

Contrail Service Orchestration Monitoring and Troubleshooting Guide

Published
2021-05-20

Release
5.1.2

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 Monitoring and Troubleshooting Guide
5.1.2

Copyright © 2021 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 | vii

Documentation and Release Notes | vii

Documentation Conventions | vii

Documentation Feedback | x

Requesting Technical Support | x

Self-Help Online Tools and Resources | xi

Creating a Service Request with JTAC | xi

1

Monitoring Contrail Service Orchestration

Monitoring Infrastructure Services and Microservices for On-premise Deployments | 2

Monitoring and Troubleshooting Overview | 2

Service and Infrastructure Monitor | 2

Kibana | 3

Accessing Kibana | 3

Setting Up the Visual Presentation of Microservice Log Files | 4

Viewing Information About Microservices | 5

Filtering Data in Kibana | 6

Troubleshooting Microservices | 6

Analyzing Performance | 6

Performing a Health Check of Infrastructure Components | 8

Backup and Restore of Contrail Service Orchestration for On-premise Deployments | 11

Backup and Restore of Contrail Service Orchestration (CSO) Databases | 11

CSO Database Backup and Restore | 11

Configuration | 13

Major Components | 13

Operations | 14

Command Usage | 15

Backup and Restore Examples | 15

Introduction to Service and Infrastructure Monitor Application for On-premise Deployments | 19

Service and Infrastructure Monitor Overview | 19

Accessing the Service and Infrastructure Monitor GUI | 20

Monitoring Network Services | 21

Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 22

Monitoring Microservices | 26

Monitoring Microservices and Their Host VMs | 27

Monitoring Physical Servers | 29

Troubleshooting Contrail Service Orchestration Issues

Troubleshooting Login Issues for Cloud-based Deployments | 32

Troubleshooting Login Issues | 32

Administration Portal IP Address Is Not Reachable | 32

Troubleshooting POPs, Tenants, and Devices Issues for On-premise Deployments | 35

Troubleshooting POPs, Tenants, and Devices Issues | 35

Failure While Creating a Hub, Site, or Tenant | 35

Base Configuration for CPE Activation | 36

Identifying Connectivity Issues for Cloud-based Deployments | 38

Identifying Connectivity Issues by Using Ping | 38

Identifying Connectivity Issues by Using Traceroute | 42

Troubleshooting Site Activation Issues for Cloud-based Deployments | 45

Troubleshooting Site Activation Issues | 45

Prerequisites to Activate a Site | 45

Site activation process is stuck in device detected state | 46

Site activation process is stuck in bootstrap state | 46

Site activation process failed in bootstrap state | 46

Site activation process failed during provisioning | 46

Troubleshooting Image, License, and Policy Deployment Issues for Cloud-based Deployments | 48

Troubleshooting Image, License, and Policy Deployment Issues | 48

- Unable to find device image version | 48
- Upgrade device image using J-Web | 49
- Unable to connect to the device | 49
- Device image version is different from the recommended version | 50
- LAN Switch image version is different from the recommended version | 50
- Policy deployment failed | 50
- No data for next-generation firewall site | 51
- No data for SD-WAN site | 51
- Traffic from Spoke Sites Are Dropped or Are Not Reaching Internet or Destination | 51
- SLA Violation-Original Link Recovered After SLA Violation | 52
- All WAN links are Up But Not All Links Are Utilized | 52

Troubleshooting CSO Installation Issues for On-premise Deployments | 53

Troubleshooting CSO Installation Issues | 53

- Salt Key Issue During CSO Installation | 53
- TimeZone Error | 55
- SSL Handshake Failure | 55
- Missing Interface on CSO VM | 56

Troubleshooting SMTP Issues for Cloud-based Deployments | 57

Troubleshooting SMTP Issues | 57

- Basic Configuration for SMTP Server | 57

Troubleshooting RBAC and OpCo Issues for On-premise Deployments | 60

Troubleshooting RBAC and OpCo Issues | 60

- Authentication Failed for the SP User, Tenant User, or OpCo User | 60
- Authorization Failed for the SP User, Tenant User, or OpCo User | 62
- Password to Onboard OpCo is Not Received or has Expired | 63

Troubleshooting Site, Device and Link Issues for On-premise Deployments | 65

Troubleshooting Site, Device, and Link Issues | 65

- Secure OAM Activation Failure | 65
- Configure SD-WAN Site Failure | 66
- Device Activation Failure | 66
- Dual-CPE Activation Failure for NFX Series Devices | 67
- Dual-CPE Activation Failure for SRX Series Devices | 67
- Link Switch Event or Performance Metrics is Not Displayed | 68
- WAN Link Performance Parameters are Not Displayed | 68
- LTE Interface Issues | 69

Recovering an Installation | 71

CSO Disaster Recovery | 71

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | vii
- Documentation Conventions | vii
- Documentation Feedback | x
- Requesting Technical Support | x

Use this guide to monitor CSO infrastructure services and microservices and troubleshoot CSO installation, login, site activation, license, and deployment-related issues.

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 viii](#) 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 viii 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>

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<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>
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.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
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 Juniper Care or Partner Support Services 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/>.

1

PART

Monitoring Contrail Service Orchestration

Monitoring Infrastructure Services and Microservices for On-premise
Deployments | **2**

Backup and Restore of Contrail Service Orchestration for On-premise
Deployments | **11**

Introduction to Service and Infrastructure Monitor Application for On-premise
Deployments | **19**

Monitoring Infrastructure Services and Microservices for On-premise Deployments

IN THIS CHAPTER

- [Monitoring and Troubleshooting Overview | 2](#)
- [Accessing Kibana | 3](#)
- [Setting Up the Visual Presentation of Microservice Log Files | 4](#)
- [Viewing Information About Microservices | 5](#)
- [Performing a Health Check of Infrastructure Components | 8](#)

Monitoring and Troubleshooting Overview

You use open-source applications for monitoring and troubleshooting infrastructure services and microservices in Contrail Service Orchestration (CSO). These applications offer a visual representation of the metrics in Contrail Service Orchestration with extensive capabilities for analyzing data and monitoring alerts. The applications used by CSO are listed below:

Service and Infrastructure Monitor

Service and Infrastructure Monitor provides a continuous and comprehensive monitoring of Contrail Service Orchestration. The application provides both a visual display of the state of the deployment and the ability to view detailed event messages.

Service and Infrastructure Monitor tracks the status of:

- Network services
- Virtualized network functions
- Microservices
- Virtual machines
- Physical servers

Kibana

The Kibana application provides a visual representation of log files. You use Kibana to view and analyze log files. You can use it to monitor:

- Network services in a central or regional POP
- Microservices in the deployment

RELATED DOCUMENTATION

[Accessing the Service and Infrastructure Monitor GUI | 20](#)

[Accessing Kibana | 3](#)

Accessing Kibana

You must log in to Kibana GUI by using Elasticsearch credentials. During CSO installation, when you run the `setup_assist.sh` script, CSO automatically generates dynamic password for all infrastructure components and displays the password on the console. You must note the passwords that are displayed on the console as they are not saved in the system.

NOTE: If you have lost or forgotten the password, you can contact the Juniper Networks Technical Assistance Center (JTAC) to obtain the new password.

To access the GUI for Kibana:

1. Using a web browser, access the URL for Kibana:

`http://ha-proxy-IP-Address:5601`

where:

ha-proxy-IP-Address—IP address of high availability (HA) proxy. Use this option to monitor the microservices.

- For a deployment without HA, use the IP address of the VM that hosts the microservices for the central POP.
- For an HA deployment, use the virtual IP address of the central or regional POP that you provided for the HA proxy when you installed CSO.

For example:

`http://192.0.2.2:5601`

2. Enter the username **admin** and the Elasticsearch password that is generated during CSO installation.

RELATED DOCUMENTATION

[Setting Up the Visual Presentation of Microservice Log Files | 4](#)

Retrieve Passwords for Infrastructure Components

Install and Configure Contrail Service Orchestration

Setting Up the Visual Presentation of Microservice Log Files

Contrail Service Orchestration includes Kibana and Logstash to view logged data for microservices in a visual format.

To set up logging in Kibana:

1. Log in to Kibana.
2. Select **Settings > Indices**.
3. Click **Create**.

This action creates the **csplogs** index file.

4. Log in as root to the installer host and access the installer directory.
5. Copy the **deploy_manager/export.json** file to a location from which you can import it to the Kibana GUI.

NOTE: Do not change the format of the JSON file. The file must have the correct format to enable visualization of the logs.

6. In the Kibana GUI, select **Settings > Objects**.
7. Click **Import**.

8. Navigate to the location of the **export.json** file that you made available in Step 5.
9. Click **Open**.
10. Confirm overwriting of any existing data.
11. Refresh the Kibana page.
12. Access the dashboard to view the logs in a visual format.

Logs appear after an end user activates a network service.

Refer to the Kibana documentation for information about viewing files in a visual format.

RELATED DOCUMENTATION

Contrail Services Orchestration (CSO) GUIs

[Monitoring and Troubleshooting Overview | 2](#)

[Viewing Information About Microservices | 5](#)

Viewing Information About Microservices

IN THIS SECTION

- [Filtering Data in Kibana | 6](#)
- [Troubleshooting Microservices | 6](#)
- [Analyzing Performance | 6](#)

When you log into Kibana, you see the Discover page, which displays a chart of the number of logs for a specific time period and a list of events for the deployment. You can filter this data to view subsets of logs and add fields to the table to find the specific information that you need. You can also change the time period for which you view events.

Filtering Data in Kibana

To filter data in Kibana:

1. Specify a high-level query in the search field to view a subset of the logs.

You can use keywords from the list of fields in the navigation bar, and specific values for parameters that you configure in Contrail Service Orchestration (CSO), such as a specific tenant name, SD-WAN policy name, job ID, job name, or a specific network service.

For example, specify the following query to view logs concerning timestamp **May 24th 2018** for the tenant name **default-tenant**.

`_exists_: May 24th 2018 AND default-tenant`

2. Select one or more fields from the left navigation bar.

For example, select **message** to show details about the message for the customer.

Troubleshooting Microservices

You can use the troubleshooting dashboard to investigate issues for the microservices.

To use the troubleshooting dashboard:

1. From the Kibana GUI, select **Dashboard > Troubleshooting**.

If the troubleshooting dashboard is not available, click the plus(+) icon in the menu bar to add a visualization. Enter **Troubleshooting** in the search bar.

The troubleshooting dashboard appears, displaying the following predefined monitoring applications:

- Log Level Vs Count

This widget shows the number of logs for each alert level.

- Status Code Vs Count

This widget shows the number of logs for each HTTP status code.

- Service App Name Vs Status Code

This widget shows a visual representation of the number of logs for each microservice analyzed by HTTP status code.

2. Click on an option, such as an alert level, in a widget to filter the data and drill down to a specific issue.

Analyzing Performance

You can use the troubleshooting dashboard to investigate issues for the microservices.

To use the troubleshooting dashboard:

1. From the Kibana GUI, select **Dashboard > Performance Analysis**.

If the performance analysis dashboard is not available, click the plus(+) icon in the menu bar to add a visualization. Enter **Performance Analysis** in the search bar.

The Performance Analysis dashboard appears, displaying the following predefined monitoring applications:

- API Vs Min/Average/Max Elapsed time

This widget shows how long an API associated with a microservice has been in use. You can view minimum, maximum, or average durations.

- Request ID Vs Timestamp

This widget shows when an API was called.

- API Vs Count

This widget shows the number of times an API has been called.

- Application Vs API

This widget shows the level of microservice use analyzed by the type of API call.

- Request ID Vs Application Vs API

This widget provides an analysis of requests by API or microservice.

2. Click on an option, such as a request identifier, in a widget to filter the data and drill down to a specific issue.

RELATED DOCUMENTATION

[Monitoring and Troubleshooting Overview | 2](#)

[Setting Up the Visual Presentation of Microservice Log Files | 4](#)

Performing a Health Check of Infrastructure Components

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

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

To check the status of infrastructure components:

1. Login to the installer VM as root.
2. Navigate to the CSO directory in the installer VM.

For example:

```
root@host:~/# cd Contrail_Service_Orchestration_4.0
```

```
root@host:~/Contrail_Service_Orchestration_4.0#
```

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_4.0#./components_health.sh central
```

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

```
root@host:~/Contrail_Service_Orchestration_4.0#./components_health.sh regional
```

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

```
root@host:~/Contrail_Service_Orchestration_4.0# ./components_health.sh
```

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

For example:

```
*****

HEALTH CHECK FOR INFRASTRUCTURE COMPONENTS STARTED IN CENTRAL ENVIRONMENT

*****

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

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

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

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

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

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

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

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

INFO      Health Check for Infrastructure Component Sim_Cluster Started
```

```

INFO      The Infrastructure Component Sim_Cluster is Healthy

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

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

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

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

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

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

```

Overall result:

The following Infrastructure Components are Healthy:

```

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

```

Backup and Restore of Contrail Service Orchestration for On-premise Deployments

IN THIS CHAPTER

- [Backup and Restore of Contrail Service Orchestration \(CSO\) Databases | 11](#)

Backup and Restore of Contrail Service Orchestration (CSO) Databases

IN THIS SECTION

- [CSO Database Backup and Restore | 11](#)
- [Configuration | 13](#)
- [Major Components | 13](#)
- [Operations | 14](#)
- [Command Usage | 15](#)
- [Backup and Restore Examples | 15](#)

This document introduces the backup and restore capabilities available in Contrail Service Orchestration (CSO). It provides an overview of the concepts, command options, and some examples of how to manage these functions. Although CSO is a GUI-based application, the backup and restore operations can only be managed from the CLI of the startup virtual machine (startup-vm). See the [“Operations” on page 14](#) for details. XYZ

CSO Database Backup and Restore

The Contrail Service Orchestration (CSO) architecture is made up of several virtual machines, each handling pieces of the workload required to make CSO function. These virtual machines store and access their working data in various databases. In order for CSO to function properly, all of the running databases must

be functioning properly. Backup and restore of this critical data is key to ensuring that your CSO installation is running at its best. Starting in CSO 4.1, full backup of all platform, op-co, tenant, and customer data can be run manually or periodically and that data can be restored from the backups when and if the need arises.

Figure 1: Backup and Restore Concept

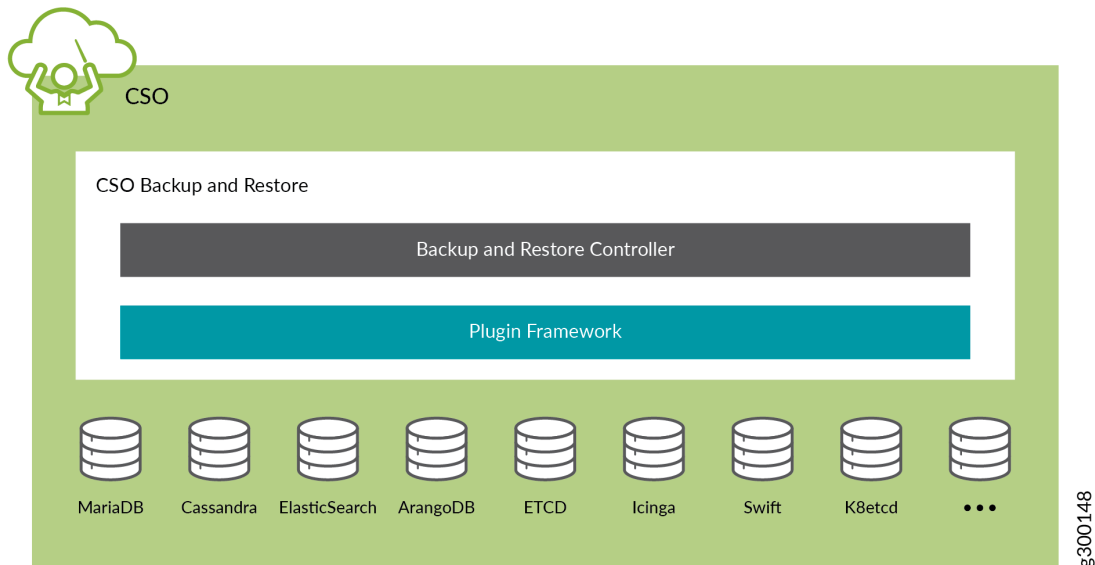


Figure 1 on page 12 shows a conceptual image of how backup and restore is implemented in CSO 4.1. The database systems that are currently backed up within the framework are: MariaDB, Cassandra, ElasticSearch, ArangoDB, Zookeeper, ETCD, Icinga, Swift, and K8ETCD. You can backup and restore data for Contrail Analytics. The system also backs up encrypted passwords, and system certificates so that restoring data from any specific backup puts CSO back into the state it was in at the time of that backup.

You can selectively restore ETCD keys. Before you restore a specific ETCD key, you must add the specific ETCD key in the `/usr/local/etc/etcd_keys.json` file.

For example,

```
"etcd":["/saml", "/routingmanager/sdwan/asnumber", "/natregions/regional",
"/csp/infra/swift/naturl"]
```

Any changes made between the last backup and the current restoration are lost. Generally, backups are made on a system-wide basis meaning that individual op-co or tenant data cannot be backed up or restored apart from the rest of the system data.

NOTE: While it is possible to backup and restore individual databases, there are risks when doing this since the restored database might not be able to fully synch with the current states of the existing databases. This is especially true if there is a long period of time between the backup and restore operations.

The backup and restore operations work with high-availability (HA). This document describes the configuration, scheduling, and operation of backup and restore procedures in CSO.

Configuration

Backup and restore are critical tasks that touch every data storage system used by CSO. Juniper relieves you of the burden of configuring backup details by automatically setting up everything needed to backup and restore CSO during the installation process. No configuration is needed.

Major Components

Although there is no major interaction between the user and the underlying components that make up the backup framework, it is helpful to know the functions that each of the components perform.

[Table 3 on page 13](#) lists the major components and a brief description of each.

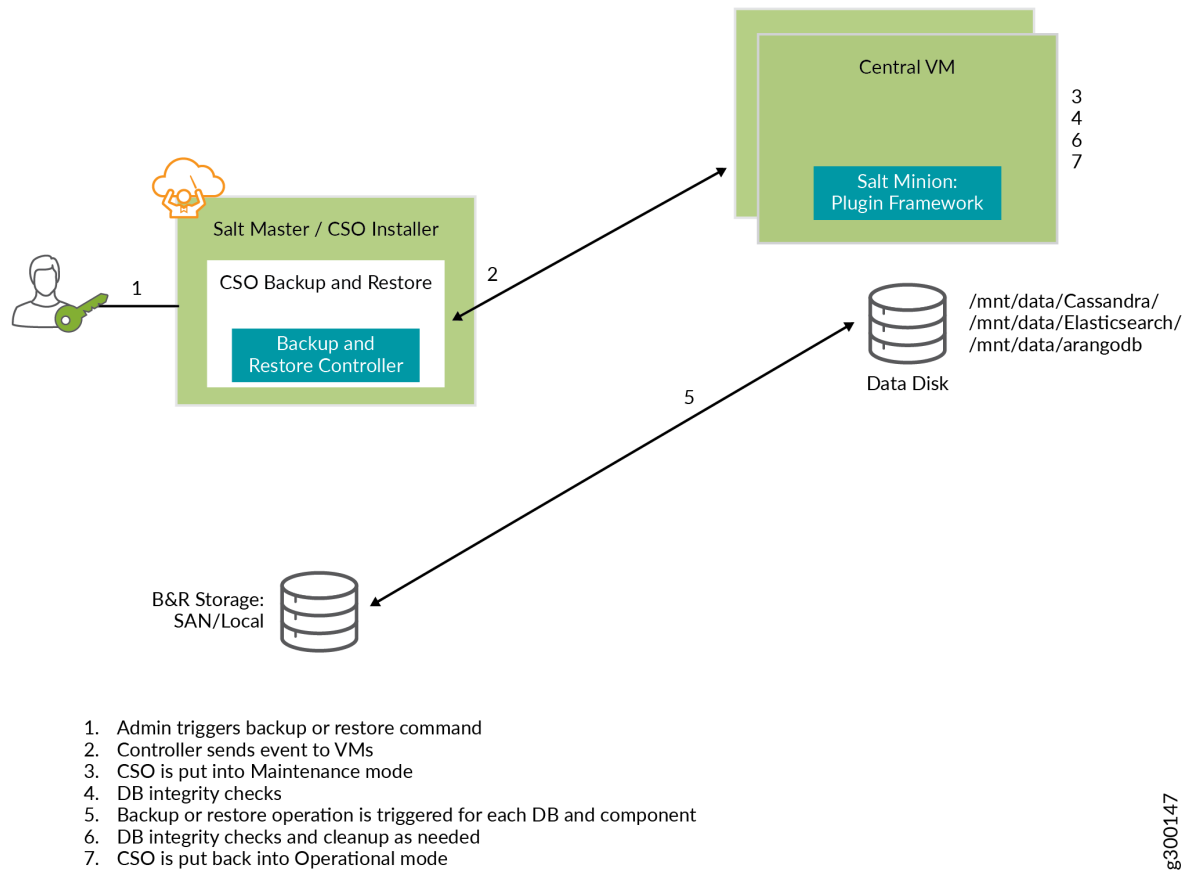
Table 3: Major Components

Component	Description
Backup and Restore Controller	<ul style="list-style-type: none">• Handles backup or restore calls from administrator. The calls are made using the <code>cso_backupnrestore</code> script that resides only on the <code>startup-vm</code>.• Communicates and delegates requests to individual plug-ins.• Manages lifecycle of backup and restore operations: pre-hook, backup and restore, and post-hook.• Salt Master
Plug-in Framework	<ul style="list-style-type: none">• Framework that allows backup and restore to deal with multiple different databases.• Allows for future inclusion of other databases.• Salt Minions
Plug-in	<ul style="list-style-type: none">• Addition of new plug-in has to adhere to standards.• All plug-ins are triggered by backup and restore controller.• Pre-hook, post hook and backup or restore operations are implemented by individual plug-ins.

Operations

All of the backup and restore operations are performed using the command line interface (CLI) of the startup-vm machine. The user in charge of the operations logs onto the startup-vm over ssh and performs any needed operations. [Figure 2 on page 14](#) shows the flow of backup and restore operations.

Figure 2: Backup and Restore Operations



For backup operations, run the `cso_backupnrestore` command on the startup-vm, using the proper arguments for backing up an individual database or all of the databases. When this happens, the backup and restore controller communicates the backup request to the individual plug-ins using the SaltStack message bus. The plug-ins that reside on the various central and regional vms receive the message and trigger the needed action.

Backups are stored in the `/backups/` directory on the startup-vm. This location cannot be changed.

For restore operations, the same `cso_backupnrestore` command is used with different options as described in [Table 4 on page 15](#) below. System stability is confirmed, and the needed restore commands are sent to the plug-ins for each database as needed. Once the restore is finished, CSO checks for system stability again, does any required cleanup and puts itself back into operational mode.

Command Usage

The CLI command used to create backups or restore files from backup is named **cso_backupnrestore**.

Options available for the **cso_backupnrestore** command are shown in [Table 4 on page 15](#). Only one of the arguments can be used with any one of the options.

Table 4: cso_backupnrestore Command Options

Option	Purpose	Arguments
-b	Specify operation (REQUIRED)	backup, restore, healthcheck, reindex, backupdetails, listbackups, scheduledbackup
-s	Specify the name of the snapshot created by backup operation or restored by restore operation.	backup name
-c	Specify which database to backup or restore [default '*'] (OPTIONAL)	For backup: only '*' is allowed. For restore: Comma separated list with any or all of: cassandra, elasticsearch, zookeeper, mariadb, etcd, arrangodb. '*' restores all databases
-r	Specify whether the restore operation is for disaster recovery or not [Default no].	yes or no
-z	Set cron job parameters for backup operation. Only valid in combination with schedulebackup argument for the -b option.	m-h-dom-mon-dow-m [-m yes] <ul style="list-style-type: none"> • m-minute (0-59) • h-hour (0-23) • dom-day of month (1-31) • mon-month (1-12) • dow-day of week (0-6)

Backup and Restore Examples

Requirements

- IP address of the startup virtual machine (startup-vm) of your CSO instance
- Root access to the startup-vm using the ssh protocol

The following commands must be run at the command line interface of the startup-vm of CSO. The location and access credentials needed to access the startup-vm in your CSO installation should be known to you or the person or group who installed CSO.

Backup

This example performs a simple backup of all CSO databases into the directory **/backup/MAR09/**

```
cso_backupnrestore -b backup -s MAR09
```

Scheduled Backup Using Cron-job

This example creates a scheduled backup that runs in maintenance mode every Sunday afternoon at 1:00 PM and stores the backup in the **/backups/DAILY/<timestamp>/** directory. The timestamp directory is created when the backup starts.

```
cso_backupnrestore -b scheduledbackup -z 0-13-*--0
```

Restore

This example restores the backup located in the **/backups/DAILY-09/2019-03-16T04/** directory.

```
cso_backupnrestore -b restore -s /backups/DAILY-09/2019-03-16T04 -r no
```

This example performs a disaster recovery restore operation from the backup located in the **/backups/DAILY-09/2019-03-16T04/** directory.

```
cso_backupnrestore -b restore -s /backups/DAILY-09/2019-03-16T04 -r yes
```

Health Check Example

This example performs a health check on the CSO installation.

```
cso_backupnrestore -b healthcheck
```

Reindex Example

This example performs a reindex of the Elasticsearch database.

```
cso_backupnrestore -b reindex
```

Restore a Specific Component Example

This example restores a specific component, ETCD.

```
cso_backupnrestore -b restore -s /backups/DAILY-09/2018-03-16T04 -c 'etcd'
```

Running Health check for a Specific Component Example

This example performs the health check of a specific component, ETCD.

```
cso_backupnrestore -b healthcheck -c etcd
```

Clear Index for Elasticsearch Example

This example clears the index for Elasticsearch.

```
cso_backupnrestore -b clear_indices
```

Repair Nodetool in Cassandra Example

This example repairs the Nodetool in Cassandra.

```
cso_backupnrestore -b nodetool_repair
```

List Available Backups Example

This example lists all backups that are available.

```
cso_backupnrestore -b listbackups
```

View Details of a Specific Backup Example

This example lists the details of a specific backup.

```
cso_backupnrestore -b backupdetails -s MAR09
```

Release History Table

Release	Description
4.1	Starting in CSO 4.1, full backup of all platform, op-co, tenant, and customer data can be run manually or periodically and that data can be restored from the backups when and if the need arises.

Introduction to Service and Infrastructure Monitor Application for On-premise Deployments

IN THIS CHAPTER

- [Service and Infrastructure Monitor Overview | 19](#)
- [Accessing the Service and Infrastructure Monitor GUI | 20](#)
- [Monitoring Network Services | 21](#)
- [Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 22](#)
- [Monitoring Microservices | 26](#)
- [Monitoring Microservices and Their Host VMs | 27](#)
- [Monitoring Physical Servers | 29](#)

Service and Infrastructure Monitor Overview

Service and Infrastructure Monitor (SIM) operates with the third-party monitoring software Icinga to provide complete monitoring and troubleshooting of the Contrail Service Orchestration (CSO) solution.

When you deploy the CSO solution, an Icinga agent is installed on servers and virtual machines (VMs), which enables Icinga to monitor data on:

- Physical servers
- VMs that host virtualized network functions (VNFs)
- VMs that host microservices

Service and Infrastructure Monitor collects events from microservices in the CSO solution, and correlates the events to provide information about network service, their component VNFs, and the VMs that host the VNFs.

All data is presented through the Icinga GUI. You use the GUI to obtain a quick visual display of the CSO solution status and more detailed lists of event messages.

RELATED DOCUMENTATION

[Monitoring Network Services | 21](#)

[Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 22](#)

[Monitoring Microservices | 26](#)

[Monitoring Microservices and Their Host VMs | 27](#)

[Monitoring Physical Servers | 29](#)

Accessing the Service and Infrastructure Monitor GUI

To access the GUI for Service and Infrastructure Monitor:

1. Using a web browser, access the URL for Service and Infrastructure Monitor:

`http://central-IP-Address:1947/icingaweb2`

central-IP-Address—IP address of the server or VM that hosts the microservices for the central point of presence (POP).

For example:

`http://192.0.2.1:1947/icingaweb2`

2. Log in with the username `icinga` and the encrypted password.

Colored squares, which may contain numbers, in the GUI provide a visual status of the CSO solution network.

- A green square indicates the number of items that are working correctly.
- A yellow square indicates the number of items with potential problems to investigate.
- A red square indicates the number of items that are not working.
- A purple square indicates the number of items with a failed connection.

The following options in the left navigation pane of the Icinga GUI are customized for the CSO solution:

- Dashboard
- Network Services
- Infrastructure

Other features in the Icinga GUI are not customized and appear in the standard Icinga GUI.

See the Icinga documentation for a general overview of the GUI and information about all non-customized features.

RELATED DOCUMENTATION

[Service and Infrastructure Monitor Overview](#) | 19

Monitoring Network Services

Service and Infrastructure Monitor displays information about network services running in the deployment. This information is related to the Network Service Overview on the dashboard, which displays information about component VNFs of network services and the VMs in which the VNFs reside. In this view, however, the focus is on the actual network service rather than its component VNFs and the VMs in which they reside.

To monitor network services:

1. In the left navigation pane, click **Network Services**.

Service and Infrastructure Monitor displays an array of network services and monitoring parameters.

2. In the array, hover over an entry to see additional information for the entry.
3. Click a colored square to see detailed information for the entry.

[Table 5 on page 21](#) shows the meaning of the monitoring parameters for network services.

Table 5: Parameters for Monitoring Network Services

Parameter	Meaning
Network Service	Name of the network service.
Network Service status	State of the network service and the time it entered that state. <ul style="list-style-type: none">• Up—operational• Down—not operational
Number of Network Functions	Number of VNFs in the service chain.

Table 5: Parameters for Monitoring Network Services (*continued*)

Parameter	Meaning
Network Function	<p>Number of network functions in a colored square that indicates the status of the instance. When you click the square you see:</p> <ul style="list-style-type: none"> • An entry for each VNF in the service chain. • The status of the host in which the VNF resides. • The IP address of the host in which the VNF resides. • The name of the VNF. • The result from the last ping the Icinga agent sent to the host, including any loss of packets, and the round trip average (RTA) travel time.
Commands	Total number of commands issued to monitor the status of the network service since it became operational.
Command Status	<p>Result of the commands issued to monitor the status of the network service. When you click the square you see:</p> <ul style="list-style-type: none"> • A list of parameters for a specific network function and its host. • The state of the parameter and how long the parameter has been in that state. • Additional details about the state of the host.

RELATED DOCUMENTATION

[Monitoring VNFs Used in Network Services and the VMs That Host the VNFs](#) | 22

Monitoring VNFs Used in Network Services and the VMs That Host the VNFs

On the dashboard, the Network Service Overview provides information about the VNFs used in network services and the VMs that host those VNFs. You can also view information about the component VNFs in a network service by clicking Monitor Network Services in the left navigation bar.

To view information about VNFs used in network services and the VMs that host the VNFs:

1. In the left navigation bar, click **Dashboard**.

The dashboard appears, displaying several arrays of information.

2. (Optional) In the Network Services Overview array, hover over a colored square in the array to see the latest event message for a specific parameter and host.
3. (Optional) In the Network Services Overview array, click a colored square to see detailed information for a specific parameter and host.
4. (Optional) In the Network Services Overview array, click an IP address to view all the event messages for a host.
5. (Optional) In the Network Services Overview array, click a parameter name to view event messages on all hosts for that parameter.

See [Table 6 on page 23](#) for information about the monitoring parameters used for VNFs and the VMs that host them.

Table 6: Parameters for Monitoring VNFs and Their Host VMs

Parameter	Meaning
left_net_interface_input_pkt_rate	Rate of traffic entering the interface that transmits data to the host.
left_net_interface_output_pkt_rate	Rate of traffic leaving the interface that transmits data to the host.
left_net_interface_stats	State of the interface that transmits data to the network host. <ul style="list-style-type: none"> • Up—operational • Down—not operational
right_net_interface1_stats	State of the interface to which the host transmits data. <ul style="list-style-type: none"> • Up—operational • Down—not operational
right_net_interface_input_packet_rate	Rate of traffic entering the interface to which the host transmits data.
right_net_interface_output_packet_rate	Rate of traffic leaving the interface to which the host transmits data.
routing_engine_ctrlplane_memusage	Percentage of the Routing Engine's control plane memory that VM is using.

Table 6: Parameters for Monitoring VNFs and Their Host VMs (*continued*)

Parameter	Meaning
routing_engine_load_average	Mean percentage of available load capacity used by the Routing Engine's control plane.
routing_engine_system_cpu	Percentage of available CPU capacity used by the Routing Engine's control plane.
<VNF>_activesessions	Number of active sessions of the VNF compared to the maximum number of sessions allowed.
<VNF>_failedsessions	Number of sessions of the VNF that VNF Manager failed to activate.
<VNF>_performance_session	Number of sessions added (ramp-up rate) for the last 60 seconds. The value does not display the total number of sessions or the number of deleted sessions.
<VNF>_performance_spu	Services processing unit (SPU), percentage of CPU capacity that handles the data plane for the security service.
check_flowd	Status of the forwarding process on the vSRX VNF. <ul style="list-style-type: none"> • Up—operational • Down—not operational
vsrx_activesessions	Number of active sessions of the vSRX VNF compared to the maximum number of sessions allowed.
vsrx_failedsessions	Number of sessions of the VNF that VNF Manager failed to activate.
vsrx_system_uptime	Amount of time since the vSRX VNF last became operational.
system_memory	Percentage of available RAM used by the vSRX VNF.
left_net_interface_status	State of the interface that transmits data to the network host. <ul style="list-style-type: none"> • Up—operational • Down—not operational
right_net_interface_status	State of the interface to which the host transmits data. <ul style="list-style-type: none"> • Up—operational • Down—not operational
right_net_interface_input_pkt_rate	Rate of traffic entering the interface to which the host transmits data.

Table 6: Parameters for Monitoring VNFs and Their Host VMs (*continued*)

Parameter	Meaning
right_net_interface_output_pckt_rate	Rate of traffic leaving the interface to which the host transmits data.
vsrx_nat_config	State of the vSRX NAT VNF. <ul style="list-style-type: none"> • Enabled—operational • Disabled—not operational
vsrx_firewall_config	State of the vSRX firewall VNF. <ul style="list-style-type: none"> • Enabled—operational • Disabled—not operational
vsrx_utm_config	State of the vSRX UTM VNF. <ul style="list-style-type: none"> • Enabled—operational • Disabled—not operational
vsrx_dpi_config	State of the DPI firewall VNF. <ul style="list-style-type: none"> • Enabled—operational • Disabled—not operational
iptables_status	State of the LxCIPtable VNF. <ul style="list-style-type: none"> • Enabled—operational • Disabled—not operational
iptables_system_uptime	Amount of time since the LxCIPtable VNF last became operational
cisco_left_interface_status	State of the interface that transmits data to the network host for the CSR-1000V VNF. <ul style="list-style-type: none"> • Up—operational • Down—not operational
cisco_right_interface_status	State of the interface to which the host transmits data for the CSR-1000V VNF. <ul style="list-style-type: none"> • Up—operational • Down—not operational
cisco_left_input_packets	Rate of traffic entering the interface that transmits data to the host for the CSR-1000V VNF.
cisco_left_output_packets	Rate of traffic leaving the interface that transmits data to the host for the CSR-1000V VNF.

Table 6: Parameters for Monitoring VNFs and Their Host VMs (*continued*)

Parameter	Meaning
cisco_right_input_packets	Rate of traffic entering the interface to which the host transmits data for the CSR-1000V VNF.
cisco_right_output_packets	Rate of traffic leaving the interface to which the host transmits data for the CSR-1000V VNF.
cisco_system-uptime	Amount of time since the Cisco CSR-1000V VNF last became operational.
cisco_activesessions	Number of active sessions of the Cisco CSR-1000V VNF compared to the maximum number of sessions allowed.

RELATED DOCUMENTATION

[Monitoring Network Services](#) | 21

Monitoring Microservices

Service and Infrastructure Monitor displays information about microservices running in each Contrail Service Orchestration (CSO) implementation. This information is related to the CSP Microservice Overview on the dashboard, which displays information about the VMs in which the microservices reside. In this view, however, the focus is on the actual microservices rather than the VMs in which they reside.

To monitor microservices:

1. In the left navigation pane, select **Infrastructure > CSP Microservices**.

Service and Infrastructure Monitor displays an array of CSP microservices and monitoring parameters.

2. (Optional) In the array, hover over an entry to see additional information for the entry.
3. (Optional) Click a colored square to see detailed information for the entry.

[Table 7 on page 27](#) shows the monitoring parameters for microservices.

Table 7: Parameters for Monitoring Microservices

Parameter	Meaning
CSP Microservice	Name of the microservice.
Microservice status	State of the microservice and the time it entered that state. <ul style="list-style-type: none"> • Up—operational • Down—not operational
Number of Instances	Number of instances of the microservice.
Instance Status	Number of microservices in a colored square that indicates the status of the instance. When you click the square you see: <ul style="list-style-type: none"> • The status of the host in which the microservice resides. • The IP address of the host in which the microservice resides. • The name of the microservice. • The result from the last ping the Icinga agent sent to the host, including any loss of packets, and the round trip average (RTA) travel time.
Monitor Commands	Total number of commands issued to monitor the status of the microservice since it became operational.
Command Status	Result of the commands issued to monitor the status of the microservice. When you click the square you see: <ul style="list-style-type: none"> • A list of parameters for a specific host. • The state of the parameter and how long the parameter has been in that state. • Additional details about the state of the host.

RELATED DOCUMENTATION

[Monitoring Microservices and Their Host VMs](#) | 27

Monitoring Microservices and Their Host VMs

On the dashboard, the CSP Microservices Overview provides information about the VMs that host microservices. The focus of the CSP Microservices Overview is the VMs that host the microservices.

To monitor microservices and their host VMs:

1. In the left navigation bar, click **Dashboard**.

The dashboard appears, displaying several arrays of information.

2. (Optional) In the CSP Microservices Overview array, hover over a colored square in the array to see the latest event message for a specific parameter and host.
3. (Optional) In the CSP Microservices Overview array, click a colored square to see detailed information for a specific parameter and host.
4. (Optional) In the CSP Microservices Overview array, click an IP address to view all the event messages for a host.
5. (Optional) In the CSP Microservices Overview array, click a parameter name to view event messages on all hosts for that parameter.

See [Table 8 on page 28](#) for information about the monitoring parameters used for VNFs and the VMs that host them.

Table 8: Parameters for Monitoring VNFs and Their Host VMs

Parameter	Meaning
check cpu usage	Percentage of unused CPU capacity
check disk IO	Status of host's input and output mechanisms for storage
check disk usage	Available storage on the VM that hosts the microservice
check elasticsearch	Number of processes associated with the database
check load average	Measure of load compared to specified values for warning and critical states
check memory usage	Percentage of RAM and swap memory used
check network usage	Percentage of network resources used
check nsdui	Availability of the Network Service Designer application
check open files	Number of open files compared to specified values for warning and critical states

Table 8: Parameters for Monitoring VNFs and Their Host VMs (*continued*)

Parameter	Meaning
check_paging_stats	Amount of data moved from RAM to swap memory compared to specified values for warning and critical states
check_socket_usage	Number of software connections compared to specified values for warning and critical states
check_contrail_api	Number of Contrail API processes
check_contrail_config	Number of Contrail configuration processes
check_contrail_control	Number of Contrail control processes
check_contrail_database	Number of Contrail database processes
check_contrail_vrouter	Number of Contrail Vrouter processes
check_contrail_vrouter_agent	Number of Contrail Vrouter agent processes
check_contrail_web	Number of Contrail web core processes
check_ifmap_server	Number of Interface for Metadata Access Points (IF-MAP) processes
check_nova_api	Number of Nova API processes

RELATED DOCUMENTATION

[Monitoring Microservices](#) | 26

Monitoring Physical Servers

Service and Infrastructure Monitor tracks the state of each physical server on which the Icinga agent is installed.

To monitor physical servers:

1. In the left navigation bar, click select **Infrastructure > CSP Bare Metal**.

Service and Infrastructure Monitor displays an array of physical servers and monitoring parameters.

2. In the array, hover over an entry to see additional information for the entry.
3. Click a colored square to see detailed information for the entry.

See [Table 9 on page 30](#) for information about the parameters.

Table 9: Parameters for Monitoring Physical Servers

Parameters	Meaning
Group Status	<p>State of the server cluster and the time when it entered that state.</p> <ul style="list-style-type: none"> • Up—Operational • Down—Not operational
Number of Servers	Number of servers in the server cluster.
Server Status	<p>Number of servers in a colored square that indicates the status of the servers. When you click the square you see:</p> <ul style="list-style-type: none"> • An entry for each server in the cluster. • The status of the server. • The IP address of the server. • The hostname of the server. • The result from the last ping the Icinga agent sent to the server, including any loss of packets, and the round trip average (RTA) travel time.
Commands	Total number of commands issued to monitor the status of the server since it became operational.
Command Status	<p>Result of the commands issued to monitor the status of the server. When you click the square you see:</p> <ul style="list-style-type: none"> • A list of parameters for a specific server. • The state of the parameter and how long the parameter has been in that state. • Additional details about the state of the server.

RELATED DOCUMENTATION

[Service and Infrastructure Monitor Overview](#) | 19

2

PART

Troubleshooting Contrail Service Orchestration Issues

Troubleshooting Login Issues for Cloud-based Deployments | **32**

Troubleshooting POPs, Tenants, and Devices Issues for On-premise
Deployments | **35**

Identifying Connectivity Issues for Cloud-based Deployments | **38**

Troubleshooting Site Activation Issues for Cloud-based Deployments | **45**

Troubleshooting Image, License, and Policy Deployment Issues for Cloud-based
Deployments | **48**

Troubleshooting CSO Installation Issues for On-premise Deployments | **53**

Troubleshooting SMTP Issues for Cloud-based Deployments | **57**

Troubleshooting RBAC and OpCo Issues for On-premise Deployments | **60**

Troubleshooting Site, Device and Link Issues for On-premise Deployments | **65**

Recovering an Installation | **71**

Troubleshooting Login Issues for Cloud-based Deployments

IN THIS CHAPTER

- [Troubleshooting Login Issues | 32](#)

Troubleshooting Login Issues

IN THIS SECTION

- [Administration Portal IP Address Is Not Reachable | 32](#)

Administration Portal IP Address Is Not Reachable

Problem

Description: The CSO Administration Portal IP address is not reachable.

Solution

Check the gateway and try to ping the gateway from the virtual machine. If the gateway is correct and if you can ping the gateway, then the issue must be due to the lab switch.

4. Check the routes and firewall policies with the help of Network Administrator.
5. For further troubleshooting, collect the logs and output results and contact Juniper Networks Technical Support team.

RELATED DOCUMENTATION

| *Resetting Your Password*

Troubleshooting POPs, Tenants, and Devices Issues for On-premise Deployments

IN THIS CHAPTER

- [Troubleshooting POPs, Tenants, and Devices Issues | 35](#)

Troubleshooting POPs, Tenants, and Devices Issues

IN THIS SECTION

- [Failure While Creating a Hub, Site, or Tenant | 35](#)
- [Base Configuration for CPE Activation | 36](#)

Failure While Creating a Hub, Site, or Tenant

Problem

Description: A failure occurred when creating a hub, site, or tenant.

Solution

- Check the job logs in the CSO Administration Portal for the task failure and the reason for the failure.
 - a. Login to the Administration Portal and select **Monitor > Jobs**
The Jobs page is displayed.
 - b. Select the failed log and click the Detailed View icon that appears before the failed log name.
The Detailed View page appears, showing the details of the job and the number of tasks associated with the job.
 - c. Click **View Logs**.

The Job status page is displayed.

- If the failure cannot be determined from the job logs, log in to Kibana and check for the logs using the job ID.

Use the Kibana dashboard <http://<central- Infra-vm-IP-Address>:5601> to view the detailed logs of hub, site, and tenant failures.

- Log in to the CSO central microservices virtual machine and execute **kubect**l get pods -n central to get the status of **tssm** and **topology** POD running on the central and regional microservices virtual machine.

```
root@centralmsvm:~# kubectl get pods -n central | grep tssm
```

csp.csp-tssm-711204925-ncjww	1/1	Running
1 18h		
csp.csp-tssm-core-407531667-x57cf	1/1	Running
1 18h		

```
root@centralmsvm:~# kubectl get pods -n central | grep topology
```

csp.csp-topology-service-3409064476-30hfr	1/1	Running	1
18h			
csp.csp-topology-service-core-1954971038-x5v0w	1/1	Running	
1 18h			

Check the status of the POD.

Execute **kubect**l logs -f pod-name -n central.. For example,

```
root@centralmsvm:~# kubectl logs -f csp.csp-tssm-core-* -n central
```

- For further troubleshooting, collect the logs and output results and contact Juniper Networks Technical Support team.

Base Configuration for CPE Activation

Problem

Description: User was unable to activate a CPE device. Specify the base configuration to activate a CPE device after loading a factory default configuration.

Solution

For Zero Touch Provisioning (ZTP) using the Juniper Networks redirect server and the Dynamic Host Configuration Protocol (DHCP) on a WAN interface (ge-0/0/0), no configuration is required from the user. The CPE activation proceeds with the factory default configuration.

If the CPE device has to be pre-staged based on customer-specific requirements such as a static IP address on WAN interfaces, using the CSO activation server as a phone-home server instead of the Juniper Networks redirect server, then execute the following additional configurations on the CPE device after the factory default configuration.

CPE-SRX [Two WAN Links]

```
set interfaces ge-0/0/0 unit 0 family inet address 192.1.1.1/29
set interfaces ge-0/0/1 unit 0 family inet address 192.1.1.2/24
set routing-options static route 0.0.0.0/0 next-hop 198.1.1.1
set security zones security-zone untrust interfaces ge-0/0/0.0
set security zones security-zone untrust interfaces ge-0/0/1.0
set security zones security-zone untrust interfaces ge-0/0/0.0 host-inbound-traffic
  system-services ssh
set system phone-home server https://regionalmsvm.englab.juniper.netset system
phone-home ca-certificate-file /root/ssl_cert.crt
set system static-host-mapping regional msvm.englab.juniper.net
```

CPE-NFX (JDM Console)

```
set system phone-home server https://CSO-regional-ms-vm-ip
set interfaces jsxe0 unit 0 family inet dhcp [or]
set interfaces jsxe0 unit 0 family inet address 192.1.1.5/29
set routing-options static route 0.0.0.0/0 next-hop 198.1.1.2
set interfaces jmgmt0 unit 0 disable << disable to avoid the default route overlap
```

You must copy the ssl_cert.crt certificate to NFX CPE device/JDM: /var/phone-home/phcd-ca.crt

Identifying Connectivity Issues for Cloud-based Deployments

IN THIS CHAPTER

- [Identifying Connectivity Issues by Using Ping | 38](#)
- [Identifying Connectivity Issues by Using Traceroute | 42](#)

Identifying Connectivity Issues by Using Ping

You can use Contrail Service Orchestration (CSO) to perform a ping operation from a device (provider hub, tenant device, CPE device, EX switch, enterprise hubs, or next-generation firewall device) to a remote host for identifying issues in connectivity with the remote host.

When you ping a remote host from a device, an Internet Control Message Protocol (ICMP) packet is sent to the remote host. By analyzing the results of the ping operation, you can identify the possible device connectivity issues between the remote host and the device.

NOTE: In Contrail Service Orchestration (CSO) Release 5.0, the following devices support ping:

- EX Series: EX2300, EX3400, EX4300, EX4600, EX4650
- NFX Series: NFX150, NFX250
- SRX Series: SRX300, SRX320, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600
- vSRX

To perform the ping operation:

1. Do one of the following:

- To initiate a ping from a provider hub device, select **Resources > Provider Hub Devices**.

The :Provider Hub Devices page appears.

- To initiate a ping from a tenant device, select **Resources > Tenant Devices**.

The Tenant Devices page appears.

2. Select a device from the list of devices displayed and click **More > Ping**.

The Ping page appears.

NOTE: You can initiate a ping from a device only when its operational status (in CSO) is Up.

3. Complete the configuration according to the guidelines provided in [Table 10 on page 39](#).

NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Ping** to initiate the ping request.

A job is created and a Ping Progress page appears. After the host sends the ping packets, the Ping Result page appears. If the ping operation is successful, the Ping Result page displays the parameters specified in [Table 11 on page 41](#).

If the ping operation fails, the Ping Result page displays an appropriate error message (such as **No response** or **No route to host**), indicating that there is an issue in the connectivity to the remote host.

Table 10: Fields on the Ping page

Field	Description
Remote Host	Enter the IPv4 address or hostname of the remote host.
Ping Request Packets	Enter the number of ping request packets to be sent to the remote host. Default: 5. Range: 1 through 300.
Advanced	

Table 10: Fields on the Ping page (*continued*)

Field	Description
Source Interface	<p>Select the source interface on the device through which you want to send the ping request to the remote host. If you do not select a source interface, ping requests are sent on all interfaces.</p> <p>To clear the selected interface, click Clear All and select another interface.</p>
Hostname Resolution	Click the toggle button to enable or disable (default) the display of hostname of the hops along the path to the remote host.
Rapid Ping	<p>Click the toggle button to enable or disable (default) sending ping requests rapidly.</p> <p>If you enable this option, the device sends a minimum of 100 ping request packets per second or sends a packet as soon as a response to the previous packet is received, whichever is greater.</p> <ul style="list-style-type: none"> • If the source device does not receive a response for 500 ms, timeout is considered. • If the source device receives a response within 500 ms, the next ping request packet is sent immediately. <p>NOTE: The ping results are displayed in a single consolidated message instead of individual messages for each ping request packet sent.</p>
Packet Fragmentation	<p>Click the toggle button to enable or disable (default) the fragmenting of ping request packets.</p> <p>If packet fragmentation is disabled, ping packets with the maximum transmission unit (MTU) greater than 1500 bytes are dropped.</p>
Packet Size (bytes)	<p>Enter the size (in bytes) of the ping request packet.</p> <p>Default: 56 bytes.</p> <p>Range:</p> <ul style="list-style-type: none"> • 1 through 1,472 bytes, if packet fragmentation is disabled. • 1 through 65,468 bytes, if packet fragmentation is enabled.
Wait Time (seconds)	<p>Enter the time (in seconds) for which the source device waits for a response to the ping request packet. The source device considers the remote host as not reachable after the wait time elapses.</p> <p>Default: 10 seconds.</p> <p>Range: 0 through 600 seconds.</p>

Table 10: Fields on the Ping page (*continued*)

Field	Description
Incoming Interface	Click the toggle button to include or exclude (default) information (on the Ping Result page) about the interface on the source device that receives the ping responses.
Routing Instance	<p>Select a specific routing instance that the ping request packets can use to reach the remote host.</p> <p>The ping result displays the information about the connectivity between the source device and the remote host based on the selected routing instance.</p> <p>To clear the selected routing instance, click Clear All and select another routing instance.</p>

Table 11: Fields on the Ping Result page

Field	Description
Packet Loss	Displays the percentage of ping packets sent for which the source device did not receive a response.
Round Trip Time Taken (in μ s)	<p>Displays the following information about the duration (in microseconds) between the time when the device sends the ping request and the time when the device receives a response from the remote host.</p> <p>Displays the following:</p> <ul style="list-style-type: none"> • Minimum: The minimum time taken to receive a response for a ping request packet. • Maximum: The maximum time taken to receive a response for a ping request packet. • Average: The average time taken to receive a response for all the ping request packets sent in a ping operation. • Standard Deviation: The variation of the round trip time from the mean round trip time.

Details

Sequence	Sequence number of all the ping request packets.
Result	Result of the ping request packets—Success or Failure.
Incoming Interface	<p>Interface on the source device on which the responses are received for the ping requests.</p> <p>This data appears if you have enabled the Incoming Interface option on the Ping page.</p>
Time Taken	Time taken (in microseconds) to receive response to a ping request packet.

Identifying Connectivity Issues by Using Traceroute

You can use Contrail Service Orchestration (CSO) to perform a traceroute operation from a device (provider hub, tenant device, CPE device, EX switch, enterprise hubs, or next-generation firewall device) to the remote host. Traceroute helps you view the path that a packet travels to reach the remote host. The result is useful in identifying the point of network failure in the path between the source device and remote host.

NOTE: In Contrail Service Orchestration (CSO) Release 5.0, the following devices support traceroute:

- EX Series: EX2300, EX3400, EX4300, EX4600, EX4650
- NFX Series: NFX150, NFX250
- SRX Series: SRX300, SRX320, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600
- vSRX

To perform traceroute operation:

1. Do one of the following:
 - To initiate traceroute from a provider hub device, select **Resources > Provider Hub Devices**.
The Provider Hub Devices page appears.
 - To initiate traceroute from a tenant device, select **Resources > Tenant Devices**.
The Tenant Devices page appears.
2. Select a device from the list of devices displayed and click **More > Traceroute**.
The Traceroute page appears.
3. Complete the configuration according to the guidelines provided in [Table 12 on page 43](#).

NOTE: Fields marked with an asterisk (*) are mandatory.

4. Click **Traceroute** to initiate the traceroute operation.

A job is created and a traceroute progress page appears. If the traceroute operation is successful, the Traceroute Result page displays the traceroute parameters specified in [Table 13 on page 44](#).

If the traceroute operation fails, the Traceroute Result page displays an appropriate error message (such as **No response** or **No route to host**).

Table 12: Fields on the Traceroute page

Field	Description
Remote Host	Enter the IPv4 address or hostname of the remote host.
Maximum Hops	<p>Specify the maximum number of network devices that a packet can pass through to reach the remote host.</p> <p>Default: 30.</p> <p>Range: 1 through 255.</p> <p>If the number of hops to reach the remote host exceeds the set value, the traceroute packet is dropped.</p>
Advanced	
Source Interface	<p>Select a source interface on the device from which you want to send the packets to the remote host.</p> <p>Click Clear All to remove the selected interface and select another interface.</p>
Hostname Resolution	Click the toggle button to enable or disable (default) the display of hostname of the hops in the path to the remote host.
Wait Time (seconds)	<p>Enter the time until which the device waits for a response from the remote host to a packet sent before considering timeout.</p> <p>Default: 10 seconds.</p> <p>Range: 0 through 86,399 seconds.</p>
Routing Instance	<p>Select a routing instance that the traceroute request packets can use to reach the remote host.</p> <p>The trace result displays the route information based on the configured routing instance type.</p> <p>To clear the selected routing instance, click Clear All and select another routing instance.</p>

Table 13 on page 44 lists the parameters on the Traceroute Result page when the traceroute operation is successful.

Table 13: Fields on the Traceroute Result page

Field	Description
Hop	Hostname or IPv4 address of the network devices that the packet passed through to reach the remote host.
Time Taken by Packet 1	Duration (in microseconds) between the time from when the source device sends a packet, and the time it received a response from the hops and the remote host.
Time Taken by Packet 2	
Time Taken by Packet 3	

Troubleshooting Site Activation Issues for Cloud-based Deployments

IN THIS CHAPTER

- [Troubleshooting Site Activation Issues | 45](#)

Troubleshooting Site Activation Issues

IN THIS SECTION

- [Prerequisites to Activate a Site | 45](#)
- [Site activation process is stuck in device detected state | 46](#)
- [Site activation process is stuck in bootstrap state | 46](#)
- [Site activation process failed in bootstrap state | 46](#)
- [Site activation process failed during provisioning | 46](#)

Prerequisites to Activate a Site

Problem

Description: User was unable to activate a site. Specify the prerequisites to activate a site.

Solution

The prerequisites to activate a site are as follows:

- Check the spoke connectivity to Internet.
- Check the firewall policies between the CPE device and the CSO. The hub or spoke must be able to communicate to CSO through ports 443 (activation), 444 (activation for small and medium deployments),

7804 (outbound-ssh), 3514(app-track logs), 514 (syslog), and 2216 (telemetry agent). See *Contrail Service Orchestration (CSO) Deployment Guide*

Site activation process is stuck in device detected state

Problem

Description: Site activation process is stuck in device detected state; how do I proceed?

Solution

Do the following:

- Verify that your device can reach the Internet.
- Verify the date and time on the device.
- Verify that the DHCP server and the device are connected to the ge-0/0/0 port.
- Reboot the device.

Site activation process is stuck in bootstrap state

Problem

Description: Site activation process is stuck in bootstrap state; how do I proceed?

Solution

If the site activation process is stuck for more than 15 minutes, then do the following:

- Verify that your network firewall allows UDP ports 500 and 4500 for the SD-WAN site.
- Verify that your network firewall allows TCP port 7804 for the next-generation firewall site.
- Reboot the device.

Site activation process failed in bootstrap state

Problem

Description: Site activation process failed in bootstrap state; how do I proceed?

Solution

Verify that the device is zeroized or running the factory-default configuration. If the device is pre-staged, then ensure that the configuration is not overlapping with the CSO stage-1 configuration. Reboot the device.

Site activation process failed during provisioning

Problem

Description: Site activation process failed during provisioning; how do I proceed?

Solution

Verify the device connectivity to the Internet. Retry the failed job in CSO. Navigate to **Monitor > Jobs**, select the failed job, and click **Retry Job**.

Troubleshooting Image, License, and Policy Deployment Issues for Cloud-based Deployments

IN THIS CHAPTER

- [Troubleshooting Image, License, and Policy Deployment Issues | 48](#)

Troubleshooting Image, License, and Policy Deployment Issues

IN THIS SECTION

- [Unable to find device image version | 48](#)
- [Upgrade device image using J-Web | 49](#)
- [Unable to connect to the device | 49](#)
- [Device image version is different from the recommended version | 50](#)
- [LAN Switch image version is different from the recommended version | 50](#)
- [Policy deployment failed | 50](#)
- [No data for next-generation firewall site | 51](#)
- [No data for SD-WAN site | 51](#)
- [Traffic from Spoke Sites Are Dropped or Are Not Reaching Internet or Destination | 51](#)
- [SLA Violation-Original Link Recovered After SLA Violation | 52](#)
- [All WAN links are Up But Not All Links Are Utilized | 52](#)

Unable to find device image version

Problem

Description: How do I find my device image version without console access to the device?

Solution

Use the J-Web interface to find the device image version.

To access the J-Web interface of the device:

1. Connect your laptop or workstation to any port (except ge-0/0/0) that is available on the device.
2. Enable DHCP on the laptop or workstation and acquire the IP address and gateway information from the device.
3. Use the gateway address (also known as the device address) in the Web browser to connect to the J-Web interface.
4. Log in with the default username **root**. As the root user, you don't need a password to log in.

The Welcome page appears displaying the device image version.

Upgrade device image using J-Web

Problem

Description: Device image version is 15.1X49-D110; how do I upgrade the device image before site onboarding?

Solution

Use the J-Web interface to upgrade the device image.

To upgrade the device image using J-Web:

1. Download the recommended image or the software version from the Juniper Networks website to your local machine.
2. Log in to the J-Web interface.
3. Select **Maintain > Software > Upload Package**.
4. Navigate to the device image file location and select the file.
5. Click **Upload and Install Package** to upgrade the device image.

Unable to connect to the device

Problem

Description: I am not able to log in to the device through the J-Web interface or through the device console. How do I proceed?

Solution

Press and hold the Reset Config button on the device for 15 seconds. Wait for two minutes for the device to restore the factory-default settings. Log in to the device as the root user (no password is required for the root user). If you are still not able to access the device, then reboot the device.

Device image version is different from the recommended version**Problem**

Description: The device image version at the site is 15.1X49D110, but the recommended image version is 15.1X49D170.x. Should I upgrade the device image manually before site onboarding?

Solution

You don't need to upgrade the device image manually before site onboarding. You can do either of the following:

- Upgrade the device image during site activation in CSO—While you are in the site configuration or onboarding workflow, select the device image from the drop-down list.

NOTE: Device image upgrade during site activation delays the site activation process.

- Upgrade the device image post site activation in CSO—Navigate to **Resources > Images**, select the image, and click **Deploy**.

LAN Switch image version is different from the recommended version**Problem**

Description: LAN switch (EX Switch) image version is lower-numbered version than the recommend version. Can I proceed with onboarding the site?

Solution

Use the recommend image version for the device.

Policy deployment failed**Problem**

Description: Policy deployment failed; how do I proceed?

Solution

Verify the device connectivity to the Internet. Retry the policy deployment.

No data for next-generation firewall site

Problem

Description: Application Visibility Monitoring page shows no data for the next-generation firewall site; how do I proceed?

Solution

Do the following:

- Verify that your network firewall allows the UDP port 514.
- Verify the application visibility monitoring page after multiple application sessions (in the time range of 3–5 minutes) traffic.
- Use an appropriate time interval for the query. For example, if you are querying for the traffic sent in the last 10 minutes, then try using a 15-minute query (minimum time interval).

No data for SD-WAN site

Problem

Description: Application visibility and WAN performance data on the Site Management page shows no data for the SD-WAN site; how do I proceed?

Solution

Do the following:

- Verify the application visibility and WAN performance data after multiple application sessions (in the time range of 3-5 minutes) traffic.
- Use an appropriate time interval for the query. For example, if you are querying for the traffic sent in the last 10 minutes, then try using a 15-minute query (minimum time interval).

Traffic from Spoke Sites Are Dropped or Are Not Reaching Internet or Destination

Problem

Description: Traffic from spoke sites are dropped or are not reaching the Internet or their specified destinations.

Solution

1. Verify the alerts for overlay or underlay connections, and check whether BGP is active.
Log in to Administration portal, and select **Monitor > Alerts and Alarm > Alerts**.
2. Check whether the firewall policies are successfully deployed to the CPE device and that the traffic or applications are matching the policies to permit the traffic to Internet or to other sites.

In Administration Portal, select **Sites > Site-Name > Policies**.

Or log in to the CPE device and verify that the next-generation firewall policies are deployed.

3. Check the routes in the default VRF route table in the CPE device.
4. Trace the route and verify the reachability from the hub to the destination. If the hub cannot reach the Internet, then verify whether the firewall and NAT policies are set up properly in the hub.
5. For further troubleshooting, collect the logs and output results and contact Juniper Networks Technical Support team.

SLA Violation-Original Link Recovered After SLA Violation

Problem

Description: The original link is recovered after a service-level agreement (SLA) violation but the application traffic does not switch back to the original link.

Solution

Applications change links only on an SLA violation, because applications are not tied to a specific link and are based on SLA type, such as path preference or link performance metrics.

All WAN links are Up But Not All Links Are Utilized

Problem

Description: All WAN links are up but not all links are being utilized.

Solution

It is possible that all SD-WAN policies can select the same WAN link if they match the SLAs. If the CPE receives a lot of matching and non-matching application traffic for SD-WAN policies, but not all WAN links are being used, then ensure the following:

1. Check that the CPE device receives multiple flows per application.
2. Check that all the WAN overlays are up (IPsec, GRE) in the CPE device and the hub device.
3. Check the SLA performance data or real-time performance monitoring (RPM) probe results in the CPE device for all links.

Log in to the Administration Portal, and select **Monitor > Applications > SLA Performance**.

Troubleshooting CSO Installation Issues for On-premise Deployments

IN THIS CHAPTER

- [Troubleshooting CSO Installation Issues | 53](#)

Troubleshooting CSO Installation Issues

IN THIS SECTION

- [Salt Key Issue During CSO Installation | 53](#)
- [TimeZone Error | 55](#)
- [SSL Handshake Failure | 55](#)
- [Missing Interface on CSO VM | 56](#)

Salt Key Issue During CSO Installation

Problem

Description: The infrastructure services for the central infrastructure virtual machine (infravm) and the central microservices virtual machine (msvm) ran successfully. However, the Contrail Analytics virtual machine (contrail-analytics-vm) failed with the following key exchange error.

```
root@cso-central-2:~#salt-key -L
```

```
Accepted Keys:
csp-central-infravm.M8DLI0.central
csp-central-msvm.M8DLI0.central
Denied Keys:
Unaccepted Keys:
```



```

Rejected Keys:
2017-06-09 13:09:37 INFO      utils.core      Minion
csp-contrailanalytics-vm.M8DLI0.central pointed to the salt master
2017-06-09 13:09:37 INFO      utils.core      RP: salt_master accepted keys are:
['csp-central-infravm.M8DLI0.central', 'csp-central-msvm.M8DLI0.central']
2017-06-09 13:09:47 ERROR      utils.core      key exchange did not go through by
reactor for csp-contrailanalytics-vm
2017-06-09 13:09:47 INFO      utils.core      Pinging minion
csp-contrailanalytics-vm.M8DLI0.central
2017-06-09 13:09:47 INFO      utils.core      timeout 5
2017-06-09 13:09:47 INFO      utils.core      calling test.ping on
csp-contrailanalytics-vm.M8DLI0.central
2017-06-09 13:09:47 INFO      utils.core      timeout 10
2017-06-09 13:09:47 INFO      utils.core      calling test.ping on
csp-contrailanalytics-vm.M8DLI0.central
2017-06-09 13:09:47 INFO      utils.core      timeout 15
2017-06-09 13:09:47 INFO      utils.core      calling test.ping on
csp-contrailanalytics-vm.M8DLI0.central
2017-06-09 13:09:47 INFO      utils.core      timeout 20
2017-06-09 13:09:47 INFO      utils.core      calling test.ping on
csp-contrailanalytics-vm.M8DLI0.central
2017-06-09 13:09:47 INFO      utils.core      timeout 25
2017-06-09 13:09:47 INFO      utils.core      calling test.ping on
csp-contrailanalytics-vm.M8DLI0.central
2017-06-09 13:09:47 INFO      utils.core      Ping to minion
csp-contrailanalytics-vm.M8DLI0.central failed
2017-06-09 13:09:48 INFO      utils.core      setting up salt agent on
csp-contrailanalytics-vm
2017-06-09 13:09:50 INFO      utils.core      Salt minion already installed on
server csp-contrailanalytics-vm
2017-06-09 13:09:50 INFO      utils.core      Pointing minion to installer host
2017-06-09 13:09:57 INFO      utils.core      attaching to master 192.168.255.53

```

Solution

1. Verify that the deployment script is running on the installer virtual machine and on the analytics virtual machine.
2. Execute the following commands:

```

root@cso-central-2:~#rm -rf /var/cache/salt*
root@cso-central-2:~#sudo apt-get remove salt-minion
root@cso-central-2:~#sudo apt-get purge salt-minion

```

TimeZone Error

Problem

Description: The following error message is displayed when installing the central microservices virtual machine.

```
ERROR    State timezone_|-time_zone_|-America/Los_Angeles_|-system on
csp-central-msvm.D2IHJM.central is in error
ERROR    Some errors while deploying the roles set(['ntp'])
ERROR    Please check the logs and correct
```

Solution

This issue occurs only when the NTP server input is not synchronized with the virtual machine. Check the NTP server details in the setup assistant.

SSL Handshake Failure

Problem

Description: SSL handshake failures are reported on both regional services (regionalmsvm1 and regionalmsvm2) in Icinga.

```
CHECK_NRPE: Error - Could not complete SSL handshake.
Check execution
Command check_nrpe
Check Source  centralmsvm2.sst.net.cn is reachable

Check execution
Command check_nrpe
Check Source  centralmsvm1.sst.net.cn is reachable

192.1.1.1      regionalmsvm1.sst.net.cn regionalmsvm1
192.1.1.2      regionalmsvm2.sst.net.cn regionalmsvm2
```

Solution

1. Edit the `/etc/nagios/nrpe.cfg` file at the regional microservices virtual machine. The **allowed_hosts** might have the central microservices virtual machine management address instead of the Operation, Administration, and Maintenance (OAM) address. Add the OAM IP address. For example:

```
allowed_hosts=192.0.2.0, 192.0.3.0
```

2. Send the **kill -HUP** command to the Nagios Remote Plugin Executor (NRPE) process.

```
root@regional-msvm:/etc/nagios#ps aux | grep nrpe
```

```
nagios      1428  0.0  0.0  23472  1204 ?        Ss   Aug14   0:38 /usr/sbin/nrpe
-c /etc/nagios/nrpe.cfg -d
root        6535  0.0  0.0  10476   888 pts/6    S+   15:40   0:00 grep --color=auto
nrpe
```

```
root@regional-msvm:/etc/nagios#kill -HUP 1428
```

Missing Interface on CSO VM

Problem

Description: When two virtual bridge interfaces(virbr0 and virbr1) on the server are mapped to eth0 and eth1, only one interface is listed in the CSO virtual machines. The second interface is not visible to the CSO virtual machines.

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

bridge name	bridge id	STP enabled	interfaces
virbr0	8000.0cc47a6efbfa	no	eth0 vnet0 vnet1 vnet2 vnet3 vnet4 vnet5 vnet6 vnet7
virbr1	8000.0cc47a6efbfb	no	eth1

Solution

1. Assign an IP address to the second interface in **/etc/network/interfaces**.
2. Set **enable_data_interface** to **true** for all virtual machines and set **data_interface** to **virbr1** for cso-host in the provision_vm.conf file.

Troubleshooting SMTP Issues for Cloud-based Deployments

IN THIS CHAPTER

- [Troubleshooting SMTP Issues | 57](#)

Troubleshooting SMTP Issues

IN THIS SECTION

- [Basic Configuration for SMTP Server | 57](#)

Basic Configuration for SMTP Server

Problem

Description: User was unable to configure the SMTP e-mail server.

Solution

1. Check the SMTP server settings.
 - SMTP server address—Check the host name or network address of the SMTP e-mail server. Typical SMTP server addresses or host names are as follows:
 - smtp.juniper.net
 - smtp.gmail.com
 - smtp.mail.yahoo.com
 - AWS

- **TLS**—Check whether Transport Layer Security (TLS) option is enabled. This setting ensures that the information is transmitted over an encrypted channel. Not all SMTP servers support encryption. If TLS option is enabled for an SMTP server that does not support TLS, then disable the TLS option.
- **Port**—Check with your e-mail service provider for the port number that the SMTP server listens to. Generally, port number 587 is used for a TLS connection and port number 25 is used for unencrypted connections.

Typical SMTP server settings are as follows:

- smtp.juniper.net—Set TLS to No and port number to 25
- smtp.gmail.com—Set TLS to Yes and port number to 587
- smtp.mail.yahoo.com—Set TLS to Yes and port number to 465 or 587

2. Check the SMTP authentication settings.

- Check whether the e-mail server requires authentication. If yes, then specify the following options.
 - From Name
 - User Name
 - Password
 - From E-mail Address

NOTE: If Gmail blocks SMTP e-mails, then log in to Gmail account, navigate to **Advanced Settings > Security > Less secure app access** and click the toggle button to turn on **Allow less secure apps** option.

3. Test SMTP settings by sending a test e-mail.

If you are unable to send a test e-mail:

- Check the SMTP server settings to see if they match the SMTP server provider's settings.
- Check authentication credentials.
- Check the SMTP server provider's security settings for SMTP (for example: Gmail blocks SMTP email unless user selects less secure app settings on their gmail account).
- Check whether there is network access from CSO to the SMTP server.
- Check whether the firewall is blocking SMTP traffic to SMTP server or whether the ports are blocked. If the server settings and authentication settings are correct, check whether the firewall is blocking

port 587 and 465 and SMTP traffic. If it is a case of the firewall blocking, then work with the network administrator to unblock ports 465, 587, and SMTP traffic.

RELATED DOCUMENTATION

| *Configuring SMTP Settings*

Troubleshooting RBAC and OpCo Issues for On-premise Deployments

IN THIS CHAPTER

- [Troubleshooting RBAC and OpCo Issues | 60](#)

Troubleshooting RBAC and OpCo Issues

IN THIS SECTION

- [Authentication Failed for the SP User, Tenant User, or OpCo User | 60](#)
- [Authorization Failed for the SP User, Tenant User, or OpCo User | 62](#)
- [Password to Onboard OpCo is Not Received or has Expired | 63](#)

Authentication Failed for the SP User, Tenant User, or OpCo User

Problem

Description: Service provider (SP) user, tenant user, or OpCo user authentication does not work.

Solution

The CSP administrator, SP administrator, or OpCo administrator must check the authentication method that is set for the SP user, tenant user, or OpCo user respectively.

1. Log in to Administration portal (Global level) and select **Administration > Authentication**.

The Authentication page appears.

2. Check the authentication method for the SP user and the tenant user.
 - If the authentication method is set as **Local**, follow these steps:

- a. Log in to central infrastructure and run the **source/etc/keystone/keystonerc** command.

```
root@centralinfravm:~# source /etc/keystone/keystonerc
```

- b. Type Openstack and press **Enter**.

```
root@centralinfravm:~# openstack
Password:
```

- c. Enter the keystone administrator password.
- d. Execute the **'user list** command and check whether the user name is listed in the keystone.

```
root@centralinfravm:~# openstack user list
```

- e. If the user name is listed, then try to set the password by using the **Forgot Password** link on the Administration Portal login page. See *Resetting Your Password*

If the user name is not available then the CSP administrator or the SP administrator must add the user in CSO. See *Add Service Provider and OpCo Users* .

- f. If the user does not receive an e-mail with the passcode, then set the password through Openstack command.

```
root@centralinfravm:~# openstack user set -password <password> <username>
```

- If the authentication method is set as **Authentication with SSO Server**, then follow these steps:
 - a. Check whether the user name is listed in the SSO server
 - b. Ensure that the SSO server SAML meta data is correct. Navigate to **Administration > Authentication > Single Sign-On Servers > Edit** to check the SAML metadata.
 - c. Ensure that the DNS name of the JCS server is correct in the client server.
 - d. Ensure that the portal URLs are mapped correctly in SSO and CSO server.
 - e. Check whether the same user name is created in CSO with same role.
 - f. Ensure that the tenant name is mapped correctly in CSO and SSO server.

- If the authentication method is set as **Authentication and Authorization with SSO Server**, then follow these steps:
 - a. Ensure that the user has created the user name in correct pattern in CSO for authentication and authorization
 - b. Ensure that the role mapping is created correctly in CSO. A mapping between the roles defined in CSO and the roles defined in external SSO or identity-provider must be provided.
 - c. Check whether the user name and role is created correctly in SSO server.
 - d. Ensure that the service provider and tenant metadata URL for SAML2 of SSO server is configured correctly.

Authorization Failed for the SP User, Tenant User, or OpCo User

Problem

Description: SP user, tenant user, or OpCo user authorization does not work. While accessing some UI pages or features, **Insufficient privileges** error message is displayed.

Solution

To resolve the issue:

- The CSP administrator, SP administrator or Opco administrator must check the user mapped role in CSO.
 1. Log in to Administration Portal and select **Administration > Roles**.
The Roles page appears.
 2. Check whether the role type is predefined or custom role, and then check the privileges assigned to that user. Select the role name and click the pencil icon to view the privileges assigned to the user.
- Check the browser console for Java Script error.
Check the privilege that is causing the error. Check whether that privilege is assigned to the user.
- Using a web browser, check the following REST API output while logging in to CSO with the user account.
 1. Access the Request URL <https://<central ms IP address >/iamsvc/get-user-capabilities>.
The output of the rest API must have the capabilities of the logged in user.
 2. Check whether the UI assigned capabilities are matched in the REST API output.
The JSON output file will list all the capabilities.

If there are any issues with privileges in the JSON output file, contact Juniper Networks Technical Support team.

Password to Onboard OpCo is Not Received or has Expired

Problem

Description: OpCo administrative user did not receive any e-mail with login credentials or OpCo administrator password has expired.

Solution

To resolve the issue:

- Access the URL for Administration Portal. Enter the user name and click **Forget Password** link on the login page to setup the new password. See *Resetting Your Password*.
- If the OpCo administrative user did not receive the e-mail, then use Openstack command to set the password for the OpCo administrator.

```
root@centralinfravm:~# source /etc/keystone/keystonerc
root@centralinfravm:~# openstack user list
root@centralinfravm:~# openstack user set -password <password> <username>
```

- Check the role assignment list to see if there is any issue with the role assignment.
 1. Log in to central infrastructure vm and execute the **source /etc/keystone/keystonerc** command.

```
root@centralinfravm:~# source /etc/keystone/keystonerc
```

2. Log in to Openstack and check the output of user list, role list and role assignment list.

```
root@centralinfravm:~# openstack user list
root@centralinfravm:~# openstack role list
root@centralinfravm:~# openstack role assignment list
```

For further troubleshooting, copy all the log files from the infra VM **/var/log/apache2** into a folder, compress the file in *.zip format and contact Juniper Networks Technical Support team.

RELATED DOCUMENTATION

[Resetting Your Password](#)

Add Service Provider and OpCo Users

Editing the Authentication Method

Configuring a Single Sign-On Server

Troubleshooting Site, Device and Link Issues for On-premise Deployments

IN THIS CHAPTER

- [Troubleshooting Site, Device, and Link Issues | 65](#)

Troubleshooting Site, Device, and Link Issues

IN THIS SECTION

- [Secure OAM Activation Failure | 65](#)
- [Configure SD-WAN Site Failure | 66](#)
- [Device Activation Failure | 66](#)
- [Dual-CPE Activation Failure for NFX Series Devices | 67](#)
- [Dual-CPE Activation Failure for SRX Series Devices | 67](#)
- [Link Switch Event or Performance Metrics is Not Displayed | 68](#)
- [WAN Link Performance Parameters are Not Displayed | 68](#)
- [LTE Interface Issues | 69](#)

Secure OAM Activation Failure

Problem

Description: After entering the activation code , the CPE device status remains in **DEVICE_DETECTED** state; the `csp.tssm_bootstrap-<site-name>` job fails or the job status remains in **In Progress** state for a long time.

Solution

Check whether CSO is reachable or not by executing the following command on the CPE device.

```
user@host > ping <cs0-ip> > source <management-ip-configured-on-loopback-interface>
```

If the ping fails, then check whether the secure OAM tunnels are up by using the following command.

```
user@host > show security ipsec inactive-tunnels
```

If the secure OAM tunnels are not up, verify the connectivity to the OAM hub.

Configure SD-WAN Site Failure

Problem

Description: The configure site operation fails for a spoke site.

Solution

1. Log in to Customer Portal and select **Sites > Site Management**.

The site status must be **Configured**. If the site status is **Configuration Failed**, then the “tssm configure sites” job must have failed.

2. Click **Monitor > Jobs** and check the job details to verify which task has failed.

If the **ship device** task has failed, then CSO has failed to push the required secure OAM tunnel configuration to the hub device.

3. Check the connectivity between CSO and the hub.
4. If there are any other failures, then go to **Sites > Site Management > Site-Name > Configure Site** and review the input provided for configuring the site.

Device Activation Failure

Problem

Description: After entering the activation code, the device status remains in **DEVICE_DETECTED** state for a long time.

Solution

After entering the activation code, the activation window must display the progress of device activation and must indicate that device has been successfully detected. If the device status remains in **DEVICE_DETECTED** state, then follow the steps listed below:

1. Log in to Customer Portal and select **Resources-> Devices**.

The Devices page appears.

2. Check the **Management Status** of the device.

If the management status is **DEVICE_DETECTED**, then the deployment of the stage-1 configuration on device has failed or device has failed to send the **BOOTSTRAP COMPLETE** notification to CSO.

3. Login into the device and verify whether the stage-1 configuration is committed on the device.
4. Verify the connectivity between CSO and the device loop back address.
5. Navigate to **Monitor > Jobs** page and verify the status of **csp.tssm_bootstrap-<site name >** job.
 - If the job is in **successful** state, then ztp job will be triggered.
 - If the job is in **in-progress** state, then the CPE device failed to establish the connection over the secure OAM tunnel.
6. If device failed to establish the connection within an hour, or if the **csp.tssm_bootstrap-<site name >** job fails, then check the bootstrap task details.
7. Once the connectivity issue is resolved, navigate to **Resources > Devices** and activate the device.

The **csp.tssm_ztp-<site name >** job must be successful state. If the job failed, check the task details verify which task has failed.

Dual-CPE Activation Failure for NFX Series Devices

Problem

Description: ZTP Job failed for dual CPE NFX Series devices.

Solution

For a site with dual CPE NFX Series devices, two ZTP jobs, namely, **csp.tssm_ztp-<site-name> _cpe0** and **csp.tssm_ztp-<site-name> _cpe1** are created. One ZTP job is created per each node.

While the jobs are still in progress and after the Gateway Router (GWR) is spawned successfully, two more jobs, namely, **form_device_cluster** are created per each node for cluster formation.

Log in to Administration Portal and select **Monitor > Jobs** to view the **form_device_cluster** job. If cluster formation fails, the **form_device_cluster** job and the **csp.tssm_ztp-<site-name> _cpe0**, **csp.tssm_ztp-<site-name> _cpe1** jobs are reported as failure.

For any cluster formation job failure, check the logs from the device at **/tmp/cluster_gwr.log**.

For further troubleshooting, collect the logs and output results and contact Juniper Networks SRE team.

Dual-CPE Activation Failure for SRX Series Devices

Problem

Description: ZTP Job failed for dual SRX Series devices

Solution

For a site with dual CPE SRX Series devices, two ZTP jobs, namely , **csp.tssm_ztp-<site-name>_cpe0** and , **csp.tssm_ztp-<site-name>_cpe1** are created. One ZTP job is created per each node.

In case of dual SRX Series devices, as a pre-requisite, the chassis cluster is already formed manually before starting the device activation. The **csp.tssm_ztp-<site-name>_cpe1** job will report success quickly, and the actual ztp progress can be tracked through the **csp.tssm_ztp-<site-name>_cpe0** job. In case of any failure, refer to ZTP job task details.

Link Switch Event or Performance Metrics is Not Displayed

Problem

Description: Link switch event is not displayed in the UI

Solution

Check whether the device is able to reach southbound load balancer VM (SBLB VM) and the time is synchronized with the NTP server.

```
root@gwr.spoke-nfx> show system uptime
Current time: 2019-03-04 15:37:46 IST
Time Source:  NTP CLOCK
System booted: 2019-02-28 15:13:49 IST (4d 00:23 ago)
Protocols started: 2019-02-28 15:13:50 IST (4d 00:23 ago)
Last configured: 2019-03-04 14:58:58 IST (00:38:48 ago) by csp
3:37PM up 4 days, 24 mins, 1 user, load averages: 0.42, 0.30, 0.26
```

Even when the link switch is successful on the device, it may not be indicated in the UI because of the missing syslog events. Link switch event in UI is indicated based on the **APPQOE_BEST_PATH_SELECTED** syslog with reason as **sla violated** that is received from CPE device.

Log in to Customer Portal and select **Monitor > Device Events** to view all the syslogs that are received from the CPE device To filter the **APPQOE_BEST_PATH_SELECTED** events, use the following query:
Event Name = APPQOE_BEST_PATH_SELECTED and **Reason = sla violated**.

WAN Link Performance Parameters are Not Displayed

Problem

Description: WAN link performance parameters, such as latency, packet loss, E2E delay, jitter, and throughput are not displayed in the UI.

Solution

Check whether the device is able to reach southbound load balancer VM (SBLB VM) and the time is synchronized with the NTP server.

```
root@gwr.spoke-nfx> show system uptime
Current time: 2019-03-04 15:37:46 IST
Time Source: NTP CLOCK
System booted: 2019-02-28 15:13:49 IST (4d 00:23 ago)
Protocols started: 2019-02-28 15:13:50 IST (4d 00:23 ago)
Last configured: 2019-03-04 14:58:58 IST (00:38:48 ago) by csp
3:37PM up 4 days, 24 mins, 1 user, load averages: 0.42, 0.30, 0.26
```

Login to Customer Portal and select **Sites > Site Management > Site-Name > WAN** tab to view the WAN link performance.

- The WAN link performance details for latency, packet loss, E2E delay, and jitter are retrieved from **APPQOE_ACTIVE_SLA_METRIC_REPORT** syslog. To filter the **APPQOE_ACTIVE_SLA_METRIC_REPORT** events, use the following query:

Event Name = APPQOE_ACTIVE_SLA_METRIC_REPORT and Site = <site-name >.

- The WAN link performance details for throughput is retrieved from **APPTRACK_ACTIVE_SLA_METRIC_REPORT** syslog. To filter the **APPTRACK_ACTIVE_SLA_METRIC_REPORT** events, use the following query:

Event Name = APPTRACK_SESSION_CLOSE and Site = <site-name>.

LTE Interface Issues

Problem

Description: LTE interface is not receiving the IP address.

Solution

- Check the data validity of the SIM using the mobile device.
- Check the LTE module connection status to ensure that there is adequate mobile signal strength.

user@host>show modem wireless network cl-1/1/0

```
LTE Connection details
Connected time: 2880
IP: 192.12.219.210
Gateway: 192.12.219.209
DNS: 192.123.123.123
IPv6: ::
Gatewayv6: ::
```



```
DNSv6: ::  
Input bps: 0  
Output bps: 0  
Bytes Received: 1952  
Bytes Transferred: 2164  
Packets Received: 10  
Packets Transferred: 20  
Wireless Modem Network Info  
Current Modem Status: Connected  
Current Service Status: Normal  
Current Service Type: PS  
Current Service Mode: LTE  
Current Band: B3  
...
```

Check the **Current Modem Status**, **Current Service Status**, **Current Service Type**, and **Current Service Mode** fields.

- For NFX150 device, ensure that the external antenna is connected properly.

Recovering an Installation

IN THIS CHAPTER

- [CSO Disaster Recovery](#) | 71

CSO Disaster Recovery

In case of any failures you can recover CSO Release 5.1.2. To recover CSO Release 5.1.2 you must have already taken a backup and saved the backup file.

To recover CSO Release 5.1.2:

1. Based on the hypervisor you are using, do one of the following:

- If you are using KVM as the hypervisor:

- a. Copy the CSO 5.1.2 backup folder to the bare metal server.
- b. From the backup folder, copy of the **topology.conf** file to the **Contrail_Service_Orchestration_5.1.2/topology/** folder.

For example:

```
cp
/root/backups/backupfordr/2020-06-19T17:27:05/config_backups/_topology.conf
/root/Contrail_Service_Orchestration_5.1.2/topology/
```

- c. Provision the VMs. For information on provisioning KVM hypervisor, see *Provision VMs on Contrail Service Orchestration Servers in CSO Installation and Upgrade Guide*.
- d. Copy the backup folder file from the bare metal server to the startupserver_1 VM.

```
user@server>scp csoVersion.tar.gz root@<startupserver_>IP>:/root/
```

- e. Log in to the startupserver_1 VM as the root user.
- f. Run the **get_vm_details.sh**script to find the IP address of the startupserver_1 VM.
- g. Access the shell of the startupserver_1 VM.
- h. Expand the installer package.

```
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.

Go to [Step 2](#).

- If you are using ESXi as the hypervisor:

- a. Copy the backup folder to the startupserver_1 VM.
- b. From the backup folder, copy of the **topology.conf** file to the **Contrail_Service_Orchestration_5.1.2/topology/** folder in the startupserver_1 VM.

For example:

```
cp
/root/backups/backupfordr/2020-06-19T17:27:05/config_backups/_topology.conf
/root/Contrail_Service_Orchestration_5.1.2/topology/
```

- c. Log in to the startupserver_1 VM as the root user.
- d. Run the get_vm_details.sh script to find the IP address of the startupserver_1 VM
- e. Access the shell of the startupserver_1 VM.
- f. Expand the installer package.

```
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.

2. Run the **deploy.sh** command.

```
root@host:~/Contrail_Service_Orchestration_5.1.2./ deploy.sh
```

Use the interactive script to create configuration files for the environment specific topology. Because you have taken a backup, the answers are prepopulated for each question in the interactive script.

NOTE: Ensure that you specify *no* for the following question:

Whether this is an upgrade setup?

3. Restore the data by using the **cso_backupnrestore** script.

```
backuppath is 5.1 backup path from above for
ex. /backups/backup512/2020-01-09T18:47:08/#cso_backupnrestore -b restore -s
backuppath -t '*' -c 'mariadb'
#cso_backupnrestore -b restore -s backuppath -t '*' -c 'zookeeper'
#cso_backupnrestore -b restore -s backuppath -t '*' -c 'elasticsearch'
#cso_backupnrestore -b restore -s backuppath -t '*' -c 'arangodb'
#cso_backupnrestore -b restore -s backuppath -t '*' -c 'icinga'
#cso_backupnrestore -b restore -s backuppath -t '*' -c 'swift'
#cso_backupnrestore -b restore -s backuppath -t '*' -c 'mariadb'
```

If the restore procedure fails for any of the above components, you must retry to restore only those components.

4. Restore Cassandra.

To restore Cassandra:

- a. Run `python /usr/local/bin/cassandra_token_update.py` to update token and cluster details.

```
root@server:~# python /usr/local/bin/cassandra_token_update.py

Enter the 5.1 backup path: /backups/drbackup0625209pm/2020-06-26T06:09:50
Please enter yes to process [yes/no]:yes
```

- b. Restore Cassandra.

```
cso_backupnrestore -b restore -s backuppath -t '*' -c 'cassandra' -r
'yes'
```

5. Synchronize the data between nodes.

```
cso_backupnrestore -b nodetool_repair
```

6. Copy the certificate from the backup folder to SDN-based load balancing (SBLB) HA Proxy.

```
salt-cp -G "roles:haproxy_confd_sblb"
/root/backups/config_backups/haproxycerts/minions/minions/csp-regional-sblb1.TNL2OQ.
regional/files/etc/pki/tls/certs/ssl_cert.pem /etc/pki/tls/certs salt-cp
-G "roles:haproxy_confd_sblb"
/root/backups/config_backups/haproxycerts/minions/minions/csp-regional-sblb1.TNL2OQ.
regional/files/etc/pki/tls/certs/ssl_cert.crt /etc/pki/tls/certs
```

7. Restart the SBLB HA Proxy.

```
salt -C "G@roles:haproxy_confd_sblb" cmd.run "service haproxy restart"
```

8. Copy the certificate from the backup folder to Central HA Proxy.

```
salt-cp -G "roles:haproxy_confd"
/root/backups/config_backups/haproxycerts/minions/minions/csp-central-lbvml.HBLGHQ.
central/files/etc/pki/tls/certs/ssl_cert.pem /etc/pki/tls/certs salt-cp -G
"roles:haproxy_confd"
/root/backups/config_backups/haproxycerts/minions/minions/csp-central-lbvml.HBLGHQ.
central/files/etc/pki/tls/certs/ssl_cert.crt /etc/pki/tls/certs
```

9. Restart the Central HA Proxy.

```
salt -C "G@roles:haproxy_confd" cmd.run "service haproxy restart"
```

10. Run the following commands on installer VM to update the Nginx certificates.

```
kubectl get secret -n central | grep cso-ingress-tls content_copyzoom_out_map
```

```
cso-ingress-tls kubernetes.io/tls 2 17d
```

```
kubectl delete secret cso-ingress-tls -n central kubectl create secret tls
cso-ingress-tls --key
/root/backups/config_backups/haproxycerts/minions/minions/csp-central-lbvml.5R8JKN.
central/files/etc/pki/tls/certs/ssl_cert.key --cert
/root/backups/config_backups/haproxycerts/minions/minions/csp-central-lbvml.5R8JKN.
central/files/etc/pki/tls/certs/ssl_cert.crt -n central
```

11. Run the **upgrade.sh** script.

12. Restore the SD-WAN and security reports.

```
cso_backupnrestore -b restore -s backuppath -t '*' -c 'swift_report' -r
'yes'
```

13. Restart all fmpm-provider-core pods by deleting the existing pods.

```
root@startupserver1:~# kubectl get pods -n central|grep fmpm-provider-core
```

```
content_copyzoom_out_map
csp.csp-fmpm-provider-core-647ff6598d-4qxxd 1/1 Running 0
2d18h
csp.csp-fmpm-provider-core-647ff6598d-94wdx 1/1 Running 0
2d18h
csp.csp-fmpm-provider-core-647ff6598d-dt6vj 1/1 Running 0
2d18h
csp.csp-fmpm-provider-core-647ff6598d-hbnw2 1/1 Running 0
2d18h
csp.csp-fmpm-provider-core-647ff6598d-mx8fn 1/1 Running 0
2d18h
csp.csp-fmpm-provider-core-647ff6598d-zd2zt 1/1 Running 0
2d18h
```

```
root@startupserver1:~# kubectl delete pod
csp.csp-fmpm-provider-core-647ff6598d-4qxxd
csp.csp-fmpm-provider-core-647ff6598d-94wdx
csp.csp-fmpm-provider-core-647ff6598d-dt6vj
csp.csp-fmpm-provider-core-647ff6598d-hbnw2
csp.csp-fmpm-provider-core-647ff6598d-mx8fn
csp.csp-fmpm-provider-core-647ff6598d-zd2zt -n central
```

14. Restore Contrail Analytics Node (CAN) database.

```
./python.sh upgrade/migration_scripts/common/can_migration.py
```

The **analyticsdb** backup files are located at

/root/backups/backup411/2020-05-21T00:43:50/central/can/can<x>

```
ssh root@<new-can-ip-[123]>
```

```
docker cp 0000/mc-* analyticsdb:/root
```

```
docker exec -it analyticsdb bash
```

```
mv /root/mc-*
```

```
/var/lib/cassandra/data/ContrailAnalyticsCql/statstablebystntagv3-c5e9b4c056f711ea8a948909f467ce30
```

#The path may be different based on uuid #The path may be different based on uuids

```
cd
```

```
/var/lib/cassandra/data/ContrailAnalyticsCql/statstablebystntagv3-c5e9b4c056f711ea8a948909f467ce30
```

```
chown -R cassandra:cassandra *
```

```
nodetool refresh -- ContrailAnalyticsCql statstablebystntagv3
```

After a successful upgrade, CSO Release 5.1.2 is functional and you can log in to the Administrator Portal and the Customer Portal.