

Contrail Service Orchestration Monitoring and Troubleshooting Guide

Published
2020-12-01

Release
4.1

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

4.1

Copyright © 2020 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 | vi

Documentation and Release Notes | vi

Documentation Conventions | vi

Documentation Feedback | ix

Requesting Technical Support | ix

Self-Help Online Tools and Resources | x

Creating a Service Request with JTAC | x

1

Monitoring Contrail Service Orchestration

Monitoring Infrastructure Services and Microservices | 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 | 11

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

CSO Database Backup and Restore | 11

Configuration | 13

Major Components | 13

Operations | 13

Command Usage | 15

Backup and Restore Examples | 16

Introduction to Service and Infrastructure Monitor Application | 18

Service and Infrastructure Monitor Overview | 18

Accessing the Service and Infrastructure Monitor GUI | 19

Monitoring Network Services | 20

Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 21

Monitoring Microservices | 25

Monitoring Microservices and Their Host VMs | 26

Monitoring Physical Servers | 28

Troubleshooting Contrail Service Orchestration Issues

Troubleshooting Login Issues | 31

Troubleshooting Login Issues | 31

Administration Portal IP Address Is Not Reachable | 31

Administration Portal User Interface Is Not Reachable | 33

Resetting the Password without E-mail Access | 34

Troubleshooting POPs, Tenants, and Devices Issues | 37

Troubleshooting POPs, Tenants, and Devices Issues | 37

Failure While Creating a Hub, Site, or Tenant | 37

Base Configuration for CPE Activation | 38

Troubleshooting Site Activation Issues | 40

Troubleshooting Site Activation Issues | 40

Prerequisites to Activate a Site | 40

Activation Failure for a Hub site | 41

Activation Failure for a Spoke Site | 43

Certificate File Location and Activation Code for an SRX300 Device | 46

Troubleshooting Image, License, and Policy Deployment Issues | 48

Troubleshooting Image, License, and Policy Deployment Issues | 48

Image Upload Failure | 48

Firewall Application Policy Deployment Failure | 49

Traffic from Spoke Sites Are Dropped or Are Not Reaching Internet or Destination | 51

Missing Data in Application Visibility Page | 51

Link Switch Does Not Happen During SLA Violation | 52

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 | 54

Troubleshooting CSO Installation Issues | 54

Salt Key Issue During CSO Installation | 54

TimeZone Error | 56

SSL Handshake Failure | 56

Missing Interface on CSO VM | 57

Troubleshooting SMTP Issues | 58

Troubleshooting SMTP Issues | 58

Basic Configuration for SMTP Server | 58

Basic Configuration for AWS CSO Installations | 60

Troubleshooting RBAC and OpCo Issues | 62

Troubleshooting RBAC and OpCo Issues | 62

Authentication Failed for the SP User, Tenant User, or OpCo User | 62

Authorization Failed for the SP User, Tenant User, or OpCo User | 64

Password to Onboard OpCo is Not Received or has Expired | 65

Troubleshooting CSO Release 4.1 Issues | 67

Troubleshooting CSO Release 4.1.0 Issues | 67

Secure OAM Activation Failure | 67

Configure Site Failure | 68

Device Activation Failure | 68

Dual-CPE Activation Failure for NFX Series Devices | 69

Dual-CPE Activation Failure for SRX Series Devices | 70

Link Switch Event or Performance Metrics is Not Displayed | 70

WAN Link Performance Parameters are Not Displayed | 71

LTE Interface Issues | 71

About the Documentation

IN THIS SECTION

- Documentation and Release Notes | vi
- Documentation Conventions | vi
- Documentation Feedback | ix
- Requesting Technical Support | ix

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 vii](#) 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 vii 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 | 2

Backup and Restore of Contrail Service Orchestration | 11

Introduction to Service and Infrastructure Monitor Application | 18

Monitoring Infrastructure Services and Microservices

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 | 19](#)

[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](#)

Generating and Encrypting Passwords for Infrastructure Components

Installing and Configuring 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

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 | 13](#)
- [Command Usage | 15](#)
- [Backup and Restore Examples | 16](#)

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 installer virtual machine (installer-vm). See the [“Operations” on page 13](#) for details.

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, and ETCD. 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.

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 can not 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 on small, medium, and large deployments both with or without 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 installer-vm. • 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 installer-vm machine. The user in charge of the operations logs onto the installer-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 installer-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 installer-vm. This location can not be changed. The storage for this location can be local to the installer-vm or it can be located on a Storage Area Network (SAN).

For restore operations, the same **cso_backupnrestore** command is used with different options as described in [Table 4 on page 15](#) below. When restoring from a backup, CSO puts itself into maintenance mode so that no changes can be made. 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
-m	Put CSO in maintenance mode prior to backup. Only valid in combination with backup argument for the -b option. [Default no]	yes or no
-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. By default, this option sets the -m option to no.	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) -m yes option overrides default and puts CSO into maintenance mode for cron-based backups.

Backup and Restore Examples

Requirements

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

The following commands must be run at the command line interface of the installer-vm of CSO. The location and access credentials needed to access the installer-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 **/bakups/DAILY/<timestamp>/** directory. The timestamp directory is created when the backup starts.

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

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
```

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

IN THIS CHAPTER

- [Service and Infrastructure Monitor Overview | 18](#)
- [Accessing the Service and Infrastructure Monitor GUI | 19](#)
- [Monitoring Network Services | 20](#)
- [Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 21](#)
- [Monitoring Microservices | 25](#)
- [Monitoring Microservices and Their Host VMs | 26](#)
- [Monitoring Physical Servers | 28](#)

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 | 20](#)

[Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 21](#)

[Monitoring Microservices | 25](#)

[Monitoring Microservices and Their Host VMs | 26](#)

[Monitoring Physical Servers | 28](#)

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](#) | 18

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 20](#) 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](#) | 21

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 22](#) 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](#) | 20

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 26](#) 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](#) | 26

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 27](#) 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](#) | 25

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 29](#) 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](#) | 18

2

PART

Troubleshooting Contrail Service Orchestration Issues

Troubleshooting Login Issues | **31**

Troubleshooting POPs, Tenants, and Devices Issues | **37**

Troubleshooting Site Activation Issues | **40**

Troubleshooting Image, License, and Policy Deployment Issues | **48**

Troubleshooting CSO Installation Issues | **54**

Troubleshooting SMTP Issues | **58**

Troubleshooting RBAC and OpCo Issues | **62**

Troubleshooting CSO Release 4.1 Issues | **67**

Troubleshooting Login Issues

IN THIS CHAPTER

- [Troubleshooting Login Issues | 31](#)

Troubleshooting Login Issues

IN THIS SECTION

- [Administration Portal IP Address Is Not Reachable | 31](#)
- [Administration Portal User Interface Is Not Reachable | 33](#)
- [Resetting the Password without E-mail Access | 34](#)

Administration Portal IP Address Is Not Reachable

Problem

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

Solution

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.

Administration Portal User Interface Is Not Reachable

Problem

Description: CSO Administration Portal IP address is reachable, but the user interface is not reachable.

Solution

Check whether the firewall in the path is blocking port 443. Also, check whether the CSO Administration Portal performance-optimized data center (POD) and other PODs are running. You can check the PODs in the Icinga Web UI or in CSO central microservices virtual machine.

- To check whether the firewall in the path is blocking port 443:

```
user@host-csp-build:~$ telnet 192.213.10.54 443
```

```
Trying 192.213.10.54...
Connected to 192.213.10.54.
Escape character is '^['.
```

- To verify the POD status in Icinga:
 1. Log in to `http://central-ms-vm-IP-Address:1947/icingaweb2`.
 2. Enter the user name **icinga** and the password that is generated during CSO installation.
 3. Select **Infrastructure > CSP Microservices > Central_Services > Central MS IP**.
The PODs are displayed along with their running status.
 4. Verify whether the Administration Portal POD status is running (indicated by green).

To verify the Administration Portal POD status in the CSO central microservices virtual machine.

1. Log in to the CSO central microservices virtual machine, and execute the following command:

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

```
csp.admin-portal-ui-2886357385-brtjg      1/1    Running    1    17h
root@centralmsvm:~#
```

Verify whether the Administration Portal POD and the service is in running state, which is indicated by 1/1. In some cases, where the microservices are clubbed together in a POD, the running status is indicated by 2/2, or 3/3.

2. Execute **kubectl get pods** to get the status of all POD running in the central microservices virtual machine. For example:

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

```
csp.admin-portal-ui-2886357385-brtjg 1/1 Running 1 17h
```

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

```
csp.admin-portal-ui-2886357385-brtjg 1/1 Running 1 17h
```

3. Check if any other POD is not in running state or if the ready state is 0/1 instead of 1/1. Then check the corresponding POD log by executing the **kubectl logs -f pod-name** command.

```
root@centralmsvm:~# kubectl logs -f csp.admin-portal-ui-2886357385-brtjg -n central
```

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

Resetting the Password without E-mail Access

Problem

Description: User is unable to log in to **cspadmin** account. The error message “Login failed. Check your username and password” is displayed. Resetting the password requires access to e-mail servers, but user does not have access to smtp servers.

Solution

The **cspadmin** user password is generated by the system for the first time. If the user has changed it from the UI, then the user must enter the changed password.

To reset the password:

1. Execute the following commands:

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

ID	Name
0e83d9a9073e44c79cdd3a51485fff8e	swift
370a6ae5cda24e60ba835a4a02b42a79	admin
569463999d4941458a102b963fb10b36	abc@juniper.net
d4507d11eb164dd8a14883e384027d7a	cloud_admin@ucpe.com
f227c8c92b0648e7bc9d271e29f53b93	abc@xyz.net
f6d6f551f7614c33b046b87c3fb123f6	cspadmin
fd38562004754526b624abf70d4b7388	retail_admin@retail.com

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

The system prompts for the password. Enter the keystone administrator password.

2. If the cspadmin account is locked, then execute the following script to unlock the account.

```
root@installervm:~# /root/cso_dl/<cso-folder>./python.sh scripts/unlock_account.py
<admin portal UI IP address> <keystone admin password>
```

3. Check whether the keystone is running in the infrastructure virtual machine (infravm) with service apache2 status.

```
root@centralinfravm:~# kubectl get pods -n central | grep iamsvc
```

4. Check whether the pod is in **Running** state. If the pod is not running, run the following command to delete the pod.

```
root@centralmsvm:~# kubectl delete pods <pod-name> -n central
```

5. Check the log message to see if the RabbitMQ connection is down. If the connection is down, then restart the RabbitMQ server.

```
root@centralinfravm:~# service rabbitmq-server restart
```

6. Navigate to msvm and run the **kubectl delete pods --all --force --grace-period=0** command. to re-establish the connection with RabbitMQ.

```
root@centralmsvm:~# kubectl delete pods --all --force --grace-period=0
```

RELATED DOCUMENTATION

| *Resetting the Password*

Troubleshooting POPs, Tenants, and Devices Issues

IN THIS CHAPTER

- [Troubleshooting POPs, Tenants, and Devices Issues | 37](#)

Troubleshooting POPs, Tenants, and Devices Issues

IN THIS SECTION

- [Failure While Creating a Hub, Site, or Tenant | 37](#)
- [Base Configuration for CPE Activation | 38](#)

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:~# kubect
```

```
l 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:~# kubect
```

```
l 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:~# kubect
```

- 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

Troubleshooting Site Activation Issues

IN THIS CHAPTER

- [Troubleshooting Site Activation Issues | 40](#)

Troubleshooting Site Activation Issues

IN THIS SECTION

- [Prerequisites to Activate a Site | 40](#)
- [Activation Failure for a Hub site | 41](#)
- [Activation Failure for a Spoke Site | 43](#)
- [Certificate File Location and Activation Code for an SRX300 Device | 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 to hub underlay reachability for IPsec/GRE tunnels or the SSH connection and vice versa.
- Check the hub or spoke to CSO (regional MS) reachability or the SSH connection and vice versa.

- Check the hub to CSO (regional MS) reachability or the SSH connection and vice versa.
- 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*

Activation Failure for a Hub site

Problem

Description: A failure occurred when activating a hub site.

Solution

1. Check the job logs in the Administration Portal for the activation failure and the reason for the failure.
 - a. Log in 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 it.
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.
2. If ZTP is enabled in the hub device template, then ensure that the hub device or image supports the phone-home feature. If the feature is not supported, then upgrade the software image.
 - If you need to disable ZTP in the hub device template, log in to Administration Portal and select **Resources > Device Templates > Template Name > Edit Device Template > Template Settings**. Disable **ZTP_ENABLED** option.
 - If you need to copy the stage-1 configuration to the hub, then log in to Administration Portal and select **Resources > Cloud Hub Devices > Stage 1 Config** and copy the configuration.
3. Check the outbound SSH connection between the hub and the regional microservices virtual machine on port 7804.

Log in to the CSO regional microservices virtual machine and execute the following command

```
root@regionalmsvm:~#netstat -anp | grep 7804
```

```
tcp        0      0 0.0.0.0:7804          0.0.0.0:*             LISTEN
```

```

    1254/haproxy
tcp      0      0 192.0.2.0:7804      192.0.3.0:7310      ESTABLISHED
1254/haproxy  >>> Spoke
tcp      0      0 192.0.2.0:7804      192.0.4.0:14632     ESTABLISHED
1254/haproxy  >>> Hub
root@regionalmsvm:~#

```

- If the outbound SSH connection is not established between the hub and the regional microservices virtual machine,
 - Ensure that TCP port 7804 is not blocked in the path.
 - Check the reachability between the hub and the regional microservices virtual machine,. Check whether the hub device can establish an SSH connection with the regional microservices virtual machine and vice versa.
 - View the detailed failure logs in the Kibana dashboard <http://regional-infra-IP-Address:5601> or log in to regional microservices virtual machine and execute the following command to view the detailed failure logs.

```
root@regionalmsvm:~#kubectl get pods -n regional | grep activation
```

```
csp.csp-activation-service-1888452022-fv1vt    1/1    Running    1    19h
```

```
root@regionalmsvm:~#kubectl logs -f csp.csp-activation-service-1888452022-fv1vt -n regional
```

- If the outbound SSH connection is established, then check if the configurations are pushed to the device.

Log in to the Administration Portal. Select **Monitor > Jobs**. Click the ZTP failure log and verify the configuration deployment task.

To view the detailed log, use the Kibana dashboard <<http://central infra ip:5601>> or log in to the central microservices virtual machine and execute the following command.

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

```
csp.csp-cms-central-2820689874-gvjbh          1/1      Running
1          19h
```

```
csp.csp-cms-central-core-2224266535-kmplk     1/1      Running
1          19h
```

```
root@centralmsvm:~# kubectl logs -f csp.csp-cms-central-core-2224266535-kmplk
-n central
```

Verify that the configurations are pushed successfully to the device.

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

Activation Failure for a Spoke Site

Problem

Description: Activation failed for a spoke site.

Solution

1. Check the job logs in the Administration Portal for the activation failure and the reason for the failure.

- a. Log in 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

2. Check the Internet reachability. If Juniper Networks redirect server is used for CPE ZTP or activation, then ensure that the CPE device can establish a connection to the Internet. The CSO activation server IP address (regional microservices virtual machine IP address for large deployments, central microservices virtual machine IP address for medium deployments and central microservices IP address for small deployments(as there are no virtual machine IP for small deployments)), activation server certificate and the CPE serial numbers are configured in the Juniper Networks redirect server.

Copy the activation server certificate for each deployments using the following command,

For large deployments, get the certificate from regional microservice virtual machine IP

```
root@regionalmsvm:~#ls -l /etc/pki/tls/certs/ssl_cert.crt
```

```
-rw-r--r-- 1 root root 1306 Dec  2 10:08 /etc/pki/tls/certs/ssl_cert.crt
root@regionalmsvm:~#
```

For medium deployments, get the certificate from central virtual machine IP

```
root@centrallbvm2# ls -l /etc/pki/tls/certs/ssl_cert.crt
```

For small deployments, get the certificate from central microservice

```
root@centralmsvm:~#ls -l /etc/pki/tls/certs/ssl_cert.crt
```

```
-rw-r--r-- 1 root root 1306 Dec 2 10:08 /etc/pki/tls/certs/ssl_cert.crt
root@centralmsvm:~#
```

NOTE: When custom generated certificates are used for CSO installation (via UI Installer) then

- The hostname for central/regional microservice virtual machine varies according to the common name of the certificate. For example, If common name available in the custom generated certificate is "cso-central-medium.englab.juniper.net" then the hostname for the centralmsvm will be "root@cso-central-medium"
- The certificate will not be in crt format. Instead convert the pem to crt format in path `/etc/pki/tls/certs/ssl_cert.pem`

3. If Juniper Networks redirect server is not used for CPE ZTP or activation, then configure the phone-home server in the CPE device and copy the certificate to the CPE device.

For large deployments

```
root@cpe-srx #show system phone-home
```

```
server https://regional-ms-ip;
ca-certificate-file /root/ssl_cert.crt;
```

For small and medium deployments

```
root@cpe-srx #show system phone-home
```

```
server https://central-ms-ip:444;
ca-certificate-file /root/ssl_cert.crt
```

4. Check the outbound SSH connection between the spoke and the microservices virtual machine on port 7804.

For large deployments

```
root@regionalmsvm:~#netstat -anp | grep 7804
```

```
tcp        0      0 0.0.0.0:7804          0.0.0.0:*              LISTEN
1254/haproxy
tcp        0      0 192.2.2.2:7804        192.3.3.3:7310          ESTABLISHED
1254/haproxy  >>> Spoke
tcp        0      0 192.2.2.2:7804        192.4.4.4:14632          ESTABLISHED
1254/haproxy  >>> Hub
root@regionalmsvm:~#
```

For small and medium deployments

```
root@centralmsvm:~#netstat -anp | grep 7804
```

```
tcp        0      0 0.0.0.0:7804          0.0.0.0:*              LISTEN
1254/haproxy
tcp        0      0 192.2.2.2:7805        192.3.3.3:7310          ESTABLISHED
1254/haproxy  >>> Spoke
tcp        0      0 192.2.2.2:7805        192.4.4.4:14632          ESTABLISHED
1254/haproxy  >>> Hub
root@centralmsvm:~#
```

- If the outbound SSH connection is not established between the spoke and the regional microservices virtual machine,
 - Ensure that TCP port 7804 is not blocked in the path.
 - Check the reachability between the spoke and the regional microservices virtual machine. The spoke device (JDM console) must establish an SSH connection with the regional microservices virtual machine.
 - View the detailed failure logs in the Kibana dashboard <http://regional infra IP-Address:5601> or log in to the regional microservices virtual machine and execute the following command.

```
root@regionalmsvm:~#kubectl get pods -n regional | grep activation
```

```
csp.csp-activation-service-1888452022-fvlvt    1/1    Running    1
19h
```

```
root@regionalmsvm:~# kubectl logs -f
csp.csp-activation-service-1888452022-fv1vt -n regional
```

- For NFX250 device, check the recommended vSRX image uploaded in CSO. Check if the vSRX image is uploaded to the CPE or NFX device. If there is any failure, then check the latency, download or upload speed between CPE device and the regional microservices virtual machine.

Log in to the Administration Portal and select **Resources > Images**.

- If the outbound SSH connection is established, then check that the configurations are pushed to the device.

Log in to the Administration Portal. Select **Monitor > Jobs**. Click the activation failure log and verify the configuration deployment task.

To view the detailed log, use the Kibana dashboard <http://central infra ilP-Address:5601> or log in to central microservices virtual machine and execute the following command.

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

```
csp.csp-cms-central-2820689874-gvjbh          1/1      Running
1          19h
csp.csp-cms-central-core-2224266535-kmplk      1/1      Running
1          19h
root@centralmsvm:~# kubectl logs -f csp.csp-cms-central-core-2224266535-kmplk
-n central
```

Verify the configuration in the log and check that the configurations are pushed successfully to the device.

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

Certificate File Location and Activation Code for an SRX300 Device

Problem

Description: User was unable to perform ZTP on an SRX300 device that acts as both an SD-WAN and a distributed CPE device. Specify the **cert** file location (to copy the certificate file from a phone-server) and the activation command.

Solution

You can paste the certificate in any directory on the system but you must reference the same location as shown in the following configuration:

```
system {  
    host-name spoke0;  
    root-authentication {  
        encrypted-password "$ABC123"; ## SECRET-DATA  
    }  
    phone-home {  
        traceoptions {  
            file phc.log size 10m;  
            flag all;  
        }  
        server https://192.1.1.9;  
        ca-certificate-file /var/ssl_cert.crt;  
    }  
}
```

You can use the **test phone-home server-authentication-code 123456** command to enter the activation code on an SRX300 device. Alternatively, you can log in to Customer Portal and enter the activation code from the **Sites > Sites Management** page.

Troubleshooting Image, License, and Policy Deployment Issues

IN THIS CHAPTER

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

Troubleshooting Image, License, and Policy Deployment Issues

IN THIS SECTION

- [Image Upload Failure | 48](#)
- [Firewall Application Policy Deployment Failure | 49](#)
- [Traffic from Spoke Sites Are Dropped or Are Not Reaching Internet or Destination | 51](#)
- [Missing Data in Application Visibility Page | 51](#)
- [Link Switch Does Not Happen During SLA Violation | 52](#)
- [SLA Violation-Original Link Recovered After SLA Violation | 52](#)
- [All WAN links are uP But Not All Links Are Utilized | 52](#)

Image Upload Failure

Problem

Description: Image upload operation failed.

Solution

1. Check the job logs in the Administration Portal for the image upload failure and the reason for the failure.
 - a. Log in to the Administration Portal and select **Monitor > Jobs**
The Jobs page is displayed.

- b. Select the log related to image upload failure and click the Detailed View icon that appears before the log.

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

2. Check latency, download or upload bandwidth, between the UI client machine(remote machine) and the central microservices virtual machine. You can use any third-party tool to check these details.
3. Try to upload the image through the CLI. You can execute the CLI configuration statement in any machine that is reachable to the central microservices virtual machine or directly in the central microservices virtual machine for a quick upload. A sample CLI configuration statement is listed below:

```
curl -v -F "imagefile=@media-vsrx-vmdisk-15.1X49-D120.3.qcow2" -H
"x-auth-token:b95980967d71474cb169443c75525caf" -F "cname=vsrx-vmdisk-15.1.qcow2"
-F "device_family=juniper-vsrx" -F "vendor=juniper" -F "major_version=1" -F
"minor_version=1" -F build_num="X53-D47.3" -F "supported_platform=NFX250" -F
"image_type=VNF_IMAGE" -k https://central-ms-vm-ip/ims-central/upload_image_file
```

4. Manually upload the image to the NFX device and update the image location in the NFX CPE device template.
5. Ensure that the image name is vsrx-vmdisk-15.1.qcow2 unless the vSRX image name has changed in the NFX device template.

Firewall Application Policy Deployment Failure

Problem

Description: The firewall application policy failed to deploy.

Solution

1. Check the job logs in the Administration Portal for the signature installation failure and the reason for the failure.
 - a. Log in to the Administration Portal and select **Monitor > Jobs**

The Jobs page is displayed.
 - b. Select the log related to the failure and click the Detailed View icon that appears before it.

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

2. Check if the CPE device is up and the outbound SSH connection is active.

Log in to Administration portal, select **Monitor > Alerts and Alarm > Alerts**

or check the outbound SSH connection in the regional microservices virtual machine.

```
root@regionalmsvm:~#netstat -anp | grep 7804
```

```
tcp        0      0 0.0.0.0:7804          0.0.0.0:*            LISTEN
1254/haproxy
tcp        0      0 192.0.0.1:7804        192.0.0.2:7310        ESTABLISHED
1254/haproxy  >>> Spoke
tcp        0      0 192.0.0.1:7804        192.0.0.3:14632       ESTABLISHED
1254/haproxy  >>> Hub
root@regionalmsvm:~#
```

3. Check that the application signature is successfully installed on the device.

In the Administration Portal, select **Administration > Signature Database**, and click **Install on device** to verify the installation status.

4. Check that the rendered configurations do not show any user input error and that they are pushed to the device. For a detailed log, check the Kibana dashboard <http://central-infra-ilP-Address:5601> or execute the following command in the central microservices virtual machine to check the rendered configuration and the deployment status.

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

```
csp.csp-cms-central-2820689874-gvjbh          1/1      Running
1          19h
csp.csp-cms-central-core-2224266535-kmplk      1/1      Running
1          19h
root@centralmsvm:~# kubectl logs -f csp.csp-cms-central-core-2224266535-kmplk
-n central
```

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

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.

Missing Data in Application Visibility Page

Problem

Description: Data is missing in the Application Visibility page.

Solution

1. Check whether the TCP connection is established between the CPE and the regional sblb virtual machine on port 3514.

```
root@regional-sblb:~#netstat -anp | grep 3514
```

```
tcp        0      0 0.0.0.0:3514          0.0.0.0:*            LISTEN
          1047/haproxy
root@regional-sblb:~#
```

Or execute the following command in the CPE device:

```
root@cpe # show security flow session | grep 3514
```

2. If the TCP connection is not established on port 3514, then check the IP connectivity between the CPE device and the regional sbld virtual machine. Ensure that TCP port 3514 is not blocked in the path.

Link Switch Does Not Happen During SLA Violation

Problem

Description: Link switch does not happen during service-level agreement (SLA) violation in bandwidth-optimized SD-WAN deployments.

Solution

1. Check that the applications match the SD-WAN policy.
2. Check that CSO or Controller recognizes the SLA violation.
Log in to the Administration Portal, and select **Monitor > Applications > SLA performance**.
3. Verify whether the CPE time is synchronized with the NTP server.
4. Click the SLA profile and ensure that the SLA performance data is correct. If it is not, then check that the violation is introduced in the appropriate link.
5. Log in to the CPE device and check the RPM result. Verify the preferred route in the SLA VRF (TC* VRF) table using the following commands.

```
root@cpe # show services rpm probe-results
root@cpe # show route table TC1-CustomerA_DefaultVPN.inet.0
```

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

IN THIS CHAPTER

- [Troubleshooting CSO Installation Issues | 54](#)

Troubleshooting CSO Installation Issues

IN THIS SECTION

- [Salt Key Issue During CSO Installation | 54](#)
- [TimeZone Error | 56](#)
- [SSL Handshake Failure | 56](#)
- [Missing Interface on CSO VM | 57](#)

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-vms) 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

IN THIS CHAPTER

- [Troubleshooting SMTP Issues | 58](#)

Troubleshooting SMTP Issues

IN THIS SECTION

- [Basic Configuration for SMTP Server | 58](#)
- [Basic Configuration for AWS CSO Installations | 60](#)

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.

Basic Configuration for AWS CSO Installations

Problem

Description: User was unable to send e-mail from Amazon Web Services (AWS) CSO installations.

Solution

AWS uses Amazon Simple Email Service (Amazon SES), which is a cloud-based e-mail service. In case of AWS installations, in addition to configuring SMTP server settings, the users must be registered with the AWS SES.

1. Register AWS SMTP email servers. Navigate to <https://docs.aws.amazon.com/ses/latest/DeveloperGuide/Welcome.html> and understand the following topics:

- Sending Email
- Receiving Email
- Controlling Access

2. Configure CSO SMTP Server settings. The information from the AWS SMTP e-mail setup must be added to CSO SMTP Server settings.

The SMTP server settings are:

- Server address—email-smtp.us-east-1.amazonaws.com
- TLS—Enabled
- Port Number—587

3. Configure SMTP authentication settings such as user name, password, and e-mail address. A sample format for e-mail address is: examplejuniperuser@juniper.net .

4. Register the users to the AWS SES. The network operator who configure the AWS SES must request the users to join the AWS SES.

The user receives an e-mail as follows:

NOTE: For brevity, only a part of the e-mail output is listed. Rest of the e-mail content has been replaced with ellipses (...).

Dear Amazon Web Services Customer,

We have received a request to authorize this email address for use with Amazon SES and Amazon Pinpoint in region US East (N. Virginia). If you requested this

verification, please go to the following URL to confirm that you are authorized to use this email address:

`https://example.com`

Your request will not be processed unless you confirm the address using this URL. This link expires 24 hours after your original verification request.

If you did NOT request to verify this email address, do not click on the link. Please note that many times, the situation isn't a phishing attempt, but either a misunderstanding of how to use our service, or someone setting up email-sending capabilities on your behalf as part of a legitimate service, but without having fully communicated the procedure first. If you are still concerned, please forward this notification to `aws-email-domain-verification@amazon.com` and let us know in the forward that you did not request the verification.

(...)

Sincerely,

The Amazon Web Services Team.

RELATED DOCUMENTATION

| *Configuring SMTP Settings*

Troubleshooting RBAC and OpCo Issues

IN THIS CHAPTER

- [Troubleshooting RBAC and OpCo Issues | 62](#)

Troubleshooting RBAC and OpCo Issues

IN THIS SECTION

- [Authentication Failed for the SP User, Tenant User, or OpCo User | 62](#)
- [Authorization Failed for the SP User, Tenant User, or OpCo User | 64](#)
- [Password to Onboard OpCo is Not Received or has Expired | 65](#)

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 the Password*

If the user name is not available then the CSP administrator or the SP administrator must add the user in CSO. See *Adding 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 the 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 the Password](#)

Adding Service Provider and OpCo Users

Editing the Authentication Method

Configuring a Single Sign-On Server

Troubleshooting CSO Release 4.1 Issues

IN THIS CHAPTER

- [Troubleshooting CSO Release 4.1.0 Issues | 67](#)

Troubleshooting CSO Release 4.1.0 Issues

IN THIS SECTION

- [Secure OAM Activation Failure | 67](#)
- [Configure Site Failure | 68](#)
- [Device Activation Failure | 68](#)
- [Dual-CPE Activation Failure for NFX Series Devices | 69](#)
- [Dual-CPE Activation Failure for SRX Series Devices | 70](#)
- [Link Switch Event or Performance Metrics is Not Displayed | 70](#)
- [WAN Link Performance Parameters are Not Displayed | 71](#)
- [LTE Interface Issues | 71](#)

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 <cso-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 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. Log in to Kibana dashboard <http://regional-infra-IP-Address:5601> to find more details regarding the job failure.

Use the request ID to query for configure site workflow. This value is available in the job details page.

5. 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 is **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.
8. Log in to Kibana dashboard <http://regional-infra-IP-Address:5601> to find more details regarding the job failure.

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**.

Log in to Kibana dashboard <http://regional-infra-IP-Address:5601> to find more details regarding the job failure.

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.

Log in to Kibana dashboard <http://regional-infra-IP-Address:5601> to find more details regarding the job failure.

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.