JUNIPER NETWORKS | Engineering Simplicity

Contrail Service Orchestration Monitoring and Troubleshooting Guide

Published
2022-11-17

RELEASE
6.3.0

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

Contrail Service Orchestration Monitoring and Troubleshooting Guide
6.3.0

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 This Guide

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

PART 1

Monitoring Contrail Service Orchestration

CHAPTER 1

Monitoring Infrastructure Services and Microservices for On-premise Deployments

IN THIS CHAPTER

Monitoring and Troubleshooting Overview

IN THIS SECTION

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 | 24

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://*NAT-IP*:5601

   where:

   NAT-IP—IP address of the NAT server. Use this option to monitor the microservices.

   For example:

   http://192.0.2.2:5601

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

RELATED DOCUMENTATION

Setting Up the Visual Presentation of Microservice Log Files | 4

Retrieve Passwords for Infrastructure Components

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" on page 4.
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

Monitoring and Troubleshooting Overview | 2
Viewing Information About Microservices | 5

Viewing Information About Microservices

IN THIS SECTION

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:

- SaltStack
- Cassandra
- MariaDB
- Swift
- Redis
- ArangoDb
- Keystone
- Elasticsearch
- Elk Elasticsearch
- Icinga
- RabbitMQ
- Etcd
- Rsyslog
- Kubernetes
- ELK Logstash
- ELK Kibana
- ZooKeeper
- Contrail Analytics
- VRR
- Microservices

To check the status of infrastructure components:

1. Log in to the startupserver_1 VM as root.

2. Navigate to the CSO directory in the startupserver_1 VM.

   For example:

```
root@host:~/# cd Contrail_Service_Orchestration_6.1.0
root@host:~/Contrail_Service_Orchestration_6.1.0#
```

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

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

```
root@startupserver_1:/opt/Contrail_Service_Orchestration_6.1.0# ./components_health.sh --
component=<component_name>
For Example:
root@startupserver_1:/opt/Contrail_Service_Orchestration_6.1.0# ./components_health.sh --
component=elasticsearch
```

   Run the following command to check the health of all the infrastructure components.

```
root@startupserver_1:/opt/Contrail_Service_Orchestration_6.1.0# ./components_health.sh
```

   After a couple of minutes, the status of each infrastructure component is displayed.

   For example:

```
INFO     Updating the mine and syncing the grains
INFO     ********************************************************************
INFO     HEALTH CHECK FOR INFRASTRUCTURE COMPONENTS STARTED IN CENTRAL ENVIRONMENT
INFO     ********************************************************************

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

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

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

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

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

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

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

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

INFO     Health Check for Infrastructure Component Rabbitmq Started
INFO     The Infrastructure Component Rabbitmq 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 Rsyslog Started
INFO     The Infrastructure Component Rsyslog is Healthy

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

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

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

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

INFO     Health Check for Infrastructure Component Icinga Started
WARNING  Icinga is not running in 1 node
INFO     The Infrastructure Component Icinga is Healthy

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

INFO     Health Check for Infrastructure Component Rabbitmq Started
INFO     The Infrastructure Component Rabbitmq 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 Elasticsearch Started
INFO     The Infrastructure Component Elasticsearch is Healthy

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

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

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

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

INFO     Overall result:
INFO     The following Infrastructure Components are Healthy:
INFO     ['Saltstack', 'Etcd', 'Mariadb', 'Keystone', 'Swift', 'Elk_Kibana', 'Rsyslog',
'Kubernetes', 'Vrr', 'Zookeeper', 'Arangodb', 'Icinga', 'Redis', 'Rabbitmq', 'Elk_Logstash',
'Elasticsearch', 'Elk_Ela
sticsearch', 'Cassandra', 'Contrail_Analytics', 'Microservices']
```

If the **components_health.sh** script detects any service as unhealthy, then it displays an error message as shown in the following sample output.

```
root@startupserver1:~/opt/Contrail_Service_Orchestration_6.1.0# ./components_health.sh
INFO     Updating the mine and syncing the grains
INFO     ********************************************************************
INFO     HEALTH CHECK FOR INFRASTRUCTURE COMPONENTS STARTED IN CENTRAL ENVIRONMENT
INFO     ********************************************************************



INFO     Health Check for Infrastructure Component Saltstack Started
INFO     Attempt: 1 - Retrying Health Check for Component Saltstack
INFO     Attempt: 2 - Retrying Health Check for Component Saltstack
ERROR    The Infra Component : Saltstack is Unhealthy


ERROR    ERROR: SaltStack is Unhealthy. Please run CSO recovery to recover SaltStack
```

To recover the service, you must run the **recovery.sh** script. See "Recovering CSO Services" on page 12.

Recovering CSO Services

IN THIS SECTION

Recovering a CSO Service

If the **components_health.sh** script detects any service as unhealthy, then you can recover the service using the **recovery.sh** script.

1. From the CSO directory, run the **recovery.sh** script:

```
root@host:~/Contrail_Service_Orchestration_6.1.0#./recovery.sh
*********************************************
This tool assists you recover your CSO setup.
*********************************************

 Following components can be recovered


1: cassandra
2: mariadb
3: etcd
4: kubernetes
5: vrr
6: saltstack
7: arangodb
8: microservices
9: icinga
10: rabbitmq
 Specify one of the component to recover (In Number):6
```

2. Specify the component (number) that needs to be recovered.

The recovery script starts the recovery process for the specified component (saltstack in this example). The following is a sample output of the messages that are displayed. A recovery completion message is displayed after the component is recovered.

```
INFO     Started recovering saltstack component at 2021-07-22 03:00:08.989767 ...
INFO     Saltstack failure recovery is initiated...
INFO     Saltstack check()
INFO     Salt Master is running
INFO     Deleting unreachable minion key csp-central-proxy_sblb2.N6RGW8.central
INFO     Deleting unreachable minion key csp-central-k8-microservices3.N6RGW8.central
INFO     Deleting unreachable minion key csp-central-k8-microservices2.N6RGW8.central
INFO     Deleting unreachable minion key csp-central-k8-infra3.N6RGW8.central
INFO     Completed recovering saltstack component at 2021-07-22 03:00:27.816847 .
INFO     Time taken to recover 0:00:18.827080
```

If the recovery.sh script detects an issue with the k8 virtual machine when recovering the kubernetes component, then it displays an error message as shown in the following sample output:

```
Kubernetes recovery failed, Please refer logs/recovery.log for more details
Failed to recover Kubernetes
Please run replace_vm for k8-master1
```

You can run the deploy.sh script to replace the k8 virtual machine.

Replacing Virtual Machines for KVM Hypervisor

You can replace only a k8 virtual machine. To replace a k8 virtual machine:

1. Run the `./deploy.sh` script from the CSO directory

```
root@host:~/Contrail_Service_Orchestration_6.1.0#./deploy.sh
Enter the number for operation to be performed:
1. Deploy CSO
2. Replace VM
0. Exit
#Your choice: [1 --> CSO Infra Deployment; 2 --> 'Replace existing VM, currently supports
only k8-master, k8-infra and k8-microservices node for replacement in KVM]:
```

2. Enter `2` at the prompt. A list of k8 virtual machines that can be replaced are displayed.

```
*******************************************
REPLACE INSTANCE
Listing the VMs that can be replaced...
[{'k8-infra1': '192.168.10.26'}, {'k8-master1': '192.168.10.34'}, {'k8-microservices1':
'192.168.10.37'}, {'k8-infra2': '192.168.10.27'}, {'k8-master2': '192.168.10.35'}, {'k8-
microservices2': '192.168.10.38'}, {'k8-infra3': '192.168.10.28'}, {'k8-master3':
'192.168.10.36'}, {'k8-microservices3': '192.168.10.39'}]
*******************************************
Please provide the IP of vm to be replaced:
```

3. Specify the IP of the k8 virtual machine, which is corrupted and must be replaced.

   The virtual machine is spawned again.

NOTE: If a power cycle occurs on the physical servers, the vrr becomes unhealthy.

The following is a sample output that shows the vrr status as unhealthy:

```
INFO     Health Check for Infrastructure Component Vrr Started
INFO     Attempt: 1 - Retrying Health Check for Component Vrr
INFO     Attempt: 2 - Retrying Health Check for Component Vrr
ERROR    The Infra Component : Vrr is Unhealthy
```

Run the recovery.sh script again to bring the vrr back online. The following is a sample output of the recovery process:

```
root@startupserver1:~/Contrail_Service_Orchestration_6.1.0# ./recovery.sh
***************
This tool assists you recover your CSO setup.
***************
 Following components can be recovered
1: contrailanalytics
2: cassandra
3: mariadb
4: etcd
5: kubernetes
6: vrr
7: saltstack
8: arangodb
9: microservices
10: icinga
11: rabbitmq

Specify one of the component to recover (In Number) : 6
INFO     Started recovering vrr component at 2021-07-22 07:38:23.504490 ...
INFO     VRR recovery is initiated...
INFO     Vrr - 192.168.10.29 is healthy
ERROR    Vrr - 192.168.10.30 is unhealthy
INFO     Recovery takes time, please be patient
INFO     VRR recovery started. Please wait...
INFO     Vrr console recovered for vrr2
INFO     VRR config sync completed successfully
INFO     Completed recovering vrr component at 2021-07-22 07:48:31.324598 .
INFO     Time taken to recover 0:10:07.820108
```

CHAPTER 2

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

IN THIS CHAPTER

Backup and Restore of Contrail Service Orchestration (CSO) Databases

IN THIS SECTION

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 19 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. A full backup of all platform, OpCo, 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

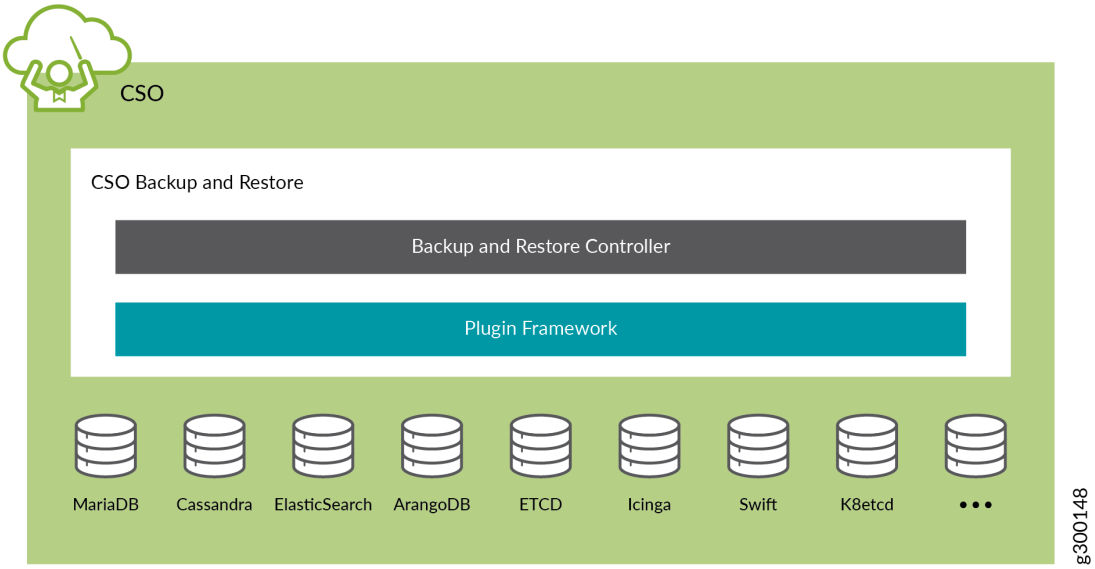

Figure 1 on page 17 shows a conceptual image of how backup and restore is implemented in CSO 4.1. The database systems that are currently backed up within the framework are: MariaDB, Cassandra, ElasticSearch, ArangoDB, Zookeeper, and ETCD. The system also backs up encrypted passwords, and system certificates 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 1 on page 18 lists the major components and a brief description of each.

Table 1: Major Components

Component	Description
Backup and Restore Controller	• Handles backup or restore calls from administrator. The calls are made using the cso_backupnrestore 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	• Framework that allows backup and restore to deal with multiple different databases. • Allows for future inclusion of other databases. • Salt Minions
Plug-in	• 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 19 shows the flow of backup and restore operations.

Figure 2: Backup and Restore Operations

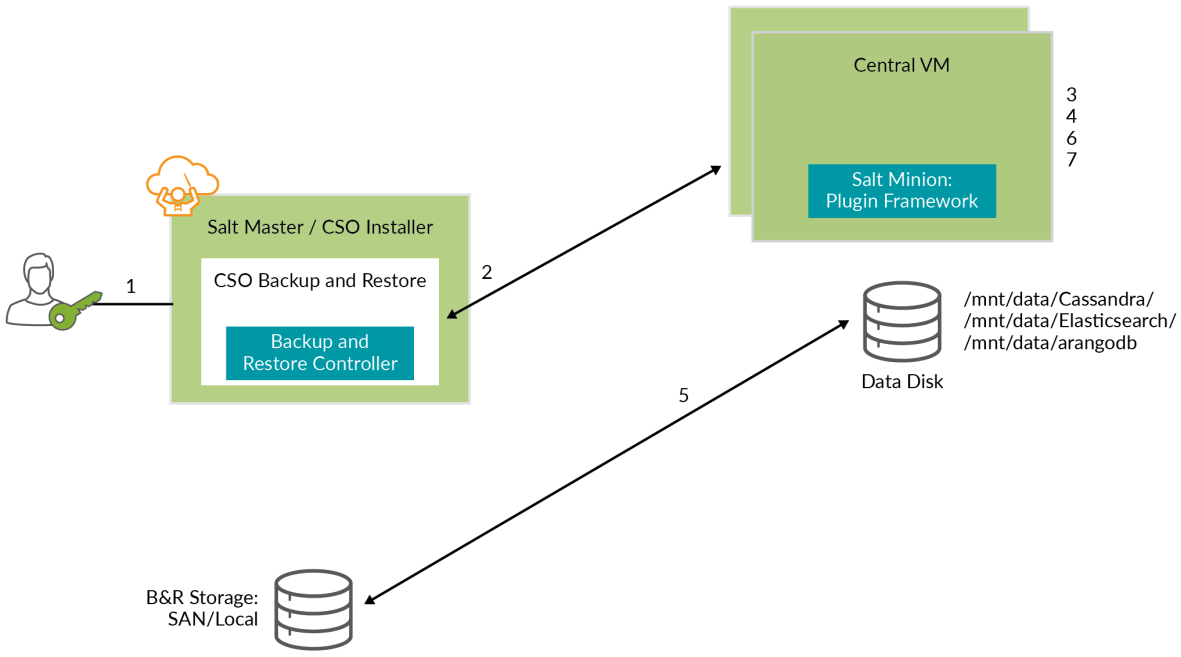

1. Admin triggers backup or restore command
2. Controller sends event to VMs
3. CSO is put into Maintenance mode
4. DB integrity checks
5. Backup or restore operation is triggered for each DB and component
6. DB integrity checks and cleanup as needed
7. CSO is put back into Operational mode

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. Ensure that there is sufficient storage space on the installer-vm or startup server to store the backup data.

For restore operations, the same `cso_backupnrestore` command is used with different options as described in Table 2 on page 20 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 2 on page 20. Only one of the arguments can be used with any one of the options.

Table 2: cso_backupnrestore Command Options

Option	Purpose	Arguments
`-b`	Specify operation (REQUIRED)	backup, restore, healthcheck, reindex, backupdetails, listbackups, scheduledbackup
`-s`	Specify the name of the snapshot created by backup operation or restored by restore operation.	backup name
`-c`	Specify which database to backup or restore [default '*'](OPTIONAL)	For backup: only '*' is allowed. For restore: Comma separated list with any or all of: cassandra, elasticsearch, zookeeper, mariadb, etcd, arrangodb. '*' restores all databases
`-r`	Specify whether the restore operation is for disaster recovery or not [Default no].	yes or no

Table 2: cso_backupnrestore Command Options *(Continued)*

Option	Purpose	Arguments
`-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] • 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
```

CHAPTER 3

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

IN THIS CHAPTER

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 | 25

Monitoring VNFs Used in Network Services and the VMs That Host the VNFs | 26

Monitoring Microservices | 31

Monitoring Microservices and Their Host VMs | 33

Monitoring Physical Servers | 35

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://*NAT-IP*:1947/icingaweb2

   NAT-IP—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 | 23

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 3 on page 25 shows the meaning of the monitoring parameters for network services.

Table 3: 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. • Up—operational • Down—not operational
Number of Network Functions	Number of VNFs in the service chain.

Table 3: Parameters for Monitoring Network Services *(Continued)*

Parameter	Meaning
Network Function	Number of network functions in a colored square that indicates the status of the instance. When you click the square you see: • 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	Result of the commands issued to monitor the status of the network service. When you click the square you see: • 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 | 26

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 4 on page 27 for information about the monitoring parameters used for VNFs and the VMs that host them.

Table 4: Parameters for Monitoring VNFs and Their Host VMs

Parameter	Meaning
left_net_interface_input_pckt_rate	Rate of traffic entering the interface that transmits data to the host.
left_net_interface_output_pckt_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. • Up—operational • Down—not operational
right_net_interface1_stats	State of the interface to which the host transmits data. • 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.

Table 4: Parameters for Monitoring VNFs and Their Host VMs *(Continued)*

Parameter	Meaning
routing_engine_ctrlplane_memusage	Percentage of the Routing Engine's control plane memory that VM is using.
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. • 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.

Table 4: Parameters for Monitoring VNFs and Their Host VMs *(Continued)*

Parameter	Meaning
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. • Up—operational • Down—not operational
right_net_interface_status	State of the interface to which the host transmits data. • Up—operational • Down—not operational
right_net_interface_input_pckt_rate	Rate of traffic entering the interface to which the host transmits data.
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. • Enabled—operational • Disabled—not operational
vsrx_firewall_config	State of the vSRX firewall VNF. • Enabled—operational • Disabled—not operational
vsrx_utm_config	State of the vSRX UTM VNF. • Enabled—operational • Disabled—not operational

Table 4: Parameters for Monitoring VNFs and Their Host VMs *(Continued)*

Parameter	Meaning
vsrx_dpi_config	State of the DPI firewall VNF. • Enabled—operational • Disabled—not operational
iptable_status	State of the LxCIPtable VNF. • Enabled—operational • Disabled—not operational
iptable_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. • Up—operational • Down—not operational
cisco_right_interface_status	State of the interface to which the host transmits data for the CSR-1000V VNF. • 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.
cisco_right_input_packets	Rate of traffic entering the interface to which the host transmits data for the CSR-1000V VNF.

Table 4: Parameters for Monitoring VNFs and Their Host VMs *(Continued)*

Parameter	Meaning
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 | 25

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 Microservervice 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 5 on page 32 shows the monitoring parameters for microservices.

Table 5: 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. • 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: • The status of the host in which the micorservice 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: • 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 | 33

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 6 on page 33 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
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

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

Parameter	Meaning
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
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 | 31

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 7 on page 35 for information about the parameters.

Table 7: Parameters for Monitoring Physical Servers

Parameters	Meaning
Group Status	State of the server cluster and the time when it entered that state. • Up—Operational • Down—Not operational
Number of Servers	Number of servers in the server cluster.
Server Status	Number of servers in a colored square that indicates the status of the servers. When you click the square you see: • 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.

Table 7: Parameters for Monitoring Physical Servers *(Continued)*

Parameters	Meaning
Command Status	Result of the commands issued to monitor the status of the server. When you click the square you see: • 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 | 23

2

PART

Troubleshooting Contrail Service Orchestration Issues

CHAPTER 4

Identifying Connectivity Issues for Cloud-based Deployments

IN THIS CHAPTER

Identifying Connectivity Issues by Using Ping

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

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

NOTE: In Contrail Service Orchestration (CSO) Release 6.1, the following devices support ping:

- NFX Series: NFX150, NFX250
- SRX Series: SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600
- vSRX

To perform the ping operation:

1. Do one of the following:
   - To initiate a ping from a provider hub device, select **Resources** > **Provider Hub Devices**.

     The :Provider Hub Devices page appears.

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

  The Tenant Devices page appears.

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

   The Ping page appears.

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

3. Complete the configuration according to the guidelines provided in Table 8 on page 39.

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

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

   A job is created and a Ping Progress page appears. After the host sends the ping packets, the Ping Result page appears. If the ping operation is successful, the Ping Result page displays the parameters specified in Table 9 on page 41.

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

Table 8: Fields on the Ping page

Field	Description
Remote Host	Enter the IPv4 address or hostname of the remote host.
Ping Request Packets	Enter the number of ping request packets to be sent to the remote host. Default: 5. Range: 1 through 300.
Advanced	
Source Interface	Select the source interface on the device through which you want to send the ping request to the remote host. If you do not select a source interface, ping requests are sent on all interfaces. To clear the selected interface, click **Clear All** and select another interface.

Table 8: Fields on the Ping page *(Continued)*

Field	Description
Hostname Resolution	Click the toggle button to enable or disable (default) the display of hostname of the hops along the path to the remote host.
Rapid Ping	Click the toggle button to enable or disable (default) sending ping requests rapidly. If you enable this option, the device sends a minimum of 100 ping request packets per second or sends a packet as soon as a response to the previous packet is received, whichever is greater. • If the source device does not receive a response for 500 ms, timeout is considered. • If the source device receives a response within 500 ms, the next ping request packet is sent immediately. **NOTE**: The ping results are displayed in a single consolidated message instead of individual messages for each ping request packet sent.
Packet Fragmentation	Click the toggle button to enable or disable (default) the fragmenting of ping request packets. If packet fragmentation is disabled, ping packets with the maximum transmission unit (MTU) greater than 1500 bytes are dropped.
Packet Size (bytes)	Enter the size (in bytes) of the ping request packet. Default: 56 bytes. Range: • 1 through 1,472 bytes, if packet fragmentation is disabled. • 1 through 65,468 bytes, if packet fragmentation is enabled.

Table 8: Fields on the Ping page *(Continued)*

Field	Description
Wait Time (seconds)	Enter the time (in seconds) for which the source device waits for a response to the ping request packet. The source device considers the remote host as not reachable after the wait time elapses. Default: 10 seconds. Range: 0 through 600 seconds.
Incoming Interface	Click the toggle button to include or exclude (default) information (on the Ping Result page) about the interface on the source device that receives the ping responses..
Routing Instance	Select a specific routing instance that the ping request packets can use to reach the remote host. The ping result displays the information about the connectivity between the source device and the remote host based on the selected routing instance. To clear the selected routing instance, click **Clear All** and select another routing instance.

Table 9: Fields on the Ping Result page

Field	Description
Packet Loss	Displays the percentage of ping packets sent for which the source device did not receive a response.

Table 9: Fields on the Ping Result page *(Continued)*

Field	Description
Round Trip Time Taken (in µs)	Displays the following information about the duration (in microseconds) between the time when the device sends the ping request and the time when the device receives a response from the remote host. Displays the following: • Minimum: The minimum time taken to receive a response for a ping request packet. • Maximum: The maximum time taken to receive a response for a ping request packet. • Average: The average time taken to receive a response for all the ping request packets sent in a ping operation. • Standard Deviation: The variation of the round trip time from the mean round trip time.
Details	
Sequence	Sequence number of all the ping request packets.
Result	Result of the ping request packets—Success or Failure.
Incoming Interface	Interface on the source device on which the responses are received for the ping requests. This data appears if you have enabled the Incoming Interface option on the Ping page.
Time Taken	Time taken (in microseconds) to receive response to a ping request packet.

Identifying Connectivity Issues by Using Traceroute

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

NOTE: In Contrail Service Orchestration (CSO) Release 6.1, the following devices support traceroute:

- NFX Series: NFX150, NFX250
- SRX Series: SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200, SRX4600
- vSRX

To perform traceroute operation:

1. Do one of the following:
   - To initiate traceroute from a provider hub device, select **Resources** > **Provider Hub Devices**.

     The Provider Hub Devices page appears.
   - To initiate traceroute from a tenant device, select **Resources** > **Tenant Devices**.

     The Tenant Devices page appears.
2. Select a device from the list of devices displayed and click **More** > **Traceroute**.

   The Traceroute page appears.
3. Complete the configuration according to the guidelines provided in Table 10 on page 43.

   NOTE: Fields marked with an asterisk (*) are mandatory.
4. Click **Traceroute** to initiate the traceroute operation.

   A job is created and a traceroute progress page appears. If the traceroute operation is successful, the Traceroute Result page displays the traceroute parameters specified in Table 11 on page 45.

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

Table 10: Fields on the Traceroute page

Field	Description
Remote Host	Enter the IPv4 address or hostname of the remote host.

Table 10: Fields on the Traceroute page *(Continued)*

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

Table 11 on page 45 lists the parameters on the Traceroute Result page when the traceroute operation is successful.

Table 11: Fields on the Traceroute Result page

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

CHAPTER 5

Troubleshooting Site Activation Issues for Cloud-based Deployments

IN THIS CHAPTER

Troubleshooting Site Activation Issues

IN THIS SECTION

Prerequisites to Activate a Site

IN THIS SECTION

Problem

Description

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

Solution

The prerequisites to activate a site are as follows:

- Check the spoke connectivity to Internet.
- Check the firewall policies between the CPE device and the CSO. The hub or spoke must be able to communicate to CSO through ports 443 (activation), 444 (activation for small and medium deployments), 7804 (outbound-ssh), 3514(app-track logs), 514 (syslog), and 2216 (telemetry agent). See Deployment Guide

Site activation process is stuck in device detected state

IN THIS SECTION

Problem

Description

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

Solution

Do the following:

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

Site activation process is stuck in bootstrap state

IN THIS SECTION

Problem

Description

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

Solution

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

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

Site activation process failed in bootstrap state

IN THIS SECTION

Problem

Description

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

Solution

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

Site activation process failed during provisioning

IN THIS SECTION

Problem

Description

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

Solution

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

CHAPTER 6

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

IN THIS CHAPTER

Troubleshooting Image, License, and Policy Deployment Issues

IN THIS SECTION

Unable to find device image version

IN THIS SECTION

Problem

Description

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

Solution

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

To access the J-Web interface of the device:

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

   The Welcome page appears displaying the device image version.

Upgrade device image using J-Web

IN THIS SECTION

Problem

Description

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

Solution

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

To upgrade the device image using J-Web:

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

Unable to connect to the device

IN THIS SECTION

Problem

Description

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

Solution

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

Device image version is different from the recommended version

IN THIS SECTION

Problem

Description

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

Solution

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

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

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

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

Policy deployment failed

IN THIS SECTION

Problem

Description

Policy deployment failed; how do I proceed?

Solution

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

No data for next-generation firewall site

IN THIS SECTION

Problem

Description

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

Solution

Do the following:

- Verify that your network firewall allows the UDP port 514.

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

No data for SD-WAN site

IN THIS SECTION

Problem

Description

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

Solution

Do the following:

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

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

IN THIS SECTION

Problem

Description

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

Solution

1. Verify the alerts for overlay or underlay connections, and check whether BGP is active.

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

2. Check whether the firewall policies are successfully deployed to the CPE device and that the traffic or applications are matching the policies to permit the traffic to Internet or to other sites.

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

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

3. Check the routes in the default VRF route table in the CPE device.

4. Trace the route and verify the reachability from the hub to the destination. If the hub cannot reach the Internet, then verify whether the firewall and NAT policies are set up properly in the hub.

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

SLA Violation-Original Link Recovered After SLA Violation

IN THIS SECTION

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

IN THIS SECTION

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**.

CHAPTER 7

Troubleshooting SMTP Issues for Cloud-based Deployments

IN THIS CHAPTER

Troubleshooting SMTP Issues

IN THIS SECTION

Basic Configuration for SMTP Server

IN THIS SECTION

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:

   a. Check the SMTP server settings to see if they match the SMTP server provider's settings.

   b. Check authentication credentials.

   c. 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).

   d. Check whether there is network access from CSO to the SMTP server.

   e. Check whether the firewall is blocking SMTP traffic to SMTP server or whether the ports are blocked. If the server settings and authentication settings are correct, check whether the firewall is blocking port 587 and 465 and SMTP traffic. If it is a case of the firewall blocking, then work with the network administrator to unblock ports 465, 587, and SMTP traffic.

RELATED DOCUMENTATION

Configuring SMTP Settings

CHAPTER 8

Recovering an Installation

IN THIS CHAPTER

CSO Disaster Recovery

In case of any failures you can recover CSO Release 6.2.0. To recover CSO Release 6.2.0 you must have already taken a backup and saved the backup file.

To recover CSO Release 6.2.0:

1. Based on the hypervisor you are using, do one of the following:
   - If you are using KVM as the hypervisor:
     a. Copy the CSO 6.2.0 backup folder to the bare metal server.
     b. From the backup folder, copy the **_topology.conf** file to the **Contrail_Service_Orchestration_6.2.0/topology/** folder.

        For example:

        ```
        cp /root/backups/backupfordr/2020-06-19T17:27:05/config_backups/_topology.conf /root/
        Contrail_Service_Orchestration_6.2.0/topology/
        ```

     c. Provision the VMs. For information on provisioning KVM hypervisor, see *Provision VMs on Contrail Service Orchestration Servers* in *CSO Installation and Upgrade Guide*.
     d. Copy the backup folder file from the bare metal server to the startupserver1 VM.

        ```
        user@server>scp -r /root/backups/backupfordr/ startupserver1:
        ```

     e. Log in to the startupserver1 VM as the root user.

f. Expand the installer package.

```
root@startupserver1:~/# tar -xvzf Contrail_Service_Orchestration_6.2.0.tar.gz
```

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

g. From the backup folder, copy the **_topology.conf** file to the **Contrail_Service_Orchestration_6.2.0/topology/** folder.

```
cp /root/backups/backupfordr/2020-06-19T17:27:05/config_backups/_topology.conf /root/
Contrail_Service_Orchestration_6.2.0/topology/
```

- If you are using ESXi as the hypervisor:

a. Copy the backup folder to the startupserver1 VM.

b. Expand the installer package.

```
root@startupserver1:~/# tar -xvzf Contrail_Service_Orchestration_6.2.0.tar.gz
```

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

c. From the backup folder, copy the **_topology.conf** file to the **Contrail_Service_Orchestration_6.2.0/topology/** folder in the startupserver1 VM.

For example:

```
cp /root/backups/backupfordr/2020-06-19T17:27:05/config_backups/_topology.conf /root/
Contrail_Service_Orchestration_6.2.0/topology/
```

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

```
root@host:~/Contrail_Service_Orchestration_6.2.0./deploy.sh
```

3. Run the following command:

```
cso_backupnrestore -b backup -s backup62new
```

4. Run the pre_disaster recovery script.

   python /usr/local/bin/pre_disaster_recovery.py

```
Enter the old backup path: /root/backups/backupfordr/2020-10-29T06:45:11:45:11
Enter the new backup path: /backups/backup62new/2020-10-30T03:47:51
COMPONENTS: ('cassandra', 'elasticsearch', 'etcd', 'arangodb', 'icinga', 'swift',
'config_backups') Start cassandra pre restore task...
Get old and new backup path for component cassandra
cassandra pre restore task successfully done
*Do you want to redeploy cassandra container to apply tokens.
*This process will delete all the existing data from cassnadra
Please enter yes to process [yes/no]:
```

   Enter **yes** at the prompt.

```
Start elasticsearch pre restore task...
Get old and new backup path for component elasticsearch
Get Elasticsearch user id for permission
Set permission for elasticsearch dir.
elasticsearch pre restore task successfully done
Start etcd pre restore task...
Get old and new backup path for component etcd
etcd pre restore task successfully done
Start arangodb pre restore task...
Get old and new backup path for component arangodb
arangodb pre restore task successfully done
Start mariadb pre restore task...
Get old and new backup path for component mariadb
mariadb pre restore task successfully done
Start icinga pre restore task...
Get old and new backup path for component icinga
icinga pre restore task successfully done
Start swift pre restore task...
Get old and new backup path for component swift
swift pre restore task successfully done
Start config_backups pre restore task...
config_backups pre restore task successfully done
Pre restore task completed for all components.
```

5. Restore the data from the new backup created in step 3 by using the **cso_backupnrestore** script.

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

#cso_backupnrestore -b restore -s backuppath -t '*' -c ‘elasticsearch' -r ‘yes’

#cso_backupnrestore -b restore -s backuppath -t '*' -c ‘arangodb' -r ‘yes’

#cso_backupnrestore -b restore -s backuppath -t '*' -c ‘icinga' -r ‘yes’

#cso_backupnrestore -b restore -s backuppath -t '*' -c ‘swift' -r ‘yes’

#cso_backupnrestore -b restore -s backuppath -t '*' -c ‘mariadb' -r ‘yes’
```

where `backuppath` is the new backup path.

If the restore procedure fails for any of the above components, you must retry to restore only those components. At times, restore of mariadb fails at the first attempt but is successful at the second attempt.

6. Synchronize the data between nodes.

```
cso_backupnrestore -b nodetool_repair
```

```
IF Cluster nodetool status is UP/Normal(UN) please proceed for nodetool repair (Y/n):
```

Enter **y** at the prompt.

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

```
salt-cp -G "roles:haproxy_confd_sblb" /root/backups/backupfordr/2020-06-19T17:27:05/
config_backups/haproxycerts/minions/minions/csp-central-proxy_sblb1.NH5XCS.central/
files/etc/pki/tls/certs/ssl_cert.pem /etc/pki/tls/certs
```

```
salt-cp -G "roles:haproxy_confd_sblb" /root/backups/backupfordr/2020-06-19T17:27:05/
config_backups/haproxycerts/minions/minions/csp-central-proxy_sblb1.NH5XCS.central/
files/etc/pki/tls/certs/ssl_cert.crt /etc/pki/tls/certs
```

8. Restart the SBLB HA Proxy.

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

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

```
salt-cp -G "roles:haproxy_confd" /root/backups/backupfordr/2020-06-19T17:27:05/
config_backups/haproxycerts/minions/minions/csp-central-proxy1.NH5XCS.central/
files/etc/pki/tls/certs/ssl_cert.pem /etc/pki/tls/certs
```

```
salt-cp -G "roles:haproxy_confd" /root/backups/backupfordr/2020-10-29T06:45:11/
config_backups/haproxycerts/minions/minions/csp-central-proxy1.NH5XCS.central/
files/etc/pki/tls/certs/ssl_cert.crt /etc/pki/tls/certs
```

10. Restart the Central HA Proxy.

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

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

```
kubectl get secret -n central | grep cso-ingress-tls
cso-ingress-tls kubernetes.io/tls                    2      17d
kubectl delete secret cso-ingress-tls -n central kubectl create secret tls cso-ingress-tls
--key /root/backups/backupfordr/2020-10-29T06:45:11/config_backups/haproxycerts/minions/
minions/csp-central-proxy1.NH5XCS.central/files/etc/pki/tls/certs/ssl_cert.key --cert /root/
backups/backupfordr/2020-10-29T06:45:11/config_backups/haproxycerts/minions/minions/csp-
central-proxy1.NH5XCS.central/files/etc/pki/tls/certs/ssl_cert.crt -n central
```

12. Deploy microservices.

```
/python.sh micro_services/deploy_micro_services.py
```

13. Reindex the elastic search.

    a. Open the csp.csp-ems-regional deployment file.

```
kubectl edit deployment -n regional csp.csp-ems-regional
```

b. Change the replicas to 2 and increase the memory from 500Mi to 2048Mi (2Gi).

c. Save the file.

d. Start the reindex process.

```
cso_backuprestore -b reindex
```

e. Using the admin token, run the following API to build the policy indices:

```
curl --location --request POST 'https://AdminPortalIP/policy-mgmt/_index' \ --header 'x-auth-token: XXXXXXX‘\ --data-raw ‘'
```

14. Create the RabbitMQ FMPM queue.

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

15. Load the data.

```
./python.sh micro_services/load_services_data.py
```

16. Synchronize the Virtual Route Reflector (VRR). Use the admin token. Do not use the cspadmin token.

a. Obtain the topo-uuid for the VRR.

```
GET: https://<IP Address>/topology-service/device
```

b. Synchronize the VRR using the POST https://<ip>/routing-manager/synchronize-vrr API.

```
{
  "input": {
    "recover_vrr": true,
    "uuid": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx”
  }
}
```

17. Restore the SD-WAN and security reports.

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

where `backuppath` is the new backup path.

18. Restart all fmpm-provider-api and fmpm-provider-core pods by deleting the existing pods.

```
root@startupserver1:~# kubectl get pods -n central|grep fmpm-provider
csp.csp-fmpm-provider-6644bc8b94-7pvfn                    1/1     Running   0          9d
csp.csp-fmpm-provider-6644bc8b94-c2psl                    1/1     Running   0          9d
csp.csp-fmpm-provider-6644bc8b94-gzkht                    1/1     Running   1          9d
csp.csp-fmpm-provider-6644bc8b94-hz8f5                    1/1     Running   0          9d
csp.csp-fmpm-provider-6644bc8b94-nsqfs                    1/1     Running   0          9d
csp.csp-fmpm-provider-6644bc8b94-rq9xq                    1/1     Running   0          9d
csp.csp-fmpm-provider-core-797f7c48c9-7nm8q               1/1     Running   0          9d
csp.csp-fmpm-provider-core-797f7c48c9-7zj67               1/1     Running   0          9d
csp.csp-fmpm-provider-core-797f7c48c9-8njsq               1/1     Running   0          9d
csp.csp-fmpm-provider-core-797f7c48c9-rh2jr               1/1     Running   0          9d
csp.csp-fmpm-provider-core-797f7c48c9-sswbg               1/1     Running   0          9d
csp.csp-fmpm-provider-core-797f7c48c9-zvhps               1/1     Running   0          9d
```

19. Delete all the pods displayed in the previous step.

```
kubectl delete pods csp.csp-fmpm-provider-6644bc8b94-7pvfn csp.csp-fmpm-provider-6644bc8b94-
c2psl csp.csp-fmpm-provider-6644bc8b94-gzkht csp.csp-fmpm-provider-6644bc8b94-hz8f5csp.csp-
fmpm-provider-6644bc8b94-nsqfs csp.csp-fmpm-provider-6644bc8b94-rq9xq csp.csp-fmpm-provider-
core-797f7c48c9-7nm8q csp.csp-fmpm-provider-core-797f7c48c9-7zj67 csp.csp-fmpm-provider-
core-797f7c48c9-8njsq csp.csp-fmpm-provider-core-797f7c48c9-rh2jr csp.csp-fmpm-provider-
core-797f7c48c9-sswbg csp.csp-fmpm-provider-core-797f7c48c9-zvhps
```

20. Restore the Contrail Analytics Node (CAN) database.

NOTE: You can restore the database only if a backup is available. CAN backup is disabled by default. To include CAN data in the backup, comment out `contrail_analytics` in the following configuration:

```
root@startupserver1:~# cat /etc/salt/master.d/backup.conf
backups:
```

```
keep: 10
timeout: 1200
path: /backups
enabled_roles:
    - cassandra
    - mariadb
    - kubemaster
    - elasticsearch
# - redis
    - icinga
    - helm_manager
# - contrail_analytics
```

To restore the CAN configuration database, run the following script:

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

To restore the CAN analytics database, perform the following steps:

The **analyticsdb** backup files are located at **/backups/daily/2021-06-07T06:46:37/central/can/contrail_analytics<x>**, where x indicates the contrail analytics node number. The value of x ranges from 1 through 3.

On all the three contrail analytics nodes:

- Copy the CAN backup files from the startupserver to each CAN VM:

  ```
  rsync -a<can-backup-files>root@<can-ip>:<created-backup-folder>
  ```

- Run the following command on the CAN VMs:

  ```
  docker cp 0000/ analytics_database_cassandra_1:/root
  ```

  ```
  docker exec -it analytics_database_cassandra_1 bash
  mv /root/mc-* /var/lib/cassandra/data/ContrailAnalyticsCql/statstablev4-
  d5b63590a7f011eba080c3eb6817d254
  ```

#The path might be different based on uuid.

```
cd /var/lib/cassandra/data/ContrailAnalyticsCql/statstablev4-
d5b63590a7f011eba080c3eb6817d254
chown -R cassandra:cassandra *
nodetool -p 7200 refresh -- ContrailAnalyticsCql statstablev4
```

After a successful upgrade, CSO Release 6.2.0 is functional and you can log in to the Administrator Portal and the Customer Portal.

CHAPTER 9

Renewing Certificates

IN THIS CHAPTER

How to Renew Certificates for CSO Components

IN THIS SECTION

You can renew or view the certificates of CSO components by using the **manage_certificate.sh** script.

NOTE: Actual output might vary from the sample output shown based on your deployment scenario.

1. Log in to the startupserver1 VM as root user.
2. Navigate to the CSO directory in the startupserver1 VM.

   For example:

```
root@startupserver1:~/# cd Contrail_Service_Orchestration_6.2.0
root@host:~/Contrail_Service_Orchestration_6.2.0#
```

3. Run the **manage_certificate.sh** script to check the status or renew the certificates of the CSO components.

```
root@startupserver1:~/Contrail_Service_Orchestration_6.2.0# ./manage_certificate.sh

************************************************************
 This tool assists you to renew CSO components certificate
************************************************************

Certificate renew sequence need to be followed:

     Kubernetes -> Haproxy -> Elasticsearch

     0: List all certificate expiry date
     1: Schedule cron for email notification

Following component's certificate can be renewed

     2: Haproxy, Nginx, Rsyslog
     3: Telemetry Agent

Select a option (In Number) :
```

NOTE: To check the options that you can use with the **manage_certificate.sh** script, enter `manage_certificate.sh -h` or `manage_certificate.sh --help`.

```
root@startupserver1:~/Contrail_Service_Orchestration_6.2.0# ./manage_certificate.sh -h
Usage:
  ./manage_certificate.sh -> to check/renew CSO components's certificate
  ./manage_certificate.sh [options]

  options:

    -c | --check     to only check and list expiry dates of CSO components
    -n | --notify    to list and send email notification with CSO components and its
expiry dates
    --cron           to schedule cron job
    -h | --help      this help
```

4. You can choose to perform any of the following tasks:

- To view the certificate expiry dates, see "How to View the Certificate Expiry Dates" on page 72.
- To schedule a cron job, see "How to Schedule a Cron Job" on page 72.
- To renew a component's certificate, see "How to Renew a Certificate" on page 74.

How to View the Certificate Expiry Dates

To list all the certificates and their expiry dates, type **0** at the prompt and press Enter. You can also view the same output by using **./manage_certificate.sh -c** or **./manage_certificate.sh --check**.

```
Select a option (In Number) : 0
INFO     Fetching certificate details...
+----------------+---------------------+----------------+-------------+
| Component Name |     Expiry Date     | Days to Expire |    Status   |
+----------------+---------------------+----------------+-------------+
|    Haproxy     | 2022-08-24 09:58:20 |      240       | Not Expired |
|     Nginx      | 2022-08-24 09:58:20 |      240       | Not Expired |
|    Rsyslog     | 2022-08-24 09:58:20 |      240       | Not Expired |
+----------------+---------------------+----------------+-------------+
```

How to Schedule a Cron Job

To schedule a cron job:

1. To schedule a cron job for e-mail notifications about certificate expiry, type **1** at the prompt and press Enter. We recommend that you configure the SMTP server information in the **/usr/local/etc/smtp_server_details.json** file before proceeding to schedule the cron job.

   You can also schedule a cron job by using **./manage_certificate.sh -n** or **./manage_certificate.sh --notify**.

   ```
   Select a option (In Number) : 1

   Is /usr/local/etc/smtp_server_details.json file configured with proper SMTP server details?
   (y/n):
   ```

   - If you did not configure an SMTP server, type **n** and press Enter.

     ```
     Is /usr/local/etc/smtp_server_details.json file configured with proper SMTP server
     details? (y/n): n

     Kindly configure /usr/local/etc/smtp_server_details.json file with proper SMTP server
     ```

```
details.
Then retry scheduling cron job
```

Configure the SMTP server and run the manage_certificate.sh again to schedule the cron job.

- If an SMTP server is configured, type **y** and press Enter.

```
Is /usr/local/etc/smtp_server_details.json file configured with proper SMTP server
details? (y/n): y

Please select the cron tab operation

     1: list
     2: create
     3: delete

Select a option (In Number) :
```

Select any of the options available. You can choose to list all the cron jobs, create a new cron job, or delete a cron job.

2. To create a cron job, type 2 at the prompt and press Enter.

   Define a schedule for the cron job using the format * * * * *, which is a set of five values (that is *Minute, Hour, Day of the Month, Month, and Day of the Week*) in a line separated by spaces. Here are a few sample schedules:

   - Every hour: 0 * * * *
   - Every Monday at 10 PM: 0 22 * * 1

   The e-mail notification contains information such as component name, certificate expiry date, number of days left for certificate expiry, and status of the certificate.

```
Select a option (In Number) : 2
Please provide a cron schedule time in below format (space separated)

* * * * *
1 2 3 4 5

1. Minute (0 - 59)
2. Hour (0 - 23)
3. Day of month (1 - 31)
4. Month (1 - 12)
```

```
5. Day of week (0 - 7) (Sunday=0 or 7)
Schedule time: 0 * * * *
INFO     Scheduling cron job: 0 * * * * cd ~/Contrail_Service_Orchestration_6.2.0 && ./
manage_certificate.sh -n > /var/log/certificate.log
INFO     Successfully scheduled cron job
INFO     Current cron tab list:

0 * * * * cd ~/Contrail_Service_Orchestration_6.2.0 && ./manage_certificate.sh -n > /var/log/
certificate.log
```

3. To delete a cron job, type 3 at the prompt and press Enter.

```
Please select the cron tab operation

     1: list
     2: create
     3: delete

Select a option (In Number) : 3
INFO     Current cron tab list:

0 * * * * cd ~/Contrail_Service_Orchestration_6.2.0 && ./manage_certificate.sh -n > /var/log/
certificate.log

Please copy-paste the cron tab line here which you wants to delete:
```

At the prompt, copy and paste the cron schedule that you want to delete and press Enter.

```
Please copy-paste the cron tab line here which you wants to delete: 0 * * * * cd ~/
Contrail_Service_Orchestration_6.2.0 && ./manage_certificate.sh -n > /var/log/certificate.log
'0 * * * * cd ~/Contrail_Service_Orchestration_6.2.0 && ./manage_certificate.sh -n > /var/log/
certificate.log' will be deleted. Do you wish to continue ? (y/n): y
INFO     Successfully deleted cron job
INFO     Current cron tab list:
```

How to Renew a Certificate

You can renew a certificate only if its status is Expired or About to Expire.

NOTE: You can renew only self-signed certificates. Third-party certificates cannot be renewed.

At the prompt that appears when you run the **manage_certificate.sh** script, type the number representing the component for which you want to renew the certificate and press Enter.

```
Following component's certificate can be renewed

     2: Haproxy, Nginx, Rsyslog
     3: Telemetry Agent
Select a option (In Number) :
```

The system checks the status of the certificate:

- If the status is Expired or About to Expire, then the certificate renewal process is initiated. After the certificate renewal, the system performs a health check.

  NOTE: When HA proxy certificate is renewed, the telemetry agent certificate for all devices provisioned on CSO is automatically renewed.

  If HA proxy certificate is renewed and if the telemetry agent renewal cannot be completed due to a failure, then you can renew the telemetry agent certificate separately. Run the **manage_certificate.sh** script and provide the number corresponding to the Telemetry Agent (3 in the sample output) to renew the certificate.

- If the status is Not Expired, then the certificate is not renewed.

  Sample output if the status of a certificate is Not Expired:

```
************************************************************
 This tool assists you to renew CSO components certificate
************************************************************

Certificate renew sequence need to be followed:

     Kubernetes -> Haproxy -> Elasticsearch

     0: List all certificate expiry date
     1: Schedule cron for email notification
```

```
Following component's certificate can be renewed

     2: Haproxy, Nginx, Rsyslog
     3: Telemetry Agent

Select a option (In Number) : 2
INFO     Started check and renew haproxy component's certificate at 2021-12-27
02:19:10.974535 ...
INFO     Checking haproxy certificate expiry date
INFO     Checking nginx certificate expiry date
INFO     Checking rsyslog certificate expiry date
INFO     Haproxy certificate is Not Expired
INFO     Nginx certificate is Not Expired
INFO     Rsyslog certificate is Not Expired
INFO     Certificate is not about to expire, So renewal is not required
INFO     Completed check and renew haproxy component's certificate at 2021-12-27
02:19:13.638765 .
INFO     Time taken to renew haproxy component's certificate : 0:00:02.664230
```