: none"> • 14 vCPUs • 64 GB RAM • 300 GB hard disk storage
csp-contrailanalytics-2	Contrail Analytics for a distributed deployment. For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform.	<ul style="list-style-type: none"> • 14 vCPUs • 64 GB RAM • 300 GB hard disk storage
csp-contrailanalytics-3	Contrail Analytics for a distributed deployment. For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform.	<ul style="list-style-type: none"> • 14 vCPUs • 64 GB RAM • 300 GB hard disk storage

Table 10: Details of VMs for a Medium Deployment (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	<ul style="list-style-type: none"> • 2 vCPUs • 4 GB RAM • 100 GB hard disk storage
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	<ul style="list-style-type: none"> • 2 vCPUs • 4 GB RAM • 100 GB hard disk storage
csp-vrr-vm1	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 10 GB hard disk storage
csp-vrr-vm2	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 10 GB hard disk storage

Use large deployments for managing approximately 6000 sites. High availability is supported for large deployments.

[Table 11 on page 25](#) shows details about the VMs for a large deployment.

Table 11: Details of VMs for a Large Deployment

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-installer-vm	—	<ul style="list-style-type: none"> • 8 vCPUs • 32 GB RAM • 400 GB hard disk storage
csp-central-infravm1	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-central-infravm2	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-central-infravm3	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-central-lbvm1	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage

Table 11: Details of VMs for a Large Deployment (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-central-lbvm2	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage
csp-central-lbvm3	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage
csp-central-msvm1	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-central-msvm2	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-central-msvm3	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-regional-infrvm1	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-regional-infrvm2	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-regional-infrvm3	Third-party applications used as infrastructure services	<ul style="list-style-type: none"> • 12 vCPUs • 48 GB RAM • 500 GB hard disk storage
csp-regional-msvm1	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-regional-msvm2	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage
csp-regional-msvm3	All microservices, including GUI applications	<ul style="list-style-type: none"> • 10 vCPUs • 38 GB RAM • 250 GB hard disk storage

Table 11: Details of VMs for a Large Deployment (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-regional-lbvm1	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage
csp-regional-lbvm2	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage
csp-regional-lbvm3	Load-balancing applications	<ul style="list-style-type: none"> • 2 vCPUs • 6 GB RAM • 100 GB hard disk storage
csp-space-vm (optional)	Junos Space Virtual Appliance and database—required only if you deploy VNFs that use this EMS	<ul style="list-style-type: none"> • 4 vCPUs • 32 GB RAM • 300 GB hard disk storage
csp-central-elkvm1	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 400 GB hard disk storage
csp-central-elkvm2	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 400 GB hard disk storage
csp-central-elkvm3	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 400 GB hard disk storage
csp-regional-elkvm1	Logging applications	<ul style="list-style-type: none"> • 2 vCPUs • 8 GB RAM • 400 GB hard disk storage
csp-regional-elkvm2	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 400 GB hard disk storage
csp-regional-elkvm3	Logging applications	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 400 GB hard disk storage
csp-regional-sblb1	Load balancer for device to FMPM microservice connectivity	<ul style="list-style-type: none"> • 2 vCPUs • 8 GB RAM • 100 GB hard disk storage

Table 11: Details of VMs for a Large Deployment (continued)

Name of VM or Microservice Collection	Components That Installer Places in VM	Resources Required
csp-regional-sblb2	Load balancer for device to FMPM microservice connectivity	<ul style="list-style-type: none"> • 2 vCPUs • 8 GB RAM • 100 GB hard disk storage
csp-vrr-vm1	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-vrr-vm2	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-vrr-vm3	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-vrr-vm4	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-vrr-vm5	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 4 vCPUs • 16 GB RAM • 300 GB hard disk storage
csp-vrr-vm6	Virtual route reflector (VRR)	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 24 vCPUs • 64 GB RAM • 700 GB hard disk storage
csp-contrailanalytics-2	Contrail Analytics for a distributed deployment For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform	<ul style="list-style-type: none"> • 24 vCPUs • 64 GB RAM • 700 GB hard disk storage
csp-contrailanalytics-3	Contrail Analytics for a distributed deployment For a centralized or combined deployment, you use Contrail Analytics in the Contrail Cloud Platform	<ul style="list-style-type: none"> • 24 vCPUs • 64 GB RAM • 700 GB hard disk storage

Table 12 on page 29 shows the ports that must be open on all CSO VMs to enable the following types of CSO communications:

- External—CSO UI and CPE connectivity
- 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 12: 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

Table 12: Ports to Open on CSO VMs (continued)

Port Number	CSO Communication Type	Port Function
5044	Internal	Beats
5543	Internal	Logstash UDP
5601	External	Kibana UI
5665	Internal	Icinga API
5666	Internal	icinga nrpe
5671	Internal	RabbitMQ SSL listener
5672	Internal	RabbitMQ client
6000	Internal	Swift Object Server
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

Table 12: Ports to Open on CSO VMs (continued)

Port Number	CSO Communication Type	Port Function
9091	Internal	xmltec-xmlmail tcp
9101	External and internal	HA proxy exporter
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 Tested for Contrail Service Orchestration on page 17](#)
- [Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 44](#)
- *Number of Sites and VNFs Supported in Contrail Service Orchestration*

CHAPTER 3

Installing Contrail Service Orchestration with GUI

- [CSO GUI Installer Overview on page 33](#)
- [Installing Contrail Service Orchestration with the GUI Installer on page 36](#)
- [Troubleshooting the CSO GUI Installer-Related Errors on page 41](#)

CSO GUI Installer Overview

The CSO installer walks you through the steps needed to install and configure CSO onto your virtual machine (VM.) It provides features such as faster downloading and installation, auto-provisioning of the CSO virtual machines, and a UI for interaction and installation status. The installer can be run on the following operating systems:

- Apple OS X and later.
- Microsoft Windows 10 and later.
- Ubuntu 14.04 and later.



NOTE: Juniper Networks recommends using the CSO installer to install and configure CSO because of its ease-of-use and is less error-prone. You can, however, use the CLI if you prefer.

The overall flow for using the CSO installer is as follows:

1. Download the CSO Downloader from the [CSO Download](#) page.
2. Launch the CSO Downloader and click on Install.
3. Log in using your Juniper Networks credentials.
4. Enter setup information, such as hypervisor and the solution to install.
5. Select your deployment size.

6. Select location to download the CSO package, such as on the existing Installer VM (IVM) or on the host or on the local machine.
7. Enter server credentials and the CSO Downloader downloads the required packages at the specified location.
8. Click **Next** to set up the Installer UI.
9. The CSO Downloader creates and deploys the necessary VMs.
10. Once the installation is successful, the CSO Installer UI is automatically launched. If not, Click on **CSO Installer GUI**.



NOTE: Use the CSO installer only for new installations. You cannot currently upgrade from a previously installed version using the CSO Downloader.

Upgrading from a previously installed version only downloads the CSO packages. It does not perform the actual upgrade. You must use the CLI to upgrade CSO. For more information, see [“Upgrading Contrail Service Orchestration Overview” on page 111](#).

The CSO Downloader and Installer has two main components that perform the steps listed above:

- Downloader—Downloads the CSO packages and creates an Installer Virtual Machine (IVM).
- Installer—Creates the necessary VMs and installs the CSO packages.

About the Downloader Component

The downloader component of the CSO installer performs the following tasks:

- Downloads the necessary CSO packages.
- Creates an IVM.
- Transfers the CSO packages to the IVM.
- Launches the Installer component web-based user interface from the IVM.

The downloader component gives you the option to install now or to install later. The install now option performs all the tasks listed above. The install later option downloads the necessary packages but does not create the IVM or transfer the packages to the IVM.

For example, if you do not know the host machine IP address and root password, you might choose to download the necessary packages now and wait to create the IVM until you have that information.

Or, you might choose install later to download the necessary packages and use the CLI to install CSO.

About the Installer Component

The installer component of the CSO installer starts automatically after the downloader component finishes, and performs the following tasks:

- Creates the required VMs.
- Installs the CSO packages in the VMs.

When running the installer component, you select the following options in addition to configuring the VMs. Each option is described in detail in the installer user interface.

- Size of the network to manage—small, medium, or large. The option you select determines the number of servers and the resource per server required. See [Table 13 on page 35](#).

Table 13: Resources per Size

	Small	Medium	Large
Approximate number of managed sites	500	3500	6000
High availability	No	Yes	Yes
Servers	1	3	7
vCPUs/server	48	48	48
RAM/server	256 GB	256 GB	256 GB
Disk space/server	Greater than the summation of VMs which are deployed on the node.		

- *Express* install or *Custom* install. The express install uses pre-defined defaults and requires less user input. Select custom install if you want full control over the installation and configuration parameters.
- Network type — CSO reachable directly or CSO behind a NAT gateway.

CSO reachable directly means you can access the managed devices and CSO without going through a NAT gateway. Here, the CSO and devices IP's are routable to each other within the enterprises private network. This topology is common in a campus environment where multiple locations are connected through VPNs as a single logical private network.

For CSO behind a NAT gateway, the managed devices are typically remote devices not residing in the data center where CSO is installed. These devices reach CSO through a NAT gateway using a public IP address exposed for the data center. This topology is common when CSO manages customer's remote or on-premises devices. For example, the branch locations of a bank or restaurant chain.

Installing Contrail Service Orchestration with the GUI Installer

If you prefer to install CSO using the CLI, following the instructions below and select the **Install Later** option to download the CSO files to your local drive. Manually transfer the CSO files to your installer VM using scp, ftp, or other similar programs. Log in to your installer VM and follow the instructions described in “[Installing and Configuring Contrail Service Orchestration](#)” on page 78.



NOTE: Upgrading from a previously installed version via GUI only downloads the CSO packages. It does not perform the actual upgrade. You must use the CLI to upgrade CSO. For more information, see “[Upgrading Contrail Service Orchestration Overview](#)” on page 111.

During IVM creation, it will use DNS information. Make sure the DNS information is correct before you begin CSO installation process.

To download and run the CSO installer:

1. From your browser, go to the [CSO download](#) page.
2. On the page that appears, click the **Software** tab and select **4.1** from the Version drop-down menu.
3. Click the CSO Downloader link corresponding to your operating system to download the file to your local drive.

The CSO Downloader will be downloaded to the local PC. During the install process, the CSO Downloader is then downloaded from the local PC to the server. If you need to start the install process over, you can save time by using the same CSO Downloader from the given location.

4. Locate the file on your local drive and launch it.
 - For Windows, double-click the executable file.
 - For Linux, enter the following command:
dpkg -i cso-downloader.deb
 - For macOS, drag the **.dmg** file to the installation window.




NOTE: When you install the CSO Downloader on macOS, you might receive an error message indicating that the application cannot be opened because it is from an unidentified developer. As a workaround, open the Security & Privacy Systems Preference. Click the **General** tab and click **Open Anyway**.

5. Install the CSO Downloader.
6. Enter your Juniper Networks support credentials and click **Next**.
7. Follow the instructions to specify the software setup, including:
 - CSO release version to install.
 - Hypervisor server type—KVM or ESXi.
 - New installation or upgrade from a previously installed version.
 - CSO solution to install—Contrail Service Orchestrator or CSO Network Service Controller.
8. Read and accept the license agreement, then click **Next**.
9. Select CSO Deployment Type—Small, or Medium, or Large.
10. Select location to download CSO package—Existing Installer VM, or Download on Host, or Local Machine

The CSO Downloader performs health check to verify connectivity, memory availability, host OS version, kernel version, disk space, and DNS server.
11. Enter the required details.

CSO downloader verifies the entered details and validates IP availability, virtual bridge, gateway ping, and subnet mask.



NOTE: Remember the password as you will need it later.
12. Click **Next** to download the CSO Installer.

Additional files are downloaded as needed. Depending on your Internet bandwidth, it might take 30 minutes or more to download the files. After the download is complete, the CSO Downloader verifies the MD5 checksum of each file.
13. Click **Next** after the download is complete to proceed with setting up the Installer UI.
14. After setting up the Installer UI, you are redirected to the Installer UI webpage automatically. On the installer welcome page, enter the IVM password you created in Step 11 and click **Login**.
15. Identify the size of the network to be managed and click either **Express** or **Custom**.

The express install uses predefined defaults and requires less user input. Whichever option you select, you can click **Back** on the next page to return to this page to select the other option.

16. If you select the *Express* install, the Express window appears.

[Figure 1 on page 38](#) shows an example of the Express window for a small install.

Figure 1: Example Express Install Window

EXPRESS (small)

VM Provisioning Details

Details needed for provisioning the VMs - Physical server IPs, and IPs for VMs

Hypervisor* ESXi

Physical Servers* ② 1 required

Hosts	IP Address/Mask	Root Password	Datastore (ESXi only)	VM Network	VM Gateway	Status
HOST 1	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check

central: 6 VMs

For small and medium-sized managed networks, all hosts belong to a common CSO cluster or region.

For large networks, you must have a minimum of two regions. By default, a central region and regional region is configured for you. You can add an additional two more regions. Note that adding more regions requires more physical hosts.

- For each host, enter the IP address with subnet mask, the root password, select the VM network (or datastores for ESXi servers) from the drop-down menu, and VM gateway.
- Select the network type—**CSO Directly Reachable** or **CSO Behind NAT**. For more information, see [“CSO GUI Installer Overview” on page 33](#).
- Enter the IP addresses for the VMs.
 - Click **Input IP Range** to add the IP addresses as a range.
 - Click **Input IP** to add a list of individual IP addresses, separated with a comma.
- (CSO Behind NAT only) Enter the central NAT gateway IP address.

This is the NAT gateway public-facing IP address.

- Enter the regional NAT gateway IP address.

Each CSO region or cluster can have a different NAT gateway or the same NAT gateway.

- Enter the NTP server IP address or FQDN name.

- Verify the default Kubernetes overlay network IP address and subnet mask and update as needed.

- Click **Install**.

The CSO installer now creates the required CSO VMs and installs services within these VMs. A status window displays the progress.

17. If you select the *Custom* install, the Custom install window appears.
 - a. For each host, enter the IP address with subnet mask, the root password, and select the VM network (or datastores for ESXi servers) from the drop-down menu, and VM gateway.
 - b. For each Contrail Analytics Node (CAN), enter the IP address with subnet mask, the root password, the VM network (or datastores for ESXi servers) from the drop-down menu, and VM gateway.
 - c. Select the network type—**CSO Directly Reachable** or **CSO Behind NAT**. For more information, see [“CSO GUI Installer Overview” on page 33](#).
 - d. Enter the IP addresses for the VMs.
 - Click **Input IP Range** to add the IP addresses as a range.
 - Click **Input IP** to add a list of individual IP addresses, separated with a comma.
 - e. (CSO Behind NAT only) Enter the central NAT gateway IP address.

This is the NAT gateway public-facing IP address.
 - f. Enter the regional NAT gateway IP address.

Each CSO region or cluster can have a different NAT gateway or the same NAT gateway.
 - g. Enter the NTP server IP address or FQDN name.
 - h. Verify the default Kubernetes overlay network IP address and subnet mask and update as needed.
 - i. For each region, enter the virtual IP address and hostname. See [Figure 2 on page 40](#).

Figure 2: Custom Install Virtual IP Address and Hostname

CUSTOM (large)

VM Provisioning Details
Details needed for provisioning the VMs - Physical server IPs, and IPs for VMs

Regions*

Hypervisor* ESXi

Physical Servers* 6 required

Hosts	IP Address/Mask	Root Password	Datastore(ESXi only)	VM Network	VM Gateway	Status
HOST 1	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check
HOST 2	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check
HOST 3	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check

Servers for Contrail Analytics Node (CAN)* 3 required

CAN Servers	IP Address/Mask	Root Password	Datastore(ESXi only)	VM Gateway	Status
CAN Server 1	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check
CAN Server 2	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check
CAN Server 3	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Enter details to check

central: 12 VMs

regional: 20 VMs

For small and medium-sized managed networks, all hosts belong to a common CSO cluster or region.

For large networks, you must have a minimum of two regions. By default, a central and regional region is configured for you. You can add an additional two more regions. Note that adding more regions requires more physical hosts.

For secure communication between devices to CSO services, digital certificates must be uploaded to CSO. These certificates are mapped to a hostname or an associated IP / Virtual IP (VIP) address. For example, two certificates can be uploaded for a single VM. One maps to its VIP and the other to its hostname. Typically, one hostname or VIP is required per region. In case of small and medium installs, there is no concept of different regions. All VMs are part of a common central region. Therefore, a single set of certificates is uploaded for this central region. In case of a large install, where the CSO VMs are replicated across multiple regions, one or more certificates for each region should be uploaded.

CSO generates certificates for VIPs and hostnames. However, these certificates are not signed by a trusted Certificate Authorities (CA). This results in "untrusted site" and "add security exception" warning messages to users. If you prefer to upload trusted CA signed certificates tied to a specific VIP or hostname, you can upload them to CSO (one or more for each region).

The regional VIP and hostname must be provided wherever high availability of services is available. For example, in medium and large installs. The small install does not provide high availability. The VIP address is used as a front-end address for a set of load-balancers of a region in the back-end. In addition, an FQDN hostname should be provided for this VIP address. For example, **csso-central.domain.net**. Note that the uploaded signed certificate should be generated for this hostname.

- j. Click **Choose Files** to locate the certificate for that specific region.
- k. Review the default component configuration settings and update as needed.
- l. If you select **External** from the Keystone Service menu, enter the following additional information:

- Keystone IP address.
- Keystone administrator password.
- External Keystone service token.

If you do not know the service token, log in to your external keystone server. View the `/etc/keystone/keystonerc` file and search for the `ES_SERVICE_TOKEN` variable. For example:

```
export OS_SERVICE_TOKEN=abcdefg1234567
```

- Keystone administrator e-mail address.
- m. Click **Install**.

The CSO installer now creates the required CSO VMs and installs services within these VMs. A status window displays the progress.



NOTE: Once the installation process is completed, click on **show CSO-passwords** link. You will need this password for the initial login.

- 18. Select launch CSO-Admin portal to login to CSO. The initial username is **cspadmin** and the initial password is the password noted in the previous step. Once you login, you will be asked to change the password.

Related Documentation

- [Hardware and Software Tested for Contrail Service Orchestration on page 17](#)
- [Removing a Previous Deployment on page 43](#)

Troubleshooting the CSO GUI Installer-Related Errors

Use the following to troubleshoot software issues with the CSO GUI Installer.

Downloader Component

If you encounter any errors with the downloader component, forward the installer log file to Juniper Networks Technical Support. You can access the application (installation) logs from within the CSO Downloader GUI. To access the logs, click the **View Logs** icon at the top-right corner of the GUI.

The installer log file is located at:

`<home>/juniper/application.log`

For example, `C:\Users\bob\juniper\application.log`.

Installer Component

The installer component screens validate user entries. For example, IP addresses are pinged to verify they are valid and reachable. However, errors can occur after you click **Install**. When this happens, you are presented the following two options:

- **Retry**—This option retries the operation that failed.
- **View Logs**—This option lets you download the installer component logs, which you can then send to Juniper Networks Technical Support for assistance.

CHAPTER 4

Installing Contrail Service Orchestration with CLI

- [Removing a Previous Deployment on page 43](#)
- [Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 44](#)
- [Setting Up the Installation Package and Library Access on page 76](#)
- [Installing and Configuring Contrail Service Orchestration on page 78](#)

Removing a Previous Deployment

You should remove a previous deployment and perform a new installation if the architecture of the VMs on the CSO server node or server has changed significantly between releases.

If you do not have previous deployment, proceed with [“Provisioning VMs on Contrail Service Orchestration Nodes or Servers” on page 44](#)

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
```

Output:

Id	Name	State
2	csp-ui-vm	running

- 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 <vm-name>
```

2. Delete the Salt minion keys.

For example:

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

Related Documentation

- [Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 44](#)

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 components. All servers and VMs for the solution should be in the same subnet.

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

The VMs created by provisioning tool have Ubuntu preinstalled.



NOTE: If you use the KVM hypervisor while installing a Distributed CPE (Hybrid WAN) or an SD-WAN solution, you must create a bridge interface on the physical server. The bridge interface should map the primary network interface (Ethernet management interface) on each CSO server node or server to a virtual interface before you create VMs. This action enables the VMs to communicate with the network.

This approach is applicable only if you are installing a Distributed CPE (Hybrid WAN) or an SD-WAN solution. It is not required for a centralized solution.

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

- Small deployment. (See [Table 9 on page 22](#))
- Medium deployment. See ([Table 10 on page 23](#))
- Large deployment. See ([Table 11 on page 25](#))

The small and medium deployments are always region-less deployment whereas the large deployment is always region-based deployment.

See “[Minimum Requirements for Servers and VMs](#)” on [page 20](#) for details of the VMs and associated resources required for each deployment.

The following sections describe the procedures for provisioning the VMs:

- [Before You Begin on page 45](#)
- [Downloading the Installer on page 45](#)
- [Creating a Bridge Interface for KVM on page 46](#)
- [Creating a Data Interface for a Distributed Deployment on page 47](#)
- [Customizing the Configuration File for the Provisioning Tool on page 49](#)
- [Provisioning VMs with the Provisioning Tool for the KVM Hypervisor on page 72](#)
- [Provisioning VMware ESXi VMs Using the Provisioning Tool on page 73](#)
- [Manually Provisioning VRR VMs on the Contrail Service Orchestration Node or Server on page 75](#)
- [Verifying Connectivity of the VMs on page 76](#)

Before You Begin

Before you begin you must:

- Configure the physical servers or node servers and nodes.
- Install Ubuntu 14.04.5 LTS as the operating system for the physical servers.
- Configure the Contrail Cloud Platform and install Contrail OpenStack if you are performing a centralized CPE deployment.

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 the [CSO Download](#) page.

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

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

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

This installer 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
```

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

A physical server or node needs Internet access to install the **libvirt-bin** package.

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 **/etc/network/interfaces** file 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
    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. Reboot the physical machine and log in as root again.
8. Verify that the primary network interface is mapped to the virbr0 interface.

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

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

Creating a Data Interface for a Distributed Deployment

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

A physical server or node needs Internet access to install **libvirt-bin** package.

To create a data interface:

1. Log in to 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 virbr1
root@host:~/# virsh brctl addif virbr1 eth1
```

3. Create a file with the name **virbr1.xml** in the **/var/lib/libvirt/network** directory.
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
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 192.0.2.1
    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
```

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 sample configuration files. For example, if the name of the installer directory is **csoVersion**.

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

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

[Table 14 on page 49](#) shows the directories that contain the sample configuration file.

Table 14: Location of Configuration Files for Provisioning VMs

Deployment	Directory for Sample Configuration File
Small deployment	confs/cso4.1.0/production/nonha/provisionvm/provision_vm_collocated_example.conf
Medium deployment	confs/cso4.1.0/production/ha/provisionvm/provision_vm_collocated_example.conf
Large deployment	confs/cso4.1.0/production/ha/provisionvm/provision_vm_example.conf

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

For example:

```
root@host:~/cspVersion# cp
confs/cso4.1.0/production/nonha/provisionvm/provision_vm_collocated_example.conf
provision_vm.conf
```

5. Open the **provision_vm.conf** file 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 (FQDNs) 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 Interdomain routing (CIDR) notation.
 - **management_interface**—Name of the Ethernet management interface, virbr0.
 - **gateway**—IP address of the gateway for the host.
 - **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. Specify the name of the data interface, such as virbr1, that you configured for a distributed deployment.
8. 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.
 - **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).
 - **vm_type**—Preinstalled OSS with VM types of baseInfra/baseMS.
 - **volumes**—Data partitions of CSO (Eg: /mnt/data:400G).
 - **base_disk_size**—OS partitions of CSO (Eg: 100G).
10. Save the file.
11. Download the ESXi-4.1.0.tar.gz file to the `/root/Contrail_Service_Orchestration_4.1.0/artifacts` folder on the Ubuntu VM.
- or
- Download the KVM-4.1.0.tar.gz file to the `/root/Contrail_Service_Orchestration_4.1.0/artifacts` folder on the Ubuntu VM.
12. Run the following command to start virtual machines:
- For KVM:

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

For ESXi:

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

The following sections show examples of customized configuration files for a small, a medium, and a large deployment.

Sample Configuration File for Provisioning VMs in a Small Deployment

```
# 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. Don't change the below server
names
server = csp-central-infravm, csp-central-msvm, csp-central-k8mastervm,
csp-installer-vm, csp-contrailanalytics-1, csp-vrr-vm, csp-regional-sblb

# 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
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
```

```
base_disk_size = 100G

[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 = 32768
vcpu = 6
enable_data_interface = false
vm_type = baseMS
volumes = swap:32G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 4096
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 8192
vcpu = 4
enable_data_interface = false
vm_type =
volumes = swap:24G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

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

```

vcpu = 8
enable_data_interface = false
vm_type =
volumes = swap:48G,/data1:1G,/data2:1G
base_disk_size = 500G

[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 = 4096
vcpu = 2
enable_data_interface = true
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 500G

[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 Medium Deployment

```

# 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-host1, cso-host2, cso-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-central-lbvm1, csp-central-lbvm2, csp-central-lbvm3, csp-installer-vm,
csp-contrailanalytics-1, csp-contrailanalytics-2, csp-contrailanalytics-3,
csp-central-elkvm1, csp-central-elkvm2, csp-central-elkvm3, csp-central-msvm1,
csp-central-msvm2, csp-central-msvm3, csp-regional-sblb1, csp-regional-sblb2,
csp-vrr-vm1, csp-vrr-vm2

# 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 = passwd
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 = passwd
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 = passwd
data_interface =

# VM Details
[csp-central-infravm1]
management_address = 192.168.1.6/24
hostname = centralinfravm1.example.net
username = root
password = passwd
local_user = infravm
local_password = passwd
guest_os = ubuntu

```

```
host_server = cso-host1
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-infravm2]
management_address = 192.168.1.7/24
hostname = centralinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-infravm3]
management_address = 192.168.1.8/24
hostname = centralinfravm3.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-installer-vm]
management_address = 192.168.1.9/24
hostname = installervm.example.net
username = root
password = passw0rd
local_user = installervm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 32768
vcpu = 4
enable_data_interface = false
vm_type =
volumes = swap:24G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-lbvm1]
management_address = 192.168.1.10/24
hostname = centrallbvm1.example.net
username = root
```



```

password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 6144
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:6G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-lbvm2]
management_address = 192.168.1.11/24
hostname = centrallbvm2.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 6144
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:6G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

# Lbvm3 is running only k8master
[csp-central-lbvm3]
management_address = 192.168.1.12/24
hostname = centrallbvm3.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 4096
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-elkv1]
management_address = 192.168.1.13/24
hostname = centralelkv1.example.net
username = root
password = passw0rd
local_user = elkv1
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 8192
vcpu = 4
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

```

```
[csp-central-elkvm2]
management_address = 192.168.1.14/24
hostname = centralelkvm2.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 8192
vcpu = 4
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-elkvm3]
management_address = 192.168.1.15/24
hostname = centralelkvm3.example.net
username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 8192
vcpu = 4
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-msvm1]
management_address = 192.168.1.16/24
hostname = centralmsvm1.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host1
memory = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-msvm2]
management_address = 192.168.1.17/24
hostname = centralmsvm2.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 38912
vcpu = 10
```

```
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-msvm3]
management_address = 192.168.1.18/24
hostname = centralmsvm3.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-regional-sblb1]
management_address = 192.168.1.19/24
hostname = sblb1.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 4096
vcpu = 2
enable_data_interface = true
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-regional-sblb2]
management_address = 192.168.1.20/24
hostname = sblb2.example.net
username = root
password = passw0rd
local_user = sblb
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 4096
vcpu = 2
enable_data_interface = true
vm_type =
volumes = swap:4G,/data1:1G,/data2:1G
base_disk_size = 500G

[csp-contrailanalytics-1]
management_address = 192.168.1.21/24
hostname = can1.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
```

```
guest_os = ubuntu
host_server = cso-host1
memory = 65536
vcpu = 14
enable_data_interface = false
vm_type =
volumes = swap:64G,/data1:1G,/data2:1G
base_disk_size = 500G
```

```
[csp-contrailanalytics-2]
management_address = 192.168.1.22/24
hostname = can2.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host2
memory = 65536
vcpu = 14
enable_data_interface = false
vm_type =
volumes = swap:64G,/data1:1G,/data2:1G
base_disk_size = 500G
```

```
[csp-contrailanalytics-3]
management_address = 192.168.1.23/24
hostname = can3.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-host3
memory = 65536
vcpu = 14
enable_data_interface = false
vm_type =
volumes = swap:64G,/data1:1G,/data2:1G
base_disk_size = 500G
```

```
[csp-vrr-vm1]
management_address = 192.168.1.24/24
hostname = vrr1.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-host2
memory = 16384
vcpu = 4
```

```
[csp-vrr-vm2]
management_address = 192.168.1.25/24
hostname = vrr2.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-host3
```

```
memory = 16384
vcpu = 4

[csp-space-vm]
management_address = 192.168.1.28/24
web_address = 192.168.1.29/24
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-host2
memory = 32768
vcpu = 4
```

Sample Configuration File for Provisioning VMs in a Large Deployment

```
# 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-central-host4, 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-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-regional-sblb1,
csp-regional-sblb2, csp-vrr-vm1, csp-vrr-vm2, csp-vrr-vm3, csp-vrr-vm4,
csp-vrr-vm5, csp-vrr-vm6

# Physical Server Details
[cso-central-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-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-central-host4]
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-central-host4
username = root
password = passw0rd
data_interface =

[cso-regional-host1]
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-host1
username = root
password = passw0rd
data_interface =

[cso-regional-host2]
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-host2
username = root
password = passw0rd
data_interface =

[cso-regional-host3]
management_address = 192.168.1.8/24
management_interface = virbr0

```

```
gateway = 192.168.1.1
dns_search = example.net
dns_servers = 192.168.10.1
hostname = cso-regional-host3
username = root
password = passw0rd
data_interface =

# VM Details

[csp-central-infravm1]
management_address = 192.168.1.9/24
hostname = centralinfravm1.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-infravm2]
management_address = 192.168.1.10/24
hostname = centralinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host4
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-infravm3]
management_address = 192.168.1.11/24
hostname = centralinfravm3.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host3
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-regional-infravm1]
management_address = 192.168.1.12/24
```

```

hostname = regionalinfravm1.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host1
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-regional-infravm2]
management_address = 192.168.1.13/24
hostname = regionalinfravm2.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host2
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-regional-infravm3]
management_address = 192.168.1.14/24
hostname = regionalinfravm3.example.net
username = root
password = passw0rd
local_user = infravm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-regional-host3
memory = 49152
vcpu = 12
enable_data_interface = false
vm_type = baseInfra
volumes = swap:48G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-space-vm]
management_address = 192.168.1.15/24
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-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 = 24576
vcpu = 4
enable_data_interface = false
vm_type =
volumes = swap:32G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-contrailanalytics-1]
management_address = 192.168.1.17/24
hostname = can1.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 65536
vcpu = 24
enable_data_interface = false
vm_type =
volumes = swap:64G,/data1:1G,/data2:1G
base_disk_size = 500G

[csp-contrailanalytics-2]
management_address = 192.168.1.18/24
hostname = can2.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 65536
vcpu = 24
enable_data_interface = false
vm_type =
volumes = swap:64G,/data1:1G,/data2:1G
base_disk_size = 500G

[csp-contrailanalytics-3]
management_address = 192.168.1.19/24
hostname = can3.example.net
username = root
password = passw0rd
local_user = canvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host4
memory = 65536
vcpu = 24
```

```
enable_data_interface = false
vm_type =
volumes = swap:64G,/data1:1G,/data2:1G
base_disk_size = 500G

[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 = 6144
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:6G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 6144
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:6G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 4096
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-regional-lbvm1]
management_address = 192.168.1.23/24
hostname = regionallbvm1.example.net
username = root
password = passw0rd
local_user = lbvm
local_password = passw0rd
```

```

guest_os = ubuntu
host_server = cso-regional-host1
memory = 6144
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:6G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 6144
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:6G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 4096
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-elkv1]
management_address = 192.168.1.26/24
hostname = centralelkv1.example.net
username = root
password = passw0rd
local_user = elkv1
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 8192
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-elkv2]
management_address = 192.168.1.27/24
hostname = centralelkv2.example.net

```

```

username = root
password = passw0rd
local_user = elkvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host2
memory = 8192
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 8192
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 8192
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 8192
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

```

```
[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 = 8192
vcpu = 2
enable_data_interface = false
vm_type =
volumes = swap:8G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-central-msvm1]
management_address = 192.168.1.32/24
hostname = centralmsvm1.example.net
username = root
password = passw0rd
local_user = msvm
local_password = passw0rd
guest_os = ubuntu
host_server = cso-central-host1
memory = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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-host4
memory = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 38912
vcpu = 10
```

```

enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 38912
vcpu = 10
enable_data_interface = false
vm_type = baseMS
volumes = swap:38G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 4096
vcpu = 2
enable_data_interface = true
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[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 = 4096
vcpu = 2
enable_data_interface = true
vm_type =
volumes = swap:4G,/mnt/data:400G,/data2:1G
base_disk_size = 100G

[csp-vrr-vm1]
management_address = 192.168.1.40/24
hostname = vrr1.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-central-host1
memory = 16384
vcpu = 4

[csp-vrr-vm2]
management_address = 192.168.1.41/24
hostname = vrr2.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-central-host2
memory = 16384
vcpu = 4

[csp-vrr-vm3]
management_address = 192.168.1.42/24
hostname = vrr3.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-central-host3
memory = 16384
vcpu = 4

[csp-vrr-vm4]
management_address = 192.168.1.43/24
hostname = vrr4.example.net
gateway = 192.168.1.1
newpassword = passw0rd
```

```
guest_os = vrr
host_server = cso-regional-host1
memory = 16384
vcpu = 4

[csp-vrr-vm5]
management_address = 192.168.1.44/24
hostname = vrr4.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-regional-host2
memory = 16384
vcpu = 4

[csp-vrr-vm6]
management_address = 192.168.1.45/24
hostname = vrr4.example.net
gateway = 192.168.1.1
newpassword = passw0rd
guest_os = vrr
host_server = cso-regional-host3
memory = 16384
vcpu = 4
```

Provisioning VMs with the Provisioning Tool for the KVM Hypervisor

If you use the KVM hypervisor for the CSO node or server, you can use the provisioning tool to create and configure the VMs for the CSO and Junos Space components.

The VMs provisioned by the tool will have pre-installed Ubuntu and 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
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 by using the provisioning tool. You must provision the VRR VM manually.

Before you begin, ensure that the maximum supported file size for a 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 with the ESXi host and run the **vmfstools -P *datastorePath*** command.

To provision VMware ESXi VMs using the provisioning tool:

1. Download the CSO Release 4.1.0 installer package from the [CSO Download](#) page to the local drive.
2. Log in as root to the Ubuntu VM with Internet access and kernel version 4.4.0-31-generic. 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_4.1.0.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_4.1.0.tar.gz**:

```
root@host:~/# tar -xvzf Contrail_Service_Orchestration_4.1.0.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_4.1.0/confs
root@host:~/Contrail_Service_Orchestration_4.1.0/confs#
```

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

For example:

```
root@host:~/Contrail_Service_Orchestration_4.1# cp
/confs/cso4.1.0/production/nonha/provisionvm/provision_vm_collocated_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 (FQDNs) 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 Interdomain Routing (CIDR) notation of the VM network. For example, 192.0.2.0/24.
 - **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**—Label 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 quotation marks..

- **datastore**—Datastore value to save all VM 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.

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

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

```
root@host:~/Contrail_Service_Orchestration_4.1.0/# ./provision_vm_ESXI.sh
```

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

For example:

```
root@host:~/Contrail_Service_Orchestration_4.1.0/# scp confs/provision_vm.conf
root@installer_VM_IP:/root/Contrail_Service_Orchestration_4.1.0/confs
```

This action of provisioning VMs using the Provisioning Tool brings up VMware ESXi VMs with the configuration provided in the files.

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

To manually provision the VRR VM:

1. Download the VRR Release 15.1R6.7 software package (.ova format) for VMware from the [Virtual Route Reflector](#) page, to a location accessible to the server.
2. Launch the VRR by 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:~/# New Password: <password>
root@host:~/# Retype New Password: <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 will fail.

Related Documentation

- [Installing and Configuring Contrail Service Orchestration on page 78](#)

Setting Up the Installation Package and Library Access

- [Copying the Installer Package to the Installer VM on page 76](#)
- [Creating a Private Repository on an External Server on page 77](#)

Copying the Installer Package to the Installer VM

After you have provisioned the VMs, perform the following steps:

1. Copy the installer package file from the central CSO server to the installer VM.



TIP: Copy the installer package from the Ubuntu VM to the installer VM for ESXi. After provisioning the VMs, copy the installer package from the central 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 by 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 **provision_vm.conf** file 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.
- 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 78](#).

- Related Documentation**
- [Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 44](#)
 - [Installing and Configuring Contrail Service Orchestration on page 78](#)

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.

- [Before You Begin on page 78](#)
- [Creating the Configuration Files on page 81](#)
- [Deploying Infrastructure Services on page 85](#)
- [Deploying Microservices on page 86](#)
- [Checking the Status of the Microservices on page 86](#)
- [Loading Data on page 87](#)
- [Taking Snapshot of the VMs on page 88](#)
- [Reverting Snapshots on page 95](#)
- [Performing a Health Check of Infrastructure Components on page 99](#)

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 44](#)).
- Copy the installer package to the installer VM and expand it. (See [“Setting Up the Installation Package and Library Access” on page 76](#))
- If you have created an installer VM using the provisioning tool, you must copy the `/Contrail_Service_Orchestration_4.1.0/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 76](#)).

The installation process uses a private repository so that you do not need Internet access during the installation.

- Determine the following information:
 - The size of deployment: small, medium, large.
 - For large deployment, 2 regions are configured by default - *central* and *regional*. You can add up to 2 additional regions with desired names. (e.g.- "Tokyo" or "East-Coast").
 - The IP address of the VM that hosts the installer.
 - The time zone for the servers in the deployment, based on the Ubuntu time zone guidelines.

The default value for this setting is the current time zone 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 the Administration Portal with the single sign-on method, enter the name of the public domain in which the CSO servers reside. Alternatively if you want to access the Administration Portal with local authentication, you need to enter 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 must use TLS unless you have an explicit reason for not encrypting data between the CPE device and CSO.

- 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:
 - 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 0–16,385, and the default range is 4000–6000.
 - The FQDN that the load balancer uses to access the installation.
 - For the small deployment, the IP address and the FQDN of the VM that hosts the HAProxy.
 - For medium or large deployments, the virtual IP address and the associated hostname that you configure for the HAProxy.
 - The replication factor for each microservice is pre-determined based on the size of the deployment.

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`:

```
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.



NOTE: Take a note of a password shown. It is later used as the initial password for the CSO Administration portal.

4. Specify the management IP address of the VM that hosts the installer file.
5. 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 105](#).
 - n—CSO is not behind NAT (default).
6. Accept the default time zone or specify the Ubuntu time zone for the servers in the topology.
7. Specify a comma-separated list of FQDNs of NTP servers.
For example: `ntp.example.net, ntp.example.com`
8. Specify whether the deployment uses high availability (HA).
 - y—Deployment uses HA.
 - n—Deployment does not use HA.
9. Specify whether the deployment has multiple regions:
 - y—Deployment uses multiple regions.

- n—Deployment uses single region.
10. Press enter if you use only one region or specify a comma-separated list of 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. For a distributed deployment, specify whether you use TLS to enable secure communication between the CPE device and CSO.
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.
 13. Specify the e-mail address of Admin User.
 14. Specify a domain name to determine how you access the Administration Portal, the main CSO GUI:
 - If you want to access the 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 the Administration portal, you specify a dummy name.
For example: *example.net*
 15. 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. We recommend that you use this default option 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.
 16. 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.
 17. Specify whether you use a common password for all CSO VMs, and if so, specify the password.

18. 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, we recommend that you create VRRs in even numbers, and you must create 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.
 - 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.

19. 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.



NOTE: The IP addresses 172.16.0.0/16 and 192.168.3.0/24 must not be used for CSO VMs.

- 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.

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 0–16,385 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 replication factor for each microservice is pre-determined based on the size of the deployment.

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

20. 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 HAProxy 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.

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

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

22. Specify the range for tunnel interface unit.

23. 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.

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

The commands shown depends on the deployment size.

- For small/medium deployments:

```
root@host:~/# ./deploy_infra_services.sh
root@host:~/# ./deploy_micro_services.sh
```

- For large deployment:

```
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 Installer VM.

2. Deploy the central infrastructure services.

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



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

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

```
root@host:~/# "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.



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.

Deploying Microservices

To deploy the microservices:

1. Log in as root to the Installer VM.
2. Deploy the central microservices.

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



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

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

```
root@host:~/# "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 with required region – central or regional.

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

NAME	READY	STATUS	RESTARTS	AGE
------	-------	--------	----------	-----

Region can only be **central** or **regional**. For any other customized region (e.g. - site1), it is listed under "regional".

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
```

NAME	READY	STATUS	RESTARTS	AGE
csp.ams-3909406435-4yb01	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 -n <region>
```

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

```
root@host:~/# kubectl get pods -n <region>
```

NAME	READY	STATUS	RESTARTS	AGE
csp.ams-4890899323-3dfd02	1/1	Running	0	1m
csp.nso-core-09009278633-fr234f	0/1	Running	0	1m

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 `/micro_services/data/inputs.yaml` file 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.

Taking Snapshot of the VMs

At the end of “Loading Data” on page 87, the following message is displayed:

Please run `snapshot.sh` to take snapshot of all the VMs. This enables to perform recovery from single node VM failures. It is a recommended step for production setups.

This utility will not take snapshot of the installer VM and the VRR VMs.



NOTE: This script needs `confs/provision_vm.conf` file to be present on the Installer VM.

The VMs on KVM host are required to shut down before taking snapshots. The VMs will reboot on running this script.

The VMs on ESXi host are not required to shut down on running this script.



CAUTION: Make sure that the `vmname` field is not empty for any of the VMs in `confs/provision_vm.conf` file. Otherwise, the script will fail.

1. On the installer VM, navigate to **Contrail_Service_Orchestration_4.1.0**

2. Run `./snapshot.sh`

Sample output:

```
2019-02-05 20:52:36 INFO      snapshot snapshot.py
=====
2019-02-05 20:52:36 INFO      snapshot snapshot.py Snapshot Start Time:
2019-02-05 20:52:36.694287
2019-02-05 20:52:36 INFO      snapshot snapshot.py
=====
2019-02-05 20:52:36 INFO      snapshot snapshot.py Going to perform CSO Health
Check
2019-02-05 20:52:54 INFO      snapshot check_infracomponent.py Health Check for
Infrastructure Component SaltStack Started
2019-02-05 20:52:55 INFO      snapshot check_infracomponent.py The Infrastructure
Component SaltStack is Healthy
```



```

2019-02-05 20:52:55 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Cassandra Started
2019-02-05 20:52:56 INFO    snapshot check_infracomponent.py The Infrastructure
Component Cassandra is Healthy

2019-02-05 20:52:56 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component ElasticSearch Started
2019-02-05 20:52:58 INFO    snapshot check_infracomponent.py The Infrastructure
Component ElasticSearch is Healthy

2019-02-05 20:52:58 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Etcd Started
2019-02-05 20:52:59 INFO    snapshot check_infracomponent.py The Infrastructure
Component Etcd is Healthy

2019-02-05 20:52:59 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component MariaDb Started
2019-02-05 20:53:01 INFO    snapshot check_infracomponent.py The Infrastructure
Component MariaDb is Healthy

2019-02-05 20:53:01 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component RabbitMQ Started
2019-02-05 20:53:04 INFO    snapshot check_infracomponent.py The Infrastructure
Component RabbitMQ is Healthy

2019-02-05 20:53:04 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component ZooKeeper Started
2019-02-05 20:53:05 INFO    snapshot check_infracomponent.py The Infrastructure
Component ZooKeeper is Healthy

2019-02-05 20:53:05 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Redis Started
2019-02-05 20:53:05 INFO    snapshot check_infracomponent.py The Infrastructure
Component Redis is Healthy

2019-02-05 20:53:05 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component ArangoDb Started
2019-02-05 20:53:07 INFO    snapshot check_infracomponent.py The Infrastructure
Component ArangoDb is Healthy

2019-02-05 20:53:07 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Sim_Cluster Started
2019-02-05 20:53:07 INFO    snapshot check_infracomponent.py The Infrastructure
Component Sim_Cluster is Healthy

2019-02-05 20:53:07 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Elk_Logstash Started
2019-02-05 20:53:08 INFO    snapshot check_infracomponent.py The Infrastructure
Component Elk_Logstash is Healthy

2019-02-05 20:53:08 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Elk_Kibana Started
2019-02-05 20:53:09 INFO    snapshot check_infracomponent.py The Infrastructure
Component Elk_Kibana is Healthy

2019-02-05 20:53:09 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Keystone Started
2019-02-05 20:53:10 INFO    snapshot check_infracomponent.py The Infrastructure
Component Keystone is Healthy

```

```

2019-02-05 20:53:10 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Swift Started
2019-02-05 20:53:12 INFO    snapshot check_infracomponent.py The Infrastructure
Component Swift is Healthy

2019-02-05 20:53:12 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Kubernetes Started
2019-02-05 20:53:17 INFO    snapshot check_infracomponent.py The Infrastructure
Component Kubernetes is Healthy

2019-02-05 20:53:17 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Contrail_Analytics Started
2019-02-05 20:53:28 INFO    snapshot check_infracomponent.py The Infrastructure
Component Contrail_Analytics is Healthy

2019-02-05 20:53:45 INFO    snapshot snapshot.py CSO health check done
2019-02-05 20:53:45 INFO    snapshot snapshot.py Calling pre-processing steps
before taking snapshot
2019-02-05 20:53:46 INFO    snapshot snapshot.py Going to free cache on the
host server cso-host
2019-02-05 20:53:46 INFO    snapshot  utils.py Testing SSH Credentials of
server cso-host
2019-02-05 20:53:47 INFO    snapshot  utils.py Testing SSH Credentials of
server cso-host
2019-02-05 20:53:47 INFO    snapshot  utils.py Going to shut down VM
csp-central-msvm
2019-02-05 20:53:48 INFO    snapshot  utils.py Shutting down centralmsvm VM
2019-02-05 20:54:18 INFO    snapshot  utils.py centralmsvm VM shut down
successfully
2019-02-05 20:54:52 INFO    snapshot  utils.py Going to shut down VM
csp-central-k8mastervm
2019-02-05 20:54:52 INFO    snapshot  utils.py Shutting down centralk8mastervm
VM
2019-02-05 20:55:12 INFO    snapshot  utils.py centralk8mastervm VM shut
down successfully
2019-02-05 20:55:57 INFO    snapshot  utils.py Going to shut down VM
csp-contrailanalytics-1
2019-02-05 20:55:57 INFO    snapshot  utils.py Shutting down canvm VM
2019-02-05 20:56:17 INFO    snapshot  utils.py canvm VM shut down successfully
2019-02-05 20:58:02 INFO    snapshot  utils.py Going to shut down VM
csp-regional-sblb
2019-02-05 20:58:02 INFO    snapshot  utils.py Shutting down regional-sblb
VM
2019-02-05 20:58:12 INFO    snapshot  utils.py regional-sblb VM shut down
successfully
2019-02-05 21:00:29 INFO    snapshot  utils.py Going to shut down VM
csp-central-infravm
2019-02-05 21:00:29 INFO    snapshot  utils.py Shutting down centralinfravm
VM
2019-02-05 21:00:49 INFO    snapshot  utils.py centralinfravm VM shut down
successfully
2019-02-05 21:01:34 INFO    snapshot  utils.py Successfully shut down all
VMs
2019-02-05 21:01:34 INFO    snapshot snapshot.py Taking snapshot of all the
VMs
2019-02-05 21:01:34 INFO    snapshot  utils.py Testing SSH Credentials of
server cso-host
2019-02-05 21:01:35 INFO    snapshot snapshot.py Preparing to take snapshot
of centralinfravm VM
2019-02-05 21:01:35 INFO    snapshot snapshot.py Took a snapshot of

```

```

centralinfravm-data2 image of centralinfravm VM at 2019-02-05/21:01:35
2019-02-05 21:01:35 INFO      snapshot snapshot.py Took a snapshot of
centralinfravm image of centralinfravm VM at 2019-02-05/21:01:35
2019-02-05 21:01:37 INFO      snapshot snapshot.py Took a snapshot of
centralinfravm-mnt_data image of centralinfravm VM at 2019-02-05/21:01:35
2019-02-05 21:01:37 INFO      snapshot snapshot.py Preparing to take snapshot
of centralmsvm VM
2019-02-05 21:01:37 INFO      snapshot snapshot.py Took a snapshot of
centralmsvm-data2 image of centralmsvm VM at 2019-02-05/21:01:37
2019-02-05 21:01:38 INFO      snapshot snapshot.py Took a snapshot of centralmsvm
image of centralmsvm VM at 2019-02-05/21:01:37
2019-02-05 21:01:39 INFO      snapshot snapshot.py Took a snapshot of
centralmsvm-mnt_data image of centralmsvm VM at 2019-02-05/21:01:37
2019-02-05 21:01:39 INFO      snapshot snapshot.py Preparing to take snapshot
of centraalk8mastervm VM
2019-02-05 21:01:40 INFO      snapshot snapshot.py Took a snapshot of
centraalk8mastervm-data2 image of centraalk8mastervm VM at 2019-02-05/21:01:39
2019-02-05 21:01:40 INFO      snapshot snapshot.py Took a snapshot of
centraalk8mastervm image of centraalk8mastervm VM at 2019-02-05/21:01:39
2019-02-05 21:01:41 INFO      snapshot snapshot.py Took a snapshot of
centraalk8mastervm-mnt_data image of centraalk8mastervm VM at 2019-02-05/21:01:39
2019-02-05 21:01:41 INFO      snapshot snapshot.py Preparing to take snapshot
of canvm VM
2019-02-05 21:01:42 INFO      snapshot snapshot.py Took a snapshot of canvm-data1
image of canvm VM at 2019-02-05/21:01:41
2019-02-05 21:01:42 INFO      snapshot snapshot.py Took a snapshot of canvm-data2
image of canvm VM at 2019-02-05/21:01:41
2019-02-05 21:01:44 INFO      snapshot snapshot.py Took a snapshot of canvm
image of canvm VM at 2019-02-05/21:01:41
2019-02-05 21:01:44 INFO      snapshot snapshot.py Preparing to take snapshot
of regional-sblb VM
2019-02-05 21:01:44 INFO      snapshot snapshot.py Took a snapshot of
regional-sblb-data2 image of regional-sblb VM at 2019-02-05/21:01:44
2019-02-05 21:01:46 INFO      snapshot snapshot.py Took a snapshot of
regional-sblb image of regional-sblb VM at 2019-02-05/21:01:44
2019-02-05 21:01:47 INFO      snapshot snapshot.py Took a snapshot of
regional-sblb-mnt_data image of regional-sblb VM at 2019-02-05/21:01:44
2019-02-05 21:01:47 INFO      snapshot snapshot.py Successfully took snapshot
of all VMS
2019-02-05 21:01:48 INFO      snapshot  utils.py Testing SSH Credentials of
server cso-host
2019-02-05 21:01:48 INFO      snapshot  utils.py Going to start VM
csp-central-infravm
2019-02-05 21:01:48 INFO      snapshot  utils.py Starting centralinfravm VM
2019-02-05 21:01:58 INFO      snapshot  utils.py centralinfravm VM started
successfully
2019-02-05 21:06:28 INFO      snapshot  utils.py Going to start VM
csp-regional-sblb
2019-02-05 21:06:28 INFO      snapshot  utils.py Starting regional-sblb VM
2019-02-05 21:06:38 INFO      snapshot  utils.py regional-sblb VM started
successfully
2019-02-05 21:06:53 INFO      snapshot  utils.py Going to start VM
csp-contrailanalytics-1
2019-02-05 21:06:53 INFO      snapshot  utils.py Starting canvm VM
2019-02-05 21:07:03 INFO      snapshot  utils.py canvm VM started successfully
2019-02-05 21:08:23 INFO      snapshot  utils.py Going to start VM
csp-central-k8mastervm
2019-02-05 21:08:23 INFO      snapshot  utils.py Starting centraalk8mastervm
VM
2019-02-05 21:08:33 INFO      snapshot  utils.py centraalk8mastervm VM started

```

```

successfully
2019-02-05 21:09:33 INFO      snapshot  utils.py Going to start VM
csp-central-msvm
2019-02-05 21:09:33 INFO      snapshot  utils.py Starting centralmsvm VM
2019-02-05 21:09:44 INFO      snapshot  utils.py centralmsvm VM started
successfully
2019-02-05 21:10:38 INFO      snapshot  utils.py Successfully started all VMs
2019-02-05 21:10:40 INFO      snapshot snapshot.py Restarting salt-minion in
csp-central-infravm VM
2019-02-05 21:10:42 INFO      snapshot snapshot.py Restarting salt-minion in
csp-central-msvm VM
2019-02-05 21:10:43 INFO      snapshot snapshot.py Restarting salt-minion in
csp-central-k8mastervm VM
2019-02-05 21:10:44 INFO      snapshot snapshot.py Restarting salt-minion in
csp-contrailanalytics-1 VM
2019-02-05 21:10:44 INFO      snapshot snapshot.py Restarting salt-minion in
csp-regional-sblb VM
2019-02-05 21:10:53 INFO      snapshot snapshot.py Going to sync salt across
all VMs; this will take a while
2019-02-05 21:10:55 INFO      snapshot snapshot.py Going to recover Swift
processes
2019-02-05 21:10:55 INFO      snapshot snapshot.py Killing Swift zombie processes
in central
2019-02-05 21:10:55 INFO      snapshot snapshot.py
2019-02-05 21:10:55 INFO      snapshot snapshot.py Restarting Swift service in
central
2019-02-05 21:11:04 INFO      snapshot snapshot.py swift-object start/running
2019-02-05 21:11:05 INFO      snapshot snapshot.py Checking MariaDB health and
calling MariaDB Recovery, if required
2019-02-05 21:11:06 INFO      snapshot check_infracomponent.py Health Check for
Component mariadb Started
2019-02-05 21:11:07 INFO      snapshot check_infracomponent.py The Components
mariadb is Healthy in central Environment

2019-02-05 21:11:28 INFO      snapshot check_infracomponent.py Health Check for
Component mariadb Started
2019-02-05 21:11:30 INFO      snapshot check_infracomponent.py The Components
mariadb is Healthy in central Environment

2019-02-05 21:11:30 INFO      snapshot snapshot.py Checking Redis health and
calling Redis Recovery, if required
2019-02-05 21:11:31 INFO      snapshot snapshot.py Going to perform CSO Health
Check
2019-02-05 21:11:49 INFO      snapshot check_infracomponent.py Health Check for
Infrastructure Component SaltStack Started
2019-02-05 21:11:50 INFO      snapshot check_infracomponent.py The Infrastructure
Component SaltStack is Healthy

2019-02-05 21:11:50 INFO      snapshot check_infracomponent.py Health Check for
Infrastructure Component Cassandra Started
2019-02-05 21:11:51 INFO      snapshot check_infracomponent.py The Infrastructure
Component Cassandra is Healthy

2019-02-05 21:11:51 INFO      snapshot check_infracomponent.py Health Check for
Infrastructure Component Elasticsearch Started
2019-02-05 21:11:52 INFO      snapshot check_infracomponent.py The Infrastructure
Component Elasticsearch is Healthy

2019-02-05 21:11:52 INFO      snapshot check_infracomponent.py Health Check for
Infrastructure Component Etcd Started

```

```

2019-02-05 21:11:54 INFO    snapshot check_infracomponent.py The Infrastructure
Component Etcd is Healthy

2019-02-05 21:11:54 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component MariaDb Started
2019-02-05 21:11:56 INFO    snapshot check_infracomponent.py The Infrastructure
Component MariaDb is Healthy

2019-02-05 21:11:56 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component RabbitMQ Started
2019-02-05 21:11:59 INFO    snapshot check_infracomponent.py The Infrastructure
Component RabbitMQ is Healthy

2019-02-05 21:11:59 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component ZooKeeper Started
2019-02-05 21:12:00 INFO    snapshot check_infracomponent.py The Infrastructure
Component ZooKeeper is Healthy

2019-02-05 21:12:00 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Redis Started
2019-02-05 21:12:00 INFO    snapshot check_infracomponent.py The Infrastructure
Component Redis is Healthy

2019-02-05 21:12:00 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component ArangoDb Started
2019-02-05 21:12:02 INFO    snapshot check_infracomponent.py The Infrastructure
Component ArangoDb is Healthy

2019-02-05 21:12:02 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Sim_Cluster Started
2019-02-05 21:12:02 INFO    snapshot check_infracomponent.py The Infrastructure
Component Sim_Cluster is Healthy

2019-02-05 21:12:02 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Elk_Logstash Started
2019-02-05 21:12:03 INFO    snapshot check_infracomponent.py The Infrastructure
Component Elk_Logstash is Healthy

2019-02-05 21:12:03 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Elk_Kibana Started
2019-02-05 21:12:04 INFO    snapshot check_infracomponent.py The Infrastructure
Component Elk_Kibana is Healthy

2019-02-05 21:12:04 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Keystone Started
2019-02-05 21:12:05 INFO    snapshot check_infracomponent.py The Infrastructure
Component Keystone is Healthy

2019-02-05 21:12:05 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Swift Started
2019-02-05 21:12:08 INFO    snapshot check_infracomponent.py The Infrastructure
Component Swift is Healthy

2019-02-05 21:12:08 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Kubernetes Started
2019-02-05 21:12:12 INFO    snapshot check_infracomponent.py The Infrastructure
Component Kubernetes is Healthy

2019-02-05 21:12:12 INFO    snapshot check_infracomponent.py Health Check for
Infrastructure Component Contrail_Analytics Started

```

```

2019-02-05 21:12:27 INFO      snapshot check_infracomponent.py Attempt: 1 -
Retrying Health Check for Component Contrail_Analytics
2019-02-05 21:13:22 INFO      snapshot check_infracomponent.py Attempt: 2 -
Retrying Health Check for Component Contrail_Analytics
2019-02-05 21:14:03 INFO      snapshot check_infracomponent.py The Infrastructure
Component Contrail_Analytics is Healthy

2019-02-05 21:14:20 INFO      snapshot snapshot.py CSO health check done
2019-02-05 21:14:20 INFO      snapshot snapshot.py ===== Snapshot
start time: 2019-02-05 20:52:36.694287 =====
2019-02-05 21:14:20 INFO      snapshot snapshot.py ===== Snapshot end
time: 2019-02-05 21:14:20.309453 =====
2019-02-05 21:14:20 INFO      snapshot snapshot.py ===== Snapshot
duration: 0:21:43.615166 =====

```

On running this script, **snapshots/** folder gets created under **CSO** folder. You can check the snapshot names of the VMs in **snapshots/** folder

Sample *snapshot.conf* file:

```

snapshot_status:
  cso_health_status_after_snapshot: success
  cso_health_status_before_snapshot: success
  post_snapshot_processes: success
  vm_backup:
    canvm:
      images:
        canvm: 2019-02-05/21:01:41_canvm_None.snap
        canvm-data1: 2019-02-05/21:01:41_canvm-data1_None.snap
        canvm-data2: 2019-02-05/21:01:41_canvm-data2_None.snap
      status: success
    centralinfravm:
      images:
        centralinfravm: 2019-02-05/21:01:35_centralinfravm_None.snap
        centralinfravm-data2: 2019-02-05/21:01:35_centralinfravm-data2_None.snap

        centralinfravm-mnt_data:
2019-02-05/21:01:35_centralinfravm-mnt_data_None.snap
      status: success
    centra1k8mastervm:
      images:
        centra1k8mastervm: 2019-02-05/21:01:39_centra1k8mastervm_None.snap
        centra1k8mastervm-data2:
2019-02-05/21:01:39_centra1k8mastervm-data2_None.snap
        centra1k8mastervm-mnt_data:
2019-02-05/21:01:39_centra1k8mastervm-mnt_data_None.snap
      status: success
    centralmsvm:
      images:
        centralmsvm: 2019-02-05/21:01:37_centralmsvm_None.snap
        centralmsvm-data2: 2019-02-05/21:01:37_centralmsvm-data2_None.snap
        centralmsvm-mnt_data: 2019-02-05/21:01:37_centralmsvm-mnt_data_None.snap

      status: success
    regional-sblb:
      images:

```

```

    regional-sblb: 2019-02-05/21:01:44_regional-sblb_None.snap
    regional-sblb-data2: 2019-02-05/21:01:44_regional-sblb-data2_None.snap
    regional-sblb-mnt_data:
2019-02-05/21:01:44_regional-sblb-mnt_data_None.snap
    status: success
    vm_snapshot: success

```

Reverting Snapshots

You can restore the parent snapshot of the VM. When you revert to a snapshot, changes made after the snapshot was taken are reverted to the snapshot point.

1. On the installer VM, navigate to **Contrail_Service_Orchestration_4.1.0**
2. Run **./revert.sh --snapshot**

This utility displays a list of all the snapshots taken:

Sample Output:

This utility will display a list of snapshots taken
List of snapshots present in the system:

```

1 : snapshot_2019-02-07_22:42:38
2 : snapshot_2019-02-05_20:52:36

```

Choose a snapshot to revert to: 1

```

2019-02-08 11:34:40 INFO      revert      revert.py
=====
2019-02-08 11:34:40 INFO      revert      revert.py Revert Start Time: 2019-02-08
11:34:40.963127
2019-02-08 11:34:40 INFO      revert      revert.py
=====
2019-02-08 11:34:40 INFO      revert      revert.py Reverting VM snapshots
2019-02-08 11:34:42 INFO      revert      utils.py Testing SSH Credentials of server
cso-host
2019-02-08 11:34:42 INFO      revert      utils.py Going to shut down VM
csp-central-msvm
2019-02-08 11:34:44 INFO      revert      utils.py Shutting down centralmsvm VM
2019-02-08 11:35:14 INFO      revert      utils.py centralmsvm VM shut down
successfully
2019-02-08 11:35:47 INFO      revert      utils.py Going to shut down VM
csp-central-k8mastervm
2019-02-08 11:35:47 INFO      revert      utils.py Shutting down centralk8mastervm
VM
2019-02-08 11:36:07 INFO      revert      utils.py centralk8mastervm VM shut down
successfully
2019-02-08 11:36:52 INFO      revert      utils.py Going to shut down VM
csp-contrailanalytics-1
2019-02-08 11:36:52 INFO      revert      utils.py Shutting down canvm VM
2019-02-08 11:37:02 INFO      revert      utils.py canvm VM shut down successfully
2019-02-08 11:38:57 INFO      revert      utils.py Going to shut down VM
csp-regional-sblb
2019-02-08 11:38:57 INFO      revert      utils.py Shutting down regional-sblb VM
2019-02-08 11:39:07 INFO      revert      utils.py regional-sblb VM shut down

```

```

successfully
2019-02-08 11:41:23 INFO      revert      utils.py Going to shut down VM
csp-central-infravm
2019-02-08 11:41:23 INFO      revert      utils.py Shutting down centralinfravm VM
2019-02-08 11:41:53 INFO      revert      utils.py centralinfravm VM shut down
successfully
2019-02-08 11:42:28 INFO      revert      utils.py Successfully shut down all VMs
2019-02-08 11:42:28 INFO      revert      revert_snapshot.py Reverting snapshots for
all the VMs
2019-02-08 11:42:28 INFO      revert      utils.py Testing SSH Credentials of server
cso-host
2019-02-08 11:42:29 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:37_centralinfravm-data2_None.snap for centralinfravm VM
2019-02-08 11:42:41 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:37_centralinfravm-mnt_data_None.snap for centralinfravm VM
2019-02-08 11:42:41 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:37_centralinfravm_None.snap for centralinfravm VM
2019-02-08 11:42:41 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:39_centralmsvm-data2_None.snap for centralmsvm VM
2019-02-08 11:42:46 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:39_centralmsvm-mnt_data_None.snap for centralmsvm VM
2019-02-08 11:42:46 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:39_centralmsvm_None.snap for centralmsvm VM
2019-02-08 11:42:46 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:41_centralk8mastervm-mnt_data_None.snap for centralk8mastervm
VM
2019-02-08 11:42:47 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:41_centralk8mastervm-data2_None.snap for centralk8mastervm VM
2019-02-08 11:42:47 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:41_centralk8mastervm_None.snap for centralk8mastervm VM
2019-02-08 11:42:47 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:42_canvm-data1_None.snap for canvm VM
2019-02-08 11:42:47 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:42_canvm-data2_None.snap for canvm VM
2019-02-08 11:42:50 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:42_canvm_None.snap for canvm VM
2019-02-08 11:42:50 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:43_regional-sblb-data2_None.snap for regional-sblb VM
2019-02-08 11:42:50 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:43_regional-sblb_None.snap for regional-sblb VM
2019-02-08 11:42:50 INFO      revert      revert_snapshot.py Reverted to snapshot
2019-02-07/22:51:43_regional-sblb-mnt_data_None.snap for regional-sblb VM
2019-02-08 11:42:51 INFO      revert      utils.py Testing SSH Credentials of server
cso-host
2019-02-08 11:42:51 INFO      revert      utils.py Going to start VM
csp-central-infravm
2019-02-08 11:42:51 INFO      revert      utils.py Starting centralinfravm VM
2019-02-08 11:43:01 INFO      revert      utils.py centralinfravm VM started
successfully
2019-02-08 11:47:31 INFO      revert      utils.py Going to start VM
csp-regional-sblb
2019-02-08 11:47:31 INFO      revert      utils.py Starting regional-sblb VM
2019-02-08 11:47:41 INFO      revert      utils.py regional-sblb VM started
successfully
2019-02-08 11:47:56 INFO      revert      utils.py Going to start VM
csp-contrailanalytics-1
2019-02-08 11:47:56 INFO      revert      utils.py Starting canvm VM
2019-02-08 11:48:06 INFO      revert      utils.py canvm VM started successfully
2019-02-08 11:49:26 INFO      revert      utils.py Going to start VM
csp-central-k8mastervm

```



```

2019-02-08 11:49:26 INFO      revert      utils.py Starting centralk8mastervm VM
2019-02-08 11:49:36 INFO      revert      utils.py centralk8mastervm VM started
successfully
2019-02-08 11:50:36 INFO      revert      utils.py Going to start VM csp-central-msvm
2019-02-08 11:50:36 INFO      revert      utils.py Starting centralmsvm VM
2019-02-08 11:50:46 INFO      revert      utils.py centralmsvm VM started
successfully
2019-02-08 11:51:41 INFO      revert      utils.py Successfully started all VMs
2019-02-08 11:51:41 INFO      revert      revert.py Waiting for CSO to boot; this
will take a while
2019-02-08 11:52:43 INFO      revert      revert.py Checking MariaDB health and
calling MariaDB Recovery, if required
2019-02-08 11:53:17 INFO      revert      check_infracomponent.py Health Check for
Component mariadb Started
2019-02-08 11:53:19 INFO      revert      check_infracomponent.py The Components
mariadb is Healthy in central Environment

2019-02-08 11:53:39 INFO      revert      check_infracomponent.py Health Check for
Component mariadb Started
2019-02-08 11:53:41 INFO      revert      check_infracomponent.py The Components
mariadb is Healthy in central Environment

2019-02-08 11:53:41 INFO      revert      revert.py Checking Redis health and calling
Redis Recovery, if required
2019-02-08 11:53:42 INFO      revert      revert.py Going to perform CSO Health Check
2019-02-08 11:54:00 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component SaltStack Started
2019-02-08 11:54:01 INFO      revert      check_infracomponent.py The Infrastructure
Component SaltStack is Healthy

2019-02-08 11:54:01 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component Cassandra Started
2019-02-08 11:54:02 INFO      revert      check_infracomponent.py The Infrastructure
Component Cassandra is Healthy

2019-02-08 11:54:02 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component ElasticSearch Started
2019-02-08 11:54:04 INFO      revert      check_infracomponent.py The Infrastructure
Component ElasticSearch is Healthy

2019-02-08 11:54:04 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component Etcd Started
2019-02-08 11:54:05 INFO      revert      check_infracomponent.py The Infrastructure
Component Etcd is Healthy

2019-02-08 11:54:05 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component MariaDb Started
2019-02-08 11:54:07 INFO      revert      check_infracomponent.py The Infrastructure
Component MariaDb is Healthy

2019-02-08 11:54:07 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component RabbitMQ Started
2019-02-08 11:54:10 INFO      revert      check_infracomponent.py The Infrastructure
Component RabbitMQ is Healthy

2019-02-08 11:54:10 INFO      revert      check_infracomponent.py Health Check for
Infrastructure Component ZooKeeper Started
2019-02-08 11:54:11 INFO      revert      check_infracomponent.py The Infrastructure
Component ZooKeeper is Healthy

```

```

2019-02-08 11:54:11 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Redis Started
2019-02-08 11:54:11 INFO      revert    check_infracomponent.py The Infrastructure
Component Redis is Healthy

2019-02-08 11:54:11 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component ArangoDb Started
2019-02-08 11:54:13 INFO      revert    check_infracomponent.py The Infrastructure
Component ArangoDb is Healthy

2019-02-08 11:54:13 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Sim_Cluster Started
2019-02-08 11:54:13 INFO      revert    check_infracomponent.py The Infrastructure
Component Sim_Cluster is Healthy

2019-02-08 11:54:13 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Elk_Logstash Started
2019-02-08 11:54:14 INFO      revert    check_infracomponent.py The Infrastructure
Component Elk_Logstash is Healthy

2019-02-08 11:54:14 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Elk_Kibana Started
2019-02-08 11:54:15 INFO      revert    check_infracomponent.py The Infrastructure
Component Elk_Kibana is Healthy

2019-02-08 11:54:15 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Keystone Started
2019-02-08 11:54:16 INFO      revert    check_infracomponent.py The Infrastructure
Component Keystone is Healthy

2019-02-08 11:54:16 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Swift Started
2019-02-08 11:54:18 INFO      revert    check_infracomponent.py The Infrastructure
Component Swift is Healthy

2019-02-08 11:54:18 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Kubernetes Started
2019-02-08 11:54:23 INFO      revert    check_infracomponent.py The Infrastructure
Component Kubernetes is Healthy

2019-02-08 11:54:23 INFO      revert    check_infracomponent.py Health Check for
Infrastructure Component Contrail_Analytics Started
2019-02-08 11:54:33 INFO      revert    check_infracomponent.py The Infrastructure
Component Contrail_Analytics is Healthy

2019-02-08 11:54:50 INFO      revert    revert.py Revert process successful
2019-02-08 11:54:50 INFO      revert    revert.py ===== Revert start
time : 2019-02-08 15:22:15.882347 =====
2019-02-08 11:54:50 INFO      revert    revert.py ===== Revert end time
: 2019-02-08 15:23:50.761300 =====
2019-02-08 11:54:50 INFO      revert    revert.py ===== Revert duration
: 0:20:10.878953 =====
2019-02-08 11:54:50 INFO      revert    revert.py
=====
2019-02-08 11:54:50 INFO      revert    revert.py Overall Revert Summary
2019-02-08 11:54:50 INFO      revert    revert.py
=====

```

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:

- SaltStack
- 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. Log in to the installer VM as root.
2. Navigate to the CSO directory in the installer VM.

For example:

```
root@host:~/# cd Contrail_Service_Orchestration_4.1.0
root@host:~/Contrail_Service_Orchestration_4.1.0#
```

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_4.1.0/components_health.sh central
```

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

```
root@host:~/Contrail_Service_Orchestration_4.1.0/components_health.sh regional
```

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

```
root@host:~/Contrail_Service_Orchestration_4.1.0 ./components_health.sh
```

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

For example:

```
*****
HEALTH CHECK FOR INFRASTRUCTURE COMPONENTS STARTED IN CENTRAL ENVIRONMENT
*****

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

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

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:

```
['SaltStack', 'Cassandra', 'ElasticSearch', 'Etcd', 'MariaDb', 'RabbitMQ',
'ZooKeeper', 'Redis', 'ArangoDb', 'Sim_Cluster', 'Elk_Logstash', 'Elk_Kibana',
'Keystone', 'Swift', 'Kubernetes', 'Contrail_Analytics']
HEALTH CHECK NOT REQUIRED FOR COLLOCATED REGION
```

Please check the logs:

Console Logs Path:

```
/root/CSP/Contrail_Service_Orchestration_4.1.0/logs/health_check_console.log
```

Error Logs Path:

```
/root/CSP/Contrail_Service_Orchestration_4.1.0/logs/health_check_error.log
```

Info Logs Path:

```
/root/CSP/Contrail_Service_Orchestration_4.1.0/logs/health_check.log
```

- Related Documentation**
- [Provisioning VMs on Contrail Service Orchestration Nodes or Servers on page 44](#)
 - [Generating and Encrypting Passwords for Infrastructure Components on page 104](#)

CHAPTER 5

Post Installation Tasks

- [Generating and Encrypting Passwords for Infrastructure Components on page 104](#)
- [Applying NAT Rules if CSO Is Deployed Behind NAT on page 105](#)
- [Applying Security Patches on page 107](#)
- [Viewing Information About Microservices on page 107](#)

Generating and Encrypting Passwords for Infrastructure Components

CSO uses an algorithm to automatically generate a dynamic password for the following infrastructure components:

- Cassandra.
- Keystone.
- MariaDB.
- RabbitMQ.
- Icinga.
- Prometheus.
- ArangoDB.
- Elasticsearch.
- ZooKeeper.

The automatically 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.

You can access the Administration portal by opening a browser with central-msvm IP address. The initial user name is **cspadmin** and the initial password is the password shown after running the `./setup_assist.sh` script under [“Installing and Configuring Contrail Service Orchestration” on page 78](#) topic.



NOTE: You must note the automatically 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.

Run the following script to retrieve password for all infrastructure components:

```
CSO>./python.sh deploy_manager/utils/decrypt_password.py
```

```
UI Password for cspadmin user: P72t9hxDVC?g
Keystone Password for admin user: piIL11XadDCFw0z
arangodb root user password: pTSWtfI8UXGmCDIV
Cassandra DB password for cassandra user: S9rVeyLcQLNgeL
Cassandra DB password for casadmin user: iz2Tj9BYM05o4XA0
Elasticsearch admin Password: A60o5f95I4hE1
Keystone Password for swift user: g6J2hPFWCfBo
Mariadb Password for root/admin user: fsHE0Z99LV9ec5ur
Mariadb Password for clusteruser user: 9KpP7oWTczYF0p
Password for prometheus user: UItspGCHJyofDL
Rabbitmq cspmq Password: VE5v59zCsQs7
```


Icinga UI Password: urgNf7nV7YGVzgIs
 Password for Zookeeper admin user: NzKiWbjoHwG



CAUTION: You can't retrieve the *cspadmin* user password. It can only be reset from the CSO Installer webpage or CLI.

Follow the steps to reset the *cspadmin* user password from the CSO Installer webpage:

1. Click *Forget Password?* on the login page of the CSO Installer webpage.
2. Receive a verification code to the registered e-mail ID.
3. Type the received verification code as the password on the CSO Installer webpage.
4. Follow the steps to reset the password.

Related Documentation

- [Installing and Configuring Contrail Service Orchestration on page 78](#)

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_assist.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 the Administration Portal, the Kibana UI, and the Rabbit MQ console.

The following procedure is only applicable in case of a NAT server. Apply similar rules if you have any external Juniper or third party firewall.

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
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 78](#)

Applying Security Patches

You can apply in-service patches to CSO microservices without having to reboot the VMs.

This feature is applicable only to microservices and is not supported for infrastructure components like Cassandra, RabbitMQ, kernel etc. The process does not impact sites or CSO workflows.

You can always revert the applied patch in case of unsuccessful execution.

You can follow the following steps to apply security patches:

- Download the tar file that contains the hotfix.
- Run patch script - **patch.sh** to apply security patches.

The script is bundled in the tar file which needs to be executed on the installer VM.

The script performs in-service patching of CSO microservices.

**Related
Documentation**

- [Contrail Service Orchestration Monitoring and Troubleshooting Guide](#)
- [Viewing Information About Microservices on page 107](#)

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.

[Table 15 on page 107](#) provides basic functions of each microservice. The list is limited to some of the external facing microservices.

Table 15: Functions of Microservices

Microservice	Description
Activation Service (Central)	Provides network activation functions to enable zero touch provisioning of devices.

Table 15: Functions of Microservices (continued)

Microservice	Description
ams	Monitors and autonomously collects data without system or human intervention.
cslm	Maintains EMS device data model for device management functions. The data model contains information like device objects, abstract configuration, device inventory object, configuration template object, device profile object, device image object etc.
Configuration Template Service	Provides configuration template management features for the CSO solution. The features include maintaining a database of config templates, template syntax validation (e.g. jinja2, python, yang rpc), template execution with input parameters using Yang RPC, and input/output validation (provided corresponding schema is given).
Device Management Service (Central)	<ul style="list-style-type: none"> Manages the lifecycle of device objects. Each device object provides an abstraction for one or more physical or virtual network devices. Provides APIs for device management.
Dataview Service (Central)	Serves the northbound applications such as portals or OSS systems, read-only data with paging, sorting and rich queries.
design-tools-central	Provides interface to Network Function Virtualization Design Tools to create config templates, VNF definitions and network service definitions.
Element Management Service (Central)	Maintains EMS device data model for device management functions. This data model contains device objects, abstract config, device inventory object, config template object, device profile object, device image object etc.
Fault and Performance Monitoring (FMPM) Collector Services	Describes APIs used by fault monitoring and performance monitoring system for collecting service check results from telemetry agents.
IAM Service	Provides identity and access management features.
IAM Service (No Authentication)	Provides identity and access management features during password recovery procedures.
Image Management Service (Central)	Provides image management functions.

Table 15: Functions of Microservices (continued)

Microservice	Description
Inventory Management Service (Central)	Provides generic inventory management functions.
Job Service	<ul style="list-style-type: none"> Provides job management functionality. Supports creation of synchronous and asynchronous jobs, track status, rack start and completion time.
Intent based Policy Management	Provides Policy and SLA profile object management service to enable software-defined WAN (SD-WAN) functions.
Policy and SLA management Service	Enables software-defined WAN (SDWAN) function.
Routing Manager Service	Provides APIs to manage routing operations such as to create VPN, interface to route-reflector, enable routing on CPE locations.
Schema Service	<ul style="list-style-type: none"> Provides highly available, persistent data store for various schemas used by CSP applications. Provides APIs to create, read, update, and delete schemas.
Shared Object Service	Varies based on type of schema.
Signature Manager Service	Manages application signatures
Template Service	<p>Provides config template of CSO management feature.</p> <p>This feature maintains database of config templates, template syntax validation (e.g - jinja2, python, yang rpc), template execution with input parameters using Yang RPC, and input/output validation ((provided corresponding schema is given).</p>
Topology Service	Provides API for modeling topologies and working with network elements like devices, hubs, spokes, policy enforcement points and other objects.
Tenant, Site and Service Manager Service	Provides APIs for tenant, site and service management.

Table 15: Functions of Microservices (continued)

Microservice	Description
VIM	Provides common APIs to create virtual networks, virtual links, instantiate VNFs, instantiate service chains for various virtual network infrastructures.

**Related
Documentation**

- *Contrail Service Orchestration Monitoring and Troubleshooting Guide*

Upgrading Contrail Service Orchestration

- [Upgrading Contrail Service Orchestration Overview on page 111](#)
- [Upgrading Contrail Service Orchestration on page 113](#)
- [Adding Virtual Route Reflectors \(VRRs\) After Upgrading on page 116](#)
- [Troubleshooting Upgrade Related Errors on page 118](#)

Upgrading Contrail Service Orchestration Overview

If your installed version is Contrail Service Orchestration (CSO) Release 3.2.1 or later, you can use a script to upgrade to CSO Release 4.1.0.



CAUTION: The upgrade process supports $n-1$ release upgrade approach. For example, if you have CSO Release 3.3 and you intend to upgrade it to CSO Release 4.1.0, you must first upgrade to CSO Release 4.0.2 and, then follow the same steps to upgrade to CSO Release 4.1.0.

We recommend that you take snapshots of your current VMs before you proceed with the upgrade process.

You can roll back to the previous CSO release if the upgrade is unsuccessful, provided you have taken snapshots of the VMs.

To upgrade to CSO Release 4.1.0, you must run the scripts that are available in the **Contrail_Service_Orchestration_4.1.0.tar.gz** file in the following order:

1. **upgrade.sh**— This script upgrades CSO software to Release 4.1.0. The **upgrade.sh** script, puts CSO in maintenance mode, takes a snapshot of all the 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 the VMs, infracomponents, and microservices are up and running, and then puts the CSO in live mode.



NOTE: Before you upgrade ensure that all ongoing jobs are stopped; otherwise, the upgrade process will fail. 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.

Upgrade to CSO Release 4.1.0 is independent of the deployment type (HA and non- HA), environment type (small, medium or large), 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.
- 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



NOTE: The upgrade process is applicable to CSO Release 3.2.1 and later.

Upgrade to CSO Release 4.1.0 has the following limitations:

- There are no changes to the device image or device configurations.
- Upgrade cannot be performed through the GUI.
- CSO Release 4.1.0 allows to separate OS and Data partitions for CSO deployments. This feature is not applicable for upgrades.

Impact of the CSO Upgrade

Table 16 on page 112 describes the impact of the CSO upgrade to Release 4.1.0.

Table 16: Impact of the CSO Upgrade

Feature	After the Upgrade
Security Management	<ul style="list-style-type: none"> • Release 4.1.0 security management-related features are supported on devices that are onboarded in Release 4.0.x or Release 3.3.x.

Table 16: Impact of the CSO Upgrade (continued)

Feature	After the Upgrade
SD-WAN	<ul style="list-style-type: none"> 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 4.1.0. For more information on AppQoE, see <i>Application Quality of Experience (AppQoE) Overview</i> in the <i>Contrail Service Orchestration User Guide</i>. Device Management functions work for Release 3.3.1 or later sites.
Cloud CPE	<ul style="list-style-type: none"> All functionalities of centralized and distributed deployments continues to work on Release 4.1.0 sites or devices that are onboarded in Release 3.3.1 or later. Multi-region support for centralized deployment is not supported on Release 4.1.0 sites or devices that are onboarded in Release 3.3.1. Device Management functions work for Release 4.1.0 sites. High availability (HA) for VRRs is not supported for sites that are created in Release 3.2.1.

Related Documentation • [Upgrading Contrail Service Orchestration on page 113](#)

Upgrading Contrail Service Orchestration

The upgrade process supports $n-1$ release upgrade approach. For example, if you have CSO Release 3.3 and you intend to upgrade it to CSO Release 4.1.0, you must first upgrade to CSO Release 4.0.x and, then follow the same steps to upgrade to CSO Release 4.1.0.

If your CSO release version is lower than 3.2.1, you have to opt for fresh installation.

We recommend that you take a snapshot of your current configuration and VMs before you proceed with upgrade process.



NOTE: The upgrade process is applicable to CSO Release 3.2.1 and later.

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 or later.
- 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_4.1.0/confs/** directory.

To upgrade to CSO Release 4.1.0:

1. Download the CSO Release 4.1.0 installer package from the [CSO Download](#) page to the local drive.
2. Log in 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_4.1.0.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_4.1.0.tar.gz,

```
root@host:~/# tar -xvzf Contrail_Service_Orchestration_4.1.0.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 4.1.0 directory in the installer VM.

```
root@host:~/# cd Contrail_Service_Orchestration_4.1.0
root@host:~/Contrail_Service_Orchestration_4.1.0#
```

(Optional) You can view the list of files in the Contrail_Service_Orchestration_4.1.0.

```
root@host:~/Contrail_Service_Orchestration_4.1.0# ls
```

The Contrail_Service_Orchestration_4.1.0.tar.gz file includes the following scripts:

- **upgrade.sh**
- **revert.sh**

6. Run the **upgrade.sh** script.



WARNING: Before you upgrade ensure that all ongoing jobs in Administration Portal and Customer Portal are stopped; otherwise, the upgrade process will fail. During the upgrade, you experience a downtime as CSO goes into maintenance mode.

This script upgrades CSO software to Release 4.1.0. 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 whether all VMs, infrastructure components, and microservices are up and running; and puts the CSO in live mode. The **upgrade.sh** script takes a snapshot of all VMs by default.



NOTE: The script does not take a snapshot of the Installer VM and Virtual Route Reflector (VRR) VM.

```
root@host:~/Contrail_Service_Orchestration_4.1.0# ./upgrade.sh
```

```
INFO      Configuration Upgrade : success
INFO      System Health Check : success
INFO      CSO Health-Check Before Upgrade : success
INFO      CSO Maintenance Mode Enabled : success
INFO      Kernel Upgrade : NA
INFO      VM Snapshot : success
INFO      Selective Infra Components Upgrade : NA
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      =====
INFO      CSO is successfully upgraded to Release
Contrail_Service_Orchestration_4.1.0
INFO      =====
```



NOTE: The password for each infrastructure component is 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, 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. All VMs on KVM are shut down while a snapshot is being taken. All VMs on VMware ESXi are up and running while a snapshot is being taken.

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_4.1.0/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.



NOTE: You can roll back to the previous release only if you have taken snapshot of VMs before starting the upgrade process.

```
root@host:~/Contrail_Service_Orchestration_4.1.0# ./revert.sh
```

```
INFO      revert      revert.py      Overall Revert Summary
INFO      revert      revert.py      =====
INFO      revert      revert.py      Revert VM Snapshots : success
INFO      revert      revert.py      Revert Kernel Upgrade : NA
INFO      revert      revert.py      Revert Salt Master Configurations : success
INFO      revert      revert.py      Post-revert processes : success
INFO      revert      revert.py      CSO Health-Check after Revert : success
INFO      revert      revert.py      Start Kubernetes pods : success
INFO      revert      revert.py      Enable CSO Services : success
INFO      revert      revert.py
=====
INFO      revert      revert.py      CSO successfully reverted to the previously
INFO      revert      revert.py      installed version.
INFO      revert      revert.py
=====
```

After a successful upgrade, CSO is functional and you can log in to Administrator Portal and Customer Portal.



NOTE: After successfully upgrading the CSO software, ensure that you download the application signatures before installing signatures on the device. This is a one-time operation after the upgrade.

Related Documentation

- [Upgrading Contrail Service Orchestration Overview on page 111](#)

Adding Virtual Route Reflectors (VRRs) After Upgrading

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 4.1.0.

To add VRRs:

1. Log in to the installer VM as root.
2. Navigate to the CSO Release 4.1.0 directory in the installer VM.

```
root@host:~/# cd Contrail_Service_Orchestration_4.1.0
root@host:~/Contrail_Service_Orchestration_4.1.0#
```

3. Run the **add_vrr.sh** script.

```
root@host:~/Contrail_Service_Orchestration_4.1.0# ./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 with CSO Release 3.2.1 or later 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 want to 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, we recommend that you create VRRs in even numbers, and you must create 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 username 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 on page 113](#)

Troubleshooting Upgrade Related Errors

This topic describes the possible errors that you might encounter while you are upgrading Contrail Service Orchestrator (CSO).

It also suggests how to resolve those errors.

- [Salt Synchronization Error on page 118](#)
- [Cache Clearance Error on page 119](#)
- [Kube-system Pod Error on page 120](#)
- [Kubernetes Node Error on page 121](#)

Salt Synchronization Error

Problem **Description:** The upgrade or revert status is displayed as **Going to sync salt...** for a considerable time while upgrading CSO to CSO Release 4.1.0 or reverting to the previously installed release.

The Salt Master on the installer VM might be unable to reach all Salt Minions on the other VMs and a salt timeout exception might occur.

Solution Based on the output of the **salt --timeout=30 '*' test.ping** command, you must restart either the Salt Master or the Salt Minion.

To resolve the error:

1. Open another instance of installer VM.
2. Run the **salt --timeout=30 '*' test.ping** command, to check whether the Salt Master on the installer VM is able to reach other VMs.

```
root@host:~/# salt --timeout=30 '*' test.ping
```

- Restart the Salt Master if the following error occurs:

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, review the output.

```
root@host:~/# salt --timeout=30 '*' test.ping
```

```
csp-regional-sblb.DB7RFF.regional:
True
csp-contrailanalytics-1.8V102D.central:
True
csp-central-msvm.8V102D.central:
True
csp-regional-k8mastervm.DB7RFF.regional:
True
csp-central-infravm.8V102D.central:
False
csp-regional-msvm.DB7RFF.regional:
False
csp-regional-infravm.DB7RFF.regional:
True
csp-central-k8mastervm.8V102D.central:
True
```

You must log in to the VM and restart the Salt Minion whether the status of the VM is False.

```
root@host:~/csp-central-infravm.8V102D.central# service salt-minion restart
```

- Rerun the **salt --timeout=30 '*' test.ping** command to verify whether the status for all VMs is True.

Cache Clearance Error

Problem **Description:** The following error might occur while upgrading CSO to CSO Release 4.1.0:
Could not free cache on host server *ServerName*

Solution You must clear the cache on the host server.

To resolve the error:

- Log in to the host server through SSH.
- To clear the cache, run the following command:

```
root@host:~/Contrail_Service_Orchestration_4.1.0# 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:** The following error might occur while upgrading to CSO Release 4.1.0:
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. Run the following command to view the status of the *kube-system* pod:

```
root@host:~/# kubectl get pods -namespace=kube-system
```

The following output is displayed:

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. Run the following command to restart *kube-proxy*

```
root@host:~/# Kubectl apply -f /etc/kubernetes/manifests/kube-proxy.yaml
```

The *kube-system* pod-related error is resolved.

Kubernetes Node Error

Problem **Description:** The following error might occur while upgrading CSO to CSO Release 4.1.0:
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. Restart the node if the status is *Not Ready* by logging in to the node through SSH and running the following command:

```
root@host:~/# service kubelet restart
```

4. Rerun the following command to check the status of the node that you restarted.

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

The Kubernetes node-related error is resolved.

Related Documentation

- [Upgrading Contrail Service Orchestration on page 113](#)

