CloudVision (CVP) Configuration Guide

Arista Networks

Version 2020.2.0

Pulse Supply
909 Ridgebrook Road.,Sparks,Maryland 21152,USA
TEL : +1-410-583-1701 FAX : +1-410-583-1704
E-mail: sales@pulsesupply.com
https://www.pulsesupply.com/datacom-systems

Headquarters	**Support**	**Sales**
5453 Great America Parkway, Santa Clara, CA 95054 Santa Clara, CA 95054 USA +1-408 547-5500	+1-408 547-5502 +1-866 476-0000	+1-408 547-5501 +1-866 497-0000
http://www.arista.com	*mailto:support@arista.com*	*mailto:sales@arista.com*

Contents

Chapter 1

Introduction to CloudVision

CloudVision™ is a turnkey solution for network-wide workload orchestration and work flow automation. It was specifically designed to complement SDN (virtualization) controller solutions that orchestrate virtual network overlays, by focusing on work flow visibility, automation tasks, and initial or ongoing network provisioning across the underlying physical network.

The CloudVision components are packaged as a virtual appliance and operate as a highly available cluster with role based privileges integrated into existing authentication tools (AAA, RADIUS, TACACS). For maximum operational flexibility, CloudVision can be managed with the interactive EOS CLI, the open eAPI for granular programmatic access, or a web-based portal interface.

CloudVision's foundation is an infrastructure service, sharing, and aggregating working state of physical switches running EOS to provide network visibility and central coordination. State from each participating EOS node is registered to CloudVision using the same publish/subscribe architecture of EOS's system database (SysDB). By communicating to each participating switch instance using a high performance binary API, CloudVision will actively synchronize state relevant to network-wide operational tasks. As an example, CloudVision's VXLAN Control Service aggregates network-wide VXLAN state for integration and orchestration with SDN controllers such as Openstack, VMWare NSX, and others.

The CloudVision web-based portal combines the most common operational tasks into a dashboard view decoupled from the underlying hardware. Workflow automation in CloudVision permits operators to execute common deployment and configuration tasks from a single visual touch point. The portal includes a turnkey solution for Arista's Zero Touch Provisioning (ZTP) and extends that from automating initial device provisioning to also include automating ongoing change controls over the operational life cycle of the device.

Using CloudVision, operators can organize devices in logical hierarchies through the use of list or configuration (config) container views for rapid categorization of device by role, type, or other specification. Configurations can be broken down into more manageable configlets that are built and stored directly on CloudVision, ready for network-wide or group-specific provisioning. The CloudVision database also keeps historical data, including a history of network state, configuration and software versions. This state can be used for taking a network-wide snapshot for change control verification of the network, helping to simplify the change management process and reduce maintenance window times.

For more information, see:

- CloudVision Portal (CVP) Overview
- CloudVision Portal (CVP) Setup
- Getting Started (CVP)

Chapter 2

CloudVision eXchange (CVX)

CloudVision eXchange (CVX) provides a single access point for real-time provisioning, orchestration and integration with third-party controllers. CVX aggregates and distributes operational state information across a set of EOS switches to support applications that provide network services.

Sections in this chapter include:

2.1 CVX Overview

A CVX deployment includes CVX and a set of CVX clients to which CVX provides services. CVX is not part of the data plane, nor does it receive data-path traffic. All CVX components exist as agents that run on EOS instances.

For more information, see:

- System Requirements
- CVX Infrastructure
- CVX Features
- CVX Clients

2.1.1 System Requirements

Certain hardware and software is required to be able to use CloudVision eXchange in your CloudVision virtual appliance implementation.

The CloudVision eXchange should be installed on a single system along with CloudVision Portal.

The following table lists the minimum hardware and software required to use CloudVision eXchange.

- System Requirements

Required Hardware

The hardware required to use the CloudVision eXchange are:

- CPU: 4 cores (base), 8 cores (recommended)
- RAM: 4G (base), 8G (recommended)
- Disk: 4G

Required Software

The software required to use the CloudVision eXchange are:

- EOS switches: Recommend 4.16.8M or later

Note: It is a best practice and highly recommended that the version of CVX should match the version running on the switches.

- CloudVision Portal: version 2016.1

(CloudVision Portal software is required if you want to use it in conjunction with CloudVision eXchange. If you plan to use only CloudVision eXchange, CloudVision Portal software is not required.)

Note: CVX supports live vMotion.

2.1.2 CVX Infrastructure

CVX provides a single integration point into network-wide services running across CVX clients. CVX is typically deployed as an EOS instance running on a VM (vEOS). The CVX infrastructure consists of a CVX instance functioning as a server and a set of CVX clients. The CVX server uses a heartbeat keepalive (KA) mechanism to maintain contact with its clients.

When de-configuring or shutting down CVX, client services should be shut down first.

2.1.3 CVX Features

CVX manages communications among the network CVX clients, and provides an integration point for services to those clients. CVX also discovers the physical network topology by aggregating topology information it receives from its client devices.

2.1.4 CVX Clients

CVX client is the agent that allows a switch to interact with a CVX server to access CVX services. Enabling the CVX client includes providing the IP address or host name of the device running CVX. The CVX client can then access services that are enabled on the CVX server.

The CVX client must be enabled to access the CVX server and the services it offers. Individual services may require additional configuration statements.

Services should be shut down or de-configured on clients before shutting down or de-configuring CVX. CVX-controlled switch features may continue to run after shutting down CVX if they are not explicitly shut down or de-configured prior to shutting down CVX.

2.2 CVX Services

CVX services are applications that run on top of the CVX infrastructure, and are accessed by CVX clients through the CVX server. All CVX services are maintained by version level; client switches negotiate the version they use when connecting to the server. This allows multiple switches that run different EOS versions to connect to the same CVX server.

The following sections briefly describe some of the services available to CVX clients through CVX:

- OpenStack Service
- VXLAN Control Service
- Hardware Switch Controller (HSC) Service
- Network Topology Service

2.2.1 OpenStack Service

The OpenStack service on CVX allows the networking component of an OpenStack deployment (also known as Neutron) to share state with CVX.

When deployed, this integration allows CVX to send state about the logical networks created in the OpenStack cloud to the CVX clients that configure the network.

More information on OpenStack software can be found in its online documentation at http://docs.openstack.org/.

2.2.2 VXLAN Control Service

The VXLAN control service allows hardware VXLAN tunnel end points (VTEPs) to share state with each other in order to establish VXLAN tunnels without the need for a multicast control plane. Configuration is required both on the client switches and in CVX.

2.2.3 Hardware Switch Controller (HSC) Service

Traffic between virtual machines which share a physical host (or between virtual machines and the rest of the network) is forwarded by virtual switches. The management and configuration of virtual switches uses the Open vSwitch Database (OVSDB) management protocol, as described in RFC 7047.

The hardware switch controller (HSC) service provides an integration point between OVSDB controllers and the VXLAN control service, allowing exchange of state information among virtual and hardware switches.

2.2.4 Network Topology Service

The network topology service gathers information from CVX clients to provide a view of the physical topology of the network. Aggregated information gathered by the network topology service is used by other CVX services, and can be viewed on the CVX server.

2.3 Deploying CVX

CloudVision Exchange (CVX) can be deployed on KVM and ESXi. The required EOS version and Aboot version vary depending on whether you are deploying CVX on KVM or ESXi.

For the detailed steps to use to deploy CVX, see:

- Deploying CVX on Kernel-based Virtual Machine (KVM)
- Deploying CVX on VMware ESXi

2.3.1 Deploying CVX on VMware ESXi

Complete the following steps to install CVX on ESXi. Once the installation is complete, you can begin the CVX configuration process.

Note: Make sure you select versions of EOS that meet the minimum requirements for CVX. The supported version is EOS (version 4.21.0or later).

Complete the following steps to install CVX.

1. Go to:http://www.arista.com.
2. Select **Support** > **Software Download**.
3. From the software download page, expand **Active Releases** > **4.21** > **EOS-4.21.0F** to download **EOS-4.21.0F.vmdk**.
4. Load the files you downloaded into a filestore location within the VMware vSphere environment.

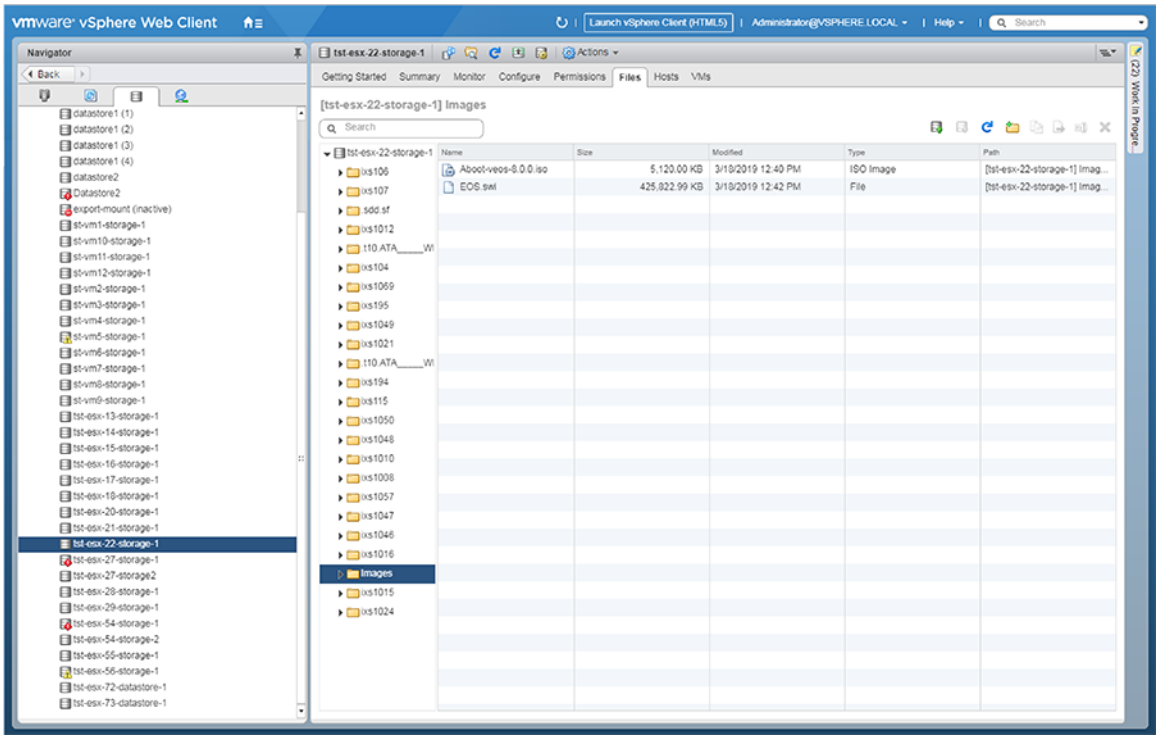

Figure 1: Loading the files into the VMware vSphere environment

5. Right-click the **filestore location** you selected, and choose **New Virtual Machine** .

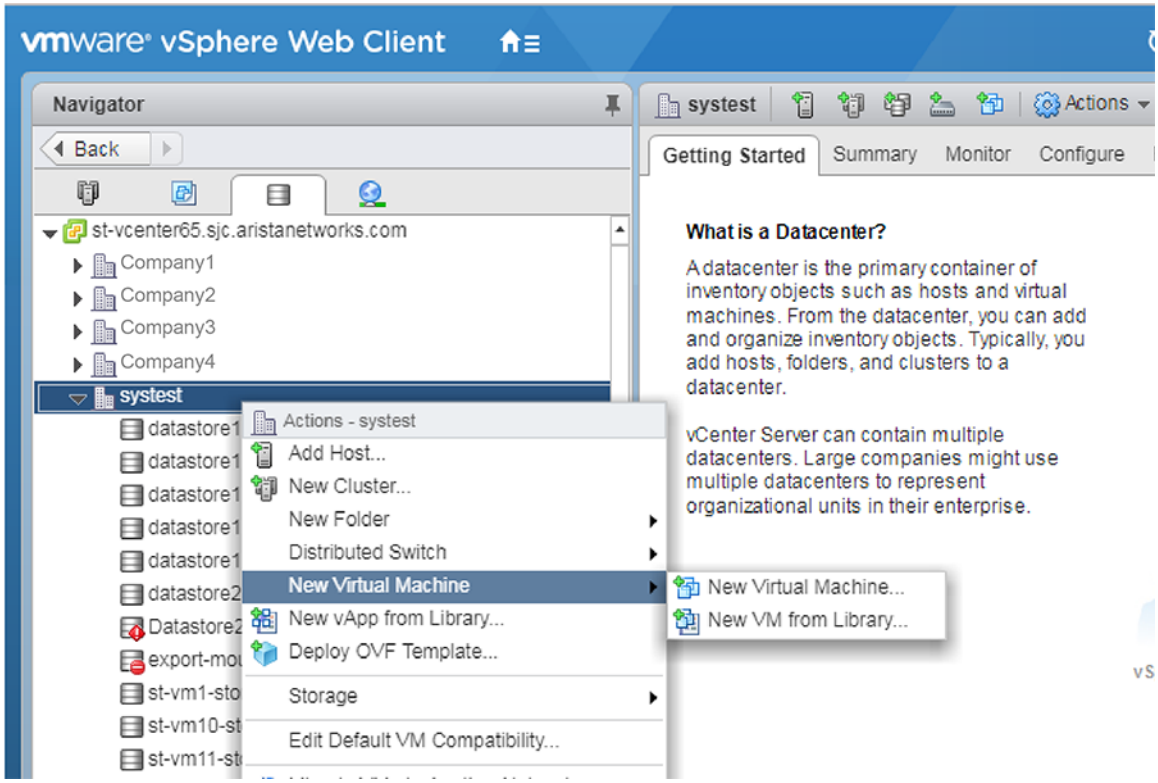

Figure 2: Selecting New Virtual Machine

The New Virtual Machine dialog appears.

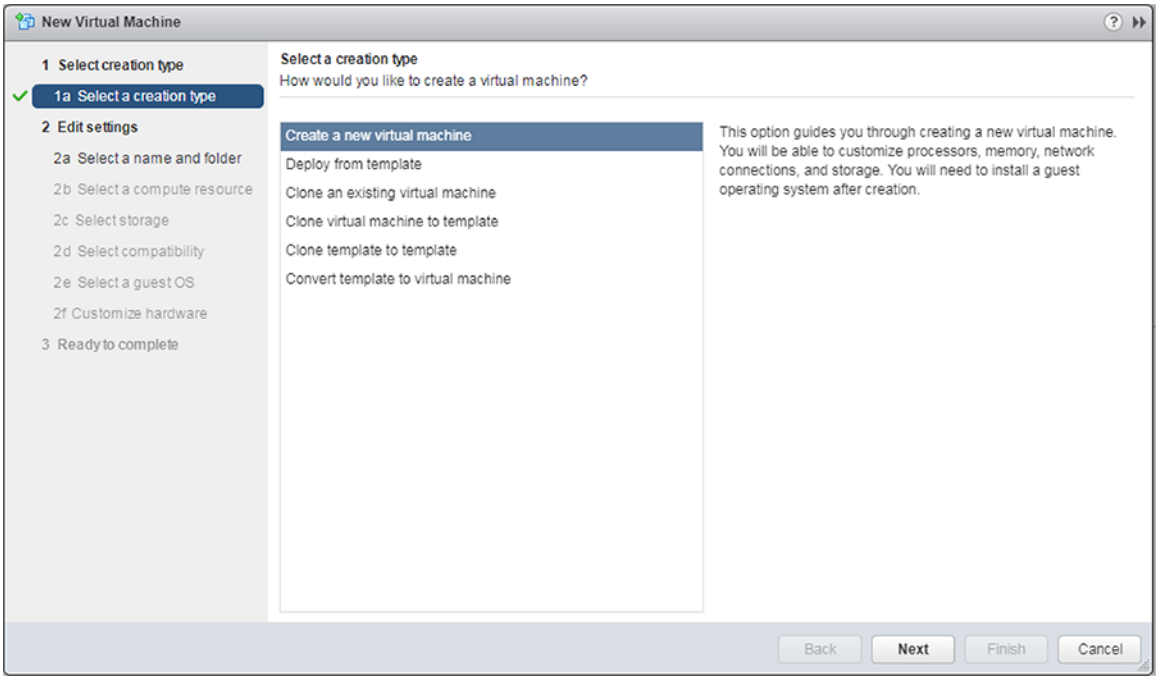

Figure 3: New Virtual Machine dialog

6. In the New Virtual Machine dialog, select **Create a new virtual machine**, and then click **Next**.

   The dialog refreshes, showing options for the new Virtual Machine.

   New Virtual Machine dialog (naming and selecting the location)

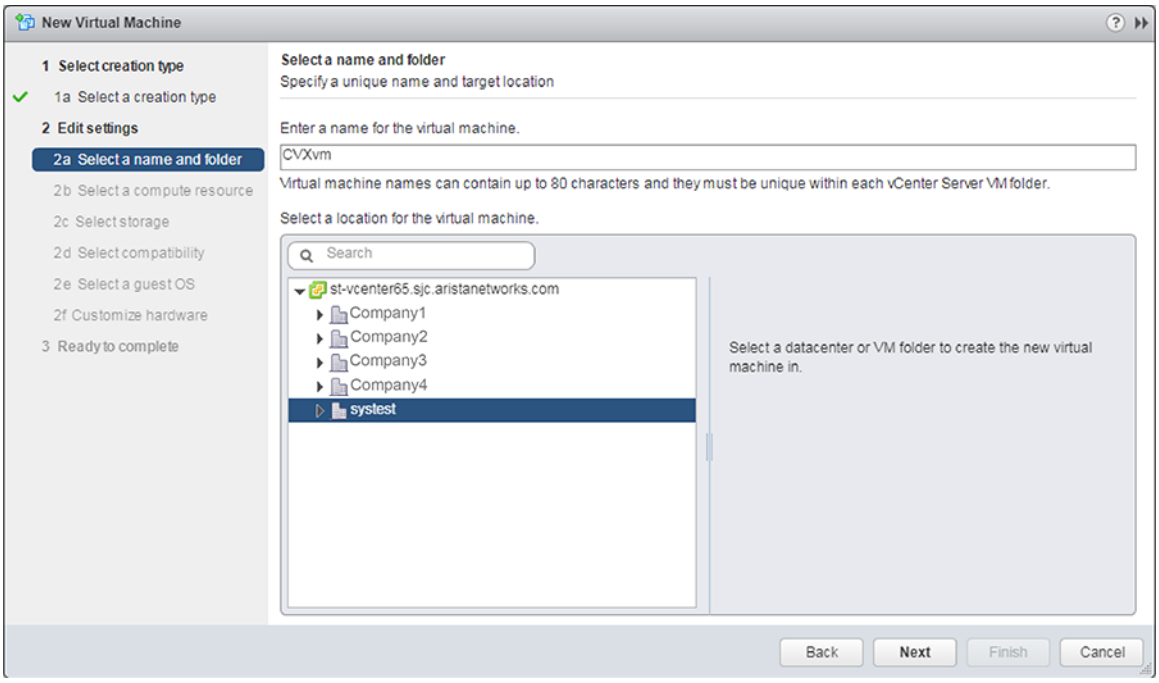

Figure 4: New Virtual Machine dialog (naming and selecting the location)

The dialog refreshes, showing options for selecting the datastore.

7. Enter a **name** for the new Virtual Machine.
8. Select a **location** for the new Virtual Machine, then click **Next**.

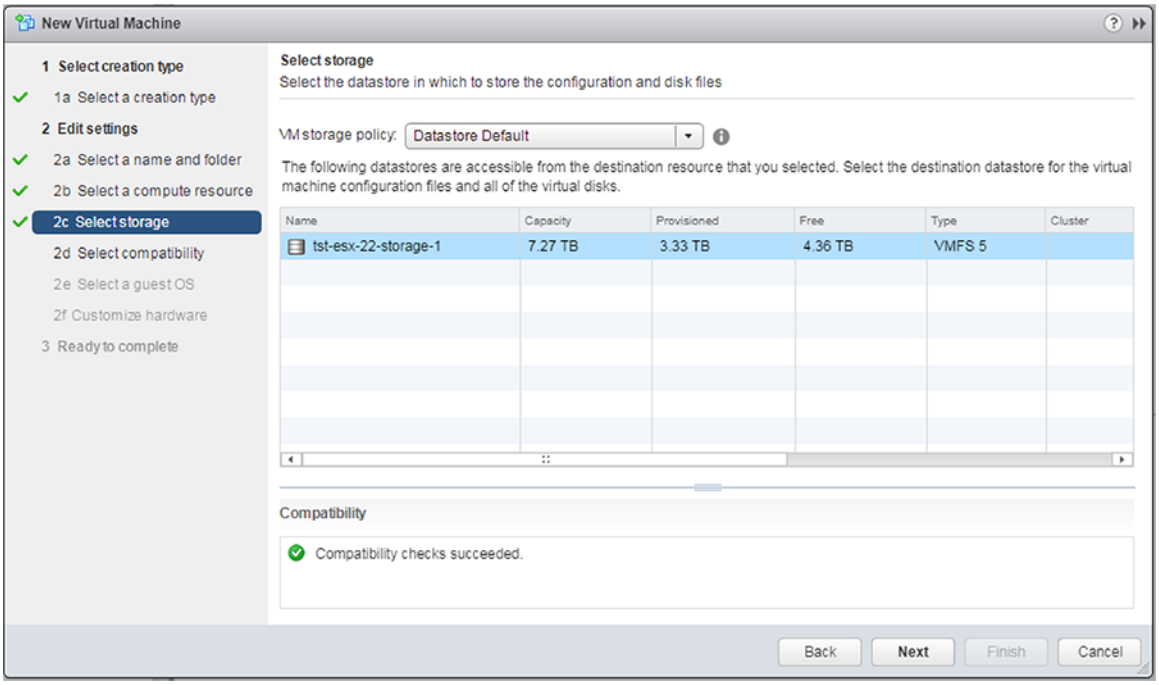

Figure 5: New Virtual Machine dialog (selecting the datastore)

9. Select the **datastore** for the new Virtual Machine configuration files and all of the virtual disks.
   - Click **Next**.

   The dialog refreshes, showing operating system selection options.
10. Click **Next**.

The dialog refreshes, showing compatibility options.

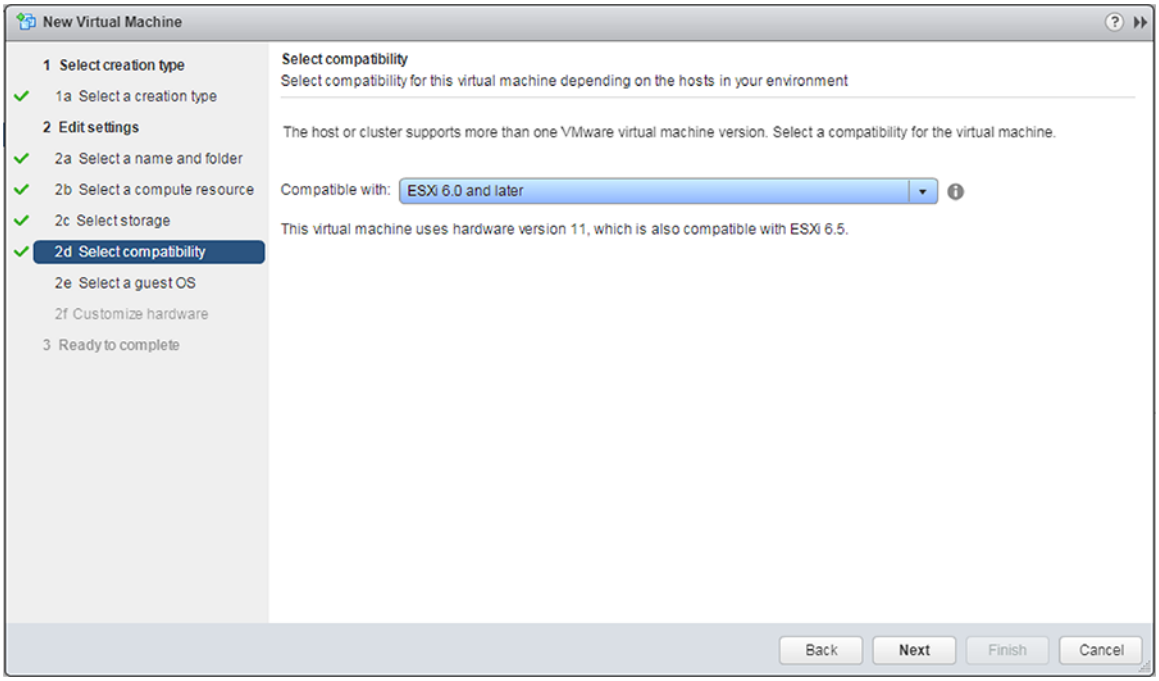

Figure 6: New Virtual Machine dialog (compatibility options)

11. Using the Compatible with menu, select the **ESXi compatibility** for the new Virtual Machine.

    Note: When adding the VMDK to ESX6, it treats this as sparse by default, whereas in ESX 5 it is thick. Converting the vEOS VMDK file from thin to thick would allow it to boot properly in ESX6: `vmkfstools -i vEOS-lab-4.18.5M.vmdk -d eagerzeroedthick vEOS-lab-4.18.5M-thick.vmdk`.

    Go to https://eos.arista.com/ and refer to the following topics for the issue and solution:

    - Tip for Arista vEOS on VMware ESX 6
    - Common Issues When Deploying CVX *4.18.2F* on vCenter 6 or 6.5

    Note: If the VM keeps rebooting and showing “This is not a bootable disk. Please insert a bootable floppy and press any key to try again", then go to https://eos.arista.com/ and refer to the **Common Issues When Deploying CVX 4.18.2F on vCenter 6 or 6.5** topic.

12. Click **Next**.

    The dialog refreshes, showing operating system selection options.

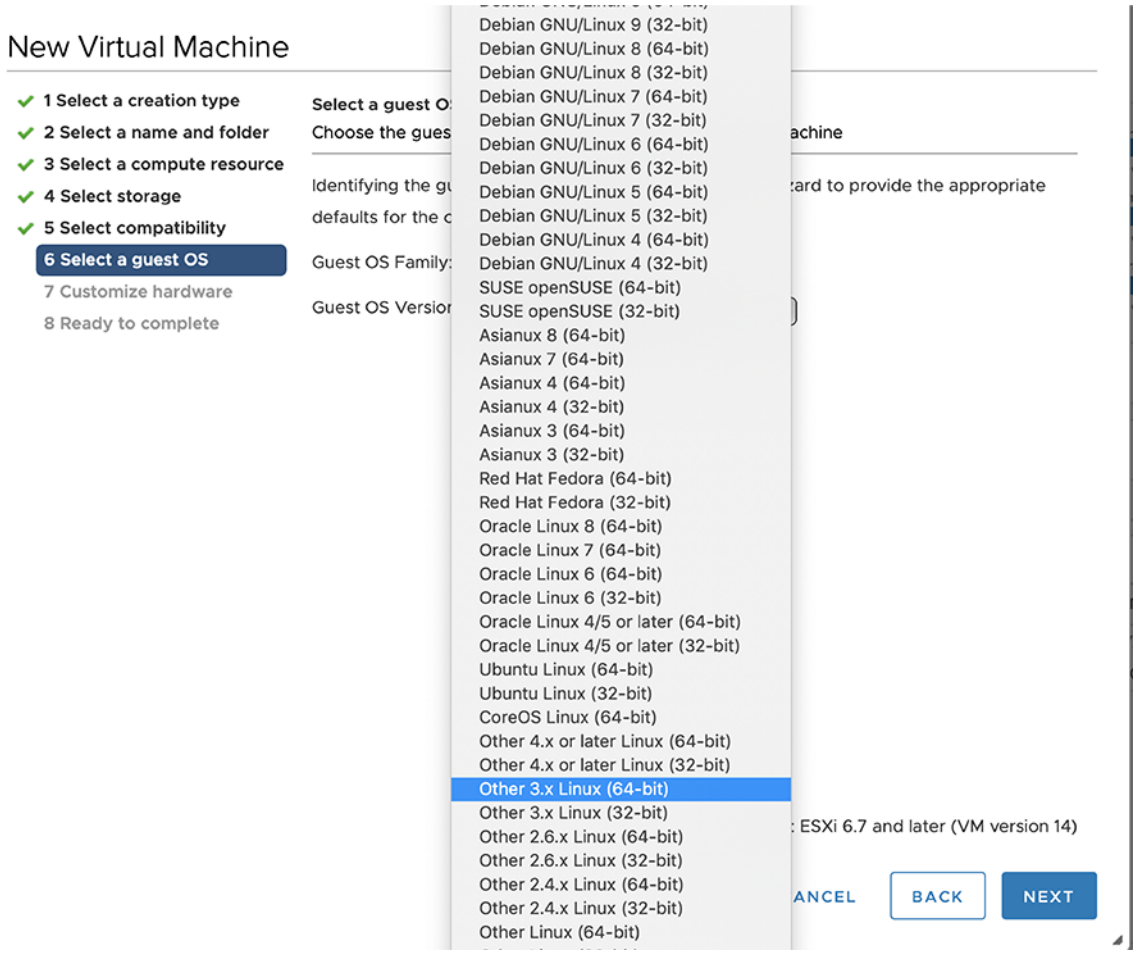

Figure 7: New Virtual Machine dialog (operating system options)

13. Using the Guest OS Family menu, choose **Linux**.
14. Using the Guest OS Version menu, choose **Other Linux (64-bit)**.
15. Click **Next**.

    The dialog refreshes, showing options for customizing hardware.

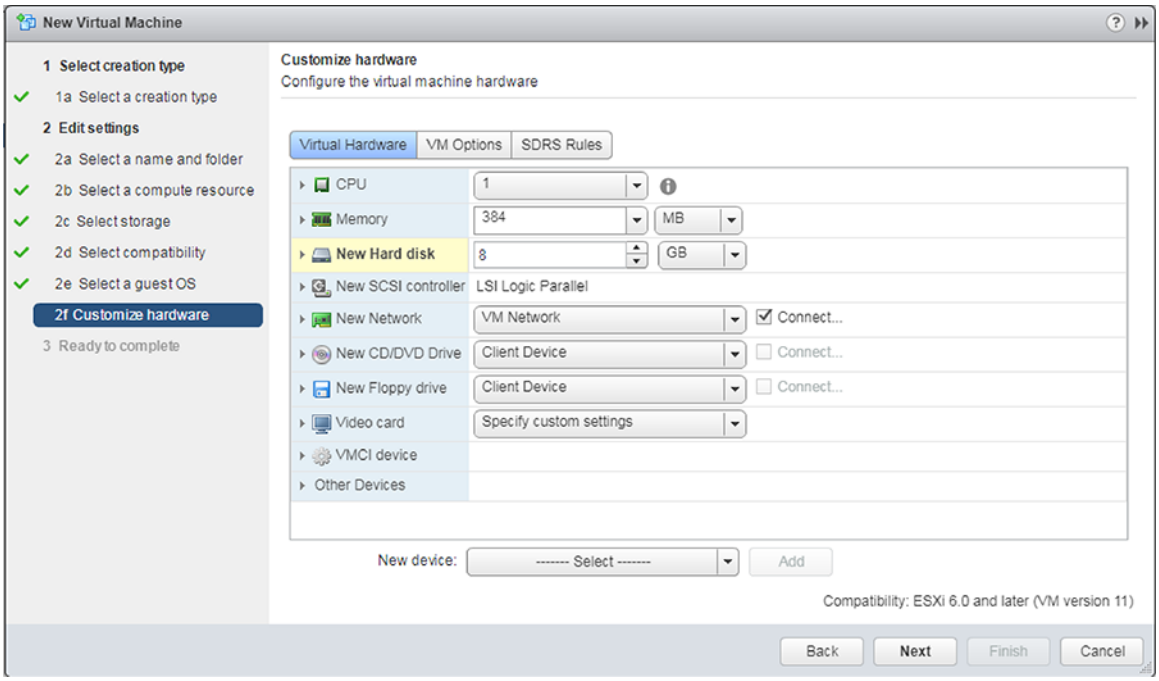

Figure 8: New Virtual Machine dialog (hardware configuration options)

16. Change the default settings for the following options:

CPU	Set to **4** (number of CPUs)
Memory	Set to **8 GB**
New Hard Disk	Delete the current setting (leave this option empty).
New Network	Specify connection to Network LAN segment with connectivity to CVX client devices (the Management LAN). Choose VMXNET3 network adapter type. This connection is used for CVX client / server communications.
Existing Hard Disk	Specify the **EOS-4.21.0F.vmdk** you downloaded in step 3.

17. (Optional) Delete the floppy drive and SCSI controller.
18. Click **Next**.

You are now ready to begin the CVX configuration (seeCVX Configuration).

2.3.2 Deploying CVX on Kernel-based Virtual Machine (KVM)

Complete the following steps to install CVX on Ubuntu/KVM. Once the installation is complete, you can begin the CVX configuration process.

Note: Make sure you select versions of EOS and Aboot that meet the minimum requirements for CVX. The supported versions are:

- EOS (version *4.16.8M* or later)
- `Aboot-veos-serial-8.0.0.iso` (located in the vEOS section of the download)

Pre-requisites

Before you begin the procedure, make sure that:

- Install **qemu-kvm**, **libvirt*,** and all related dependencies using yum (RHEL7/CentOS7) and apt-get (Ubuntu).

- Two bridges are configured for use by the KVM VM, and that you have the names of the bridges. (Steps are included in the procedure to add bridges, if they are not already configured.)

  Note: The bridges must be configured to persist (**brctl** commands do no persist across reboots). You can use Network Manager (or another application available to you) to complete this configuration.
- You have both **generateXmlForKvm.py** and **cvpTemplate.xml**. They are required to complete the procedure. You can find them in the CVP tarball for Ubuntu.

Complete the following steps to install CVX.

1. Download the Aboot and EOS files from: https://www.arista.com/en/support/software-download/.
2. Use **sudo su** to acquire superuser privileges, which are required to complete some of the installation steps.
3. Confirm that KVM is running on the server by entering the following command:

   ```
   virsh -c qemu:///system listAb
   ```

   The command output should match this example:

   ```
   Id    Name     State
     ------------------------------
     $
   ```
4. If the output does not look correct (previous step) go to for additional assistance: https://help.ubuntu.com/community/KVM/Installation.
5. Use the following command to convert the **vmdk** file to qcow2: `qemu-img convert EOS_4_16_8M.vmdk -O qcow2 EOS.qcow2`

   Note: Step **6** and **7** are required if you do not already have 2 bridges defined in different subnets. If the bridges exist, go directly to step **8**.
6. Use **brctl** to add bridges for the KVM VM to use (br1 and br2 can be any names you choose).

   ```
   brctl addbr br1
   brctl addbr br2
   ```

   ifconfig can be used to identify Ethernet ports to be bridged. Once you identify the ports, add them to the bridges.

   Example:

   ```
   brctl addif br1 enx803f5d086eae
   ```
7. Confirm that the bridges are up using **brctl show**.
   - Enter: ifconfig br1 up
   - And: ifconfig br2 up

   Note: The following step uses a number of input parameters (the number required vary depending on your server setup). To ensure the command executes successfully, we recommend that you type it into a scratch pad and edit as needed before typing it into the Linux Terminal.
8. Use the following command to generate **cvx.xml**, which will be used to setup the CVX VM.
   `generateXmlForKvm.py`

   Example:

   ```
   python generateXmlForKvm.py -n cvx --device-bridge br1 --cluster-bridge
    br2 -e /usr/bin/kvm -i cvpTemplate.xml -c /home/myname/Downloads/Abo
   ot-veos-serial-8.0.0.iso -x /home/myname/Downloads/EOS.qcow2 -b 8192 -p
    2 -t

    -n cvx: VM name.
   ```

```
--device-bridge br1: This is the name you gave the bridge - br1 or
anything else.
--cluster-bridge br2: Cluster bridge if clustering servers.
-i cvpTemplate.xml: Path to XML file input template.
-k: VM ID number used by virsh. If not entered, a random number is
assigned.
-b 8192: 8G of RAM.
-p 2: # of CPU cores.
-c: Path to Aboot file.
-x: Path to qcow2 file created in step 3.
-t: This parameter indicates the file defined by -x is for CVX.
-e '/usr/bin/kvm': Ubuntu path to KVM.
(for RHEL KVM this is: -e 'usr/libexec/qemu-kvm')
-o: XML file used by virsh to define the KVM VM.
```

9. Run the following commands:

```
virsh define cvx.xml
virsh start cvx
virsh console cvx
```

10. (Optional) To configure CVX to start automatically, enter:

```
virsh autostart cvx
```

You are now ready to begin the CVX configuration (seeCVX Configuration).

2.4 CVX Configuration

CVX, its clients, and its services, are independently configured. These sections describe configuration processes for each:

- Ports Used by CVX
- CVX Server Configuration
- CVX Client Configuration
- CVX Client Services Configuration

2.4.1 Ports Used by CVX

CVX uses the following ports:

- Controller database (Controllerdb): Port 9979
- Client-server out-of-band connection: Port 50003
- CVX cluster peer out-of-band connection: Port 50004

Note: All of these connections are TCP.

2.4.2 CVX Server Configuration

Enabling CVX on the CVX Server

CVX parameters for the server infrastructure are configured in **CVX configuration** mode. CVX configuration mode is not a group-change mode; running-config is changed when commands are entered, and exiting the mode does not modify running-config. The `cvx` command places the switch in **CVX configuration** mode.

CVX is disabled by default. The `no shutdown (CVX)` command enables CVX on the switch.

Example

These commands enter CVX-configuration mode and enable CVX.

```
switch(config)#cvx
switch(config-cvx)#no shutdown
switch(config-cvx)#
```

CVX Heartbeat Configuration

CVX synchronizes with its client devices by exchanging heartbeat signals. The heartbeat transmission frequency and timeout period determine when a client's access to the server is disrupted.

The interval between heartbeat messages that the server transmits is specified by the `heartbeat-interval (CVX)` command. The CVX timeout period is specified by the `heartbeat-timeout (CVX)` command. When CVX does not receive a subsequent heartbeat message from a CVX client before the timeout expiry, the server discontinues CVX services to that client.

Best practices dictate that CVX and its client applications configure identical heartbeat interval and heartbeat timeout values.

Example

These commands configure a CVX heartbeat interval of 30 seconds and a server heartbeat timeout period of 90 seconds.

```
switch(config-cvx)#heartbeat-interval 30
switch(config-cvx)#heartbeat-timeout 90
switch(config-cvx)#
```

Disabling CVX on the CVX Server

Note: Before disabling or de-configuring CVX on the CVX server, CVX client services should be explicitly disabled or shut down. Failure to disable or de-configure services prior to disabling or de-configuring CVS may result in CVX features continuing to run after CVX shutdown.

When disabling the CVX service, service VXLAN configuration may be retained or erased. Be sure to disable or shut down client services prior to disabling the CVX service.

Examples

- These commands shut down the CVX service while retaining the CLI configuration for service VXLAN.

```
localhost(config)#cvx
localhost(config-cvx)#service vxlan
localhost(config-cvx-vxlan)#shutdown
```

- These commands shut down the CVX service and also erase service VXLAN CLI configuration.

```
localhost(config-cvx-vxlan)#
localhost(config)#cvx
localhost(config-cvx)#no service vxlan
```

2.4.3 CVX Client Configuration

This section describes the CVX client configuration and commands that enable CVX services. Most commands for the configuration of the CVX client infrastructure are accessed in Management-CVX configuration mode.

- Enabling CVX on the CVX Client

CVX client parameters are configured in **Management-CVX configuration** mode. Management-CVX configuration mode is not a group-change mode; running-config is changed when commands

are entered, and exiting the mode does not modify running-config. The `management cvx` command places the switch in **Management-CVX configuration** mode.

CVX client is disabled by default. The `no shutdown (Management-CVX)` command enables CVX client on the switch.

For the CVX network topology service to create an inventory of all CVX clients, ensure that LLDP is enabled on each client switch using the `lldp run` command.

Example

These commands enter **Management-CVX-configuration** mode and enable the CVX client.

```
switch(config)#lldp run
switch(config)#management cvx
switch(config-mgmt-cvx)#no shutdown
switch(config-mgmt-cvx)#
```

- **CVX Client Heartbeat Configuration**

A CVX client synchronizes and maintains contact with CVX by exchanging heartbeat signals. The heartbeat transmission frequency and timeout period define when communication with CVX will be considered down.

The interval between heartbeat messages that the CVX client transmits is configured by the `heartbeat-interval (Management-CVX)` command.

The CVX client timeout period is specified by the `heartbeat-timeout (Management-CVX)` command. When a CVX client does not receive a subsequent heartbeat message from CVX within this timeout period, the client assumes that services provided by CVX are no longer available.

Best practices dictate that a CVX client's heartbeat interval and heartbeat timeout values are identical to those of the CVX server to which it connects.

Example

- This command configures a CVX client heartbeat interval of 30 seconds and client timeout period of 90 seconds.

```
switch(config-mgmt-cvx)#heartbeat-interval 30
switch(config-mgmt-cvx)#heartbeat-timeout 90
switch(config-mgmt-cvx)#
```

- Connecting the CVX Client to a Server

The `server host (Management-CVX)` command identifies the location of the CVX server that the client accesses. The `source-interface (Management-CVX)` command specifies the interface from which the client derives the IP address it uses as the source in CVX packets that it transmits. And the `no shutdown (Management-CVX)` command enables CVX on the client switch.

Example

- These commands configure the switch as a CVX client, connecting to a CVX server at IP address *10.1.1.14* and using IP address *10.24.24.1* as the source address for its outbound packets.

```
switch(config)#interface loopback 5
switch(config-if-Lo5)#ip address 10.24.24.1/24
switch(config-if-Lo5)#management cvx
switch(config-mgmt-cvx)#server host 10.1.1.14
switch(config-mgmt-cvx)#source-interface loopback 5
switch(config-mgmt-cvx)#no shutdown
switch(config-mgmt-cvx)#
```

2.4.4 CVX Client Services Configuration

Switches running EOS must be configured as CVX clients to access the network services running on CVX. Individual services may require additional configuration.

Refer to the following for information regarding the services available to a CVX client.

- Configuring OpenStack Service
- Configuring VXLAN Control Service
- Configuring Hardware Switch Controller Service (HSC)
- Configuring Network Topology Service

2.4.4.1 Configuring OpenStack Service

The OpenStack service is enabled from CVX-OpenStack configuration mode, which is accessed by the **`service openstack`** command. The **`no shutdown (CVX-OpenStack)`** command enables CVX OpenStack services on the CVX server. Additional configuration is necessary to deploy OpenStack; http://docs.openstack.org/.

Example

- These commands enable the CVX-OpenStack service.

```
switch(config-cvx)#service openstack
switch(config-cvx-openstack)#no shutdown
switch(config-cvx-openstack)#
```

2.4.4.2 Configuring VXLAN Control Service

The VXLAN control service is enabled on CVX by the **`no shutdown (CVX-VXLAN)`** command and on the client switches by enabling CVX and configuring the VXLAN as a controller client. When VXLAN control service is enabled, CVX functions as a VXLAN controller for its clients.

For information about configuring VXLAN on the client switch, see the VXLAN chapter of the *User Manual* .

Examples

- These commands enable VXLAN control service on the CVX server.

```
switch(config-cvx)#service vxlan
switch(config-cvx-vxlan)#no shutdown
switch(config-cvx-vxlan)#
```

- These commands enable VXLAN Control Service on the CVX client. (This example assumes that the VXLAN has already been configured on the client switch. For information about configuring VXLAN, see the VXLAN chapter of the *User Manual*).

```
switch(config)#interface vxlan 1
switch(config-if-Vx1)#vxlan controller-client
```

2.4.4.3 Configuring Hardware Switch Controller Service (HSC)

The hardware switch controller (HSC) service is enabled on the CVX server by the `no shutdown (CVX-HSC)` command.

Certificate Requirements for CVX Interoperability with VMware NSX *6.2.2* and higher

The certificate type needs to be changed from MD5 to SHA512 for use with VMware NSX *6.2.2*. Complete the following steps to make the change.

1. At the EOS prompt of CVX, use the following commands.

```
switch(config)#cvx
switch(config-cvx)#service hsc
switch(config-cvx-hsc)#shut
```

2. Acquire superuser privileges and edit the default.

```
switch(config)#bash
switch(config)#sudo su
switch(config)#vi /usr/bin/ovs-pki
```

3. Find and replace **default_md** with **sha512 (**from **md5)**

```
default_md =md5
default_md =sha512
```

4. Delete all files and folders from `/persist/secure/openvswitch/`

```
cd /persist/secure/openvswitch/
bash-4.1#sudo rm -r *
```

5. Generate the new certificate.

```
[admin@CVX ~]$ exit
logout
CVX(config-cvx-hsc)#no sh
CVX(config-cvx-hsc)#end
```

6. Verify the change using the command:

```
CVX# show nsx status
```

Example

- These commands enable the CVX-HSC service.

```
switch(config)#cvx
switch(config-cvx)#no shutdown
switch(config-cvx)#service hsc
switch(config-cvx-hsc)#no shutdown
```

The HSC service sends flood lists to each VTEP through CVX. Some controllers (such as VMware NSX's Service Nodes) implement replication nodes for head-end replication of unknown packets. For these controllers, BUM packets should be sent to a single replication node (send-to-any replication), and the flood list sent by the HSC service is a list of replication nodes. Other controllers (such as Nuage VSP) require each VTEP to perform its own head-end replication. For these, BUM packets should be sent to every known VTEP, and the flood list sent by the HSC service is the list of VTEPs.

The default behavior is to use a send-to-any replication list of VTEPs. If the required behavior is send-to-all replication of, use the all option of the `vtep (CVX-HSC)` command.

Example

- This command configures the CVX-HSC service to use send-to-any replication.

```
switch(config-cvx-hsc)#vtep flood list type all
switch(config-cvx-hsc)#
```

Note: HSC also makes use of the VXLAN control service; ensure that VXLAN control service is enabled and properly configured (see VXLAN Control Service for details).

HSC also requires a connection to an OVSDB controller. Configure the IP address or host name of the controller using the `manager` command.

Example

- This command configures the CVX-HSC service to connect to an OVSDB controller at IP address 192.168.2.5, using the default port 6632.

```
switch(config-cvx-hsc)#manager 192.163.2.5
switch(config-cvx-hsc)#
```

Having established a connection to the OVSDB controller, the HSC service will publish the inventory of switches managed by CVX to OVSDB. For the inventory to succeed, LLDP must be enabled on each CVX client switch with the `lldp run` command.

Note: LLDP is enabled by default on Arista switches.

Example

- This command enables LLDP.

```
switch(config)#lldp run
switch(config)#
```

2.4.4.4 Configuring Network Topology Service

A network topology agent runs on each Arista switch whether or not the switch is connected to a CVX server. It requires no configuration. The network topology service on the CVX server is also enabled by default and requires no configuration.

To view the aggregated topology information, use the **`show network physical-topology`** command on the switch running the CVX server instance.

Examples

- This command displays all visible hosts.

```
switch#show network physical-topology hosts
Unique Id            Hostname
-------------------- ------------------------------
001c.7385.be69       cvx287.sjc.aristanetworks.com
0000.6401.0000       cvc1
0000.6402.0000       cvc2
0000.6403.0000       cvc3
0000.6404.0000       cvc4
bcf6.85bd.8050       dsj14-rack14-tor1
```

- This command displays all connections in the topology.

```
switch#show network physical-topology neighbors
cvx287.sjc.aristanetworks.com
Interface          Neighbor Intf      Neighbor Host
------------------ ------------------ ----------------------
Ethernet1          Ethernet7          cvc4
Ethernet2          Ethernet7          cvc2
Ethernet9          Ethernet7          cvc1
Ethernet10         Ethernet7          cvc3
Management1        27                 dsj14-rack14-tor1

OUTPUT OMITTED FROM EXAMPLE
  dsj14-rack14-tor1

Interface          Neighbor Intf      Neighbor Host
```

```
------------------ ------------------ ------------------------
27                 Management1        cvx287.sjc.aristanetwork
```

2.5 CVX Secure out-of-band Connection

This feature adds support for securing out-of-band connection between CVX server and CVX clients by SSL/TLS transport protocol. SSL/TLS is an application-layer protocol that provides secure transport between client and server through a combination of authentication, encryption and data integrity. SSL/TLS uses certificates and private-public key pairs to provide this security. We will use the term SSL to mean SSL/TLS.

By default, CVX server and CVX clients communicate over insecure transport (there is no authentication and encryption between CVX server and CVX clients). This poses the possibility of security risks, such as communicating with untrusted CVX server and CVX clients, or eavesdropping CVX server/client communications. This feature can be used to secure the out-of-band connection between CVX server and CVX clients.

Note: The CVX client-server out-of-band connection uses port 50003. The CVX cluster peer out-of-band connection uses port 50004. These are TCP ports.

For more information, see:

- Configuring the CVX Secure out-of-band Connection
- Show Commands
- Troubleshooting

2.5.1 Configuring the CVX Secure out-of-band Connection

This feature uses SSL certificate and key management infrastructure for managing certificates, keys and SSL profiles. For more information regarding this infrastructure see *SSL Certificate and Key Management* in the *Arista User's Guide*.

1. On CVX server, copy the server certificate and key and also the CA certificate to verify CVX clients.

```
switch(config)#!Copy the PEM encoded certificate and RSA key files for
 CVX server
switch(config)#!Lets call them server.crt and server.key
switch(config)#copy <url> certificate:server.crt
switch(config)#copy <url> sslkey:server.key
switch(config)#!Copy the PEM encoded CA certificate to verify the
 certificate of CVX clients.Lets call it ca.crt
switch(config)#copy <url> certificate:ca.crt
```

2. On CVX server, configure SSL profile with the certificates and key as below. Lets call the SSL profile as "serverssl".

```
switch(config)#management security
switch(config-mgmt-security)#ssl profile serverssl
switch(config-mgmt-sec-ssl-profile-serverssl)#certificate server.crt
 key server.key
switch(config-mgmt-sec-ssl-profile-serverssl)#!You can trust multiple
 CA certificates
switch(config-mgmt-sec-ssl-profile-serverssl)#trust certificate ca.crt
```

Note: If you are using intermediate certificates to build a 'Chain of Trust' (such as **server.crt -> intermediate1.crt -> intermediate2.crt -> ca.crt**), then you need to configure the intermediate certificates as part of the SSL profile using the following commands:

```
switch(config-mgmt-sec-ssl-profile-serverssl)#chain certificate
 intermediate1.crt
switch(config-mgmt-sec-ssl-profile-serverssl)#chain certificate
 intermediate2.crt
```

3. On CVX server, configure to use the “serverssl” SSL profile. With this configuration, the CVX server starts listening on a secure port. The CVX server will continue to listen on the default port. i.e., the CVX server will accept connections from CVX clients over both SSL and default non-SSL transports. During a SSL negotiation, the CVX server will authenticate itself to the CVX clients by presenting ‘server.crt’ and it verifies the authenticity of the CVX client by checking if the CVX client certificate is signed by the trusted certificate “ca.crt”.

```
switch(config)#cvx
switch(config-cvx)#ssl profile serverssl
```

4. On CVX client, copy the client certificate and key and also the CA certificate to verify CVX server.

```
switch(config)#!Copy PEM encoded certificate and RSA key files for CVX
 client
switch(config)#!Lets call them client.crt and client.key
switch(config)#copy <url> certificate:client.crt
switch(config)#copy <url> sslkey:client.key
switch(config)#!Copy PEM encoded CA certificate used to verify the
switch(config)#!certificate of CVX server. Lets call it ca.crt
switch(config)#copy <url> certificate:ca.crt
```

Note: If you are using intermediate certificates to build a 'Chain of Trust' (such as **client.crt -> intermediate1.crt -> intermediate2.crt -> ca.crt**), then you need to configure the intermediate certificates as part of the SSL profile using the following commands:

```
switch(config-mgmt-sec-ssl-profile-clientssl)#chain certificate
 intermediate1.crt
switch(config-mgmt-sec-ssl-profile-clientssl)#chain certificate
 intermediate2.crt
```

5. On CVX client, configure SSL profile with the certificates and key as below. Lets call the SSL profile as “clientssl”.

```
switch(config)#management security
switch(config-mgmt-security)#ssl profile clientssl
switch(config-mgmt-sec-ssl-profile-clientssl)#certificate client.crt
 key client.key
switch(config-mgmt-sec-ssl-profile-clientssl)#!You can trust multiple
 CA certificates
switch(config-mgmt-sec-ssl-profile-clientssl)#trust certificate ca.crt
```

6. On CVX client, configure to use the SSL profile – “clientssl”. With this configuration, the CVX client will connect to the secure port of the CVX server over SSL transport. During SSL negotiation, the CVX client will authenticate itself to the CVX server by presenting ‘client.crt’ and it verifies the authenticity of the CVX server by checking if the CVX server certificate is signed by the trusted certificate ‘ca.crt’.

```
switch(config)#management cvx
switch(config-mgmt-cvx)#ssl profile clientssl
```

2.5.2 Show Commands

For information regarding show commands of SSL certificate, key and profile, please refer to *SSL Certificate and Key Management*.

To show the SSL profile status on CVX server, use the `show cvx` command.

```
switch#show cvx

CVX Server
 Status: Enabled
 UUID: beb19142-dfaa-11e4-b996-001c73105347
 Heartbeat interval: 20.0
 Heartbeat timeout: 60.0
 SSL profile: serverssl
  Status: Enabled
```

The “Enabled” SSL status means that the SSL profile is enabled for CVX server and the CVX clients can connect to CVX server over SSL transport. If there are any errors, then the status will show “Disabled” and the reason will be listed. In ‘Disabled’ state, the CVX clients wont be able to connect to CVX server over SSL transport.

To show the SSL connection status of CVX clients on CVX server, use the show cvx connections command.

```
switch#show cvx connections

Switch 00:1c:73:10:53:48
 Hostname: sq302
 Status: up
Last heartbeat sent: 0:00:04 ago
Last heartbeat received: 0:00:10 ago
Clock offset: -0.00201620385865
Out-of-band connection: SSL secured
In-band connection: Not secured (SSL not supported)
```

The out-of-band connection shows as “SSL secured”, which means that the CVX client has connected to CVX server over SSL transport. The in-band connection is another connection between CVX server and CVX client. The SSL is not yet supported for this connection and hence it shows as ‘SSL not supported’. There is already some level of protection for the in-band connection. The CVX server and CVX client opens up the access to in-band connection only if the out-of-band connection is successful. Since the out-of-band connection is configured to use SSL, the in-band connection access is granted only for authentic CVX client and CVX server.

To show SSL profile status and connection status on CVX client, use the `show management cvx` command.

```
switch#show management cvx

CVX Client
 Status: Enabled
 Last connected time: 2015-04-14 11:16:19
 Connection status: Connected
  Out-of-band connection: SSL secured
  In-band connection: Not secured (SSL not supported)
 Negotiated version: 2
 Controller UUID: 0e7dee2e-e2cf-11e4-880f-001c73105347
 Controller: 127.0.0.1
  Last heartbeat sent: 0:00:00 ago
  Last heartbeat received: never
  Clock offset: 0.0
```

```
SSL profile: clientssl
 Status: Enabled
```

The “Enabled” SSL status means that the SSL profile is enabled and the CVX client can connect to CVX server over SSL transport. If there are any errors, then the status will show as “Disabled” and the reason will be listed. In Disabled state, the CVX client won’t be able to connect to the CVX server.

Similar to the CVX server, the out-of-band connection shows as “SSL secured” and the SSL is not yet supported for in-band connection.

The possible reasons for ‘Disabled’ SSL status on CVX server and CVX client are:

- **SSL profile does not exist:** If the SSL profile configured under CVX server/client is not configured under ‘management security’, you will see this message. Configure the SSL profile with required certificates and key under ‘management security’.
- **Invalid SSL profile:** If the SSL profile configured under CVX server/client is in ‘invalid’ state, you will see this message. Check `show management security ssl profile <name>` command to see the errors on the SSL profile and fix them.
- **Trusted certificates not configured in SSL profile:** If the SSL profile configured under CVX server/client does not have trusted certificates configured, you will see this message. Please configure trusted CA certificates in the SSL profile.
- **Certificate not configured in SSL profile:** If the SSL profile configured under CVX server/client does not have certificate key pair configured, you will see this message. Please configure certificate and key pair in the SSL profile.

Diffie-Hellman parameters not yet ready: When EOS is booted, a Diffie-Hellman parameters file is auto generated by the system if one does not exist. This Diffie-Hellman parameters file is used for symmetric key exchange during SSL negotiation. Only the CVX server uses this file and hence this message can be seen only on `show cvx` command output. If the file is not yet generated, you will see this message. When the file is ready, this message automatically goes away and the SSL profile will become ‘Enabled’.

2.5.3 Troubleshooting

Check `show cvx` on the CVX server and see if the SSL profile is in “Enabled” state. If it’s in “Disabled” state, check the reason listed and fix it.

Check “show management cvx” on CVX client and see if SSL profile is in “Enabled” state. If it’s in “Disabled” state, check the reason listed and fix it.

2.6 CVX High Availability

CVX provides high availability by enabling you to use multiple (redundant) CVX Controllers in the same cluster. Each Controller in the cluster has its own dedicated machine so that if a Controller fails, the failure is isolated to a single machine.

Within a cluster, one of the Controllers is a primary (leader), and the other Controllers are backup (follower) Controllers. If the primary Controller fails, one of the backup Controllers automatically assumes the role of the primary Controller.

CVX high availability does not prevent or compromise the detection of software failures or link failures that may cause Controllers to be unreachable on the network.

The configuration that is required to ensure CVX is set up for high availability involves:

- Configuring the CVX cluster.
- Configuring the CVX clients.

For more information, see:

- CVX Clusters
- Handling of CVX Controller Failures
- CVX Support for EOS Failure Modes
- Client Interaction
- Service Agents Interaction
- Leader Election

2.6.1 CVX Clusters

CVX clusters are sets of CVX Controllers (usually 3 Controllers). Within a cluster, each Controller runs on its own dedicated machine, and all of the Controllers run the same version of CVX. Each Controller in the cluster functions as either the primary (leader) Controller, or a backup (follower) Controller.

One of the CVX Controllers is elected by the group of Controllers to be the primary Controller. Once a Controller is elected to be the primary, the other Controllers in the cluster are automatically assigned the role of backup Controllers. Cluster members maintain an out-of-band connection amongst themselves, which is used for the leader election protocol.

CVX Controllers in a cluster that are not the primary Controller always function as backup Controllers. Within the same cluster, only one CVX Controller can assume the role of a primary at any time.

For more information, see:

- Required Number of Controllers to Support High Availability
- Cluster Configuration Options

2.6.1.1 Required Number of Controllers to Support High Availability

A cluster must have enough Controllers so that in the case of a failure of the primary Controller, there are enough remaining Controllers for the election process to be completed. The election process is used by clusters to select a new primary Controller in the case of failure.

Note: The number of Controllers for a cluster is **3** (one primary and two backup Controllers).

Examples

In a cluster with only **two** Controllers (one primary and one backup), a simple majority of backup Controllers does not exist after a failure of the primary Controller. A simple majority of two backup Controllers is required for the leader election process.

Related Topics

- Cluster Configuration Options
- Handling of CVX Controller Failures
- CVX Support for EOS Failure Modes
- Client Interaction
- Service Agents Interaction
- Leader Election

2.6.1.2 Cluster Configuration Options

You can configure the cluster for high availability using either of the following modes:

- Cold followers mode - Only the Controllerdb of the primary (leader) CVX Controller mounts from the client switches.
- Warm followers mode - The Controllerdb of every (all) CVX Controllers in the cluster mount from the client switches.

Advantages and disadvantages of the modes

The advantage of the warm follower mode is that if the primary CVX Controller fails, the switchover to the new primary is faster than a switchover in cold follower mode. The reason for this is that the state of the new primary does not have to be rebuilt from scratch. The disadvantage of the warm follower mode is that serialization from the switch is slower compared to cold follower mode.

Related Topics

- Required Number of Controllers to Support High Availability
- Handling of CVX Controller Failures
- CVX Support for EOS Failure Modes
- Client Interaction
- Service Agents Interaction
- Leader Election

2.6.2 Handling of CVX Controller Failures

CVX Controllers can fail because of hardware or software faults. Because EOS agents are designed to be software fault-tolerant, an agent that fails is automatically restarted and resumes operation statefully. The most recent saved state in Sysdb for the agent is used to restore the state of the agent.

Unlike software failures, hardware failures are not handled by EOS. CVX handles hardware failures through the use of redundant backup (follower) CVX Controllers that run on their own dedicated machine. Within a cluster, any backup Controller can assume the role of the primary (leader) Controller.

Note: In the event of a network partition, the partition with a majority of the Controllers elects a leader from its Controllers, and the minority partition relinquishes any leadership it might have had.

Related Topics

- CVX Clusters
- CVX Support for EOS Failure Modes
- Client Interaction
- Service Agents Interaction
- Leader Election

2.6.3 CVX Support for EOS Failure Modes

CVX supports both EOS failure modes that apply when a CVX Controller fails. The EOS failure modes are:

- Fail-stop
- Fail-recover

Because CVX supports both EOS failure modes, a failed CVX Controller can rejoin the cluster if the following failures occur:

- A crash of the agent or machine running CVX.
- The CVX controller or dedicated machine it runs on is removed (partitioned) from the cluster.

Related Topics

- Handling of CVX Controller Failures
- CVX Clusters
- Client Interaction
- Service Agents Interaction
- Leader Election

2.6.4 Client Interaction

Client switches maintain an out-of-band connection to all members of the cluster. The connection is used to determine liveness and for communications. The connection is also used to signal a change in leadership (switchover) to the client switches. Switchovers that are changes in leadership within a cluster are executed similarly to CVX Graceful Reboot switchovers.

The ControllerClient agent on the switch is responsible for maintaining liveness with the Controllers and for exchanging metadata. The ControllerClient agent registers with all cluster members. Each Controller's ControllerStatus has an additional flag to record whether the Controller is a leader within the cluster.

If there is more than one leader, the switch automatically waits until only one Controller is designated as the leader in the cluster. Once a single Controller is designated as the leader, the switch executes a graceful switchover to the new leader Controller.

Related Topics

- Handling of CVX Controller Failures
- CVX Clusters
- CVX Support for EOS Failure Modes
- Client Interaction
- Service Agents Interaction
- Leader Election

2.6.5 Service Agents Interaction

One change to Service Agents is required to support CVX high availability. Service Agents must be modified to include the leader flag (this flag identifies the leader CVX Controller in the cluster). On a leader switchover, Service Agents are deactivated on the old leader Controller and activated on the new leader Controller. The client switches will perform a graceful switchover to the new leader Controller.

Related Topics

- Handling of CVX Controller Failures
- CVX Clusters
- CVX Support for EOS Failure Modes
- Client Interaction
- Leader Election

2.6.6 Leader Election

Leader election is an internal, system-run process that is essential to CVX high availability. The leader election process is used to safely elect a new leader Controller within a cluster following the failure of the current leader Controller, or a network configuration change that results in the loss of the current leader Controller in the cluster.

The leader election process is designed to ensure stability of leader Controllers within clusters. The process is based on an algorithm that provides the mechanism for the backup (follower) Controllers to elect (by consensus), the new leader Controller in the cluster.

2.6.7 Configuring CVX Clusters for High Availability

Configuring CVX clusters for high availability is a simple process that involves pointing each cluster member to the other cluster members using the peer host command. The objective of this task is to successfully register each cluster member with the other cluster members. Successful registration of

the cluster members with each other ensures that the members can communicate with each other to elect a new leader member if the original leader member fails.

Once you complete the process, the cluster members will be successfully registered with each other. In addition, the cluster members will automatically elect a leader member and assign the 'leader' to that member. The non-leader members are automatically assigned the role of 'follower'.

Requirements

The requirements for setting up clusters for high availability are:

- The number of CVX Controllers in a cluster is **3**.
- An **odd number** of CVX instances (CVX Controllers) are required to form a cluster.

Note: If an even number of CVX Controllers are configured in a cluster, a CVX instance will automatically refuse to participate in the cluster.

- All cluster members must point to each other. This is essential for clusters to operate normally. (The steps required to complete this task are included in the following procedure.)

Procedure

Note: This procedure provides configuration examples for each step. The 'example' cluster used throughout the procedure contains 3 cluster members (named *cvs1*, *cvs2*, and *cvs3*). The IP addresses of the cluster members are:

- *cvs1* (*10.0.0.1*)

- *cvs2* (*10.0.0.2*)

- *cvs3* (*10.0.0.3*).

Complete the following steps to configure clusters for high availability.

1. Using the peer host command, configure one of the cluster members to point to every other cluster member.

   This example shows the configuration of cluster member ***cvs1*** to point to the other cluster members (*cvs2* and *cvs3*).

```
cvs1(config-cvx)#peer host 10.0.0.2 (connects cvs1 to cvs2)
cvs1(config-cvx)#peer host 10.0.0.3 (connects cvs1 to cvs3)
```

2. Use the sh cvx command to check the **Mode** and **Peer registration state** status values for cluster member *cvs1*. The status values should be:

   - **Mode** = *Cluster*
   - **Peer registration state** = *Connecting*

   

   Note: Mode automatically changes from "Standalone" to "Cluster" when configuring a CVX cluster. This is because the presence of multiple CVX "peers" causes the Mode to change to "Cluster".

   Peer registration state remains in "Connecting" status after you configure the first cluster member. This is because the two peers must register with each other for the registration of the two members to be successful.

3. Using the peer host command, configure peer cluster member *cvs2* to point to every other cluster member.

   This example shows the configuration of cluster member ***cvs2*** to point to the other cluster members (*cvs1* and *cvs3*).

```
cvs2(config-cvx)#peer host 10.0.0.1 (connects cvs2 to cvs1)
cvs2(config-cvx)#peer host 10.0.0.3 (connects cvs2 to cvs3)
```

4. Use the sh cvx command to check the **Peer registration state** settings for *cvs1*. This is done to verify that peers *cvs1* and *cvs2* are successfully registered with each other.

```
cvs1(config-cvx)#sh cvx
```

Example

This example shows the output of the sh cvx command for *cvs1*. The **Peer registration state** setting of “Registration Complete” for peer *cvs2 indicates a successful registration between cvs1 and cvs2*.

```
cvs1(config-cvx)#sh cvx

CVX Server
 Status: Enabled
 UUID: 6c208fba-7324-11e5-8fef-1d98cdd3b27a
 Mode: Cluster
 Heartbeat interval: 20.0
 Heartbeat timeout: 60.0
 Cluster Status
  Name: default
  Role: Standby
  Leader: 10.0.0.2
  Peer timeout: 10.0
  Last leader switchover timestamp: 0:00:03 ago
  Peer Status for 10.0.0.3
   Peer registration state: Connecting
   Peer service version compatibility : Version mismatch
  Peer Status for 10.0.0.2
   Peer Id : 02-01-63-02-00-00
   Peer registration state: Registration complete
   Peer service version compatibility : Version ok
```

5. Using the peer host command, configure peer cluster member *cvs3* to point to every other cluster member.

This example shows the configuration of cluster member ***cvs3*** to point to the other cluster members (*cvs1* and *cvs2*).

```
cvs3(config-cvx)#peer host 10.0.0.1 (connects cvs3 to cvs1)
cvs3(config-cvx)#peer host 10.0.0.2 (connects cvs3 to cvs2)
```

6. Use the sh cvx command to check the **Peer registration state** settings for *cvs1*. This is done to verify that peers *cvs1* and *cvs3* are successfully registered with each other.

cvs1(config-cvx)#sh cvx

Example

This example shows the output of the sh cvx command for *cvs1*. The **Peer registration state** setting of “Registration Complete” for peer *cvs3 indicates a successful registration between cvs1 and cvs3*.

```
cvs1(config-cvx)#sh cvx

CVX Server
 Status: Enabled
 UUID: 6c208fba-7324-11e5-8fef-1d98cdd3b27a
 Mode: Cluster
 Heartbeat interval: 20.0
 Heartbeat timeout: 60.0
 Cluster Status
  Name: default
```

```
Role: Standby
Leader: 10.0.0.2
Peer timeout: 10.0
Last leader switchover timestamp: 0:05:37 ago
Peer Status for 10.0.0.3
 Peer Id : 02-01-63-03-00-00
 Peer registration state: Registration complete
 Peer service version compatibility : Version ok
Peer Status for 10.0.0.2
 Peer Id : 02-01-63-02-00-00
 Peer registration state: Registration complete
 Peer service version compatibility : Version ok
```

Next Steps

You are now ready to configure the CVX clients for high availability (see Configuring CVX Clients for High Availability).

2.6.8 Configuring CVX Clients for High Availability

Configuring CVX clients for high availability is a simple process that involves pointing each CVX client to every CVX cluster member using the server host command. The objective of this task is to successfully establish connections between each CVX client and every CVX cluster member. The connections are essential to ensure that the CVX clients are aware of the current status of each cluster member.

Note: If a CVX client is not pointing to every cluster member, or if it is pointing to a CVX instance (Controller) that is not part of the cluster, the client may not be aware of leadership changes in the cluster, or may become confused about which cluster member is currently the leader. Either of these scenarios can result in unexpected errors.

Once you complete the process, the CVX clients will have established connections with each cluster member (the Connection status for each Controller should be ‘Established’). In addition, the clients will be aware of which CVX instance (Controller) is currently the leader in the cluster.

Procedure

Note: This procedure provides configuration examples for each step. The ‘example’ CVX client used throughout the procedure is named cvc1. The IP addresses of the cluster members are *10.0.0.1* (*cvs1*), *10.0.0.2* (*cvs2*), and *10.0.0.3* (*cvs3*).

Complete the following steps to configure CVX clients for high availability.

1. Using the server host command, configure each of the CVX clients to point to every cluster member.

   This example shows the configuration of client ***cvc1*** to point to all of the cluster members (the addresses of the cluster members are *10.0.0.1*, *10.0.0.2*, and *10.0.0.3*).

```
cvc1(config-mgmt-cvx)#server host 10.0.0.1 (connects cvc1 to cluster
 member 10.0.0.1)
cvc1(config-mgmt-cvx)#server host 10.0.0.2 (connects cvc1 to cluster
 member 10.0.0.2)
cvc1(config-mgmt-cvx)#server host 10.0.0.3 (connects cvc1 to cluster
 member 10.0.0.3)
```

2. Use the sh man cvx command to check the status of client ***cvc1***.

   The Connection status for each cluster member should be “Established”. In addition, the client is also aware that cluster member *10.0.0.3* is the current Master.

```
cvc1(config-mgmt-cvx)#sh man cvx
```

```
CVX Client
 Status: Enabled
 Source interface: Inactive (Not configured)
 Controller cluster name: default
  Controller status for 10.0.0.1
   Connection status: established
    Out-of-band connection: Not secured
    In-band connection: Not secured (SSL not supported)
   Negotiated version: 2
   Controller UUID: 6c208fba-7324-11e5-8fef-1d98cdd3b27a
   Last heartbeat sent: 0:00:07 ago
   Last heartbeat received: 0:00:07 ago
 Controller status for 10.0.0.3
  Master since 0:03:34 ago
  Connection status: established
   Out-of-band connection: Not secured
   In-band connection: Not secured (SSL not supported)
 Negotiated version: 2
 Controller UUID: c64954b8-7324-11e5-9f33-51f8b016cae8
 Last heartbeat sent: 0:00:14 ago
 Last heartbeat received: 0:00:14 ago
Controller status for 10.0.0.2
 Connection status: established
  Out-of-band connection: Not secured
  In-band connection: Not secured (SSL not supported)
 Negotiated version: 2
 Controller UUID: 6a0dbf2c-7324-11e5-94f3-ff17a8a1cdc8
 Last heartbeat sent: 0:00:05 ago
 Last heartbeat received: 0:00:05 ago
```

2.7 CVX VIP

CVX VIP provides the virtual IP address that actively follows the master controller of the CVX cluster.

The virtual IP address of the CVX HA Cluster is configured on a macvlan interface setup on top of a physical management interface of the master controller. The virtual IP and virtual MAC needs to be provided by the customer as part of the controller configuration. This information is available to all controllers as each cluster member has to be configured manually by the user on all controllers.

The macvlan interface created should be designated as `Management0`. `Management0` is currently used for the ManagementActive interface on modular switches. Without explicit configuration of VIP and VMAC, CVX VIP functionality will not work in the CVX HA cluster.

Customers can pick the VMAC from a pool of MAC addresses reserved for use with CVX clusters. The OUI pool, 00:1C:73:00:00:AA # 00:1C:73:00:00:FF has been reserved for this purpose.

The macvlan interface is setup if all of the following conditions are met:

- VMAC is configured by the user
- The controller instance is a leader
- There are more than one controller instances
- The controller is not being run on a modular system

2.7.1 Configuring vip

All CLI commands applicable to the management interface of the controller will be allowed on `Management0`, with the exception of layer 1 / phy level commands. So auto-negotiation or flow control can#t be configured on the `Management0` interface. Instead these commands can only be run on the physical management interfaces. This makes sense as the phy-level configuration really depends on what the interface is physically wire

To configure VMAC/VIP

```
CVX(config)#interface management 0
CVX(config-if-Ma0)# mac-address 00:1C:72:00:00:FF
CVX(config-if-Ma0)# ip address 10.0.0.2
```

2.7.2 Data Replication

At EOS boot time, SSH host keys and Diffie-Hellman parameters are automatically generated and persistently stored on each controller. Multiple SSL profiles / keys / certificates might also be created and used by various agents on the controllers. Since these information contribute to the identity of the master, they will need to follow the master controller for all time.

In case of a controller switchover, the newly elected master controller will need to use the same SSH host keys & SSL profiles / keys / certificates to retain its identity and prevent any kind of network security alarms from being tripped. For example, if an SSH client notices that the host key has changed, it will normally flag an error warning the user of a possible man-in-the-middle type attack. Hence, this data will be replicated from the master to slaves.

2.7.3 SSH Host Key Tagging

SSH host keys are tagged with the chassis MAC address to deal with key regeneration issues when a supervisor module is moved from one chassis to another. This behavior will cause regeneration issues if we replicate the SSH host keys across the cluster resulting in the key fingerprint seen by management tools to be different.

To mitigate this, in addition to the chassis MAC address, the host keys would now be tagged with VMAC of the CVX HA cluster. If CVX VIP and VMAC are configured, SshHostKeysAgent will not regenerate keys if tagged VMAC and configured VMAC are the same, even if there is a mismatch between the chassis MAC and tagged MAC.

2.8 Upgrading CVX

You can upgrade CVX from a previous version to the current version by performing a few simple tasks. You can use the following procedure to upgrade any previous version of CVX to the current version.

- **Requirements**

Make sure you follow these requirements during the upgrade process.

- If you have CVP, CVX and client switches in your environment, make sure you upgrade each component in the following order:
- Upgrade CVP first
- Upgrade the CVX cluster.
- Upgrade the client switches. The reason for this is to ensure backward compatibility.
- You must upgrade the CVX cluster before you upgrade the client switches.
- If the CVX cluster is a 3 node cluster, make sure that only one node of the cluster is down at any one time during the upgrade process. (The order in which you upgrade the nodes does not matter.)

Pre-requisites

Before you begin the upgrade, make sure that:

- You perform a backup to ensure that you can restore data if needed.
- You download the latest version of CVX from Arista#s Software Download page (https://www.arista.com/en/support/software-download).

Complete the following steps to upgrade CVX.

1. Login to the cluster to be upgraded. (You can login to any node.)
2. Upgrade the node. You must deploy a new image to perform the upgrade.
3. Wait for the node you are upgrading to rejoin the cluster. Once the node has rejoined, go to the next step. (The node automatically rejoins the cluster as a follower node.)
4. Repeat steps **1 through 3** to upgrade the two remaining nodes one node at a time. It does not matter the order in which you upgrade the remaining nodes.

2.9 CVX Command Descriptions

CVX Server Commands

- cvx
- heartbeat-interval (CVX)
- heartbeat-timeout (CVX)
- port (CVX)
- show cvx
- shutdown (CVX)

CVX Client Commands

- management cvx
- heartbeat-interval (Management-CVX)
- heartbeat-timeout (Management-CVX)
- server host (Management-CVX)
- source-interface (Management-CVX)
- shutdown (Management-CVX)

CVX OpenStack Commands

- name-resolution force (CVX-OpenStack)
- name-resolution interval (CVX-OpenStack)
- service openstack
- shutdown (CVX-OpenStack)

CVX VXLAN Control Service Commands

- resync-period
- service vxlan
- shutdown (CVX-VXLAN)
- vtep (CVX-VXLAN)

CVX Hardware Switch Controller (HSC) Commands

- manager
- ovsdb-shutdown
- shutdown (CVX-HSC)
- vtep (CVX-HSC)

CVX Network Topology Service Commands

- lldp run
- show network physical-topology

Related Topics

- CVX Overview
- CVX Services
- Deploying CVX
- CVX Configuration
- CVX Secure out-of-band Connection
- CVX High Availability

2.9.1 cvx

CVX (CloudVision eXtension) aggregates and shares status across a network of physical switches running EOS. CVX services provide visibility and coordinate activities across a network of switches that are configured as CVX clients.

The **cvx** command enters CVX configuration mode. CVX configuration mode is not a group-change mode; **running-config** is changed immediately upon entering commands. Exiting CVX configuration mode does not affect **running-config**. The **exit** command returns the switch to global configuration mode.

The **no cvx** and **default cvx** commands restore all CVX server defaults by deleting all CVX configuration mode statements from running-config.

Command Mode

Global Configuration

Command Syntax

```
cvx
no cvx
default cvx
```

Commands Available in CVX Configuration Mode

- port (CVX)
- service openstack
- service vxlan
- shutdown (CVX)
- heartbeat-interval (CVX)
- heartbeat-timeout (CVX)

Example

- These commands enter CVX-configuration mode and display the CVX configuration.

```
switch(config)#cvx
switch(config-cvx)#show active all

 cvx
  shutdown
  port 9979
  heartbeat-interval 20
  heartbeat-timeout 60
  no service vxlan
  service openstack
   shutdown
   name-resolution interval 21600
switch(config-cvx)#
```

2.9.2 heartbeat-interval (CVX)

The **heartbeat-interval** command configures the interval between heartbeat messages that the switch sends as a CVX server. Heartbeat messages are part of the keepalive mechanism between CVX and the CVX clients to which it connects.

The **no heartbeat-interval** and **default heartbeat-interval** commands restore the heartbeat interval to the default setting by removing the **heartbeat-interval** command from **running-config**.

Command Mode

CVX Configuration

Command Syntax

```
heartbeat-interval period
no heartbeat-interval
default heartbeat-interval
```

Parameters

- **period** Interval duration (seconds). Value ranges from 5 through 60. Default value is 20.

Related Commands

- **cvx**
- **heartbeat-timeout (CVX)**

Guidelines

Heartbeat messages flow independently in both directions between CVX and clients. When a client stops receiving heartbeat messages from the server within a specified period, the client assumes that the CVX server is no longer functioning.

Best practices dictate that CVX and its client applications configure identical heartbeat interval values.

Examples

- This command configures a CVX server heartbeat interval of 30 seconds:

```
switch(config)#cvx
switch(config-cvx)#heartbeat-interval 30
switch(config-cvx)#
```

2.9.3 heartbeat-interval (Management-CVX)

The heartbeat-interval command configures the interval between heartbeat messages that the switch sends as a CVX client. Heartbeat messages are part of the keepalive mechanism between the CVX client and the CVX server to which it connects.

The no heartbeat-interval and default heartbeat-interval commands revert the heartbeat interval to the default setting by removing the heartbeat-interval command from running-config.

Command Mode

Mgmt-CVX Configuration

Command Syntax

heartbeat-interval period

no heartbeat-interval

default heartbeat-interval

Parameters

- period: Interval duration (seconds). Value ranges from 5 through 60. Default value is 20.

Guidelines

Heartbeat messages flow independently in both directions between CVX and clients. When the server stops receiving heartbeat messages from a client within a specified period, the server assumes that the device it is no longer functioning as a CVX client.

Best practices dictate that the CVX client's heartbeat interval value is identical to that of its CVX server.

Related Commands

- management cvx places the switch in **Mgmt-CVX configuration** mode.
- heartbeat-timeout (Management-CVX) specifies the CVX client timeout interval.

Examples

- These commands configure a CVX client heartbeat interval of 30 seconds:

```
switch(config)#management cvx
switch(config-mgmt-cvx)#heartbeat-interval 30
switch(config-mgmt-cvx)#
```

2.9.4 heartbeat-timeout (CVX)

The **heartbeat-timeout** command specifies the CVX timeout period. When a CVX server does not receive consecutive heartbeat messages from a CVX client within the heartbeat timeout period, the server discontinues providing CVX services to the client device. The default timeout period is 60 seconds.

The **no heartbeat-timeout** and **default heartbeat-timeout-timeout** commands restore the heartbeat timeout to the default setting by removing the heartbeat-timeout command from running-config.

Command Mode

CVX Configuration

Command Syntax

heartbeat-timeout period

no heartbeat-timeout

default heartbeat-timeout

Related Commands

- cvx places the switch in CVX configuration mode.
- heartbeat-interval (CVX) specifies the CVX heartbeat interval.

Parameters

- period: heartbeat timeout interval (seconds). Value ranges from 15 to 10800. Default value is 60.

Guidelines

Best practices dictate that CVX and its client applications configure identical heartbeat timeout values.

Examples

- These commands set the CVX timeout period to 90 seconds.

```
switch(config)#cvx
switch(config-cvx)#heartbeat-timeout 90
switch(config-cvx)#
```

2.9.5 heartbeat-timeout (Management-CVX)

The heartbeat-timeout command specifies the CVX client timeout period. When a CVX client does not receive consecutive heartbeat messages from a CVX server within the period specified by this command, the client assumes that its connection to CVX is disrupted. The default timeout period is 60 seconds.

The no heartbeat-timeout and default heartbeat-timeout-timeout commands restore the CVX client heartbeat timeout to the default setting by removing the heartbeat-timeout command from running-config.

Command Mode

Mgmt-CVX Configuration

Command Syntax

heartbeat-timeout period

no heartbeat-timeout

default heartbeat-timeout

Parameters

- period: heartbeat timeout interval (seconds). Value ranges from 15 to 10800. Default value is 60.

Guidelines

Best practices dictate that the CVX client's heartbeat timeout value is identical to that of its CVX server.

Related Commands

- management cvx places the switch in Mgmt-cvx configuration mode.
- heartbeat-interval (Management-CVX) specifies the CVX client heartbeat interval.

Examples

- These commands set the CVX client timeout period to 90 seconds.

```
switch(config)#management cvx
switch(config-mgmt-cvx)#heartbeat-timeout 90
switch(config-mgmt-cvx)#
```

2.9.6 lldp run

The **`lldp run`** command enables LLDP on the Arista switch.

Command Mode

Global Configuration

Command Syntax

```
lldp run

no lldp run

default lldp run
```

Examples

- This command enables LLDP globally on the Arista switch.

```
switch(config)# lldp run
switch(config)#
```

- This command disables LLDP globally on the Arista switch.

```
switch(config)# no lldp run
switch(config)#
```

2.9.7 management cvx

The **`management cvx`** command places the switch in **mgmt-CVX configuration** mode to configure CVX client parameters.

Mgmt-CVX configuration mode is not a group-change mode; **running-config** is changed immediately upon entering commands. Exiting mgmt-CVX configuration mode does not affect running-config. The **`exit`** command returns the switch to global configuration mode.

The **`no management cvx`** and **`default management cvx`** commands delete all mgmt-CVX configuration mode statements from running-config.

Command Mode

Global Configuration

Command Syntax

```
management cvx
no management cvx
default management cvx
exit
```

Commands Available in Mgmt-CVX Configuration Mode

- `heartbeat-interval (Management-CVX)`
- `heartbeat-timeout (Management-CVX)`
- `server host (Management-CVX)`
- `source-interface (Management-CVX)`
- `shutdown (Management-CVX)`

Examples

- This command places the switch in mgmt-CVX configuration mode:

```
switch(config)#management cvx
switch(s1)(config-mgmt-cvx)#
```

- This command returns the switch to global management mode:

```
switch(config-mgmt-cvx)#exit
switch(config)#
```

2.9.8 manager

The **manager** command configures the IP address of the OVSDB controller for the HSC service, allowing CVX to connect to the controller.

The **no manager** and **default manager** commands remove the HSC manager configuration from **running-config**.

Command Mode

CVX-HSC Configuration

Command Syntax

```
manager ip_address [port]
```

Parameters

ip_address IP address of the HSC manager.

port connection port. Values range from 1 to 65535; default value is 6632.

Related Commands

- **service hsc** places the switch in CVX-HSC configuration mode.

Example

- These commands point the HSC service to a controller at IP address *192.168.2.5* using the default port *6632*.

```
switch(config)#cvx
switch(config-cvx)#service hsc
switch(config-cvx-hsc)#manager 192.163.2.5
switch(config-cvx-hsc)#
```

2.9.9 name-resolution force (CVX-OpenStack)

The **name-resolution force** command initiates an OpenStack controller function that communicates with the OpenStack Keystone and Nova services to update names of VMs and tenants mapped by the local OpenStack instance.

The OpenStack controller accesses the Keystone and Nova services in response to various triggering events (such as the creation of a new tenant, network or VM), and also at a regular interval configured by the **name-resolution interval (CVX-OpenStack)** command (default interval 6 hours). The **name-resolution force** command is used to force an immediate update without waiting for a triggering event.

Command Mode

CVX-OpenStack Configuration

Command Syntax

```
name-resolution force
```

Related Commands

- **service openstack** places the switch in CVX-OpenStack configuration mode.
- **name-resolution interval (CVX-OpenStack)** sets the interval for automatic Keystone updates.

Example

- These commands update the OpenStack instance immediately with data from the Keystone service.

```
switch(config)#cvx
switch(config-cvx)#service openstack
switch(config-cvx-openstack)#name-resolution force
switch(config-cvx-openstack)#
```

2.9.10 name-resolution interval (CVX-OpenStack)

The name-resolution interval command specifies the period between consecutive requests that the OpenStack controller sends to the Keystone service for VM and tenant name updates. Keystone is OpenStack's authentication and authorization service.

The default period is 21600 seconds (6 hours).

The name-resolution force (CVX-OpenStack) command performs an immediate update, as opposed to waiting for the periodic update.

Command Mode

CVX-OpenStack Configuration

Command Syntax

name-resolution interval period

Parameters

- period: Keystone identity service polling interval (seconds).

Related Commands

- service openstack places the switch in **CVX-OpenStack** configuration mode.

Example

- These commands set the name resolution interval period at five hours.

```
switch(config)#cvx
switch(config-cvx)#service openstack
switch(config-cvx-openstack)#name-resolution interval 18000
switch(config-cvx-openstack)#
```

2.9.11 ovsdb-shutdown

The **ovsdb-shutdown** command shuts down the OVSDB server.

The **no ovsdb-shutdown** and **default ovsdb-shutdown** commands enable the OVSDB server by removing the ovsdb-shutdown commandfrom running-config.

Command Mode

CVX-HSC Configuration

Command Syntax

```
ovsdb-shutdown
```

```
no ovsdb-shutdown
```

```
default ovsdb-shutdown
```

Related Commands

- **service hsc** places the switch in **CVX-HSC configuration** mode.

Example

- These commands shut down the OVSDB server used by the HSC service.

```
switch(config)#cvx
switch(config-cvx)#service hsc
switch(config-cvx-hsc)#ovsdb-shutdown
switch(config-cvx-hsc)#
```

2.9.12 port (CVX)

The **port** command specifies the TCP port number the CVX server listens on. The default port number is **9979**.

The **no port** and **default port** commands restore the default port number by removing the port statement from **running-config**.

Command Mode

CVX Configuration

Command Syntax

`port` *port_number*

`no port`

`default port`

Parameters

- *port_number*: TCP port number. Value ranges from 1 to 65535.

Related Commands

- cvx places the switch in **CVX configuration** mode.

Example

- These commands configure **9500** as the CVX server port.

```
switch#config
switch(config)#cvx
switch(config-cvx)#port 9500
switch(config-cvx)#
```

- These commands restore the default port (9979) as the CVX server port.

```
switch(config-cvx)#no port
switch(config-cvx)#
```

2.9.13 resync-period

The **resync-period** command configures the grace period for completion of synchronization between the VXLAN control service and clients after a CVX restart. Arista recommends leaving the grace period set to its default of 300 seconds.

The **no resync-period** command disables VXLAN control service graceful restart. The **default resync-period** command resets the grace period to its default of 300 seconds.

Command Mode

CVX-VXLAN Configuration

Command Syntax

`resync-period` *seconds*

`no resync-period`

`default resync-period`

Parameters

- *seconds*: synchronization grace period in seconds. Values range from 30 to 4800; default is 300.

Examples

- These commands reset the VXLAN control service synchronization grace period to 300 seconds.

```
switch(config)#cvx
switch(config-cvx)#service vxlan
switch(config-cvx-vxlan)#default resync-period
switch(config-cvx-vxlan)#
```

2.9.14 server host (Management-CVX)

The **`server host`** command configures the IP address or host name of the CVX server to which the CVX client device connects. The configuration of this address is required for the switch to function as a CVX client. By default, no CVX host address is specified.

The **`no server host`** and **`default server host`** commands remove the CVX host address assignment by removing the server host statement from running-config.

Command Mode

Mgmt-CVX Configuration

Command Syntax

```
server host host

no server host

default server host
```

Parameters

- host: IPv4 address (in dotted decimal notation) or FQDN host name of the CVX server.

Related Commands

- `management cvx` places the switch in **Mgmt-CVX configuration** mode.

Examples

- This command specifies *10.1.1.14* as the address of the server to which the CVX client connects.

```
switch(config)#management cvx
switch(config-mgmt-cvx)#server host 10.1.1.14
switch(config-mgmt-cvx)#
```

2.9.15 service hsc

The **service hsc** command enters **CVX-HSC configuration** mode where the hardware switch controller (HSC) service is enabled and configured.

CVX-HSC configuration mode is not a group change mode; **running-config** is changed immediately upon entering commands. Exiting **CVX-HSC configuration** mode does not affect **running-config**. The **exit** command returns the switch to global configuration mode.

Command Mode

CVX Configuration

Command Syntax

service hsc

Commands Available in CVX-HSC Configuration Mode

- manager
- ovsdb-shutdown
- shutdown (CVX-HSC)
- vtep (CVX-HSC)

Related Commands

- **cvx** places the switch in CVX configuration mode.

Example

- These commands enter **CVX-HSC configuration** mode.

```
switch(config)#cvx
switch(config-cvx)#service hsc
switch(config-cvx-hsc)#
```

2.9.16 service openstack

The service openstack command places the switch in CVX-OpenStack configuration mode.

In order to integrate Arista switches into an OpenStack managed cloud network, OpenStack needs to interact with CVX to configure and maintain VLANs on appropriate physical switch ports that connect to hosts where the VMs reside.

CVX-OpenStack configuration mode is not a group change mode; **running-config** is changed immediately upon entering commands. Exiting **CVX-OpenStack** configuration mode does not affect **running-config**. The **exit** command returns the switch to global configuration mode.

Command Mode

CVX Configuration

Command Syntax

```
service openstack
```

Commands Available in CVX-OpenStack Configuration Mode

- `name-resolution force (CVX-OpenStack)`
- `name-resolution interval (CVX-OpenStack)`
- `shutdown (CVX-OpenStack)`

Related Commands

- **cvx** places the switch in CVX configuration mode.

Example

- These commands places the switch in **CVX-OpenStack** configuration mode.

```
switch(config)#cvx
switch(config-cvx)#service openstack
switch(config-cvx-openstack)#
```

2.9.17 service vxlan

The service vxlan command enters **CVX-VXLAN configuration** mode where the VXLAN control service is enabled and configured.

CVX-VXLAN configuration mode is not a group change mode; ***running-config*** is changed immediately upon entering commands. Exiting **CVX-VXLAN configuration** mode does not affect running-config. The **exit** command returns the switch to global configuration mode.

Command Mode

CVX Configuration

Command Syntax

```
service vxlan
```

Commands Available in CVX-VXLAN Configuration Mode

- `resync-period`
- `shutdown (CVX-VXLAN)`
- `vtep (CVX-VXLAN)`

Related Commands

- **cvx** places the switch in CVX configuration mode.

Example

- These commands enter **CVX-VXLAN configuration** mode.

```
switch(config)#cvx
switch(config-cvx)#service vxlan
switch(config-cvx-vxlan)#
```

2.9.18 show cvx

The **show cvx** command displays the enable status and current configuration of CVX.

Command Mode

EXEC

Command Syntax

```
show cvx
```

Example

- This command displays status and configuration of CVX.

```
switch(config)#cvx
 cvx
 no shutdown
 heartbeat-interval 30
 heartbeat-timeout 90
switch(config-cvx)#dis
switch>show cvx
CVX Server
 Status: Enabled
 UUID: 75ce27ce-cc04-11e4-a404-233646319a2c
 Heartbeat interval: 30.0
 Heartbeat timeout: 90.0
switch>
```

2.9.19 show network physical-topology

The **show network physical-topology** command displays the network topology discovered through CVX.

Command Mode

EXEC

Command Syntax

`show network physical-topology` *hosts*|*neighbors*

Parameters

- *hosts* Displays all hosts visible in the topology.
- *neighbors* Displays all connections in the network topology. Table is sorted by host name, and can be optionally filtered by host.

Example

- This command displays all visible hosts.

```
switch#show network physical-topology hosts

Unique Id            Hostname
-------------------- ------------------------------
001c.7385.be69       cvx287.sjc.aristanetworks.com
0000.6401.0000       cvc1
0000.6402.0000       cvc2
0000.6403.0000       cvc3
0000.6404.0000       cvc4
bcf6.85bd.8050       dsj14-rack14-tor1
```

- This command displays all connections in the topology.

```
switch#show network physical-topology neighbors

cvx287.sjc.aristanetworks.com

Interface          Neighbor Intf      Neighbor Host
------------------ ------------------ ----------------------
Ethernet1          Ethernet7          cvc4
Ethernet2          Ethernet7          cvc2
Ethernet9          Ethernet7          cvc1
Ethernet10         Ethernet7          cvc3
Management1        27                 dsj14-rack14-tor1

OUTPUT OMITTED FROM EXAMPLE
dsj14-rack14-tor1

Interface          Neighbor Intf      Neighbor Host
------------------ ------------------ ----------------------
27                 Management1        cvx287.sjc.aristanetwork
```

2.9.20 shutdown (CVX)

The **shutdown** command, in cvx mode, disables or enables the switch as a CVX server. By default, CVX is disabled on the switch.

The **no shutdown** command enables the switch as a CVX server. The shutdown and **default shutdown** commands disable the switch as a CVX server by removing the **no shutdown** command from running-config.

Note: Be sure to de-configure or shut down all CVX client services before disabling CVX; failure to do so may result in CVX client services continuing to run after CVX has been disabled.

Command Mode

CVX Configuration

Command Syntax

```
shutdown
no shutdown
default shutdown
```

Related Commands

- cvx places the switch in CVX configuration mode.

Example

- These commands enable the switch as a CVX server.

```
switch#config
switch(config)#cvx
switch(config-cvx)#no shutdown
switch(config-cvx)#
```

- This command disables CVX on the switch.

```
switch(config-cvx)#shutdown
switch(config-cvx)#
```

2.9.21 shutdown (CVX-HSC)

The **shutdown** command, in **CVX-HSC configuration** mode, disables or enables the CVX hardware switch controller (HSC) service on the switch. HSC is disabled by default.

When a CVX server enables HSC, its clients (hardware VTEPs) are able to share state to establish VXLAN tunnels without the need for a multicast control plane. Configuration is also required on the client switches.

The **no shutdown** command enables the HSC service; the **shutdown** and **default shutdown** commands disable the HSC service.

Command Mode

CVX-VXLAN Configuration

Command Syntax

```
shutdown
no shutdown
default shutdown
```

Related Commands

- **service hsc** places the switch in CVX-HSC configuration mode.

Example

- These commands enable the HSC service.

```
switch(config)#cvx
switch(config-cvx)#service hsx
switch(config-cvx-hsc)#no shutdown
switch(config-cvx-hsc)#
```

- These commands disable the HSC service.

```
switch(config)#cvx
switch(config-cvx)#service hsx
switch(config-cvx-hsc)#shutdown
switch(config-cvx-hsc)#
```

2.9.22 shutdown (CVX-OpenStack)

The **shutdown** command, in **cvx-openstack** configuration mode, disables or enables CVX-OpenStack on the switch. CVX-OpenStack is disabled by default.

When a CVX server enables OpenStack services, its clients are accessible to the OpenStack network controller (Neutron). Integrating Arista switches into an OpenStack-managed cloud network requires OpenStack to interact with CVX to configure and maintain VLANs on appropriate physical switch ports that connect to the hosts where the VMs reside.

The **no shutdown** command enables CVX-OpenStack. The **shutdown** and **default shutdown** commands disable CVX-OpenStack by removing the corresponding no shutdown command from running-config.

Command Mode

CVX-OpenStack Configuration

Command Syntax

```
shutdown

no shutdown

default shutdown
```

Related Commands

- **service openstack** places the switch in **CVX-OpenStack configuration** mode.

Example

- These commands enable CVX-OpenStack.

```
switch(config)#cvx
switch(config-cvx)#service openstack
switch(config-cvx-openstack)#no shutdown
switch(config-cvx-openstack)#
```

- These commands disable CVX-OpenStack.

```
switch(config-cvx-openstack)#
switch(config-cvx-openstack)#shutdown
switch(config-cvx-openstack)#
```

2.9.23 shutdown (CVX-VXLAN)

The **shutdown** command, in ***CVX-VXLAN configuration*** mode, disables or enables the CVX VXLAN control service on the switch. VXLAN control service is disabled by default.

When a CVX server enables VXLAN control service, its clients (hardware VTEPs) are able to share state to establish VXLAN tunnels without the need for a multicast control plane. Configuration is also required on the client switches.

The **no shutdown** command enables the VXLAN control service. The **shutdown** and **default shutdown** commands disable the VXLAN control service.

Command Mode

CVX-VXLAN Configuration

Command Syntax

```
shutdown

no shutdown

default shutdown
```

Related Commands

- **service vxlan** places the switch in CVX-VXLAN configuration mode.

Example

- These commands enable VXLAN control service.

```
switch(config)#cvx
switch(config-cvx)#service vxlan
switch(config-cvx-vxlan)#no shutdown
switch(config-cvx-vxlan)#
```

- These commands disable VXLAN control service

```
switch(config)#cvx
switch(config-cvx)#service vxlan
switch(config-cvx-vxlan)#shutdown
switch(config-cvx-vxlan)#
```

2.9.24 shutdown (Management-CVX)

The **shutdown** command, in ***mgmt-cvx*** mode, disables or enables CVX client services on the switch. CVX services are disabled by default.

The **no shutdown** command enables CVX client services. The **shutdown** and **default shutdown** commands disable CVX client services by removing the corresponding no shutdown command from running-config.

Command Mode

Mgmt-CVX Configuration

Command Syntax

```
shutdown
```

```
no shutdown
```

```
default shutdown
```

Related Commands

- `management cvx` places the switch in **Mgmt-cvx configuration** mode.

Example

- These commands enable CVX client services.

```
switch(config)#management cvx
switch(config-mgmt-cvx)#no shutdown
switch(config-mgmt-cvx)#
```

- This command disables CVX client services.

```
switch(config-mgmt-cvx)#shutdown
switch(config-mgmt-cvx)#
```

2.9.25 source-interface (Management-CVX)

The **source-interface** command specifies the interface from where the IPv4 address is derived for use as the source for outbound CVX packets that the switch sends as a CVX client. There is no default source interface assignment.

The **no source-interface** and **default source-interface** commands remove the source interface assignment for the CVX client by deleting the source-interface statement from ***running-config***.

Command Mode

Mgmt-CVX Configuration

Command Syntax

`source-interface` **INT_NAME**

`no source-interface`

`default source-interface`

Parameters

INT_NAME: Interface type and number. Options include:

- **ethernet** *e_num*: Ethernet interface specified by e_num.
- **loopback** *l_num*: Loopback interface specified by l_num.
- **management** *m_num*: Management interface specified by m_num.
- **port-channel** *p_num*: Port-Channel Interface specified by p_num.
- **vlan** *v_num*: VLAN interface specified by v_num.

Related Commands

- management cvx places the switch in **Mgmt-CVX configuration** mode.

Example

- These commands configure the CVX client to use the IP address *10.24.24.1* as the source address for its outbound packets.

```
switch#config
switch(config)#interface loopback 5
switch(config-if-Lo5)#ip address 10.24.24.1/24
switch(config-if-Lo5)#exit
switch(config)#management cvx
switch(config-mgmt-cvx)#source-interface loopback 5
switch(config-mgmt-cvx)#
```

2.9.26 vtep (CVX-HSC)

The HSC service sends flood lists to each VTEP through CVX. Some controllers (such as VMware NSX's Service Nodes) implement replication nodes for head-end replication of unknown packets. For these controllers, BUM packets should be sent to a single replication node (send-to-any replication), and the flood list sent by the HSC service is a list of replication nodes. Other controllers (such as Nuage VSP) require each VTEP to perform its own head-end replication. For these, BUM packets should be sent to every known VTEP, and the flood list sent by the HSC service is the list of VTEPs.

The default behavior is to use a send-to-any replication list of VTEPs. If the required behavior is send-to-all replication of, use the all option of the `vtep` command in **CVX-HSC configuration** mode.

Command Mode

CVX-HSC Configuration

Command Syntax

vtep flood list type alllany

no vtep flood list type

default vtep flood list type

Parameters

- all: send-to-all replication; flood list is the list of VTEPs.
- any: send-to-any replication; flood list is a list of replication nodes. This is the default setting.

Related Commands

- service hsc places the switch in **CVX-HSC configuration** mode.

Example

- These commands configure the HSC to use send-to-all replication.

```
switch(config)#cvx
switch(config-cvx)#service hsc
switch(config-cvx-hsc)#vtep flood list type all
switch(config-cvx-hsc)#
```

2.9.27 vtep (CVX-VXLAN)

The OVSDB management protocol includes provisions for control-plane MAC learning, which allows MAC addresses to be distributed among VTEPs without using the data plane. Some controllers (such as VMware NSX) take advantage of this facility; others (such as Nuage VSP) do not. By default, CVX uses control-plane MAC learning.

To switch to data plane MAC learning, use the **vtep** command in CVX-VXLAN configuration mode, as shown below.

Command Mode

CVX-VXLAN Configuration

Command Syntax

`vtep mac-learning` **control-plane|data-plane**

Related Commands

- **service vxlan** places the switch in **CVX-VXLAN configuration** mode.

Example

- These commands configure CVX to use data-plane MAC address learning.

```
switch(config)#cvx
switch(config-cvx)#service vxlan
switch(config-cvx-vxlan)#vtep mac-learning data-plane
switch(config-cvx)#
```

Chapter 3

Macro-Segmentation Service (CVX)

Arista MSS is designed as a service in CloudVision that provides the point of integration between individual vendor firewalls or a firewall manager and the Arista network fabric. MSS provides flexibility on where to place the service devices and workloads. It is specifically aimed at Physical-to-Physical (P-to-P) and Physical-to-Virtual (P-to-V) workloads.

Sections in this chapter include:

3.1 Overview

The advent of contemporary networking features such as mobile applications and the Internet of Things (IoT) bring in additional security challenges that are unprotected by legacy infrastructure. These security breaches cannot be handled by installing a firewall at the Internet edge. Arista's MSS addresses the security breach issue, besides securing access, protecting critical data and end-user privacy.

Arista MSS is designed as a service in CloudVision that provides the point of integration between individual vendor firewalls or a firewall manager and the Arista network fabric. MSS provides flexibility on where to place the service devices and workloads. It is specifically aimed at Physical-to-Physical (P-to-P) and Physical-to-Virtual (P-to-V) workloads.

MSS components include:

- Arista leaf-spine switch fabric
- Arista CloudVision
- Vendor firewall attached to a spine or service leaf switches. Different vendor firewalls can be attached to different switches to enhance scalability.

The above component topology allows for consistency in application deployment, scale, manageability, and easier scalability of the network and service layers.

3.1.1 Benefits

MSS provides the following key benefits:

- Enhanced security between any physical and virtual workloads in the data center.
- The automatic and seamless service insertion ability of MSS eliminates manual steering of traffic for a workload or a tenant.
- Security policies are applied to the host and application throughout the network.

- MSS is flexible since there are no proprietary frame formats, tagging, or encapsulation.

3.1.2 Terminology

The following terms related to MSS are used to describe the MSS feature:

- **Intercept Switch/VTEP**: TOR switch and VXLAN tunnel end-point connected to host from which traffic is intercepted. In the topology diagram, Intercept-1 and Intercept-2 are intercept switches.
- **Service Switch/VTEP**: TOR switch and VXLAN tunnel end-point connected to a firewall. In the topology diagram, Service-1 is the service switch.
- **Service VNI**: VXLAN tunnel created to redirect intercepted traffic to the service device (mapped to locally significant service VLAN).
- **Original VNI**: Original VNI traffic (mapped to Original VLAN).
- **VXLAN**: Virtual eXtensible LAN - a standards-based method of encapsulating Layer 2 traffic across a Layer 3 fabric.
- **CVX**: Arista CloudVision eXchange (CVX) is a part of CloudVision and is a virtualized instance of the same Extensible Operating System (EOS) that runs on physical switches. It functions as a point of integration between customer firewalls or firewall policy managers and the Arista network in order to steer traffic to the firewall.

3.1.3 Usage Scenarios

The following usage scenarios describe a few major security challenges in today's data center that are successfully handled by MSS.

1. Securing server-server traffic.

   This scenario provides information about the role of MSS in securing network traffic between physical-to-physical (P-to-P) and physical to virtual (P-to-V) servers. Prior to MSS, network infrastructure devices followed the "firewall sandwich" setup where firewalls were placed in line between the security zones. This setup would impact scalability and performance of the servers.

   Using MSS, this restriction on firewall placement is reduced. Firewalls are now attached to a service leaf switch in the network fabric and they still protect hosts without concern about their physical location. The following topology demonstrates the usage scenario.

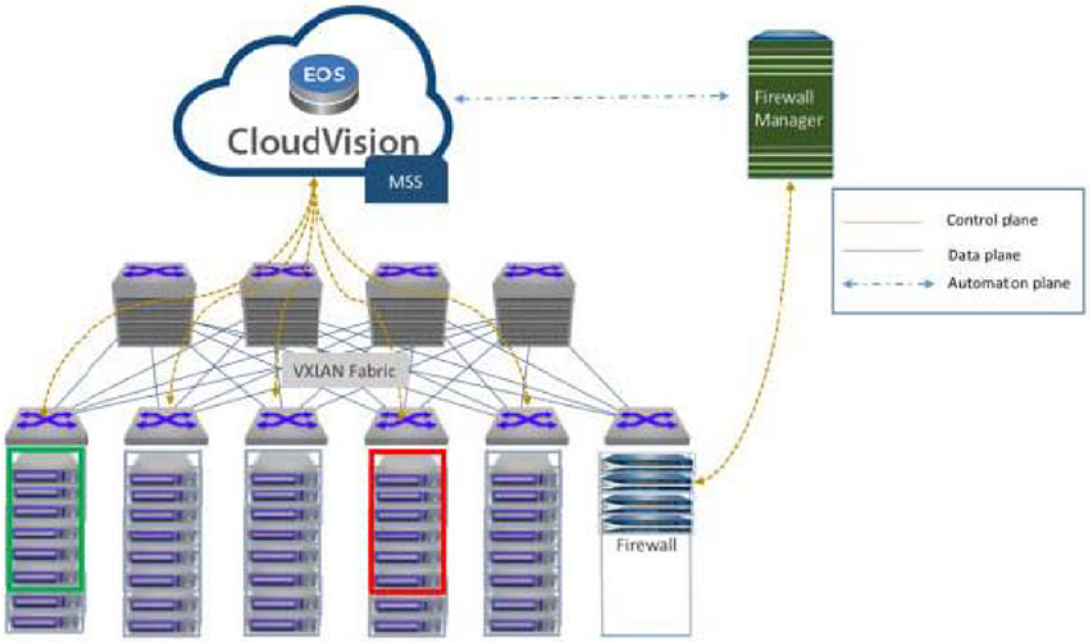

Figure 9: Securing server-server traffic

2. Monitoring and securing management traffic.

   This usage scenario demonstrates how MSS successfully monitors and secures management interfaces in the data center.

   The modern data center caters to managing the application, storage, virtualization, network, analytics and other layers. With virtualization, the hypervisor management also needs to be secured to prevent unwanted access to a hypervisor management interface. In the event of a rogue access, Arista's MSS protects management interfaces. The explicitly allowed hosts can gain access through a jump host or administrator end-user computing instances. The following topology diagram illustrates the role of MSS in a data center.

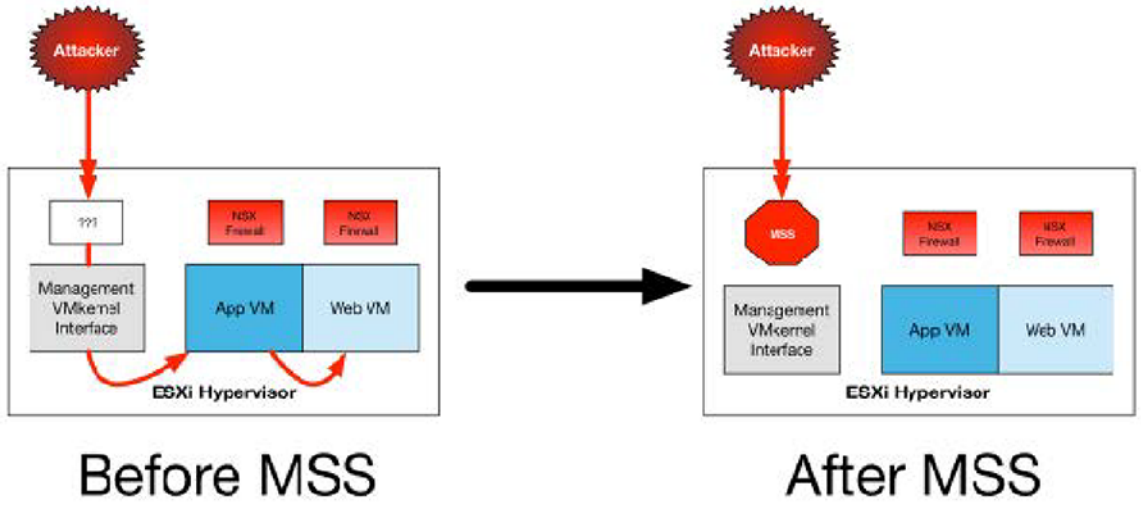

Figure 10: Monitoring and Securing management traffic

3.2 How MSS Works

The following steps provide information about how MSS works as a service in the data center.

1. MSS is enabled on the CloudVision eXchange (CVX) and the Arista switches are configured to stream their active state to CVX. This allows CVX to build a database of hosts and firewalls attached to the network and also to identify physical ports and IP addresses. CVX is also configured to communicate and synchronize policies from a vendor's firewall.
2. CVX sends a request to the firewall or firewall manager to provide information about the security policies which are tagged for MSS usage. Accordingly, it will determine where traffic needs to be intercepted.
3. CVX applies an intercept to steer the network traffic and pushes the intercept rules to the intercept switches where the server or applications are located.

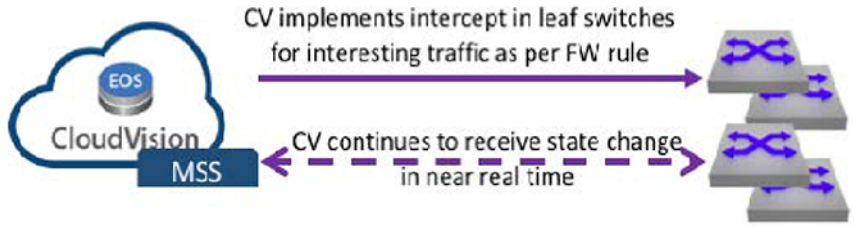

Figure 11: CVX intercept

4. The leaf switch starts sending intercepted traffic to the service leaf when the intercept has been applied to the leaf switch.

Figure 12: Leaf switches intercept

5. Traffic is forwarded completely unmodified to the firewall after it enters the service leaf where the firewall is attached. Based on the configuration policy, the firewall applies the required actions such as inspection, log, allow, or deny.
6. The service leaf switch sends the inspected traffic to its final destination or to the destination based on the firewall policy.

3.3 Configuration

The following sections provide detailed information about MSS configuration, system requirements, recommendations, and limitations.

The traffic flow below is an example of a typical MSS deployment with a 3-tiered application. The goal of this design is to limit access between hosts in the following zones: web-untrust, app-untrust, db-untrust, web-trust, app-trust, and db-trust.

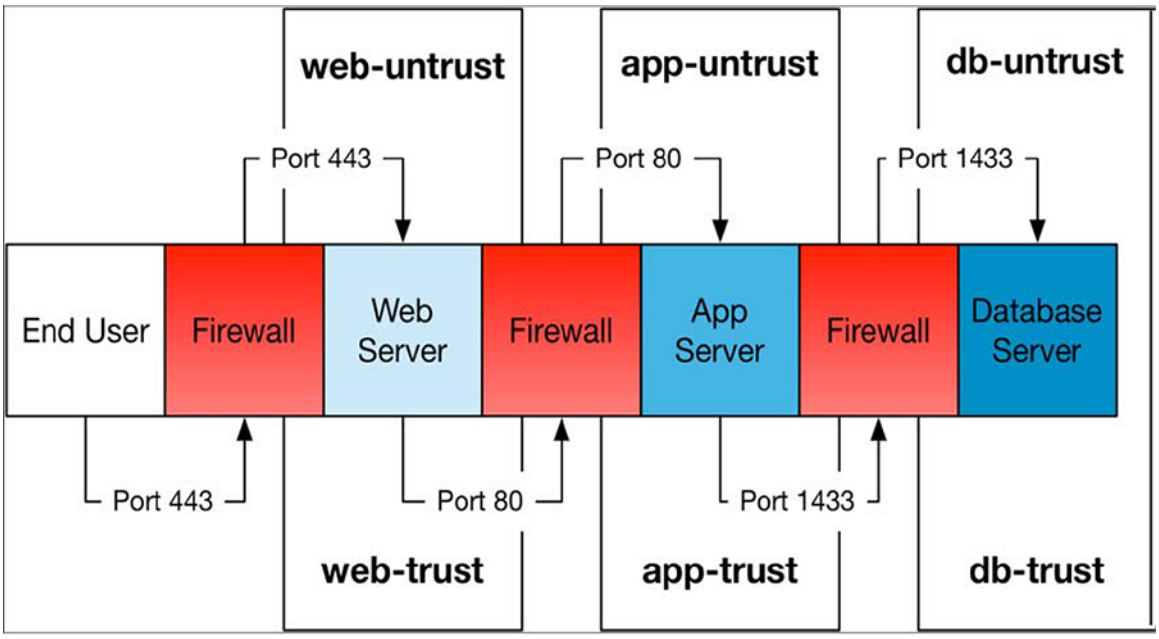

Figure 13: Traffic flow in an MSS deployment

End users in the untrust zone access the web server through the TCP/443 port. Traffic flows through the active firewall to the web server interface in the web-untrust security zone. The web server interface in the web-trust security zone accesses the application server interface in the app-untrust security zone through port TCP/80 after traversing the firewall. From there, the application server interface in the app-trust security zone accesses the database through TCP/1433 in the db-untrust zone.

The following physical topology indicates the MSS setup.

The hosts are attached to a pair of intercept leaf switches. A firewall is connected to a service leaf switch using a pair of physical interfaces with a subinterface per zone or vWire.

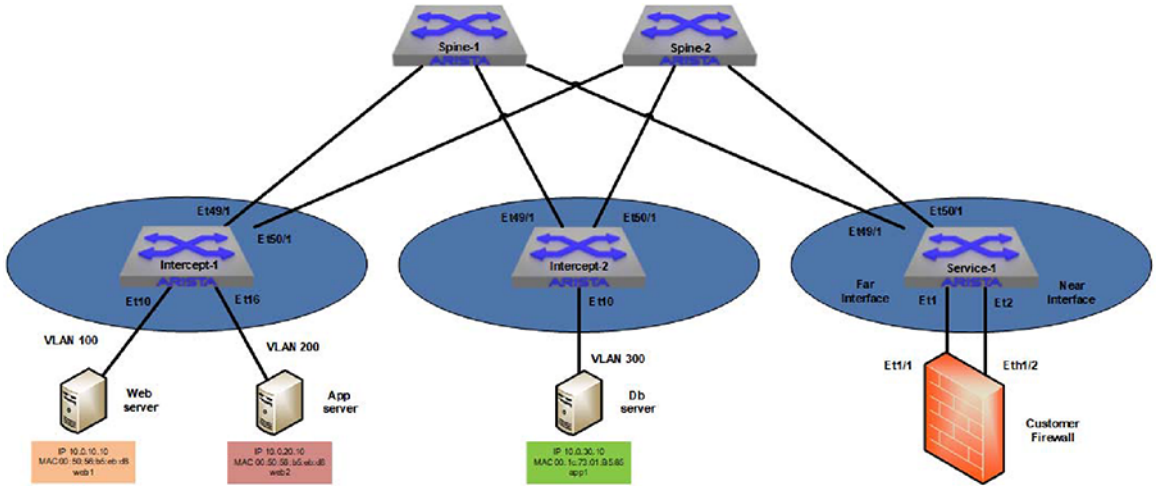

Figure 14: Physical topology of the MSS

- System Requirements
- Recommendations and Limitations
- Configuring MSS

3.3.1 System Requirements

The system requirements to effectively run MSS are listed below.

- Arista CloudVision eXchange (CVX)
- Arista 7050X, 7050X2, 7060X, and 7060X2 series top of rack (TOR) switches
- Connected to the hosts to intercept traffic from the firewall devices
- Connected to and monitored by CVX
- The network must be a VXLAN-enabled fabric with CVX running the VXLAN Control Service (VCS)
- Link Layer Discovery Protocol (LLDP) should be enabled on the firewall interfaces attached to the Arista TOR switches. Note that static mapping can be configured if required.

3.3.2 Recommendations and Limitations

TOR and CVX Switches

- Service switches should be dedicated exclusively to firewalls and not to host connectivity.
- In the event of an entry time-out, the server ARP entries are not re-learned on the service VTEP.

Firewall

- The firewall policy name must not have any whitespace character in the name. As an example, “PCI policy” is an unacceptable policy name. An acceptable name would be “PCI_policy”.
- When High Availability firewalls are used in the system, all links to switches must be port channels and a Multi-Chassis Link Aggregation (MLAG) bow-tie configuration should be used.

3.3.3 Configuring MSS

These sections describe steps to configure MSS.

3.3.3.1 Deploying CVX

Deploy CloudVision and configure the Arista TOR switches to connect to it. A CVX cluster of 3 instances with host names of cvx01, cvx02, and cvx03 are configured as an example.

Note: As a best practice, always deploy the CV in a HA cluster with a minimum of three instances.

3.3.3.2 Enabling the VXLAN Control Service on CVX

Enable the VXLAN Control Service (VCS) on every CVX instance after the three Arista CVX instances have been deployed and the TOR switches are configured to be managed by them.

VCS allows hardware VXLAN Tunnel End Points (VTEPs) to share state with each other in order to establish VXLAN tunnels without the need for a multicast control plane.

Example

CVX instance cvx01

```
cvx01(config-cvx)#service vxlan
cvx01(config-cvx-vxlan)#no shutdown
```

Similarly, VCS is enabled on the cvx02 and cvx03 devices.

3.3.3.3 Configuring the Access switches and the Service switch ports

Configure the switch ports that are connected to the hosts, whose traffic should be steered to the firewalls and the service switch ports which are connected to the firewalls.

Access switch configuration

The switch ports connected to the hosts, whose traffic needs to be intercepted, need to be configured as 802.1q trunks with the VLAN that is mapped to the VNI requiring interception. Unique VLAN IDs are configured for each tier of the application.

Access switch (intercept-1)

```
intercept-1# configure
intercept-1(config)# interface et10
intercept-1(config-if-Et10)# description web server
intercept-1(config-if-Et10)# switchport mode trunk
intercept-1(config-if-Et10)# switchport trunk allowed vlan 100

intercept-1(config)# interface et16
intercept-1(config-if-Et16)# description app server
intercept-1(config-if-Et16)# switchport mode trunk
intercept-1(config-if-Et16)# switchport trunk allowed vlan 200
```

Access switch (intercept-2)

```
intercept-2# configure
intercept-2(config)# interface et10
intercept-2(config-if-Et1)# description db server
intercept-2(config-if-Et1)# switchport mode trunk
intercept-2(config-if-Et1)# switchport trunk allowed vlan 300
```

Note: For untagged traffic, configure a native VLAN on the port using the **switchport trunk native vlan** command.

Service switch (service-1)

```
service-1# configure
service-1(config)#interface port-channel 10
service-1(config-if-Po10)# description Far Interface
service-1(config-if-Po10)# switchport mode trunk
service-1(config-if-Po10)# switchport trunk allowed vlan none
service-1(config-if-Po10)# spanning-tree bpdufilter enable

service-1(config)#interface port-channel 20
service-1(config-if-Po20)# description Near Interface
service-1(config-if-Po20)# switchport mode trunk
service-1(config-if-Po20)# switchport trunk allowed vlan none
service-1(config-if-Po20)# spanning-tree bpdufilter enable
```

Note: Dynamically mapped VLANs are not shown in the switch port configuration. You can view them by running the `show vlan` command on the switch once a policy is applied.

3.3.3.4 Enabling DirectFlow on access switches and service switches

Arista MSS uses DirectFlow to intercept traffic while the VxLAN is used to carry tunnel traffic from the intercepted host to the firewall and back. DirectFlow should be enabled on every intercept switch as well as the service switches.

Switch service-1

```
service-1# configure
service-1(config)# directflow
service-1(config-directflow)# no shutdown
```

Switch intercept-1

```
intercept-1# configure
intercept-1(config)# directflow
intercept-1(config-directflow)# no shutdown
```

Switch intercept-2

```
intercept-2# configure
intercept-2(config)# directflow
intercept-2(config-directflow)# no shutdown
```

3.3.3.5 Enabling VXLAN routing on the TOR switches

CVX uses Address Resolution Protocol (ARP) to determine where intercept hosts are physically located in the network. VXLAN routing should be configured on every TOR switch that will be intercepting traffic to ensure that CVX is aware of every host ARP entry.

The following configuration shows the routing configuration for each tier of the application, but not the entire VXLAN configuration. For more information on how to configure VXLAN and VXLAN routing, refer to the VXLAN section of the *Arista EOS Configuration Guide*.

Switch intercept-1

```
intercept-1# configure
intercept-1(config)# ip routing
intercept-1(cofig)# interface vlan100
intercept-1(config-if-Vl100)# ip address virtual 10.0.10.254/24
intercept-1(config)# interface vlan200
intercept-1(config-if-Vl200)# ip address virtual 10.0.20.254/24
intercept-1(config)# interface vlan300
intercept-1(config-if-Vl300)# ip address virtual 10.0.30.254/24
```

Switch intercept-2

```
intercept-2# configure
intercept-2(config)# ip routing
intercept-2(cofig)# interface vlan100
intercept-2(config-if-Vl100)# ip address virtual 10.0.10.254/24
intercept-2(config)# interface vlan200
intercept-2(config-if-Vl200)# ip address virtual 10.0.20.254/24
intercept-2(config)# interface vlan300
intercept-2(config-if-Vl300)# ip address virtual 10.0.30.254/24
```

Switch service-1

```
service-1# configure
service-1(config)# ip routing
service-1(cofig)# interface vlan100
service-1(config-if-Vl100)# ip address virtual 10.0.10.254/24
service-1(config)# interface vlan200
service-1(config-if-Vl200)# ip address virtual 10.0.20.254/24
service-1(config)# interface vlan300
service-1(config-if-Vl300)# ip address virtual 10.0.30.254/24
```

3.3.3.6 Configuring MSS on CVX

This step enables configuring Arista MSS on CVX. The topology diagram depicts three CVX instances in a cluster and the configuration is the same for every instance. The active and standby vendor firewalls are configured. If Panorama is used, only Panorama should be configured.

Example

In the example, the primary vendor firewall has a DNS name of *fw-ha-node-1*. The standby firewall has a DNS name of *fw-ha-node-2*. The username and password are set as *admin*.

CVX instance cvx01

```
cvx01# configure
cvx01(config)# cvx
cvx01(config-cvx)# no shutdown
cvx01(config-cvx)# service mss
cvx01(config-cvx-mss)# no shutdown
cvx01(config-cvx-mss)# vni range 20000-30000
cvx01(config-cvx-mss)# dynamic device-set panfw1
cvx01(config-cvx-mss-panfw1)# tag Arista_MSS
cvx01(config-cvx-mss-panfw1)# type palo-alto firewall
cvx01(config-cvx-mss-panfw1)# state active
cvx01(config-cvx-mss-panfw1)# device fw-ha-node-1
cvx01(config-cvx-mss-panfw1-fw-ha-node-1)# username admin password 0
 admin
```

CVX instance cvx02

```
cvx02# configure
cvx02(config)# cvx
cvx02(config-cvx)# no shutdown
cvx02(config-cvx)# service mss
cvx02(config-cvx-mss)# no shutdown
cvx02(config-cvx-mss)# vni range 20000-30000
cvx02(config-cvx-mss)# dynamic device-set panfw1
cvx02(config-cvx-mss-panfw1)# tag Arista_MSS
cvx02(config-cvx-mss-panfw1)# type palo-alto firewall
cvx02(config-cvx-mss-panfw1)# state active
cvx02(config-cvx-mss-panfw1)# device fw-ha-node-1
cvx02(config-cvx-mss-panfw1-fw-ha-node-1)# username admin password 0
 admin
```

CVX instance cvx03

```
cvx03# configure
cvx03(config)# cvx
cvx03(config-cvx)# no shutdown
cvx03(config-cvx)# service mss
cvx03(config-cvx-mss)# no shutdown
cvx03(config-cvx-mss)# vni range 20000-30000
cvx03(config-cvx-mss)# dynamic device-set panfw1
cvx03(config-cvx-mss-panfw1)# tag Arista_MSS
cvx03(config-cvx-mss-panfw1)# type palo-alto firewall
cvx03(config-cvx-mss-panfw1)# state active
cvx03(config-cvx-mss-panfw1)# device fw-ha-node-1
cvx03(config-cvx-mss-panfw1-fw-ha-node-1)# username admin password 0
 admin
```

3.3.3.7 Configuring the Firewall

Three policies are created in addition to the default implicit deny policy for inter-zone traffic. The implicit deny ensures that the inter-zone traffic is not allowed unless a policy explicitly allows for it.

The first policy “untrust_to_web1” is from the untrust zone to the web1 zone, that allows HTTPS traffic from anywhere to the web server web.

The third policy “web2_to_app1” is from the web2 zone to the app1 zone that allows HTTP traffic between the web server web and the application server app.

The fifth policy “app2_to_db1” is from the app2 zone to the db1 zone that allows database traffic on port TCP/1433 between the application server app and the database server db.

The second, fourth, and sixth policies prevent the firewall to drop a session for which does not see the initial connection to the protected resource. This could happen if the protected resource has not sent any traffic previous to this point.

Refer to the following images for more clarity about the above policies and interface configuration.

				Source		Destination		Rule Usage					
	Name	Tags	Type	Zone	Address	Zone	Address	Hit Count	Last Hit	First Hit	Application	Service	Action
1	vl 150 allow	offload	universal	vl 1200	150.0.1.0/30	phy int	160.0.5.1/32	0	-	-	any	UDP_dst_2000-1	Allow
2	test	offload	universal	vl 1200	170.0.1.0/30	any	any	0	-	-	any	application-default	Allow
3	vl 150 drop	offload	universal	vl 1200	150.0.1.0/30	any	any	0	-	-	any	udp_10000 UDP_dst_2000-1	Drop
4	vl 150 offload	offload	universal	vl 1200	150.0.1.0/29	any	any	0	-	-	any	application-default	Allow
5	vl 170-171	mss	universal	vl 1200	170.0.1.0/30	phy int	171.0.7.0/30	0	-	-	any	udp 17100	Allow
6	vl 150	mss	universal	vl 1200	150.0.1.0/27	any	any	0	-	-	any	application-default	Allow
7	intrazone-default	none	intrazone	any	any	(intrazone)	any	0	-	-	any	any	Allow
8	interzone-default	none	interzone	any	any	any	any	160	2019-03-11 13:38:35	2019-02-11 20:03:33	any	any	Allow

Figure 15: Firewall policy configuration

Interface	Interface Type	Management Profile	Link State	IP Address	Virtual Router	Tag	VLAN / Virtual-Wire	Security Zone	Features	Comment
ethernet1/1	Aggregate (ae1)			none	none	Untagged	none	none		Po 100 - near interface
ethernet1/2	Aggregate (ae1)			none	none	Untagged	none	none		po 100 - near interface
ethernet1/3	Aggregate (ae2)			none	none	Untagged	none	none		po 101 - far interface
ethernet1/4	Aggregate (ae2)			none	none	Untagged	none	none		po 101 - far interface
ethernet1/5	Layer2			none	none	Untagged	none	none		
ethernet1/6	Virtual Wire			none	none	Untagged	none	none		
ethernet1/7	Tap			none	none		none	none		
ethernet1/8	Layer2			none	none	Untagged	none	none		
ethernet1/9	Layer2			none	none	Untagged	none	none		
ethernet1/10	Aggregate (ae7)			none	none	Untagged	none	none		
ethernet1/11	Aggregate (ae8)			none	none	Untagged	none	none		
ethernet1/12	Aggregate (ae7)			none	none	Untagged	none	none		
ethernet1/13	Aggregate (ae8)			none	none	Untagged	none	none		
ethernet1/14	Tap			none	none		none	none		
ethernet1/15	Virtual Wire			none	none	Untagged	none	none		
ethernet1/16	Virtual Wire			none	none	Untagged	none	none		
ethernet1/17	Layer3	allow ping		none	none	Untagged	none	none		
ethernet1/17.1200	Layer3	allow ping		199.2.0.2/21 199:2::100:2/64	default	1200	none	vl 1200		
ethernet1/17.1201	Layer3	allow ping		199.2.8.2/21	default	1201	none	vl 1201		

Figure 16: Firewall interface configuration

Create a rule that Arista MSS will use to intercept and redirect traffic and add a firewall policy with the default “Arista_MSS” tag as shown in the example above. MSS intercepts all traffic from endpoints identified in policies that match the tag values configured in CVX. The firewall will apply all rules (tagged or untagged) to all traffic.

Note: LLDP should always be enabled on the firewall interfaces attached to the service switches. To minimize reconvergence time on the network changes, reduce the LLDP transmit interval and hold time multiples on the firewall, while keeping the LLDP hold time above the LLDP timer configured on the connected Arista switches.

Alternatively, the **device interface map** command can be used on CVX to manually map a device to Arista switch interfaces. To map multiple devices, add a mapping entry for each device.

```
dynamic device-set fw1
device dc-firewall-1
map device-interface ethernet1/1 switch 00:1c:73:7e:21:bb interface
 Ethernet1
map device-interface ethernet1/2 switch 00:1c:73:7e:21:bb interface
 Ethernet9
```

3.4 MSS Commands

Configuration Commands

- dynamic device-set
- exception device
- group
- service mss
- state
- tag
- type palo-alto

CVX Show Commands

- dynamic device-set
- exception device
- group
- service mss
- state
- tag
- type palo-alto
- show service mss dynamic device-set
- show service mss policy
- show service mss status
- show service mss zone

3.4.1 dynamic device-set

The **dynamic device-set**command configures a device such as a firewall to communicate with the MSS in the MSS configuration mode.

The no dynamic device-set command removes a previously configured device from the MSS configuration and returns to the CVX mode.

Command Mode

MSS Configuration

Command Syntax

`dynamic device-set` *device-set_name*

`no dynamic device-set` *device-set_name*

Parameters

- device-set_name a unique name for the device set.

Example

- This example creates a set of firewalls with the name “panfw1”.

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# no shutdown
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# no shutdown
cvx(config-cvx-mss)# vni range 30000-40000
cvx(config-cvx-mss)# dynamic device-set panfw1
cvx(config-cvx-mss-panfw1)#
```

Note: The **vni range** command configures a range of VXLAN Network Identifiers (VNI) that MSS uses to tunnel traffic to the firewall. If VNI range is not configured, the default VNIs in the range of 1 to 16777214 are used.

3.4.2 exception device

The **exception device** command bypasses or continues redirecting traffic to service device such as a firewall if the service device control-plane API is unreachable after initial policies have been processed.

The no exception device command.

Command Mode

MSS Configuration

Command Syntax

`exception device unreachable` [bypass | redirect]

`no exception device unreachable` [bypass | redirect]

`default exception device unreachable bypass`

Parameters

- device: service device in the device set.
- unreachable: service device control-plane API is unreachable.
- bypass: bypass the service device.
- redirect: continue redirecting traffic to the service device.

Example

- This example redirects traffic to the service device.

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# no shutdown
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# no shutdown
cvx(config-cvx-mss)# vni range 30000-40000
cvx(config-cvx-mss)# dynamic device-set fw
cvx(config-cvx-mss-fw)# device firewall-dc7
cvx(config-cvx-mss-fw)# username admin password 7 PKigsmo3IcnW5rqoZXWQ
cvx(config-cvx-mss-fw)# state active
cvx(config-cvx-mss-fw)# type palo-alto firewall
cvx(config-cvx-mss-fw)# exception device unreachable redirect
```

3.4.3 group

The **group** command configures the Panorama device group name to be used with MSS.

The no group command removes the group from the MSS configuration when the Panorama firewall manager is used.

See the type palo-altocommand for more information about the firewall manager.

Command Mode

Device-set mode

Command Syntax

`group` *group_name*

`no group` *group_name*

Parameters

- *group_name* the name of the group.

Example

- This command configures the group name as *mssDevices*.

```
cvx(config)#cvx
cvx(config-cvx)#service mss
cvx(config-cvx-mss)#dynamic device-set pano2
cvx(config-cvx-mss-pano2)#type palo-alto panorama
cvx(config-cvx-mss-pano2)#device myPanorama
cvx(config-cvx-mss-pano2-myPanorama)#group mssDevices
```

3.4.4 service mss

The service mss command enters the MSS configuration sub-mode.

The no service mss command exits the MSS configuration mode and returns to the CVX mode.

Command Mode

CVX Configuration

Command Syntax

```
service mss
```

```
no service mss
```

```
default service mss
```

Example

- This example enables MSS on CVX and enters the MSS config mode.

Note: The **`no shutdown`** command enables MSS on the CloudVision eXchange (CVX).

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# no shutdown
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# no shutdown
```

3.4.5 state

The **state** command configures device set as active or disabled or suspended state.

The no state command disables the previously configured state of the device set.

Command Mode

MSS Configuration

Command Syntax

`state` [active | shutdown | suspend]

`no state`

Parameters

- active: the active state of the device set. Policy monitoring and network traffic redirection are enabled.
- shutdown: the disabled state of the device set. Policy monitoring and network traffic redirection is stopped.
- suspend: the suspended state of the device set. Policy monitoring is suspended but there is no change in the existing traffic redirection.

Example

- This output example configures the device set state as “active”.

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# no shutdown
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# no shutdown
cvx(config-cvx-mss)# vni range 30000-40000
cvx(config-cvx-mss)# dynamic device-set panfw1
cvx(config-cvx-mss-panfw1)# tag Arista_MSS
cvx(config-cvx-mss-panfw1)# type palo-alto firewall
cvx(config-cvx-mss-panfw1)# state active
```

3.4.6 tag

The **tag** command specifies the tag or tags that MSS searches when it is reading the security policy from the firewall or firewall manager in the dynamic device-set configuration mode. You can specify more than one tag as well.

The no tag command removes the tag from the MSS configuration.

Note: The tag specified should always match with the firewall policy tags in the vendor firewall policy for the MSS to read the policy and set up the intercept.

Command Mode

MSS Configuration

Command Syntax

`tag` *tag_name*

`no tag`

`default tag` Arista_MSS

Parameters

- *tag_name*: a unique name for the tag.

Examples

- This command specifies the tag with the name Arista_MSS.

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# no shutdown
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# no shutdown
cvx(config-cvx-mss)# vni range 30000-40000
cvx(config-cvx-mss)# dynamic device-set panfw1
cvx(config-cvx-mss-panfw1)# tag Arista_MSS
```

- This command specifies multiple tags with names *mss1, mss2,* and *mss3*.

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# no shutdown
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# no shutdown
cvx(config-cvx-mss)# vni range 30000-40000
cvx(config-cvx-mss)# dynamic device-set panfw1
cvx(config-cvx-mss-panfw1)# tag mss1 mss2 mss3
```

3.4.7 type palo-alto

The **`type palo-alto`** command configures the firewall type to be used in the MSS configuration.

The no type palo-alto command disables the firewall type from the MSS configuration.

Command Mode

MSS Configuration

Command Syntax

`type palo-alto` [firewall | panorama]

`no type palo-alto`

Parameters

- firewall: the Palo Alto Networks firewall.
- panorama: the Palo Alto Networks Panorama firewall manager.

Example

- This command configures the Palo Alto Networks firewall type.

```
cvx# configure
cvx(config)# cvx
cvx(config-cvx)# service mss
cvx(config-cvx-mss)# dynamic device-set panfw1
cvx(config-cvx-mss-panfw1)# type palo-alto firewall
```

3.4.8 show service mss dynamic device-set

The **show service mss dynamic device-set**command displays detailed information about a specific service device set. Information such as device group members, high availability, network, resource details are displayed.

Note: Interfaces from multiple switches can be placed in the same zone by the device.

Command Mode

EXEC

CVX Configuration

Command Syntax

`show service mss dynamic device-set` *device_set_name* [device *device_name* [group-members | high-availability | neighbors | network | policies | resources]]

Parameters

- *device_set_name* defines the device set name.
- device *device name* defines the service device properties such as the DNS hostname or IP address of the service device.
- **group members** lists device-group members for an aggregation manager.
- high-availability displays service device high availability information.
- **neighbors** displays the service device's ethernet interface neighbor information.
- **network** displays the service device's network interface information.
- **policies** displays the list of policies read from service device that have the MSS tag.
- **resources** displays the service device's system resource information.

Related Commands

- show service mss status
- show service mss policy

Examples

- This command displays information about interfaces that are placed in a zone by the "device1".

```
switch# show service mss zone
Source: static
------------------------------------------------------------
Device: device1
```

- This command displays information about interfaces that are placed in a zone by the "device1".

```
switch# show service mss zone
Source: static
------------------------------------------------------------
Device: device1
Zone: zone1
Switch: 00:00:00:00:00:01
Hostname: switch1.arista.com
Interfaces:
Ethernet1/1
Allowed VLAN: 1000-1010
Port-Channel2/1:
Allowed VLAN: 1000-2000
Switch: 00:00:00:00:00:02
Hostname: switch2.arista.com
Interfaces:
Ethernet10/1
```

```
Allowed VLAN: 1000-1010
Zone: zone2
Switch: 00:00:00:00:00:01
Hostname: switch1.arista.com
Interfaces:
Ethernet10/1
Allowed VLAN: 1000-1010
Ethernet 20/1
Allowed VLAN: 1000-2000
```

3.4.9 show service mss policy

The **show service mss policy**command displays generic information about the configuration and operational state of the macro-segmentation service (MSS) policies on a device.

Command Mode

EXEC

CVX Configuration

Command Syntax

`show service mss policy` [[device *device_name*] [name *policy-name*] [source (static | *plugin_name*)]]

Parameters

- device *device name* defines the service device name.
- **name** *policy-name* the filter policy name.
- source the source of the policy.
- **static** the policy configured using the command line interface.
- *plugin_name* the service device type.

Related Commands

- show service mss status
- show service mss zone

Example

- This command displays information about the MSS policy “policy1” enabled on the device.

```
cvx# show service mss policy name policy1
Source Device Policy Config Status
------  --------  ------  -------  -------------  -------------
vendor  Firewall  pan100  policy1  Enabled        Initialized
```

The "Config" column indicates the configuration state of a policy. The different states are: *Enabled*, *dry run*, and *disabled* states.

The "Status" column indicates the operational state of a policy. The different status types are *initialized*, *pending*, *initializing*, *active*, *reinitializing*, *dry-run Complete*, and *deactivating*.

3.4.10 show service mss status

The **show service mss status** command displays the status of a macro-segmentation service (MSS) on the device.

Command Mode

EXEC

CVX Configuration

Command Syntax

```
show service mss status
```

Related Commands

- show service mss policy
- show service mss zone

Examples

- This command displays the MSS status on the device as “Enabled”.

```
switch# show service mss status
State: Enabled
Service VNIs: 1500-1600,1800,1900-2000
```

- This command displays the MSS status on the device as “Disabled”.

```
switch# show service mss status
State: Disabled
Service VNIs: 1-16777214
```

3.4.11 show service mss zone

The **show service mss zone**command displays information about the interfaces that are placed in a single zone by the service device. Along with the show service mss policy command, we can use this command to identify issues with the policy configuration.

Interfaces from multiple switches can be placed in the same zone by the device.

Command Mode

EXEC

CVX Configuration

Command Syntax

show service mss zone [[device *device_name*]|[name *zone_name*]| [source (static | *dynamic_source*)]]

Parameters

- device *device name* defines the service device properties.
- **name** *policy-name* the filter zone name.
- source the source of the zone.
- **static** the zone configured using the command line interface.
- *dynamic_source* the service device type.

Related Commands

- show service mss status
- show service mss policy

Examples

- This command displays information about interfaces that are placed in a zone by the “device1”.

```
switch# show service mss zone
Source: static
------------------------------------------------------------
Device: device1
Zone: zone1
Switch: 00:00:00:00:00:01
Hostname: switch1.arista.com
Interfaces:
Ethernet1/1
Allowed VLAN: 1000-1010
Port-Channel2/1:
Allowed VLAN: 1000-2000
Switch: 00:00:00:00:00:02
Hostname: switch2.arista.com
Interfaces:
Ethernet10/1
Allowed VLAN: 1000-1010
Zone: zone2
Switch: 00:00:00:00:00:01
Hostname: switch1.arista.com
Interfaces:
Ethernet10/1
Allowed VLAN: 1000-1010
Ethernet 20/1
Allowed VLAN: 1000-2000
```

Chapter 4

CloudVision Portal (CVP) Overview

CloudVision Portal (CVP) is the web-based GUI for the CloudVision platform.

The Portal provides a turnkey solution for automating network operations, including network device provisioning, compliance, change management, and network monitoring. It communicates southbound to Arista switches via eAPI and has open standard APIs northbound for integration with 3rd-party or in-house service management suites.

Figure 17: CloudVision Portal (CVP) overview shows CloudVision as the network control point between the physical infrastructure (network layer) and the layer of service management.

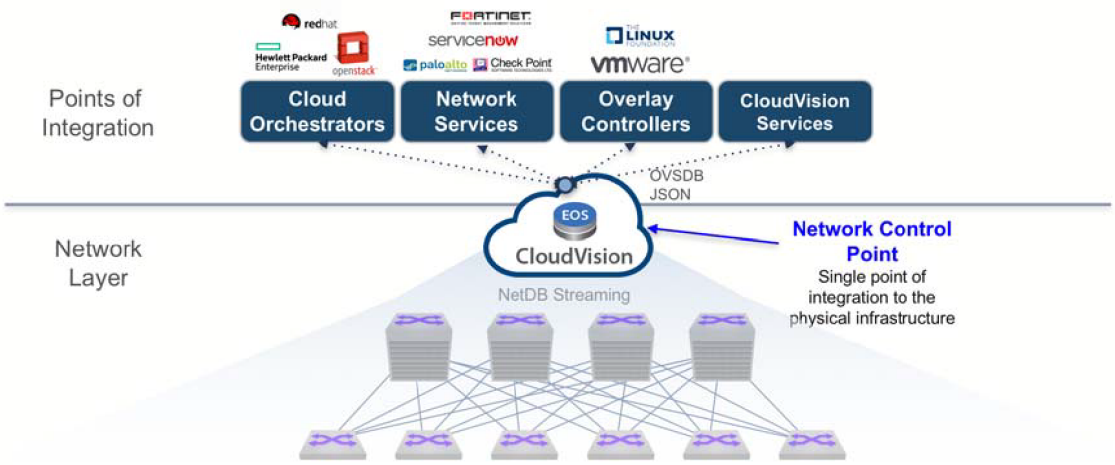

Figure 17: CloudVision Portal (CVP) overview

Sections in this chapter include:

- CVP Virtual Appliance
- CloudVision WiFi
- CVP Cluster Mechanism
- System Requirements
- Key CVP Terms

4.1 CVP Virtual Appliance

The CVP virtual appliance is a packaged ova file that consists of Base OS packages, Hadoop, HBase, Apache Tomcat, JAVA jdk and the CVP web application.

The virtual appliance can be deployed as a stand alone (singlenode) or a cluster of three (multi-node). The different deployment options will be discussed later on in this section, but for production

deployments it is recommended that the cluster option is chosen. The single VM instance is recommended for testing purposes as it provides a simpler setup and requires less resources.

- CVX and CVP

4.1.1 CVX and CVP

Certain CVP features leverage CVX. For the 2017.1 features, CVP is not dependent on any functionality provided by CVX, so deploying CVX along with CVP is recommended but not required.

You can register CVX with CVP in one of two ways:

- By provisioning CVX and then manually registering it in CVP.
- By ZTP booting CVX with CVP.

Note: CVX does not boot into ZTP mode by default, since it is a Virtual Machine (VM). Setting it up and then registering it manually with CVP is the recommended option.

The CVP appliance is shipped as a single OVA file which can be run on any x86 hypervisor. The hypervisors listed below have been tested and confirmed to work with the CVP appliance.

Hypervisor	Version
VMware ESX	5.5
Linux RHEL	6.5-7.0

Related topics:

- System Requirements
- Key CVP Terms

4.2 CloudVision WiFi

The CloudVision WiFi (CVW) service is available as a container on the Arista CloudVision platform. Once you activate the CVW service, you can configure, monitor, troubleshoot, and upgrade Arista WiFi access points using the cognitive CVW UI.

Figure 18: CVW Architecture provides a conceptual overview of the Arista CVW solution.

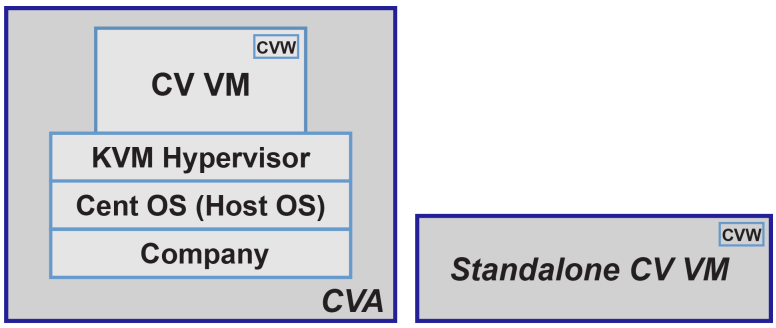

Figure 18: CVW Architecture

CVW is containerized within the CV whether it's CVA (CV on a CV appliance) or a standalone CV VM. The CVW service runs on both single-node CV and CV cluster. In case of a CV cluster, CVW operates as a single logical instance in High Availability mode (HA-mode).

- CVW HA Mode Operation
- Key Features of CVW on CV
- Capacity of CVW on CV

4.2.1 CVW HA Mode Operation

When setting up CVW for the first time, it must be enabled on all the nodes of a cluster. Once CVW is enabled, then at boot time, the CVW service on the primary node automatically becomes the Active instance, and the one on the secondary node becomes the Standby instance. The HA failover and recovery mechanisms work exactly as expected. That is, if the primary node goes down, the CVW instance on the secondary node becomes active. When the primary node rejoins the cluster, a split-brain recovery kicks in and re-elects the new active and standby containers.

4.2.2 Key Features of CVW on CV

Except for OS and kernel processes, the CVW service on CV runs all the application processes required to manage Arista WiFi and wireless intrusion prevention system (WIPS). Some key features of the CVW service are as follows:

- CVW uses ports 3851 and 161 (both UDP) for all CV communication with external entities. These ports need to be opened in your network.
- CVW consists of two key components:
  - **wifimanager**, the server that manages the WiFi network.
  - **aware**, the cognitive WiFi UI of the server.

4.2.3 Capacity of CVW on CV

The table below shows the number of access points (APs) that a CVW container supports for the given CPU, RAM, and hard disk settings. The CPU and RAM values displayed in this table are the default settings for a DCA-200 device; the actual capacity may vary based on deployment, environment, and load.

Table 1: Capacity of CVW on CV

Setting	Up to 5000 APs
CPU	8 Core
RAM	32 GB
Hard Disk	250 GB

4.3 CVP Cluster Mechanism

CVP consists of distributed components such as Zookeeper, Hadoop/HDFS and HBase. Zookeeper provides consensus and configuration tracking mechanism across a cluster. Hadoop/HDFS is a distributed and redundant data store while HBase is a distributed key/value store. Running these services in a reliable fashion on multiple nodes require a quorum mechanism which is subject to limitations imposed by that mechanism.

- CVP Cluster and Single Node Failure Tolerance

4.3.1 CVP Cluster and Single Node Failure Tolerance

In absence of a quorum or a quorum leader, each node assumes itself to be the cluster leader in a three-node cluster leading to chaos and even data corruption. This leads to the quorum constraint for CVP cluster where only single node failure can survive. For example, a single node is allowed to form a cluster in a three-node cluster. In such cases, if cluster nodes cannot communicate with each other, all three nodes assume itself to be the lone survivor and operate accordingly. This is called a split-brain scenario where the original three-node cluster has split into multiple parts.

In real scenarios, assume only two nodes are active after a reboot and they failed to connect with each other. As no quorum is required, each node elects itself as the cluster leader. Now two clusters are formed where each cluster captures different data. For example, devices can be deleted from one cluster but not from the other. Device status is in compliance in one cluster but not on the other, etc. Additionally, services that store zookeeper configuration now has two copies with different data. Consequently, there is no effective way to reconcile the data when these nodes re-establish communication.

Let's consider HBase component in CVP. HBase is a distributed key-value store and splits its data across all cluster nodes. Let's assume that one node splits off from other two. If a single node can form a cluster, this single node forms one cluster and the other two together forms another cluster. It means that there are 2 HBase masters. That is the process which keeps track of metadata for all key/value pairs in HBase. In other words, HBase creates two independent sets of metadata which can even frustrate manual reconciliation. In essence, distributed infrastructure pieces must meet mandatory quorum requirements and which in turn means we cannot survive more than a single node failure.

Another reason to not tolerate dual node failures in a three-node CVP cluster is that all nodes are not made the same and total capacity of the cluster is more than what a single node can handle. Some services might be configured to run only on two of the three nodes and will fail when attempted to run on another. The total configured capacity of CVP cluster is 2 times that of a single node. That means in a three-node cluster, two nodes will have the capacity to run everything but one node cannot. Hence in a cluster of three CVP nodes, the cluster can survive only one CVP node failure.

4.4 System Requirements

The CloudVision Portal is deployed as a virtual or physical appliance. For redundancy, three CloudVision Portal nodes are configured as a cluster. Each VM must be configured to meet the minimum system requirements.

Table 2: Minimum System Requirements

Required Hardware	
Lab Deployment (< 25 devices)	**Production Deployment**
The **minimum** hardware requirements to use CloudVision Portal in a lab environment are: • CPUs: 16 cores • RAM: 22 GB • Disk: 1 TB (use RPM installer) • Disk Throughput: 20MB/s	The **recommended** hardware required to use CloudVision Portal in a production environment (3 node cluster) are: • CPUs: 28 cores • RAM: Recommended 52 GB • Disk: 1 TB • Disk Throughput: 40MB/s

Note: For production deployments, information about device scale is available in the release specific version of the product release notes. For more information on throughput, refer to Troubleshooting and Health Checks .

Table 3: required Software Versions

Required Software Versions

The software versions compatible with CloudVision Portal are:
• EOS switches: EOS version 4.18.1F, or above
• EOS license: Z license
• CVP license: Full subscription license
• One of the following browsers:
  • Mozilla Firefox, version 39+
  • Google Chrome, version 44+ (recommended)
• TerminAttr version 1.6.1 or later

Note: CloudVision Portal does not support live vMotion for multi-node clusters. If the Hypervisor environment is set up for live vMotion, it has to be disabled for the CVP VMs.

Related topics:

- Key CVP Terms
- CVP Virtual Appliance

4.5 Key CVP Terms

Make sure you are familiar with the following key CloudVision Portal (CVP) terms. These terms are used throughout this guide to describe the various CVP features, and the CVP user interface contains icons that represent each of the key terms.

Icon	Term	Definition
	Device	Devices managed by the CloudVision Portal.
	Container	Containers are a logical entity used to group network devices, and define a hierarchy to which user configuration can be applied.
	Device	Devices define the subset of available devices.
	Configlet	Configlets define a subset of a device's configuration.
	Image	Images define the software running on a given device.
	Label	Labels are arbitrary tags defined by the user and applied to devices for identification and filtering purposes.

Icon	Term	Definition
	Notification	Notifications are system messages providing the list of on-going, completed and canceled activities that are not tracked by tasks.
	Task	Tasks are work orders for taking an action against a given device.
N/A	Export to CSV	Downloads the table in csv format to your local drive. **Note:** Replaces hyphen (-) with **N/A** where hyphen indicates empty data. Replaces cells using the **(unknown)** string with empty cells where **(unknown)** indicates data missing due to an error(s).

Related topics:

- CVP Virtual Appliance
- System Requirements

Chapter 5

CloudVision Portal (CVP) Setup

CloudVision Portal (CVP) can be run on ESX or KVM hypervisors. Before you can begin using the CVP, you must complete the CVP setup process which, involves the following:

1. Deploying CVP
2. Configuring CVP

Sections in this chapter include:

- Deploying CVP OVA on ESX
- Deploying CVP on KVM
- Set Up CVW on CV
- Shell-based Configuration
- Shell Reconfiguration of Single-node, Multi-node Systems
- ISO-based Configuration
- Certificate-Based TerminAttr Authentication

There are two different deployment procedures. One for deploying CVP on ESX, and one for deploying CVP on KVM. After you complete the deployment procedures, you then configure CVP. The deployment procedures are:

- Deploying CVP OVA on ESX
- Deploying CVP on KVM

There are two configuration methods for the CloudVision Portal (CVP): shell-based and ISO-based. Both of these methods eliminate the need to directly modify system and CVP configuration files. This simplifies the setup process and reduces the potential for issues.

The configuration methods enable you to configure CVP in both single-node systems and multi-node systems. The configuration methods are:

- Shell-based Configuration (recommended)
- ISO-based Configuration

Note: Reconfiguration is limited to certain parameters on a deployed CVP multi-node cluster.

5.1 Deploying CVP OVA on ESX

Deploying the CVP OVA file should be the first step in any setup. After the CVP OVA file is deployed, you can chose between the two configuration methods for CloudVision Portal (CVP).

Pre-requisites:

Use of the Deploy OVF Template requires the VMware Client Integration plugin, which is not supported by the Chrome browser after versions 42.

1. The OVA file can be deployed as a VM in a VMware environment by using the drop menu under the Actions heading and selecting **Deploy the OVA template**.

Note: For multi-node setups, the following steps must be completed 3 times to launch 3 VMs (once for each VM).

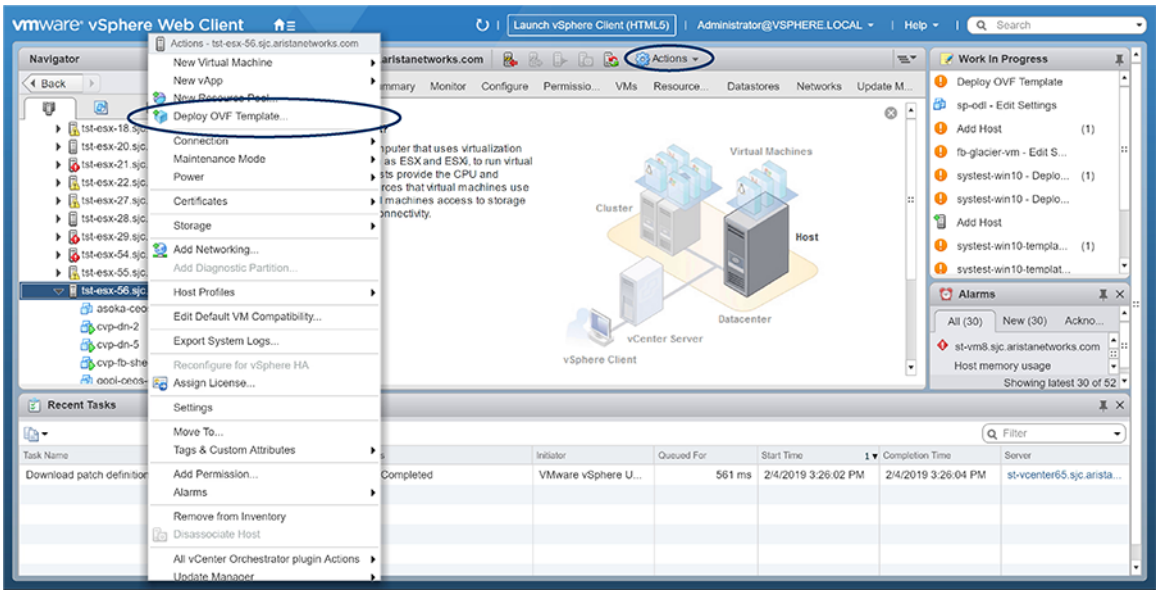

Figure 19: Deploy the OVA template

2. Having selected the Deploy OVF Template option, VCenter will prompt for the location of the OVA file; this can be either on a local hard disk, network share, or Internet URL. The location of the OVA file should be entered or selected.

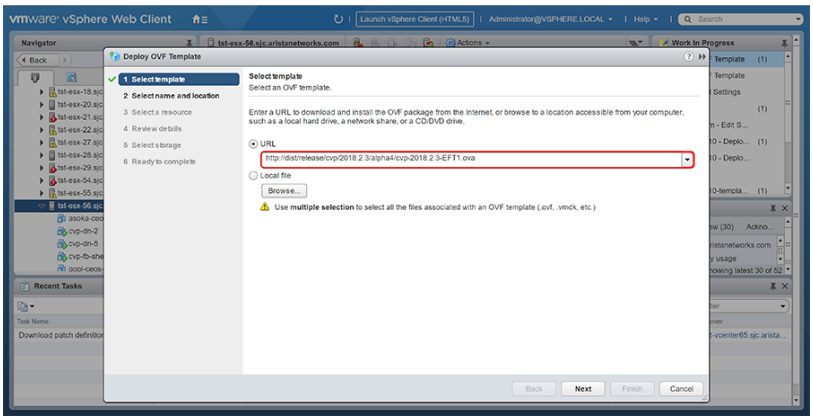

Figure 20: Location of the OVA file

3. Click **Next** to go to the next task.
4. Review the OVA template details (Figure 21: Review OVA template details).

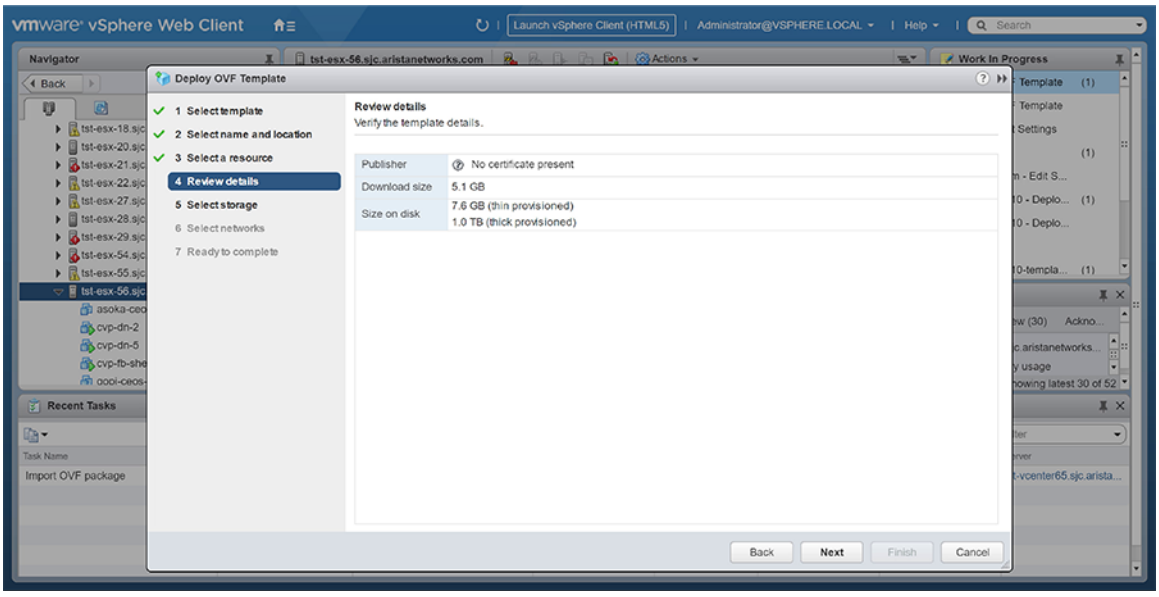

Figure 21: Review OVA template details

5. Click **Next** to go to the next task.
6. Type the name for the OVA file in the **Name** field and select the folder for the OVA file (Figure 22: Select name and folder location for OVA file).

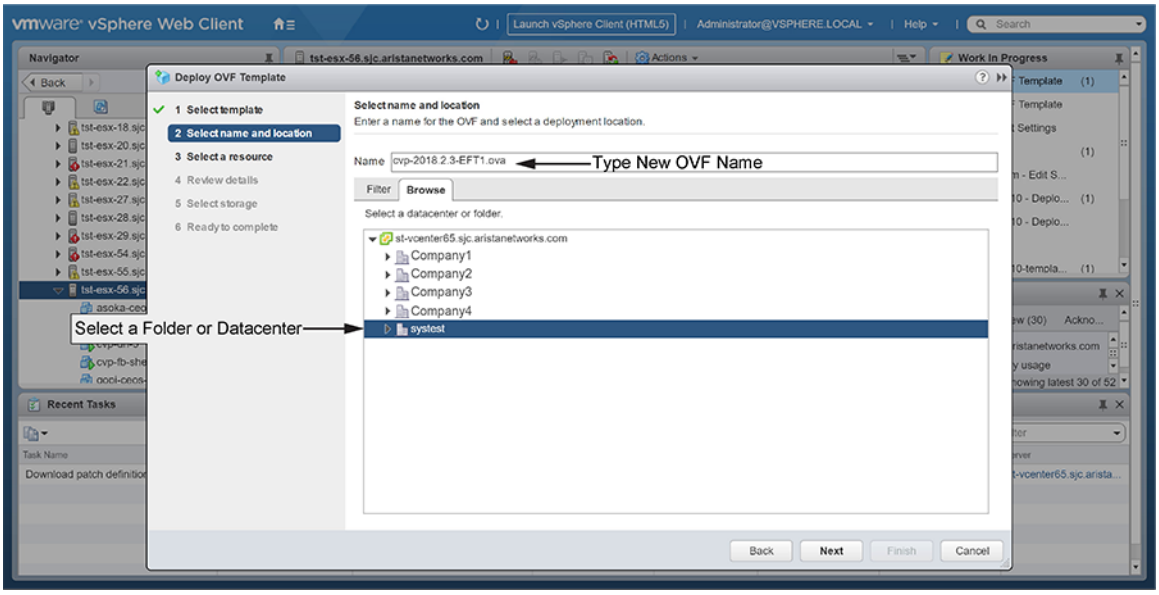

Figure 22: Select name and folder location for OVA file

7. Click **Next** to go to the next task.
8. Select the resource where you want the deployed template (OVA file) to be run (Figure 23: Select the resource).

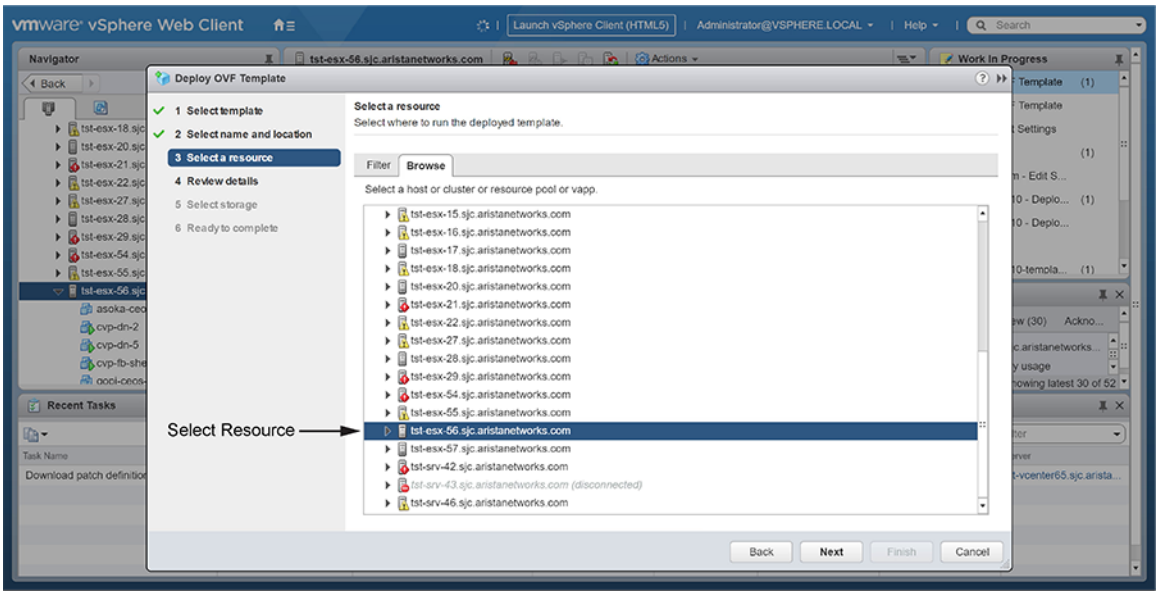

Figure 23: Select the resource

9. Click **Next** to go to the next task.
10. Select the location where you want the files for the deployed template to be stored. You can choose the virtual disk format (Figure 24: Select the destination storage).

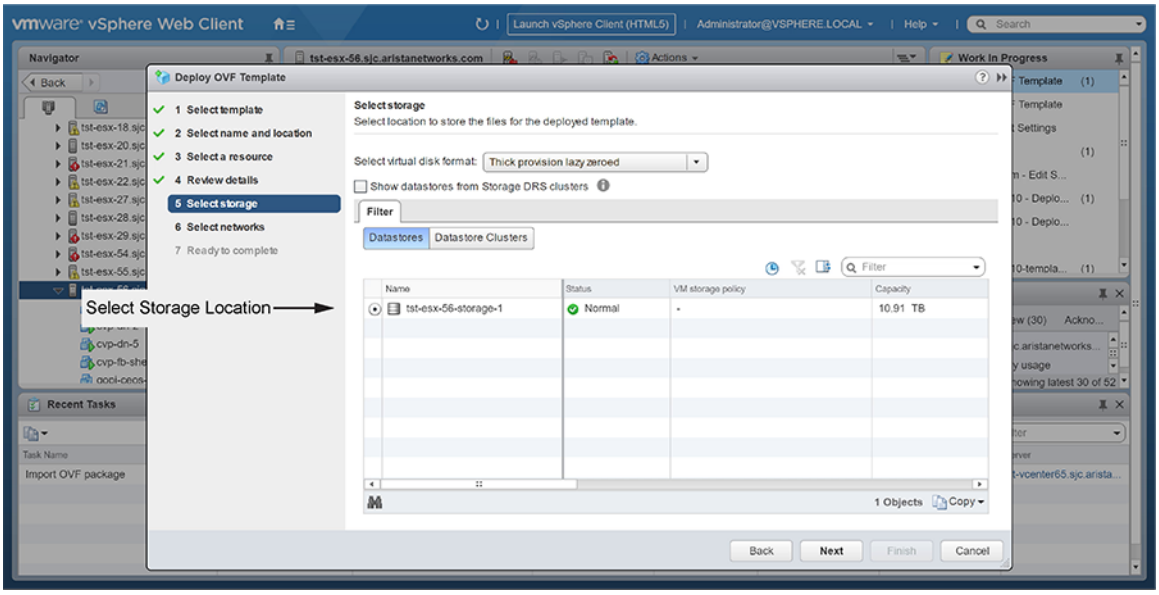

Figure 24: Select the destination storage

11. Click **Next** to go to the next task.
12. Setup the networks that the deployed template should use (Figure 25: Setup the networks).

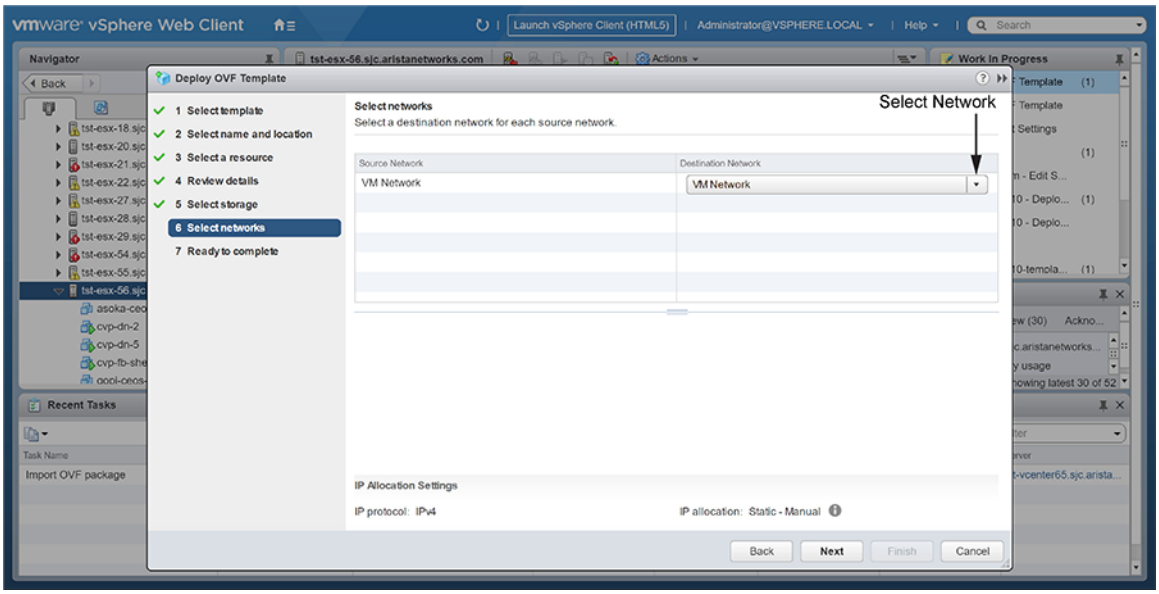

Figure 25: Setup the networks

13. Click **Next**.

    VCenter loads the OVA and displays the configuration settings (Figure 26: Select the “Finish” button to accept these settings).

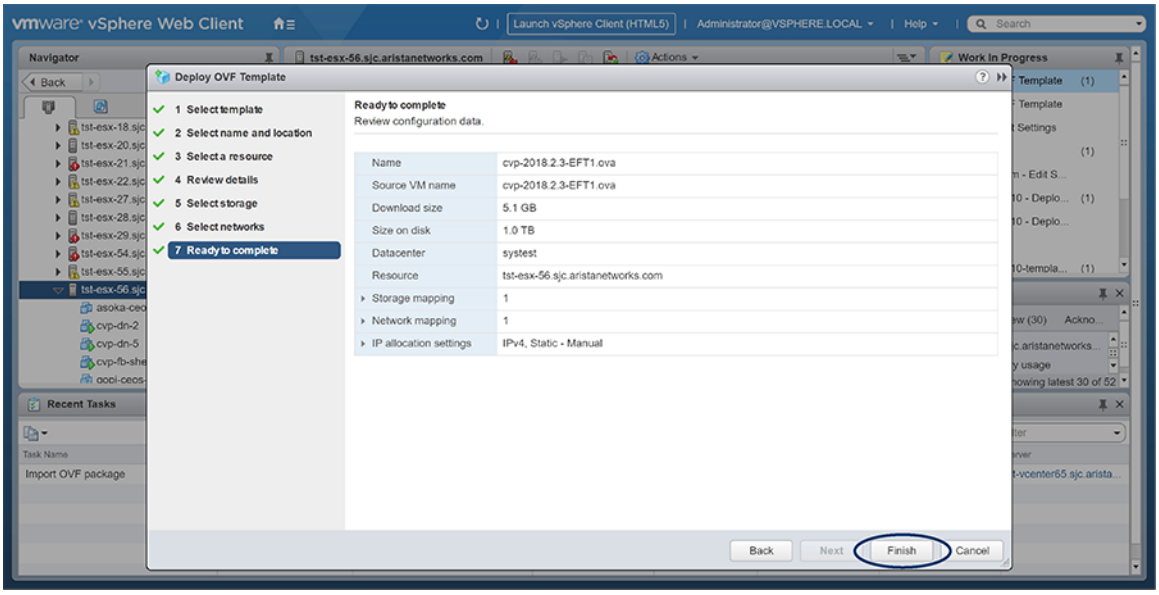

Figure 26: Select the “Finish” button to accept these settings

14. Review the configuration settings, and click **Finish** to accept and save the configuration.

    VCenter begins to deploy the virtual appliance. Once the appliance is deployed, you can configure the CVP application.

Related topics:

- Deploying CVP on KVM
- Shell-based Configuration
- ISO-based Configuration

5.2 Deploying CVP on KVM

In standard KVM environments, deploying a CVP VM involves the following tasks:

- Downloading and extracting the CVP KVM tarball (.tgz archive)
- Creating Virtual Bridge and Network Interface Cards (NIC)
- Generating the XML file that defines the CVP VM
- Defining and Launching the CVP VM

Once you complete these tasks, you can configure the CVP VM.

5.2.1 Downloading and extracting the CVP KVM tarball (.tgz archive)

The first task in the deployment process involves downloading and extracting the CVP KVM tarball. The tarball is a .tgz archive that contains:

- The CVP VM
- Disk images for the CVP application
- The files used to configure CVP VM.

You download the tarball to the host server that is configured for KVM. The files contained in the .tgz archive include:

	Filename	Description
1	disk1.qcow2	VM disk image for the CVP application.
2	disk2.qcow2	Data disk image for the CVP application.
3	cvpTemplate.xml	A template for creating the XML file for libvirt domain specification.
4	generateXmlForKvm.py	A script for generating the CVP VM definition XML based on the XML template.
5	createNwBridges.py	A script for creating the network interfaces for the CVP VM.

Complete the following steps to download and extract the CVP VM .tgz archive:

1. Go to the Arista software downloads webpage and download the CVP VM tarball (`cvp-<version>-kvm.tgz`) to the host server set up for KVM.
2. Extract the tarball (`cvp-<version>-kvm.tgz`).

   Figure 27: Extracting the cvp kvm .tgz archive shows an example of extracting the `CVP KVM .tgz` archive.

```
[arastra@kvm1 vms]# cd cvpTests
[arastra@kvm1 cvpTests]# ls
cvp-2018.2.2-kvm.tar
[arastra@kvm1 cvpTests]# tar -xvf cvp-2018.2.2-kvm.tar
addIsoToVM.py
createNwBridges.py
cvpTemplate.xml
disk1.qcow2
disk2.qcow2
generateXmlForKvm.py
[arastra@kvm1 cvpTests]#  ./generateXmlForKvm.py -n cvpTests --device-bridge br1 -k 1 -i cvpTemplate.xml -o qemuout.xml
-x '/home/arastra/vms/cvpTests/disk2.qcow2' -y '/home/arastra/vms/cvpTests/disk1.qcow2' -b 22528 -p 8 -e '/usr/libexec/q
emu-kvm'
SUCCESS: XML output is in qemuout.xml
[arastra@kvm1 cvpTests]# virsh define qemuout.xml
Domain cvpTests defined from qemuout.xml

[arastra@kvm1 cvpTests]# virsh start cvpTests
Domain cvpTests started

[arastra@kvm1 cvpTests]# virsh console cvpTests
Connected to domain cvpTests
Escape character is ^]
```

Figure 27: Extracting the cvp kvm .tgz archive

5.2.2 Creating Virtual Bridge and Network Interface Cards (NIC)

The second task in deploying CVP for KVM involves creating the bridges and interfaces that provide network connectivity for the CVP VM. You use the `CreateNwBridges.py` script you extracted in the previous task to create the required bridges and interfaces.

Note: If the required network interfaces for CVP already exist, you do not have to complete this task. Go directly to Generating the XML file that defines the CVP VM

You have the option of deploying CVP with either two bridge interfaces or a single bridge interface.

- Two interfaces (the cluster bridge interface and the device bridge interface).
- Single interface (the device bridge interface).

Complete the following steps to create the network interfaces for CVP KVM connectivity:

1. (Optional) Use the `./createNwBridges.py --help` command to view a list of all the parameters available in the script.
2. Use the `./createNwBridges.py` to create the device bridge (or bridges) and interfaces needed.

   shows an example of creating a single device bridge for a single-node deployment.

```
[arastra@kvm1 ~]# ./createNwBridges.py --device-bridge br1 --swap-device-nic-ip --gateway 172.31.0.1
WARNING: You are trying to pull IP address from NIC and apply it to the bridge. This may cause the network connectivity to  to be adversely affected.
Do you want to continue [Y/n] ?Y
SIOCADDRT: File exists
[arastra@kvm1 ~]# brctl show
bridge name     bridge id               STP enabled     interfaces
br1             8000.0cc47a71d958       no              eno1
                                                        vnet0
                                                        vnet1
                                                        vnet2
                                                        vnet3
br2             8000.000000000000       no
br3             8000.000000000000       no
br4             8000.000000000000       no
docker0         8000.0242b8f54337       no
virbr0          8000.5254001f0bd5       yes             virbr0-nic
virbr1          8000.525400c022d4       yes             virbr1-nic
[arastra@kvm1 ~]#
```

Figure 28: Creating a device bridge (single node deployment)

3. (Optional) Use the `brctl show` command to verify that the bridges were successfully created.
4. (Optional) Use the `ifconfig` command to verify that the IP addresses have been allocated. In this example the one IP address for the br1 bridge.

   The following image shows an example of verification of bridge creation and IP address allocation. In this example, a bridge br1 was created, and one IP address has been allocated for the bridge.

```
[arastra@kvm1 cvpdTest]#
[arastra@kvm1 cvpdTest]# ifconfig | head -25
br1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 172.31.6.78  netmask 255.255.0.0  broadcast 172.31.255.255
        inet6 fd7a:629f:52a4:7777:ec4:7aff:fe71:d958  prefixlen 64  scopeid 0x0<global>
        inet6 fe80::ec4:7aff:fe71:d958  prefixlen 64  scopeid 0x20<link>
        ether 0c:c4:7a:71:d9:58  txqueuelen 0  (Ethernet)
        RX packets 858749171  bytes 8616556645135 (7.8 TiB)
        RX errors 0  dropped 11937095  overruns 0  frame 0
        TX packets 633183283  bytes 44547378278 (41.4 GiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

br2: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet6 fe80::20ec:b0ff:fe61:b6e3  prefixlen 64  scopeid 0x20<link>
        ether 22:ec:b0:61:b6:e3  txqueuelen 0  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 8  bytes 648 (648.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

br3: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet6 fe80::20f8:8bff:fed5:ad6b  prefixlen 64  scopeid 0x20<link>
        ether 22:f8:8b:d5:ad:6b  txqueuelen 0  (Ethernet)
        RX packets 0  bytes 0 (0.0 B)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 8  bytes 648 (648.0 B)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
[arastra@kvm1 cvpdTest]# route -n
Kernel IP routing table
Destination     Gateway         Genmask         Flags Metric Ref    Use Iface
0.0.0.0         172.31.0.1      0.0.0.0         UG    0      0        0 br1
169.254.0.0     0.0.0.0         255.255.0.0     U     1004   0        0 br1
169.254.0.0     0.0.0.0         255.255.0.0     U     1005   0        0 br2
169.254.0.0     0.0.0.0         255.255.0.0     U     1006   0        0 br3
169.254.0.0     0.0.0.0         255.255.0.0     U     1007   0        0 br4
172.17.0.0      0.0.0.0         255.255.0.0     U     0      0        0 docker0
172.31.0.0      0.0.0.0         255.255.0.0     U     0      0        0 br1
192.168.0.0     0.0.0.0         255.255.255.0   U     0      0        0 virbr1
[arastra@kvm1 cvpdTest]#
```

Figure 29: Verification of bridge creation and IP address allocation

5.2.3 Generating the XML file that defines the CVP VM

The third task in deploying CVP for KVM involves generating the XML file that you use to define the CVP VM. You use generateXmlForKvm.py script and the cvpTemplate.xml file you extracted previously to generate the XML file you use to define the CVP VM.

The `cvpTemplate.xml` file is a template that defines wildcard values that are filled by the other parameters that are specified when you execute the script.

Complete the following steps to generate the XML file:

1. (Optional) Use the `./generateXmlForKvm.py --help` command to view a list of all the parameters available in the script.
2. Run the `./generateXmlForKvm.py` script using the XML template (`cvpTemplate.xml`) as one of the inputs.

   shows an example of an XML being generated that can be used to define a CVP VM named cvpTest. The generated XML file is named qemuout.xml.

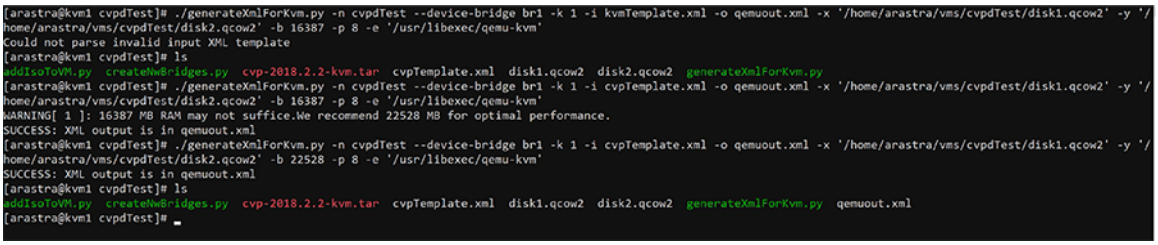

Figure 30: Generation of XML file used to define CVP VM

5.2.4 Defining and Launching the CVP VM

The last task in deploying CVP for KVM is to define and launch the CVP VM. You use the XML file you generated in the previous task to define the CVP VM.

Complete the following steps to define and launch the CVP VM:

1. Run the `virsh define` command to define the CVP VM (specify the generated XML file).
2. Run the `virsh start` command to launch the newly defined CVP VM.
3. Run the `virsh console` command to attach (connect) to the CVP VM console.

Defining and Launching the CVP VM shows an example of the use of the commands to define and launch a CVP VM named cvpTest. The XML file used to define the CVP VM is named qemuout.xml.

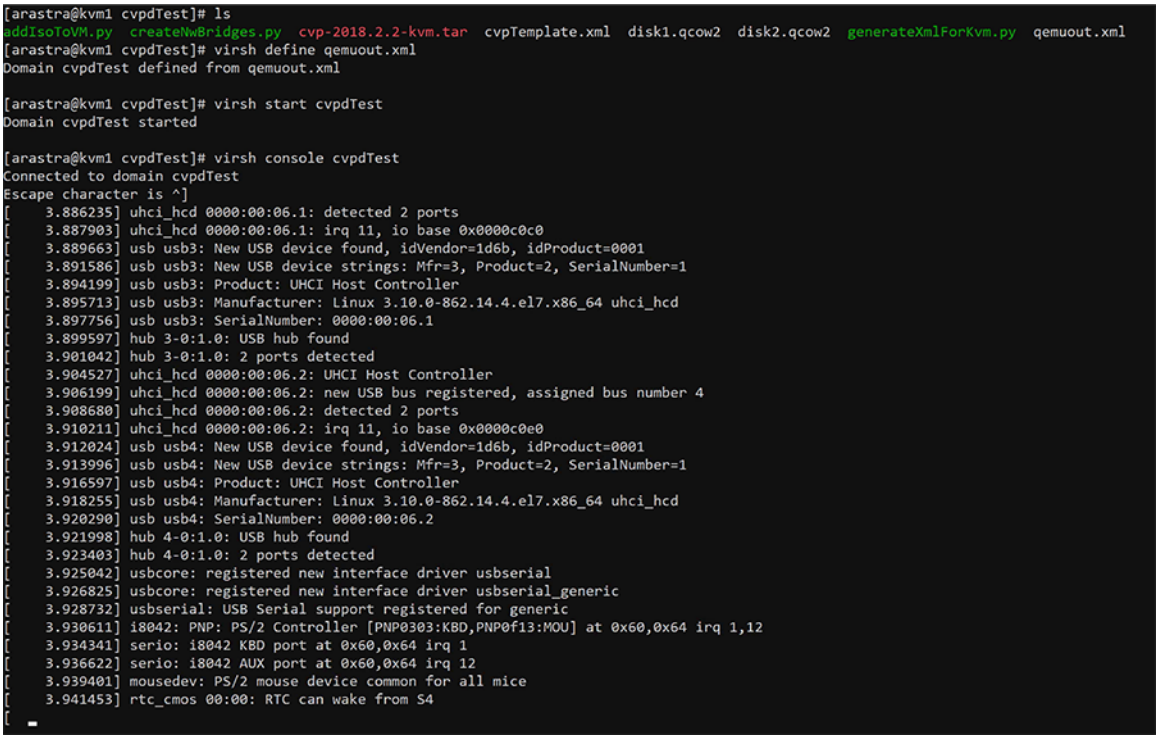

Figure 31: Defining and Launching the CVP VM

You can now login as cvpadmin and complete the configuration of the CVP application. See Configuring a Single-Node CVP Instance using CVP Shell for the steps used to complete the configuration.

Related topics:

- Shell-based Configuration
- ISO-based Configuration

- #unique_132

5.3 Set Up CVW on CV

This section describes the process to:

- Setup CVW on a Standalone CV
- Set Up CVW on a CV Cluster

5.3.1 Setup CVW on a Standalone CV

CVW is disabled by default.

To enable CVW, perform the following steps:

1. Log in to the CV admin shell via the cvpadmin user.
2. Enter **e** to edit the settings. The CV configuration wizard is launched.

   Note: If you are setting up CV for the first time, you need to enter the values for all the settings (DNS, IP addresses, etc.) in the configuration wizard. Refer to the Shell-based Configuration for information on these settings. If you have already set up or just upgraded CV, and you only want to enable CVW, go to Step 3.
3. Set the **CloudVision WiFi Enabled** option to **Yes**.

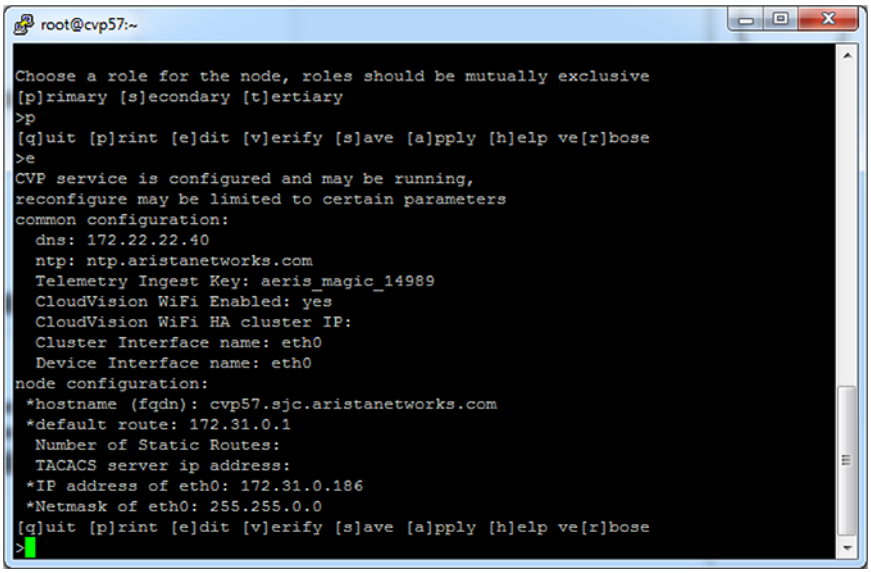

Figure 32: Setup CVW on a Standalone CV

4. Once the cursor is at the bottom of the configuration wizard, enter a to apply the configuration changes.

5.3.2 Set Up CVW on a CV Cluster

A few important points about the CVW service in a cluster deployment:

- CVW is disabled by default.
- For a CV cluster, you first need to Figure 33: Enable CVW on Primary Node and then Set Up CVW on Secondary and Tertiary Nodes.

  Note: The CVW service runs only on the primary and secondary nodes, but you need to apply the configuration changes to all the nodes, including the tertiary node. The CVW

service starts on both nodes only after the setup on all the nodes (including the tertiary node) of the cluster has been completed.

- The CV configuration wizard consists of two parts (Figure 33: Enable CVW on Primary Node):
  - **common configuration**: Settings common to all the nodes in the cluster (For example, DNS and services such as CVW).
  - **node configuration**: Settings specific to a node (For example, Hostname and IP settings).

5.3.2.1 Enable CVW on Primary Node

To enable CVW on the primary node, perform the following steps:

1. Log in to the CV admin shell via the **cvpadmin** user.
2. Enter **e** to edit the settings. The CV configuration wizard is launched.

   Note: If you are setting up CV for the first time, you need to enter the values for all the settings (those belonging to the common configuration as well as the node configuration). Refer to Shell-based Configuration and Shell Reconfiguration of Single-node, Multi-node Systems for information on these settings. If you have already set up or just upgraded CV, and you only want to enable CVW, go to Step 3.
3. You can optionally assign a **CloudVision WiFi HA Cluster IP**.

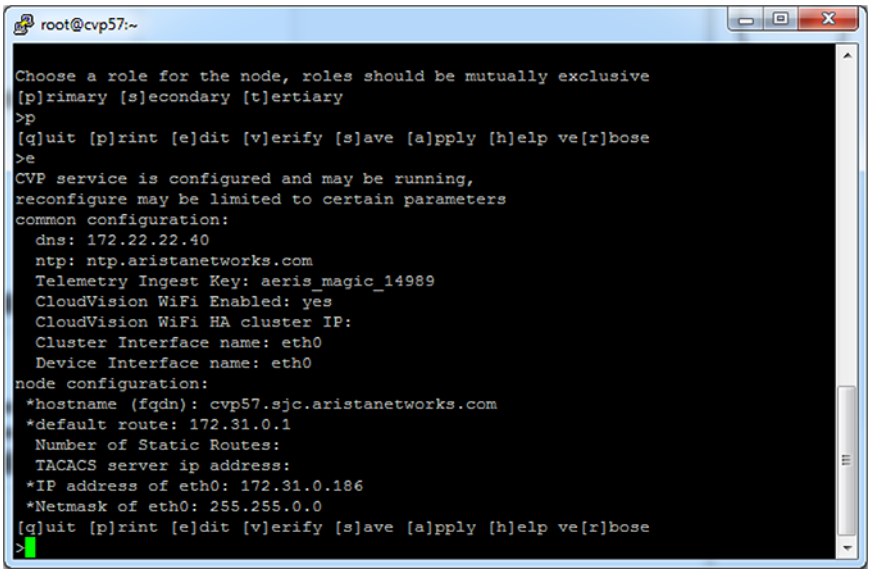

Figure 33: Enable CVW on Primary Node

   Note: CloudVision WiFi in HA mode configures an optional IP address, known as HA cluster IP that is automatically assigned to the active node in a cluster. Ensure that the HA Cluster IP address is different from the IP addresses of the actual device and cluster interfaces; but belongs to the same subnet as the Device Interface IP addresses of primary and secondary nodes. If HA cluster IP is not configured, IP addresses of both primary and secondary nodes must be configured on access points.
4. Set the **CloudVision WiFi Enabled** option to **Yes**.

5.3.2.2 Set Up CVW on Secondary and Tertiary Nodes

To set up CVW on the secondary and tertiary nodes, perform the following steps on the respective nodes:

1. Log in to the CV admin shell via the **cvpadmin** user.
2. Enter **e** to edit the settings. The CV configuration wizard is launched.

Note: The **Shell-based Configuration** settings are not editable on the secondary and tertiary nodes. If you are setting up CV for the first time, you need to enter the values for all the Shell Reconfiguration of Single-node, Multi-node Systems settings. If you have already set up or just upgraded CV, and you only want to enable CVW, go to Step 3.

3. Press **Enter** until the cursor reaches the bottom of the configuration wizard, past all the settings.
4. Once the cursor is at the bottom of the configuration wizard, enter **a** to apply the configuration changes.

Note: Whether **CloudVision WiFi Enabled** is set to **Yes** or **No**, applying the configuration changes causes the secondary and tertiary nodes to update their settings based on the primary node. This will start the CVW service on the primary and secondary nodes.

5.4 Shell-based Configuration

The shell-based configuration can be used to set up either a single-node CVP instance or multi-node CVP instances. The steps you use vary depending on whether you are setting up a single-node instance or a multi-node instance.

Cluster and device interfaces

A cluster interface is the interface that is able to reach the other two nodes in a multi-node installation. A device interface is the interface used by managed devices to connect to CVP. The ZTP configuration file is served over this interface. These two parameters are optional and default to eth0. Configuring these two interfaces is useful in deployments where a private network is used between the managed devices and a public-facing network is used to reach the other two cluster nodes and the GUI.

- Configuring a Single-Node CVP Instance using CVP Shell
- Configuring Multi-node CVP Instances Using the CVP Shell

5.4.1 Configuring a Single-Node CVP Instance using CVP Shell

After initial bootup, CVP can be configured at the VM's console using the CVP config shell. At points during the configuration, you must start the network, NTPD, and CVP services. Starting these services may take some time to complete before moving on to the next step in the process.

Pre-requisites:

Before you begin the configuration process, make sure that you:

- Launch the VM (see Deploying CVP OVA on ESX, or Deploying CVP on KVM.)

To configure CVP using the CVP config shell:

1. Login at the VM console as **cvpadmin**.
2. Enter your configuration and apply it (see the following example).

   In this example, the root password is not set (it is not set by default). In this example of a CVP shell, the bold text is entered by the **cvpadmin** user.

Note: To skip static routes, simply press enter when prompted for number of static routes.

```
localhost login: cvpadmin
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
Enter a command
[q]uit [p]rint [s]inglenode [m]ultinode [r]eplace [u]pgrade
>s
Enter the configuration for CloudVision Portal and apply it when done.
```

```
Entries marked with '*' are required.

common configuration:
  dns: 172.22.22.40
  ntp: ntp.aristanetworks.com
  Telemetry Ingest Key: magickey
  Cluster Interface name: eth0
  Device Interface name: eth0
node configuration:
 *hostname (fqdn): cvp80.sjc.aristanetworks.com
 *default route: 172.31.0.1
  Number of Static Routes: 1
  Route for Static Route #1: 1.1.1.0
  Netmask for Static Route #1: 255.255.255.0
  Interface for Static Route #1: eth0
  TACACS server ip address:
 *IP address of eth0: 172.31.0.168
 *Netmask of eth0: 255.255.0.0
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>v
Valid config format.
Applying proposed config for network verification.
saved config to /cvpi/cvp-config.yaml
Running : cvpConfig.py tool...
[  189.568543] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[  189.576571] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[  203.860624] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[  203.863878] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[  204.865253] Ebtables v2.0 unregistered
[  205.312888] ip_tables: (C) 2000-2006 Netfilter Core Team
[  205.331703] ip6_tables: (C) 2000-2006 Netfilter Core Team
[  205.355522] Ebtables v2.0 registered
[  205.398575] nf_conntrack version 0.5.0 (65536 buckets, 262144 max)
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[  206.856170] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[  206.858797] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[  206.860627] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[  207.096883] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[  211.086390] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[  211.089157] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[  211.091084] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[  211.092424] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
[  211.245437] warning: `/bin/ping' has both setuid-root and effective
 capabilities. Therefore not raising all capabilities.
Warning: External interfaces, ['eth1'], are discovered under /etc/
sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are correct.
Otherwise, actions from the CVP shell may fail.

Valid config.
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
```

5.4.2 Configuring Multi-node CVP Instances Using the CVP Shell

Use this procedure to configure multi-node CVP instances using the CVP shell. This procedure includes the steps to set up a primary, secondary, and tertiary node, which is the number of nodes required for redundancy. It also includes the steps to verify and apply the configuration of each node.

The sequence of steps in this procedure follow the process described in the basic steps in the process

Pre-requisites:

Before you begin the configuration process, make sure that you:

- Launch the VM (see Deploying CVP OVA on ESX, or Deploying CVP on KVM.)
- Login to the VM console for each of the three(3) nodes (login as **cvpadmin** on each node).

Complete the following steps to configure multi-node CVP instances:

1. Login at the VM console for the primary node as **cvpadmin**.
2. At the **cvp installation mode** prompt, type **m** to select a multi-node configuration.
3. At the prompt to select a role for the node, type **p** to select primary node.

   Note: You must select primary first. You cannot configure one of the other nodes before you configure the primary node.
4. Follow the CloudVision Portal prompts to specify the configuration options for the primary node. (All options with an asterisk (*) are required.) The options include:

   - Root password (*)
   - Default route (*)
   - DNS (*)
   - NTP (*)
   - Telemetry Ingest Key
   - Cluster interface name (*)
   - Device interface name (*)
   - Hostname (*)
   - IP address (*)
   - Netmask (*)
   - Number of static routes
   - Route for each static route
   - Interface for static route
   - TACACS server ip address
   - TACACS server key/port
   - IP address of primary (*) for secondary/tertiary only

   Note: If there are separate cluster and device interfaces (the interfaces have different IP addresses), make sure that you enter the hostname of the cluster interface. If the cluster and device interface are the same (for example, they are “eth0”), make sure you enter the IP address of “eth0” for the hostname.
5. At the following prompt, type **v** to verify the configuration.

   ```
   [q]uit, [p]rint, [e]dit, [v]erify, [s]ave, [a]pply, [h]elp ve[r]bose.
   ```

   If the configuration is valid, the system shows a “Valid config” status message.
6. Type “**a**” to apply the configuration for the primary node and wait for the line “Waiting for other nodes to send their hostname and ip” with spinning wheel.

   The system automatically saves the configuration as a YAML document and shows the configuration settings in pane 1 of the shell.)

7. When `Waiting for other nodes to send their hostname and ip` line is printed by the primary node, go to the shell for the **secondary** node, and specify the configuration settings for the **secondary** node (All options with an asterisk (*) are required, including primary node IP address)
8. At the following prompt, type **v** to verify the configuration.

```
[q]uit, [p]rint, [e]dit, [v]erify, [s]ave, [a]pply, [h]elp ve[r]bose.
```

   If the configuration is valid, the system shows a “Valid config” status message.
9. At the **Primary’s root password** prompt, type (enter) the password for the primary node, and then press **Enter**.
10. Go to the shell for the **tertiary** node, and specify the configuration settings for the node. (All options with an asterisk (*) are required.)
11. At the following prompt, type **v** to verify the configuration.

```
[q]uit, [p]rint, [e]dit, [v]erify, [s]ave, [a]pply, [h]elp ve[r]bose.
```

    If the configuration is valid, the system shows a “Valid config” status message.
12. At the **Primary IP** prompt, type the IP address of the primary node.
13. At the **Primary’s root password** prompt, press **Enter**.

    The system automatically completes the CVP installation for all nodes (this is done by the primary node). A message appears indicating that the other nodes are waiting for the primary node to complete the CVP installation.

    When the CVP installation is successfully completed for a particular node, a message appears in the appropriate pane to indicate the installation was successful. (This message is repeated in each pane.)
14. Go to shell for the primary node, and type **q** to quit the installation.
15. At the cvp login prompt, login as **root**.
16. At the **[root@cvplogin]#** prompt, switch to the **cvp** user account by typing **su cvp**, and then press **Enter**.
17. Run the `cvpi status all` command, and press **Enter**.

    The system automatically checks the status of the installation for each node and provides status information in each pane for CVP. The information shown includes some of the configuration settings for each node.

    For more information about the process, see:

    - Rules for the Number and Type of Nodes
    - The Basic Steps in the Process
    - The CVP Shell
    - Examples

5.4.2.1 Rules for the Number and Type of Nodes

Three nodes are required for multi-node CVP instances, where a node is identified as either the primary, secondary, or tertiary. You define the node type (primary, secondary, or tertiary) for each node during the configuration.

5.4.2.2 The Basic Steps in the Process

All multi-node configurations follow the same basic process. The basic steps are:

1. Specify the settings for the nodes in the following sequence (you apply the configuration later in the process):
   - Primary node
   - Secondary node

- Tertiary node

2. Verify and then apply the configuration for the **primary** node. (During this step, the system automatically saves the configuration for the primary node as a YAML document. In addition, the system shows the configuration settings.)

   Once the system applies the configuration for the primary node, the other nodes need to send their hostname and IP address to the primary node.
3. Verify and then apply the configuration for the **secondary** node.

   As part of this step, the system automatically pushes the hostname, IP address, and public key of the secondary node to the primary node. The primary node also sends a consolidated YAML to the secondary node, which is required to complete the configuration of the secondary node.
4. The previous step (verifying and applying the configuration) is repeated for the **tertiary** node. (The automated processing of data described for the secondary node is also repeated for the tertiary node.)

   Once the configuration for all nodes has been applied (steps 1 through 4 above), the system automatically attempts to complete the CVP installation for all nodes (this is done by the primary node). A message appears indicating that the other nodes are waiting for the primary node to complete the CVP installation.
5. You quit the installation, then login as root and check the status of CVP.

   The system automatically checks the status and provides status information in each pane for the CVP service.

5.4.2.3 The CVP Shell

For multi-node configurations, you need to open 3 CVP consoles (one for each node). Each console is shown in it's own pane. You use each console to configure one of the nodes (primary, secondary, or tertiary).

The system also provides status messages and all of the options required to complete the multi-node configuration. The status messages and options are presented in the panes of the shell that correspond to the node type.

Figure 34: CVP Console Shells for Multi-node Configurations shows three CVP Console shells for multi-node configurations. Each shell corresponds to a CVP Console for each node being configured.

Figure 34: CVP Console Shells for Multi-node Configurations

5.4.2.4 Examples

The following examples show the commands used to configure (set up) the primary, secondary, and tertiary nodes, and apply the configurations to the nodes. Examples are also included of the system output shown as CVP completes the installation for each of the nodes.

- Primary Node Configuration
- Secondary Node Configuration
- Tertiary Node Configuration
- Verifying the Primary Node Configuration and Applying it to the Node
- Verifying the Tertiary Node Configurations and Applying them to the Nodes
- Waiting for the Primary Node Installation to Finish
- Waiting for the Secondary and Tertiary Node Installation to Finish

5.4.2.4.1 Primary Node Configuration

This example shows the commands used to configure (set up) the primary node.

```
localhost login: cvpadmin
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
Enter a command
[q]uit [p]rint [s]inglenode [m]ultinode [r]eplace [u]pgrade
>m
Choose a role for the node, roles should be mutually exclusive
[p]rimary [s]econdary [t]ertiary
>p

Enter the configuration for CloudVision Portal and apply it when
 done.
Entries marked with '*' are required.

common configuration:
  dns: 172.22.22.40
  ntp: ntp.aristanetworks.com
  Telemetry Ingest Key: magickey
  Cluster Interface name: eth0
  Device Interface name: eth0
node configuration:
 *hostname (fqdn): cvp57.sjc.aristanetworks.com
 *default route: 172.31.0.1
  Number of Static Routes:
  TACACS server ip address:
 *IP address of eth0: 172.31.0.186
 *Netmask of eth0: 255.255.0.0
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>
```

5.4.2.4.2 Secondary Node Configuration

This example shows the commands used to configure (set up) the secondary node.

```
localhost login: cvpadmin
Changing password for user root.
```

```
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
Enter a command
[q]uit [p]rint [s]inglenode [m]ultinode [r]eplace [u]pgrade
>m
Choose a role for the node, roles should be mutually exclusive
[p]rimary [s]econdary [t]ertiary
>s

Enter the configuration for CloudVision Portal and apply it when
 done.
Entries marked with '*' are required.

common configuration:
  dns: 172.22.22.40
  ntp: ntp.aristanetworks.com
  Telemetry Ingest Key: magickey
  Cluster Interface name: eth0
  Device Interface name: eth0
 *IP address of primary: 172.31.0.186
node configuration:
 *hostname (fqdn): cvp65.sjc.aristanetworks.com
 *default route: 172.31.0.1
  Number of Static Routes:
  TACACS server ip address:
 *IP address of eth0: 172.31.0.153
 *Netmask of eth0: 255.255.0.0
>
```

5.4.2.4.3 Tertiary Node Configuration

This example shows the commands used to configure (set up) the tertiary node.

```
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
Enter a command
[q]uit [p]rint [s]inglenode [m]ultinode [r]eplace [u]pgrade
>m
Choose a role for the node, roles should be mutually exclusive
[p]rimary [s]econdary [t]ertiary
>t

Enter the configuration for CloudVision Portal and apply it when
 done.
Entries marked with '*' are required.

common configuration:
  dns: 172.22.22.40
  ntp: ntp.aristanetworks.com
  Telemetry Ingest Key: magickey
  Cluster Interface name: eth0
  Device Interface name: eth0
 *IP address of primary: 172.31.0.186
node configuration:
  hostname (fqdn): cvp84.sjc.aristanetworks.com
 *default route: 172.31.0.1
  Number of Static Routes:
```

```
  TACACS server ip address:
 *IP address of eth0: 172.31.0.213
 *Netmask of eth0: 255.255.0.0
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>
```

5.4.2.4.4 Verifying the Primary Node Configuration and Applying it to the Node

This example shows the commands used to verify the configuration of the primary node and apply the configuration to the node.

```
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>v
Valid config format.
Applying proposed config for network verification.
saved config to /cvpi/cvp-config.yaml
Running : cvpConfig.py tool...
[ 8608.509056] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[ 8608.520693] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[ 8622.807169] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[ 8622.810214] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[ 8624.027029] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[ 8624.030254] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[ 8624.032643] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[ 8624.238995] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[ 8638.294690] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[ 8638.297973] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[ 8638.300454] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[ 8638.302186] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
[ 8638.489266] warning: `/bin/ping' has both setuid-root and
 effective capabilities. Therefore not raising all capabilities.
Warning: External interfaces, ['eth1'], are discovered under /
etc/sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are
 correct.
Otherwise, actions from the CVP shell may fail.

Valid config.
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>
```

5.4.2.4.5 Verifying the Tertiary Node Configurations and Applying them to the Nodes

This example shows the commands used to verify the configurations of the tertiary nodes and apply the configurations to the nodes.

```
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>v
Valid config format.
Applying proposed config for network verification.
saved config to /cvpi/cvp-config.yaml
Running : cvpConfig.py tool...
[ 9195.362192] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[ 9195.365069] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[ 9195.367043] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[ 9195.652382] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[ 9209.588173] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[ 9209.590896] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[ 9209.592887] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[ 9209.594222] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[ 9210.561940] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[ 9210.564602] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[ 9224.805267] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[ 9224.808891] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[ 9224.811150] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[ 9224.812899] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Warning: External interfaces, ['eth1'], are discovered under /
etc/sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are
 correct.
Otherwise, actions from the CVP shell may fail.

Valid config.
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>
```

5.4.2.4.6 Waiting for the Primary Node Installation to Finish

These examples show the system output shown as CVP completes the installation for the primary node.

- Waiting for primary node installation to pause until other nodes send files

```
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>a
Valid config format.
saved config to /cvpi/cvp-config.yaml
Applying proposed config for network verification.
saved config to /cvpi/cvp-config.yaml
Running : cvpConfig.py tool...
[15266.575899] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15266.588500] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[15266.591751] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15266.672644] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[15280.937599] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15280.941764] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[15280.944883] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15280.947038] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[15282.581713] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15282.585367] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[15282.588072] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15282.948613] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[15296.871658] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15296.875871] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[15296.879003] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15296.881456] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
Warning: External interfaces, ['eth1'], are discovered under /etc/
sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are correct.
Otherwise, actions from the CVP shell may fail.

Valid config.
Running : cvpConfig.py tool...
[15324.884887] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15324.889169] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[15324.893217] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15324.981682] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[15339.240237] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15339.243999] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[15339.247119] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15339.249370] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[15340.946583] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15340.950891] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[15340.953786] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15341.251648] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
```

```
[15355.225649] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15355.229400] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[15355.232674] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15355.234725] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
Waiting for other nodes to send their hostname and ip
\
```

- Waiting for the primary node installation to finish

```
Waiting for other nodes to send their hostname and ip
-
Running : cvpConfig.py tool...
[15707.665618] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15707.669167] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[15707.672109] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15708.643628] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[15722.985876] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15722.990116] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[15722.993221] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15722.995325] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
[15724.245523] Ebtables v2.0 unregistered
[15724.940390] ip_tables: (C) 2000-2006 Netfilter Core Team
[15724.971820] ip6_tables: (C) 2000-2006 Netfilter Core Team
[15725.011963] Ebtables v2.0 registered
[15725.077660] nf_conntrack version 0.5.0 (65536 buckets, 262144 max)
Stopping: ntpd
Running : /bin/sudo /sbin/service ntpd stop
Running : /bin/sudo /bin/systemctl is-active ntpd
Starting: ntpd
Running : /bin/sudo /bin/systemctl start ntpd.service
-
-
Verifying configuration on the secondary node
Verifying configuration on the tertiary node
Starting: systemd services
Starting: cvpi-check
Running : /bin/sudo /bin/systemctl start cvpi-check.service
Starting: zookeeper
Running : /bin/sudo /bin/systemctl start zookeeper.service
Starting: cvpi-config
Running : /bin/sudo /bin/systemctl start cvpi-config.service
Starting: cvpi
Running : /bin/sudo /bin/systemctl start cvpi.service
Running : /bin/sudo /bin/systemctl enable zookeeper
Running : /bin/sudo /bin/systemctl start cvpi-watchdog.timer
Running : /bin/sudo /bin/systemctl enable docker
Running : /bin/sudo /bin/systemctl start docker
Running : /bin/sudo /bin/systemctl enable kube-cluster.path
Running : /bin/sudo /bin/systemctl start kube-cluster.path
Waiting for all components to start. This may take few minutes.
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.36
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.56
[15843.549983] FS-Cache: Loaded
[15843.645979] FS-Cache: Netfs 'nfs' registered for caching
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.10
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.53
[15904.022085] hrtimer: interrupt took 4615311 ns
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.06
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 16.96
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 3.31
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 3.56
```

```
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 4.03
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 6.39
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.80
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 22.44
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.62
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 10.90
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.34
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 11.25
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 5.13
CVP installation successful
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>
```

5.4.2.4.7 Waiting for the Secondary and Tertiary Node Installation to Finish

This example shows the system output displayed as CVP completes the installation for the secondary and tertiary nodes.

```
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>a
Valid config format.
saved config to /cvpi/cvp-config.yaml
Applying proposed config for network verification.
saved config to /cvpi/cvp-config.yaml
Running : cvpConfig.py tool...
[15492.903419] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15492.908473] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[15492.910297] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15493.289569] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[15507.118778] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15507.121579] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[15507.123648] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15507.125051] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[15508.105909] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15508.108752] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[15522.301114] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15522.303766] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[15522.305580] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15522.306866] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Warning: External interfaces, ['eth1'], are discovered under /
etc/sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are
 correct.
Otherwise, actions from the CVP shell may fail.
```

```
Valid config.
Running : cvpConfig.py tool...
[15549.664989] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15549.667899] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[15549.669783] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15550.046552] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[15563.933328] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15563.937507] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[15563.940501] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15563.942113] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Stopping: network
Running : /bin/sudo /sbin/service network stop
Running : /bin/sudo /bin/systemctl is-active network
Starting: network
Running : /bin/sudo /bin/systemctl start network.service
[15565.218666] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15565.222324] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[15565.225193] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15565.945531] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[15579.419911] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15579.422707] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[15579.424636] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15579.425962] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Running : cvpConfig.py tool...
[15600.608075] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15600.610946] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[15600.613687] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15600.986529] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[15615.840426] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15615.843207] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[15615.845197] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15615.846633] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
[15616.732733] Ebtables v2.0 unregistered
[15617.213057] ip_tables: (C) 2000-2006 Netfilter Core Team
[15617.233688] ip6_tables: (C) 2000-2006 Netfilter Core Team
[15617.261149] Ebtables v2.0 registered
[15617.309743] nf_conntrack version 0.5.0 (65536 buckets, 262144
 max)
Stopping: ntpd
Running : /bin/sudo /sbin/service ntpd stop
Running : /bin/sudo /bin/systemctl is-active ntpd
Starting: ntpd
Running : /bin/sudo /bin/systemctl start ntpd.service
Pushing hostname, ip address and public key to the primary node
Primary's root password:
```

```
Transferred files
Receiving public key of the primary node
-
Waiting for primary to send consolidated yaml
-
Received authorized keys and consolidated yaml files
Running : /bin/sudo /bin/systemctl start cvpi-watchdog.timer
Running : cvpConfig.py tool...
[15748.205170] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[15748.208393] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000
 Mbps
[15748.210206] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[15748.591559] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes
 ready
[15752.406867] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[15752.409789] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000
 Mbps
[15752.412015] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[15752.413603] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes
 ready
Stopping: zookeeper
Running : /bin/sudo /sbin/service zookeeper stop
Running : /bin/sudo /bin/systemctl is-active zookeeper
Stopping: cvpi-check
Running : /bin/sudo /sbin/service cvpi-check stop
Running : /bin/sudo /bin/systemctl is-active cvpi-check
Stopping: ntpd
Running : /bin/sudo /sbin/service ntpd stop
Running : /bin/sudo /bin/systemctl is-active ntpd
Starting: ntpd
Running : /bin/sudo /bin/systemctl start ntpd.service
Starting: cvpi-check
Running : /bin/sudo /bin/systemctl start cvpi-check.service
Starting: zookeeper
Running : /bin/sudo /bin/systemctl start zookeeper.service
Running : /bin/sudo /bin/systemctl enable docker
Running : /bin/sudo /bin/systemctl start docker
Running : /bin/sudo /bin/systemctl enable kube-cluster.path
Running : /bin/sudo /bin/systemctl start kube-cluster.path
Running : /bin/sudo /bin/systemctl enable zookeeper
Running : /bin/sudo /bin/systemctl enable cvpi
Waiting for primary to finish configuring cvp.
-
Please wait for primary to complete cvp installation.
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>
```

5.5 Shell Reconfiguration of Single-node, Multi-node Systems

The configuration of single-node systems and multi-node systems can be reconfigured using the CVP shell, even after the installation is complete. The reconfiguration process brings down the applications and CVPI for a brief period of time until reconfiguration is complete.

- Single-node Shell Reconfiguration
- Multi-node Shell Reconfiguration

5.5.1 Single-node Shell Reconfiguration

The process for reconfiguring a single-node system is based on the process used to complete the initial installation. You can change any of the configuration settings during the reconfiguration.

Note: The system must be in healthy state before reconfiguration is attempted.

To change an existing single-node configuration, do the following:

1. Follow the same steps you use for an initial single-node, shell-based install (see Configuring a Single-Node CVP Instance using CVP Shell).
2. When prompted with the message **Are you sure you want to replace config and restart? yes/no:** enter **yes**, and then press **Enter**. (Make sure there are no configuration errors.)

   This system automatically completes the configuration.

5.5.2 Multi-node Shell Reconfiguration

The process for reconfiguring a multi-node system is based on the process used to complete the initial installation. Just like initial installations, you can only edit the configuration of the node you are logged into.

- Configurable and Read-only Parameters
- #unique_152
- Example of Primary Node Reconfiguration
- Procedure

5.5.2.1 Configurable and Read-only Parameters

You can change some, but not all of the configuration settings during the reconfiguration. The configuration parameters you cannot change are read-only after the initial configuration.

The configurable and read-only parameters are:

- Configurable parameters
  - default route (gateway)
  - dns
  - ntp
  - aeris ingest key
  - TACACS server IP address
  - TACACS server key/port
  -
- Read-only parameters
  - Cluster interface name
  - Device interface name
  - hostname (fqdn)
  - ip address
  - netmask
  - Number of static routes
  - Route for each static route
  - Interface for static route
  - Primary IP address (use current primary ip address)

Note: The cluster must be in healthy state before reconfiguration is attempted. Also, do not edit `cvp-config.yaml` directly. Make sure you use the shell-based install to reconfigure it.

5.5.2.3 Example of Primary Node Reconfiguration

```
localhost login: cvpadmin
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
Enter a command
[q]uit [p]rint [s]inglenode [m]ultinode [r]eplace [u]pgrade
>m
Choose a role for the node, roles should be mutually exclusive
[p]rimary [s]econdary [t]ertiary
>p
...
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>e
CVP service is configured and may be running,
reconfigure may be limited to certain parameters
common configuration:
  dns: 172.22.22.40
  ntp: ntp.aristanetworks.com
  Telemetry Ingest Key: modified_ingest_key_for_telemetry <-- modified
 key
  Cluster Interface name: eth0
  Device Interface name: eth0
node configuration:
 *hostname (fqdn): cvp57.sjc.aristanetworks.com
 *default route: 172.31.0.1
  Number of Static Routes:
  TACACS server ip address:
 *IP address of eth0: 172.31.0.186
 *Netmask of eth0: 255.255.0.0
>v
Valid config format.
Using existing settings for new proposed network verification.
Warning: External interfaces, ['eth1'], are discovered under /etc/
sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are correct.
Otherwise, actions from the CVP shell may fail.

Valid config.
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>a
Valid config format.
saved config to /cvpi/cvp-config.yaml
Using existing settings for new proposed network verification.
Warning: External interfaces, ['eth1'], are discovered under /etc/
sysconfig/network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are correct.
Otherwise, actions from the CVP shell may fail.

Valid config.
Are you sure you want to replace config and restart? yes/no: no
```

5.5.2.4 Procedure

To change an existing multi-node configuration, do the following:

1. Follow the same steps you use for an initial multi-node, shell-based install (see #unique_155).

2. When prompted with the message **Are you sure you want to replace config and restart? yes/no:** enter **yes**, and then press **Enter**. (Make sure there are no configuration errors.)

   Note: You will also be prompted for primary node ip address and root passwords during reconfiguration.

5.6 ISO-based Configuration

The ISO-based configuration can be used to set up either a single-node or multi-node CVP instance(s). Before configuring and starting CVP, the following tasks must be completed.

Quick Start Steps:

- Launch the VM (seeDeploying CVP OVA on ESX or Deploying CVP on KVM).
- Create a YAML Document
- Feed the YAML File into the geniso.py Tool
- Map ISO to VM's CD-ROM Drive
- Verify the host name, reachability of the name server, and VM connectivity.

5.6.1 Create a YAML Document

Create a YAML document describing the node(s) (one or three) in your CVP deployment. When creating a YAML document, the following should be considered:

- The version field is required and must be 2.
- The "dns" and "ntp" entries are lists of values.
- The "dns", and "ntp" parameters are optional, but recommended to use.

  Note: The parameters, which are the same for all nodes, can be specified only once in the "common" section of the YAML. For example, "default_route" can be specified only once in the common section and not three times, once for each node.

Example:

The following example of a YAML document shows the use of separate (different) interfaces for cluster and device-facing networks. These parameters are explained in the previous section. For a single-node deployment, remove the sections for "node2" and "node3" (assuming all nodes are on the same subnet and have the same default route).

```
>cat multinode.yaml
version: 2
common:
 aeris_ingest_key: magickey
 cluster_interface: eth0
 default_route: 172.31.0.1
 device_interface: eth0
 dns:
 - 172.22.22.40
 ntp:
 - ntp.aristanetworks.com
node1:
hostname: cvp6.sjc.aristanetworks.com
interfaces:
eth0:
 ip_address: 172.31.3.236
 netmask: 255.255.0.0
 vmname: cvp6

node2:
   vmname: cvp9
```

```
   hostname : cvp9.sjc.aristanetworks.com
   interfaces:
      eth0:
         ip_address: 172.31.3.239
         netmask: 255.255.0.0
      eth1:
         ip_address: 10.0.0.2
         netmask: 255.255.255.0
node3:
   vmname: cvp10
   hostname: cvp10.sjc.aristanetworks.com
   interfaces:
      eth0:
         ip_address: 172.31.3.240
         netmask: 255.255.0.0
      eth1:
         ip_address: 10.0.0.3
         netmask: 255.255.255.0
```

5.6.2 Feed the YAML File into the *geniso.py* Tool

Once you have created the YAML file, you are ready to feed it into the tool so that you can generate the ISO files for the CVP nodes. The root password can be provided at the command line or prompted from the user. If password is empty, no password will be set for root.

Note: The `geniso.py` tool is provided by `cvp-tools-1.0.1.tgz` which can be found at https://www.arista.com/en/support/software-download. The package also contains a README file with more details and requirements for `geniso.py`.

Complete the following steps:

1. Run the `yum install mkisofs` command.
2. Feed the YAML document into the `geniso.py` tool.

   The system generates the ISO files for the nodes using the input of the YAML document.

   Example:

   - In this example, you are prompted for the root password.

```
> mkdir tools
> tar zxf cvp-tools-1.0.1.tgz -C tools
> cd tools

...<edit multinode.yaml>...

> ./geniso.py -y multinode.yaml
Please enter a password for root user on cvp
Password:
Please re-enter the password:
Building ISO for node1 cvp1: cvp.iso.2015-11-04_00:16:23/node1-cvp1.
iso
Building ISO for node2 cvp2: cvp.iso.2015-11-04_00:16:23/node2-cvp2.
iso
Building ISO for node3 cvp3: cvp.iso.2015-11-04_00:16:23/node3-cvp3.
iso
```

3. In case of using KVM as a hypervisor in a multi-node setup, copy the following ISO files to the corresponding nodes:

   - SCP node2's ISO to node 2

```
[root@localhost cvp]# scp node2-cvp-appliance-2.iso root@172.28.1
61.44://data/cvp/
```

```
root@172.28.161.44's password:
node2-cvp-appliance-2.iso

100%  360KB  57.5MB/s   00:00
```

- SCP node3's ISO to node 3

```
[root@localhost cvp]# scp node3-cvp-appliance-3.iso root@172.28.1
61.45://data/cvp/
root@172.28.161.45's password:
node3-cvp-appliance-3.iso

100%  360KB  54.7MB/s   00:00
```

Note: The script has to be run on one machine only. This generates three ISO images which contains the same ssh keys, thus allowing the nodes to send files without a password. If the script is run individually on each node, it result in images containing different ssh keys and the deployment process fails, until the user manually adds the ssh keys in `~/.ssh/authorized_keys`.

5.6.3 Map ISO to VM's CD-ROM Drive

You can map the ISO to the VM's CD-ROM drive through either ESXi or KVM.

5.7 Certificate-Based TerminAttr Authentication

Arista/EOS switches use TerminAttr for streaming network data to CVP. Each TerminAttr connection must be authenticated using either shared keys or certificate. The certificate-based TerminAttr authentication provides the following additional security features:

- Eliminates the shared key from the switch's configuration
- Uniquely authenticates each TerminAttr connection between the switch and CVP

Note: Third party devices can use only the shared key authentication. The minimum required version of TerminAttr to use this feature is *v1.6.1*.

The following sections describes configuring devices with certificate-based TerminAttr authentication:

- Enabling Certificate-Based TerminAttr Authentication
- Reboarding Existing Devices
- Re-ZTP On-Boarded Devices
- Switching the Authentication from Shared Keys to Certificates
- Switching the Authentication from Certificates to Shared Keys

5.7.1 Enabling Certificate-Based TerminAttr Authentication

When on-boarding a device through Zero Touch Provisioning (ZTP) or direct import, the certificate-based TerminAttr authentication uses a temporary token to enroll client certificates from CVP. The SYS_TelemetryBuilderV3 generates the TerminAttr configuration that uses certificate-based TerminAttr authentication.

Note: By default, CVP authenticates TerminAttr connections using shared keys.

Perform the following steps to enable certificate-based TerminAttr authentication:

1. In CloudVision portal, click the gear icon at the upper right corner of the page.

   The system displays the Settings screen.

2. Under the Cluster Management pane, enable **Device authentication via certificates** using the toggle button.

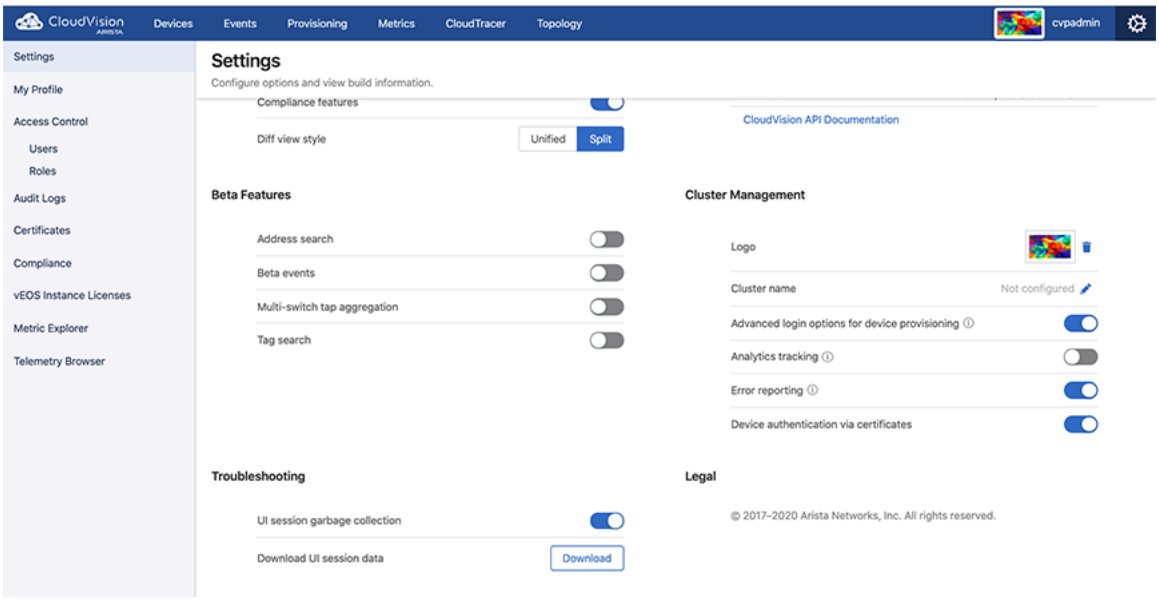

Figure 35: Enable Device Authentication via Certificates

5.7.2 Reboarding Existing Devices

You must reboard a device when the certificate-based TerminAttr authentication fails due to missing or invalid client certificates.

Perform the following steps to reboard devices:

1. In CloudVision portal, click the **Devices** tab.

   The system displays the Inventory screen.

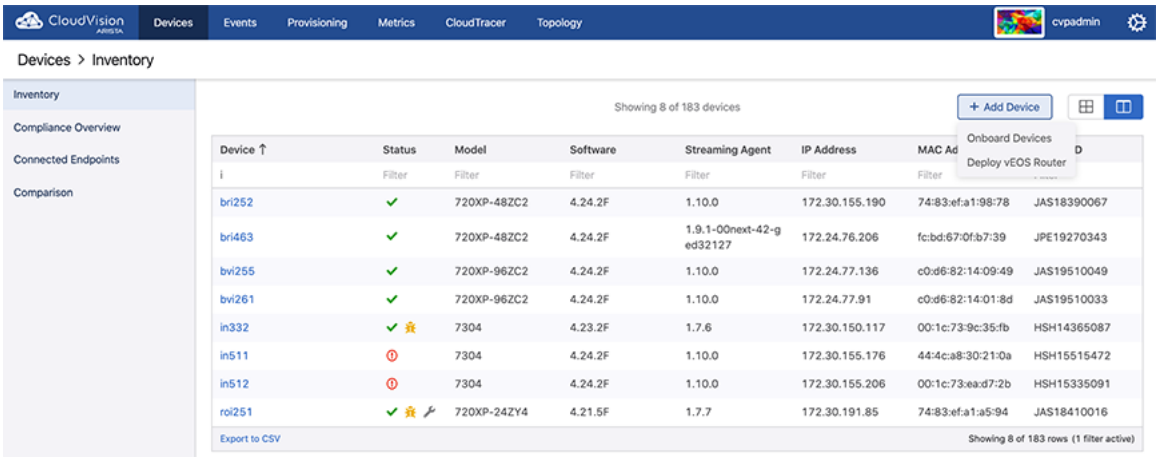

Figure 36: Inventory Screen

2. Select **Onboard Devices** from the **Add Device** drop-down menu at the upper right corner of the **Inventory** screen.

   The system displays the Onboard Devices pop-up window.

3. Click the **Existing Device Registration** tab at the lower end of the **Onboard Devices** pop-up window.

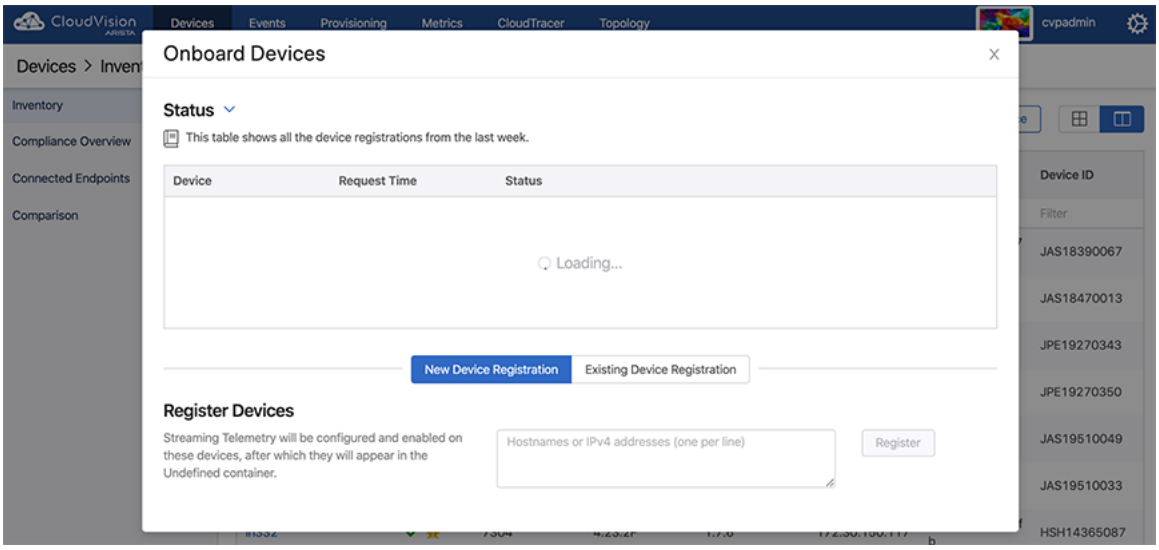

Figure 37: Existing Device Registration Tab

Note: To view all devices, disable the **Show only inactive devices** option using the toggle button.

4. Select the required device.
5. Click **Register n Device(s)** where *n* is the count of selected devices.

   The system refreshes the selected device with new certificates, returns to the last provisioning state, and resumes streaming to CVP.

5.7.3 Re-ZTP On-Boarded Devices

Manual intervention is required to re-ZTP on-boarded devices after enabling the certificate-based TerminAttr authentication. This prevents unauthorized or malicious software from spoofing previously on-boarded devices.

Perform the following steps to re-ZTP devices:

1. In CloudVision portal, click the **Devices** tab.

   The system displays the Inventory screen.
2. Select Re-ZTP Devices from the Add Device drop-down menu at the upper right corner of the Inventory screen.

   The system displays the Re-ZTP Devices pop-up window.

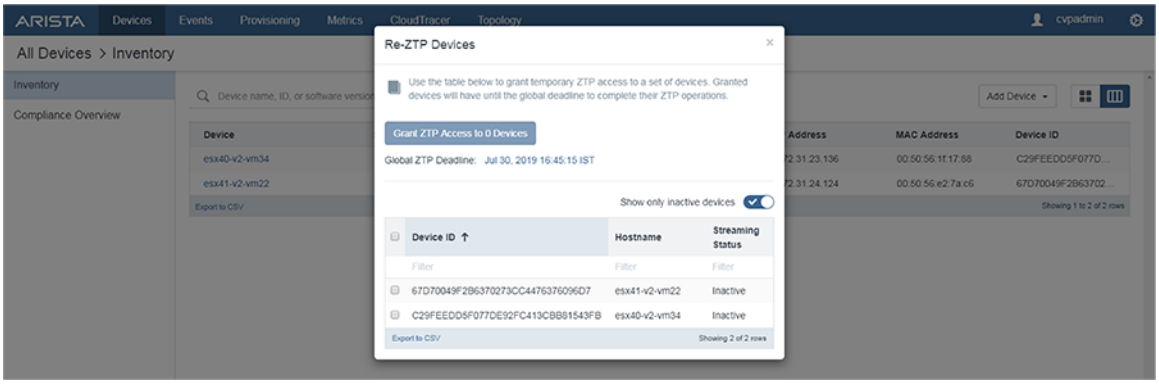

Figure 38: Re-ZTP Devices Pop-Up Window

Note: To view all devices, disable the Show only inactive devices option using the toggle button.

3. Select the required device.
4. (Optional) Click the time next to Global ZTP Deadline and configure the preferred time to re-ZTP selected devices.
5. Click **Grant ZTP Access to *n* Device(s)** where n is the count of selected devices.

   Devices must complete their re-ZTP before the enrollment window closes.

Chapter 6

Getting Started (CVP)

The login screen is displayed when you first connect to the application using a web browser.

The CloudVision Portal (CVP) application is accessible after the CVP service has been started on the appliance. The login screen is displayed when you first connect to the application using a web browser. JavaScript must be enabled in the browser for the web application to work.

Sections in this chapter include:

6.1 Accessing the CVP Login Page

1. To access the login page, point your browser to the CloudVision Portal (http://HOSTNAME or https://HOSTNAME).The system opens the CVP login page.

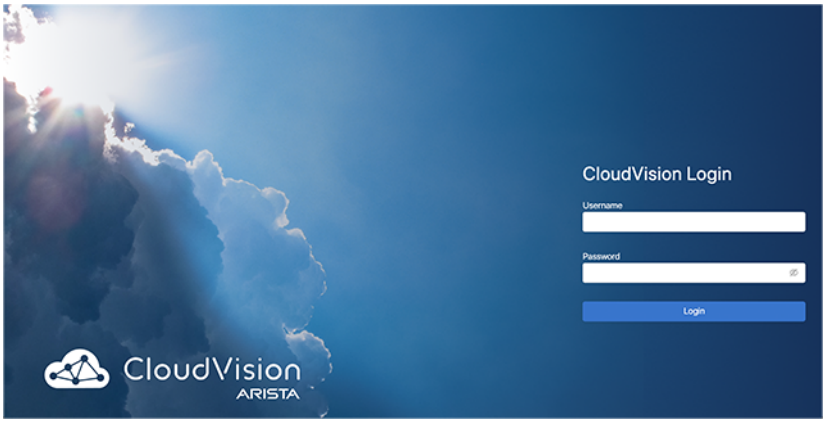

Figure 39: CVP Login Page

2. Enter login credentials in the CVP login section.

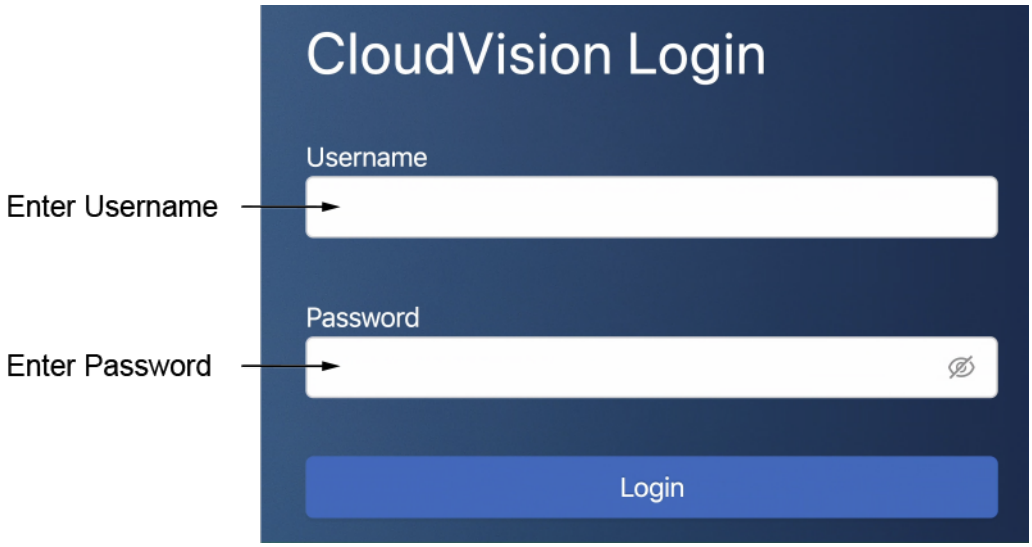

Figure 40: Login Section

Note:

The username and passwords required will depend on the authentication method and accounts previously set up. Login using the username and password created when CVP was installed. If you chose the local authentication and authorization options, login initially using *cvpadmin* for the username and password.

3. Click **Login**. The system opens the CVP home page.

6.2 Accessing the Home Page

All features like Devices, Events, Provisioning, Metrics, CloudTracer, Topology, Inventory, and Compliance are displayed on the home panel. A service dashboard scroller also exists to the right of the screen.

Note: You must have required privileges to access a switch.

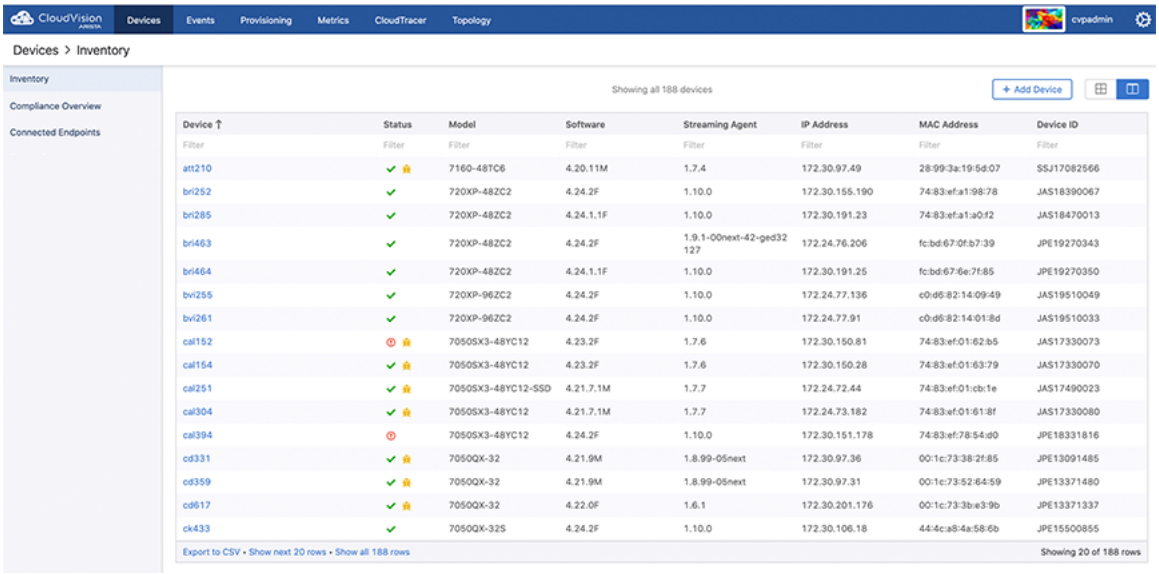

Figure 41: Home Page

The home page provides the following selections.

- **Devices**: View all devices across multiple topologies.
- **Events**: View multiple events on multiple devices.
- **Provisioning**: Hierarchical tree structure of the network is maintained here. All the configuration and image assignment to the network switches are made via this module.
- **Metrics**: View multiple metrics across multiple devices. Select at least one metric and one device to begin.
- **CloudTracer**: CloudTracer metrics across multiple devices or hosts. Select at least one metric and one device or host to begin.
- **Topology**: View the location of devices in individual topologies.

6.3 Customizing the Home Screen and Dashboard Logo

CloudVision enables you to customize the visible options and dashboard logo shown on the home page. You change the visible options and dashboard logo by customizing them from the Settings page.

By default, no dashboard logo is selected. The image you select for the logo appears in the dashboard next to the notifications icon.

Note: Note Any image you select for either the Home screen background or dashboard logo must not exceed 200 KB for each image. In addition, the images must JPG, PNG, or GIF.

Complete the following steps to customize the visible and dashboard logo:

1. Login to CVP.
2. Click the gear icon at the upper right corner of the page.

3. Click **Settings** in the left menu.
4. Select the required options provided under **Basic Settings**, **Beta Features**, **Cluster Management**, and **Troubleshooting** sections.

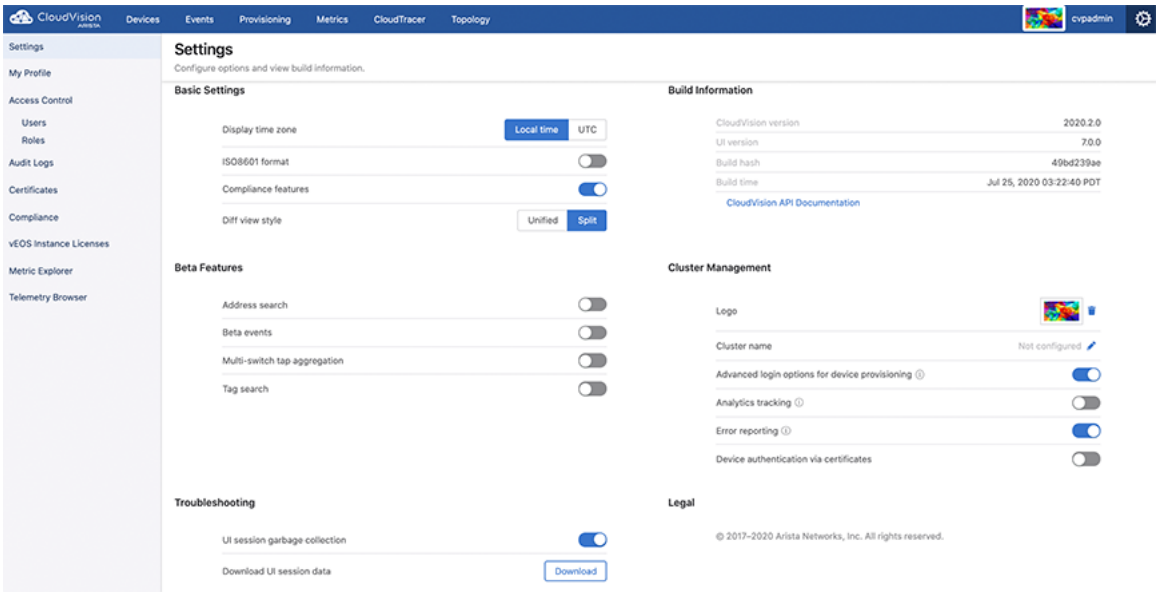

Figure 42: Default Settings for Home Page and Dashboard Logo

5. To customize the dashboard logo, perform the following steps:
   - Click the image box next to the logo field.
   - In the Upload logo dialog, Click **Select file**.
   - Navigate to the desired image, and click **Open**. (The imported image is displayed next the Select file box.)
   - Click **Upload**.

6.4 Accessing CloudVision Wifi

You can access the CloudVision Wifi (CVW) service via either the CLI Access or the UI Access.

CLI Access

To log in to the wifimanager container using CLI, run the `/cvpi/apps/wifimanager/bin/wifimanager.sh cli 2>/dev/null` command on the primary or the secondary node.

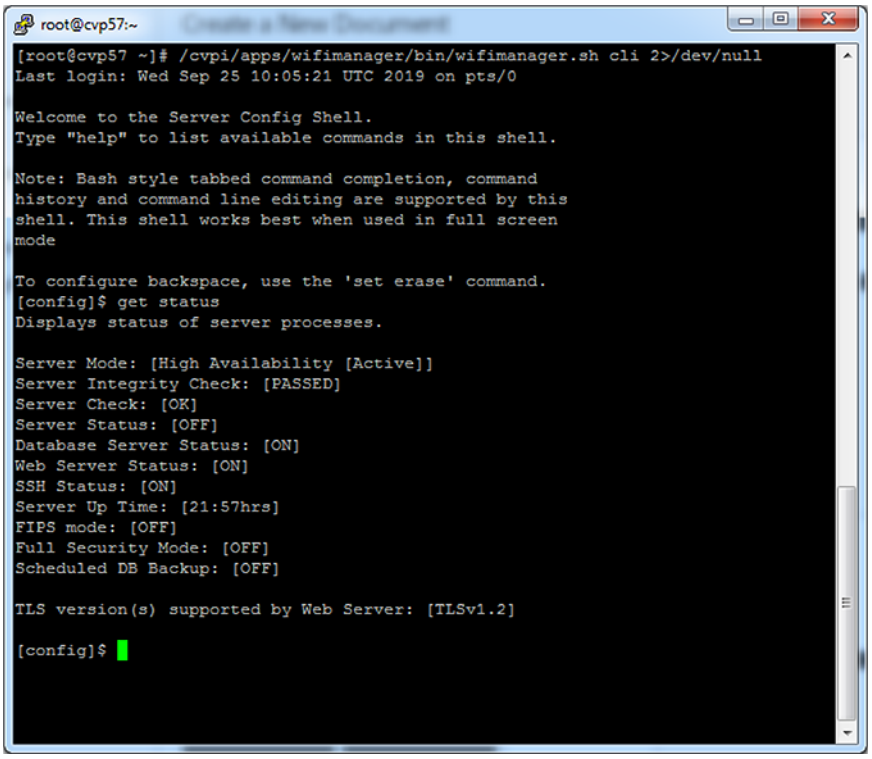

Figure 43: CLI Access

You can now run wifimanager commands. See the Wifimanager CLI Commands for a list of wifimanager CLI commands and their descriptions.

UI Access

The URL to access the wifimanager UI is **http(s)://<*CVP-IP*>/wifi/wifimanager** is where CVP-IP refers to the actual CloudVision Portal (CVP) IP/domain name.

The URL to access the cognitive Wifi UI is **http(s)://<*CVP-IP*>/wifi/aware** where ***CVP-IP*** refers to either the actual CVP IP or domain name.

For example, if the IP address of CVP is *10.12.3.4*, then the URL to access the wifimanager UI is https://10.12.3.4/wifi.wilimanager and the cognitive Wifi UI is https://10.12.3.4/wifi/aware.

You can access CVW UI by clicking on the **WiFi** tab in the CVP UI, or you can access it directly using the URLs of either wifimanager UI or Wifi UI.

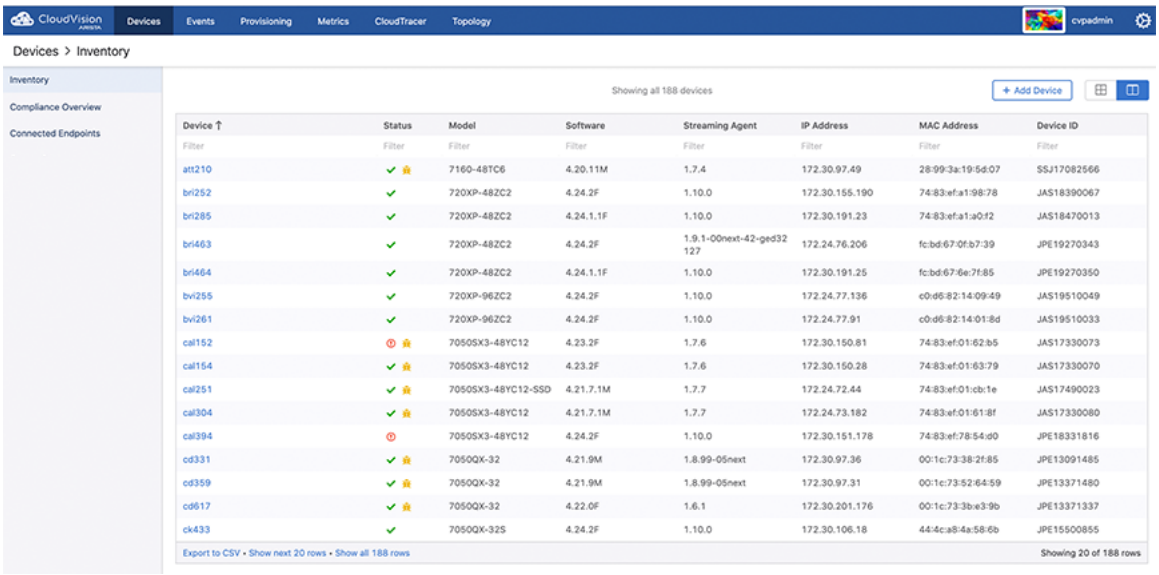

Figure 44: UI Access

When you access the UI for the first time, you need to apply the CVW service license.

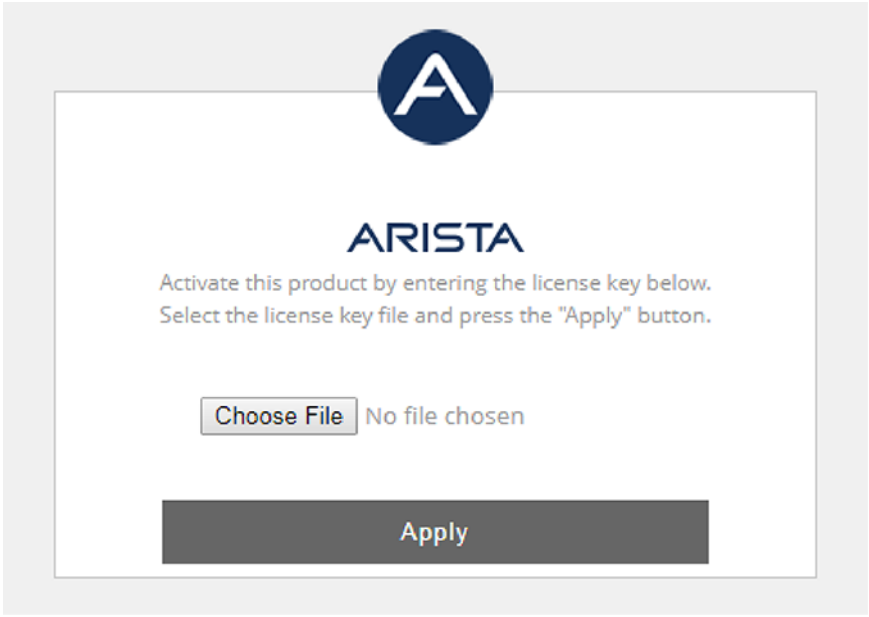

Figure 45: CVW Service License

Note:

- For the license file, please contact Arista Technical Support at http://support-wifi@arista.com.
- Use the `ifconfig` command on the CV root shell to get the eth0 MAC addresses of the primary and secondary CV servers (you need not access the wifimanager CLI for this). You need to include both these MAC addresses when you email support to request a license. One license is generated for the two (primary and secondary) MAC addresses.

Once you apply the license, you must log in to the CVW UI using the following default credentials:

Username: **admin**

Password: **admin**

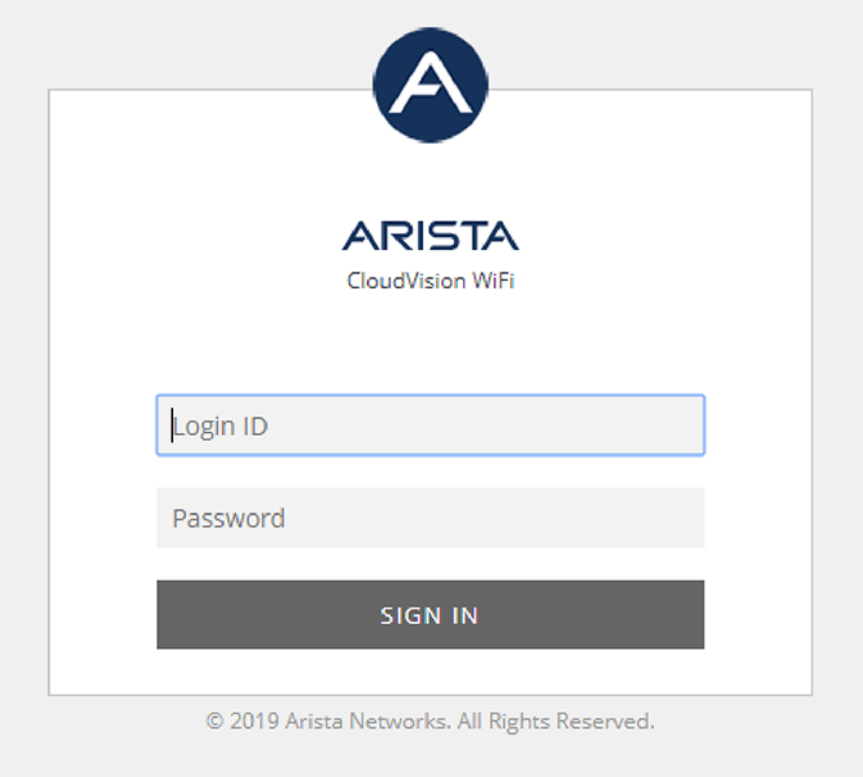

Figure 46: CVW Login Page

You can then change the password and add other users.

Note: You can now also connect Arista access points to the server.

6.5 Key CVW Operations and Directories

CVW is containerized as a service on CV. See the Wifimanager CLI Commands section for a list of wifimanager CLI commands and their descriptions.

For details on how to configure, monitor, and troubleshoot WiFi using CloudVision WiFi, see the CloudVision WiFi User Guide on the Arista WiFi Support Portal at https://support.wifi.arista.com/. You can access the portal from the WiFi - Support Portal tile on your dashboard. For details and credentials to access the portal, contact support-wifi@arista.com.

CVPI Commands for CVW

The following table lists the operations you can perform on wifimanager and the corresponding CVPI commands used.

Table 4: CVPI Commands

Operation	CVPI Command
start	cvpi start wifimanager
stop	cvpi stop wifimanager
status	cvpi status wifimanager
restart	cvpi restart wifimanager

Operation	CVPI Command
reset	cvpi reset wifimanager
backup	cvpi backup wifimanager
restore	cvpi restore wifimanager </path/to/backup/file>
debug	cvpi debug wifimanager

Note: The backup restore fails if the user running the restore command does not have access to the path where the backup file is stored.

The restart command restarts the wifimanager service, whereas the **reset** command resets wifimanager settings and data to factory default values. The **debug** command generates a debug bundle containing log files and configuration files that can be used to troubleshoot issues.

The following table lists the operations you can perform on aware and the corresponding CVPI commands used.

Table 5: Aware CVPI Commands

Operation	CVPI Command
start	cvpi start aware
stop	cvpi stop aware
status	cvpi status aware

6.5.1 Wifimanager Directories

CVW stores its data in docker volumes that reside under the **/data/wifimanager** directory on the CV. The following table lists the important wifimanager directories and the information they contain.

Table 6: Contents of wifimanager Directories

Directory on CV	Contains
/data/wifimanager/log/glog	Application logs
/data/wifimanager/data/conf	Configuration files
/data/wifimanager/data/data	System data files/directories
/data/wifimanager/data/instances	Customer data files/directories
/data/wifimanager/data/pgsql_data	Postgres data
/data/wifimanager/log/slog	System logs
/data/wifimanager/backup	On-demand backups

6.6 Wifimanager CLI Commands

The following table provides the list of wifimanager CLI commands and their descriptions.

Table 7: Wifimanager CLI Commands

Command	Description
db backup	Backs up the database to the specified remote server.
db clean	Cleans up resources without disrupting services.
db restore	Restores the database from a previous backup on a remote server.
db reset	Resets the database to factory defaults but maintains network settings.
get cert	Generates a self-signed certificate.
get openconfig mode	Displays current OpenConfig mode.
get cors	Displays the current status of CORS support.
get certreq	Generates a Certificate Signing Request.
get db backup info	Displays scheduled DB backup information.
get debug	Creates a debug information tarball file. This file can be used for debugging.
get debug verbose	Creates a basic debug information tarball.
get debug ondemand	Displays the debug information.
get device upgrade bundles	Displays information about device upgrade bundles available in the local repository.
get device repo config	Displays configuration (Mode and Hostnames) for repositories that store upgrade bundles and device capability information.
get idle timeout	Displays the current idle timeout value. A value of 0 indicates no timeout.
get integrity status	Checks the integrity of critical server components.
get ha	Displays High Availability (HA) Pair configuration and service status.
get lldp	Displays the LLDP configuration.
get remote logging	Displays the remote logging configuration.
get log config	Displays the logger configuration.
get log level gui	Displays log levels of GUI modules.
get log level aruba	Displays the log level of Aruba Mobility Controller Adapter module.
get log level wlc	Displays the log level of the Cisco WLC Adapter module.
get log level msmcontroller	Displays the log level of HP MSM Controller Integration.

Command	Description
get msmcontroller cert	Generates a self-signed certificate for HP Adapter.
get msmcontroller certreq	Generates a Certificate Signing Request for HP Adapter.
get access address	Shows access IP Address/Hostname of this server.
get server config	Displays complete server configuration.
get server cert	Uploads server certificate to a remote host.
get server check	Runs a server consistency check and displays results. If any fatal item fails, a failure result is recorded.
get server tag	Displays the custom tag set by the user.
get serverid	Displays the server ID.
get sensor debug logs	Uploads AP debug logs to the specified upload URL.
get sensor list	Displays the list of APs.
get sensor reset button	Displays the state of the AP's pinhole reset button.
get status	Displays the status of server processes.
get ssh	Displays the SSH server status.
get version	Displays the version and build of all the server components.
get packet capture	Captures packets on Public and HA/Management network interface(s).
set scan config	Modify AP background scanning parameters.
set openconfig mode	Enable/disable OpenConfig mode.
set cert	Installs a signed SSL certificate.
set cors	Enables or disables CORS support.
set dbserver	Starts/stops database server.
set db backup info	Sets scheduled DB backup information.
set device capability	Updates the device capability information.
set device upgrade bundles	Upload/delete device upgrade bundles in the local repository.
set device repo config	Sets configuration (Mode and Hostnames) for repositories that store upgrade bundles and device capability information.
set erase	Configures the backspace key.
set ha dead time	Changes the Dead Time of High Availability (HA) service.

Command	Description
set ha link timeout	Sets the timeout in seconds to signal Data Sync Link failure.
set idle timeout <*timeout-in-minutes*>	Sets the idle timeout for the command shell. A value of 0 disables the idle timeout.
set lldp	Sets LLDP configuration.
set remote logging	Sets remote logging configuration.
set log config	Sets the configuration of the logger.
set log level gui	Sets log levels of GUI modules.
set log level aruba	Sets the log level of Aruba Mobility Controller Adapter Module.
set log level wlc	Sets log level of Cisco WLC Adapter Module.
set log level msmcontroller	Sets log level of HP MSM Controller Integration.
set msmcontroller cert	Installs a signed SSL certificate for HP Adapter.
set loginid case sensitivity	Toggles login ID case sensitivity.
set server	Starts/stops application server.
set server discovery	Changes server discovery settings on given AP(s).
set server tag	Configure a custom tag for files generated by this server.
set access address	Sets access IP Address/Hostname of the server.
set serverid	Sets server ID.
set ssh	Starts/stops SSH access to the server.
set communication passphrase	Sets the communication passphrase used for AP-server authentication and to encrypt the communication between APs and the server.
set communication key	Sets the communication key used for AP-server authentication and to encrypt the communication between APs and the server.
set communication key default	Resets the communication key used for AP-server authentication and to encrypt the communication between APs and the server.
set sensor legacy authentication	This allows/disallows APs running on versions lower than 6.2 to connect to the server.
set sensor reset button	Sets the state of the AP's pinhole reset button (select AP models only).
set smart device oui	Add, remove MAC OUI's for specific smart device type IDs.
set webserver	Starts/stops web server.
set wlc mapper	Manage Cisco WLC Custom Mapper file.

Command	Description
exit	Exits the config shell session.
ping <*Hostname/IP Address*>	Ping a host.
reset locked gui	Unlocks Graphical User Interface (GUI) account for the "admin" user.
reset password gui	Sets Graphical User Interface (GUI) password for the "admin" user to factory default value.
upload db backup	Uploads successful DB backup(s) to an external server.
application signature update	Updates app visibility signature.

Chapter 7

General Customizations

CloudVision Portal (CVP) enables you to customize the grid columns of CVP graphical user interface (GUI) pages. You can customize the grid columns of all CVP GUI grids.

CVP also enables you to easily paginate (navigate) through the pages of the grids of the GUI. The pagination controls are available in all grids.

- Column Customization
- Pagination Controls

7.1 Column Customization

CloudVision Portal (CVP) enables you to customize the columns of the grids of CVP graphical user interface (GUI) pages. You can customize columns of any grid of the CVP GUI.

You use the **Columns Settings** dialog to customize the columns of the active grid. You can open the **Columns Settings** dialog by clicking the column customization icon, which is available of every page of the GUI.

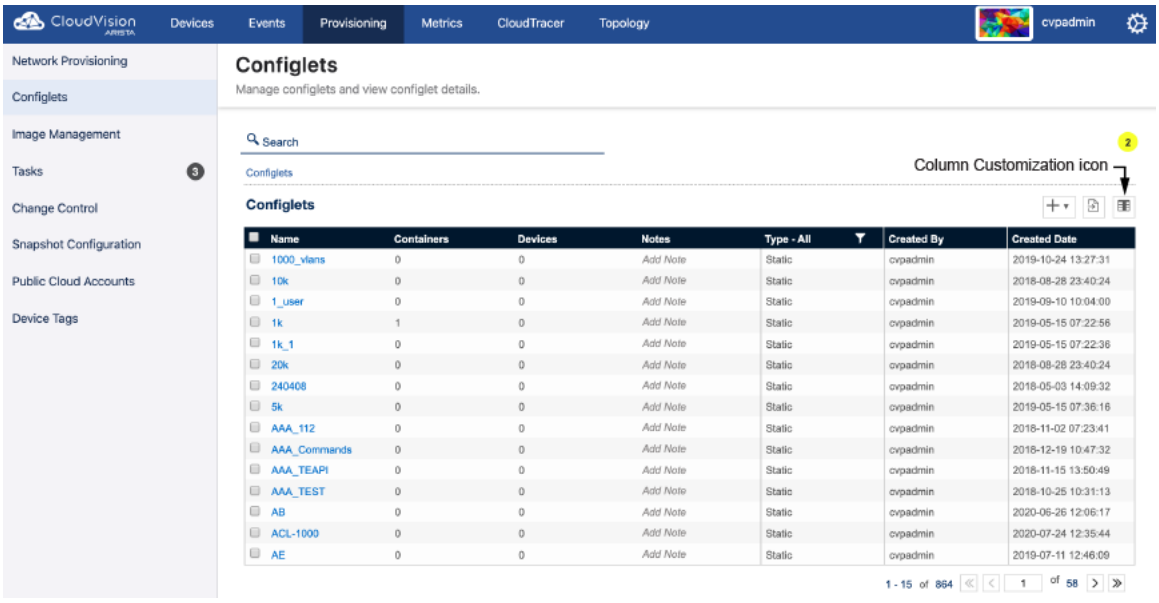

Figure 47: Configlet Management page

Complete these steps to customize grid columns.

1. Go to a page that has the grid you want to customize.
2. Click the column customization icon.

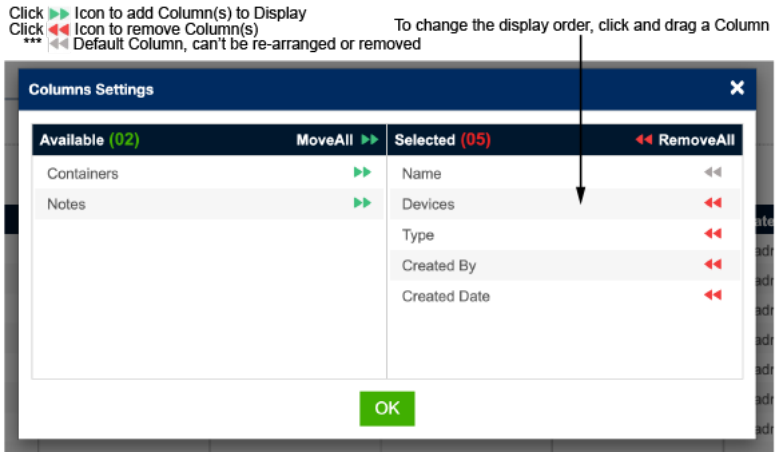

Figure 48: Column Settings dialog

3. Use the arrow icons to rearrange the columns of the grid as needed.
4. Once you are done rearranging the grid columns, click **OK** to save the changes.

7.2 Pagination Controls

The pagination controls you use to navigate through the pages of grids are available for each grid. The controls enable you to:

- Go to the previous page of the grid
- Go to the next page of the grid
- Go to the first page of the grid
- Go to the last page of the grid
- Go to directly to a specific page

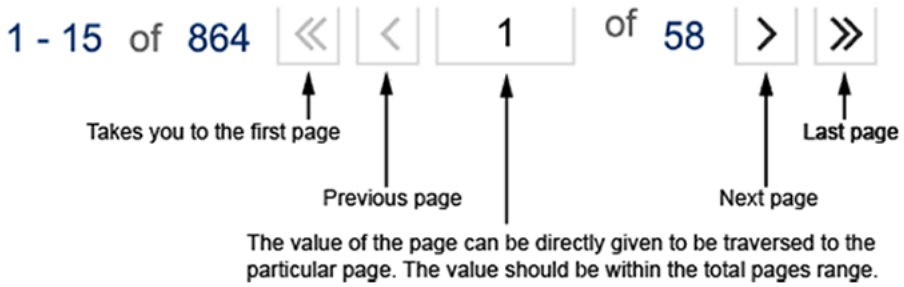

Figure 49: Pagination controls of the CVP GUI grids

Chapter 8

Device Management

CloudVision Portal (CVP) provides a powerful, event-driven, streaming analytics platform that enables you to monitor the state of all devices currently managed by CVP.

By configuring devices to stream device-state data to CVP, you can manage all of the devices in your current inventory of devices to gain valuable insights into the state of your devices, including real-time updates about changes in device state.

The device inventory is comprised of all devices that you have imported into CVP. After a device is imported into CVP, it can be configured and monitored using the various CVP modules.

8.1 Requirements

Make sure you review the software and hardware requirements for deploying and using the Telemetry platform before you begin deploying the platform.

System Requirements

Note: If you upgraded from a previous version of CVP, you must verify that all of the CVP node VMs on which you want to enable Telemetry have the required resources to use Telemetry. See *Resource Checks* for details on how to check CVP node VM resources and perform any modifications needed to increase the current CVP node VM resources.

8.2 Limitations

The following table lists the current limitations of the Telemetry platform. Review the limitations to ensure you do not inadvertently attempt configurations that exceed the limitations.

Table 8: CVP Telemetry Platform Limitations

Limitations

Maximum number of devices	This represents the total number of devices currently configured to stream Telemetry data.
Device-state data	Streaming of LANZ data is not enabled by default. You must enable it on devices.
Secret configuration	If "enable secret" is configured, the secret must be the same as the Cloudvision user's password.

8.3 Features

The list the current supported and unsupported Telemetry platform features are provided in the following topics:

- Supported Features
- Unsupported Features

8.3.1 Supported Features

The CVP Telemetry Supported Features table lists the supported features. Review the supported features to ensure you are aware of the features available to you to monitor devices using Telemetry data.

Table 9: CVP Telemetry Supported Features

	Supported Feature
Real-time monitoring of devices	The Telemetry platform provides interfaces for viewing real-time updates about changes in device state as well as events. You can also view trends in device-state metrics and queries of historical device-state data.
Instant state change updates	Changes in the state of a device are instantly streamed to CVP.
Full state change data	All changes in device-state are captured and streamed to CVP for viewing. Types of device-state include: • All SysDB state (except state under /Sysdb/cell/*). • All SMASH tables. • Process and kernel data (for example, CPU and memory usage). • System log messages
Analytics engine	The Telemetry platform provides a robust analytics engine that aggregates the streamed device-state data across devices, monitors device state, and generates events to indicate issues. It also normalizes data so it is easier for other applications to use.
Telemetry events	Device-state and system environment event types are streamed to CVP: • Informational (updates about changes in device state). • Warning (for example, unsupported EOS version on a device) • Errors (data discards or input errors on interfaces, and more). • Critical (system environment issues such as overheating).

High performance database	The Telemetry platform utilizes a high performance Hbase database to store device-state data, including events. Data is stored in compressed format without a loss of resolution. • The data storage capacity is approximately: • 60480 records worth of raw data per path • 1 week of 10 second aggregated data • 6 weeks of 60 second aggregated data • 30 weeks worth of 5 minute aggregated data
Disk space protection	To prevent telemetry data from consuming too much disk space in the CVP cluster, the Telemetry platform automatically blocks the ingest port for the entire cluster if disk usage exceeds **90%** on any node of the cluster. Once the ingest port is blocked, it remains blocked until disk usage drops below **80%** on all nodes in the cluster.
Data management	To ensure that the most relevant data is given priority, the Telemetry platform provides automated data management, including: • Maximum time limit on stored device-state data (1 month). • Current and the most recent device-state updates are always stored (given priority over older state updates). Periodic clean-up jobs are executed weekly (Saturday at 11:00 P.M.). Old device-state data is purged.
Command support	Several commands are provided for: • Checking status of the Telemetry components. • Enabling and disabling of Telemetry platform components. • Starting and stopping Telemetry components. • Viewing the debug log for Telemetry components. • Troubleshooting the Telemetry components, including checking to see that logs are being created for the component. • To display granular information on disk space usage of telemetry data and delete telemetry data selectively.

8.3.2 Unsupported Features

The CVP Telemetry Unsupported Features table lists the unsupported features. Review the limitations to ensure you do not inadvertently attempt to configure or use unsupported Telemetry features.

Table 10: CVP Telemetry Unsupported Features

	Unsupported Feature
Streamed device-state data	Flexroute is not supported.

8.4 Telemetry Platform Components

Arista's streaming Telemetry platform consists of a set of components, all of which are essential to the proper operation of the platform.

The components of the Telemetry platform are:

- NetDB State Streaming Component
- CloudVision Analytics Engine Component
- REST and Websocket based APIs are available to programatically get data from the CloudVision Analytics Engine. Contact your Arista Sales Engineer for more information.

8.4.1 NetDB State Streaming Component

The NetDB State Streaming component is an agent that runs on Arista switches. It is the Telemetry platform component that streams device-state data from devices to the CloudVision Analytics Engine, which is the back-end component of platform.

8.4.2 CloudVision Analytics Engine Component

The CloudVision Analytics Engine is the back-end component of the Telemetry platform. It is a set of processes that run on CVP. Collectively, the processes perform the following operations:

- Receives all of the device-state data streamed by the NetDB State Streaming component from devices that have been configured to stream device-state data.
- Runs automated data analysis on the device-state data received from the NetDB State Streaming component. The analytics processes aggregate the device-state data across devices, monitor device state, and generate events if something goes wrong. The processes also normalize data so it is easier for other applications to use.
- Stores all of the streamed device-state data received from the NetDB State Streaming component, and then makes the stored data available in CloudVision.
- Provides CloudVision Analytics Engine Viewer, which is referred to as the Aeris Browser. You use it to directly view device-state data received from devices that have been configured to stream device-state data. The Aeris Browser enables you to view raw device-state data.
- REST and Websocket based APIs are available to programatically get data from the CloudVision Analytics Engine. Contact your Arista Sales Engineer for more information.

8.5 Supplementary Services: Splunk

For more information on the requirements for CVP to manage Splunk extensions on EOS devices, go to https://www.arista.com/en/support/software-download and download the PDF from **Extensions > Splunk > AristaTelemetry.pdf**.

Related topics:

- Requirement
- Installation
- Quick Start

8.5.1 Requirement

EOS 4.15.2 or later is required.

8.5.2 Installation

You can access the Splunk Telemetry App directly from CVP by completing the following steps. From your browser.

1. Copy the RPM to and install it on the switch.

```
show extensions
```

```
Name Version/Release Status RPMs
```

2. Install the Splunk Universal Forwarder RPM on EOS.

```
copy <source>/splunkforwarder-6.1.4-233537.i386.rpm extension:
extension splunkforwarder-6.1.4-233537.i386.rpm
```

3. Install the AristaAppForSplunk on EOS.

```
copy <source>/AristaAppForSplunk-1.3.2.swix extension:
extension AristaAppForSplunk-1.3.2.swix
```

Note: Extensions must be installed on all supervisors.

Restart the SuperServer agent.

```
(config)# agent SuperServer shutdown
(config-mgmt-api-http-cmds)# no agent SuperServer shutdown
```

4. Verify the extensions are loaded.

```
show extensions
Name Version/Release Status RPMs
---------------------------------------- ------------------------
 ------ ----
AristaAppForSplunk-<version>.swix <version>/1.fc14 A, I 3
splunkforwarder-6.1.4-233537.i386.rpm 6.1.4/233537 A, I 1
EosSdk-1.7.0-4.15.2F.i686.rpm 1.7.0/2692966.gaevanseoss A, I 1
A: available | NA: not available | I: installed | NI: not installed |
 F: f
```

8.5.3 Quick Start

1. Use the configuration to enable forwarding to the Splunk indexer. This assumes that a username/password and eAPI have been configured for the AristaAppForSplunk extension previously.

```
daemon SplunkForwarder
 exec /usr/bin/SplunkAgent
 no shutdown
```

2. Configure and turn on the desired indexes for data collection. The credentials must match 'username <name> secret <passphrase>' configured on the switch.

```
option eapi_username value <username>
 option eapi_password value 7 <encrypted-password>
 option eapi_protocol value https
```

3. Turn on desired indexes for data collection.

```
option index-inventory value on
 option index-interface-counters value on
 option index-lanz value on
 option index-topology value on
 option index-syslog value on
 option index-data value <index-name
```

4. Configure Splunk server IP and destination port.

```
option splunk-server value <Server-IP:Port>
```

5. Start Splunk data forwarding.

```
option shutdown value off
```

8.6 Architecture

Telemetry Platform Architecture shows the architecture of the Telemetry platform, including all of the platform components and the data path of the streamed device-state data.

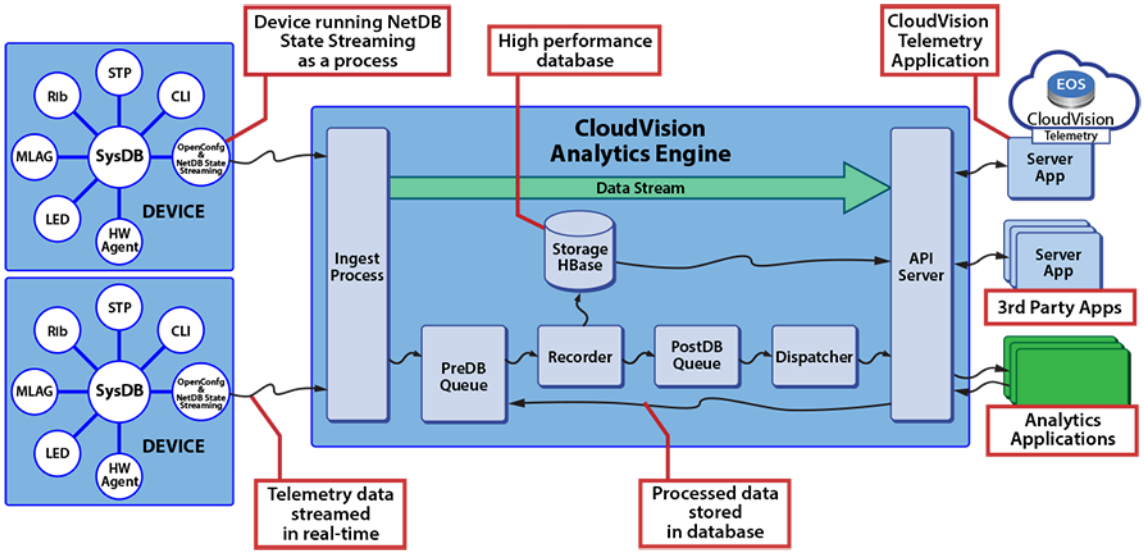

Figure 50: Telemetry Platform Architecture

8.7 Accessing the Telemetry Browser Screen

You can access the CloudVision Telemetry Browser screen directly from CVP by completing the following steps. Open your browser.

1. Point your browser to the CVP IP address or hostname.
2. Login to CVP.

   The CVP Home screen appears.

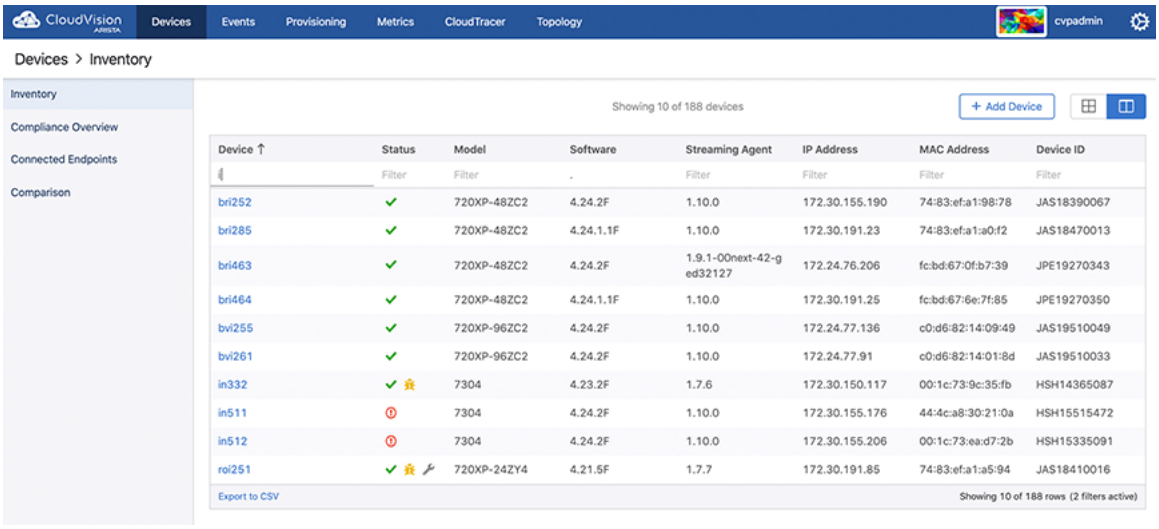

Figure 51: CVP Home Screen

3. Click the gear icon at the upper right corner of the screen.

Figure 52: Gear Icon

4. Click Telemetry Browser in the left pane.

   The system opens the Telemetry Browser screen that allows exploring the raw data stored in CVP telemetry.

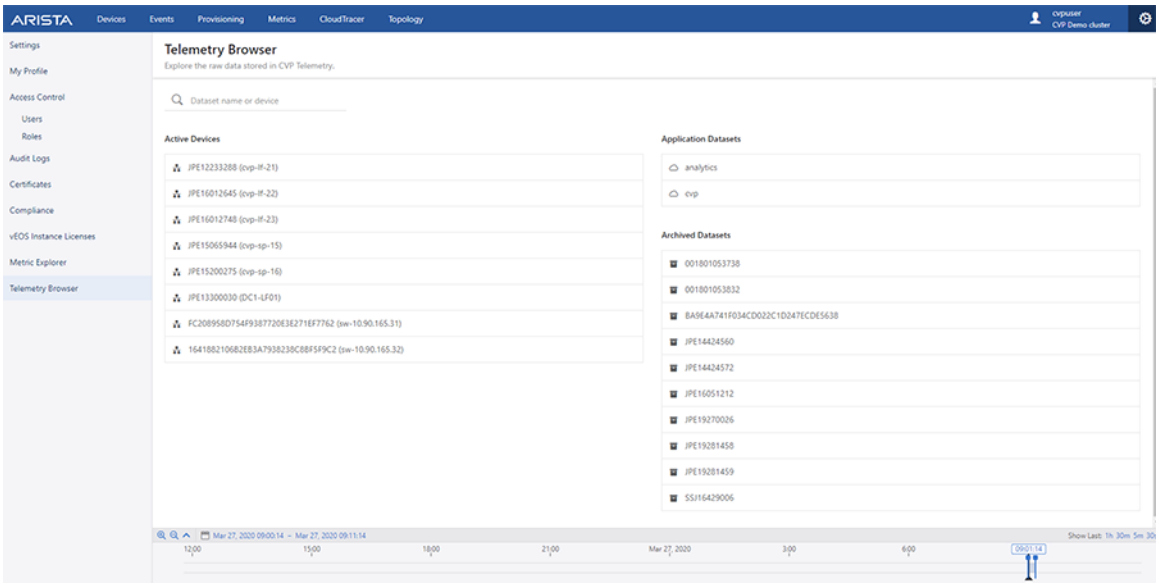

Figure 53: CloudVision Telemetry Browser Screen

8.8 Viewing Devices

You can quickly view information about devices that are currently configured to stream device-state data to CVP. Starting with *2018.2.0*, the inventory management screen is available under Devices in the CVP user interface.

Related topics:

- Tiles View
- Tabular View

8.8.1 Tiles View

The tiles view allows search by device hostname, serial number, or EOS version. The screen updates to show all of the devices currently configured to stream device-state data to CVP. For each device, the name and the version of the EOS image are shown on the Devices screen.

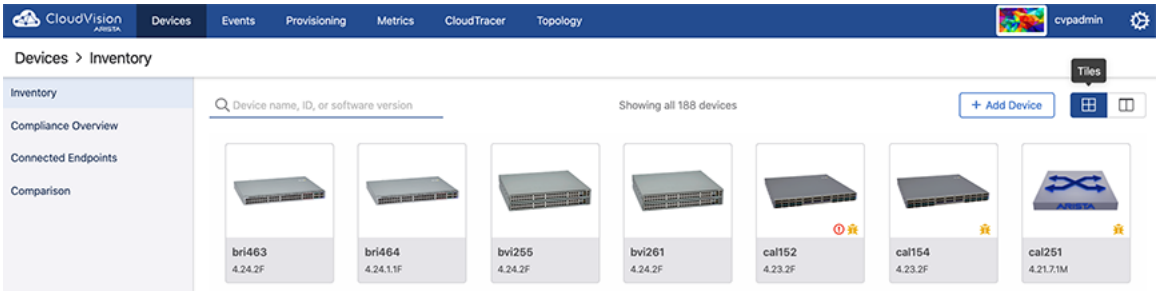

Figure 54: Viewing Devices (View Showing all Devices)

8.8.2 Tabular View

The tabular view lists device status, model, software, TerminAttr agent, IP address, MAC address, and serial number. You can search for devices based on device hostname, serial number, or EOS version.

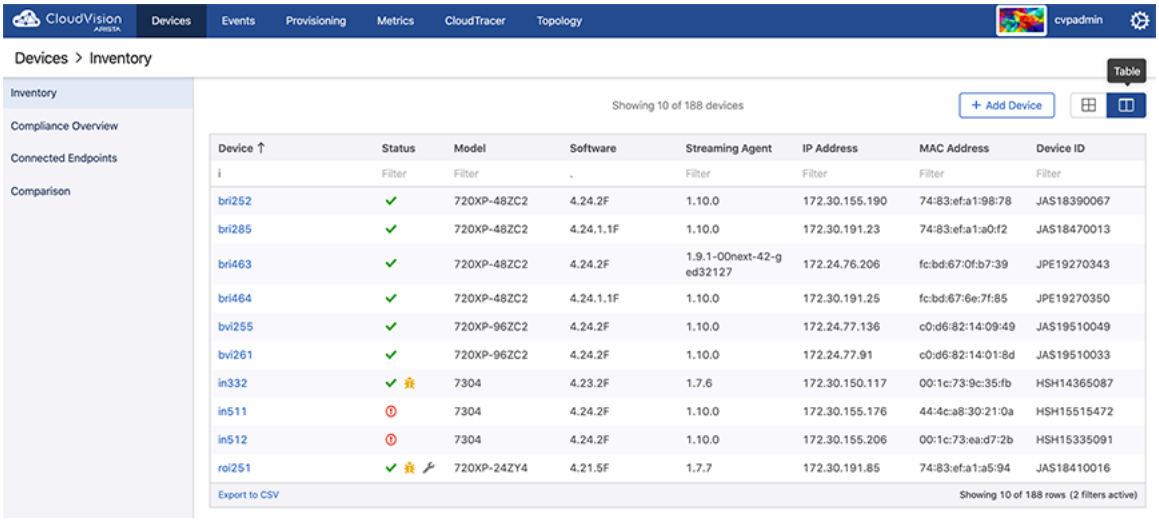

Figure 55: Device Inventory

8.9 Viewing Device Details

From the Inventory screen, you can quickly drill down to view details about a particular device by clicking the device icon. In tabular view, click the device name to view the corresponding device details.

The screen refreshes to show the device-state data streamed from the device to CVP.

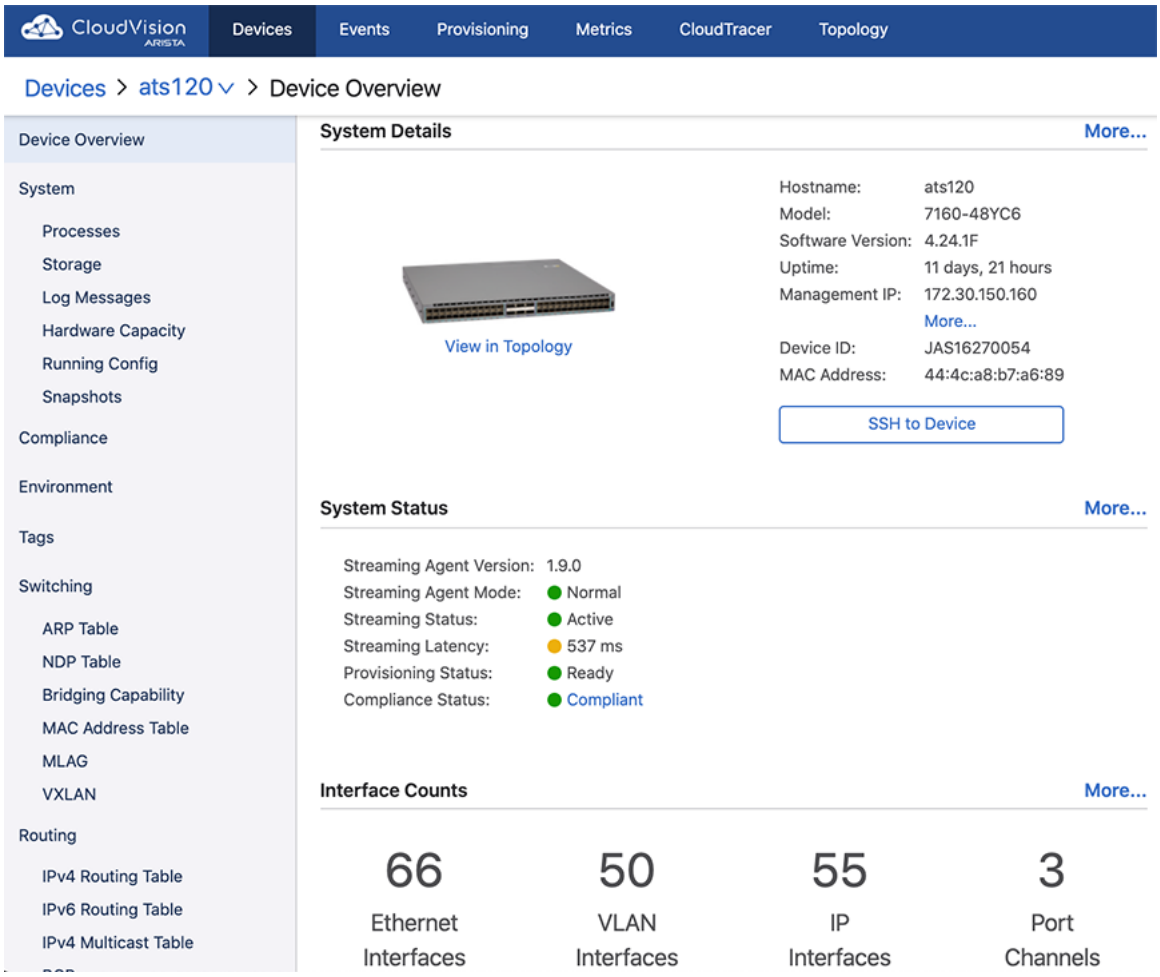

Figure 56: Viewing Devices Details (Single Device)

Device details include the information on overview, system, compliance, environment, switching, routing, and interfaces.

Related topics:

- Device Overview
- System Information
- Compliance
- Environment Details
- Switching Information
- Routing Information
- Status of Interfaces

8.9.1 Device Overview

The Device Overview section provides an overview of system details, telemetry status, and interface counts. Click **More** to reach corresponding sections for detailed information.

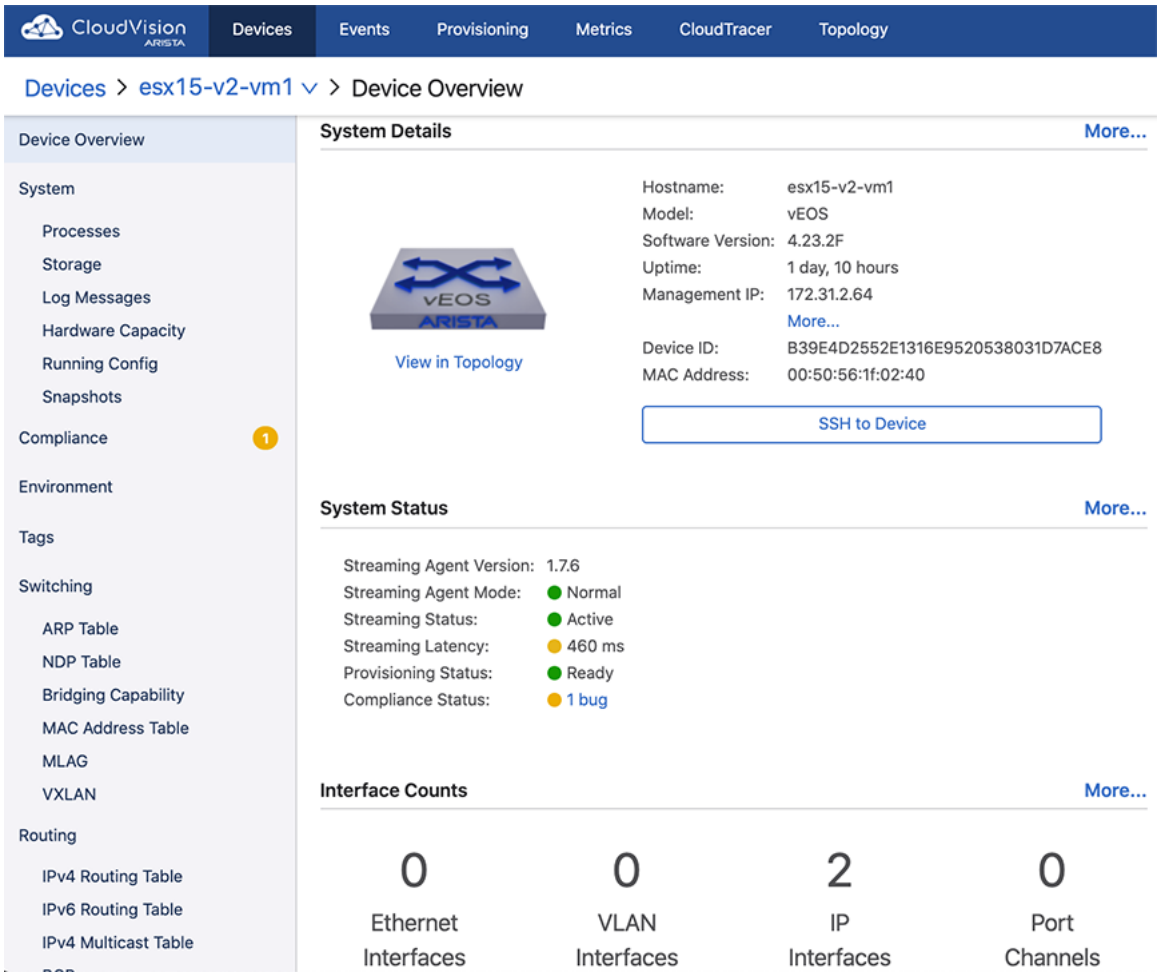

Figure 57: Device Overview Section

The Historical Comparison sub-section provides the information on EOS version, 5-minute CPU load average, MLAG status, IPv4 attached routes, IPV4 learned routes, configured BGP, IPv6 attached routes, IPV6 learned routes, and MAC addresses learned.

The system displays only Device Overview and System information for third-party devices.

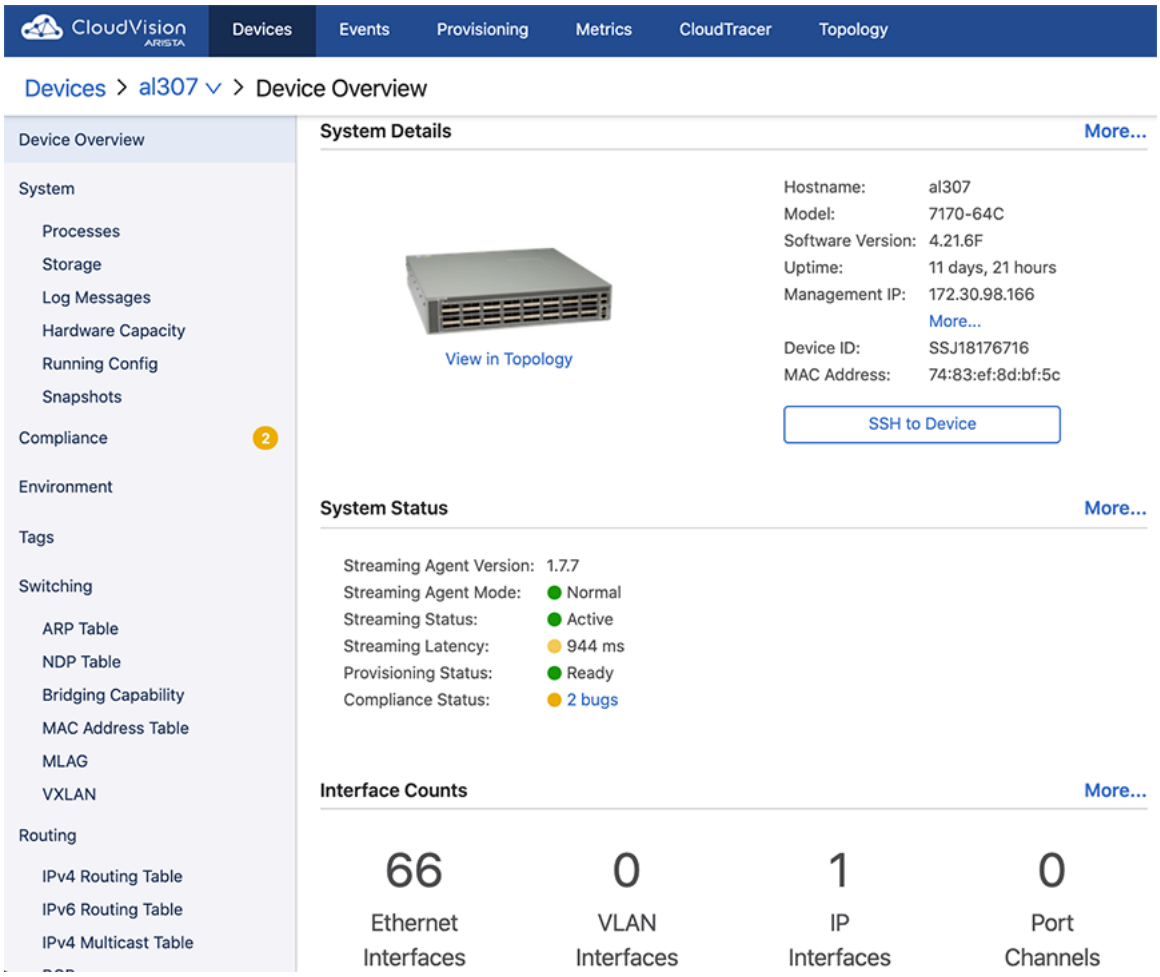

Figure 58: Third-Party Device Overview

8.9.2 System Information

The System section provides an overview of device details, telemetry status, and PTP status.

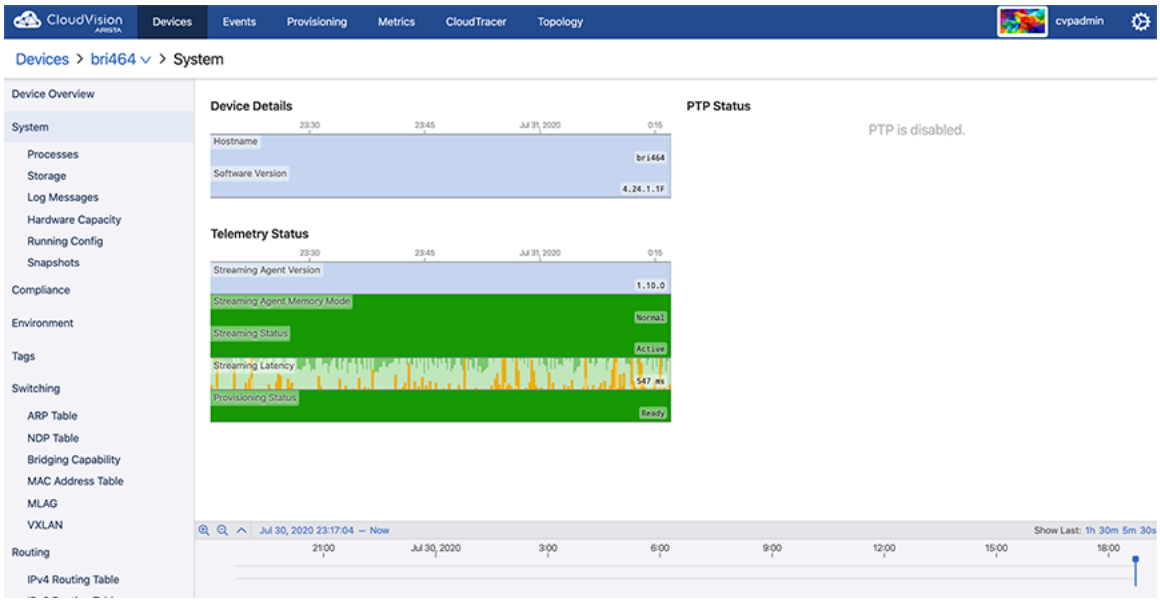

Figure 59: System Section

Sub-sections provide information on processes, storage, log messages, hardware capacity, running config, and snapshots.

8.9.3 Compliance

The Compliance section provides information on vulnerability to known bugs.

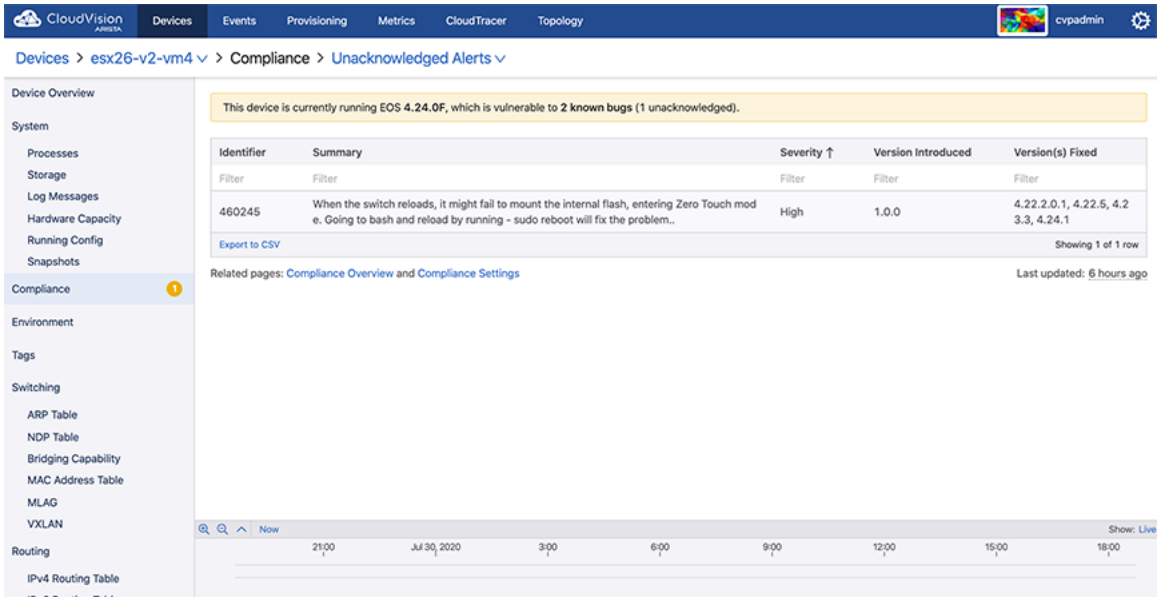

Figure 60: Compliance Section

8.9.4 Environment Details

The Environment section provides statistics on temperature, fan speeds, and output power.

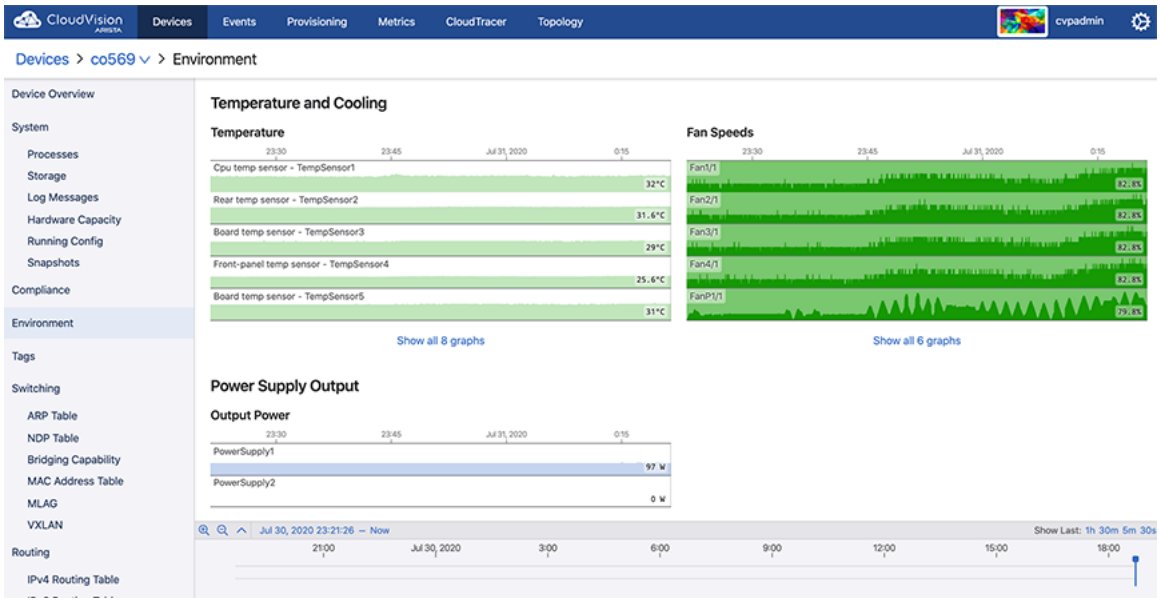

Figure 61: Environment Section

8.9.5 Switching Information

The Switching section provides the count of VLANs in which MAC address learning is enabled, count of total VLANs, count of configured VLANs, and detailed information on configured VLANs.

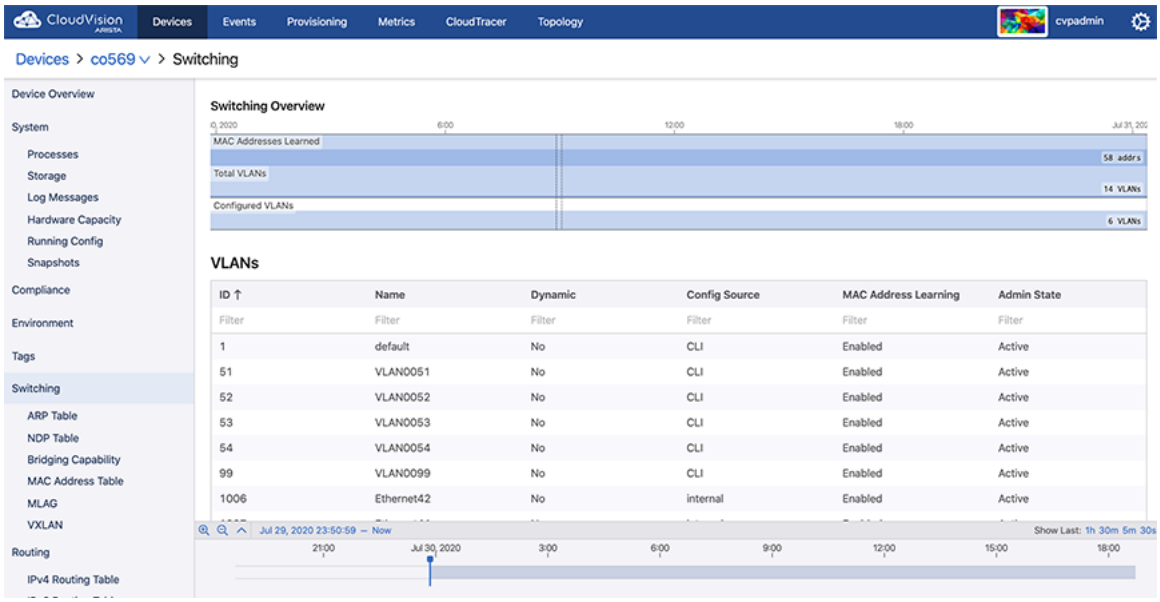

Figure 62: Switching Section

Sub-sections provide switching data like ARP table, NDP table, bridging capability, MAC address table, MLAG, and VXLAN.

8.9.6 Routing Information

The Routing section provides statistics on IPV4 route count by type, IPv6 route count by type, and routing statistics by VRF.

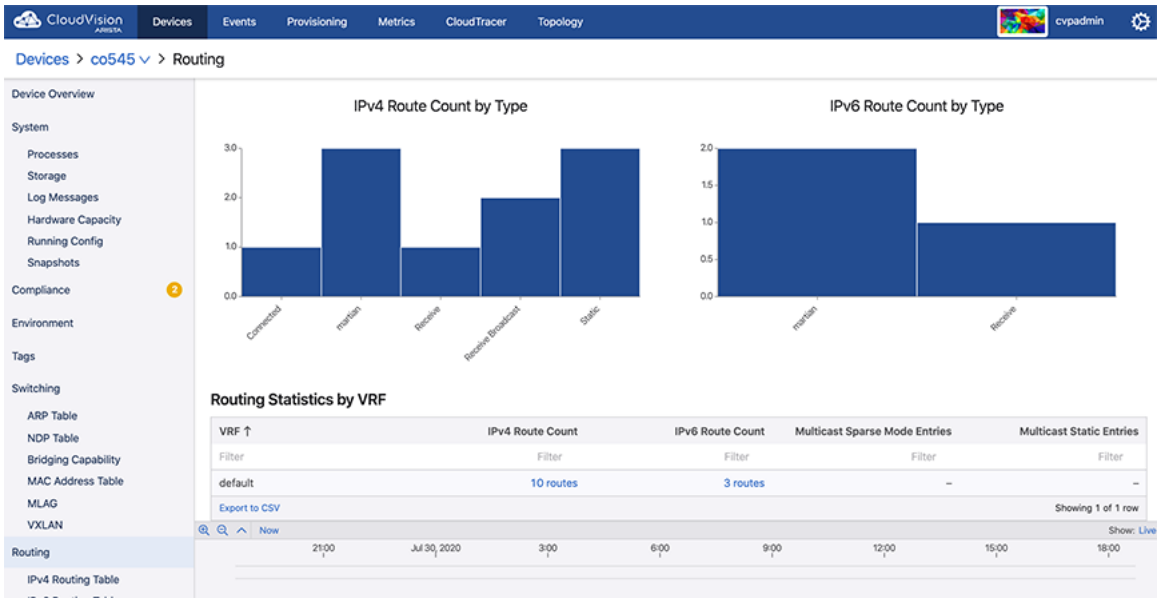

Figure 63: Routing Section

Sub-sections provide routing data like IPv4 and IPv6 routing tables, routing table changes, multicast data like sparse mode PIM and static, and BGP information.

8.9.7 Viewing Traffic Flows

CloudVision's traffic flows analyze the network traffic routed through a device. You can drill down into the details of network flow activity using bar charts, stacked time series graphs, heat-maps, and tables of usage statistics. See Accessing the Traffic Flows Screen.

To view the data on traffic flows, you must enable traffic flow tracking in devices to get data. See Enabling Traffic Flow Tracking.

8.9.7.1 Enabling Traffic Flow Tracking

Enabling flow tracking on a device allows CloudVision to provide a detailed breakdown of the forwarded network traffic. Traffic flow tracking is enabled through either of the following methods:

- Enable sFlow Sampling on a Device
- Enable Hardware Based IPFIX Flow Tracking

Enable sFlow Sampling on a Device

Arista switches provide a single sFlow agent instance that samples ingress traffic from all Ethernet and port channel interfaces.

Run the following commands to enable sFlow sampling on a device:

```
switch(config)#sflow sample <sampling rate>
switch(config)#sflow polling-interval <polling interval>
switch(config)#sflow destination 127.0.0.1
switch(config)#sflow source-interface <source interface>
switch(config)#sflow run
```

Note: The device must have a TerminAttr, version 1.6.0 or higher to stream out the sampled flow data.

sFlow monitors a random sample of packets at the configured sampling rate. Reported bandwidth and packet measurements are scaled up using the sampling rate to provide estimates of actual bandwidth usage and packet counts.

Enable Hardware Based IPFIX Flow Tracking

Arista switches also allow exporting flow information using the IPFIX format. This device supports hardware based IPFIX flow tracking.

Run the following commands to enable hardware based IPFIX flow tracking:

```
switch(config)#flow tracking hardware
switch(config)#!
switch(config)#tracker <tracker name>
switch(config)#record export on inactive timeout <inactive timeout>
switch(config)#record export on interval <interval>
switch(config)#record format ipfix standard timestamps counters
switch(config)#!
switch(config)#exporter <exporter name>
switch(config)#collector <loopback interface ip>
switch(config)#local interface <loopback interface>
switch(config)#template interval <interval>
switch(config)#no shutdown
switch(config)#exit
switch(config)#interface <interface>
switch(config)#flow tracker hardware <tracker name>
switch(config)#no shutdown
```

8.9.7.2 Accessing the Traffic Flows Screen

On the CloudVision portal, navigate to **Devices** > *Device_Name* > **Traffic Flows** to view the Traffic Flows screen. See the figure below.

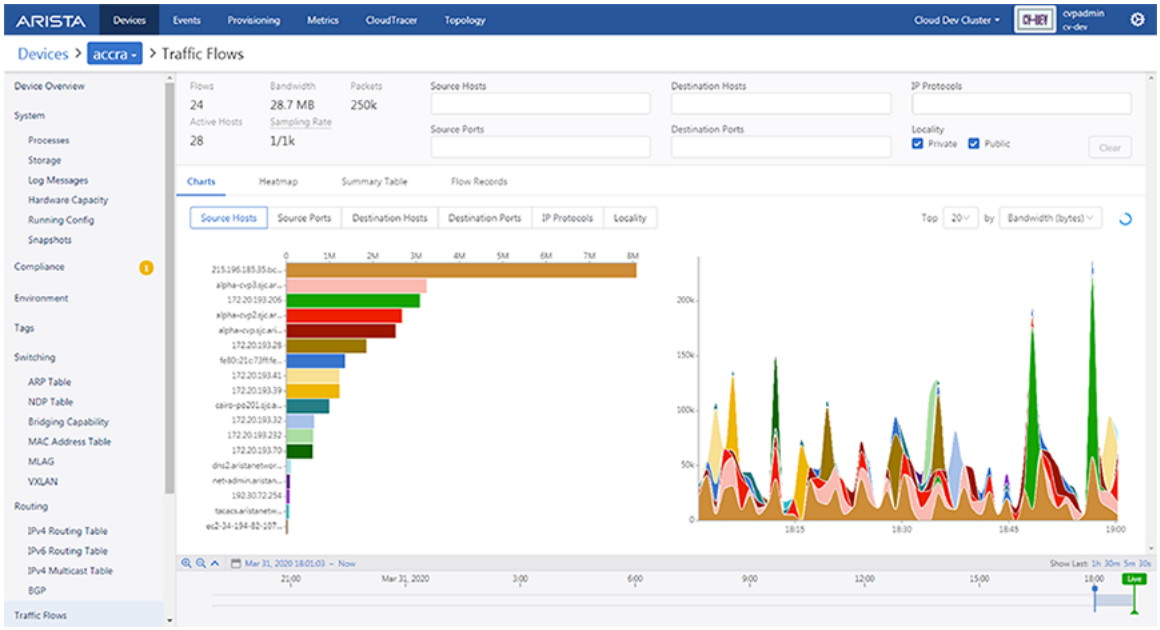

Figure 64: Traffic Flows Screen

This screen displays the summary of flows, bandwidth, packets, active hosts, and sampling rate. Provide the following details to view custom information of traffic flows:

- **Source Hosts** autocomplete field - Hostnames, IP addresses, or subnets in CIDR notation of the source host
- **Destination Hosts** autocomplete field - Hostnames, IP addresses, or subnets in CIDR notation of the destination host
- **Source Ports** autocomplete field - Port numbers or service names of the source port

- **Destination Ports** autocomplete field - Port numbers or service names of the destination port
- **IP Protocols** autocomplete field - IP protocols
- **Locality** - Select **Public** and **Private** checkboxes to view traffic flows of corresponding networks
- Clear button - Clears all specified filters
- Topology icon - Click to view the Topology Flows screen.
- Display options - Select any of the following display types:
  - Charts
  - Heatmap
  - Summary Table
  - Flow Records

Charts View

The **Charts** display option presents the summary of traffic flows in charts. See the figure below.

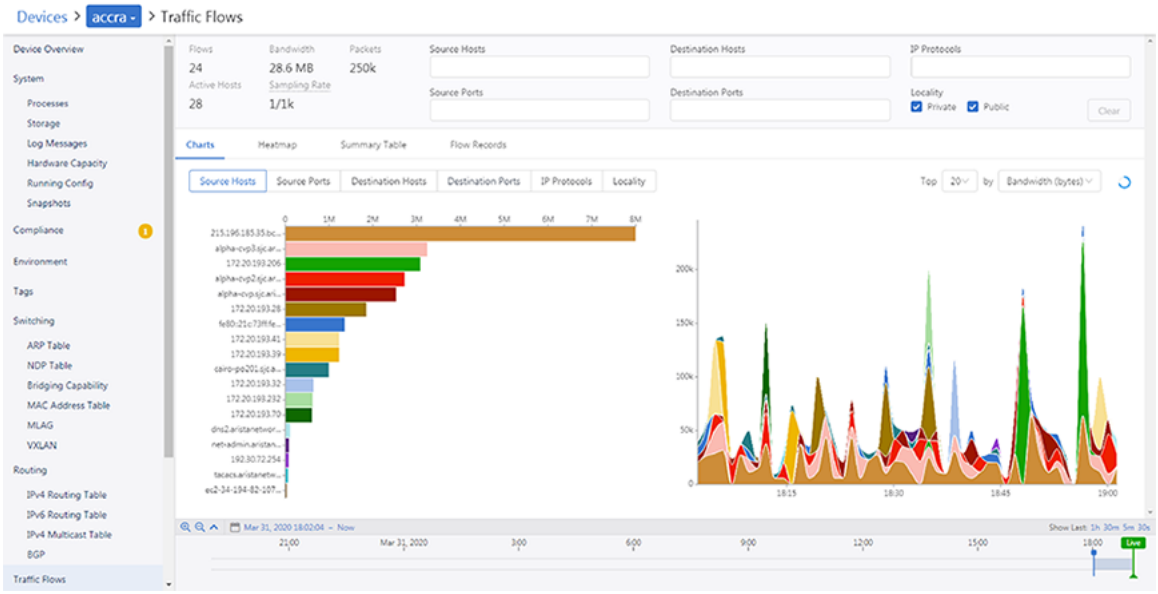

Figure 65: Traffic Flow Summary in Charts

The traffic flow data is displayed based on selected breakdown. Options include:

- Source Hosts
- Source Ports
- Destination Hosts
- Destination Ports
- IP Protocols
- Locality
- As per your selection from the top n drop-down menu, the top n items are displayed for each break down.
- Sort By drop-down menu - Select the required method to measure traffic. Options include:
  - Bandwidth (bytes)
  - Packets
  - Flow Count
- Refresh icon - Provides countdown in seconds to refresh the traffic flow data.

  Note: The data in live mode gets updated every 30 seconds.
- The following information is provided for each break down:

- Bar charts that display the total usage over the time period for items
- Stacked time series graphs that provide the following information:
  - The rate of usage vs. time

    Note: This information is provided only when the Sort By option is either Bandwidth (bytes) or Packets.
  - The number of flows active vs. time

    Note: This information is provided only when the Sort By option is Flow Count.
- Clicking on a bar in the bar chart or a time series in the stacked graph sets the clicked-on item as a filter wherever it is possible. For example, hosts or ports of source and destination.

Heatmap View

The **Heatmap** display option presents the summary of traffic flows in a heatmap. See the figure below.

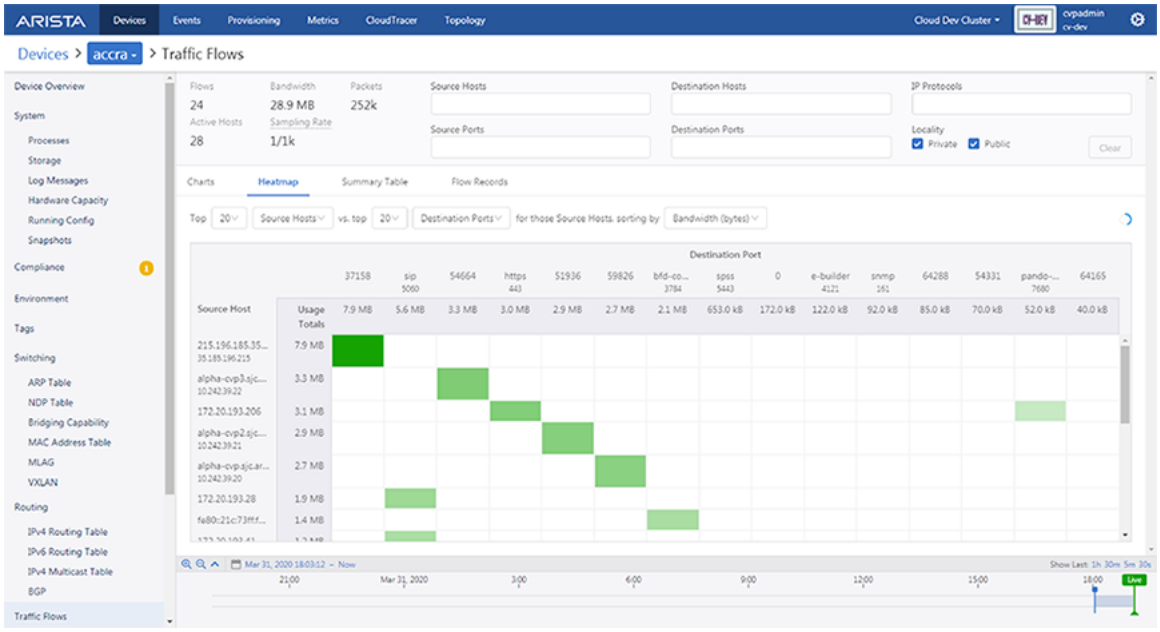

Figure 66: Traffic Flow Summary in Heatmap

The heatmap plots two breakdowns against each other. Options include:

- Top-n dropdown menu - As per your selection from the top n drop-down menu, the top n items are displayed for each break down.

  Note: The system provides multiple options under the top n drop-down.
- Source hosts
- Source ports
- Destination hosts
- Destination ports
- IP protocols
- Localities
- Sort By drop-down menu - Select the required method to measure traffic. Options include:
  - Bandwidth (bytes)
  - Packets
  - Flow Count

For example, the user selects top 20 source hosts vs. top 20 destination hosts. The system displays the top 20 destination hosts that communicated with any of those top 20 source hosts.

Each pairing of source host and destination host is shown as a cell in the grid. Cells are displayed in various shades of green based on their usage. The higher the usage, the darker the green shade.

Note: The system displays an empty cell if there is no usage.

Summary Table View

The **Summary Table** display option presents the summary of traffic flows in a table. See the figure below.

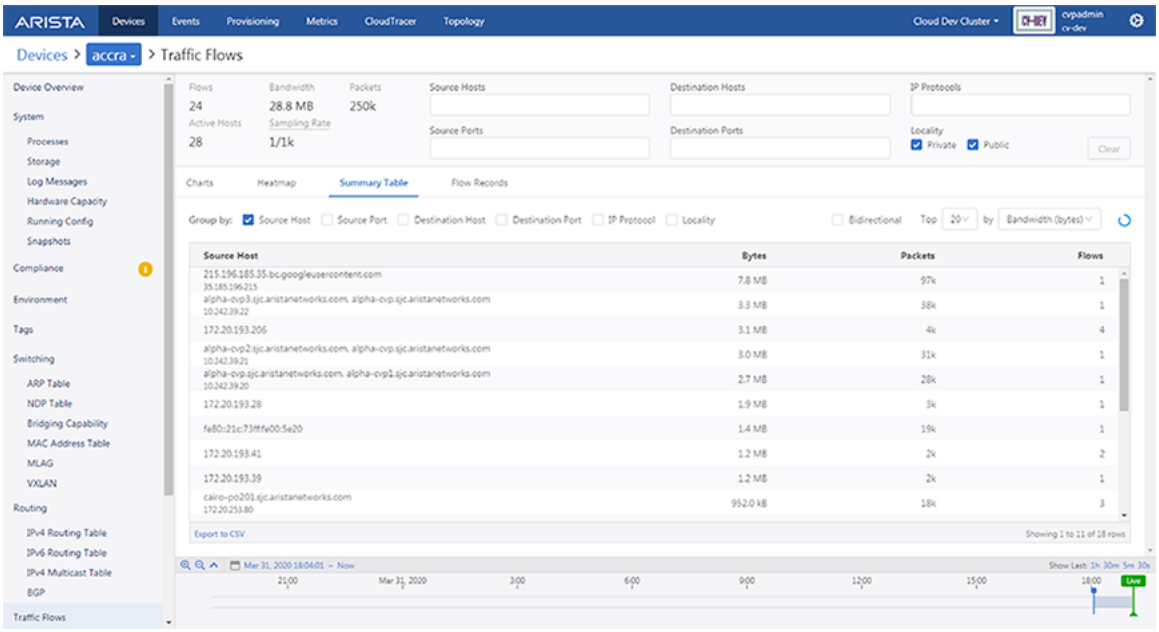

Figure 67: Traffic Flow Summary in Table

The traffic flow data is grouped based on selected segregation. Options include:

- Source host
- Source port
- Destination host
- Destination port
- IP protocol
- Locality
- Bidirectional
- Top-n dropdown menu - The top *n* items are displayed as per your selection from the top *n* drop-down menu.
- Sort By drop-down menu - Select the required method to measure traffic. Options include:
  - Bandwidth (bytes)
  - Packets
  - Flow Count

If multiple options are selected in the **Group By** field, the table displays a summary of usage statistics that is broken down according to the selected criteria. The summary can be sorted by bytes, packets, or flows in descending order.

Flow Records View

The **Flow Records** display option presents the record of all traffic flows in a tabular format. See the figure below.

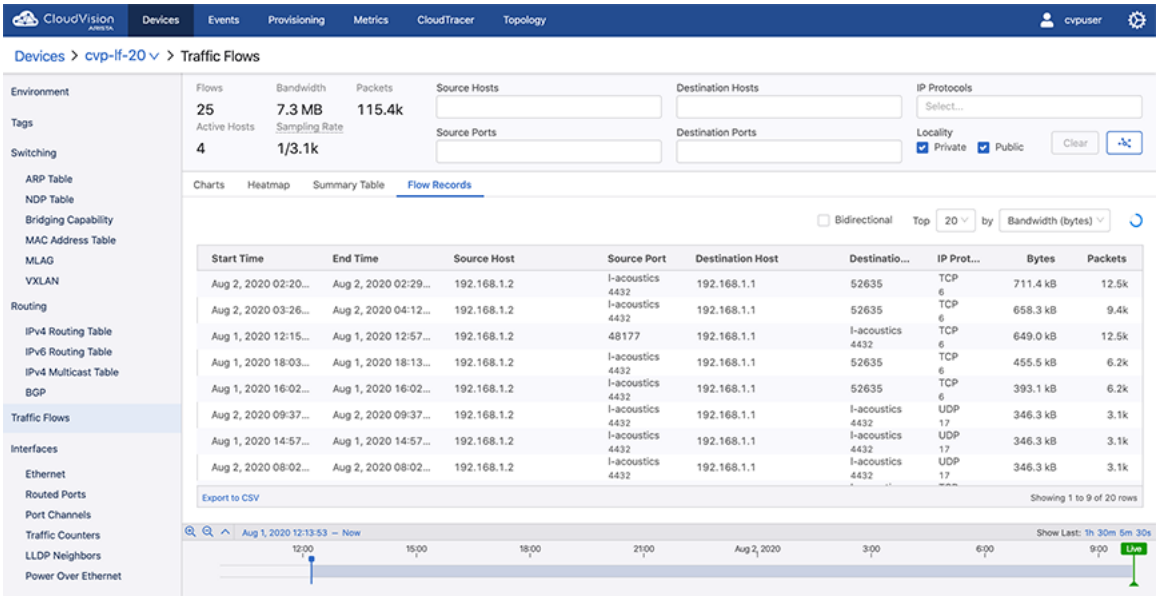

Figure 68: Traffic Flow Record

The traffic flow data is grouped based on selected segregation. Options include:

- Bidirectional
- Top-n dropdown menu - The top *n* items are displayed as per your selection from the top *n* drop-down menu.
- Sort By drop-down menu - Select the required method to measure traffic. Options include:
  - Bandwidth (bytes)
  - Packets
  - Newest

8.9.8 Status of Interfaces

The Interfaces section provides status of Ethernet interfaces, VLAN interfaces, IP interfaces, and port channels.

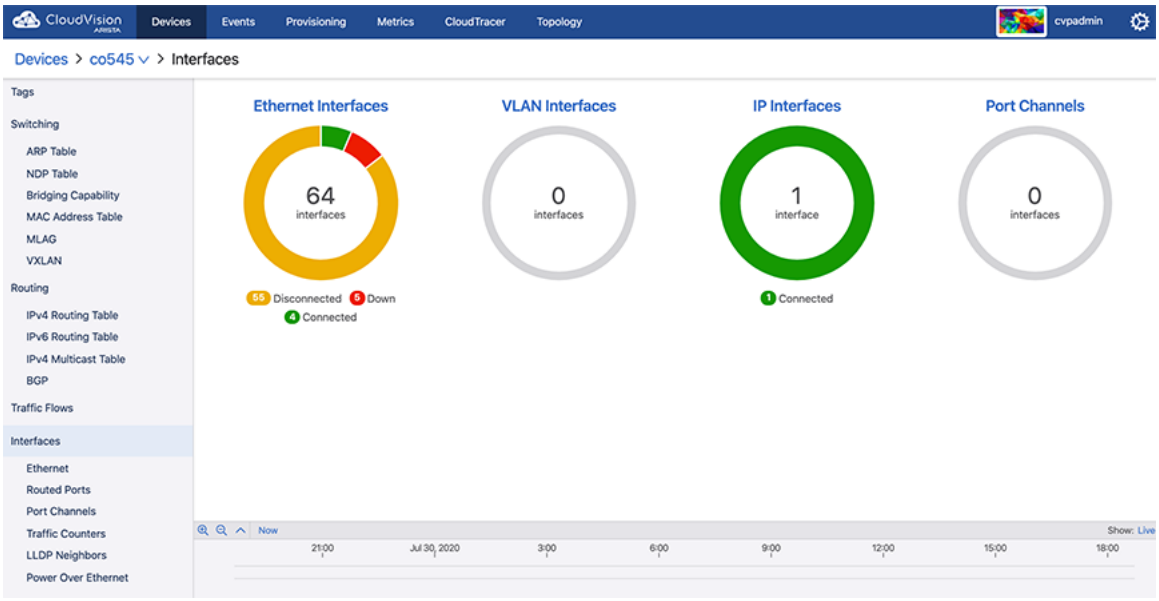

Figure 69: Interfaces Section

Sub-sections provide detailed information on Ethernet interfaces, routed ports, port channels, traffic counters, LLDP neighbors, and Power Over Ethernet.

8.9.8.1 Power Over Ethernet

Power Over Ethernet (PoE) is a technology for delivering electrical power along with network data over physical Ethernet connections. Some benefits of PoE are provided below:

- Reduces the need of extension cables and additional outlets
- Provides a reliable power source on difficult terrain
- Prevents data transmission hiccups
- Substantial reductions in space usage, cost, and time

In CloudVision, the Power Over Ethernet screen provides a summary of all interfaces along with information on each interface.

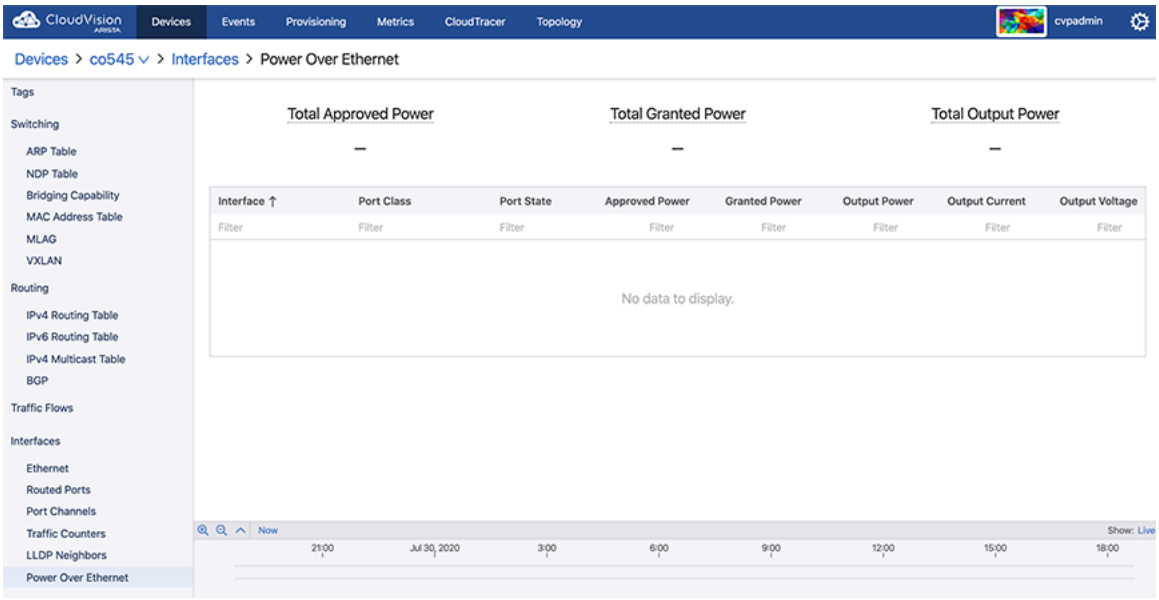

Figure 70: Power Over Ethernet Screen

The Power Over Ethernet screen displays the following information:

- Summary of All Interfaces
  - Total Approved Power - Sum of the approved maximum power amounts configured for each Ethernet port
  - Total Granted Power - Sum of the approved power amounts minus power loss to transmission over Ethernet cables
  - Total Output Power - Sum of actual power amounts delivered to each powered Ethernet device
- Information on Individual Interfaces
  - Interface - Interface name
  - Port Class - Maximum power in watts (W)
  - Port State - Operational status of a PoE device connected to the port
  - Approved Power - Configured maximum power output in watts (W) for the interface
  - Granted Power - Maximum power available to the device
  - Output Power - Power drawn by the device
  - Output Current - Current available on the PoE link in milliamps (mA)
  - Output Voltage - Voltage available over the PoE link in volts (V)

Note: PoE metrics are also available in the Metrics Explorer and can be built into custom metrics dashboards. Data on individual interfaces is available under the Interfaces metric type. Aggregate data totals of each device are available under the Devices metric type. See Accessing Metrics .

8.10 Viewing Connected Endpoints

Connected Endpoints are identified by DHCP collector. By default, the DHCP collector is enabled in TerminAttr. You must enable it on VLANs where you would like to identify connected endpoints. See Enabling DHCP Collector.

Once it is enabled, the Connected Endpoints summary screen provides information on all connected endpoints. See Accessing the Connected Endpoints Summary Screen.

Enabling DHCP Collector

As of TerminAttr v.1.6.0, the ECO DHCP Collector is enabled by default and listens on 127.0.0.1:67 for UDP traffic. Add 127.0.0.1 as an IP helper address on VLANs to capture device identification.

```
switch(config)# interface vlan100
switch(config-if-Vl100)# ip helper-address dhcp_server_address
switch(config-if-Vl100)# ip helper-address 127.0.0.1
switch(config-if-Vl100)# exit
switch(config)# ip dhcp snooping
switch(config)# ip dhcp snooping information option
switch(config)# ip dhcp snooping vlan 100
```

Accessing the Connected Endpoints Summary Screen

On the CloudVision portal, navigate to **Devices** > **Connected Endpoints** to view the Connected Endpoints Summary screen.This screen provides the classified summary of all endpoints along with the detailed information of each endpoint. See the figure below.

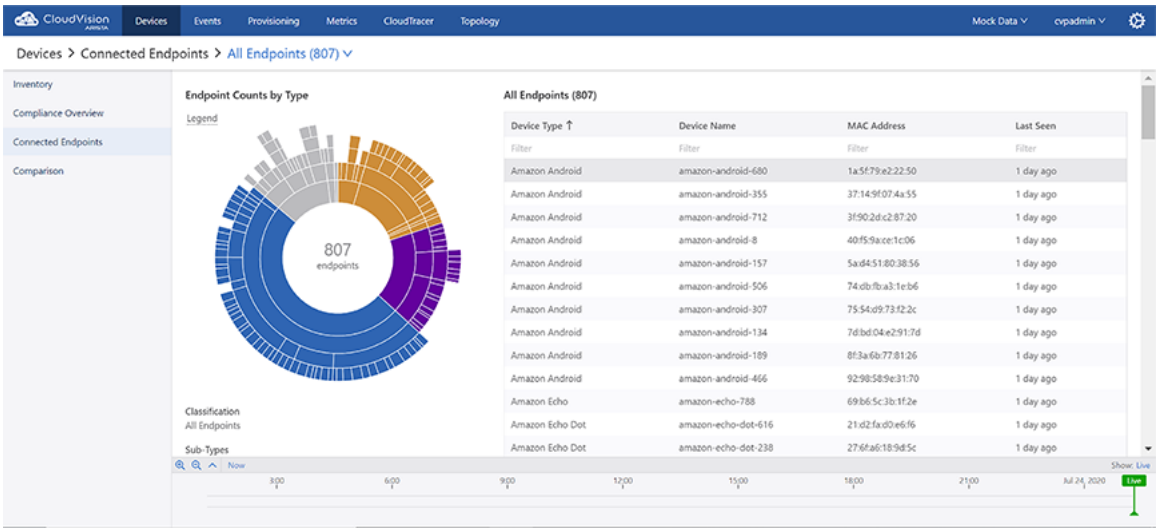

Figure 71: Connected Endpoints Summary Screen

Note: To reset to all endpoints, click the refresh icon (next to selected endpoint in breadcrumbs) that is displayed after selecting a particular endpoint.

This screen provides the following functionalities:

- Classification drop-down menu - Click and select the required classification.
- Endpoints Counts by Type pane - This pane provides a summary of the selected classification through the following groups:
  - Legend - Hover the cursor on Legend to view color classifications used for various categories.
  - Sunburst graph - Provides the summarized view of all endpoints in various categories, hierarchies, and counts.

    Note: Clicking on a category sets the appropriate category as the new active classification.
  - Classification - Displays selected classification in bread crumbs

    Note: Clicking a breadcrumb link sets the appropriate classification as the new active classification.
  - Sub-Types (Optional) - Displays the count of sub-types under classification

    Note: Clicking a sub-type link sets the appropriate sub-type as the new active classification
- All selected classification Endpoints pane - This pane provides the specified information of each endpoint in selected classification under the following categories:
  - Device Type
  - Device Name
  - MAC Address
  - Last Seen

8.11 Assigning Tags

Perform the following steps to assign a tag to a device:

1. On CVP, click **Provisioning** > **Device Tags**.

   The system displays the Device Tags screen.

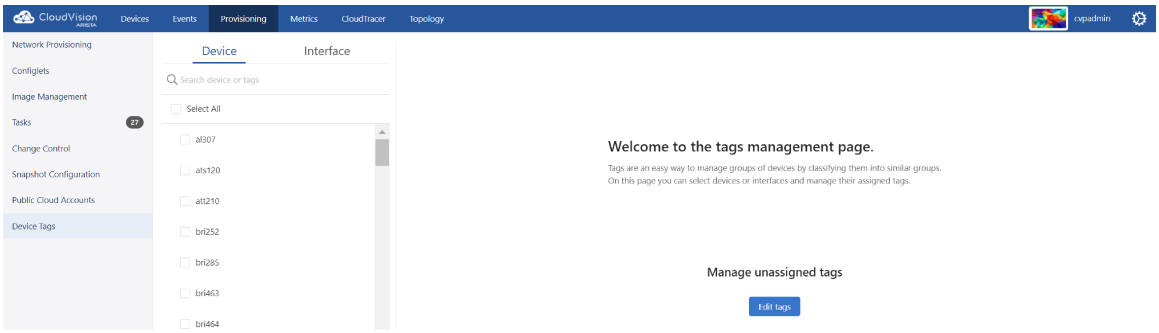

Figure 72: Device Tags Screen

Note: To assign tags to interfaces, click the **Interface** tab.

2. On the main panel, select device(s) that you want to create tag for.

   Note: Tags should be of the form <*label*>: <*value*>. For example, owner: Bill.

3. Select required devices.

   The system displays the **Assigned tags** panel.

   Note: Optionally, use the search bar for searching required devices.

4. Under **User Tags** > **Add or create tags**, provide the required information in the **Type the label then the value separated by a colon** field.

   Note: The **System Tags** panel displays tags automatically created by CVP.

5. Click **Create and Assign**.

   The new tag is displayed under **Manage assigned tags**.

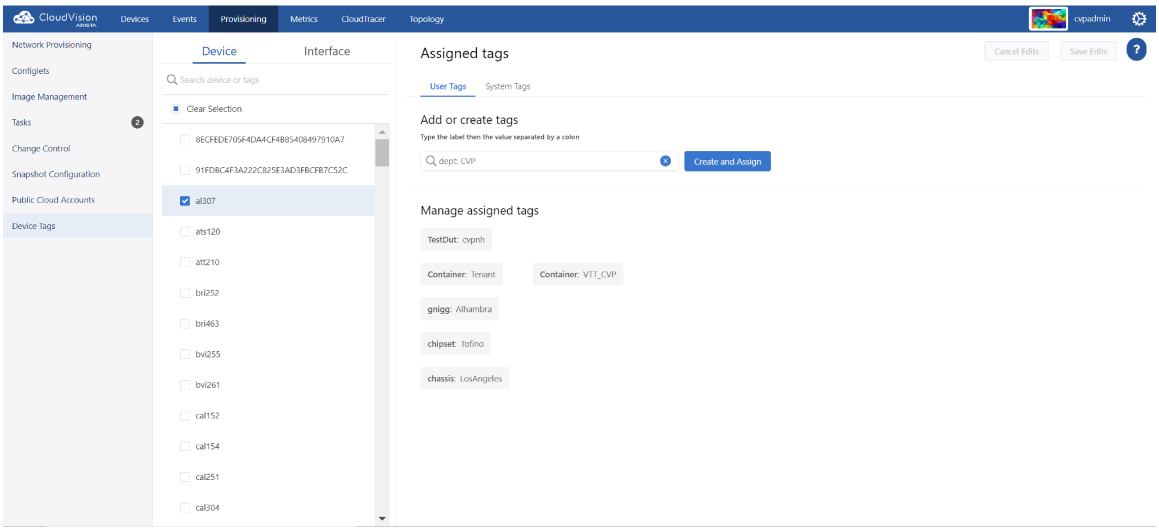

Figure 73: Create and Assign

Note: To delete a tag, click on the inessential tag > the minus sign > **Save edits**.

8.11.1 Adding or Removing Tags from Multiple Devices

Perform the following steps to either add or remove a tag that is available in selected multiple device(s):

1. On the main panel of the device tags screen, select required devices.
2. Click the desired tag.

The system pops up plus and minus signs beneath the tag.

3. Click either the plus sign to add this tag to all selected devices, or the minus sign to remove it from all selected devices.

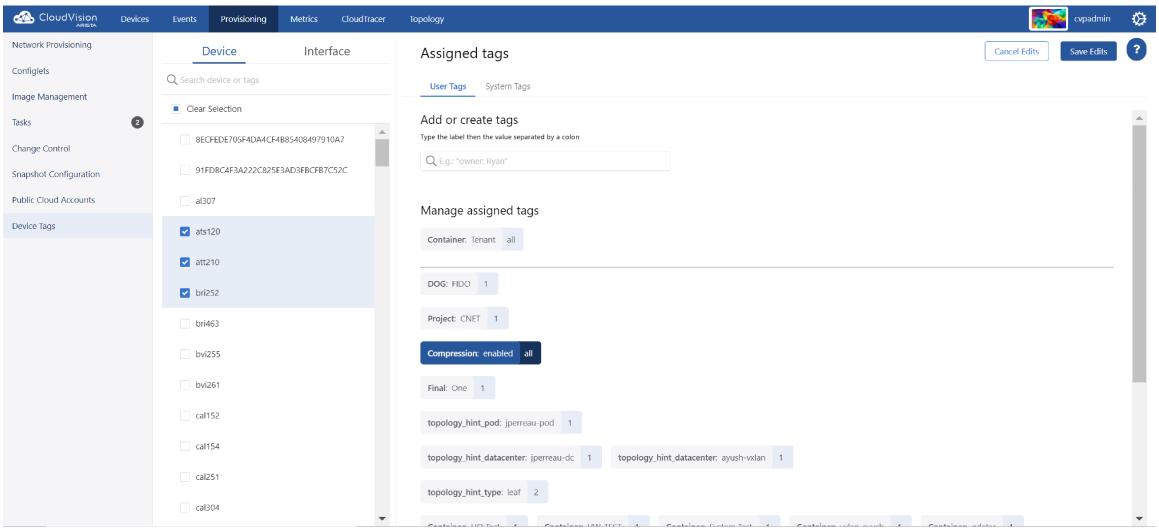

Figure 74: Adding a Tag to Selected Devices

4. Click **Save edits**.

8.11.2 Managing Unassigned Tags

Perform the following steps to manage unassigned tags:

1. On the main panel of the device tags screen, click **Edit tags**.

   The system displays all unassigned tags.

2. Click the inessential tag.

   The system displays the inessential tag in red.

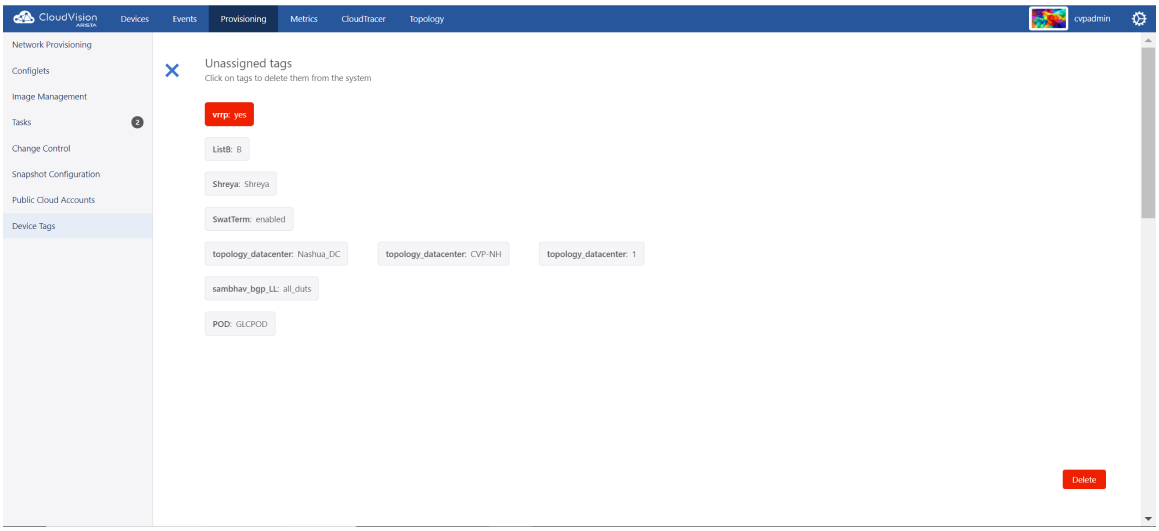

Figure 75: Delete Unassigned Tags

3. Click **Delete**.

   The system deletes the tag from CVP.

8.12 Accessing Metrics

The Metrics application creates customizable dashboards consisting of multiple metrics across various datasets in different views. You can quickly view metrics data gathered from devices configured for streaming telemetry data to CVP.

Related topics:

- Metrics Summary Screen
- Creating Dashboards
- Editing Dashboards
- Editing Views

8.12.1 Metrics Summary Screen

On the CloudVision portal, click the Metrics tab to view the Metrics screen. This screen consists of the Dashboards tab and the Explorer tab.

8.12.1.1 Dashboards Tab

The Dashboards summary screen lists existing dashboards along with other options.

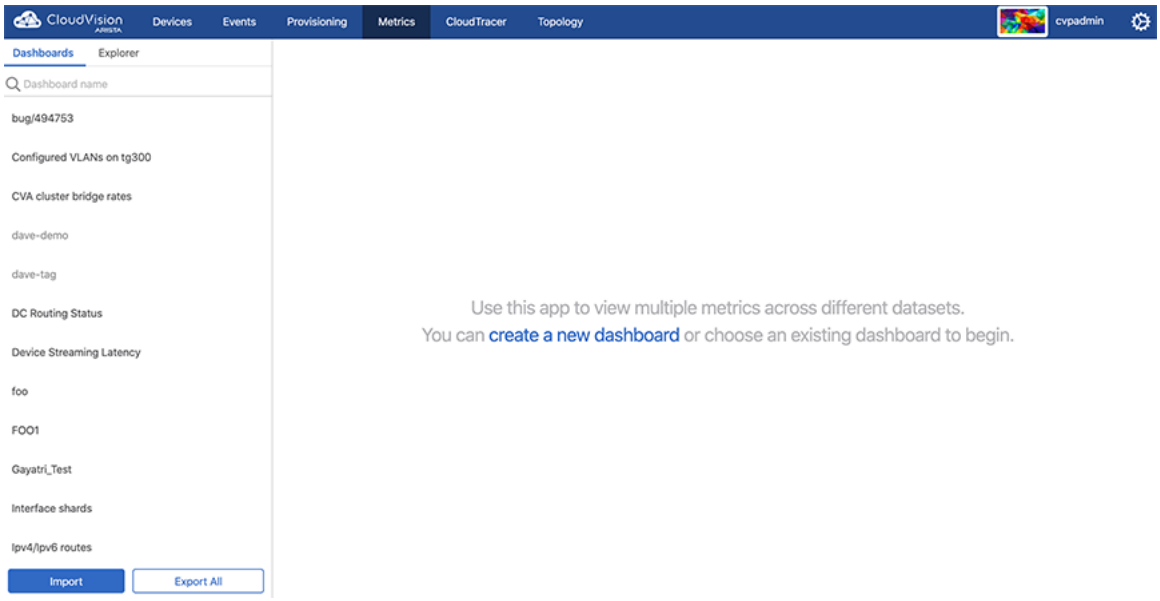

Figure 76: Dashboards Screen

Left Pane

The left pane provides the following options:

- Dashboard name search field - Perform a search of dashboard names
- List of current dashboards - Hover the cursor on a dashboard to view a vertical ellipsis button on the right end of the corresponding pane. Click on the ellipsis button to get the following options:
  - **Add a View** - Click to add a new view based on chosen metrics
  - **Delete** - Click to delete the corresponding dashboard

Right Pane

The right pane provides the **create a new dashboard** option.

8.12.1.2 Explorer Tab

The initial **Explorer Summary** screen does not display any data.

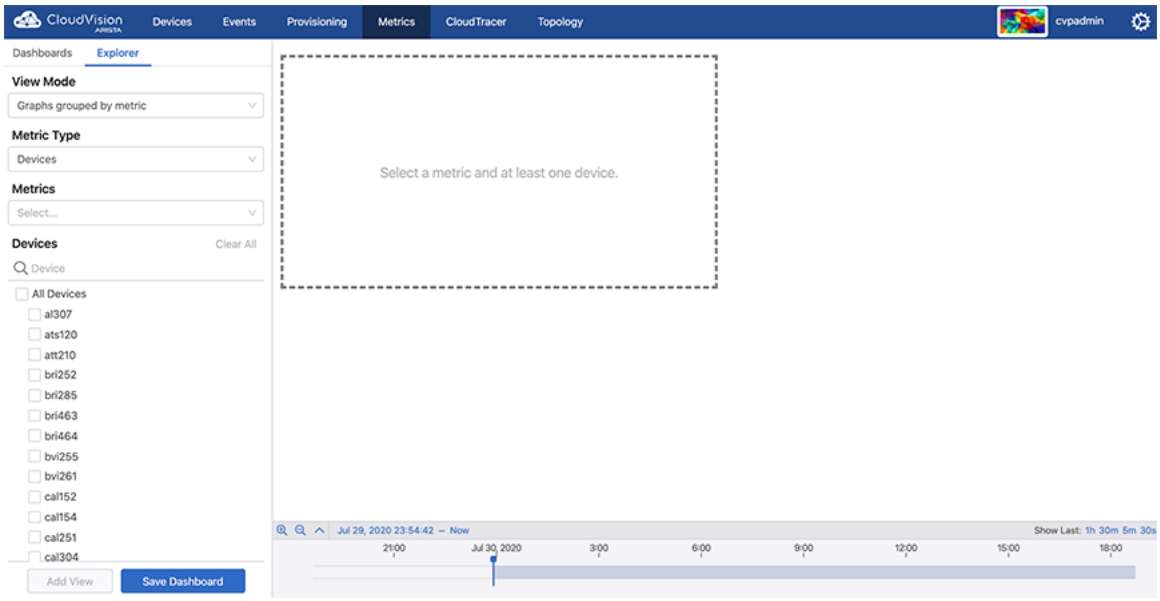

Figure 77: Explorer Screen

To view metrics data, you must either select an existing dashboard from the Dashboards tab or provide the following information in the left pane of **Explorer** screen:

- **View Mode** - Select the **View** mode. Options include:
  - Graphs grouped by dataset - Displays multiple metrics for appropriate metric type
  - Graphs grouped by metric - Displays one metric for multiple entities in appropriate metric type
  - Table - Displays multiple metrics for multiple entities in appropriate metric type
  - Aggregate - Displays grouped metric values for multiple entities in appropriate metric type
- **Metric Type** - Select the metrics type (Devices, Interfaces, Analytic processes, or CloudTracer connections)
- **Metrics** - Select the required option based on appropriate view mode and metric type
- **Devices/Interfaces/Analytics/Connection**s
  - Search field - Perform a search of specified entities
  - List of datasets - Select one or more dataset; or dataset groups

  Note: The field name differs based on the selected combination of **View Mode** and **Metric Type**.
- **Clear All** - Click to clear the selection of all datasets
- **Add View** - Click to add a new view
- **Save Dashboard** - Click to save the current dashboard
- **Dotted box** - Indicates the view that is currently being edited

8.12.2 Creating Dashboards

Perform the following steps to create a dashboard:

1. Under the **Dashboards** tab on the **Metrics** screen, click **create a new dashboard** in right pane.
   The system displays the **Explorer** screen.
2. Provide the appropriate information in available User Interface (UI) elements in the left pane.
   The system creates a view based on the information provided and displays it in the right pane.

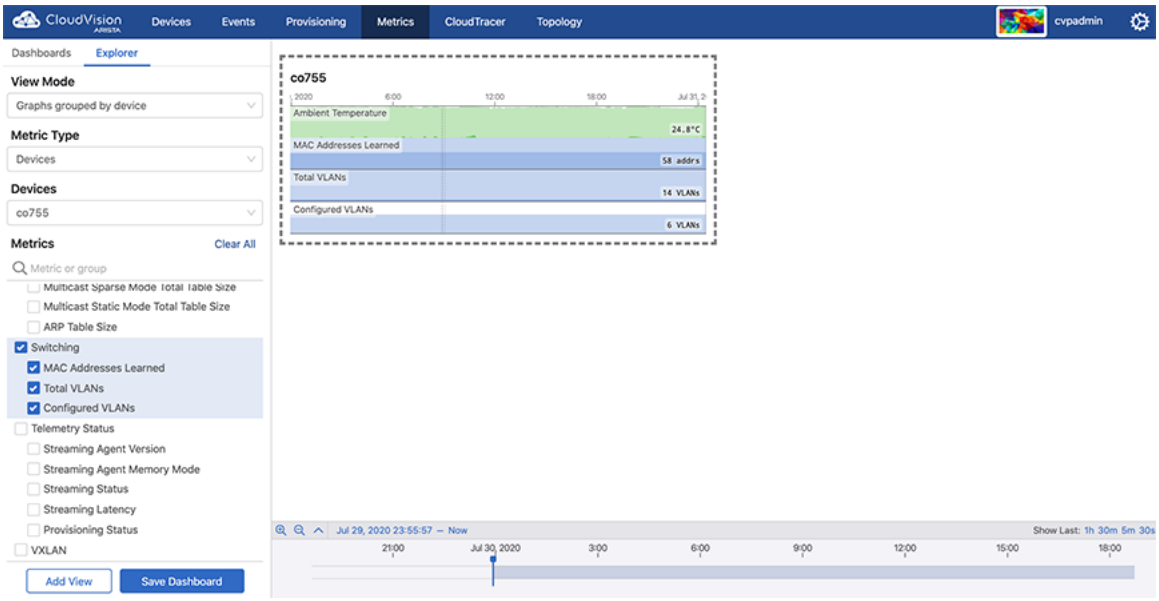

Figure 78: Explorer Screen with View

Note: To create a new view, click **Add View** at the lower end of the left pane. To edit an existing view, refer to Editing Views.

3. Click **Save Dashboard**.

The system displays the Save Dashboard dialog box.

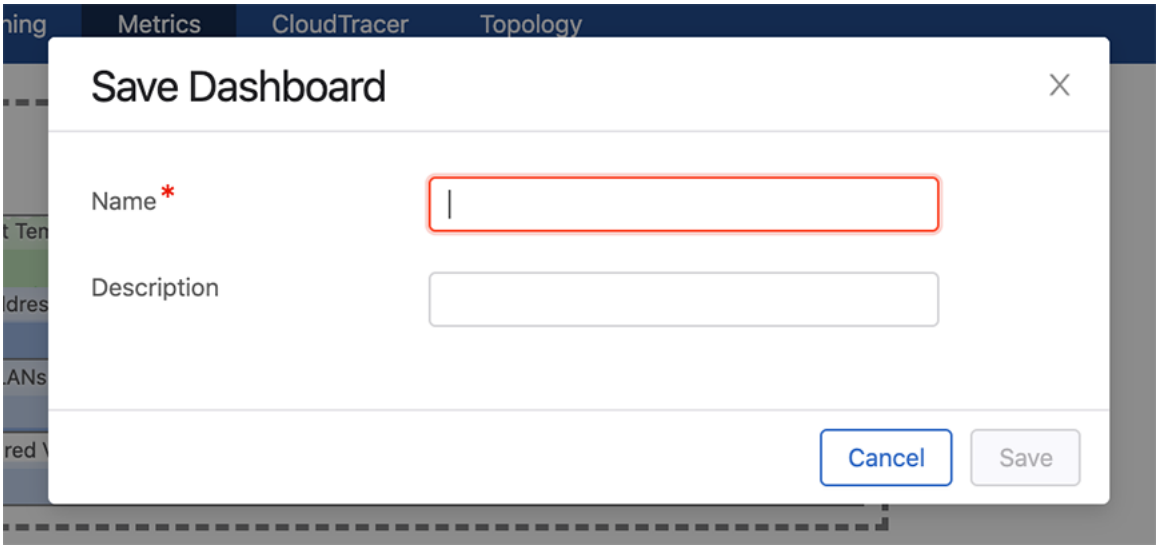

Figure 79: Save Dashboard Dialog Box

4. Type a name in the **Name** field.
5. (Optional) Type a description in the **Description** Field.
6. Click **Save**.

Note: If you create a dashboard with a name that already exists, the system displays a 'Save & Overwrite' warning through the **Confirm** dialog box.

8.12.3 Editing Dashboards

Perform the following steps to edit a dashboard:

1. On the CloudVision portal, click the **Metrics** tab.

The system displays the **Metrics** screen with the list of current dashboards on the left pane.

Note: Alternatively, you can either add a view in an existing dashboard or delete a dashboard by hovering the cursor on the corresponding dashboard and selecting the appropriate option.

2. On the left pane of **Dashboards** screen, click the required dashboard.

   The system displays the dashboard details screen.

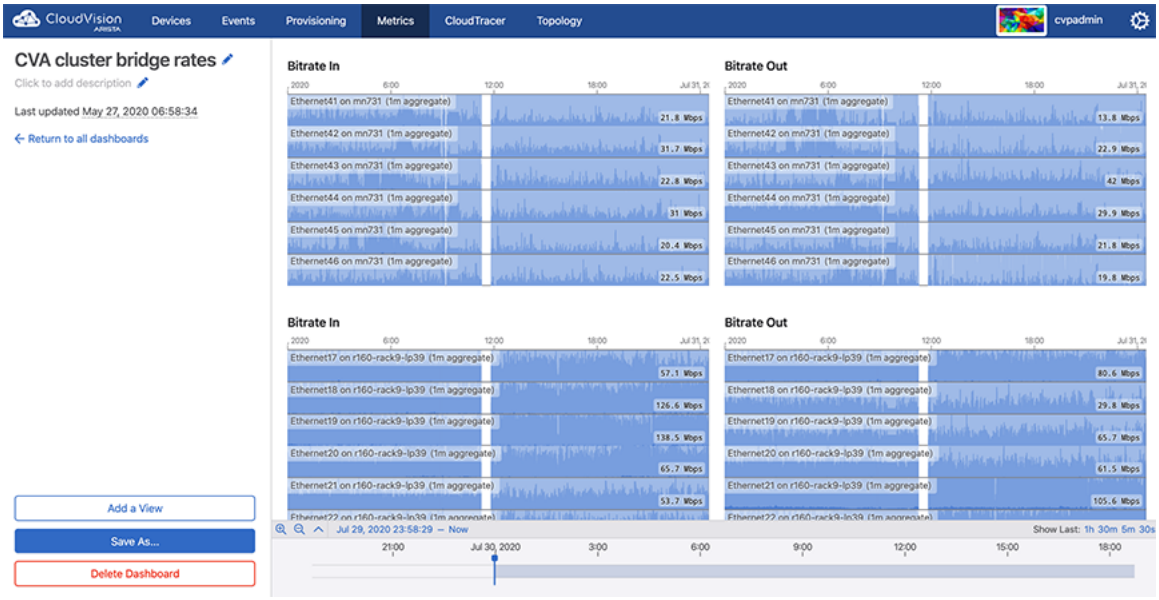

Figure 80: Dashboard Details Screen

3. Perform the following actions in the left pane:

   - Click the dashboard name to edit it and press **Enter**.

     Note: Alternatively, click the edit icon under **METRIC DASHBOARD** to edit the dashboard name. Type the new name and press **Enter**.

   - Click the dashboard description to edit it and press **Enter**.

     Note: Alternatively, click the edit icon under **DESCRIPTION** to edit the dashboard description. Type the new description and press **Enter**.

   - Click **Add a View** to add a new view.

     Note: To edit an existing view, refer to Editing Views. To delete the current dashboard, click **Delete Dashboard** and then click **Remove** on the Confirm dialog box.

4. Click **Save As**.

   The system displays the **Save Dashboard** dialog box.

   Note: Alternatively, you can edit the dashboard name and description in the **Save Dashboard** dialog box.

5. Click **Save**.

   Note: If required, select another dashboard from the Change dashboard drop-down menu. Alternatively, you can select another dashboard from the list under RECENTLY VIEWED. The system displays up to five dashboards under RECENTLY VIEWED.

8.12.4 Editing Views

Perform the following steps to edit a view:

1. On the CloudVision portal, click the **Metrics** tab.

The system displays the **Metrics** screen with the list of current dashboards on the left pane.

2. On the left pane of **Dashboards** screen, click the required dashboard.

   The system displays the **Dashboard details** screen.

3. On the right pane, hover the cursor on the required view pane.

   The system displays editable options at the right end of the pane.

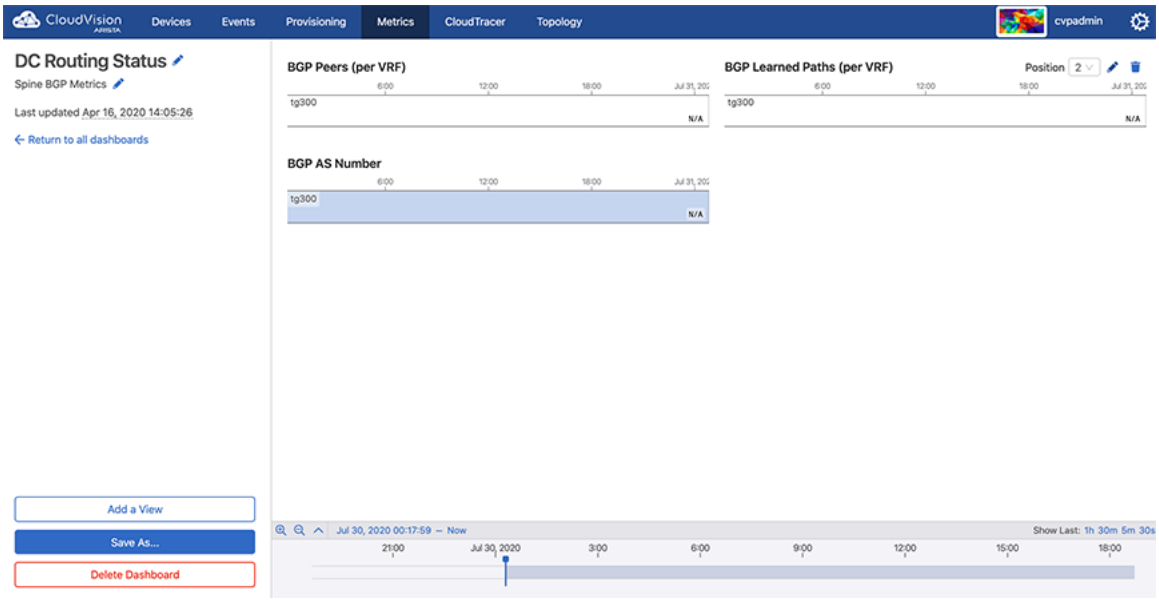

Figure 81: View Edit Options

Note: To delete a view, click the appropriate trash icon and then click **OK** on the confirm dialog box.

4. Select the desired sequence from the **Position** drop-down menu.
5. Click the **Edit** icon.

   The system displays editing options in the left pane.

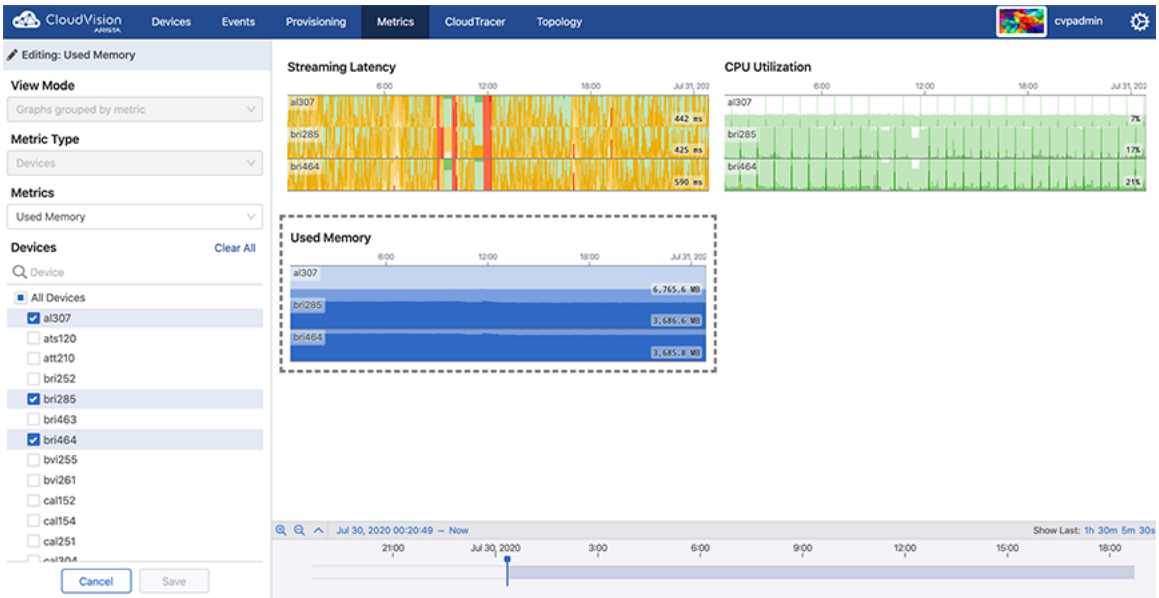

Figure 82: Metrics Editing Options

6. Provide desired changes in the **Editing View** pane.
7. Click **Save**.

Note: If you are editing a view while creating a dashboard, click **Done** at the lower end of the left pane.

8.13 Topology View

You can view the network hierarchy for the devices and subnetwork in real-time. The topology view is available for devices running on LLDP including Arista switches and connected neighbors.

Related topics:

- Setup
- Overlays
- Custom Topology Views
- Changing the Node Type
- Nodes and Features

8.13.1 Setup

You can customize the topology by completing the following steps.

1. Click the **Topology** tab to view your network.
2. To enter layout hints, click on a device in the topology view and then click on the layout tab.
   Following example shows the detail of a device.

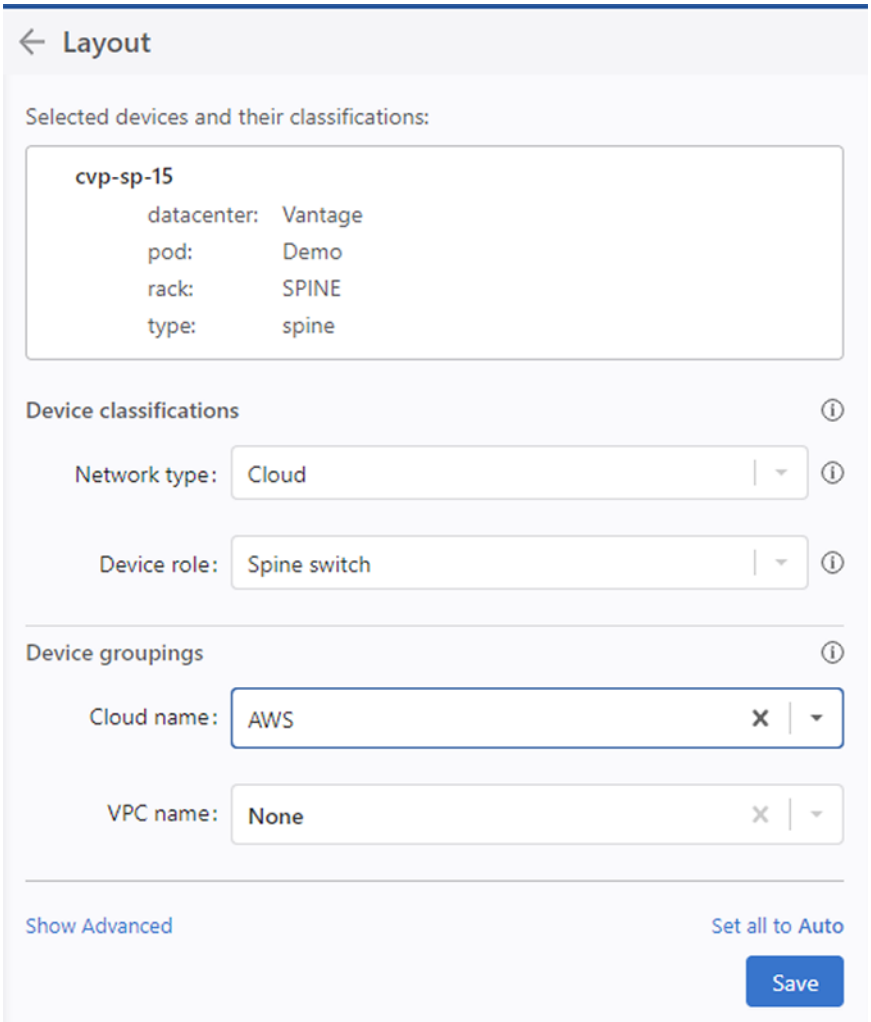

Figure 83: CVP Detail Layout

8.13.2 Overlays

You can superimpose link-level metrics overlay onto the network topology. Use the Layers Panel to view these overlays and color-codes based on the severity of that metric. Following are the overlays supported in this release.

The following table lists the Overlays supported in this release.

Table 11: Supported Overlays

Overlay	Description
Bandwidth Utilization	Shows the bitrate as a percentage of the speed of the link. It uses the maximum bitrate in either direction on the link, averaged out over a one-minute window. Light green indicates a small percent of the link is being used, while darker greens indicate higher usage. Beyond 80% utilization, the links show up in yellow or red.
Traffic Throughput	Shows the bitrate of a link as an absolute number. Darker blues indicate higher utilization.

Error Rates	Show if either end of a link is registering input or output errors (for example, CRC Errors). It uses a one-minute window, and displays severity in increasingly dark reds.
Discard Rates	Indicate that a link is dropping packets, likely due to congestion. Links discarding more packets in a one-minute window are shown in darker red.
None	Turns off all colors.

8.13.3 Custom Topology Views

From the **Topology** tab, you can perform the following steps to customize a view:

1. To move a rack to a different pod use the **Pod** field. For example, the switch called cv-demo-sw3 is set to be in a pod 1.

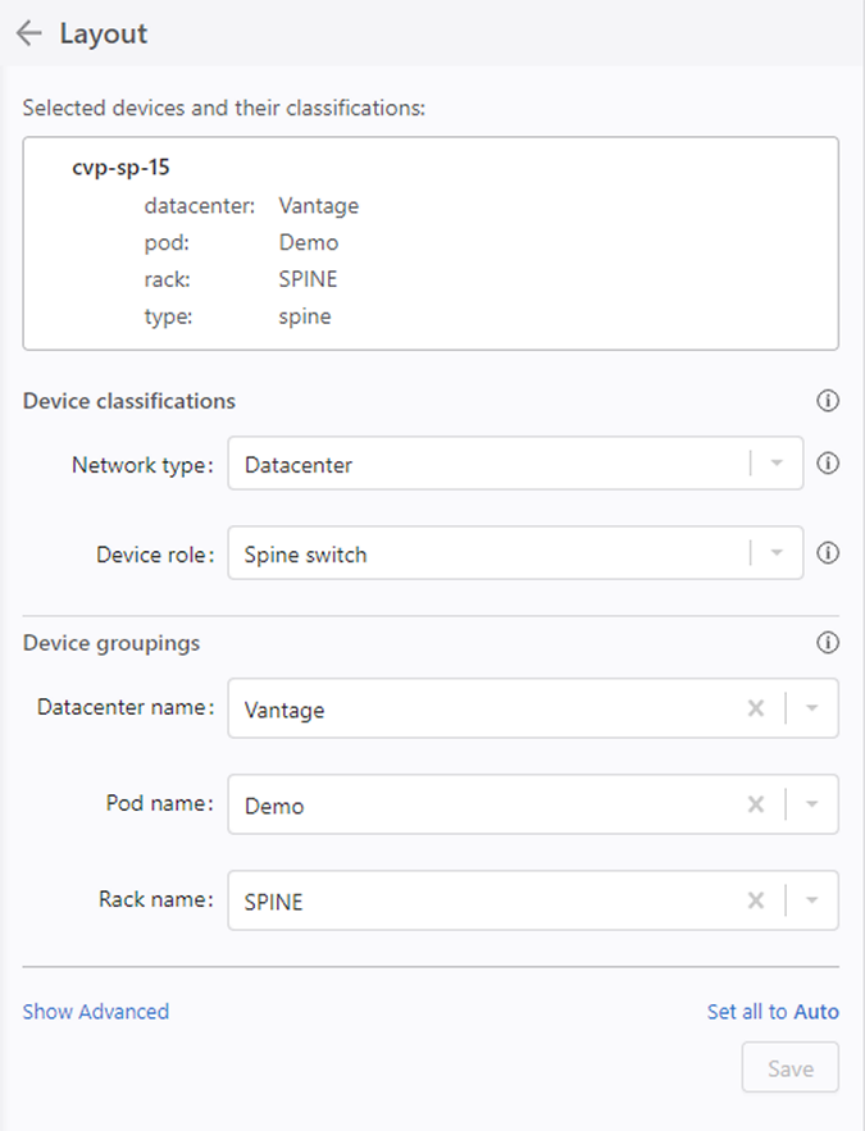

Figure 84: User Layout Hints

2. To setup the pod or rack names, apply a layout hint for switch with alternate name or pod hint for the spine switch to rename the pod. Following example shows the top-of-rack switch cv-demo-sw3 default name change via the rack layout hint.

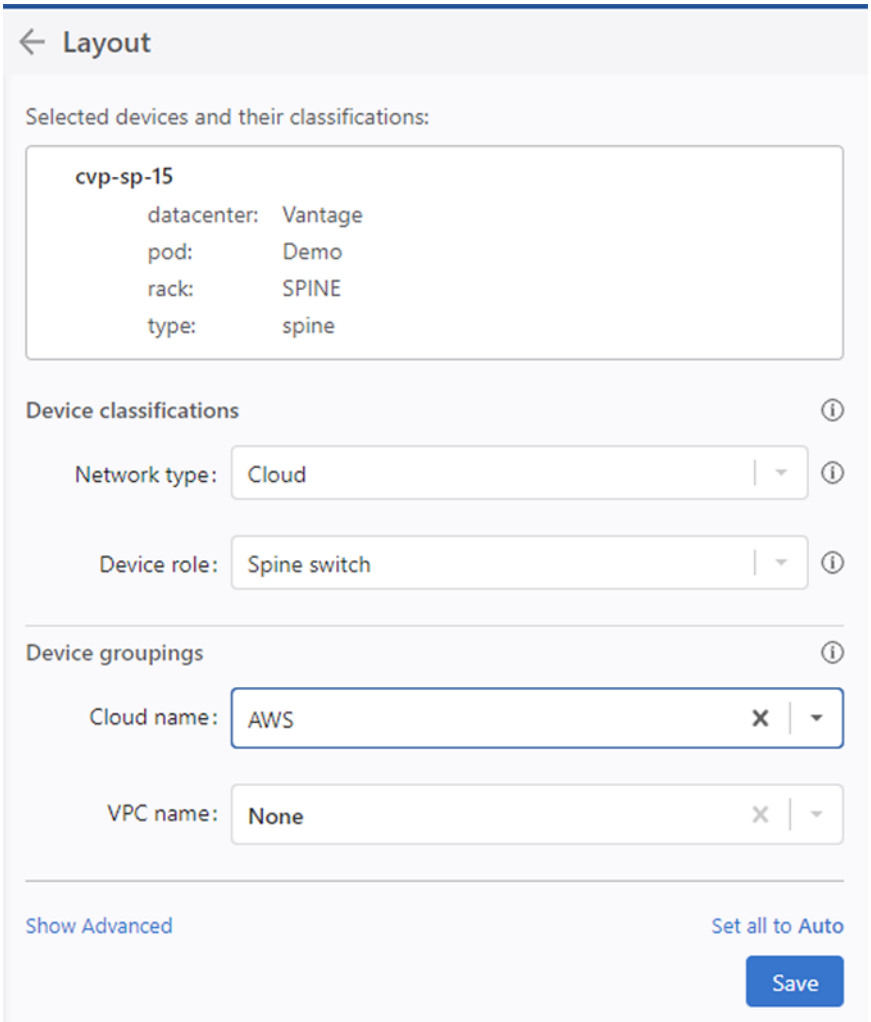

Figure 85: Device Details in Layout

8.13.4 Changing the Node Type

The following table lists the node types supported by the Topology view.

Table 12: Supported Node Type

Node Type	Description
Edge Device	The device is an edge device, for example, leading to the Internet or another network, or a similar function device.
Core Switch	The device is at the core level switch (above spines) or similar function device.
Spine Switch	The device is a pod level (spine or aggregation) switch or similar function device.
Leaf Switch	The device is a top of rack switch or similar function device.
Endpoint Device	The device is a server or similar endpoint device.

Setting the **Node Type** layout hint gives the **Topology** view of the type of device selected. Selecting **skip auto-generating** forces the auto tagger to ignore the device and not assign or modify any of the hints.

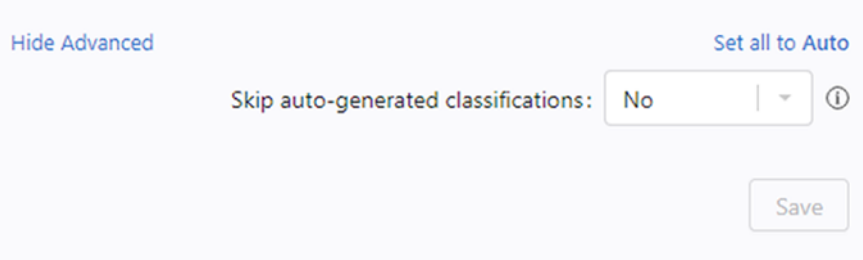

Figure 86: Changing Node Type

8.13.5 Nodes and Features

Nodes are arranged in clusters. To expand a cluster, click on the representative **Cluster-node**. To collapse a cluster, click on the minus (**-**) icon.

You can select various overlays on the graph for color coding links.

To see details about a node and its neighbors, click on the **Node**. You can also see the immediate neighbors of the device and the metrics related to particular physical links between devices by clicking **Neighbors List**.

8.14 Accessing Events

You can access the following events screens:

- Events Summary Screen
- Event Details Screen

Related topices:

- Events Summary Screen
- Event Details Screen
- Configuring Event Generations
- Managing Events
  - Disabling All Events of the Selected Type
  - Disabling All Events of the Selected Type with Exceptions
- Acknowledging Events
- Configuring Notifications
  - Configuring Status
  - Configuring Platforms
  - Configuring Receivers
  - Configuring Rules

8.14.1 Events Summary Screen

The events summary screen displays all events, and configures alerts and event generation. To view this screen, click **Events** on the CloudVision portal.

The **Events** screen provides the following information and functionalities:

- Left Pane
  - A search field for events, devices, and interfaces

- Buttons to perform a search based on severity levels (info, warning, error, and critical)
- A toggle button to add and remove acknowledged events from search results
- The count of events from search results
- A button that allows you to display new events and also provides their count A list of events (the most recent are shown at the top of the list)

- Right Pane
  - The count of all events and devices from search results
  - The time frame from which events are selected
  - Devices that have reported the most events and errors (shown in the **Most Active Devices** pane)
  - Most common events (shown in the **Most Common Events** pane)
  - Count of each error type from device errors (shown in the **Event Severities** pane)
  - A chronological history of all errors (shown at the bottom of the screen)

Click the **Events** tab to view all events.

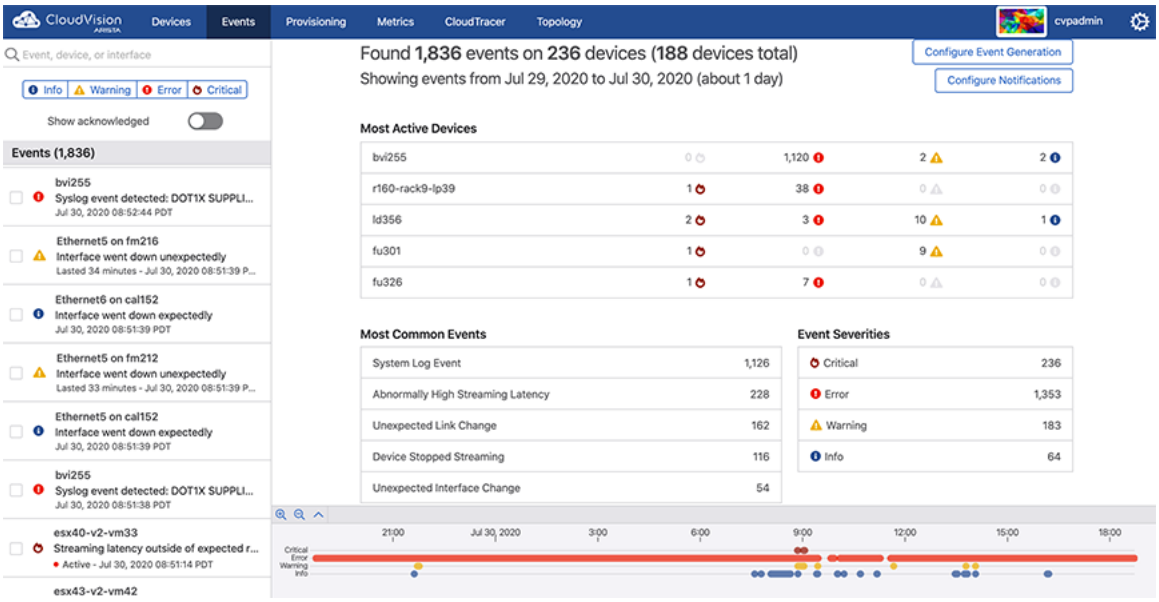

Figure 87: Events Summary Screen

8.14.2 Event Details Screen

An event details screen displays appropriate event details, acknowledges the event, and configures event generation. To view this screen, click one of the events listed in the left pane from the **Events** screen.

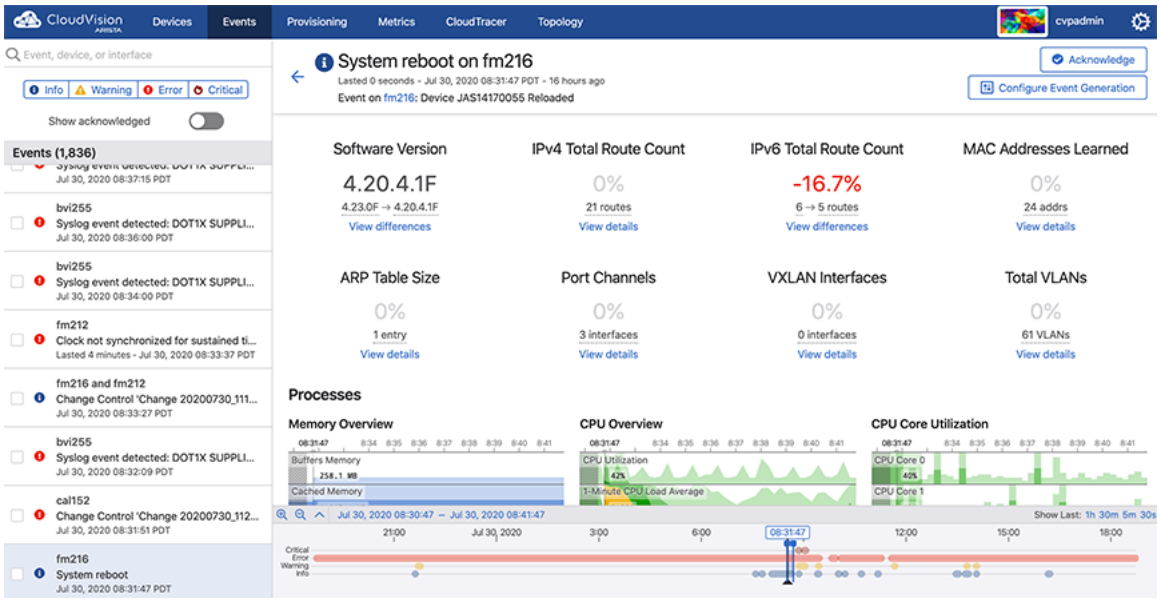

Figure 88: Event Details Screen

This screen provides the following information and functionalities in the right pane:

- Left arrow to return to the events summary screen
- Warning of the event
- Time when event details were captured
- Hover the cursor on the event name. The system displays a popup window with event details.

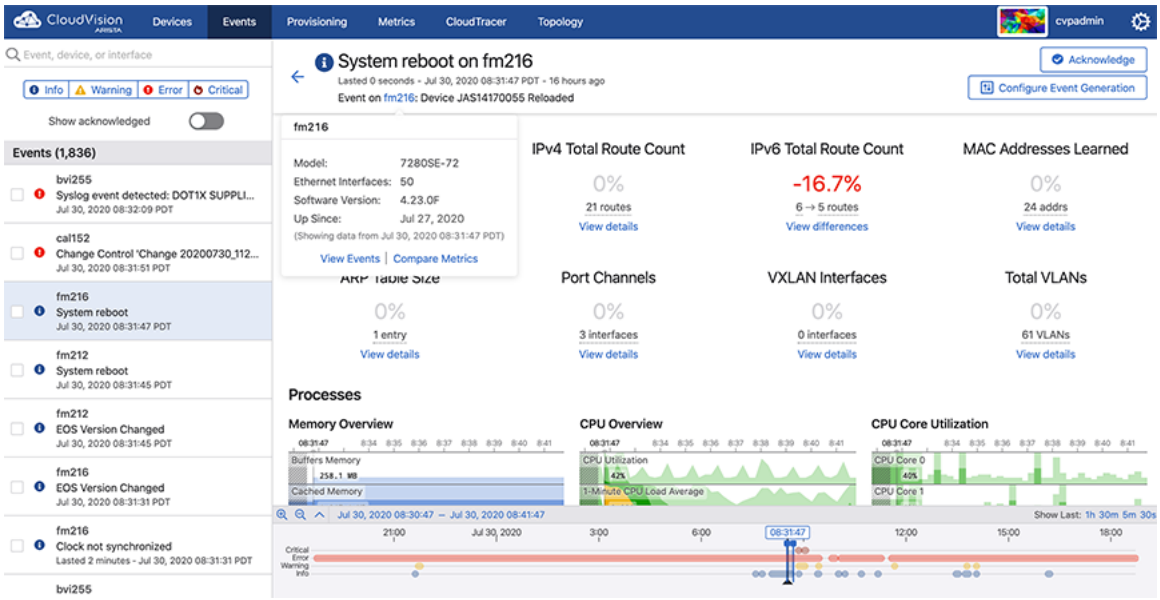

Figure 89: Event Name Popup Window

The popup window provides the following options:

- Click **View Events** to view search results with the same event name.

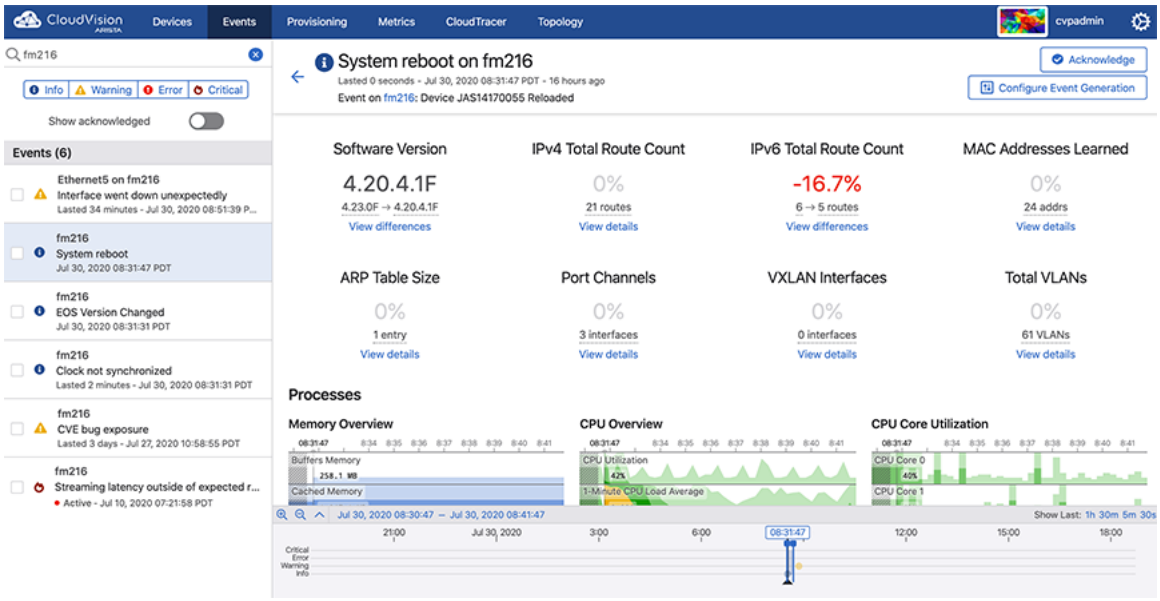

Figure 90: Search Results with the Same Event Name

- Click **Compare Metrics** to navigate to the **Explorer** tab in Metrics app.
- Hover the cursor on the event name. The system displays a popup window with device details in that location.

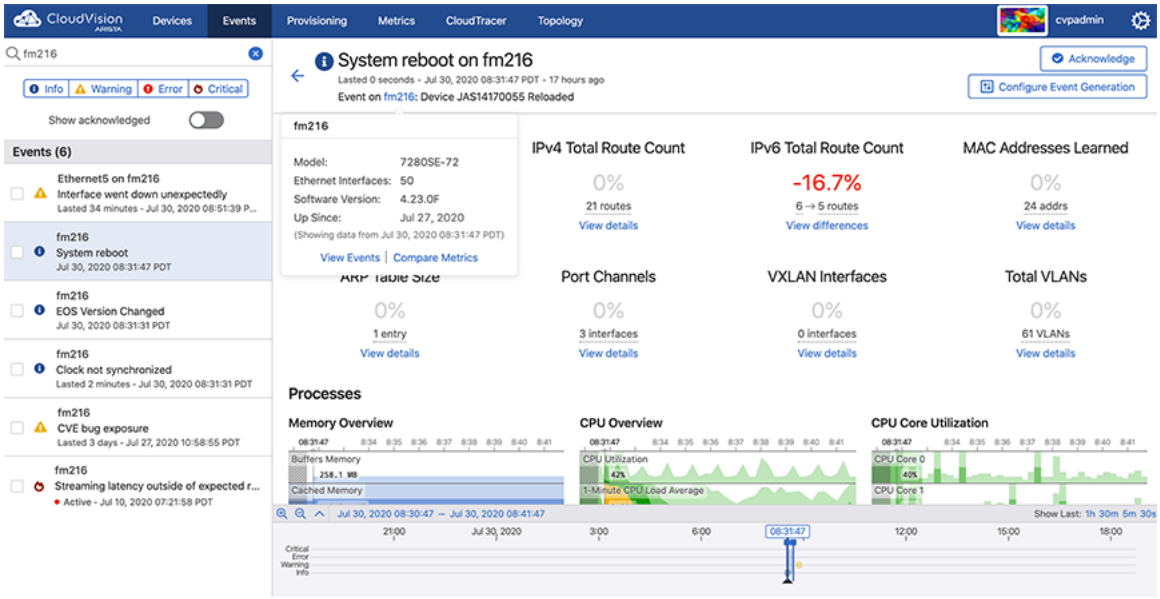

Figure 91: Location Name Popup Window

The popup window provides the following options:

- Click **View Events** to view search results with the same location name.

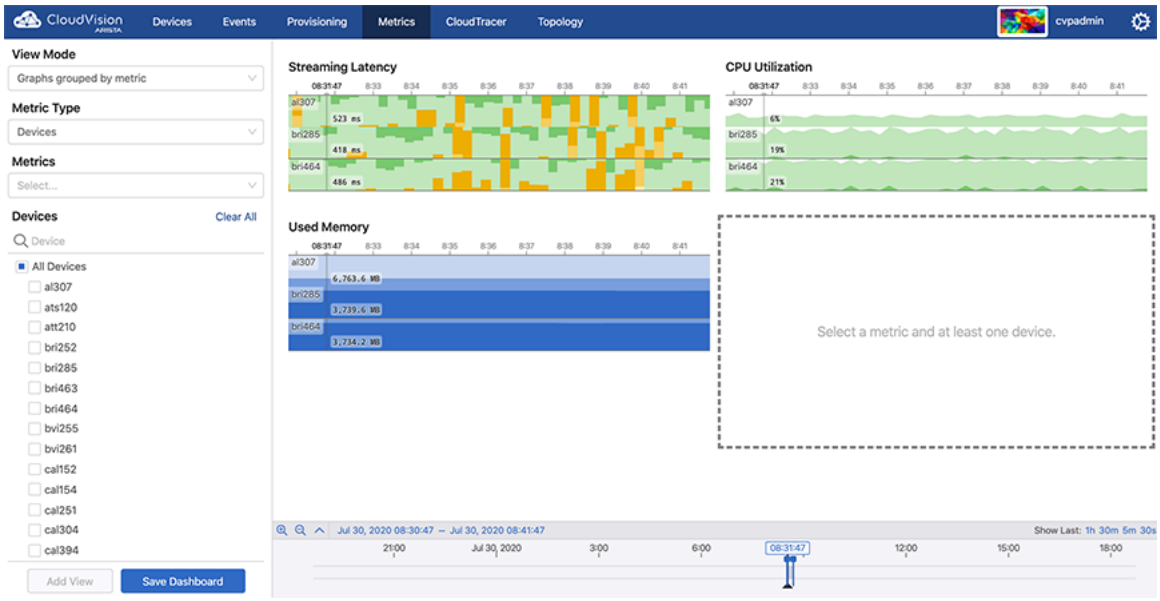

Figure 92: Search Results with the Same Location Name

- Click **Compare Metrics** to navigate to the **Explorer** tab under **Metrics**.
- The **Acknowledge** button to acknowledge the appropriate event.
- The **Configure Event Generation** button to configure the generation of appropriate event.
- Metric details of the event
- A chronological history of all errors (shown at the bottom of the screen)

8.14.3 Configuring Event Generations

Configuring events customizes the prerequisites of an event.

Perform the following steps to configure the settings for generating events:

1. On the CloudVision portal, click the **Events** tab. The system displays the **Events** screen.
2. Click **Configure Event Generation** at the upper right corner of the **Events** section. The system displays the **Generation Configuration** screen with all configurable events listed in the left pane.

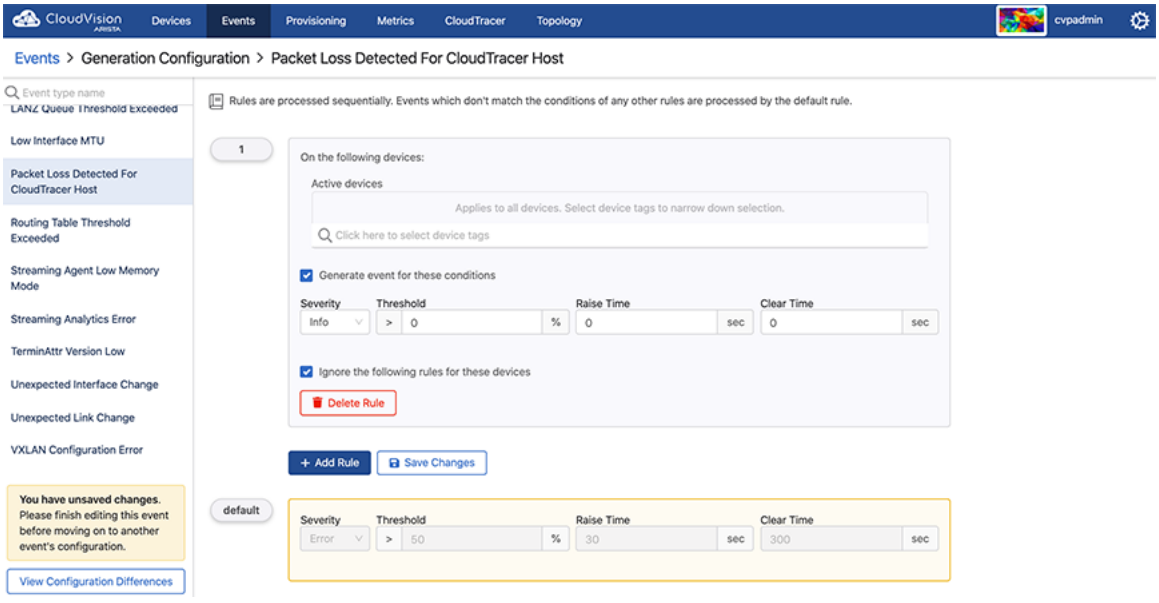

Figure 93: Generation Configuration Screen

Note: Alternatively, you can go to an event details screen and click **Configure Event Generation** to configure rules for generating events.

3. Click the required event in the left pane.
4. Click **Add Rule** in the lower end of right pane. A new **Condition** pane is displayed on the screen.

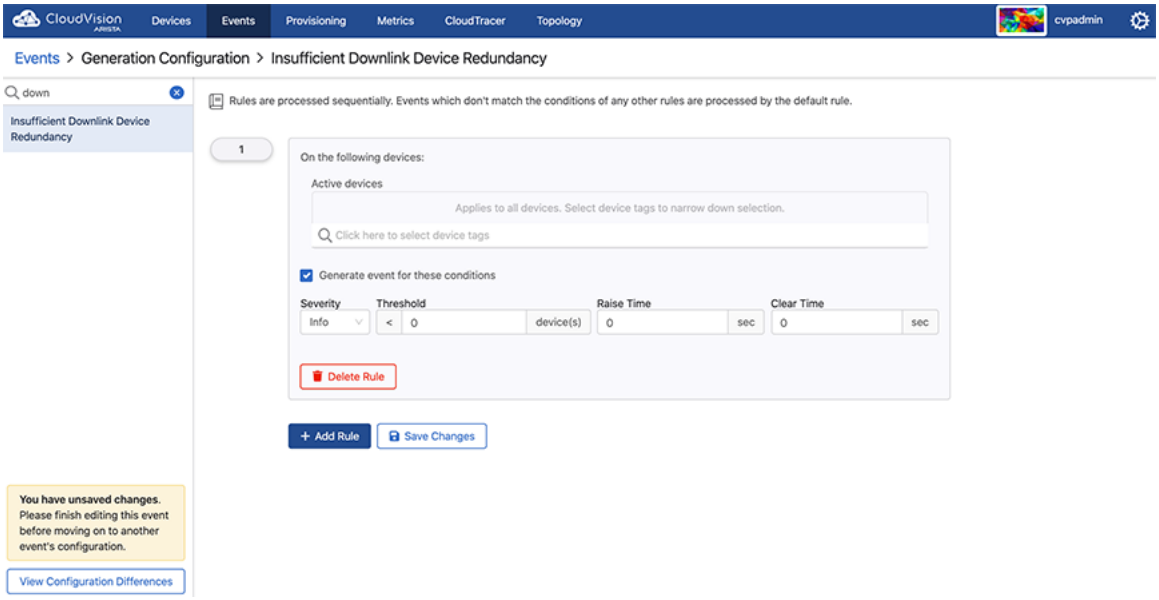

Figure 94: Add Rule Pane in Generation Configuration

5. In the **Condition** pane, click on the search field. The system displays the list of configured devices tags.

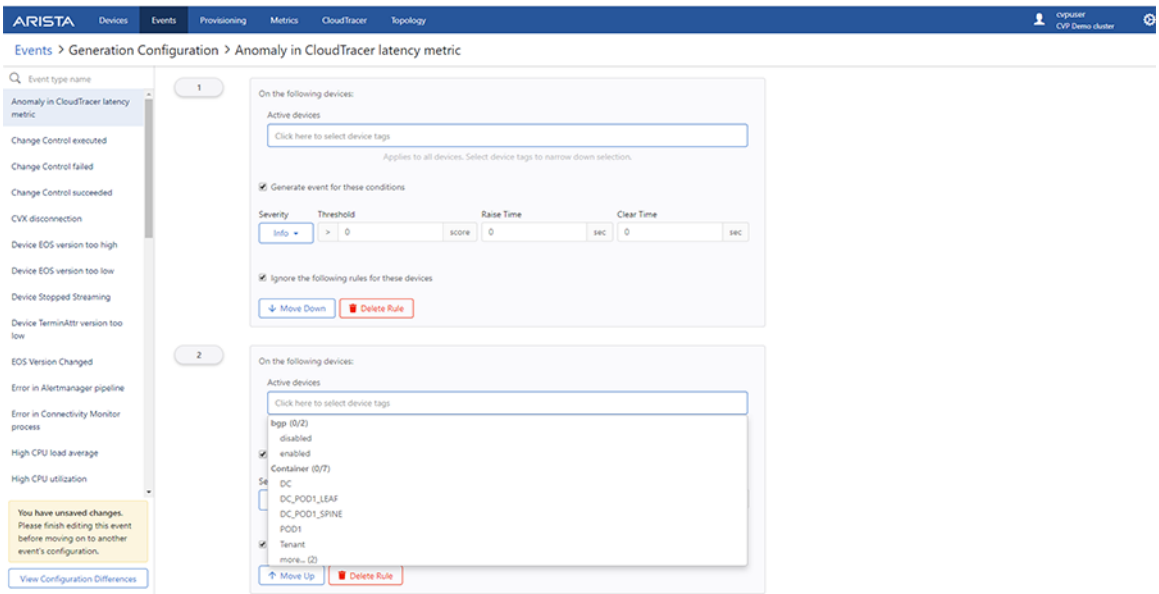

Figure 95: List of Configured Device Tags

Note: Alternatively, you can type the required device tag in the search field for a quick search.

6. Select preferred devices tags from the displayed list.

Note: After you have selected the device, the system displays the count of matched devices. The rule is applicable to all devices when you do not select any device tag.

7. Click on the **Interfaces** search field (available only for interface events).

The system displays the list of configured interface tags. See Figure 96: List of Configured Interface Tags.

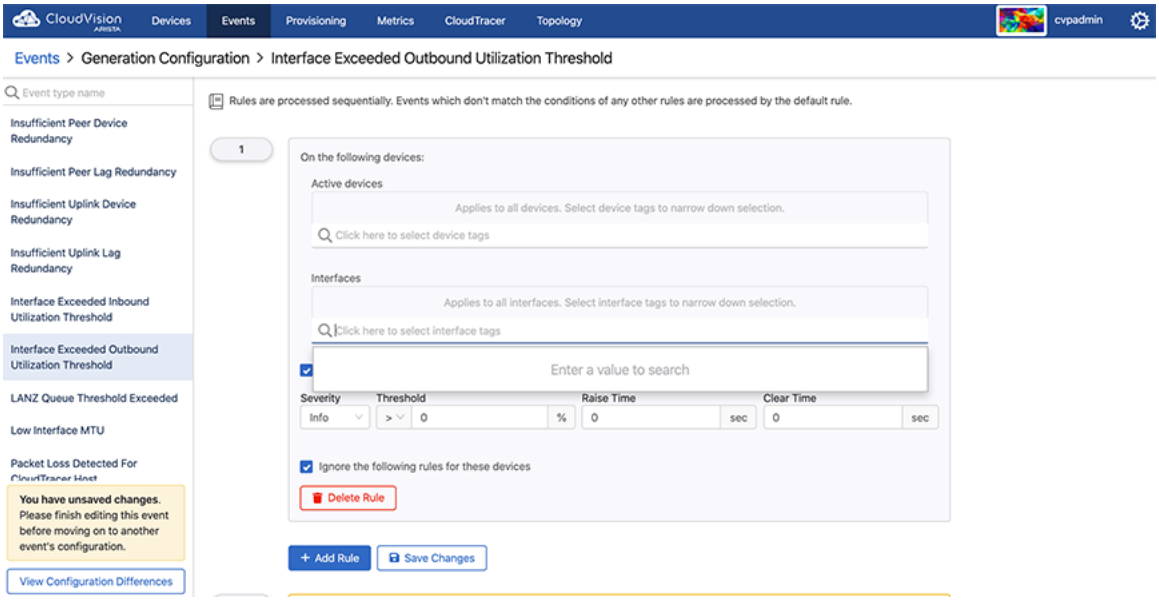

Figure 96: List of Configured Interface Tags

8. Select preferred interface tags from the displayed list.

   Note: After you have selected an interface tag, the system displays the count of matching interfaces. The rule is applicable to all interfaces when you do not select any interface tag.

9. Provide the following criteria required to generate events:

   - **Severity** - Select the severity type from the drop-down menu. Options include **Info**, **Warning**, **Critical**, and **Error**.
   - **Threshold** (applicable only to threshold events) - Type the threshold value.
   - **Raise Time** - Type the preferred wait time (seconds) to create an event after reaching the threshold limit.
   - **Clear Time** - Type the precise time (seconds) to delete an event after the current value goes below the threshold limit.

   Note: Select the **Stop generating events** and checking rules checkbox if you do not want to apply further rules for selected tags. If no tags are selected, further rules are not applicable to any device.

10. Click **Move up** if you prefer to move this rule up in the priority list.

    Note: Rules are processed sequentially. The default rule is applied only when an event does not match any other rules. Click **Delete** rule to delete the corresponding rule. Click **Move down** in configured rules to move the corresponding rule down in the priority list.

11. Click **Save** in the left pane.

    Note: Click **View Configuration Differences** in the lower left pane to view differences in event configurations.

8.14.4 Configuring Event Generations

Configuring events customizes the prerequisites of an event.

Perform the following steps to configure the settings for generating events:

1. On the CloudVision portal, click the **Events** tab. The system displays the **Events** screen.
2. Click **Configure Event Generation** at the upper right corner of the **Events** section. The system displays the **Generation Configuration** screen with all configurable events listed in the left pane.

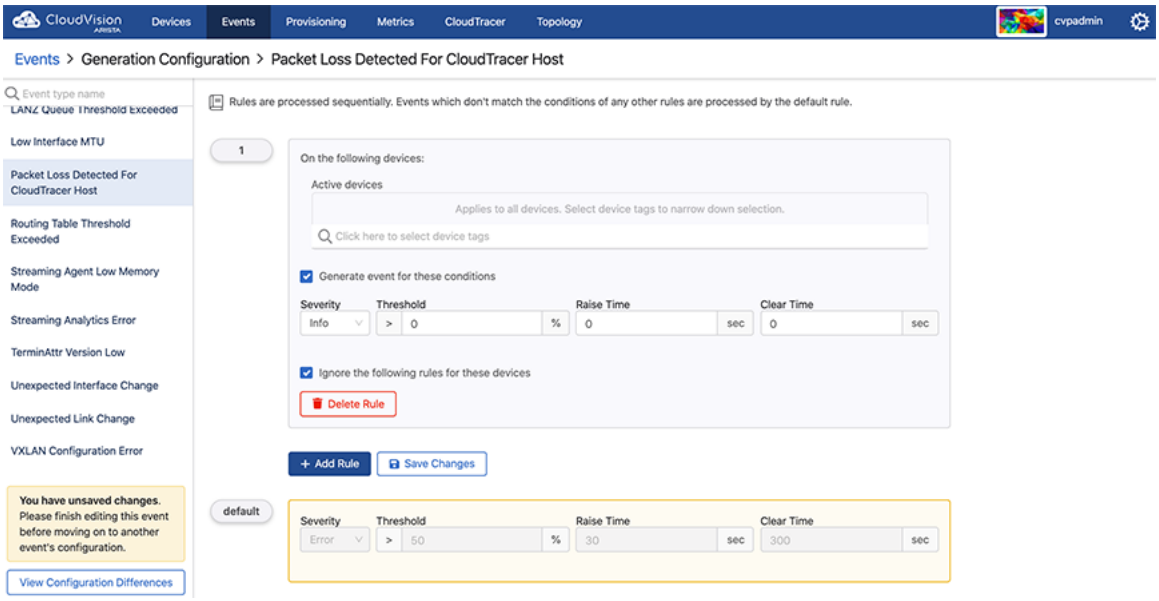

Figure 97: Generation Configuration Screen

Note: Alternatively, you can go to an event details screen and click **Configure Event Generation** to configure rules for generating events.

3. Click the required event in the left pane.
4. Click **Add Rule** in the lower end of right pane. A new **Condition** pane is displayed on the screen.

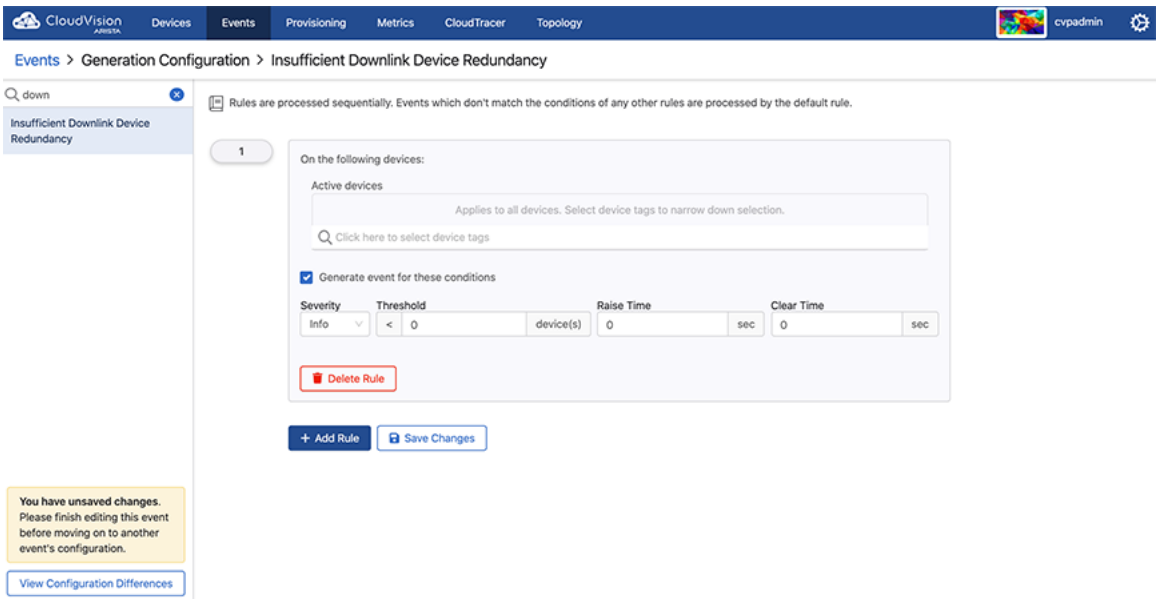

Figure 98: Add Rule Pane in Generation Configuration

5. In the **Condition** pane, click on the search field. The system displays the list of configured devices tags.

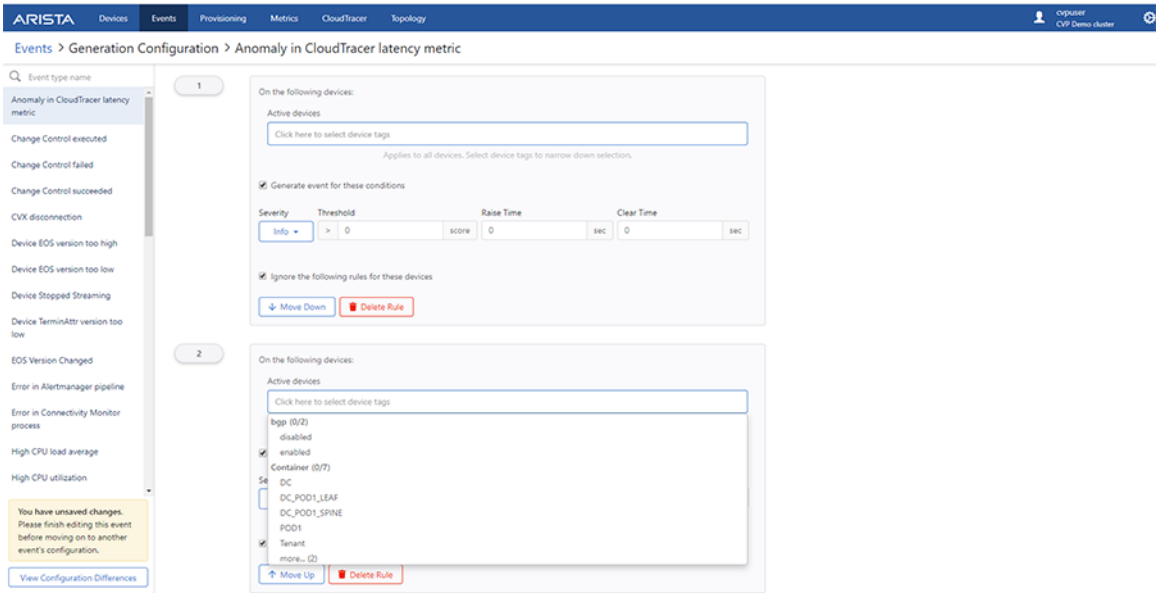

Figure 99: List of Configured Device Tags

Note: Alternatively, you can type the required device tag in the search field for a quick search.

6. Select preferred devices tags from the displayed list.

   Note: After you have selected the device, the system displays the count of matched devices. The rule is applicable to all devices when you do not select any device tag.

7. Provide the following criteria required to generate events:

   - **Severity** - Select the severity type from the drop-down menu. Options include **Info**, **Warning**, **Critical**, and **Error**.
   - **Threshold** (applicable only to threshold events) - Type the threshold value.
   - **Raise Time** - Type the preferred wait time (seconds) to create an event after reaching the threshold limit.
   - **Clear Time** - Type the precise time (seconds) to delete an event after the current value goes below the threshold limit.

   Note: Select the **Stop generating events** and checking rules checkbox if you do not want to apply further rules for selected tags. If no tags are selected, further rules are not applicable to any device.

8. Click **Move up** if you prefer to move this rule up in the priority list.

   Note: Rules are processed sequentially. The default rule is applied only when an event does not match any other rules. Click **Delete** rule to delete the corresponding rule. Click **Move down** in configured rules to move the corresponding rule down in the priority list.

9. Click **Save** in the left pane.

   Note: Click **View Configuration Differences** in the lower left pane to view differences in event configurations.

8.14.5 Managing Events

You can manage an event by customizing event rules differently. Refer to the following examples:

- Disabling All Events of the Selected Type
- Disabling All Events of the Selected Type with Exception

8.14.5.1 Disabling All Events of the Selected Type

Perform the following steps to disable all events of the selected type:

1. Navigate to the **Generation Configuration** screen.
2. Click the required event type in the left pane.
3. In the right pane, Click the **+ Add Rule** button.

   Note: Retain only one rule with no values defined. To disable the event only for selected datasets, select appropriate devices tags in the **Devices** field.
4. Select the **Stop generating events** and checking rules checkbox.

   The system disables all events of the selected event type.

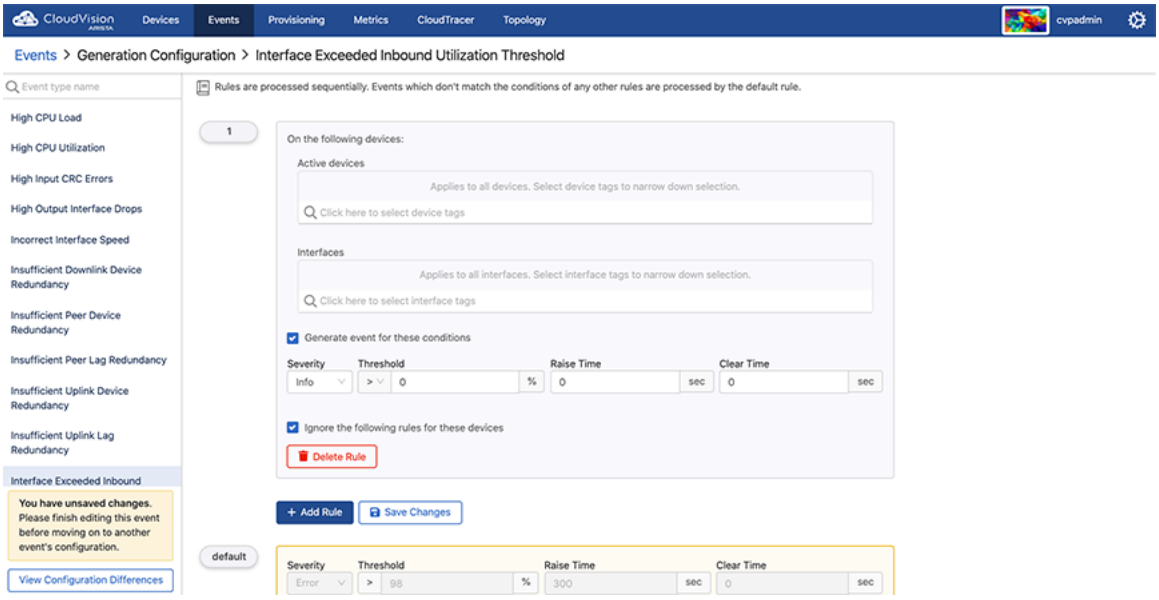

Figure 100: Disable All Events of the Selected Type

5. Click **Save** in the left pane.

8.14.5.2 Disabling All Events of the Selected Type with Exception

Perform the following steps to disable all events of the selected type with exceptions:

1. Navigate to the **Generation Configuration** screen.
2. Click the required event type in the left pane.
3. In the right pane, Click the **+ Add Rule** button.
4. In the **Conditions** pane, provide the device tags that you still want to generate an event for. The system creates rule 1.

   Note: If you need devices with different conditions, add another rule by repeating steps 3 and 4
5. Click the **+ Add Rule** button.
6. In the appropriate **Conditions** pane, select the Stop generating events and checking rules checkbox. The system creates rule 3.

   Note: If you skip steps 5 and 6, the system applies default rules to all device tags except the ones that are defined in rules 1 and 2.

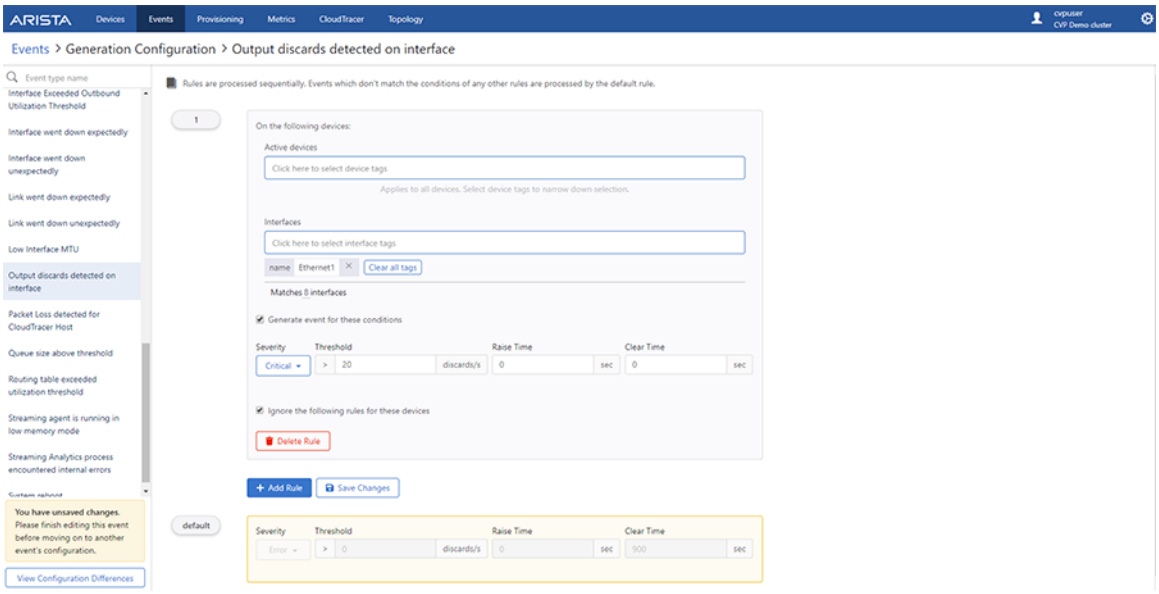

Figure 101: Disable All Events of the Selected Type with Exception

The system disables all events of the selected type except the ones that are defined in rules 1 and 2.

8.14.6 Acknowledging Events

Acknowledging an event confirms that you are aware of the corresponding event and its consequences. By default, acknowledged events are hidden and do not send alerts.

Perform the following steps to acknowledge an event:

1. Click the **Events** tab. The system displays the **Events** screen.
2. Select preferred event(s) in the side panel.
3. Click **Acknowledge *n*** in the upper right corner of the side panel.

Note: ***n*** represents the count of selected events.

The system displays the **Acknowledgment Event** window.

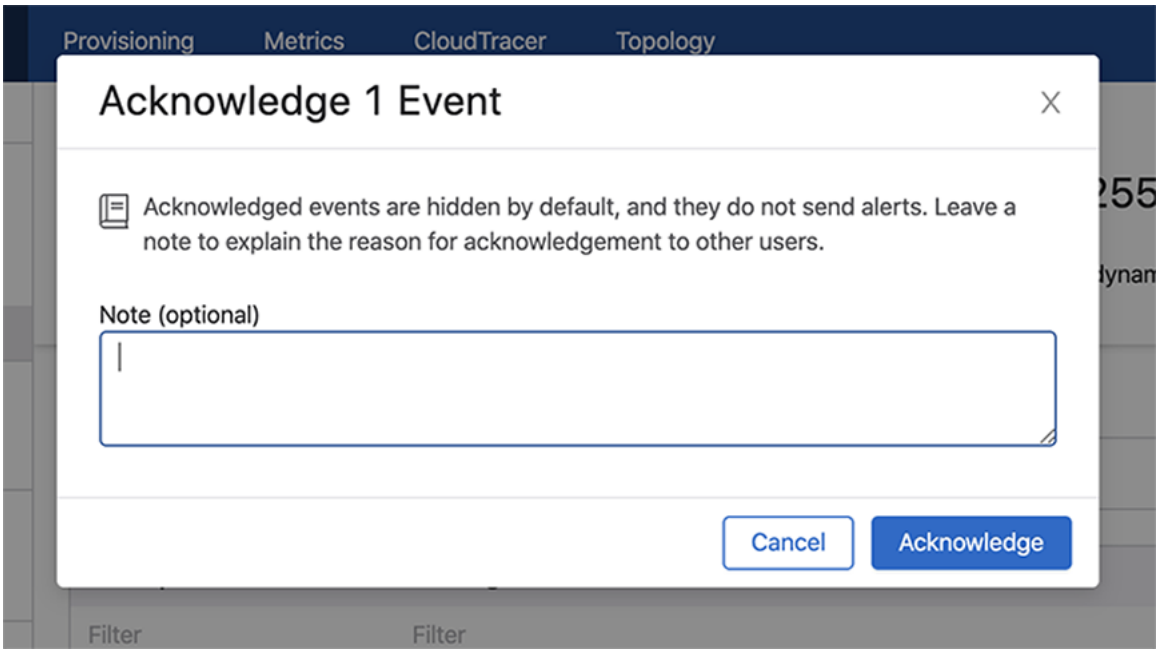

Figure 102: Acknowledgment Event Pop-Up

4. (Optional) Type a note for other users explaining the reason for the acknowledgment.
5. Click **Acknowledge *n* events** where ***n*** represents the count of selected events.

> **Note:** For acknowledged events, the system replaces the **Acknowledge** button with **Un-Acknowledge** button. To undo the acknowledgment activity, Click **Un-Acknowledge** in the side panel of the acknowledged event.

8.14.7 Configuring Notifications

The event alerting system sends notifications for CVP events as they alert operating platforms that you have set up. Once you have customized the topology view for your network, provide the required information to configure the monitoring of notifications.

Perform the following steps to configure event alerts:

1. Click the **Events** tab.
2. Click **Configure Notifications** at the upper right corner of the Events section. The system displays the Notification Configuration screen.
3. Configure the following entities:
   - Configuring Status
   - Configuring Platforms
   - Configuring Receivers
   - Configuring Rules
4. Click **Save** in the left pane

8.14.7.1 Configuring Status

The **Status** section configures monitoring the health of notification system.

Perform the following steps to configure the notification criteria:

1. Click **Status**. The system displays the **Status** screen.

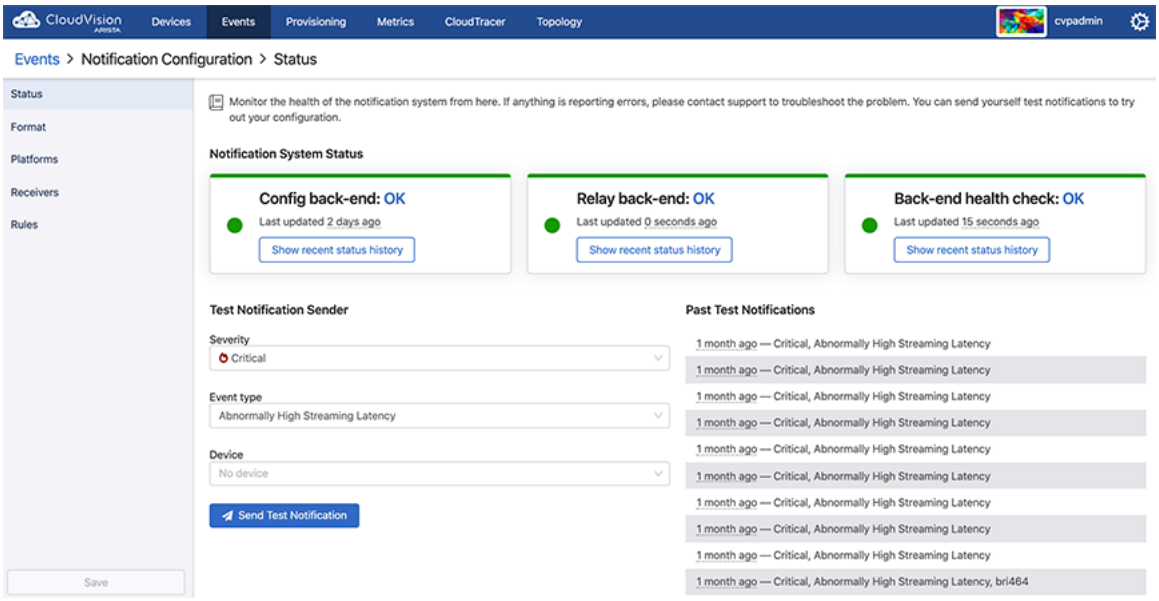

Figure 103: Status Screen of Notification Configuration

2. On the **Test Alert Sender** pane, provide the required criterion in **Severity**, **Event type**, and **Device** drop-down menus.
3. If required, click **Send Test Notification** to verify current configuration.

8.14.7.2 Configuring Platforms

The Platforms section specifies what platforms will receive notifications.

Perform the following steps to configure preferred platforms:

1. Click **Platforms**. The system displays the **Platforms** screen.

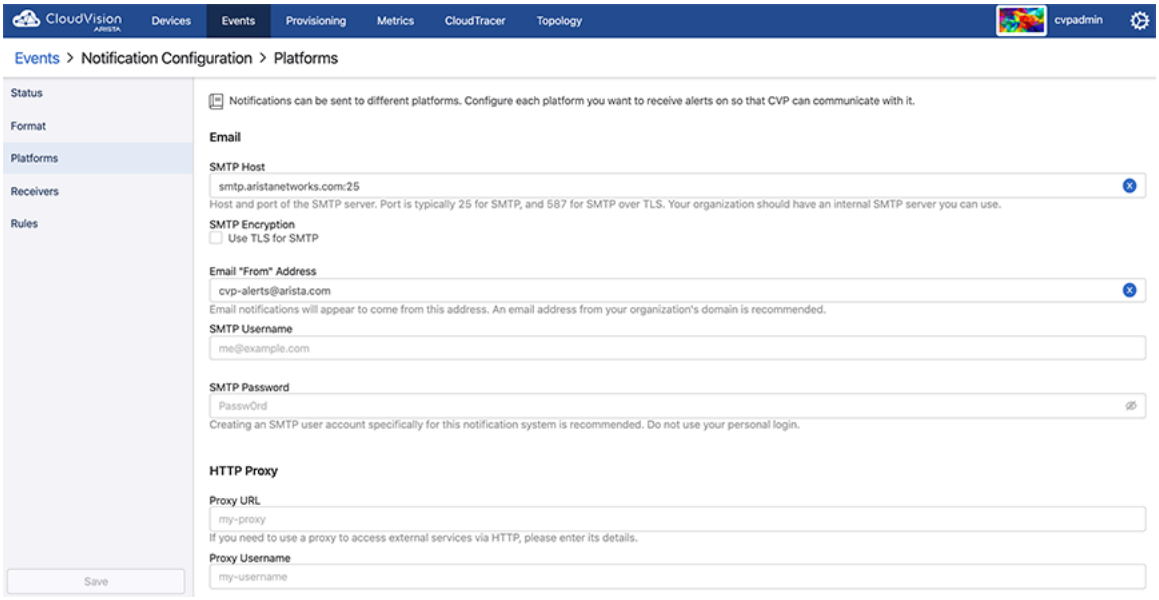

Figure 104: Platforms Screen of Notification Configuration

2. Configure any of the following platforms through which you prefer to receive notifications from CVP:
   - **Email**

     Provide the following information to receive email notifications:

- Type your SMTP server's hostname and port number separated by a colon in the **SMTP Hos**t field.

  Note: Typically, the port numbers of SMTP and SMTP over TLS are 25 and 587.

- Select the **Use TLS for SMT**P checkbox if you prefer to encrypt notifications received from and sent to the SMTP server.
- Type the email address that you prefer to display as a sender in the **Email "From" Address** field.

  Note: We recommend an email address with the domain of your organization.

- Type the username of your SMTP account in the **SMTP Username** field.
- Type the password of your SMTP account in the **SMTP Password** field.

- **Slack**

  Create a custom integration through the "Incoming WebHooks" Slack application and type the "Webhook URL" in the **Slack Webhook URL** field.

- **VictorOps**

  - In your **VictorOp**s settings, add a new alert integration for "Prometheus" and type the "Service API Key" in the **VictorOps API Key** field.
  - If required, type a custom API URL in the **VictorOps API URL** field.

- **PagerDuty**

  If required, type a custom API URL in the **PagerDuty URL** field.

- OpsGenie

  - Create an "API" integration for your OpsGenie team and type the API key in the **OpsGenie API Key** field.
  - If required, type a custom API URL in the **OpsGenie API URL** field.

- WeChat

  - Type your WeChat credentials in the **WeChat API Secret** field.
  - Type your WeChat corporate ID in the **WeChat Corporate ID** field.
  - If required, type a custom API URL in the **WeChat API URL** field.

8.14.7.3 Configuring Receivers

The Receivers section configures a receiver for each preferred team to send notifications and link receivers to notification platforms.

Perform the following steps to add new receivers:

1. Click **Receivers**. The system displays the Receivers screen.

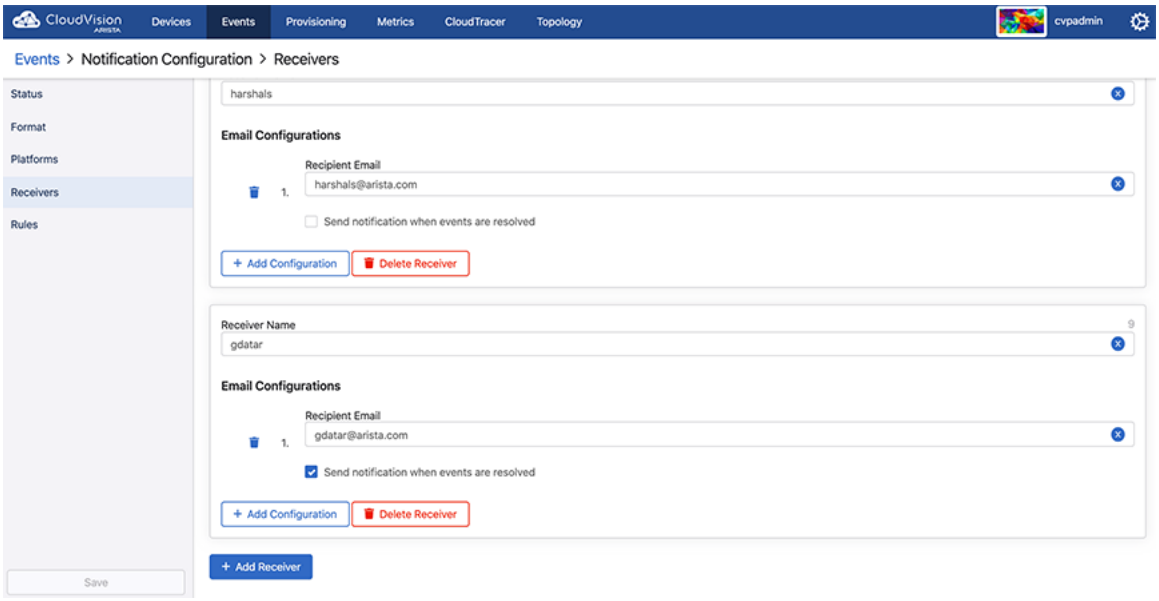

Figure 105: Receivers Screen of Notification Configuration

2. Click **Add Receivers** at the end of the screen.
3. Type receiver's name in the **Receiver Name** field.

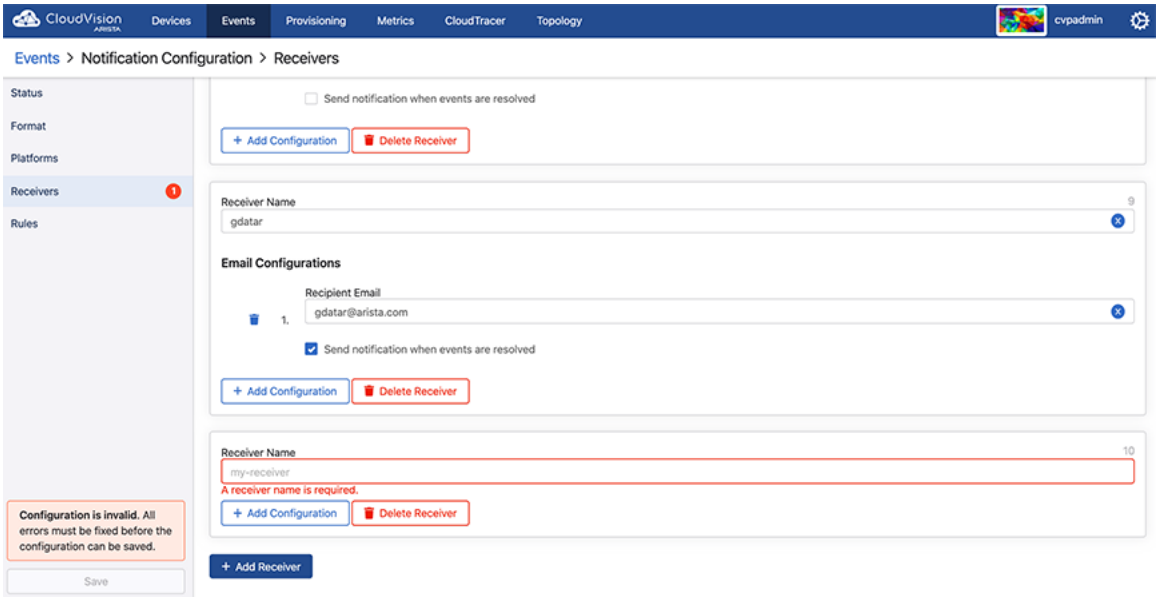

Figure 106: Add Receiver Pane

4. Click the **Add Configuration** drop-down menu.
5. Select any of the options in following table and provide the required information to link alert receivers with alerting platforms.

Table 13: Configuration Options

Configuration Options	Required Information
Add Email Configuration	• Type recipient's email address in the **Recipient Email** field. • If required, select the **Send alert when events are resolved** checkbox.

Add VictorOps Configuration	• Type a routing key in the **Routing Key** field. • If required, select the **Send alert when events are resolved** checkbox.
Add PagerDuty Configuration	• Type a routing key in the **Integration Key** field. • If required, select the **Send alert when events are resolved** checkbox.
Add OpsGenie Configuration	Select the **Send alert when events are resolved** checkbox.
Add Slack Configuration	• Type a channel in the **Channel** field. • If required, select the **Send alert when events are resolved** checkbox.
Add WeChat Configuration	Select the **Send alert when events are resolved** checkbox.
Add Pushover Configuration	• Type a recipient's user key in the **Recipient User Key** field. • Type a pushover API token in the **Application API Token** field. • If required, select the **Send alert when events are resolved** checkbox.
Add Webhook Configuration	• Type the URL where you prefer to post event alerts in the **Target URL** field. • If required, select the **Send alert when events are resolve**d checkbox

Note: Click the recycle bin icon at the right end of corresponding fields if you prefer to delete that configuration. Click **Delete Receiver** next to **Add Configuration** if you prefer to delete the corresponding receiver.

8.14.7.4 Configuring Rules

The Rules section customizes notifications that are sent to receivers.

Perform the following steps to add a new rule:

1. Click **Rules**. The system displays the Rules screen.

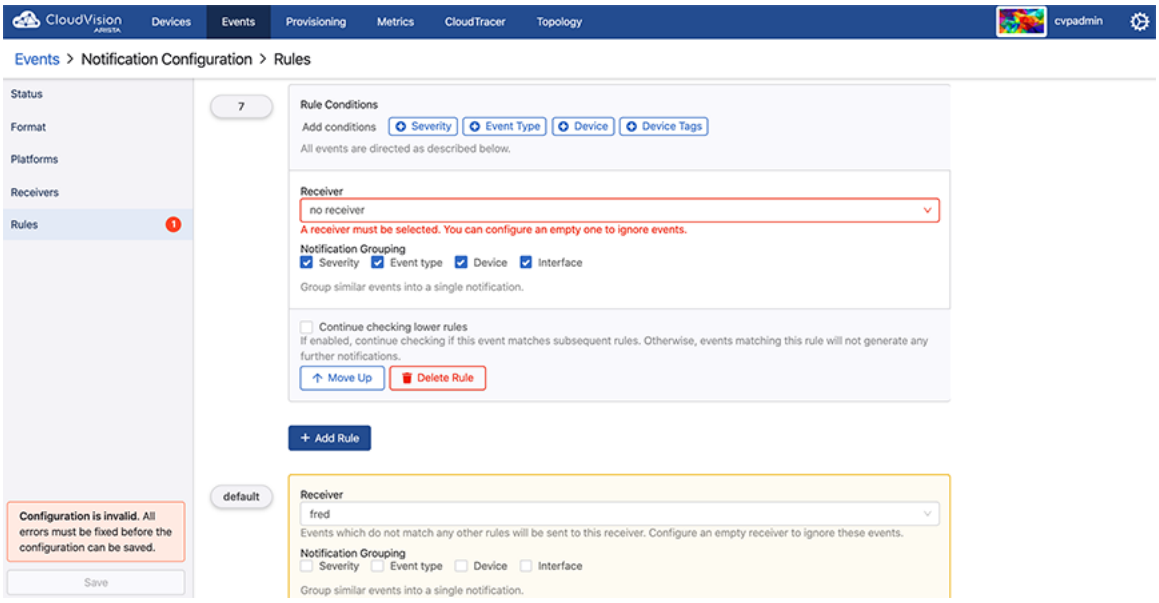

Figure 107: Rules Screen of Notification Configuration

2. Click **Add Rules**. A new Rules Conditions pane is displayed on the screen.

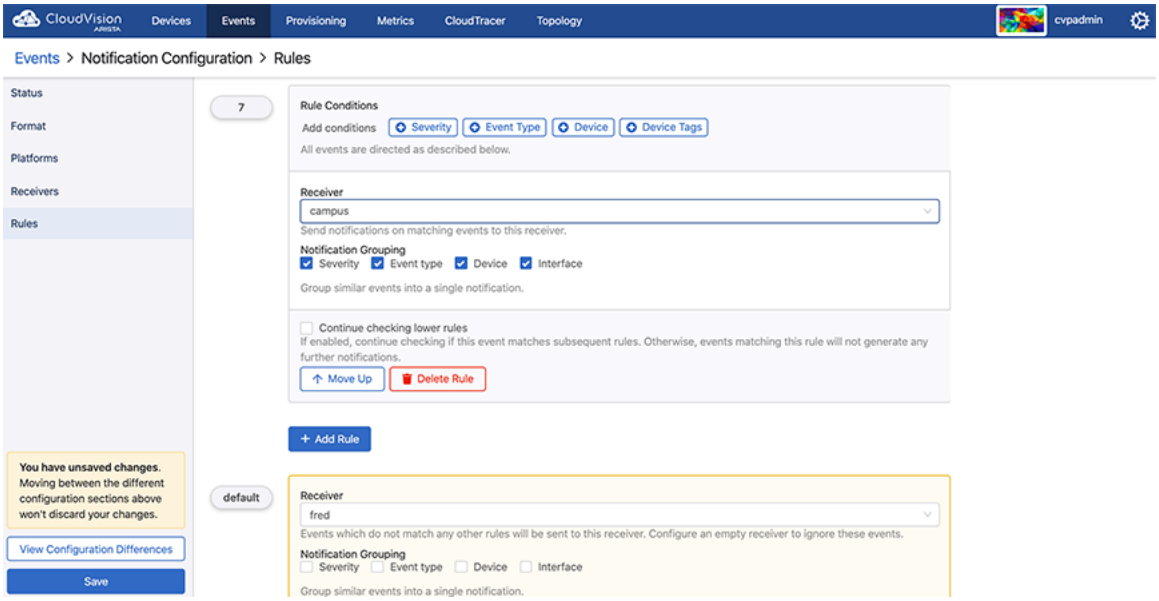

Figure 108: Rule Conditions Pane

3. Next to **Add Conditions**, click **Severity**, **Event Type**, **Device**, and **Device Tags** to provide the criteria that are used for monitoring the health of the alerting system.

   Note: Click **Remove** at the end of a field to delete that configuration.

4. Select the required receiver from the **Receiver** drop-down menu.
5. Select required checkboxes among Severity, Event Type, Device, and Interface to group similar events into a single alert.
6. Select the **Continue checking lower rules** checkbox to continue checking for alerts if this event matches subsequent rules.
7. Click **Move up** if you prefer to move this rule up in the priority list.

Note: Rules are processed sequentially. The default rule is applied only when an event does not match any other rules. Click **Delete rule** to delete the corresponding rule. Click **Move down** in configured rules to move the corresponding rule down in the priority list.

8.15 Troubleshooting

A number of commands are provided with the Telemetry platform that you can use to troubleshoot the Telemetry platform components. The types of troubleshooting you can perform using the Telemetry platform commands are:

- General Troubleshooting
- Troubleshooting the NetDB State Streaming Agent
- Checking the Status of the Ingest Port

8.15.1 General Troubleshooting

Telemetry commands are provided that enable you to troubleshoot the Telemetry platform components. By default, debug log files are available for all of the Telemetry platform components, which you can view using Telemetry commands. You can also use standard CVP commands to check the status of Telemetry components and applications.

8.15.1.1 Viewing Debug Log Files

You can view debug log files for all platform components in a single log file, or for a particular platform component.

Note: To use the commands, you must login as **cvp** user. You must also login as **cvp** user to execute `su cvp`.

To view debug log files for all platform components in a single log file

Use the `cvpi logs all` command.

To view the location of debug log files for a particular platform component

Use the `cvpi logs <component>` command.

You must specify the component using the name of the component as it is specified in the component's yaml file definition.

To create a zip archive (.tgz) containing debugging information

Use the `cvpi debug` command.

This command creates a .tgz archive on each CVP node that contains debugging information. The archive is automatically saved to the /data/debug directory on each node. Files need to be collected manually.

8.15.1.2 Checking CVPI Status

You can use commands to check status of the Telemetry components and applications, and to check the status of the entire CVP environment.

To check the status of CVPI

Use the `cvpi status all` command.

This command checks the status of CVPI, including the Telemetry components and applications.

To check the status of CVP environment

Use the `cvpi check all` command.

This command runs a check to ensure that the CVP environment is setup correctly. In a multi-node setup, it checks to make sure that the nodes can communicate with to each other and have the same environments and configuration.

8.15.2 Troubleshooting the NetDB State Streaming Agent

The Telemetry platform component provides commands you can use to troubleshoot issues you may encounter with the installation or performance of the NetDB State Streaming Agent.

The commands enable you to:

- Inspect the agent's configuration
- Restart the agent
- View the agent's logs

8.15.2.2 Restart the agent

Run the following commands to toggle the shutdown attribute:

```
switch (config-daemon-TerminAttr)# shutdown
switch (config-daemon-TerminAttr)# no shutdown
```

8.15.3 Checking the Status of the Ingest Port

The Telemetry platform automatically blocks the ingest port for the entire CVP cluster if the disk usage on any node of the cluster exceeds 90%. This feature prevents the potential for telemetry data to consume too much disk space in the CVP cluster.

You can easily check to see if the ingest port is blocked using the cvpi status ingest-port command.

Example

```
[cvp@cvp109 bin]$ cvpi status ingest-port
[ingest-port:status] Executing...
[ingest-port:status] FAILED

COMPONENT    ACTION   NODE    STATUS     ERROR


ingest-port   status    primary  NOT RUNNING  command: Error running  '/
cvpi/bin/ingest-port.sh status'...
ingest-port       status       secondary     NOT RUNNING     command:
 Error running '/cvpi/bin/ingest-port.sh status': exit status 1
ingest-port       status       tertiary      NOT RUNNING     command:
 Error running '/cvpi/bin/ingest-port.sh status': exit status 1
[cvp@cvp109 bin]$
```

Chapter 9

Device Comparison Application

To gain valuable insights into the state of your devices, such as state changes and comparison with another device, you can manage your inventory for real-time status updates.

The device comparison application gives information about the configuration running on the devices, the VXLAN table, MAC addresses of the devices, IPv4 and IPv6 routing tables, etc.

9.1 Comparison Dashboard

The Comparison Dashboard from the Device tab explores the difference between devices or changes that happened to devices over time. You can compare devices in the following categories:

- Two devices: Two devices at current time with live updates
- Two times: The state of a single device at two chosen times
- Advanced: Two devices at two chosen times

9.1.1 Accessing the Comparison Browser Screen

You can access the Cloud Vision Telemetry Browser screen directly from CVP by completing the following steps. Open your browser.

1. Point your browser to the CVP IP address or hostname.
2. Login to CVP. The CVP Home screen appears.
3. Click **Devices**.
4. Click **Comparison**.

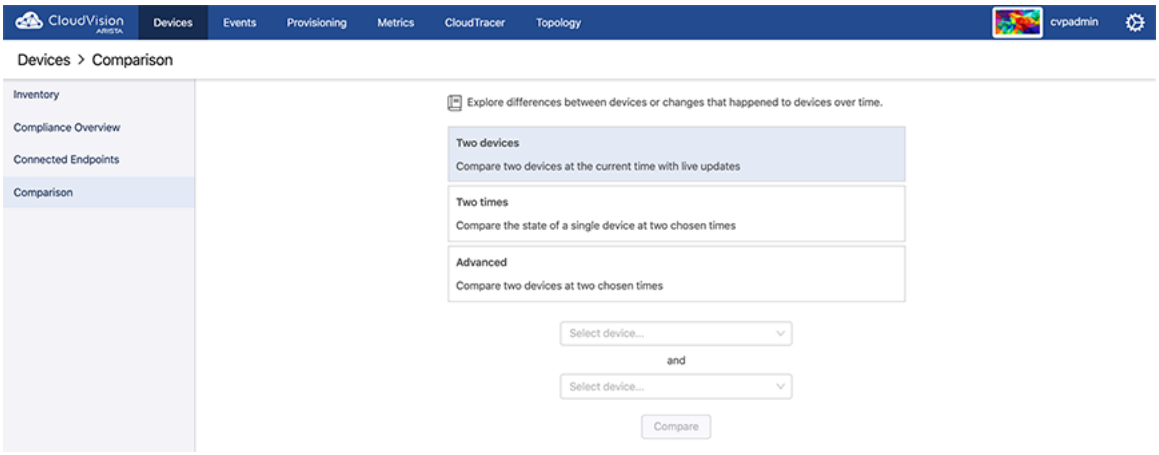

Figure 109: Start page for comparison of devices

For a particular device with two chosen times, select the Two times option.

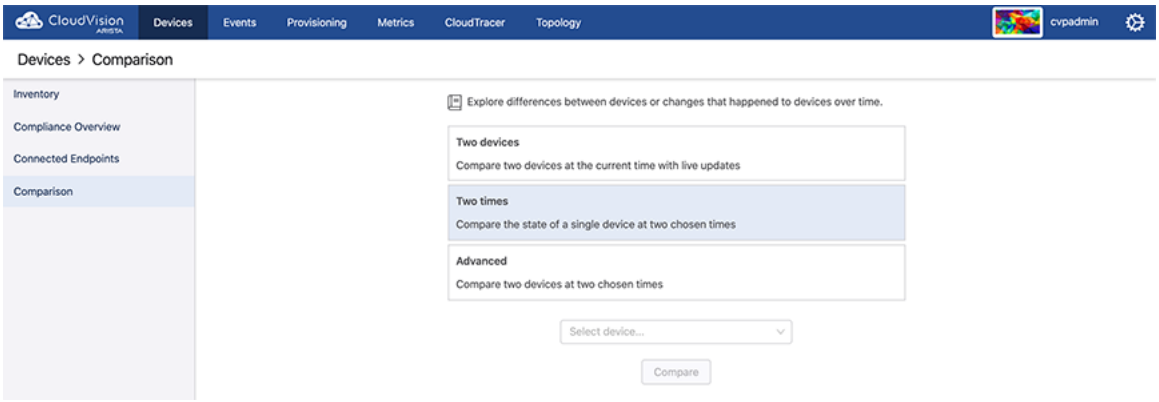

Figure 110: Comparison of device at two chosen times

Comparing two devices at two chosen times, select the Advanced option:

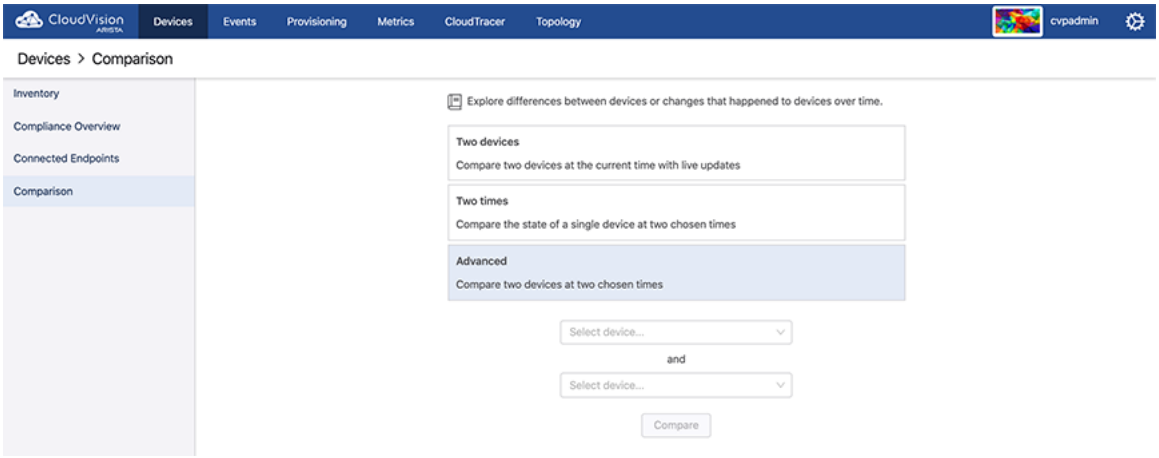

Figure 111: Comparison of device advanced

9.2 Running Configuration

To compare the data for the Running configuration for different devices, select **Running Config**. You have an option for current time comparison or chosen times comparison.

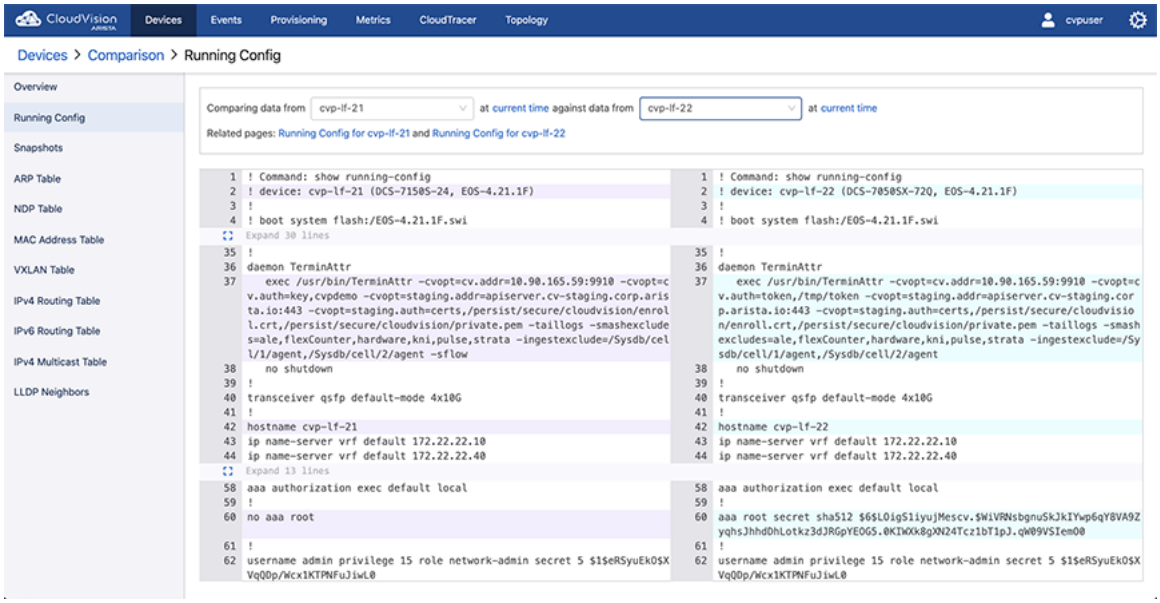

Figure 112: Comparison of Running configuration for two devices

- Supported Snapshots

9.2.1 Supported Snapshots

All Snapshots give the list of snapshots, its capture time and its last executioner in the following figure.

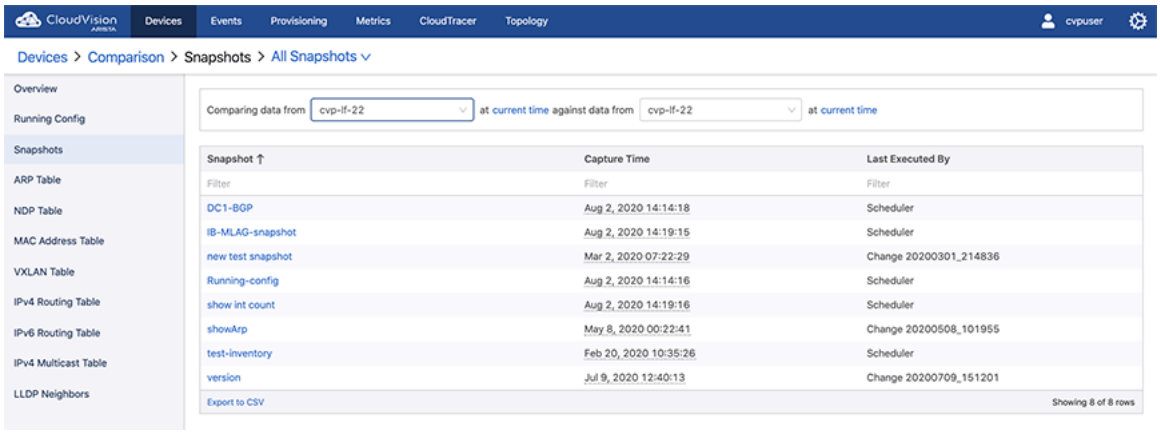

Figure 113: All Snapshots options

9.3 Snapshots

On the CloudVision portal, navigate to **Devices > Comparison** to **Snapshots** to view the snapshot for the device.

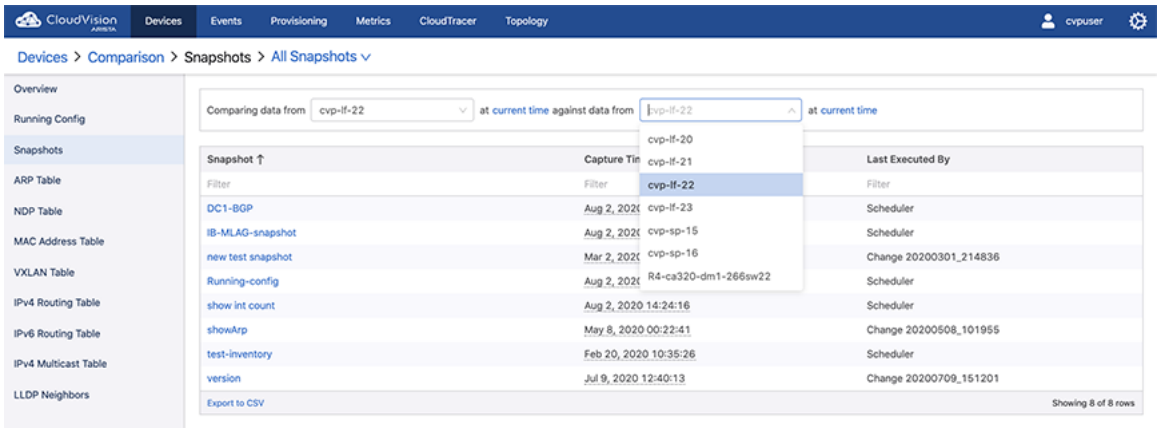

Figure 114: Comparing snapshots

The screen provides the following functionalities:

- All Snapshots: Displays all current snapshots options
- Snapshots Filter: Select the required snapshot filter

9.4 ARP Table

On the Cloud Vision portal, navigate to **Devices > Comparison** to ARP Table to view the information about ARP. Arista's device comparison platform for ARP table compares data between two devices at the same time and at different time settings.

You can compare the following:

- Device's IP Address
- Device's MAC Address
- Interface

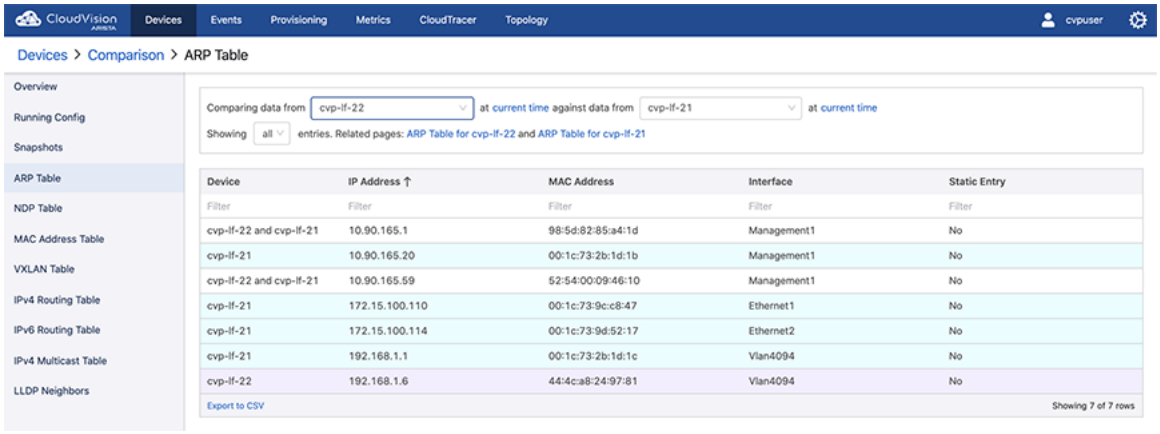

Figure 115: Comparing ARP table

9.5 Comparing NDP Table

On the Cloud Vision portal, navigate to **Devices > Comparison** to NDP Table to view the information about NDP. Arista's device comparison platform for NDP table compares data between two devices at the same time and at different time settings.

The components of the comparison are as follows:

- Device's IP Address
- Device's MAC Address
- Interface
- Static entry

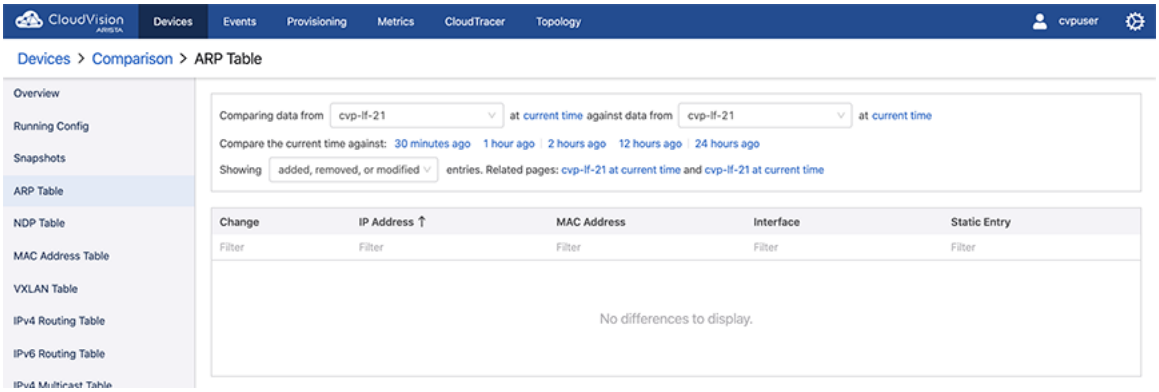

Figure 116: Comparing NDP table

You can compare the status at the current time against the following times:

- 30 minutes
- 1 hour
- 2 hours
- 12 hours and
- 24 hours ago.

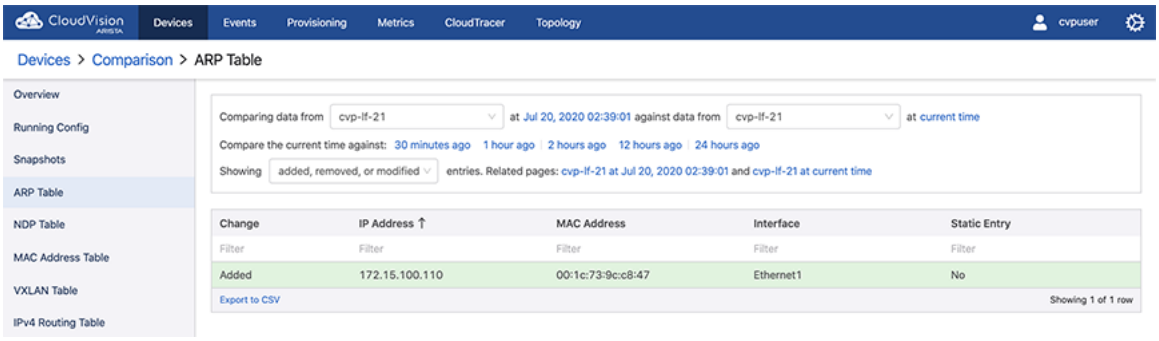

Figure 117: Comparing same device for NDP table for different times

9.6 MAC Address Table

On the Cloud Vision portal, navigate to **Devices > Comparison** to MAC AddressTable to view the information about MAC addresses for the devices. Arista's device comparison platform for MAC Address table compares data between two devices at the same time and at different time settings.

The components of the comparison are as follows:

- VLAN
- Device's MAC Address
- Type of the VLAN
- Port
- Number of moves on the Port
- Timing for last movement

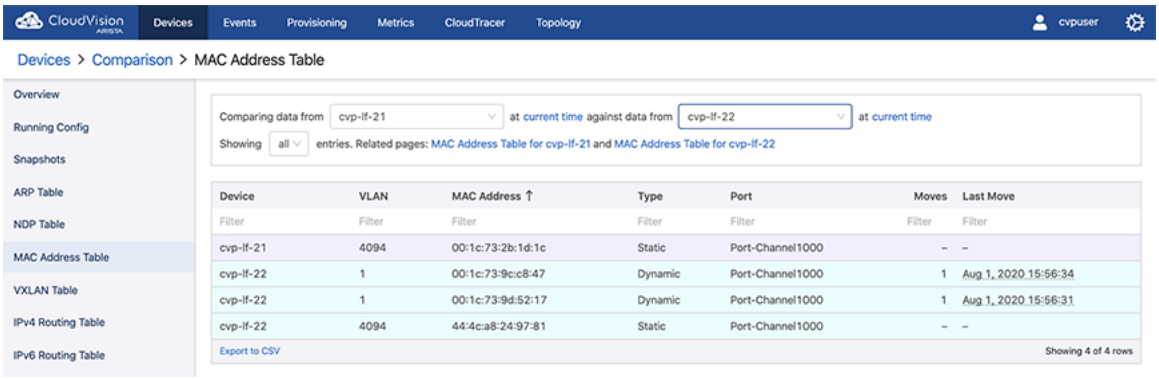

Figure 118: Comparing MAC Address table for current time for two devices

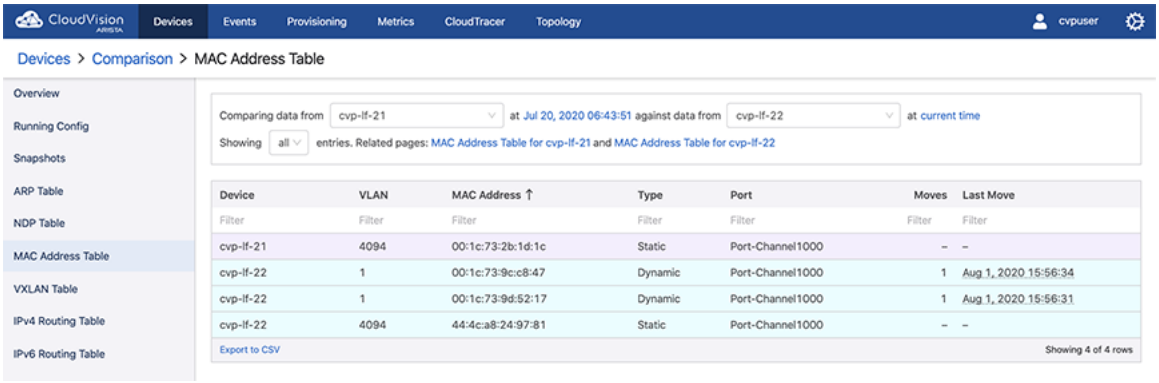

Figure 119: Comparing MAC Address table for different times for two devices

You can compare the status at the current time against the following times:

- 30 minutes
- 1 hour
- 2 hours
- 12 hours and
- 24 hours ago.

Status is shown by added, removed and modified entries.

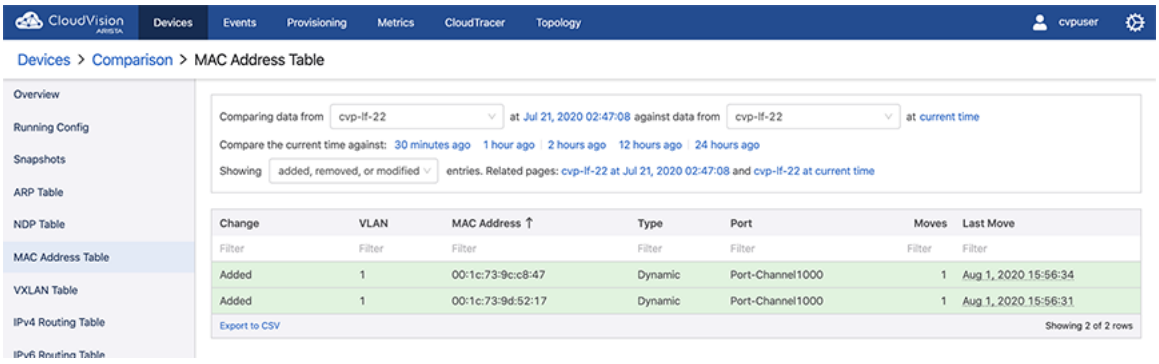

Figure 120: Comparing same device for different times and status

To show all entries for the devices, Click ALL.

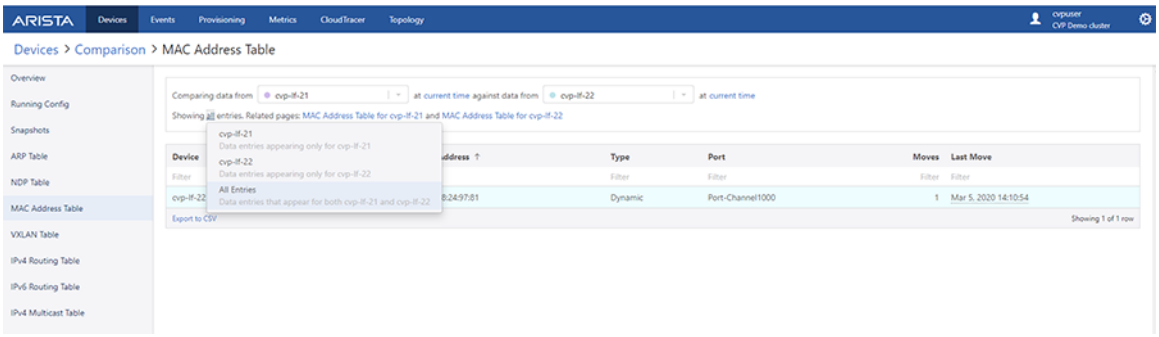

Figure 121: Showing all entries for the Devices for MAC Address table

9.7 VXLAN Table

On the Cloud Vision portal, navigate to **Devices > Comparison** to VXLAN Table to view the information about MAC addresses for the devices.

The components of the comparison are as follows:

- VLAN VNIs
- VXLAN MAC Address

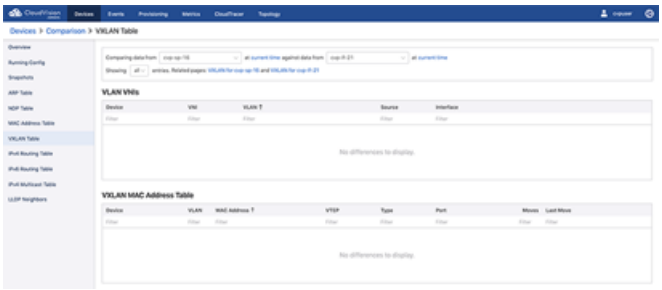

Figure 122: Comparing VXLAN table for current time for two devices

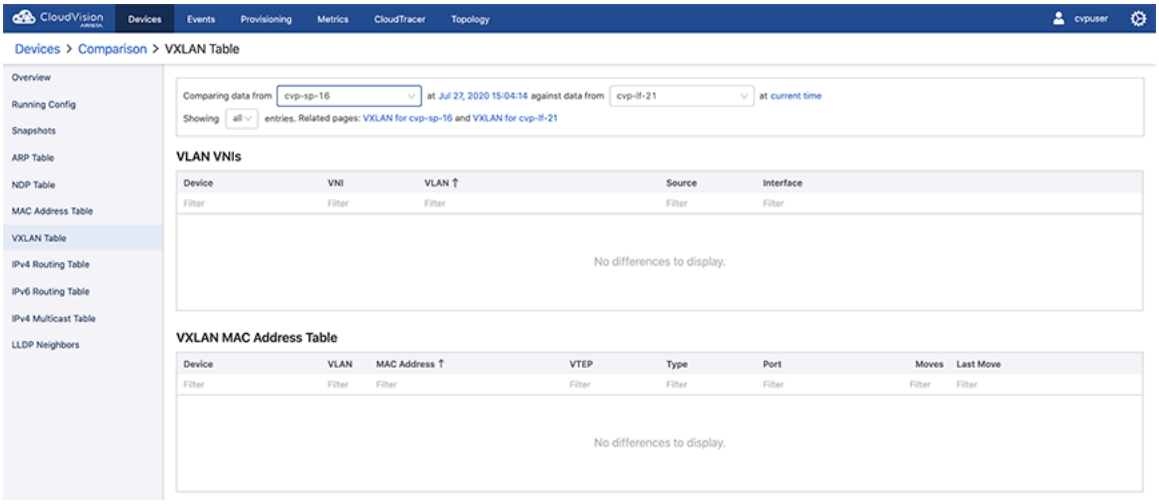

Figure 123: Comparing VXLAN table for different times for two devices

You can compare the status at the current time against the following times:

- 30 minutes

- 1 hour
- 2 hours
- 12 hours and
- 24 hours ago.

Status is shown by added, removed and modified entries.

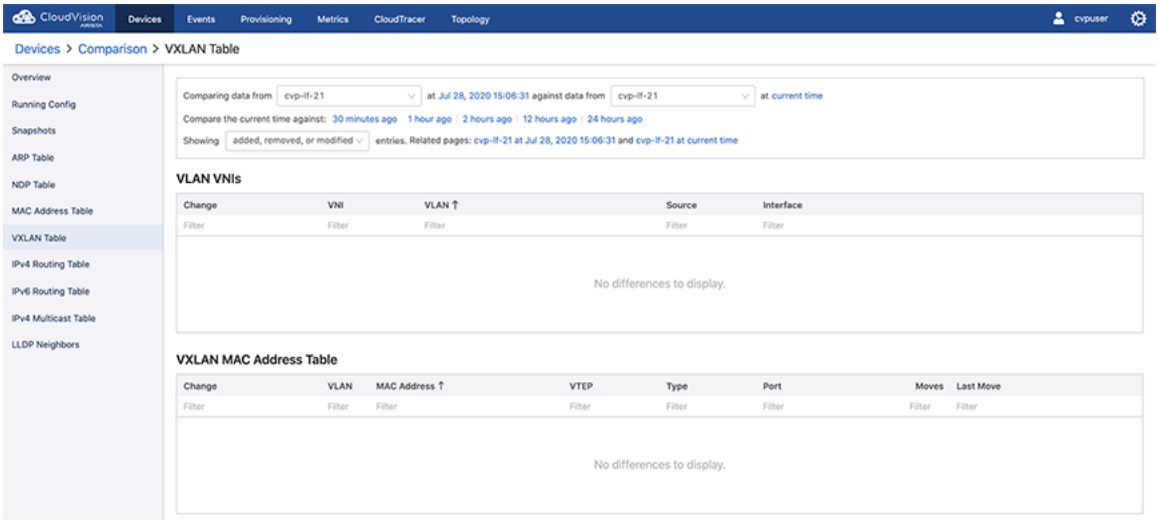

Figure 124: Comparing same device for different times and status

To show all entries for the devices, Click ALL.

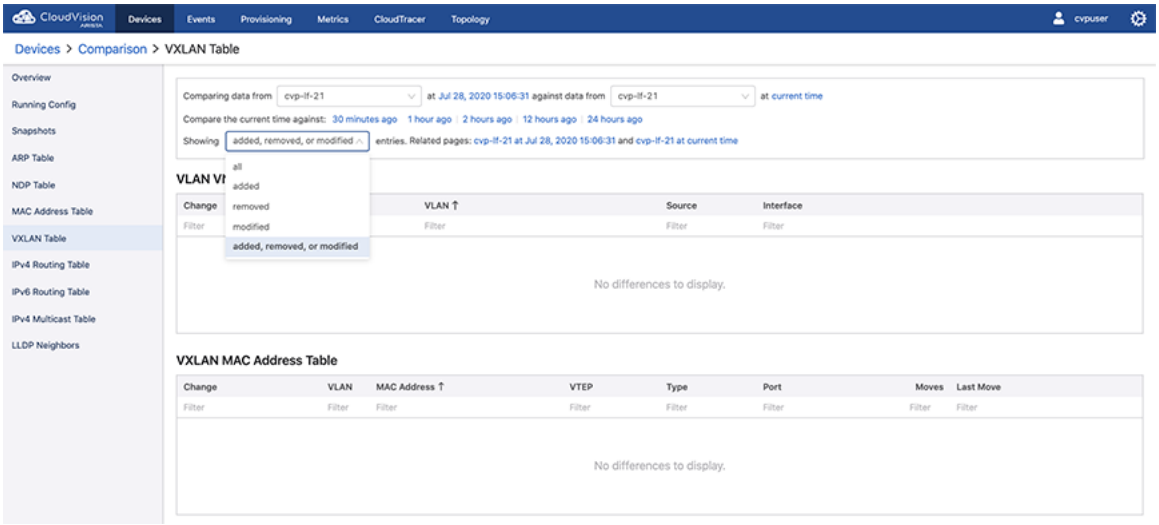

Figure 125: Showing all entries for the Devices for VXLAN table

9.8 Viewing Device IPv4 Routing Table

From the Comparison screen, you can quickly drill down to view details about IPv4 Routing from different devices. In tabular view, click the device names to compare the corresponding device details.

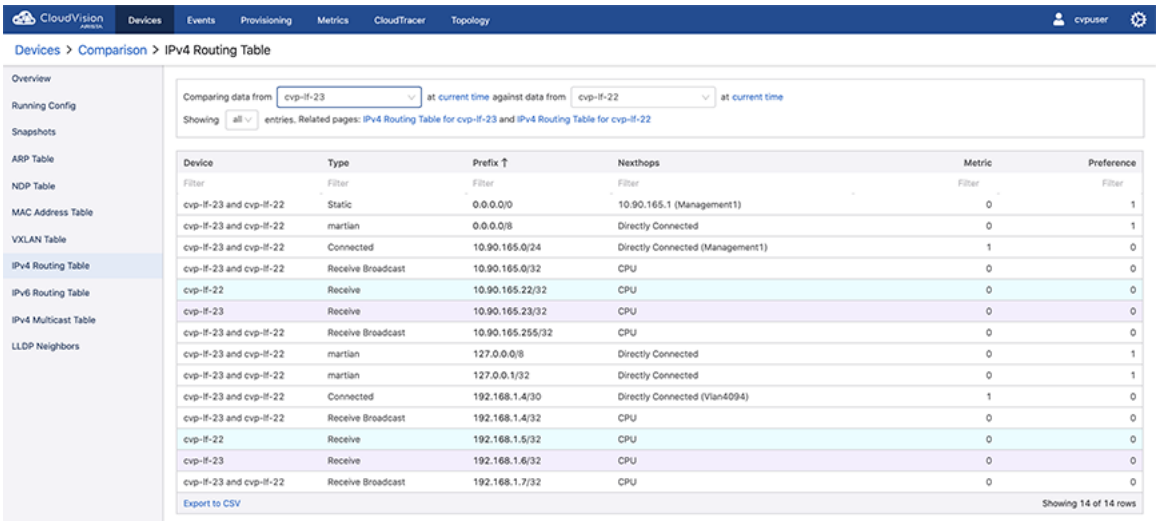

Figure 126: Comparing IPv4 routing table for different devices

The screen refreshes to show the status, IP address and functions it does for Nexthop. Status is generally shown by Static, Martian, Connected, Receive and Receive Broadcast.

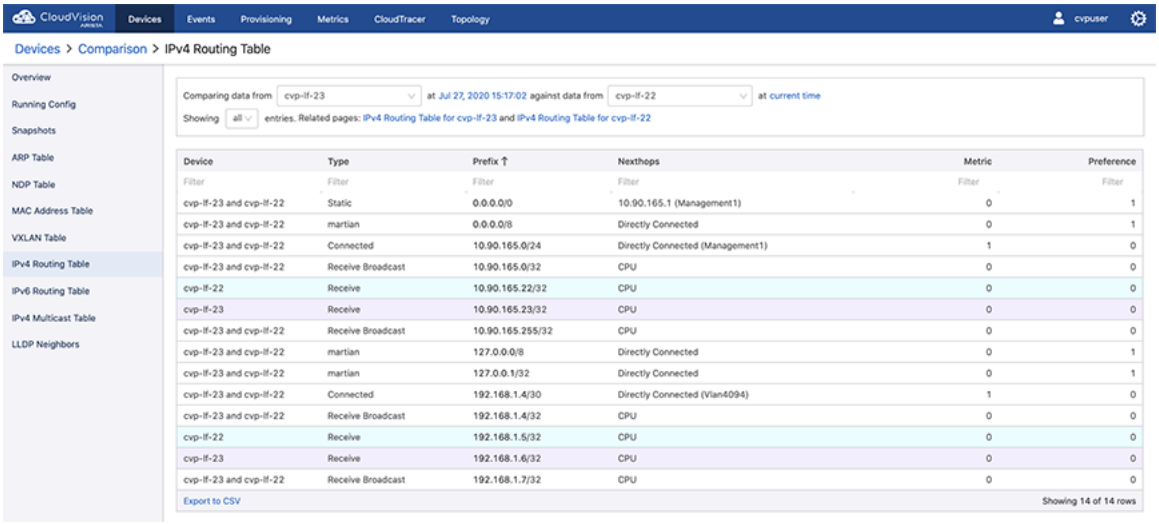

Figure 127: Comparing IPv4 Routing table for different times for two devices

You can compare the status at the current time against the following times:

- 30 minutes
- 1 hour
- 2 hours
- 12 hours and
- 24 hours ago.

Status is shown by added, removed and modified entries.

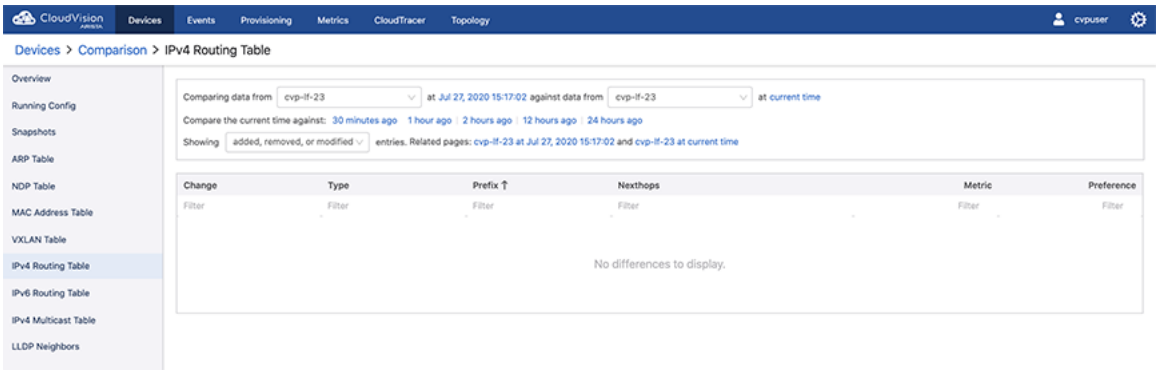

Figure 128: Comparing same device for different times and status

9.9 Viewing Device IPv6 Routing Table

From the Comparison screen, you can quickly drill down to view details about IPv6 Routing from different devices. In tabular view, click the device names to compare the corresponding device details.

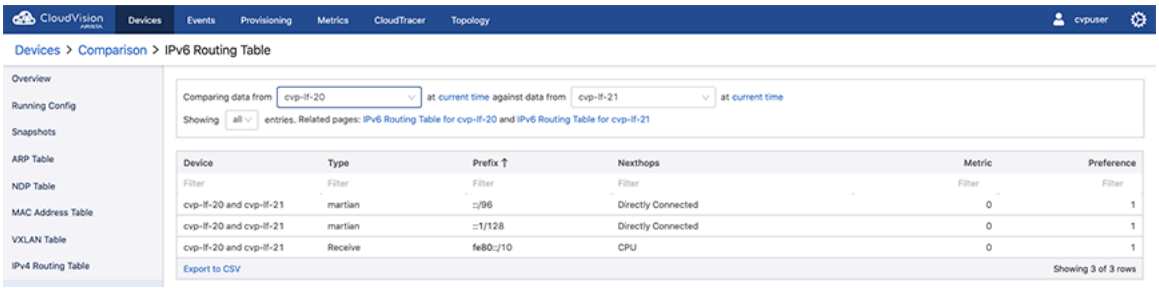

Figure 129: Comparing IPv6 routing table for different devices

The screen refreshes to show the status, IP address and functions it does for Nexthop. Status is generally shown by Static, Martian, Connected, Receive and Receive Broadcast.

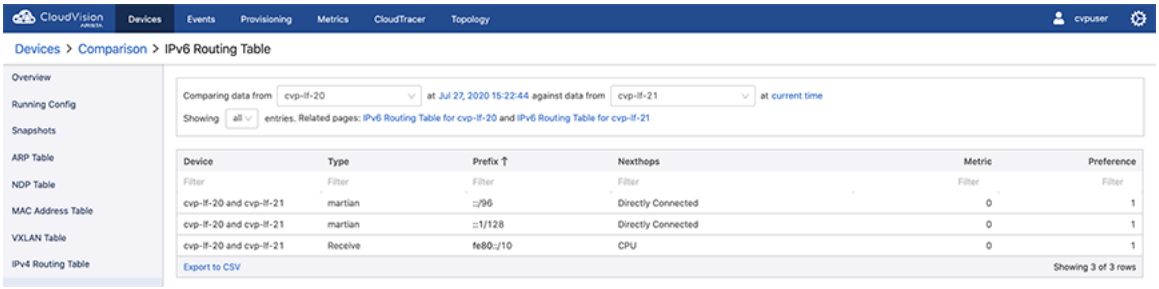

Figure 130: Comparing IPv6 Routing table for different times for two devices

You can compare the status at the current time against the following times:

- 30 minutes
- 1 hour
- 2 hours
- 12 hours and
- 24 hours ago.

Status is shown by added, removed and modified entries.

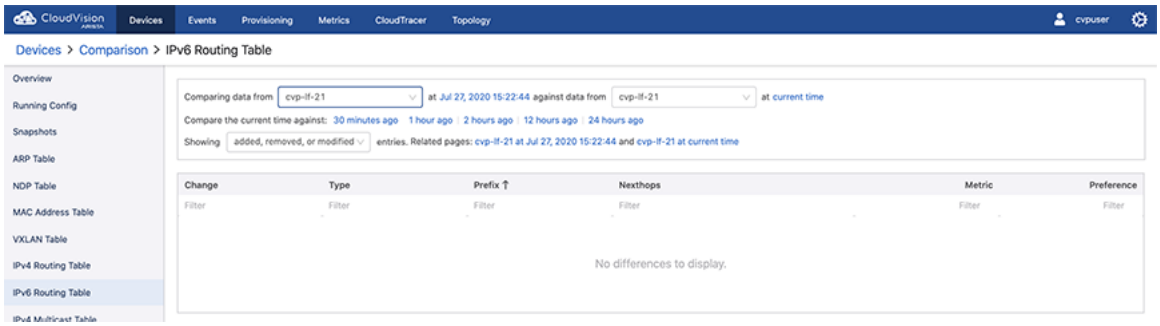

Figure 131: Comparing same device for different times and status

9.10 Comparing IPv4 Multicast Table

On the Cloud Vision portal, navigate to **Devices > Comparison to IPv4 Multicast Table** to view the information about Multicast. Arista's device comparison platform for IPv4 Multicast table compares data between two devices at the same time and at different time settings.

The components of the comparison are as follows:

- Sparse Mode PIM
- Static

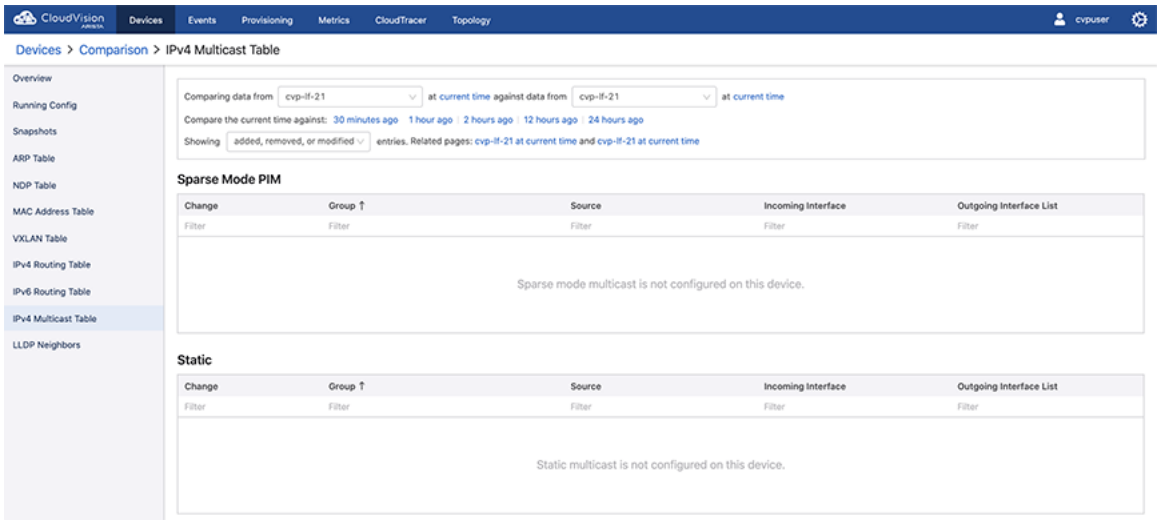

Figure 132: Comparing IPv4 Multicast table

You can compare the status at the current time against the following times:

- 30 minutes
- 1 hour
- 2 hours
- 12 hours and
- 24 hours ago.

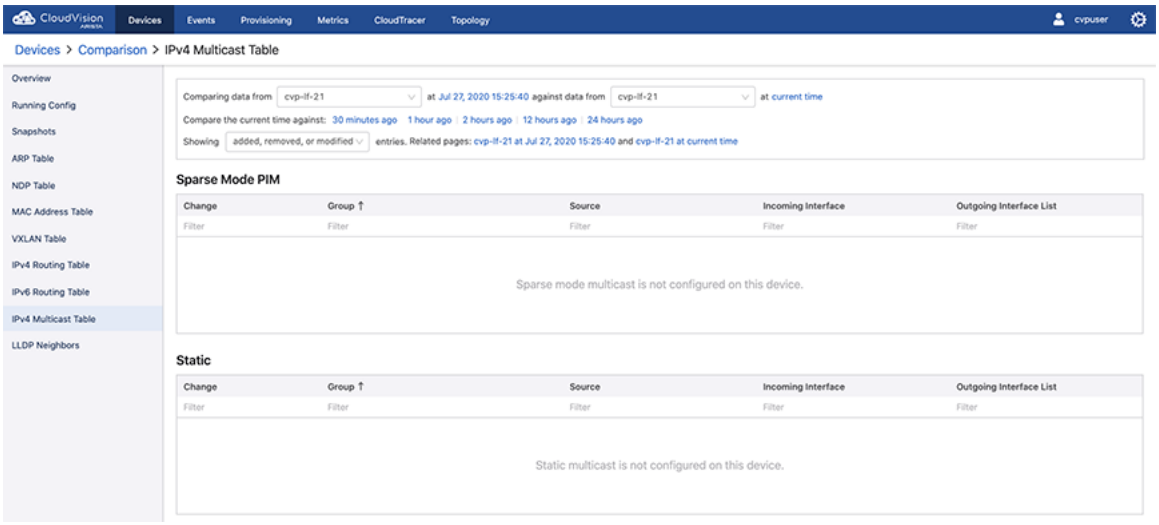

Figure 133: Comparing same device for IPv4 Multicast table for different times

Chapter 10

Network Compliance (CVP)

The **Compliance Dashboard** from the **Inventory** tab presents the number of devices and their compliance status. The three categories shown are:

- Bug Exposure
- Security Advisories
- Configuration and Software Image

Sections in this chapter include:

10.1 Compliance Dashboard

The **Compliance Dashboard** provides a real-time summary view of image, configuration and security compliance for all managed devices. The assessment uses bug details published on https://www.arista.com and leverages the network wide database to compute the exposure based on hardware and software versions. The CVP *2020.2.0* release comes packaged with a file named ‘AlertBase.json’ which contains information about software defects and security vulnerabilities. See the figure below.

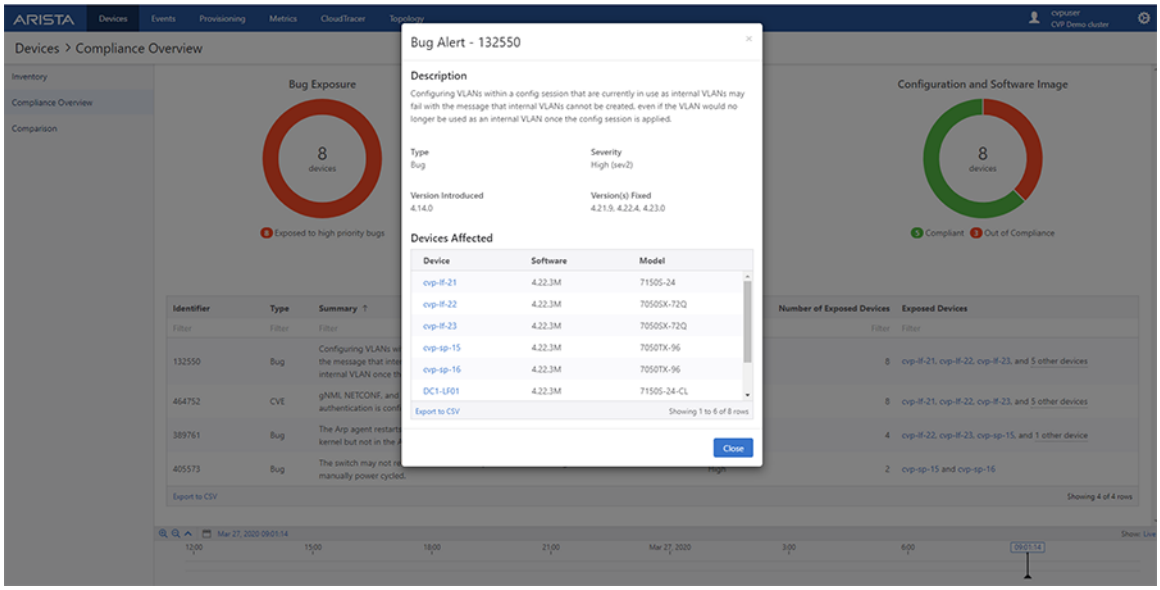

Figure 134: Compliance Dashboard

The Compliance Dashboard screen displays graphical and tabular presentation of bugs alerts.

Note: You can filter bug alerts using **All Alerts**, **Unacknowleged Alerts**, and **Acknowledged Alerts** dropdown options available next to breadcrumbs.

The compliance dashboard table consists of **Bug Alerts** and **Device Configuration** tabs.

Bug Alerts

The **Bug Alerts** tab provides the following information:

- **Identifier**: Bug number for issues tracked.

  Note: The checkmark next to identifier ID signifies acknowledged bugs.

- **Type**: Identifies the type of bug. Security vulnerabilities are tracked by type **CVE**. Software defects are tracked by type **Bug**. This field can be used to filter on either of these types.
- **Summary**: Provides a description of the software defect/security vulnerability.
- **Severity**: Calls out the severity of the software defect.
- **Device Count**: Lists the number of devices impacted by the tracked issue.

  Note:

  - If a device is acknowledged in tracked issues, this count is decreased by one.
  - If the bug is acknowledged, CVP displays zero.
  - Unacknowledged actions undo these results.

- **Exposed Devices**: Lists the names of devices impacted by the software defect or security vulnerability.

  Note:

  - If a device is acknowledged in tracked issues, CVP does not list its name.
  - If a bug is acknowledged, CVP displays **None**.
  - Unacknowledged actions undo these results.
  - CVP generates events for CVE bugs that are exposed on device(s). These events last until the bug either is resolved on the device or is acknowledged.

Click the listed bug alert to view more details from the corresponding **Bug Alert -** *Identifier ID* pop-window. See the figure below.

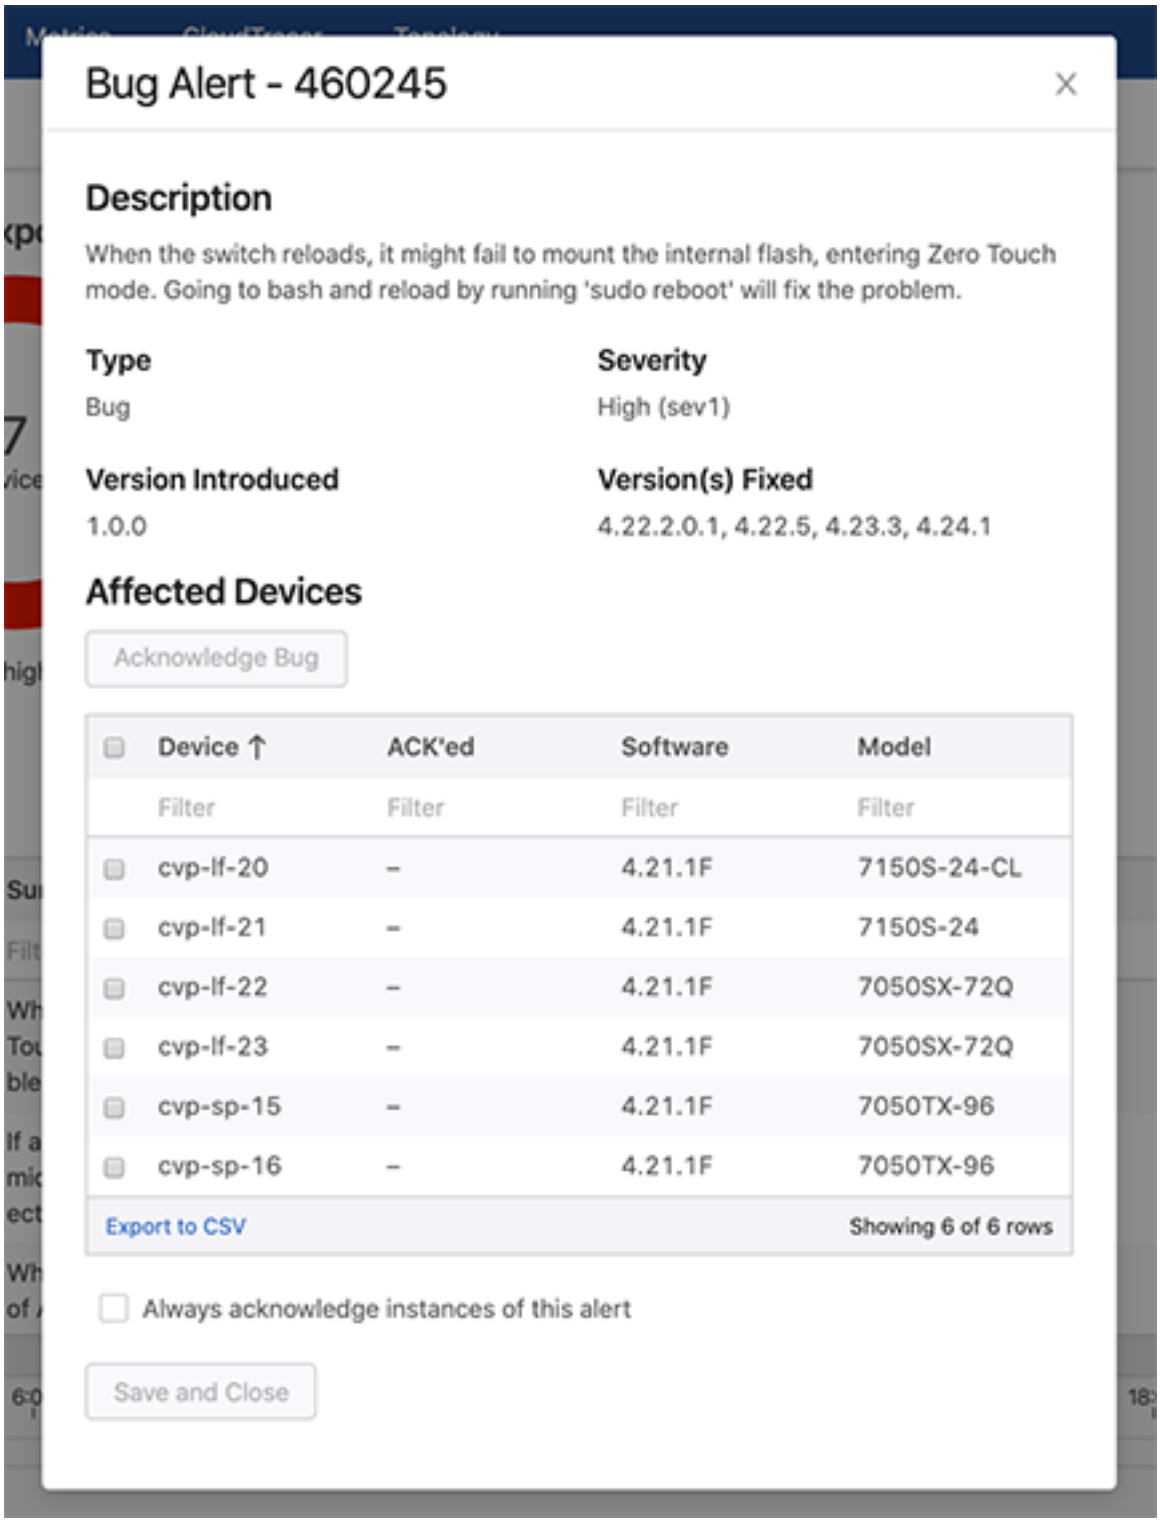

Figure 135: Bug Alert Pop-Up Window

You can fix listed bugs through one of the following ways:

- Upgrading your device to versions mentioned under **Version(s) Fixed**
- Installing the hotfix available at https://www.arista.com/en/support/advisories-notices as either a part of an image bundle or directly using the EOS CLI.

Note: You can search for hotfixes via identifier IDs.

Click the **Acknowledge Bug on *n* Device(s) and Close** button to hide the corresponding bug from bug info in selected devices.

Note:

- ***n*** presents the count of selected devices.
- (Optional) Provide reasons for acknowledgement in the text box.
- To undo the acknowledgement, reopen the bug to select acknowledged devices and click the **Unacknowledge Bug on *n* Device(s) and Close** button.

To acknowledge a bug for all current and future devices, select **Always acknowledge instances of this alert** checkbox and click **Save and Close** button.

Note:

- (Optional) Provide reasons for acknowledgement in the text box.
- To undo the acknowledgement, reopen the bug, unselect the checkbox, and click **Save and Close**.

The list of software defects and security vulnerabilities affecting a device are also available in the device view under the Compliance section.

Note: A checkmark is displayed next to an Identifier ID when either the bug is acknowledged or the current device is acknowledged for the correpoding bug.

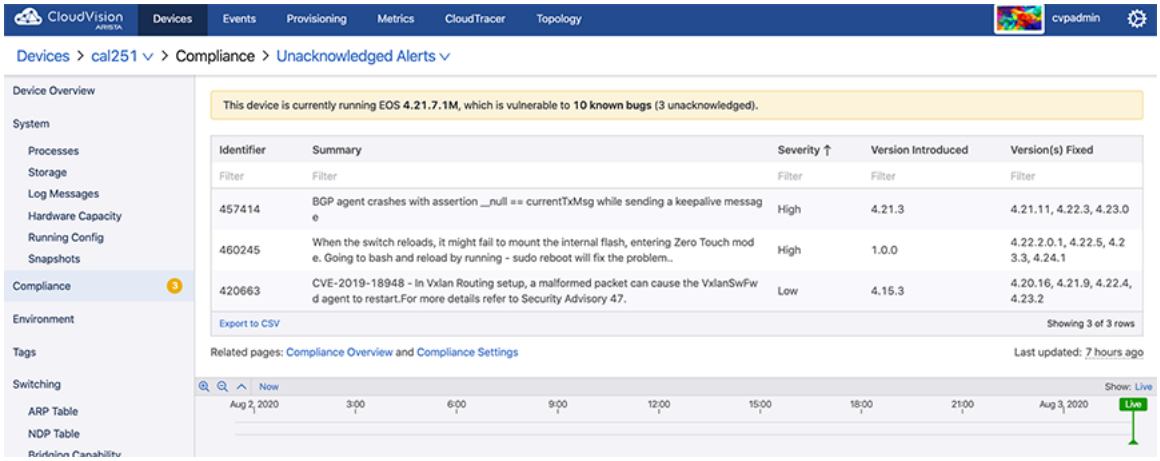

Figure 136: Compliance Section Showing Status of Bugs

Device Configuration

The **Device Configuration** tab displays the following information:

- **Device** - Lists the hostnames of devices.

  Note: Clicking on a device name opens the **Running Configuration** screen.

- **Status** - Displays the device status on configuration compliance.

  Note: CVP tracks out of sync status for configuration, image, and extensions.

- **Last Compliance Check** - Displays the timestamp of last compliance check.

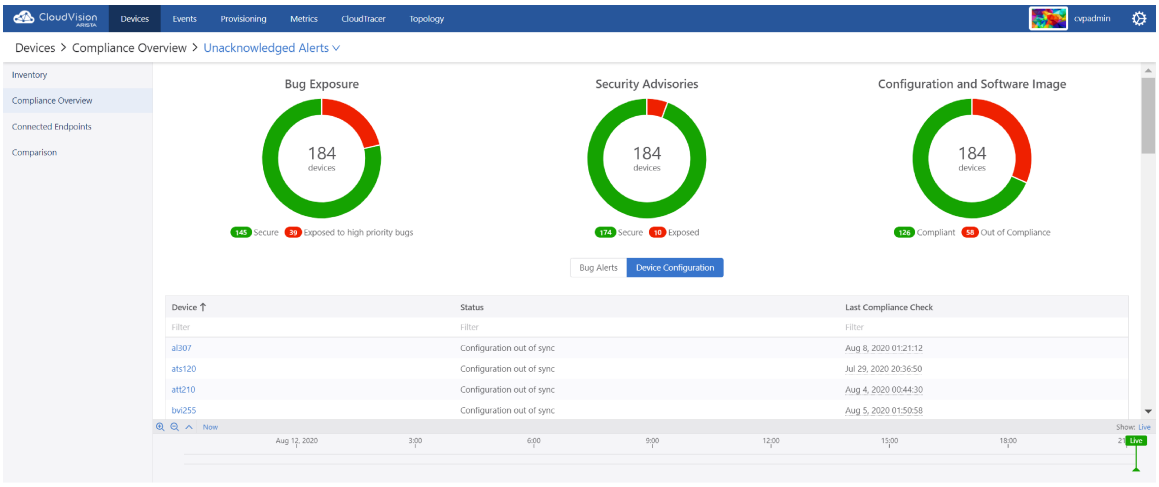

Figure 137: Device Configuration Tab

10.2 Print Compliance Dashboard

Perform the following steps to print the Compliance dashboard:

1. Select **Print** from the browser menu.
   CVP displays the Print pop-up window. See the figure below.

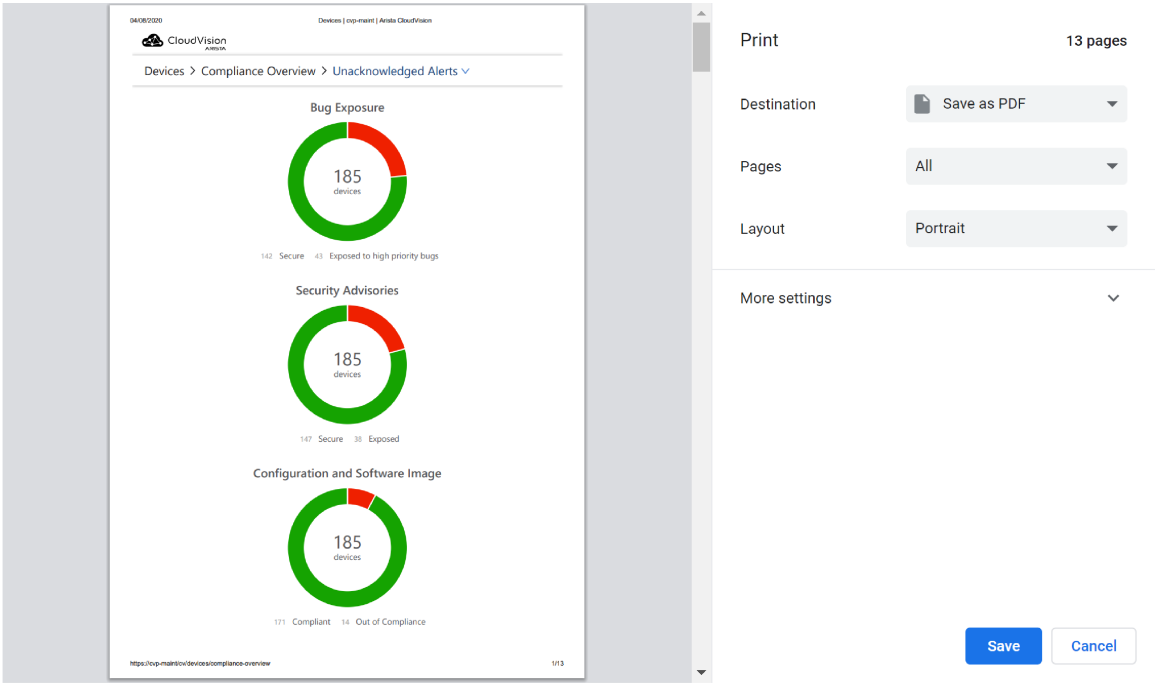

Figure 138: Print Pop-Up Window

2. Select your printer from the **Destination** dropdown menu to print the screen.

   Note: To save a print-friendly version of the screen, select **Save as PDF** from the **Destination** dropdown menu. This PDF contains all rows of the compliance table.

3. Click **Save**.

10.3 Setup for Automatic Sync of Compliance Bug Database

In order to keep the bug database up to date and receive real-time assessments on exposure to software defects and security vulnerabilities, an automated sync can be configured between CVP and https://www.arista.com using a token-based authentication and proxy URL.

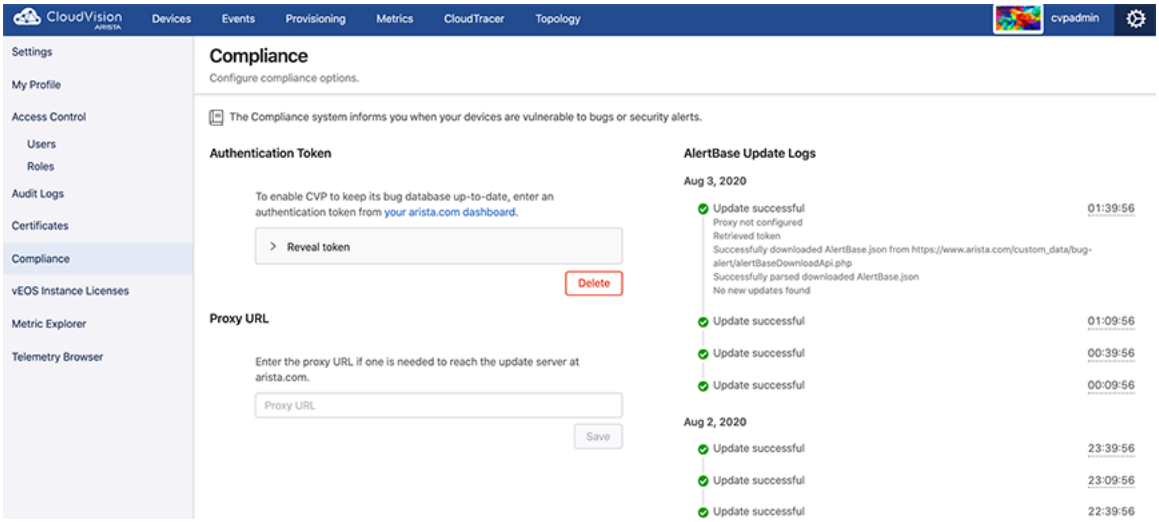

Figure 139: Configuring Compliance Settings

The Compliance screen has a compliance section that accepts the following information:

- An authentication token generated by www.arista.com to enable CVP to keep its bug database up-to-date.
- Proxy URL to reach the update server at www.arista.com.

This token is generated per user and can be obtained from the user profile screen under the Portal Access section on www.arista.com.

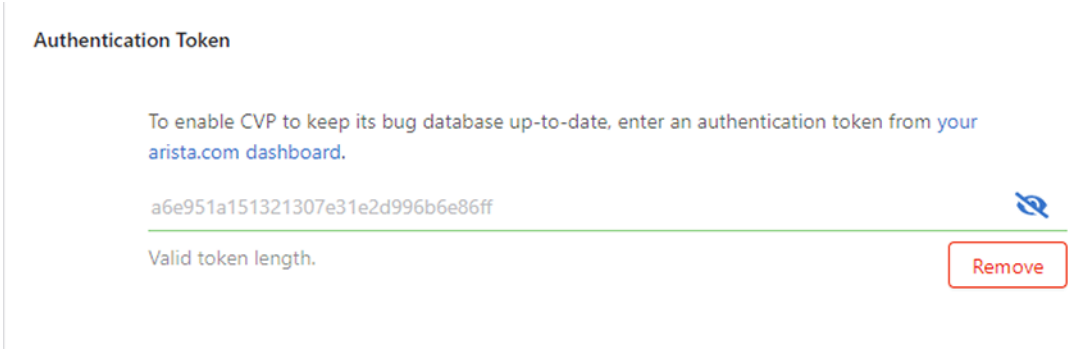

Figure 140: Compliance Portal Access

When this token is provided in the Compliance settings screen, it allows CVP to download the latest version of the https://www.arista.com/en/login file that is available on the Software downloads page.

Note: To leverage automatic updates of the compliance bug database, connectivity to www.arista.com should be ensured from the CVP VM.

The version and release date of the compliance bug database in use can be viewed in the **Settings** screen under **Telemetry Browser > analytics > BugAlerts > update**.

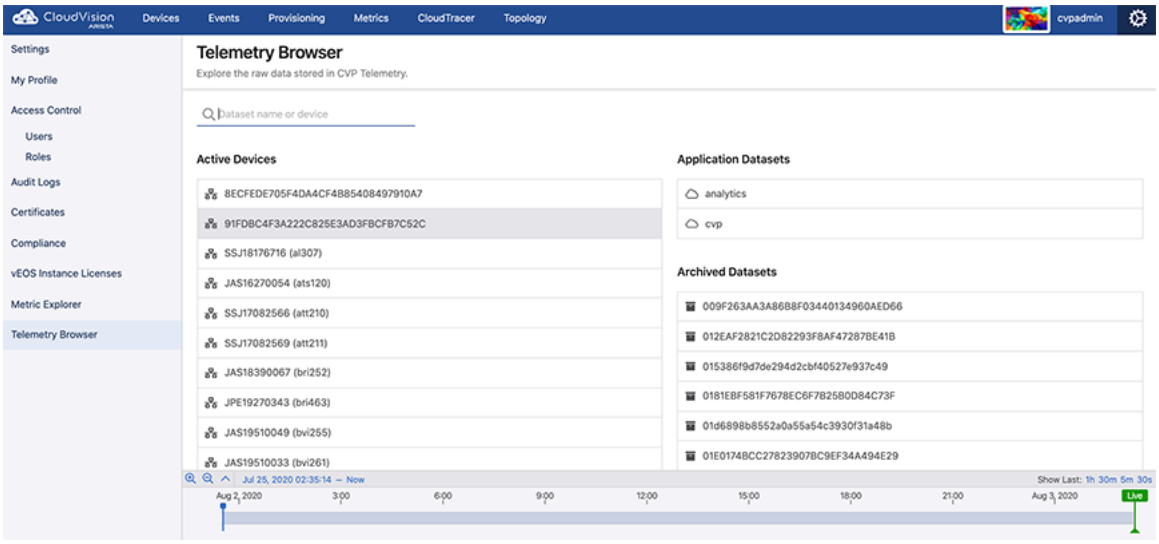

Figure 141: Telemetry Browser Screen

Chapter 11

Network Provisioning (CVP)

The Network Provisioning Screen presents a hierarchical view of the network configuration.

It is not a network topology; it is a configuration tree view. The switches at the bottom of the tree inherit the configuration specified in the containers above them as well as the configuration that is specific to them. The containers and switches all have sub menus that are accessed by right mouse clicking on them. The main features of the screen are described below.

Note: Switches that have been added to the network from new will ZTP boot using generic details from CVP and appear in the Undefined container.

11.1 Network Provisioning View

The topology view of the Network Provisioning screen is a tree structure that consists of containers and devices. This view represents the current groupings of devices (devices grouped by container) as well individual devices.

By default, two types of containers are available in the topology view.

- **Tenant**: Top-most container.
- **Undefined**: Container for all devices that have registered themselves with the CloudVision Portal using Zero Touch Provisioning (ZTP) and are awaiting configuration. Undefined containers are shown in the view in a different color than defined containers.

The example shown below includes:

- One tenant container (there is always only one tenant container).
- Three containers under the tenant container (one of the three is an undefined container).
- Seven devices (one is under the undefined container, and 6are grouped under the container named Vantage-DC (6)).

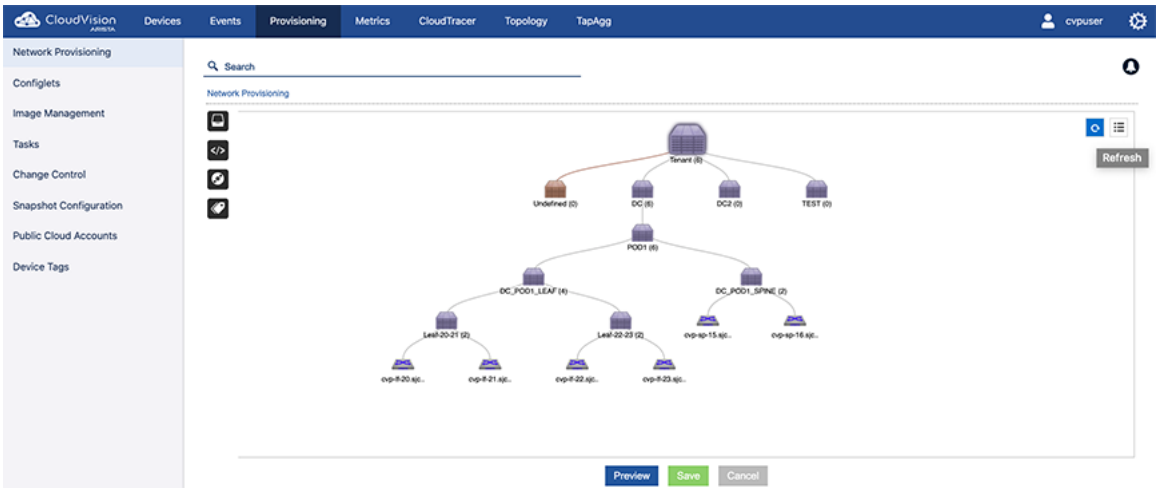

Figure 142: Network provisioning view showing tree structure

Note: Different color icons are used to indicate that devices have compliance alerts or access alerts.

For more information, see:

- Network Provisioning Screen Options
- Changing Between Network Provisioning View and List View

Related topics:

- Container Level Actions (Create, Rename, Delete)
- Device-level Actions
- Viewing Containers and Devices

11.1.1 Network Provisioning Screen Options

The following options are available from the **Network Provisioning** screen.

- **Device Management** Lists all the switches that reside below the selected container level, these could belong to the selected container or reside in containers within the selected container.
- **Configlet Management** Lists the configlets associated with the selected container or if a switch is selected all of the configlets applied to it both directly and inherited.
- **Image Management** Lists the EOS or vEOS software image associated with a container or switch. Switches below the container selected will be loaded with this image.
- **Label Management** Lists the system or custom labels associated with the selected container or switch.
- **Refresh and Listview** Refresh the current screen to show any updates or changes to the switches or devices. Listview changes the display from **Topology View** and displays the switches in a list.
- **Containers** Containers are the basic logical construct of the topology view. They are used to used group devices and to apply configurations and deploy images to the device groups.

Container Right Click Options:

- **Show From Here** Changes the display to show only the containers and switches below the selected container.
- **Expand / Collapse** toggles between shrinking or growing the tree topology below the selected container.
- **Show All Devices** Lists the switches that are associated with that specific container. The container turns blue if it contains more than five switches and will only display 25 of the total number of switches in the topology structure.
- **Container: Add / Delete** Create or remove a container that from the selected container.

- **Device: Add / Manage** Add a device to the selected container or manage the switches already associated with the container. The manage option displays a list of switches which can be selected by enabling the tick box on the left-hand side. The selected switches can then be moved to another container, reset (returned to a ZTP boot state and associated with the undefined container), or removed from CVP completely.
- **Manage: Configlet / Image Bundle** Allocate or remove a configlet or Image to or from a switch or container.
- **View Config** View the configuration created from the combined configlets. At the container level this shows the combined configlet configuration associated with that container.
- **Check Compliance** - To initiate a compliance check on all devices under the container.
- **Reconcile** - To initiate configuration reconcile on all devices under the container.

Device Right Click Options:

- **Manage: Configlet / Image Bundle** Allocate or remove a configlet or Image to or from a switch or container.
- **Labels** Lists / assigns the user created labels associated with the selected switch.
- **View Config** View the configuration created from the combined configlets. At the switch level the entire configuration that will be applied to the switch is shown.
- **Check Compliance** Compares the current running configuration on the switch against the designed configuration in CVP. If they are out of sync the device change to an orange color.
- **Move** Allows a user to move a switch from one container to another.
- **Factory Reset** Erases the configuration on the switch then ZTP boots it. This will return it to the undefined container on the provisioning screen.
- **Remove** Removes the switch from CVP. This stops CVP making changes to it and tracking its configuration. The switch is left running with its current configuration on it.
- **Replace** - To perform a Zero Touch Replacement (ZTR) of the selected device.

Related topics:

- Changing Between Network Provisioning View and List View
- Container Level Actions (Create, Rename, Delete)
- Device-level Actions
- Viewing Containers and Devices

11.1.2 Changing Between Network Provisioning View and List View

Click the icons to toggle between the topology view and the list view of the Network Provisioning screen.

Changing to List View

Click the **List** icon for a list view.

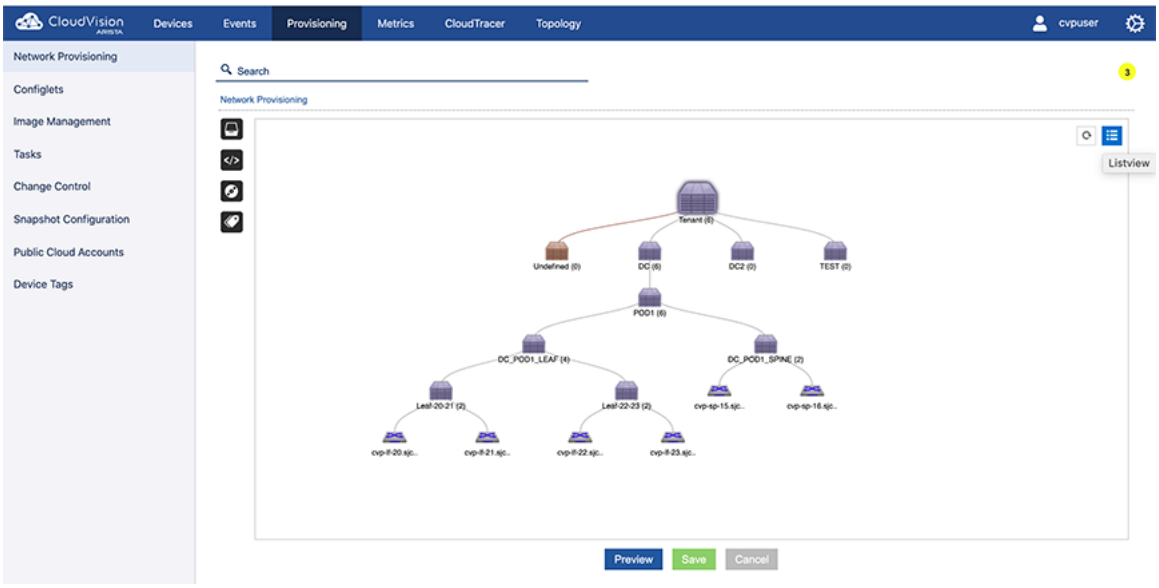

Figure 143: Changing to List View

Changing to Topology View

Click the **Topology** icon for a topology view.

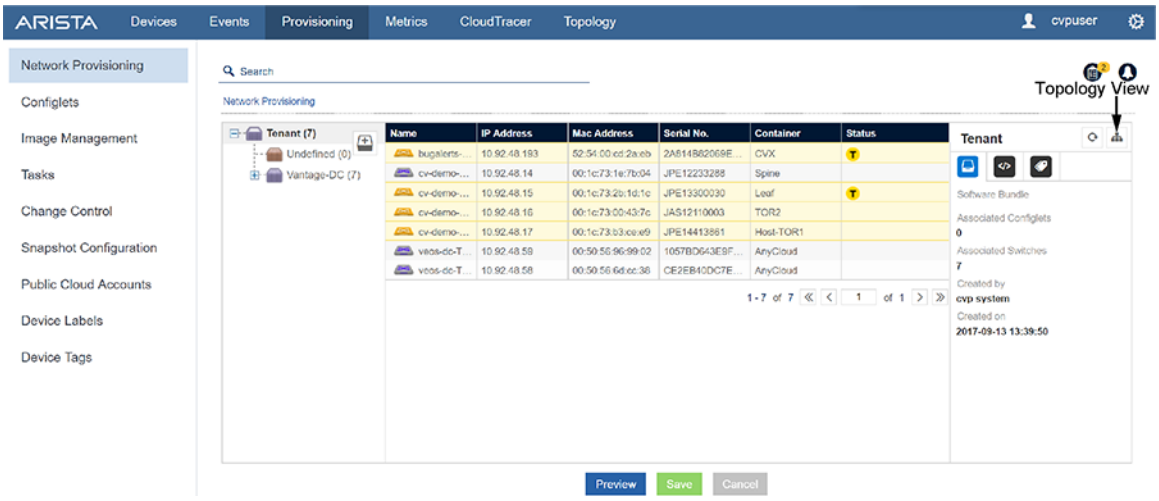

Figure 144: Changing to Topology View

Related topics:

- Network Provisioning Screen Options
- Container Level Actions (Create, Rename, Delete)
- Device-level Actions
- Viewing Containers and Devices

11.2 Container Level Actions (Create, Rename, Delete)

Containers are a logical entity used to group network devices and to define a hierarchy to which configurations can be applied. When you apply a configlet to a container, the configlet is automatically applied to all of the devices in the container's hierarchy.

Simple container implementations:

- Create a container for every datacenter.
- Within each datacenter container, create a container for every POD (leaf-spine deployment).
- Add devices that belong to each POD to the POD container. Tenant: Top-most container.

For details on how to create, rename, and delete containers, see:

- Creating a Container
- Deleting a Container
- Renaming a Container

Related topics:

- Device-level Actions
- Viewing Containers and Devices
- Device Compliance
- Notifications for Container-level Compliance Checks and Reconciles

11.2.1 Creating a Container

To create a container:

1. Select a parent container (the container to which you want to add a new container).
2. Right-click the container and choose **Add > Container**. The **New Container** dialog appears:

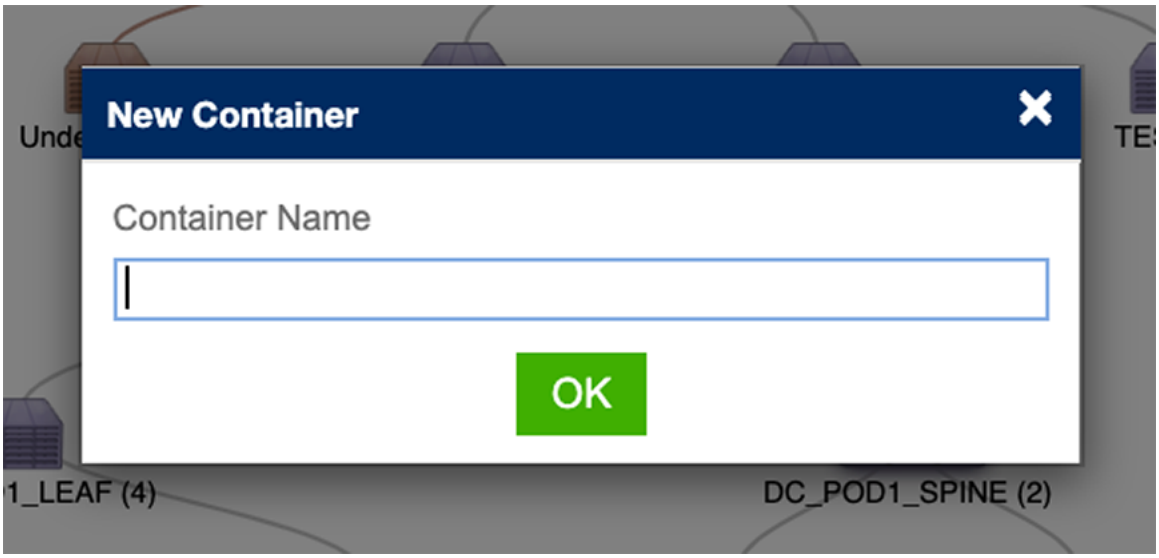

Figure 145: New Container Dialog

3. Enter the name of the new container and select **OK** to create the container.
4. Click **Save** to apply the changes.

Related topics:

- Device-level Actions
- Viewing Containers and Devices
- Device Compliance

11.2.2 Deleting a Container

Note: Only empty containers can be deleted.

1. Locate the container to be deleted.
2. Right-click the container and click **Remove**.

Related topics:

- Device-level Actions
- Viewing Containers and Devices
- Device Compliance

11.2.3 Renaming a Container

To rename a container in a topology:

1. Double-click the name field of the container to open the name field editor.
2. Enter a new, unique name for the container and click **Enter** to rename the container.

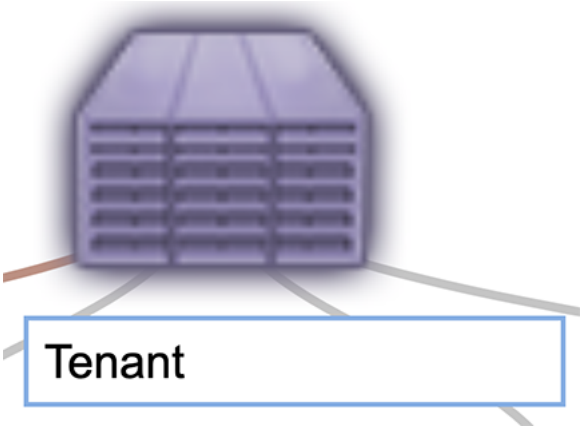

Figure 146: Rename Container

Related topics:

- Device Bootstrap Process
- Device-level Actions
- Viewing Containers and Devices
- Device Compliance

11.3 Device Bootstrap Process

The device bootstrap process is a process that automatically makes un-provisioned devices available for configuration through CVP. Un-provisioned devices automatically boot up in Zero Touch Provisioning mode and register themselves with the CloudVision Portal (CVP). Once they are registered with CVP, devices become available for configuration in the Undefined Container.

1. Un-provisioned devices boot into Zero Touch Provisioning mode and send out a DHCP request.
2. The DHCP server then assigns the device an IP Address and returns a URL pointing to the CloudVision portal in the bootfile-name option. The URL to specify is http://IPADDRESS/ztp/bootstrap.
3. The device executes this bootstrap script and registers itself with the CloudVision Portal. At this point, the device is available in the Undefined Container.

   You can now add the device to the destination container of your choice and apply the correct image and configuration to the device.

Related topics:

- Device-level Actions
- Viewing Containers and Devices

- Device Compliance

11.4 Device-level Actions

CloudVision Portal (CVP) enables you to provision devices as needed based on your current networking requirements. Some examples of the types of actions you can perform include:

- Adding devices (use this action to add devices from the undefined container to defined containers)
- Moving devices (used this action to move devices from one defined container to another defined container)
- Removing devices (removing devices from the CVP topology)
- Reset devices
- Replace devices

For details on the steps you use to perform these device level actions, see:

- Adding Devices (from Undefined Container)
- Deploying vEOS Routers
- Registering Devices
- Moving Devices from one Container to Another Container
- Removing a Device from a Container
- Device Factory Reset
- Replacing Switches Using the ZTR Feature
- Managing Configurations
- Configuration Validation
- Using Hashed Passwords for Configuration Tasks
- Reconciling Configuration Differences
- Managing EOS Images Applied to Devices
- Rolling Back Images and Configurations
- Device Labels
- Viewing Containers and Devices
- Device Compliance
- Notifications for Container-level Compliance Checks and Reconciles
- Global Search
- Management IP

When resetting a device:

- The device will be removed from the parent container.
- The running configuration of the device will be flushed.
- Device will reboot with ZTP mode enabled.
- Device will be identified under undefined container.

There are three options you can use to move devices. They are:

- Option 1:
- Option 2:
- Option 3:

Option 1:

1. Locate the device.
2. Right-click the device and choose **Factory Reset**.

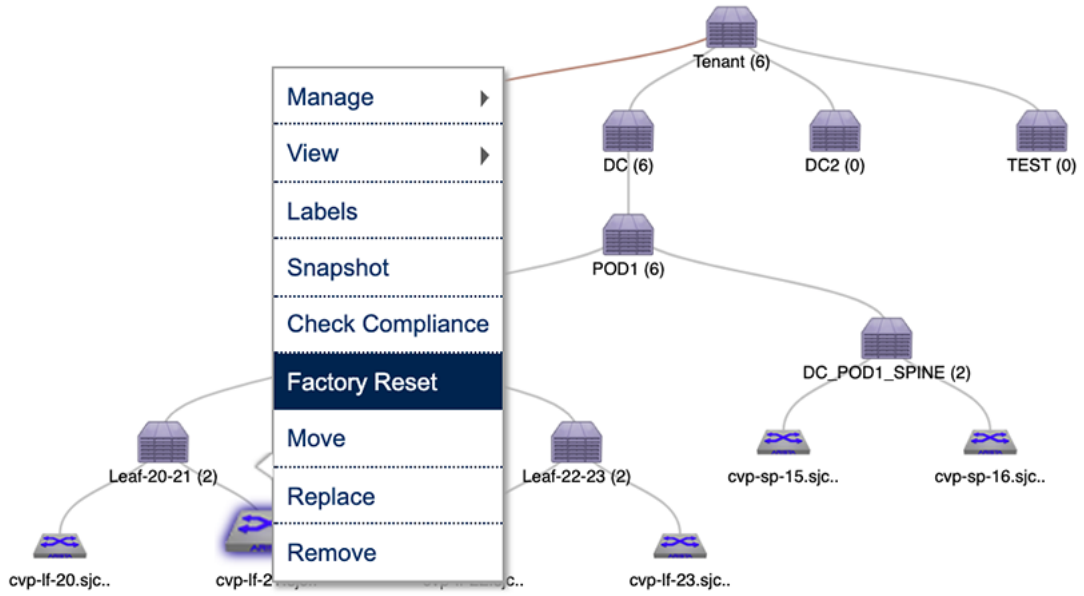

Figure 147: Resetting the Device (option 1)

Option 2:

1. Locate the parent container.
2. Right-click the container and choose **Show All Devices.** This will list all the devices under the container.

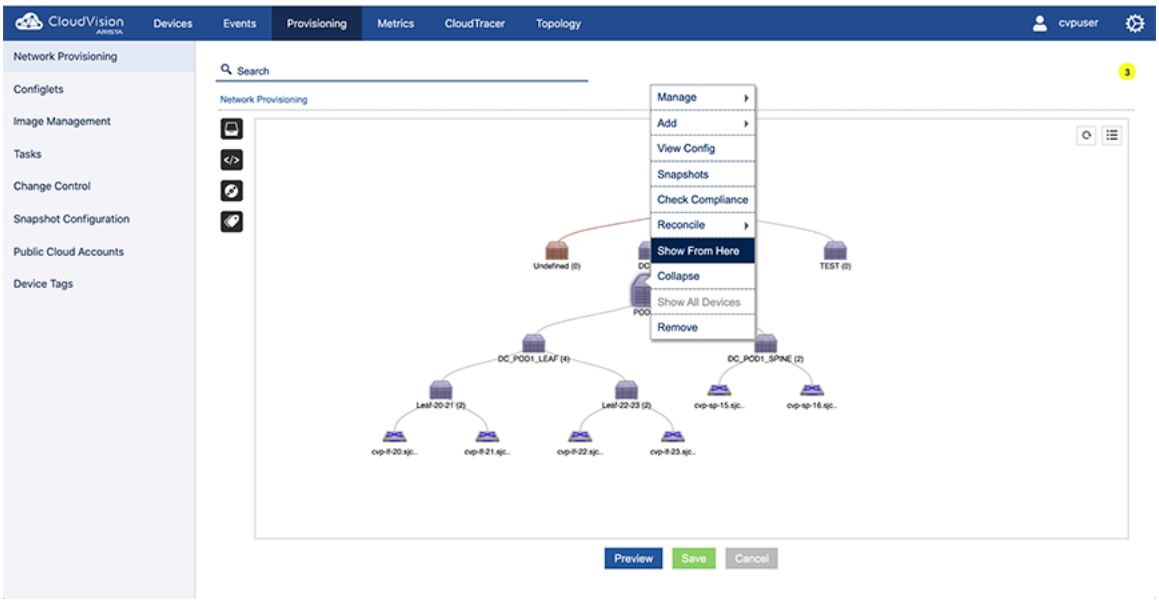

Figure 148: Showing all devices during factory reset (option 2)

3. Right-click the device and choose **Factory Reset**.

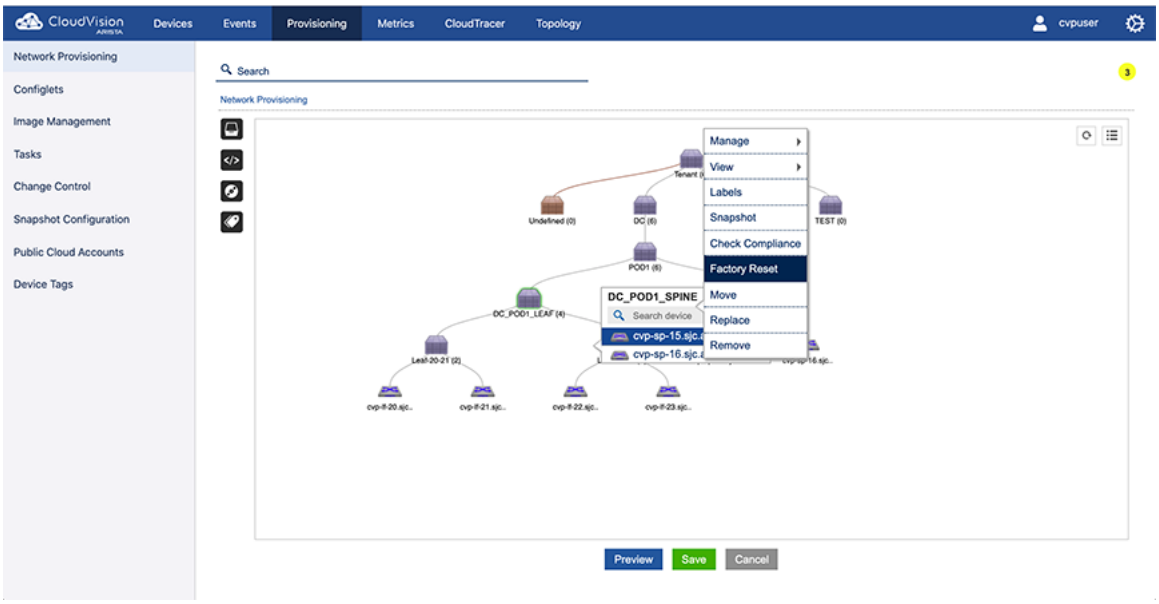

Figure 149: Resetting the device (option 2)

Option 3:

1. Locate the parent container.
2. Right-click the container and choose **Manage > Device**. This will load the inventory of all the child devices under the container.
3. Select the checkbox of the device to be reset, and click the reset icon.

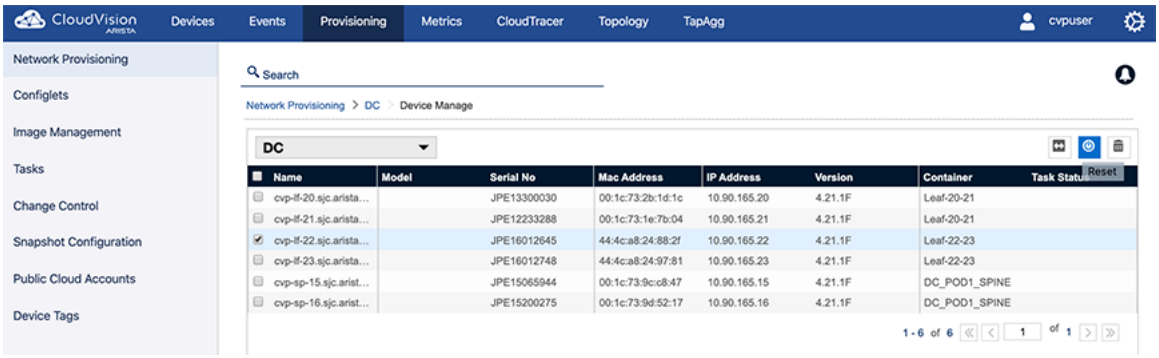

Figure 150: Selecting the device and resetting it (option 3)

On saving the session, a task will be spawned to reset the selected device.

11.4.1 Adding Devices (from Undefined Container)

Adding devices from the undefined container is the most common method for adding devices to a container in the CVP topology. This method involves adding devices that are not part of the hierarchy of devices to defined containers in the CVP topology. Containers that receive the added devices are called destination containers.

Complete the following steps to add a device from the undefined container to a destination container:

1. Locate the container to which you want to add a device.
2. Right-click the container and choose **Add > Device**. The current inventory of undefined devices for the selected container appears.

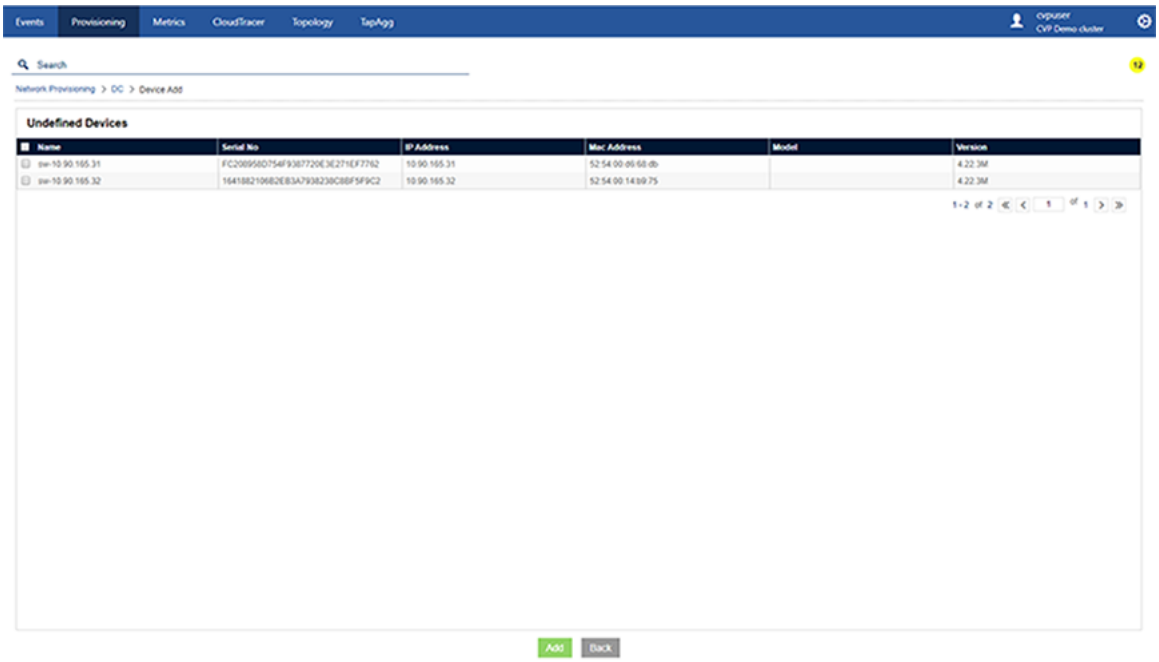

Figure 151: Adding a device

3. Select the device and click **Add**.
4. Save the session.
5. Execute the **Device Add** task using the **Task Management** module to add the device to destination container.

11.4.2 Deploying vEOS Routers

CVP deploys and provisions vEOS routers from cloud and datacenter to Amazon Web Services (AWS) and Microsoft Azure. Based on the requirement in vEOS deployment, configlets are assigned for push EOS configuration along with deployment parameters such as AWS Virtual Private Cloud (VPC), subnets, and security groups.

Note: When CVP is deployed behind NAT devices, the vEOS telemetry configuration needs to be updated. You can view telemetry data coming from the deployed device when you configure the public IP address of CVP.

Related Topics:

- Prerequisites
- Adding IPsec and vEOS Licenses
- Adding AWS to Public Cloud Accounts
- Deploying the vEOS Router to AWS
- Adding Microsoft Azure to Public Cloud Accounts
- Deploying a vEOS Router to Microsoft Azure

11.4.2.1 Prerequisites

The prerequisites to deploy vEOS routers within a cloud are:

- vEOS version *4.21.1.1F* or later
- *CVP 2018.2.0*
- vEOS license
- Cloud (AWS/Microsoft Azure) credentials

- vEOS deployment parameters including VPC within which the vEOS has to be deployed, subnets and security groups associated with vEOS
- IP connectivity from deployed vEOS to CVP

11.4.2.2 Adding IPSec and vEOS Licenses

The addition of an IPSec license is optional based on the deployment.

Perform the following steps to add IPSec and vEOS licenses:

1. Click the gear icon at the upper right corner of the CVP. The system displays the **Settings** screen.
2. Click **EOS Feature Licenses** in the left pane. The system displays the **EOS Feature Licenses** screen.

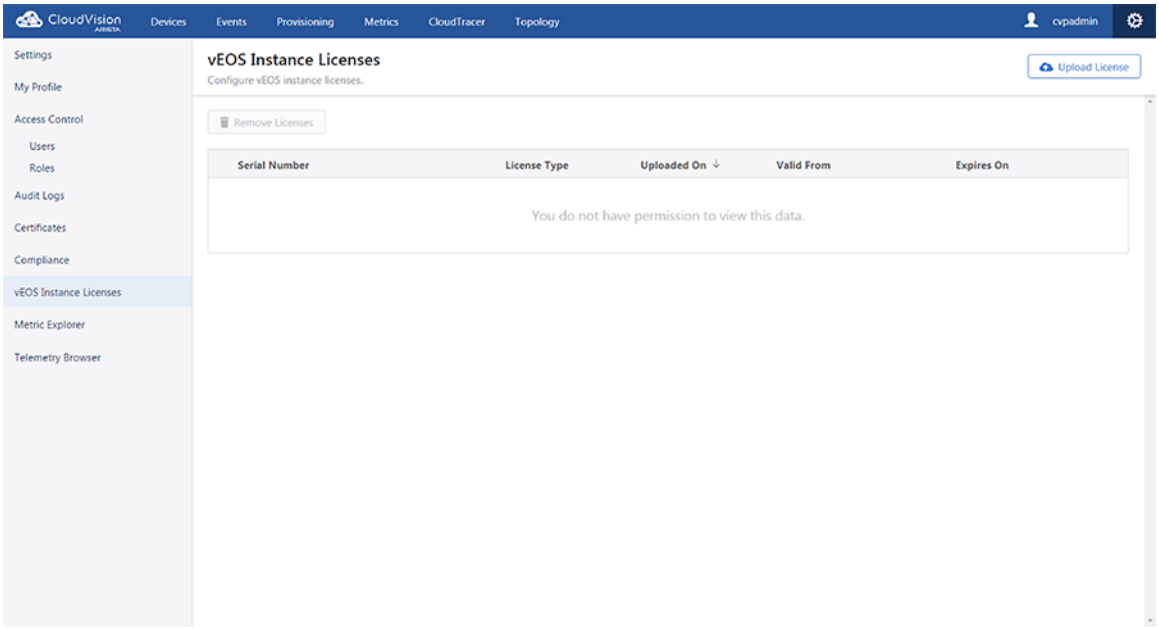

Figure 152: EOS Feature Licenses Screen

3. Click **Add License** in the right pane. The system displays the **Add License** window.

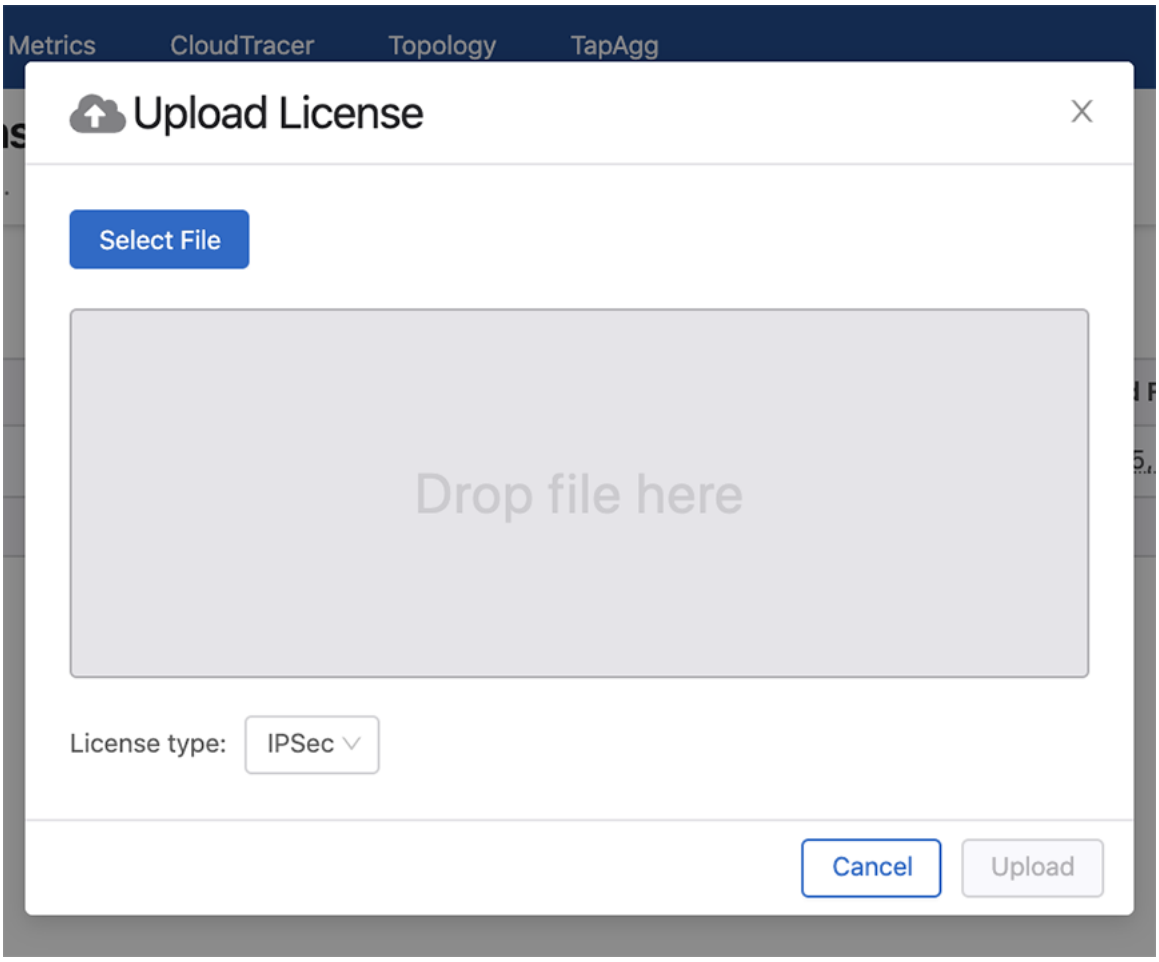

Figure 153: Add License Window

4. Click **Select license file**. The system displays the Windows Explorer.
5. Navigate to the required location and select the license.
6. Click **Open**.
7. Select the required option from the **License type** drop-down menu.
8. Click **Upload**. The system lists uploaded licenses in the **EOS Feature Licenses** screen.

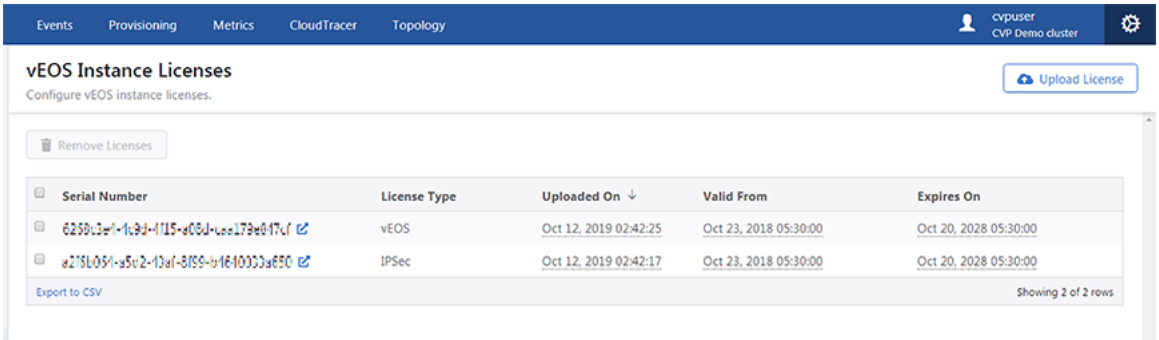

Figure 154: Licenses Listed in EOS Feature Licenses Screen

11.4.2.3 Adding AWS to Public Cloud Accounts

AWS Security Token Service (STS) is required when adding an AWS account to public cloud accounts.

AWS STS gives CVP temporary access to your AWS environment with proper permissions. This allows CVP to deploy the vEOS router and related resources in your AWS VPC.

CVP calls certain AWS APIs to query VPC information and creates a vEOS router Virtual Machine (VM) in VPC. It needs an AWS IAM (Identity and Access Management) role with permissions as listed in the code below .

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeRegions",
                "ec2:DescribeVpcs",
                "ec2:DescribeImages",
                "ec2:DescribeAddresses",
                "ec2:DescribeKeyPairs",
                "ec2:DescribeAvailabilityZones",
                "ec2:DescribeSubnets",
                "ec2:DescribeSecurityGroups",
                "ec2:DescribeNetworkInterfaces",
                "ec2:CreateNetworkInterface",
                "ec2:ModifyNetworkInterfaceAttribute",
                "ec2:DetachNetworkInterface",
                "ec2:DeleteNetworkInterface",
                "ec2:AllocateAddress",
                "ec2:AssociateAddress",
                "ec2:DisassociateAddress",
                "ec2:ReleaseAddress",
                "ec2:RunInstances",
                "ec2:TerminateInstances"
            ],
            "Resource": "*"
        }
    ]
}
```

Note: You receive the STS token after the IAM role is created.

Perform the following steps to add a AWS account to public cloud accounts:

1. Click **Provisioning**. The system displays the **Network Provisioning** screen.
2. Click **Public Cloud Accounts** in the left pane. The system displays the **Public Cloud Accounts** screen.

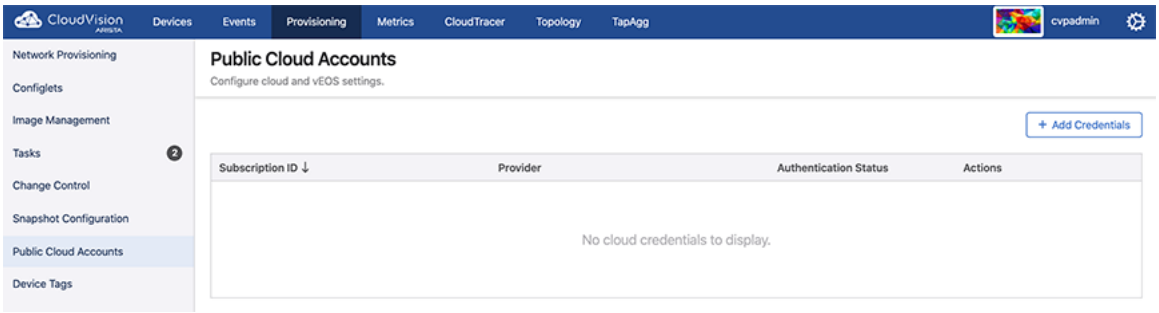

Figure 155: Public Cloud Accounts Screen

3. Click **Add Credentials** in the upper right corner of the right pane. The system displays the **Add Credentials** window.
4. Select **Amazon Web Services** from the **Provider** drop-down menu.

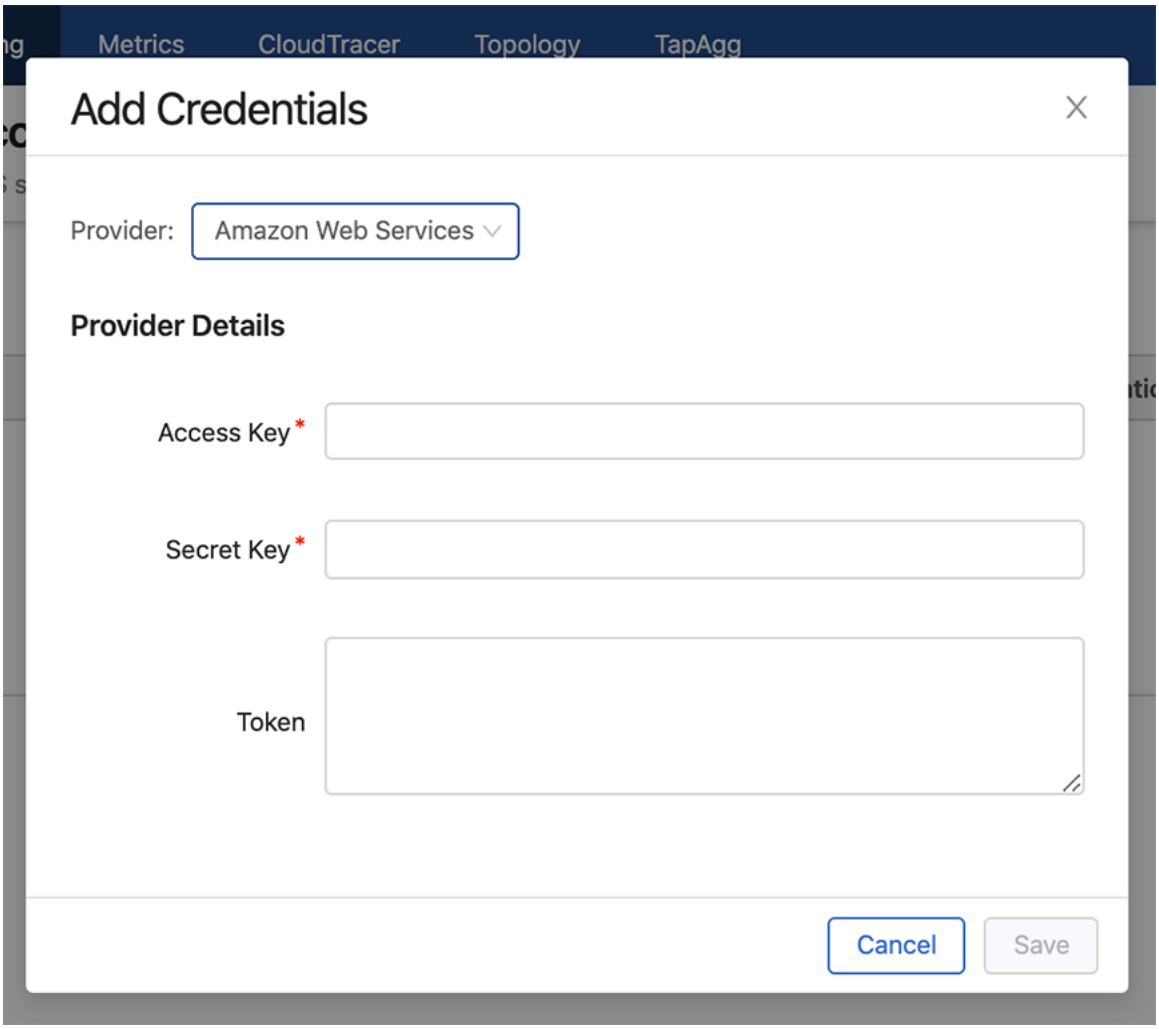

Figure 156: Add Credentials Window for AWS

5. On the **Provider Details** pane, provide the access key, secret key, and token details in the corresponding fields.
6. Click **Save**. The system displays the configured AWS account in the **Public Cloud Accounts** screen.

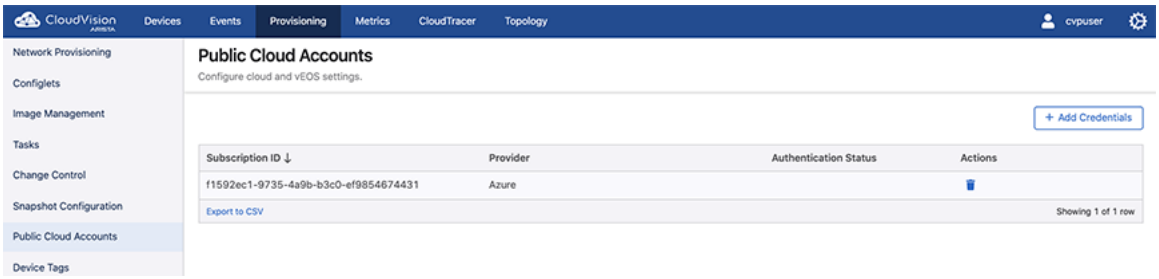

Figure 157: AWS Configured in Public Cloud Accounts

11.4.2.4 Deploying the vEOS Router to AWS

Perform the following steps to deploy the vEOS router to AWS:

1. Click **Devices**. The system displays the Inventory screen.
2. Click the **Add Devices** drop-down menu at the upper right corner of the right pane.
3. Select **Deploy vEOS Router**. The system displays the **Deploy vEOS Router** window.

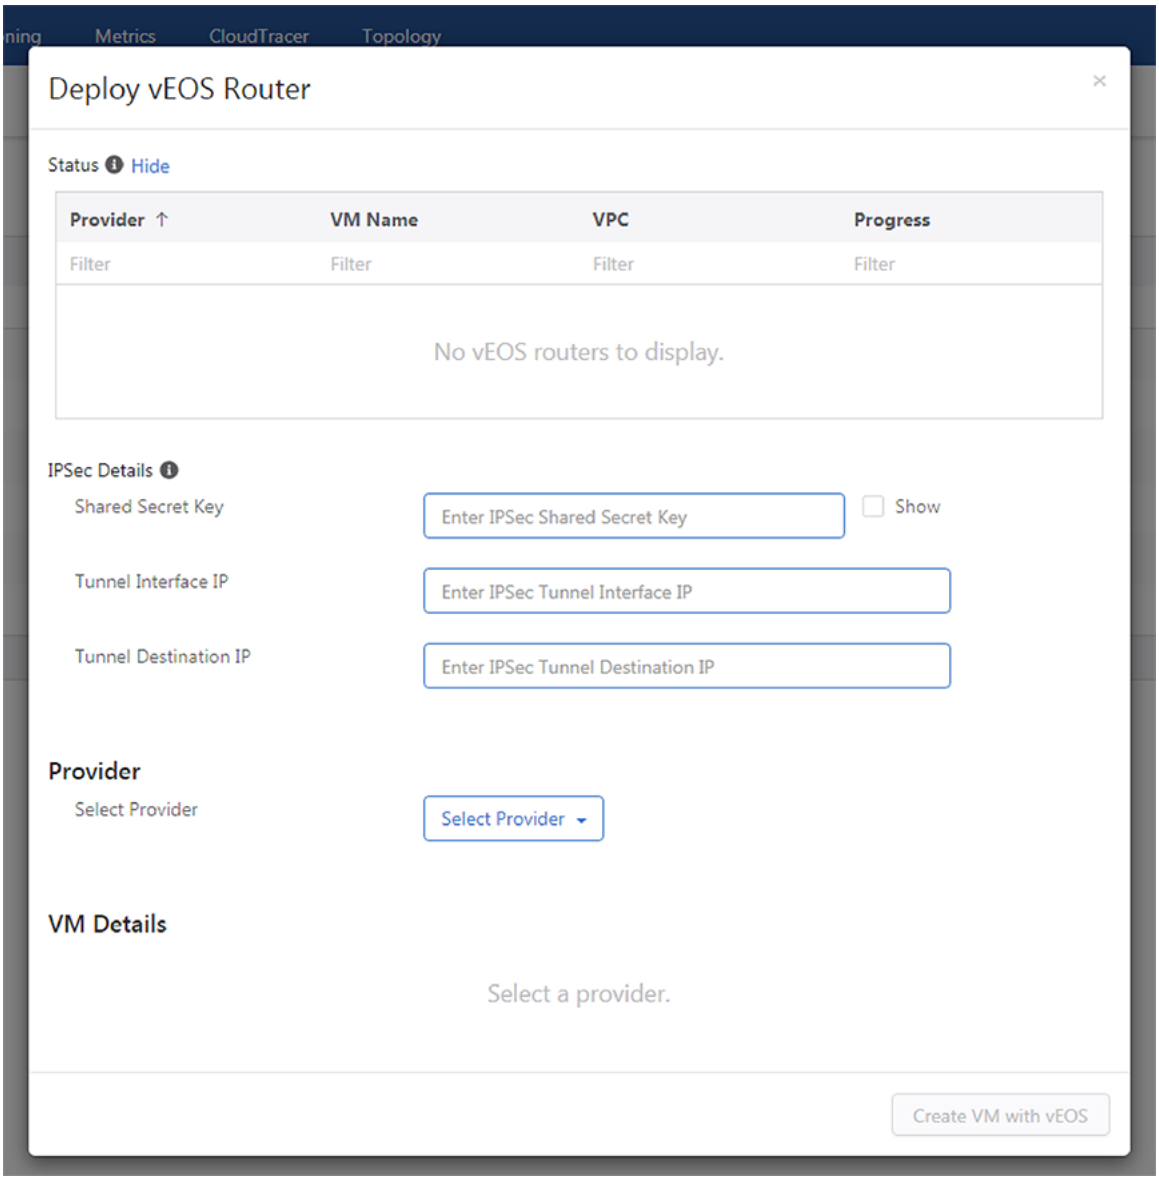

Figure 158: Deploy vEOS Router Window

4. Provide the following IPSec details in the appropriate fields:
   - **Shared Secret Key (optional)** - Pre-shared key for IPSec profile
   - **Tunnel Interface IP (optional)** - IP address under tunnel interface
   - **Tunnel#1 Destination IP (optional)** - Peer's (tunnel destination) IP address
5. Click the **Select Provider** drop-down menu and select **AWS**.

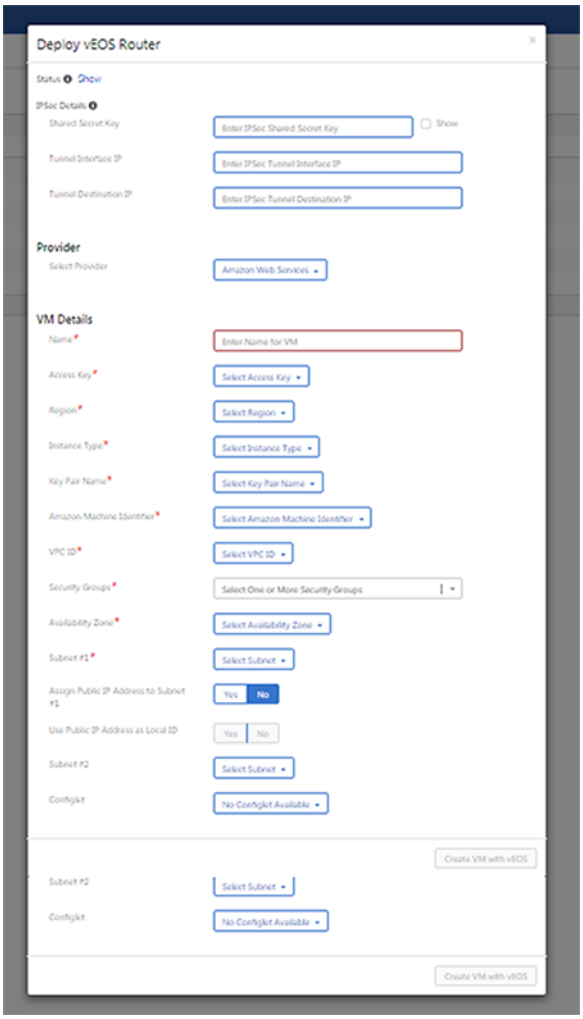

Figure 159: VM Details for AWS

6. Provide the following VM details in the appropriate fields:
   - **Name** - The name of the vEOS router instance
   - **Access Key** - The access key used in the public cloud account
   - **Region** - The region that the vEOS router will be deployed in
   - **Instance Type** - The type of vEOS router that the instance will run on
   - **Key Pair Name** - The Elastic Compute Cloud (EC2) keypair used to log in to the vEOS router
   - **Amazon Machine Identifier** - The vEOS AMIs on the AWS marketplace
   - **VPC ID** - The VPC that the vEOS router will be deployed to
   - **Security Group** - The security group that will be associated with the vEOS interface
   - **Availability Zone** - The availability zone that vEOS will be deployed in
   - **Subnet #1** - The first subnet that vEOS puts Ethernet1 in
   - **Assign Public IP Address to Subnet #1** - Select Yes if you need a public IP address assigned to the vEOS router; otherwise, select No
   - **Use Public IP Address as Local ID** - The public IP address of the vEOS router

     Note: The system displays the public IP address of the vEOS router after the VM is created.
   - **Subnet #2 (optional)** - The second subnet that vEOS puts Ethernet2 in
   - **Configlet (optional)** - The configlet to configure vEOS once it is active
7. Click **Create VM with vEOS**. The system displays the status of vEOS deployment under the **Progress** column on the **Status** pane.

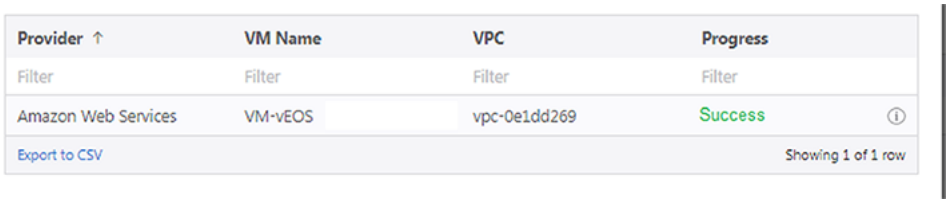

Figure 160: Status of vEOS Deployment to AWS

You can also check the VM deployment process on your AWS Portal. Hover the mouse over the corresponding information icon to view detailed information about the vEOS router deployment. After the successful deployment of the vEOS router to AWS, you can use your AWS SSH Privacy Enhanced Mail (PEM) key to login to vEOS.

Note: To make CVP manage vEOS routers, register this device using the instructions in Registering Devices. Ensure that the AWS security group associated with vEOS router VM has an ingress rule of allowing TCP port 9910 from CVP's IP address. You must configure AWS for the vEOS router to function as a VPC gateway using the instructions in Using vEOS Router on the AWS Platform.

11.4.2.5 Adding Microsoft Azure to Public Cloud Accounts

You need a subscription ID, a tenant ID, a client ID, and client server details in order to an azure account to public cloud accounts.

To get these details, you must create an application in the Azure active directory and assign proper permissions to CVP for authentication with Microsoft Azure environment to make API calls. CVP uses a few APIs to create a vEOS router. Therefore, you must add a "contributor" role to the resource group that has either Virtual Network Protocol (VNET) or the whole subscription.

Perform the following steps for adding the Microsoft Azure account to public cloud accounts:

1. Click **Provisioning**. The system displays the **Network Provisioning** screen.
2. Click **Public Cloud Accounts** in the left pane. The system displays the **Public Cloud Accounts** screen.
3. Click **Add Credentials** in the upper right corner of the right pane. The system displays the **Add Credentials** window.

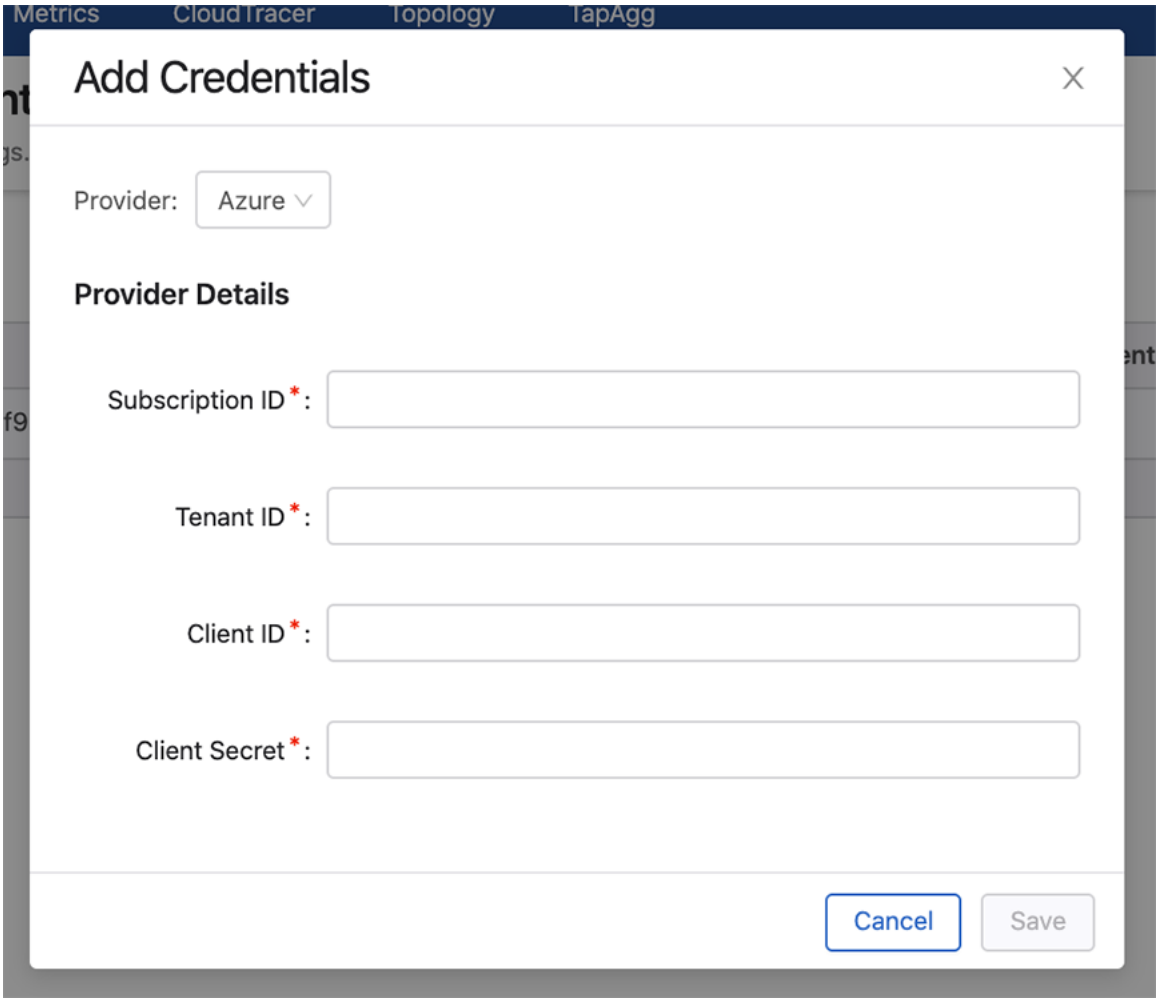

Figure 161: Add Credentials Window for Microsoft Azure

4. Select **Azure** from the **Provider** drop-down menu.
5. Under the **Provider Details** pane, provide the subscription ID, tenant ID, client ID, and client server details in the appropriate fields.
6. Click **Save**. The system displays the configured Microsoft Azure account in the **Public Cloud Accounts** screen.

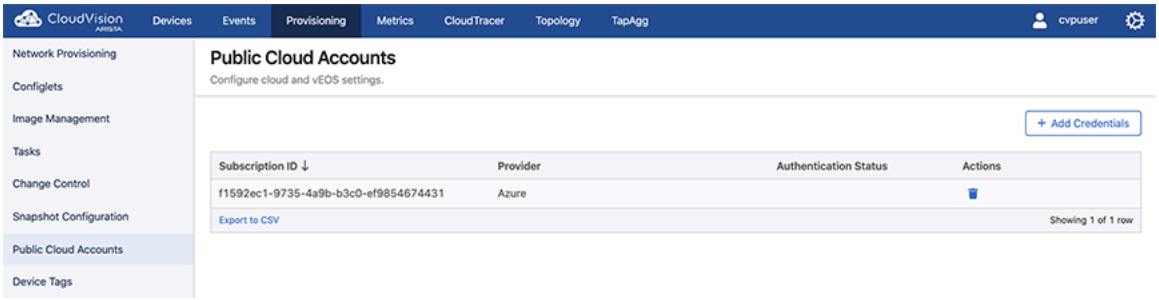

Figure 162: Microsoft Azure Configured in Public Cloud Accounts

11.4.2.6 Deploying a vEOS Router to Microsoft Azure

Perform the following steps to deploy a vEOS router to the Azure VNET:

1. Click **Devices**. The system displays the **Inventory** screen.
2. Click the **Add Devices** drop-down menu at the upper right corner of the right pane.

3. Select **Deploy vEOS Router**. The system displays the **Deploy vEOS Router** window.
4. Provide the following IPSec details in the appropriate fields:
   - **Shared Secret Key** (optional) - Pre-shared key for IPSec profile
   - **Tunnel Interface IP** (optional) - IP address under tunnel interface
   - **Tunnel#1 Destination IP** (optional) - Peer's (tunnel destination) IP address
5. Select **Azure** from the **Select Provider** drop-down menu.

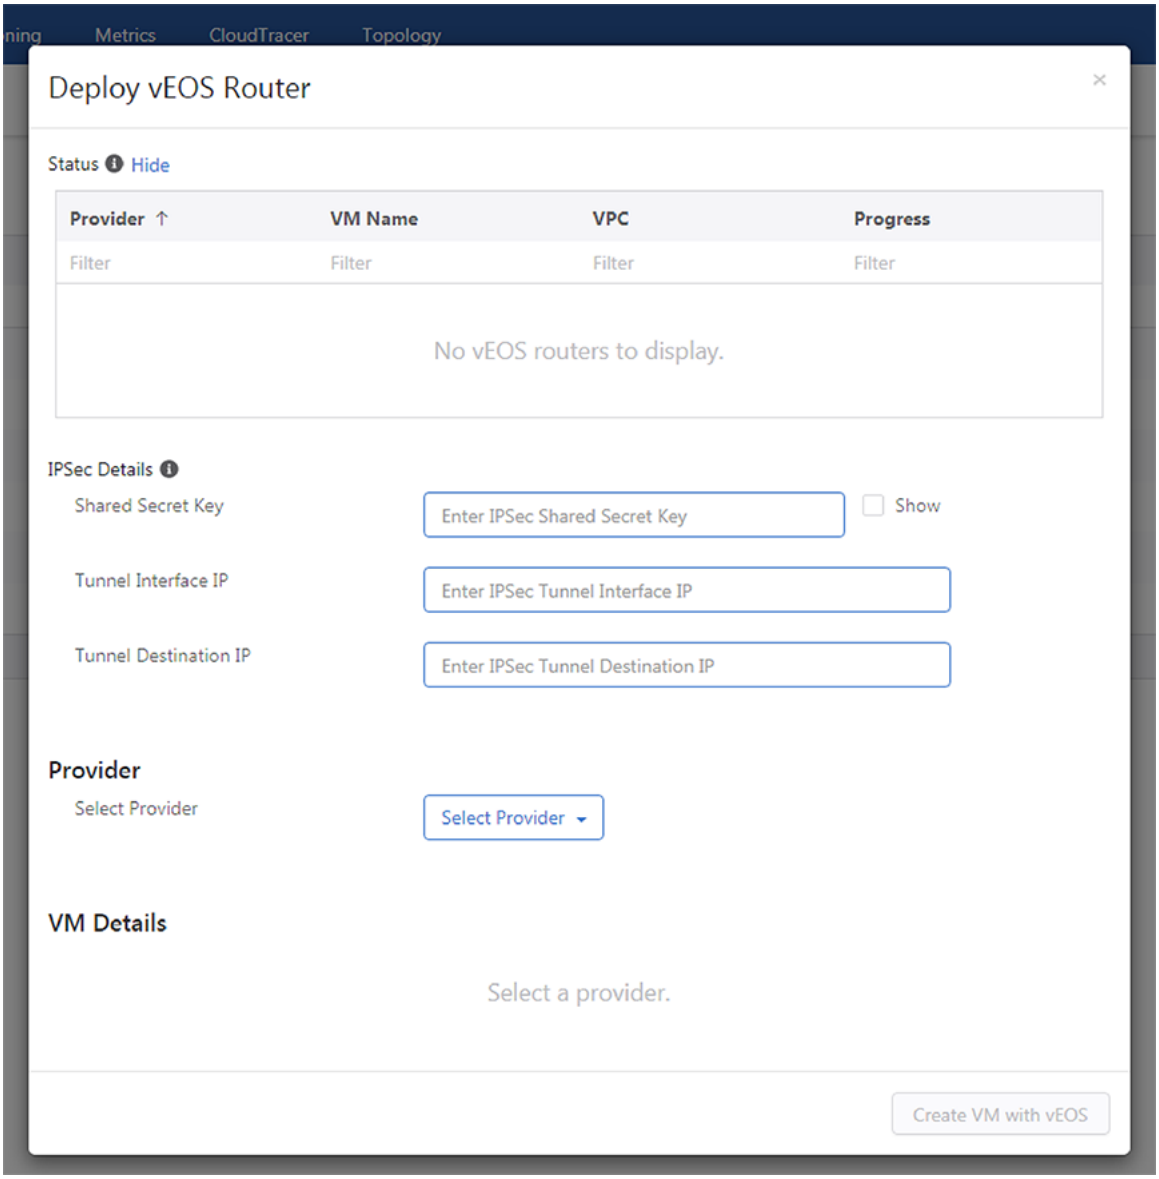

Figure 163: VM Details for Microsoft Azure

6. Provide the following VM details in the appropriate fields:
   - **Name** - The name of the vEOS router instance.
   - **Subscription ID** - The subscription that the vEOS router will be deployed to.
   - **Instance Size** - The size of vEOS router that the instance will run on.
   - **Resource Group** - The resource group that the vEOS router will be deployed to.
   - **Location** - The Azure region that contains the VNET.
   - **Security Group** - The network security group that will be associated with the vEOS interface.
   - **Virtual Network** - The VNET that vEOS will be deployed in.
   - **Subnet #1** - The first subnet that vEOS puts Ethernet1 in.

- **Assign Public IP Address to Subnet #1** - Select Yes if you need a public IP address assigned to vEOS router, else select No.
- **Use Public IP Address as Local ID** - The public IP address of vEOS Router.

  Note: The system displays the public IP address of vEOS router after the VM is created.
- **Subnet #2** - The second subnet that vEOS puts Ethernet2 in.
- **Configlet** - The configlet to configure vEOS once it is up.
- **EOS Image** - The vEOS images on Azure marketplace.

7. Click **Create VM with vEOS**. The system displays the status of vEOS deployment under the Progress column in the Status pane.

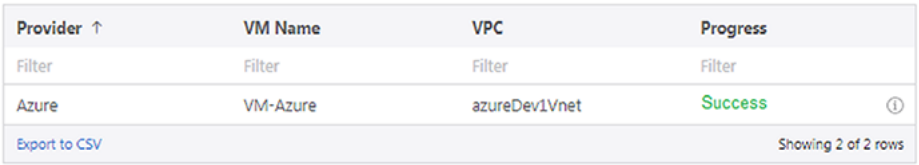

Figure 164: Status of vEOS Deployment to Microsoft Azure

You can also check the VM deployment process on your Microsoft Azure Portal. Hover the mouse over the corresponding information icon to view detailed information about the vEOS router's deployment. It contains the initial login credentials you can use to login to vEOS router, you can change the credentials after logging into the device.

Note: To make CVP manage vEOS routers, register this device using the instructions in Registering Devices. Ensure that the Azure network security group associated with vEOS router VM has an ingress rule of allowing TCP port 9910 from CVP's IP address. You must configure Microsoft Azure for the vEOS router to function as VNET gateway using the instructions in **Using the vEOS Router on Microsoft Azure**.

11.4.3 Registering Devices

Registering is the method used for adding devices to CVP. As a part of registering devices, CloudVision automatically enables streaming of the registered devices' state to the cluster by installing and configuring the TerminAttr agent. Newly registered devices are always placed under an undefined container.

Note: Manual installation or configuration of streaming telemetry is not required prior to registration.

Complete the following steps to register devices with CVP:

1. Navigate to the **Inventory** screen.
2. Click the **Add Device** drop-down menu and select **Register Existing Device**. The **Device Registration** pop-up window appears.

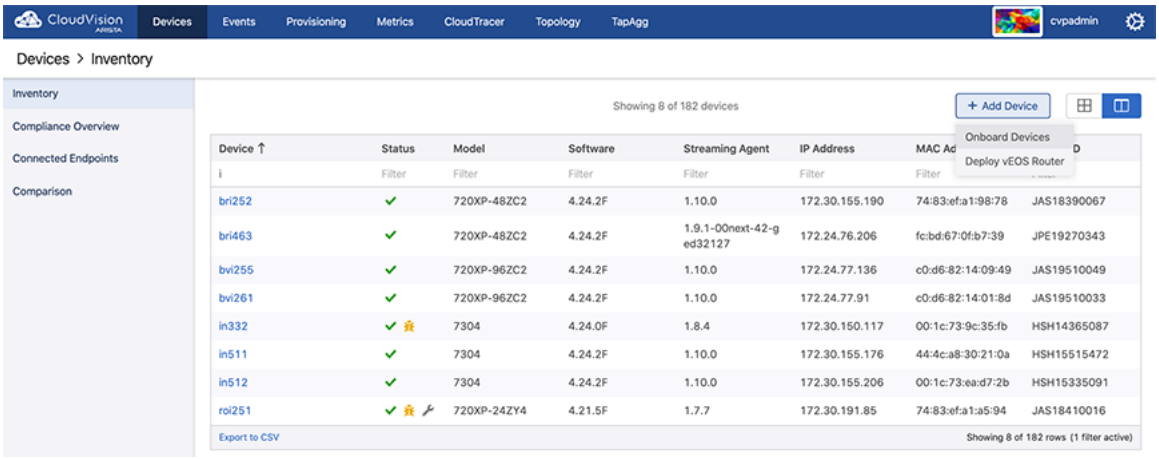

Figure 165: Add Device for Registration

3. Enter the host name or IPv4 addresses of the device(s) to be registered; and click **Register**.

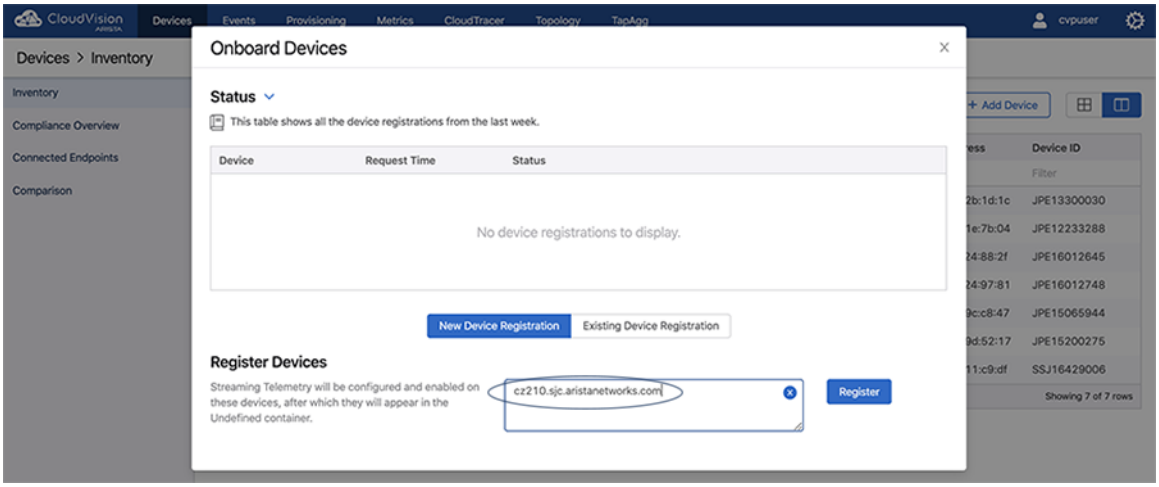

Figure 166: Selecting Device for Registering

The following figures show the device registration status through the registration process.

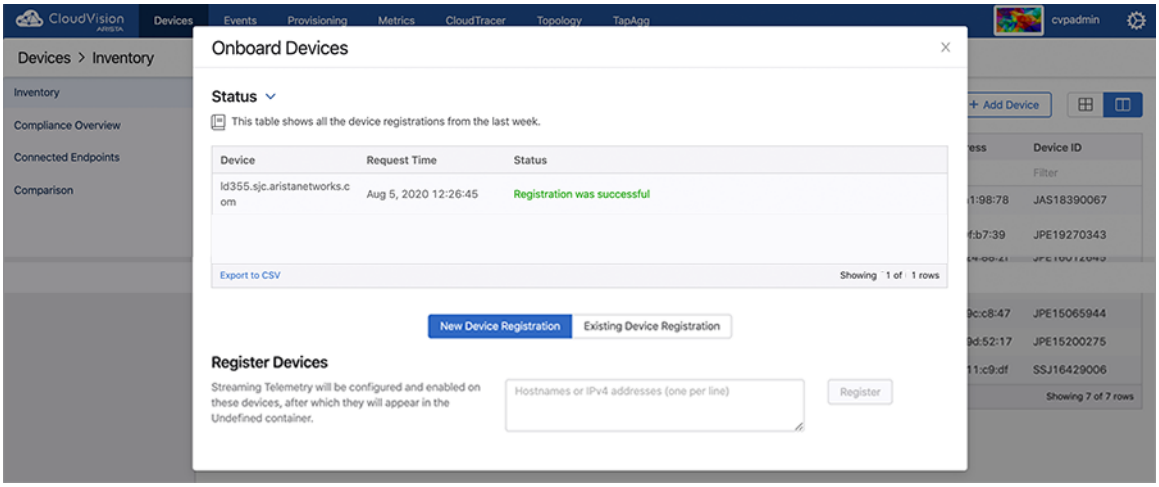

Figure 167: Registration Status

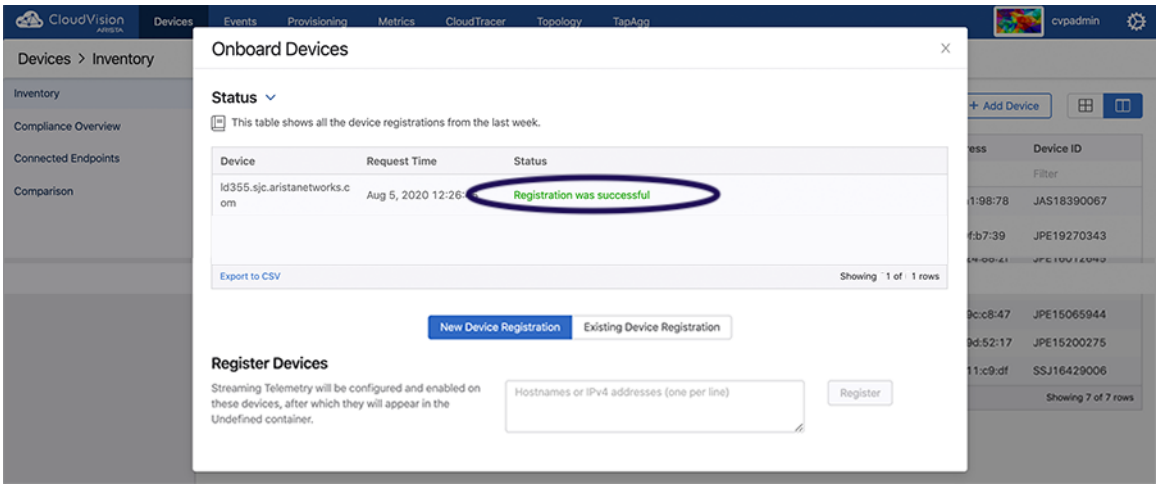

Figure 168: Registration Successful

The newly registered devices are now shown in the inventory.

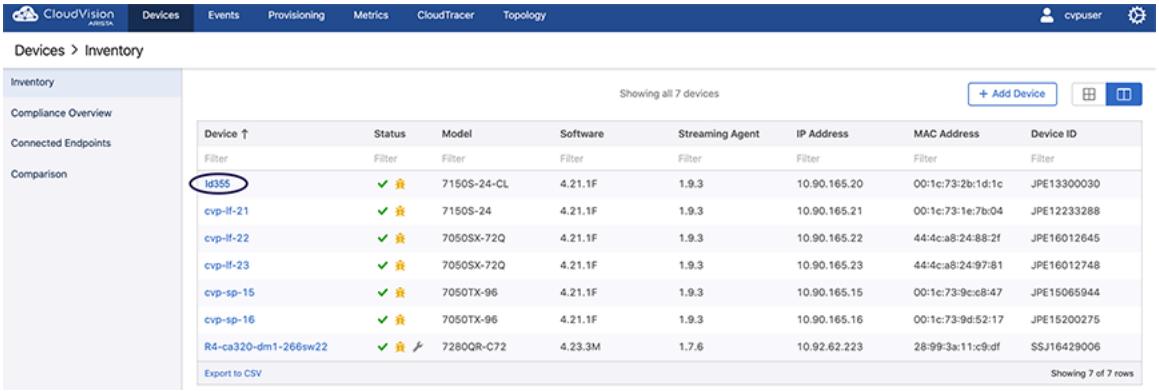

Figure 169: List of Registered Devices

The newly registered devices are shown in the undefined container in the **Network Provisioning** view.

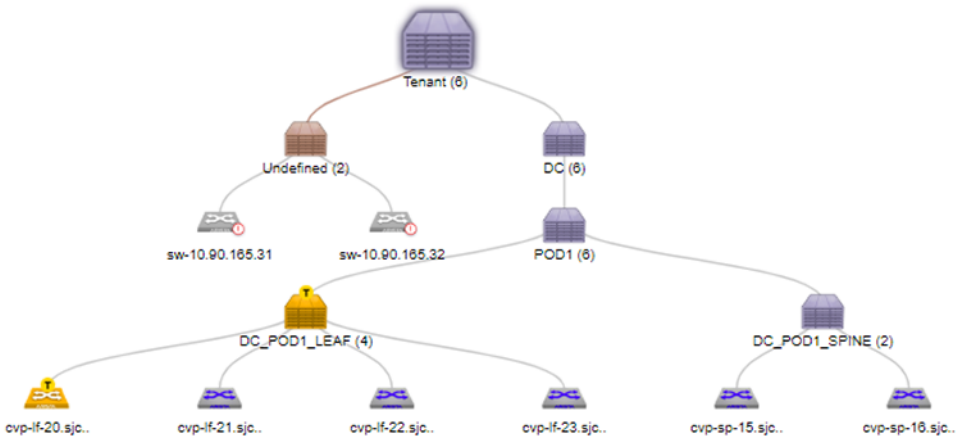

Figure 170: Registered Devices in the Network Provisioning View

11.4.4 Moving Devices from one Container to Another Container

Moving devices from one defined container to another is a method you can use to add devices to a container in the CVP topology. You use this method when you want to add devices to a container, and the device you want to add is currently under another container in the CVP topology. This method involves locating the device to be moved, and then moving it to the destination container. Containers that receive the imported devices are called destination containers.

There are three options you can use to move devices. They are:

- Option 1
- Option 2
- Option 3

11.4.4.1 Option 1

1. Locate the device.
2. Right-click the device and choose **Move**.

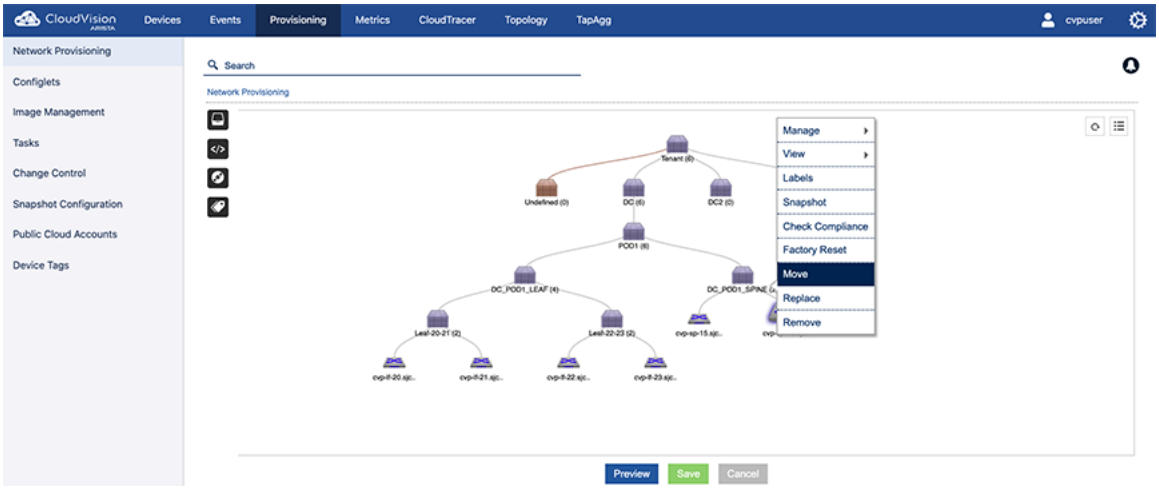

Figure 171: Selecting the device to be moved (option 1)

3. Select the destination container from the drop-down menu.
4. Save the session to move the device to the destination container.

11.4.4.2 Option 2

1. Locate the container that has the device you want to move.
2. Right-click the container and choose **Show All Devices**. This will load the inventory of all the devices under the container.
3. Locate the device to be moved.
4. Right-click the device and choose **Move**. After moving there will be a "T" icon to indicate the move has been tasked. (The task won't automatically be executed.)

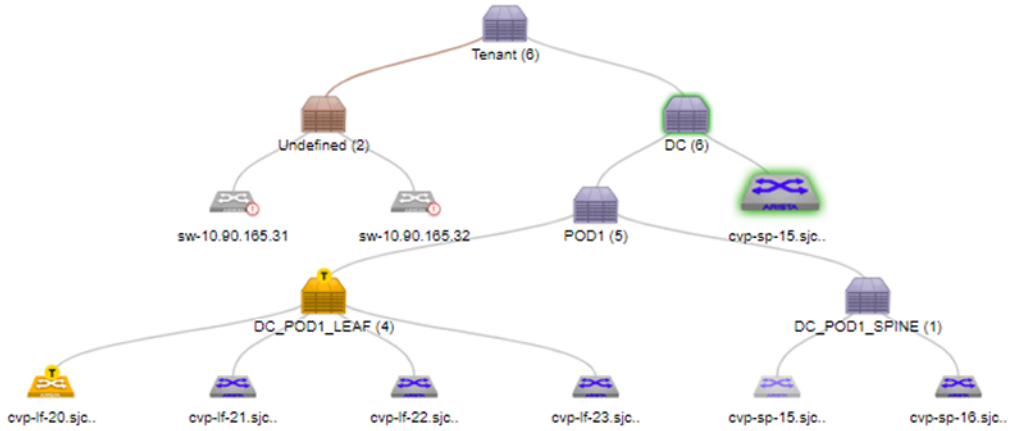

Figure 172: Device with pending move task (option 2)

5. Go to Tasks and explicitly execute the move task. After the task has been executed, the "T" icon is removed.

11.4.4.3 Option 3

1. Locate the container that has the device you want to move.
2. Right-click the container and choose **Manage > Device**. This will load the inventory of all the devices under the container.
3. Select the device to be moved and click **<–>** to choose the destination container.
4. From the popup menu, select the destination container and click **OK**. This will provision a move for the device

11.4.5 Removing a Device from a Container

A device can be removed from a container. Removing a device from the container will:

- Remove the device from parent container.
- Clear all information about the device in the CloudVision Portal.
- Stop any monitoring of the device.

There are three options you can use to remove devices. They are:

- Option 1
- Option 2
- Option 3

11.4.5.1 Option 1

1. Locate the device.
2. Right-click the device and choose **Remove**.

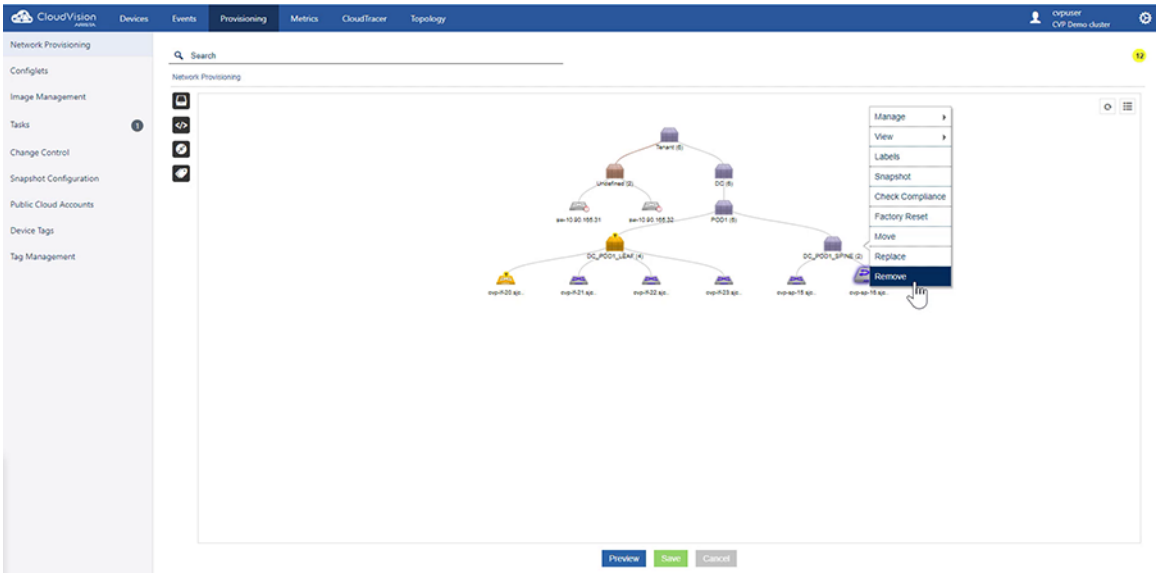

Figure 173: Removing a device (option 1)

11.4.5.2 Option 2

This option is available only for topology views.

1. Locate the parent container.
2. Right-click the container and choose **Show All Devices**. All the devices under the container are listed.

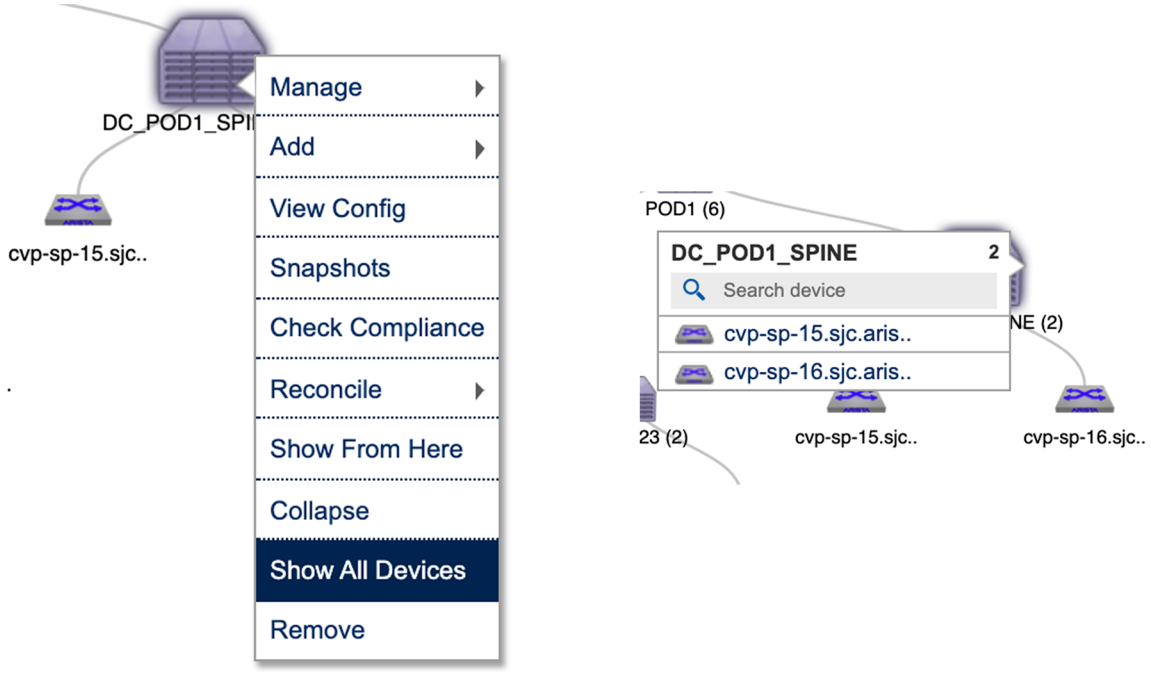

Figure 174: Selecting the device to be removed (option 2)

3. Select the device you want to remove.
4. Right-click the device and choose **Remove**. The device is removed from the Network Provisioning view.

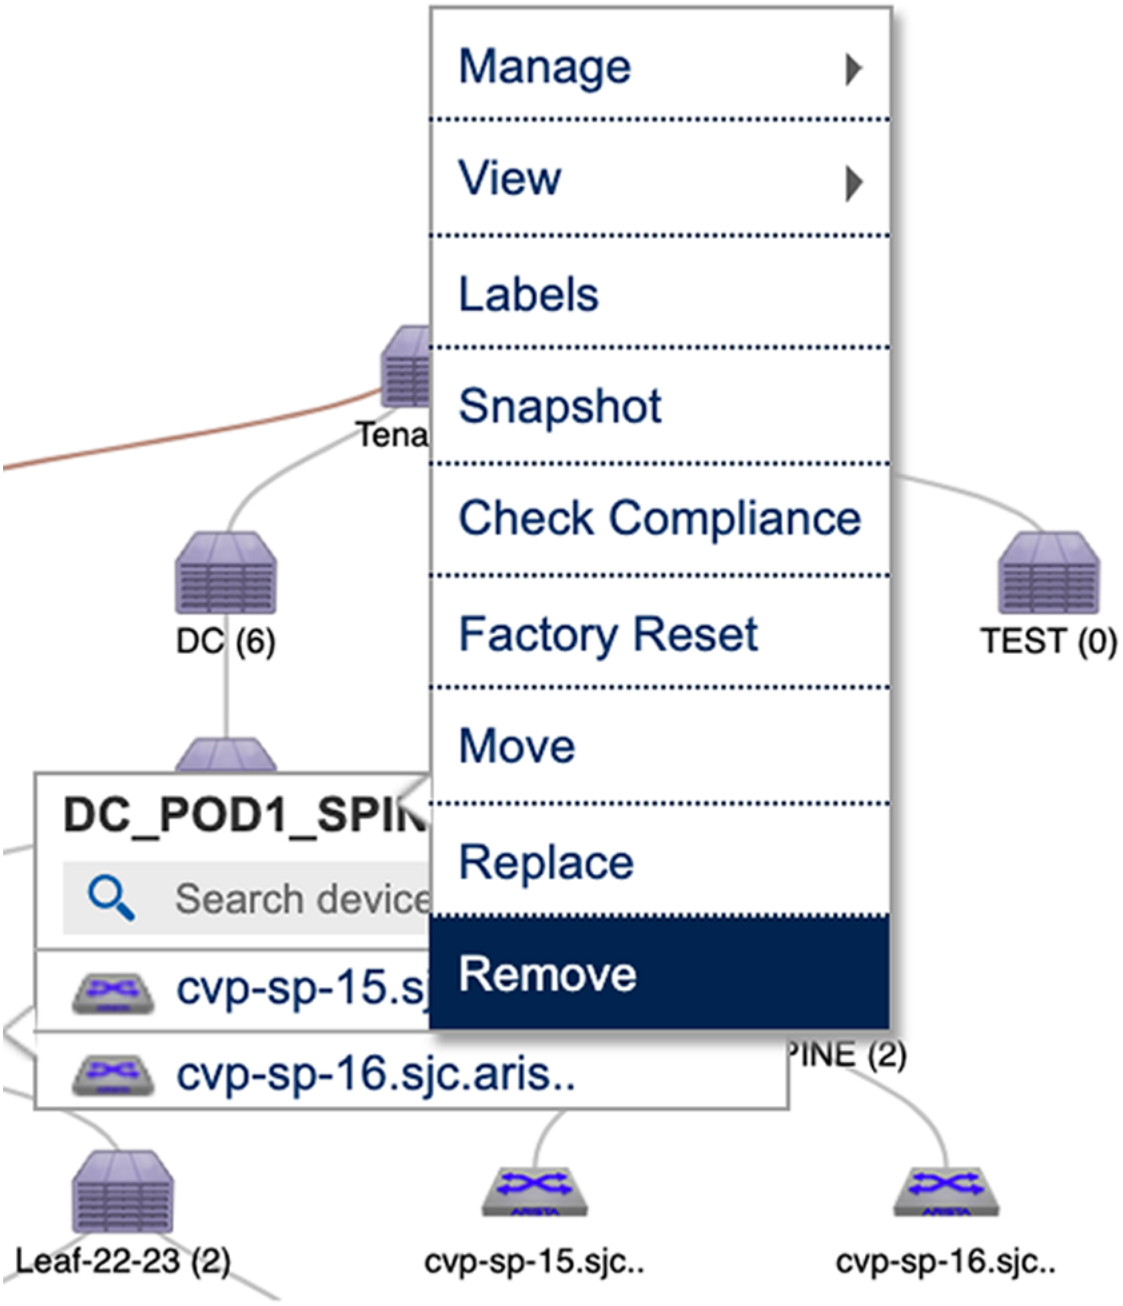

Figure 175: Removing the device (option 2)

11.4.5.3 Option 3

This option is available only for the list view of the Network Provisioning screen.

1. Locate the parent container.
2. Right-click the container and choose **Manage > Device**. This will load the inventory of all the child devices under the container.

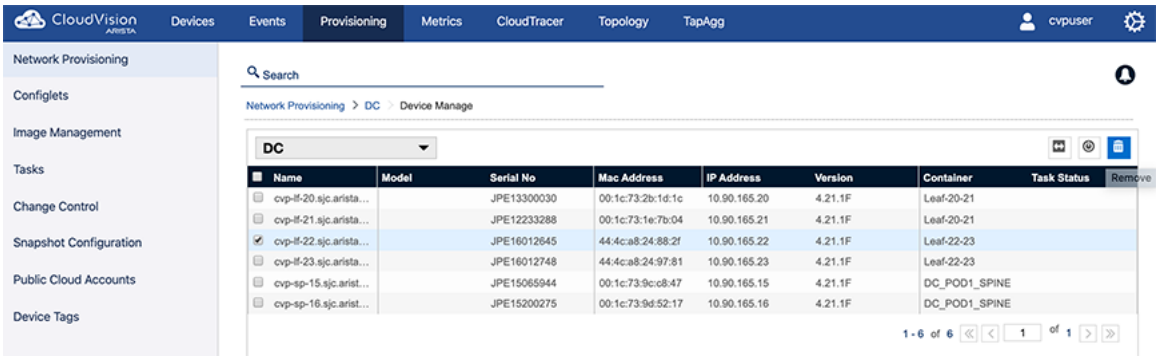

Figure 176: Remove device from the container (option 3)

3. Select the device you want to remove and then click **Remove**. On saving the session, a task will be spawned to reset the selected device.

11.4.6 Device Factory Reset

When resetting a device:

- The device will be removed from the parent container.
- The running configuration of the device will be flushed.
- Device will reboot with ZTP mode enabled.
- Device will be identified under undefined container.

There are three options you can use to move devices. They are:

- Option 1
- Option 2
- Option 3

11.4.6.1 Option 1

1. Locate the device.
2. Right-click the device and choose **Factory Reset**.

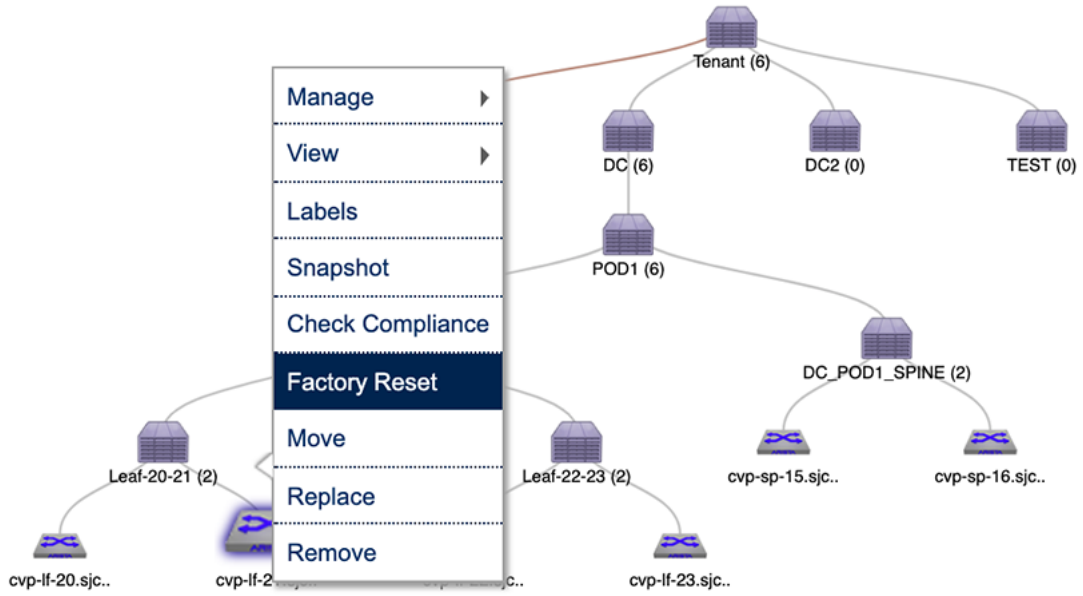

Figure 177: Resetting the device (option 1)

11.4.6.2 Option 2

1. Locate the parent container.
2. Right-click the container and choose **Show All Devices**. This will list all the devices under the container.

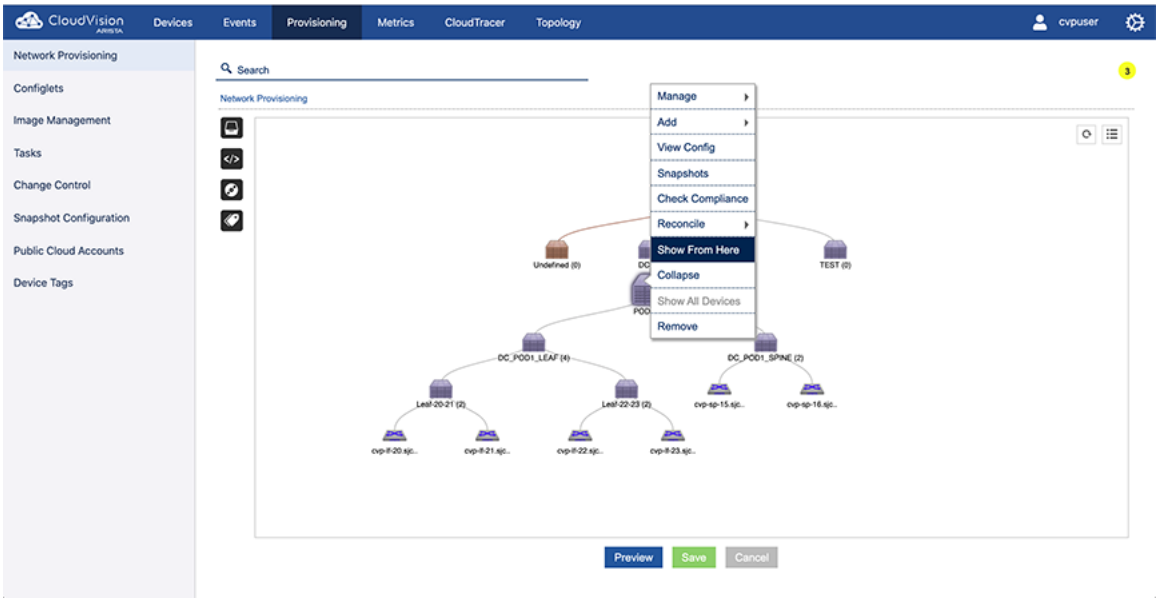

Figure 178: Showing all devices during factory reset (option 2)

3. Right-click the device and choose **Factory Reset**.

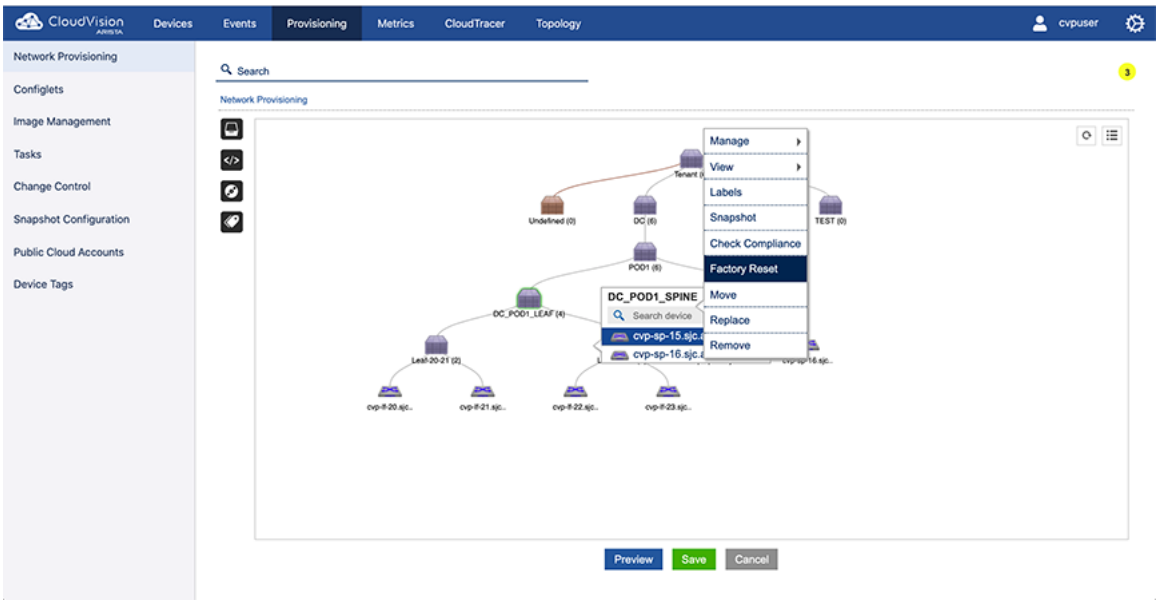

Figure 179: Resetting the device (option 2)

11.4.6.3 Option 3

1. Locate the parent container.
2. Right-click the container and choose **Manage > Device**. This will load the inventory of all the child devices under the container.
3. Select the checkbox of the device to be reset, and click the **reset** icon. On saving the session, a task will be spawned to reset the selected device.

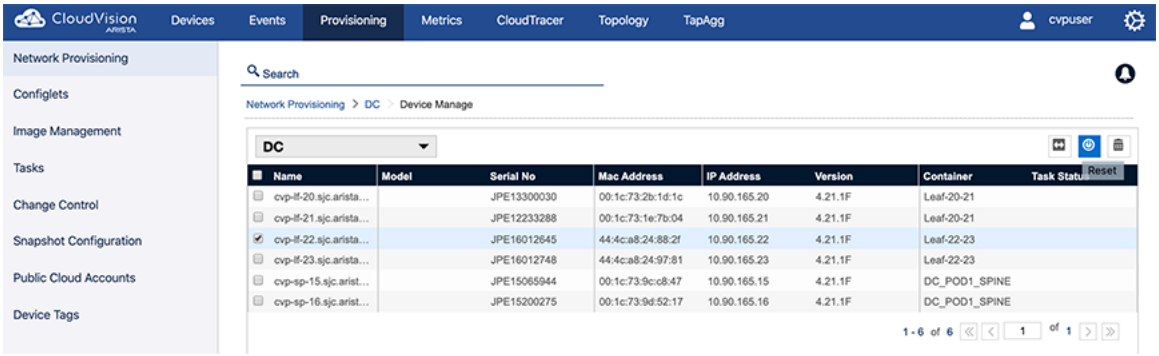

Figure 180: Selecting the device and resetting it (option 3)

11.4.7 Replacing Switches Using the ZTR Feature

The Zero Touch Replacement (ZTR) feature enables you to replace switches without having to configure the new switch. When you replace a switch using this feature, the new switch assumes the identity (IP), image, and configuration of the old switch. You use the Network Provisioning screen to replace switches using the (ZTR) feature.

Pre-requisites: Before you can begin the process to replace a switch using ZTR, make you must complete the following steps:

1. Make sure that the old switch is physically powered down and is not physically connected to the network.
2. Physically connect the new switch to the network exactly as the old switch was connected.

3. Power on the new switch.
4. Make sure the new switch comes up using ZTP, and that it shows up in the undefined container as an available resource.

Complete these steps to replace a switch using ZTP:

1. Go to the **Network Provisioning** screen.
2. Right-click on the old switch, and select **Replace**. This initiates ZTR, and opens the **Undefined Device** screen.

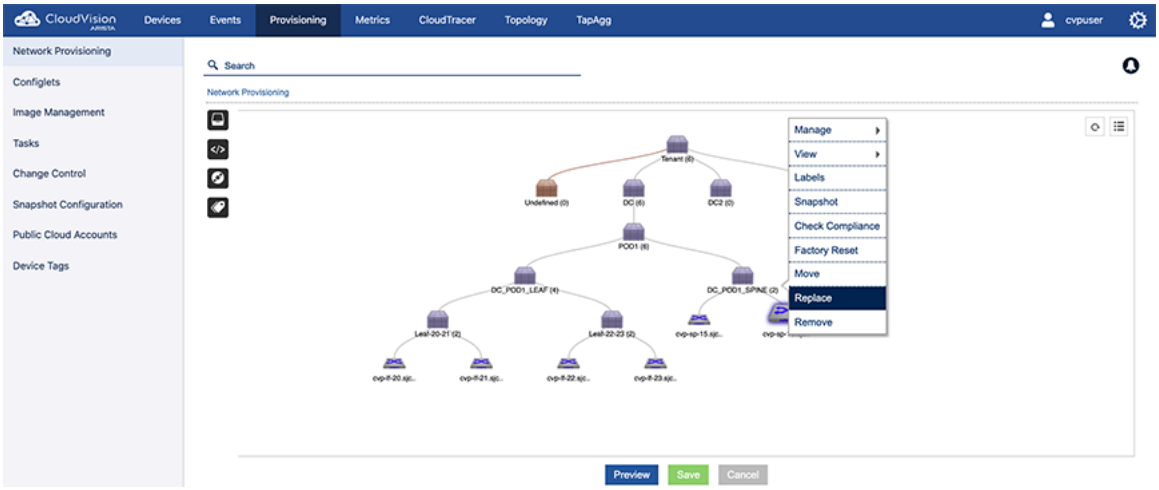

Figure 181: Selecting the switch to be replaced

3. Select the new switch by checking the checkbox next to the Serial No. column, and then click **Replace**.

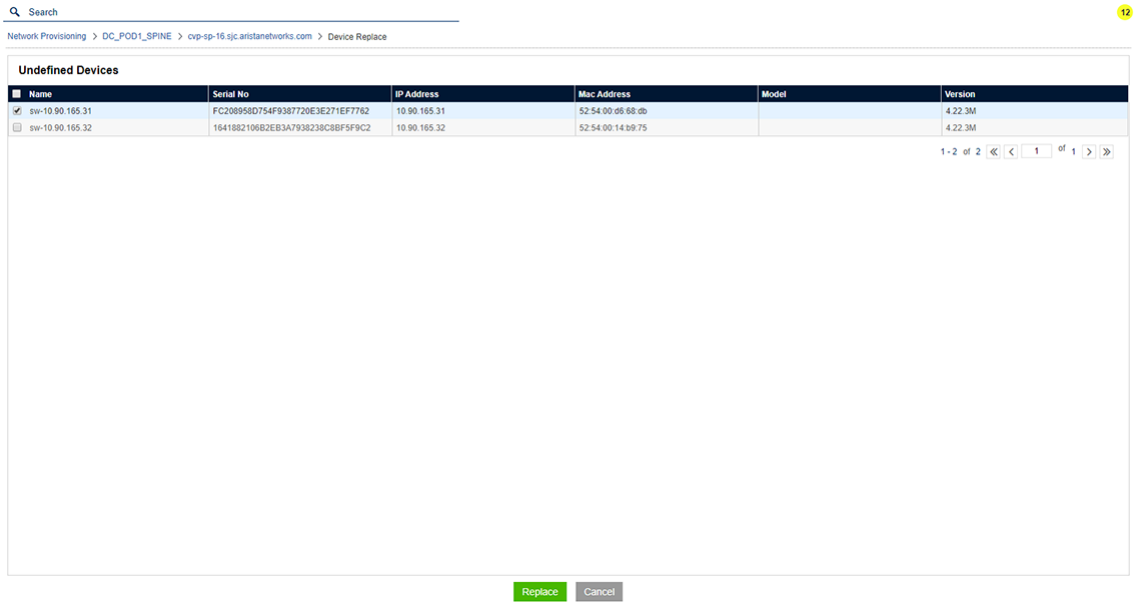

Figure 182: Selecting the new device and replacing the old device

4. In the Network Provisioning screen, click **Save**. A task icon "**T**" shows on the old switch, indicating that a task to replace it has been scheduled. Also, an "**R**" icon shows on the new switch, indicating that it is the replacement switch for a scheduled ZTR task.

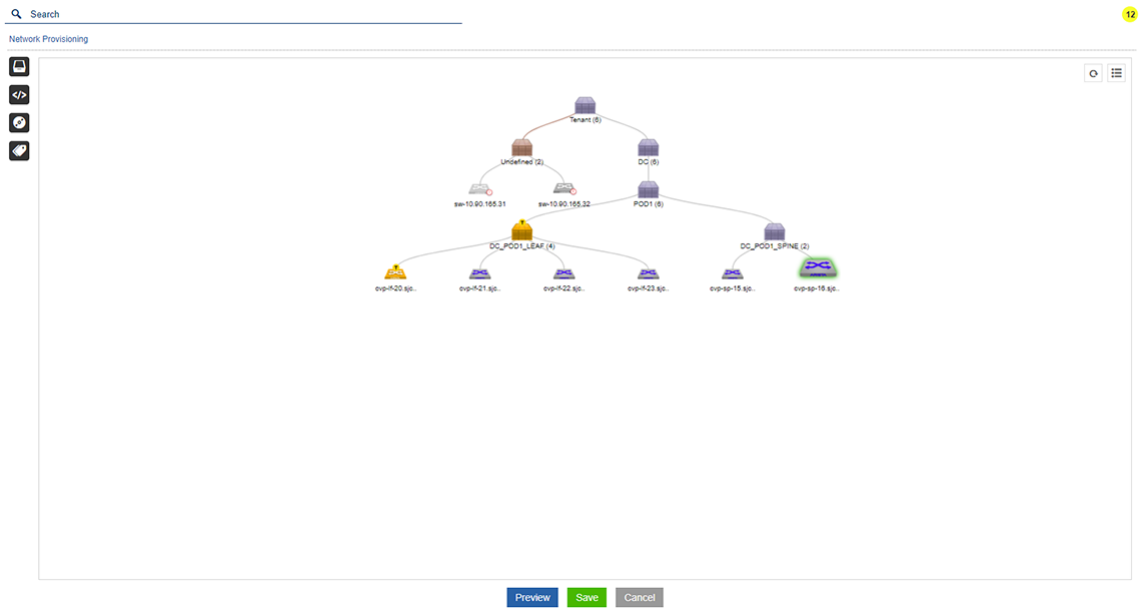

Figure 183: Topology view showing device with pending replace task

5. Go to the **Tasks** screen.
6. Select the task and click the play icon to execute the task.

   While the task is executing, you can open the logs for the task to view how ZTR manages the replacement. ZTR first pushes the old switch's image and configuration to the new replacement switch, and then initiates the reboot.

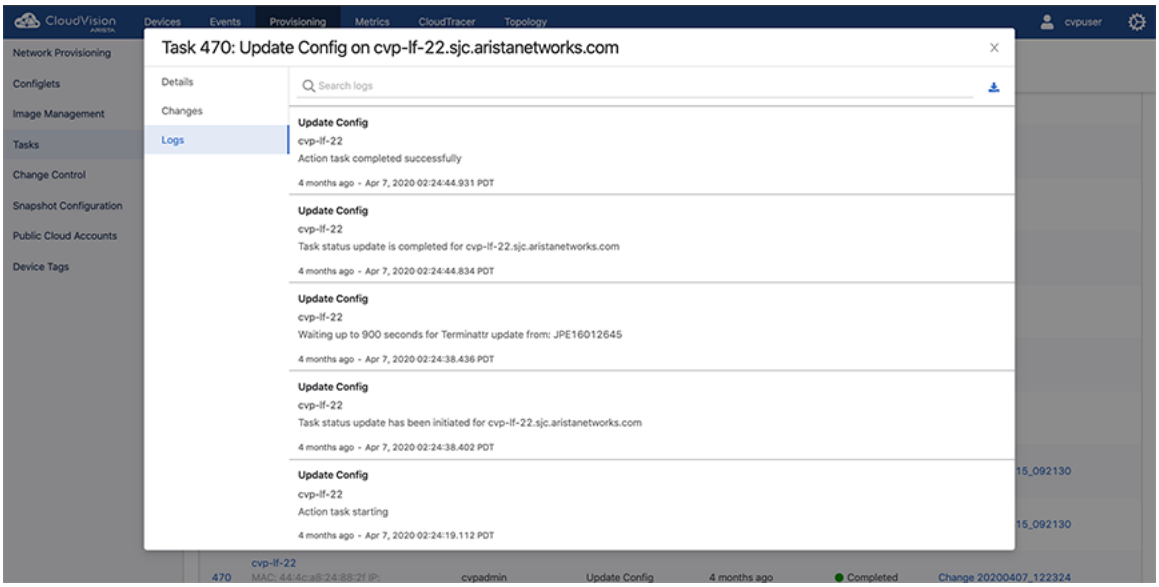

Figure 184: Task log showing processing of device replacement

11.4.8 Managing Configurations

CloudVision Portal (CVP) enables you to manage configurations by assigning configurations to containers and to devices. Configurations that you assign to containers are applied to all devices under the container's hierarchy. CVP also enables you to easily view the configuration currently assigned to containers and devices.

- Applying Configurations to Containers

- Viewing the Configuration Applied to Devices
- Applying Configurations to a Device

11.4.8.1 Applying Configurations to Containers

Applying configurations to containers involves adding Configlets to containers or removing Configlets from containers.

Adding Configlets

1. Locate the container.
2. Right-click the container and choose **Manage >Configlet**. This will open the window display the inventory of configlets.
3. Select the configlet and click **Update**. This will provision configlet add for the container and all the devices under it.

Removing Configlets

To remove the configlet inventory from a container.

1. Locate the container.
2. Right-click the container and choose **Manage>Configlet** .
3. Remove the configlets.
4. Click **Update**.

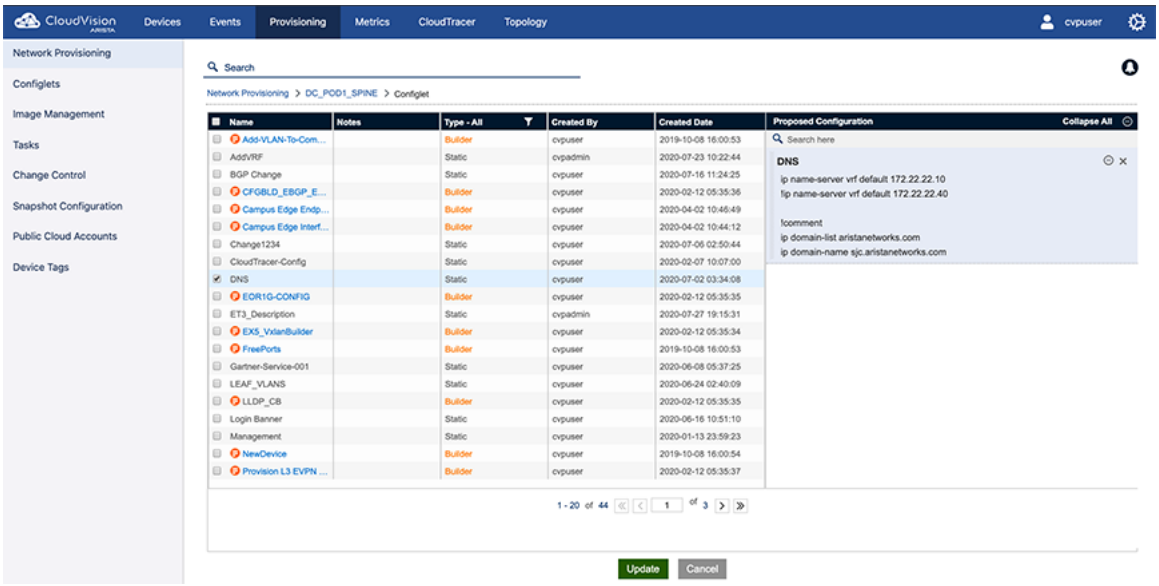

Figure 185: Remove the configlet and select Update

11.4.8.2 Applying Configurations to a Device

Applying configurations to devices involves adding Configlets to devices.

Note: When you update a device configuration using configlets, CVP replaces the entire device configuration with the Designed Configuration for the device. For new devices with pre-existing configurations added into CVP, you must explicitly perform a one-time reconciliation to save the desired device-specific running configuration in CVP. If you do not, that configuration may be lost, or the configuration update task may fail (see Reconciling Device Configurations at the Device Level).

Adding Configlets

1. Select the device and choose **Manage > Configlets**.

This loads the configlet inventory screen.

2. Select the configlets.

   You are required to validate the configuration.

3. To validate the configurations, select **Validate**.

   The validation screen will be loaded.

4. Select **Save** to propose a Config Assign action.

   When saving the session, this will spawn a Config Assign task.

11.4.8.3 Viewing the Configuration Applied to Devices

CloudVision Portal (CVP) enables you to use the **Network Provisioning** screen to view the configuration (Confliglets) currently assigned to devices. When you view the Configlets, you can also see which Configlets are inherited from Containers, and which are applied directly to the device.

Complete the following steps to view the Configlets applied to a device.

1. Go to the **Network Provisioning** screen.
2. Make sure you are using the topology view, not the list view.
3. Click on the device in the topology.
4. Click the Configlet icon.

   The Configlets applied to the device are listed in a drop-down list.

   - If a Configlet is inherited from a Container to which the device belongs, the Container icon appears in front of the Configlet name.
   - If a Configlet is directly applied to the device, no Container icon is shown next to the Configlet name.

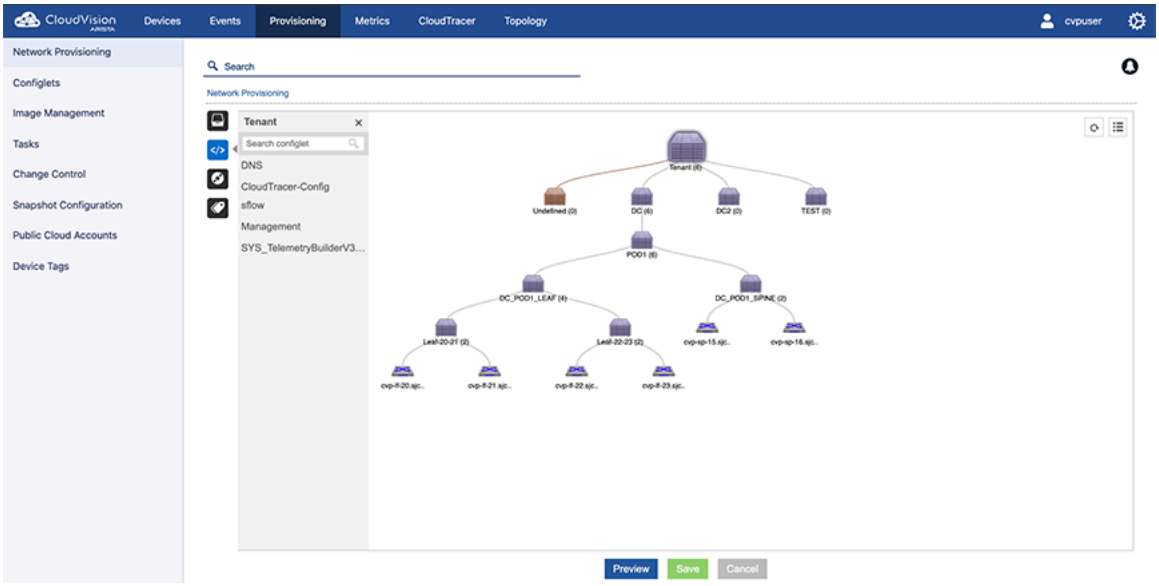

Figure 186: Viewing the Configlets applied to a device

11.4.8.4 Rolling Back Configurations Assigned to a Device

CloudVision's Network Rollbacks feature enables you to restore a previous configuration to devices. You can apply the rollback to all the devices in a container, or to single devices. When you rollback a container or device, you select the date and time for the rollback and whether you want to rollback the configuration or EOS image (or both).

See Rolling Back Images and Configurations for details.

11.4.9 Configuration Validation

The validation screen consists of three panes.

- Pane 1: Shows the proposed configuration.
- Pane 2: Shows the designed configuration. (This shows how a resulting running configuration will look like after successful configuration push.)
- Pane 3: Shows the current running configuration of a device.

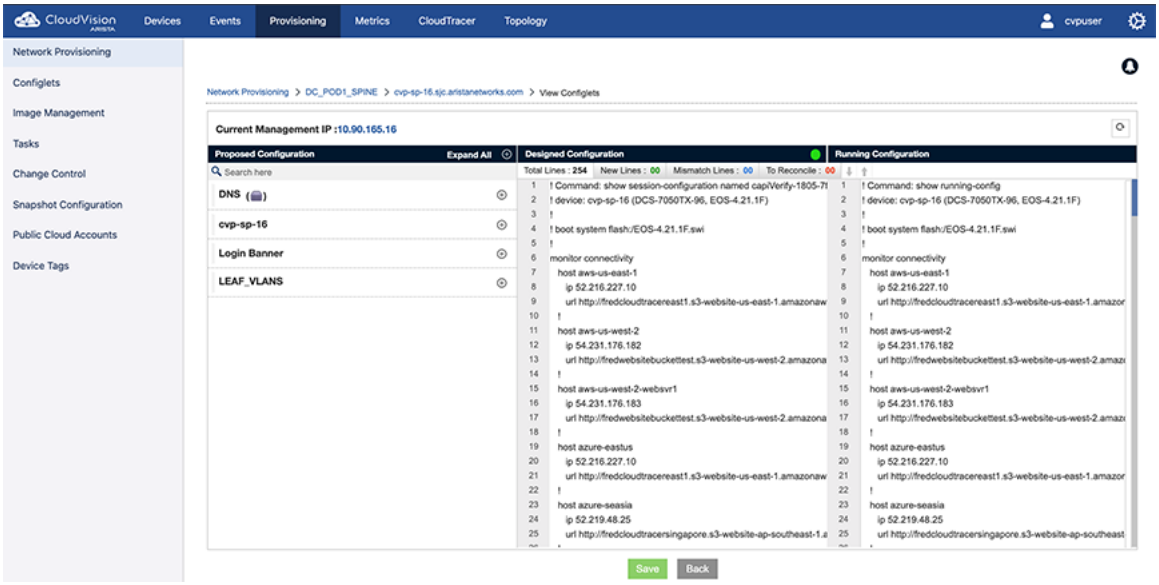

Figure 187: Validating your configurations

11.4.10 Using Hashed Passwords for Configuration Tasks

Some EOS commands take a password or a secret key as a parameter. There are usually two ways of passing EOS command parameters:

- As plain text.
- As a hashed string.

Note: Because EOS always returns the hashed version of the command in its running configuration, using the plain text version of commands in Configlets results in the following issues:

- CVP shows that there are configuration differences that need reconciling, even if there are none.
- Compliance checks show devices to be out of compliance.

To avoid these issues, you should use the hashed version of EOS commands in Configlets (for example, use `ntp authentication-key 11 md5 7 <key>` instead of `ntp authentication-key 11 md5 0 <key>`). Using the hashed versions of commands also keeps the real password hidden.

11.4.11 Reconciling Configuration Differences

CloudVision enables you to reconcile differences between the designed (managed) configuration and running configuration on devices so that CVP is maintaining the full configuration of each device.

Related topics:

- Key Terms
- Reconciling Device Configurations at the Device Level
- Reconciling Device Configuration Differences at the Container Level

11.4.11.1 Key Terms

Reconcilable differences	Configuration differences between the designed configuration and the running configuration, which do not conflict with the configuration in any configlets, other than the reconcile configlet.
Reconcile configlet	A specially marked device configlet that is system generated and used to store reconcilable differences in order for the designed configuration to match the running configuration.

Reconciling device configuration differences does not require a task, because there is no configuration to be pushed out to the device. Reconcilable differences are only adjusted in the reconcile configlet, to match the running configuration. Because of this, there is no task pushed to change the running configuration.

When you reconcile device configuration differences, you add the reconcilable differences found in the running configuration to the reconcile configlet of the designed configuration.

For details on reconciling device configuration differences, see:

- Reconciling Device Configurations at the Device Level
- Reconciling Device Configuration Differences at the Container Level

11.4.11.2 Reconciling Device Configurations Differences at the Container Level

CloudVision enables you to reconcile device configuration differences for all devices under the hierarchy of a selected container, instead of having to initiate this device by device.

Note: The designed configurations of devices in the container that do not have reconcilable differences are not changed.

For devices that have reconcilable differences, the lines or commands on the device that are not present in the designed configuration are pulled into the reconcile configlet for that device in one of two ways:

- Using the existing reconcile configlet that is specific to that device.
- Creating a new reconcile configlet that is specific to that device. This is done when there is no existing reconcile configlet specific for the device. The system automatically creates a unique name for the configlet.

A green checkmark beside the configlet indicates it as the reconcile configlet for the device.

Complete the following steps to reconcile device configuration differences for a container:

1. Go to the **Network Provisioning** screen.
2. Locate the container in the topology where you want to reconcile the configurations of all devices under that container hierarchy.
3. Right-click the container, hover the cursor on Reconcile, and click either **Reconcile All** or **Reconcile New**.

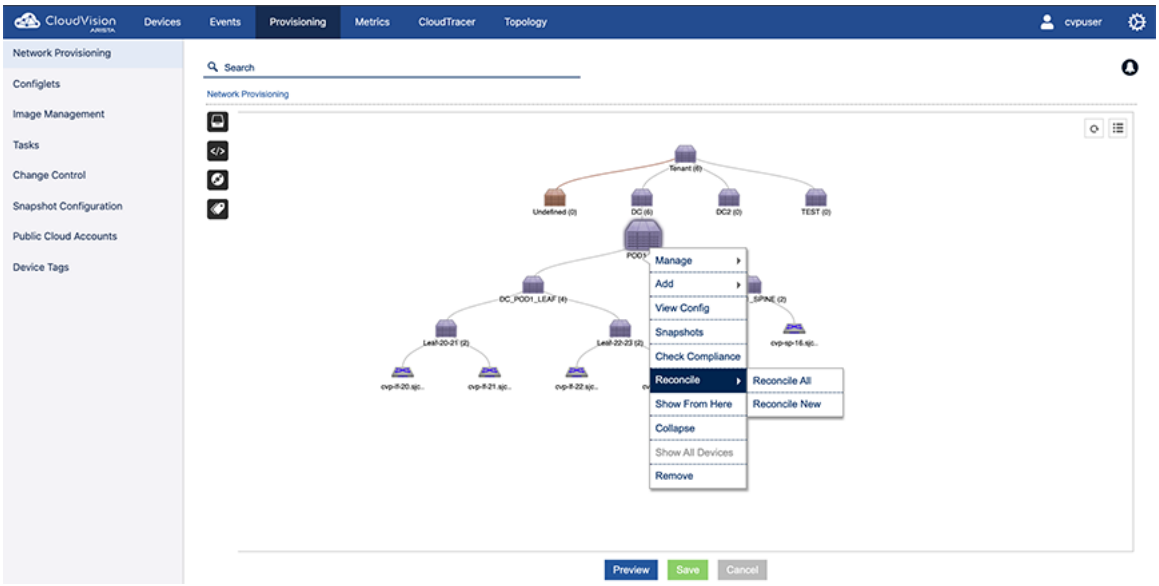

Figure 188: Device configuration reconciliation at the container level

The **Reconcile New** option reconciles only the configuration lines that exist on the device, but not in the designed configuration.

The **Reconcile All** option reconciles new lines and also lines that differ in designed and running configurations. This usually brings the device into compliance because the resulting designed configuration will be identical to running configuration. However, there can be cases where in spite of reconciling device configuration lines, the designed configuration may not end up identical to running configuration. In these cases, no changes are made to the reconcile configlet. Arista recommends to go through the device-level reconcile process (See Reconciling Device Configurations at the Device Level), and select the desired lines.

Note: The bell icon in the upper right corner turns yellow to indicate unread notifications.

4. (Optional) To view the notification for the reconciliation, click the bell icon. The notification list appears showing the container-level configuration reconciliation, and any other unread notifications.

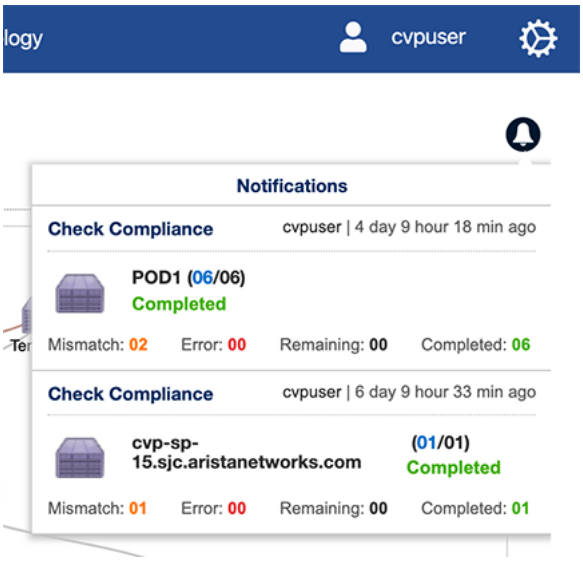

Figure 189: List of unread notifications

11.4.11.3 Reconciling Device Configurations at the Device Level

CloudVision enables you to reconcile device configuration differences at the device level (specific, individual devices). Configuration differences at the device level occur when there are reconcilable differences in the running configuration of the device.

The **Configuration Validation** screen shows details of the configuration differences. When the system identifies a reconcilable difference, the Reconcile option becomes available, and the extra reconcilable configuration is listed in a text editor on the screen.

Reconcile Configlets

You use a type of configlet called a reconcile configlet to reconcile device configuration differences at the device level. A reconcile configlet is a configlet for a single specific device, and is explicitly marked as the reconcile configlet for that device. The reconcile configlet for a device contains the additional running configuration for that device.

Note: There is only one reconcile configlet for any device. It is the only configlet that contains the additional running configuration for the device.

Every time a device-level or a container-level reconcile is performed, the reconcile configlet for each device included in the reconcile action is modified to include the extra running configuration.

To reconcile device level configuration, perform the following steps:

1. If required, select additional lines from running configuration to reconcile.
2. Click the blue **Reconcile** button to add the reconcilable configuration in the running configuration to the reconcile configlet of the designed configuration.

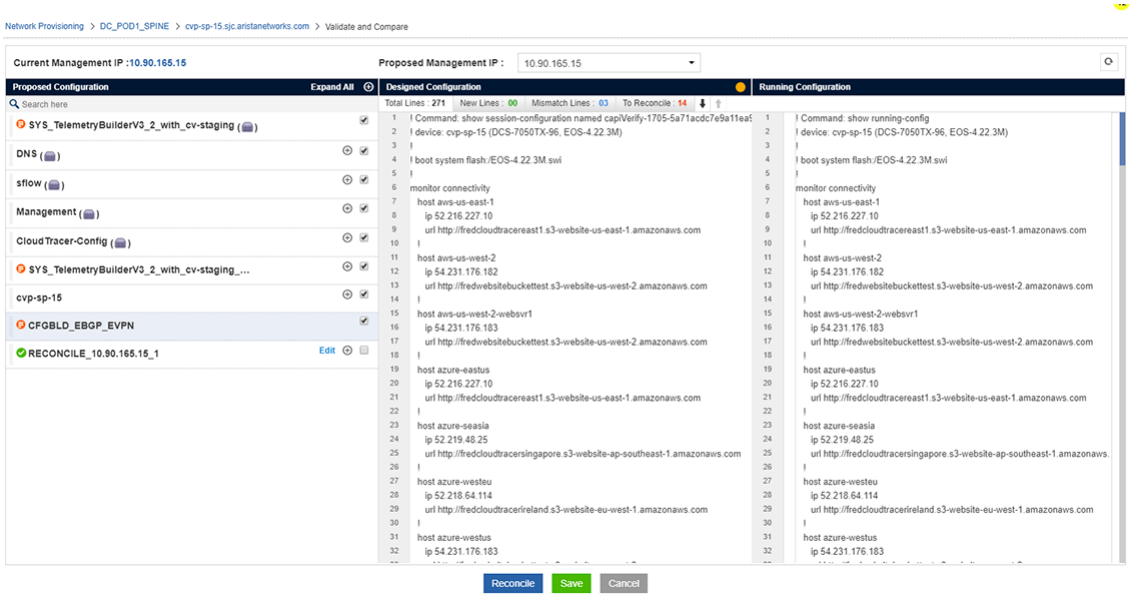

Figure 190: Configuration validation screen showing device-level configuration differences

3. (Optional) Click **Edit** next to the configlet name to edit or rename the reconciled configlet.
4. (Optional) Click the reconcile disk icon next to the configlet name to save the reconciled configlet with the extra commands present in the running configuration.

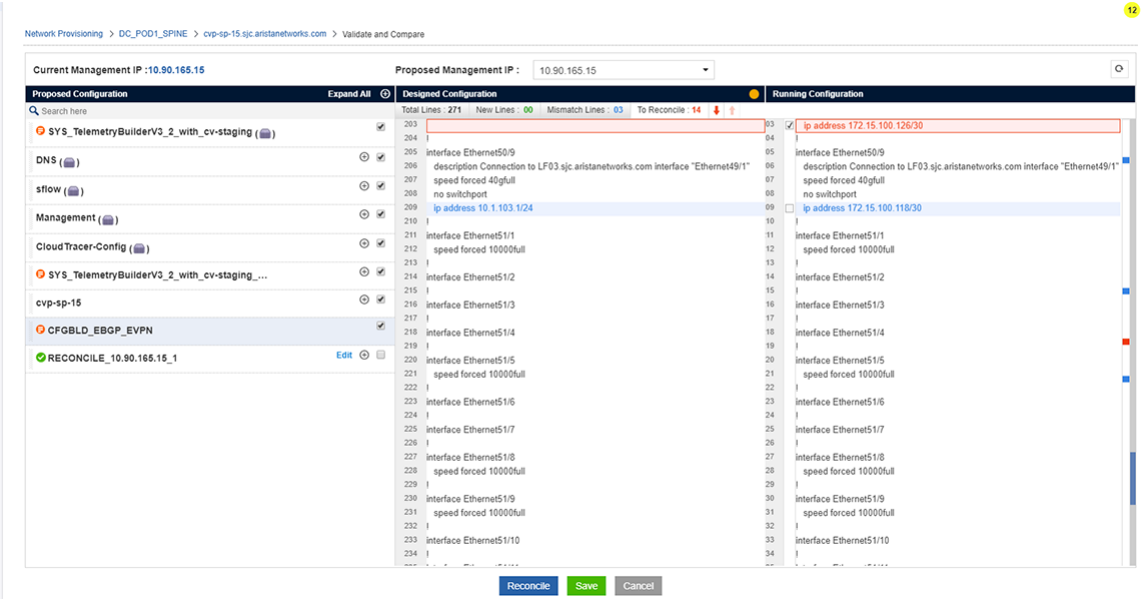

Figure 191: Reconcile Disk icon

Note: CVP will not execute pushing a configuration that causes CVP to lose connectivity with the device if the management interface or IP is missing in the configuration. When the task is executed, it will fail.

5. Click **Save**.

11.4.12 Managing EOS Images Applied to Devices

CloudVision enables you to efficiently manage the EOS images of devices by assigning image bundles to containers or devices in the current CloudVision network topology. An image bundle assigned to containers are automatically applied to all devices under that container.

The image bundle you want to apply must already exist in the set of current EOS image bundles.

The following tasks are involved in managing the EOS image bundles assigned to devices:

- Applying an Image Bundle to a Container
- Viewing the Image Bundle Assigned to Devices
- Applying an Image Bundle to a Device
- Setting up an Image Bundle as the default for ZTP
- Rolling Back Configurations Assigned to a Device

11.4.13 Rolling Back Images and Configurations

CloudVision's Network Rollbacks feature enables you to restore a previous EOS image and configuration to containers and devices. You can apply the rollback to all the devices in a container, or to single devices. When you rollback a container or device, you select the date and time for the rollback and whether you want to rollback the EOS image or configuration (or both).

CloudVision supports rollback to any previous point in time irrespective of captured snapshots. However, rollback is possible to a point that is far beyond the CloudVision Cluster update to *2018.2.0* only when your devices are upgraded to TerminAttr 1.4+ long before that.

Note: To help you select the desired rollback destination day and time, you can compare the image and running configuration differences between current and rollback times of all effected devices. The potential destination rollback date and time in the comparison is based on the destination rollback date and time you select.

11.4.13.1 Rolling Back Container Images and Configurations

Complete the following steps to apply a network rollback in containers:

1. Go to the **Network Provisioning** screen.
2. Right-click on the container you want to rollback, and then choose **Manage > Network Rollback**.

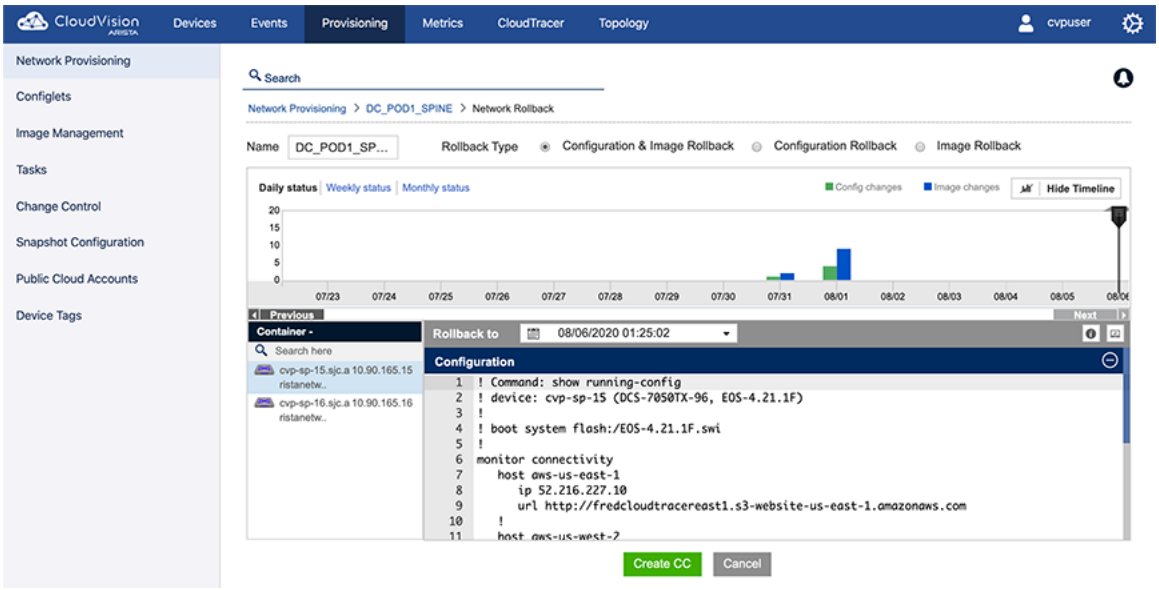

Figure 192: Network Rollback Screen

3. Using the Rollback Type: options near the top of the screen, select the type of rollback. The options are:
   - Configuration & Image Rollback (both the configuration and EOS image are rolled back)
   - Configuration Rollback (only the configuration is rolled back)
   - Image Rollback (only the EOS image is rolled back)
4. Either drag the vertical slider on the timeline to the desired date and select the time for rollback; or use the Rollback to menu for selecting rollback date and time (directly above the configuration pane on the left side).
5. Click the telemetry icon (directly above the configuration pane on the right side) for viewing the running configuration differences between current and rollback times.
6. If required, change the destination date and time for the rollback.
7. Click **Create CC** to create a Change Control (CC) record for the network rollback. CloudVision automatically creates a rollback task for each device in the rollback; and makes them part of CC.

Note: Rollback Change Controls are automatically assigned a unique name. You can rename the Change Control record by editing the Change Control record. Once the Change Control is created, it can be executed like any other Change Control.

11.4.13.2 Rolling Back Device Images and Configurations

Complete the following steps to apply a rollback in devices:

1. Go to the **Network Provisioning** screen.
2. Right-click on the device you want to rollback, and then choose **Manage > Rollback**.

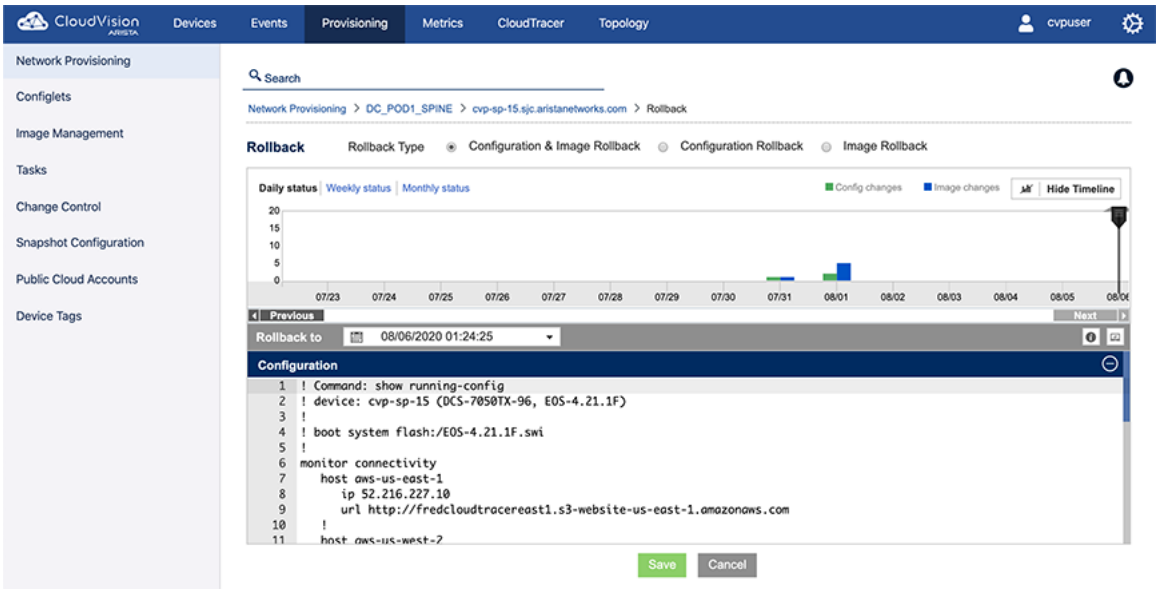

Figure 193: Device Rollback Screen

3. Using the **Rollback Type**: options near the top of the screen, select the type of rollback. The options are:
   - Configuration & Image Rollback (both the configuration and EOS image are rolled back)
   - Configuration Rollback (only the configuration is rolled back)
   - Image Rollback (only the EOS image is rolled back)
4. Either drag the vertical slider on the timeline to the desired date and select the time for rollback; or use the **Rollback to** menu for selecting rollback date and time (directly above the **configuration** pane on the left side).
5. Click the telemetry icon (directly above the **configuration** pane on the right side) for viewing the running configuration differences between current and rollback times.

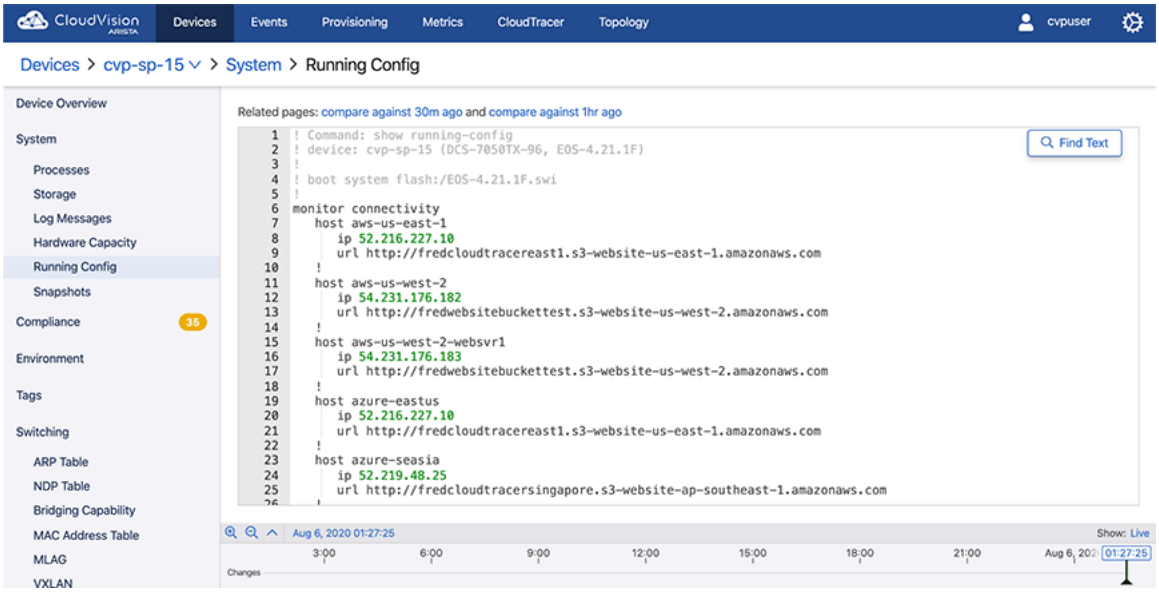

Figure 194: Differences in Running Configuration

The **Unified** tab displays running configuration differences in a single window with differences highlighted. The **Split** tab displays running configurations in different windows with differences highlighted.

6. If required, change the destination date and time for the rollback.
7. Click **Save** to create a task for the device rollback.

11.4.13.3 Rolling Back Configurations Assigned to a Device

CloudVision's Network Rollbacks feature enables you to restore a previous configuration to devices. You can apply the rollback to all the devices in a container, or to single devices. When you rollback a container or device, you select the date and time for the rollback and whether you want to rollback the configuration or EOS image (or both).

See Rolling Back Images and Configurations for details.

11.4.14 Device Labels

A label is simply defined as Text Tags. There are two types of label:

- System labels: Assigned automatically by the system.
- Custom labels: Defined and assigned by the user.
  - Users can assign custom labels to devices from the **Network Provisioning** screen.
  - A device can be tagged with one or more custom labels.
  - Labels can be used to filter the devices in the **Network Provisioning** screen.

11.4.14.1 System Labels

System labels are defined by the system and are automatically applied to and removed from devices based on the following characteristics of that device:

- Software version
- Software bundle
- Product model and family
- Assigned configlet name
- DANZ enabled
- MLAG enabled
- Parent container name

Note: System labels cannot be modified or removed by the user.

11.4.14.2 Custom Device Labels

You can create custom device labels and assign them to devices. The device labels you assign to a device show on the **Network Provisioning** screen next to the device.

11.4.14.2.1 Assigning an Existing Label to a Device

Complete these steps to assign an existing label to a device.

1. Select the device to be labeled.
2. Right-click the device and choose **Labels**.

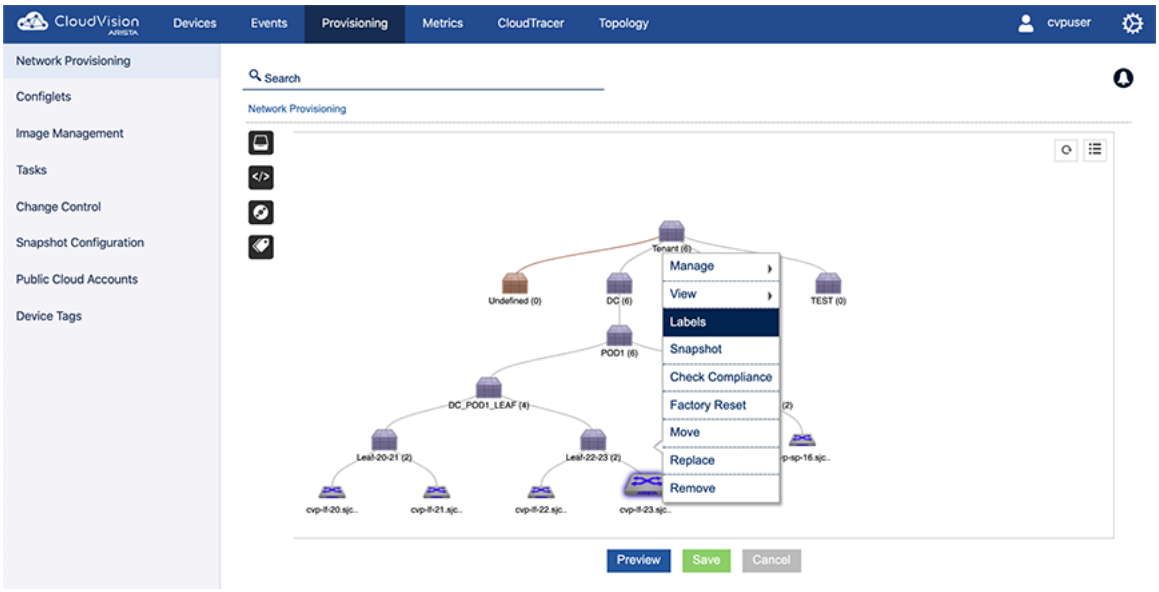

Figure 195: Choose Labels

The **Assign Label** pop-up menu appears, showing the available device labels.

3. Select the label to be applied and click **Save**.

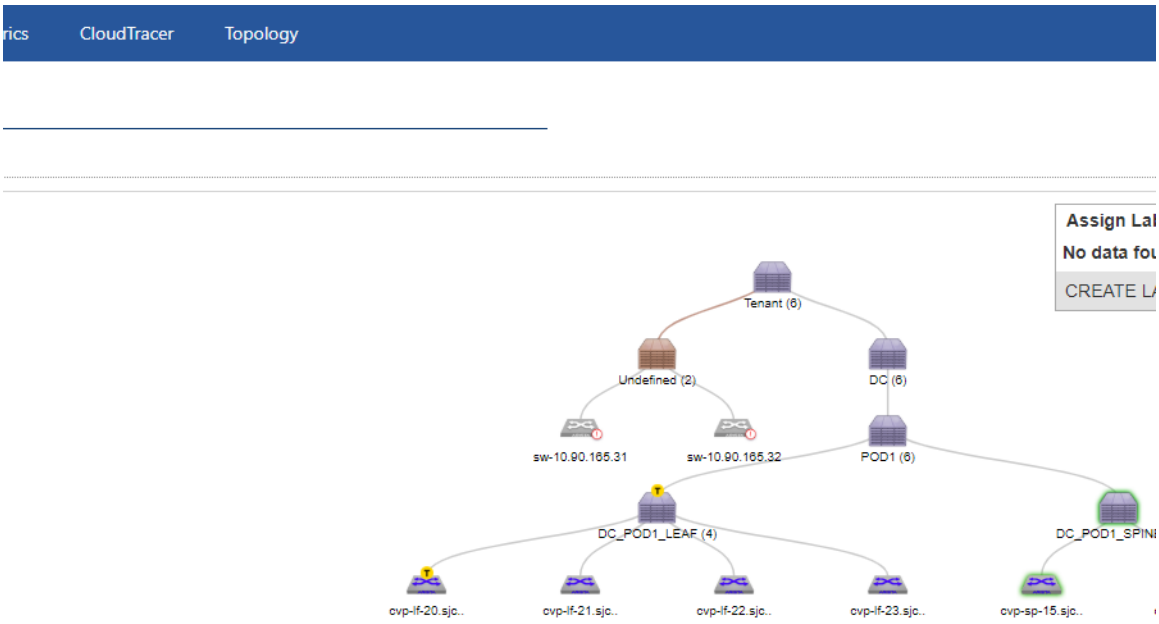

Figure 196: Assign Label

The selected label will be applied to the device.

11.4.14.2.2 Creating a Custom Label for a Device

Complete these steps to create a new, custom label to a device.

1. Select the device for which you want to create a new, custom label.
2. Right-click the device and choose **Labels.**

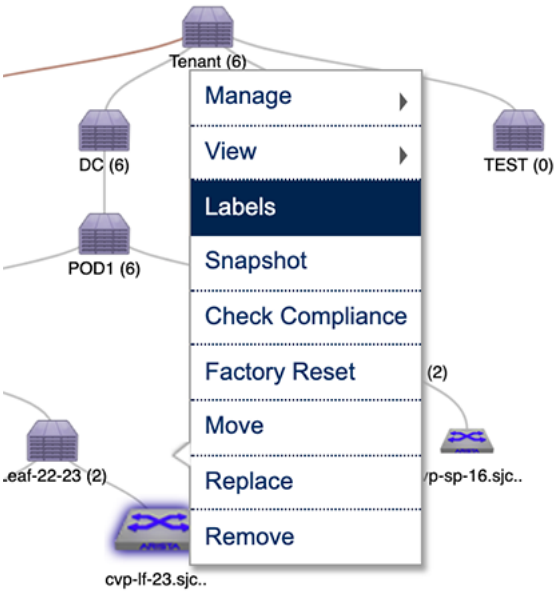

Figure 197: Choose Labels

The Assign Label pop-up menu appears, showing the available device labels.

3. In the pop-up menu, click on **CREATE LABEL**.

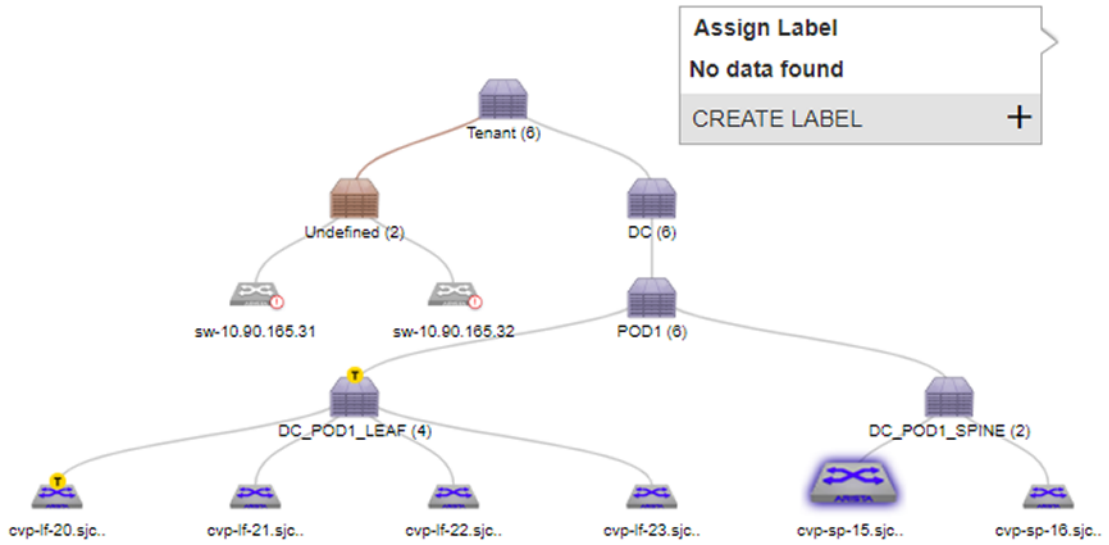

Figure 198: Create label Pop-up

The Create Label dialog appears.

4. Type the new, custom label for the device, then click **Save**.

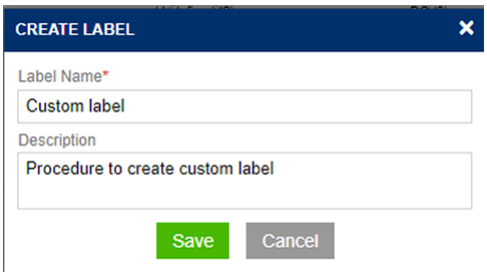

Figure 199: Create Label

The new label is created and is assigned to the device.

11.4.14.3 Left Pane Behavior in Network Provisioning View

The left pane in the topology view is used to display information on the resources assigned to a given device or container.

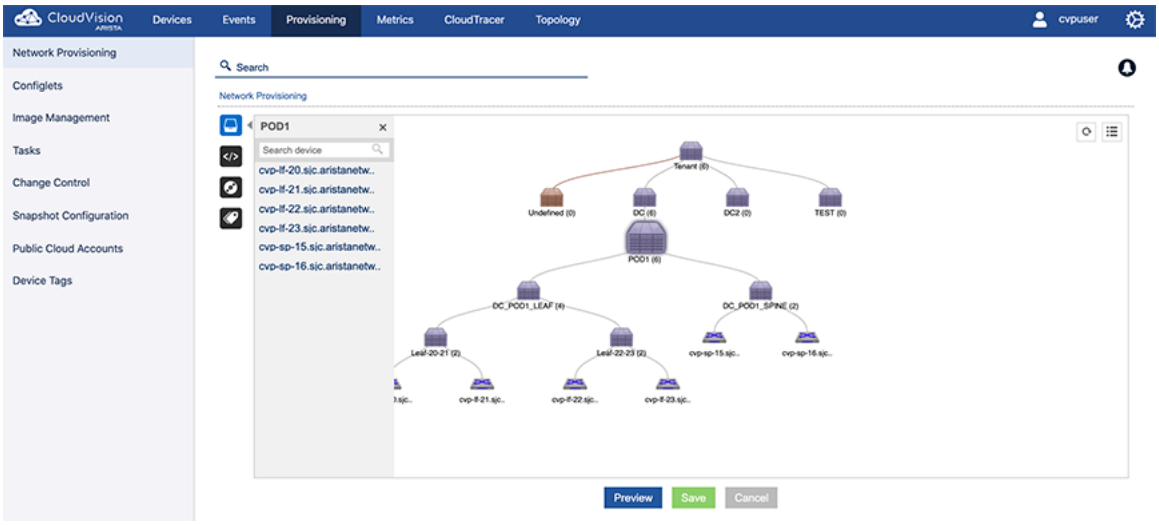

Figure 200: Left pane view

Opening and Closing the Left Pane

1. Double click the container or device to open the left pane.
2. Click the "**X**" button to close it.

11.4.14.4 Right Pane Behavior in List View

Similar to the left pane in the Network Provisioning View, the right pane in the List view is used to display information regarding resources assigned to the given device or container.

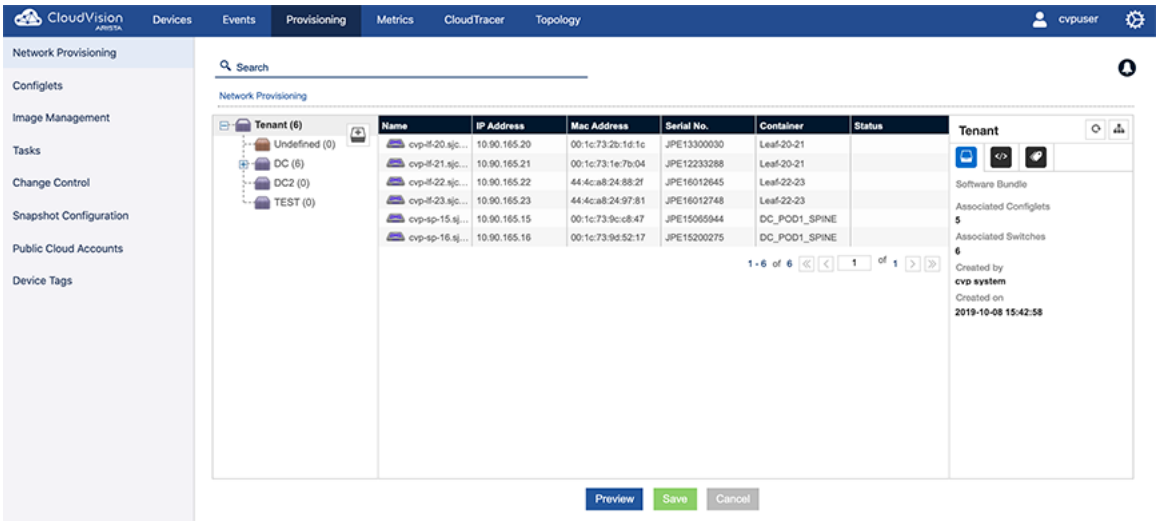

Figure 201: Right pane view

Right Pane - Summary Screen

The summary screen gives the information on Device model, IP address, MAC address and its EOS version.

For a Container, it displays the total number of devices associated, container “creation date” and “created by” details.

11.4.15 Viewing Containers and Devices

The Network Provisioning screen provides you with various options that enable you to easily control the topology view so that you can view containers and devices based on your needs.

The options you use are:

- **Expand / Collapse** (see Expanding and Collapsing Containers).
- **Show From Here** (see Show From Here).
- **Show Full Topology** (see Show Full Topology).

CloudVision Portal uses color coded icons to indicate compliance or access issues with devices. See Device Compliance Status Indicators and Device Access Alerts for more information.

11.4.15.1 Expanding and Collapsing Containers

Containers can be expanded and collapsed within the Network Provisioning topology view so that you can change the view as needed based on your needs.

You use the **Show From Here** and **Show Full Topology** options to expand or collapse containers shown in the **Network Provisioning** screen.

The **Expand and Collapse** option is only available for the **Network Provisioning** view. It is not available for the List view.

The default view mode for containers is expanded. When you choose **Expand/Collapse** option for a container, one of the following occurs, depending on the current view mode:

- A container currently in expanded (normal) view is collapsed.
- A container currently in collapsed mode is returned to expanded view mode (the default).

Complete these steps to expand or collapse a container view from the **Network Provisioning** screen.

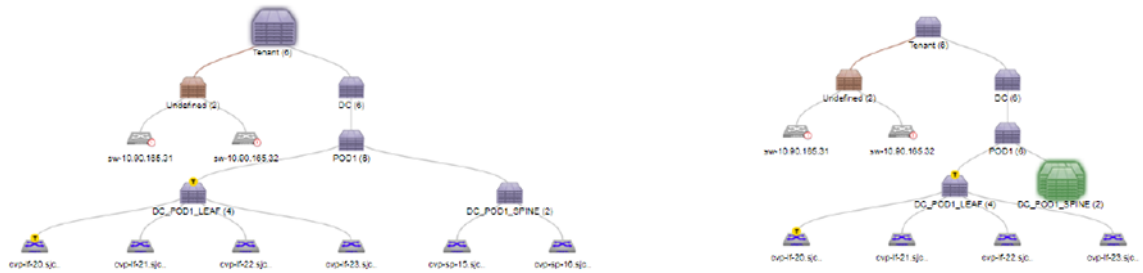

Figure 202: Expanded and collapsed view of a container

1. Select a container.
2. Right-click it and select the **Expand/Collapse** option.

11.4.15.2 Show From Here

The **Show From Here** option displays the topology with the selected container as the root. The hierarchy above the selected container will be hidden from the view allowing the user to only focus on the chosen container and the tree below it.

1. Select a container.
2. Right click **Show From Here** to display the option. The hierarchy from the selected container will be displayed.

11.4.15.3 Show Full Topology

The **Show Full Topology** option allows the user to get back to the full topology view. This option will be enabled for a particular container once the user uses the show from here option on it.

1. Select a container.
2. Right-click **Show Full Topology** to view the option.

11.4.16 Device Compliance

In CloudVision Portal (CVP), devices have a compliance status which indicates whether the running configuration and image of a device is different from the designed (managed) configuration and image for the device.

The possible device compliance statuses are:

- **Compliant:** Devices in which the running configuration and image are identical to the designed configuration and image for the device.
- **Non-compliant:** Devices in which the running configuration or image are different from the designed configuration or image for the device

CVP enables you to check devices to determine if they are non-compliant. It also provides device compliance status indicators so you can easily identify non-compliant devices and the functionality required to bring non-compliant devices into compliance. One process used to resolve the difference in running and designed configuration is referred to as reconciling.

For more information, see:

- Device Compliance Status Indicators
- Device Compliance Checks
- Reconciling Configuration Differences

11.4.16.1 Device Compliance Status Indicators

CloudVision Portal (CVP) provides device compliance status information in both the **Network Provisioning** screen and the **Inventory** screen (list view).

- Network Provisioning Screen Compliance Status Indicators
- Representation Under “Show All Devices”

11.4.16.1.1 Network Provisioning Screen Compliance Status Indicators

The **Network Provisioning** screen (topology view) utilizes color coding to indicate the presence of compliance alerts on devices. A compliance alert on a device indicates that the running configuration or image is different from the designed configuration or image for the device. This feature enables you to easily see if a device has a compliance alert.

In addition to using color codes for device icons, CVP also uses color codes for container icons to indicate that a device within the container has a compliance alert. If a device within a container has an active alert, the container inherits the alert color of the device. For example, if a device within a container has a configuration mismatch, the container inherits the alert color used to indicate a configuration mismatch.

This feature enables you to easily see if a device within a container has an alert, even if the device is not visible. It also prevents you from having to open a container to see if a device within it has an alert.

Note: Containers only inherit the alert color of a device if the device is directly underneath the container in the hierarchy. If the device is not directly underneath the container in the hierarchy, the container does not show the alert notification color of the device.

For descriptions of the color codes used to indicate compliance status, see:

- Device Icon Compliance Status Color Codes
- Container Icon Compliance Status Color Codes

Device Icon Compliance Status Color Codes

The color of the device icon indicates the compliance status of the device. This table lists and describes the device icon color codes:

Icon	Description
	**Gray** The compliance status is normal (no compliance alert).
	**Orange (no task)** The device has a configuration mismatch (the running configuration or image are different from the designed configuration or image for the device). No task to resolve the mismatch is associated with the device.
	**Orange (with task)** The device has a configuration mismatch (the running configuration or image are different from the designed configuration or image for the device). A task to resolve the mismatch is associated with the device.

See Representation Under “Show All Devices” for how this status is shown when using the **Show All Devices** option.

Container Icon Compliance Status Color Codes

The figure below shows a container that has a device within it that has an alert. In this example, the alert color is yellow, which indicates one of the following:

- A device within the container has a configuration mismatch.

- A device within the container has a configuration mismatch, and there is a task associated with the device to resolve the mismatch.

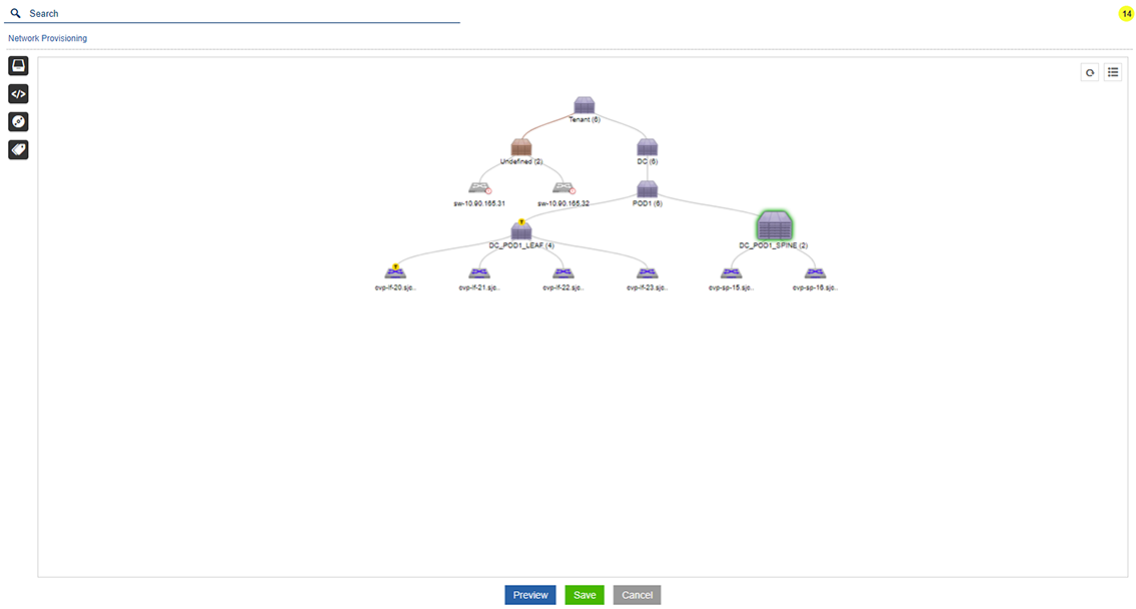

Figure 203: Container showing alert color

11.4.16.12 Representation Under “Show All Devices”

The image below shows the representation of device compliance status information for devices that are only visible by accessing “**Show all devices**. The statuses shown are the same as those shown using device icons in the topology view (see Device Compliance Status Indicators).

Name	IP Address	Mac Address	Serial No.	Container	Status
cvp-lf-20.sjc.aristan...	10.90.165.20	00:1c:73:2b:1d:1c	JPE13300030	DC_POD1_LEAF	T
cvp-lf-21.sjc.aristan...	10.90.165.21	00:1c:73:1e:7b:04	JPE12233288	DC_POD1_LEAF	
cvp-lf-22.sjc.aristan...	10.90.165.22	44:4c:a8:24:88:2f	JPE16012645	DC_POD1_LEAF	
cvp-lf-23.sjc.aristan...	10.90.165.23	44:4c:a8:24:97:81	JPE16012748	DC_POD1_LEAF	
cvp-sp-15.sjc.arista...	10.90.165.15	00:1c:73:9c:c8:47	JPE15065944	DC_POD1_SPINE	
cvp-sp-16.sjc.arista...	10.90.165.16	00:1c:73:9d:52:17	JPE15200275	DC_POD1_SPINE	

1 - 6 of 6 1 of 1

Figure 204: Show All Devices display of device compliance status

11.4.16.13 Representation in List View

The image below shows the representation of device compliance status information when using the **List View**. The statuses shown are the same as those shown using device icons in the **Topology** view.

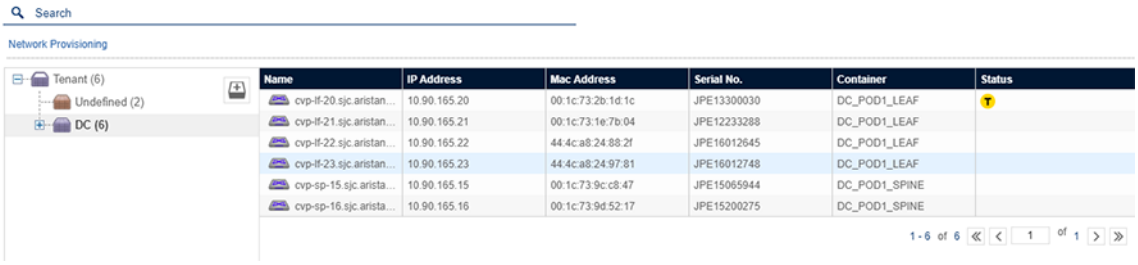

Figure 205: List View display of device compliance status

11.4.16.1.4 Removing Compliance Indicators

The **Network Provisioning** screen shows non-compliance whenever these is a mismatch between the running configuration or image and designed configuration or image of devices in the topology. Compliance indicators are removed from the display only when there is no configuration mismatch.

To remove compliance indicators, you may need to do one or more of the following:

- Run a compliance check on any devices on which there is no configuration mismatch.
- Run a task to bring any non-compliant devices into compliance.
- Reconcile the configuration of any devices that have a configuration mismatch.

11.4.16.2 Device Compliance Checks

CloudVision Portal (CVP) enables you to see if devices are non-compliant by performing compliance checks at the device level and at the container level.

11.4.16.2.1 Running container-level compliance checks

When you run a compliance check at the level of the container, CloudVision Portal (CVP) compares the running configuration and image on each device in the container against the designed configuration and image for each device in the container. You run the check from the **Network Provisioning** screen.

When you start the compliance check, a message at the top of the screen indicates that the check has started. When the check is done, non-compliant devices are indicated on the screen using colors (yellow or red). In addition, CVP automatically generates a notification that a compliance check has been completed. You can view the notification for the compliance check by clicking the bell icon on the screen.

Complete these steps to run a container-level compliance check:

1. Make sure the **Network Provisioning** tab is selected.
2. On the **Network Provisioning** screen, locate the container to be checked for compliance.
3. Right-click on the container and choose **Check Compliance**.

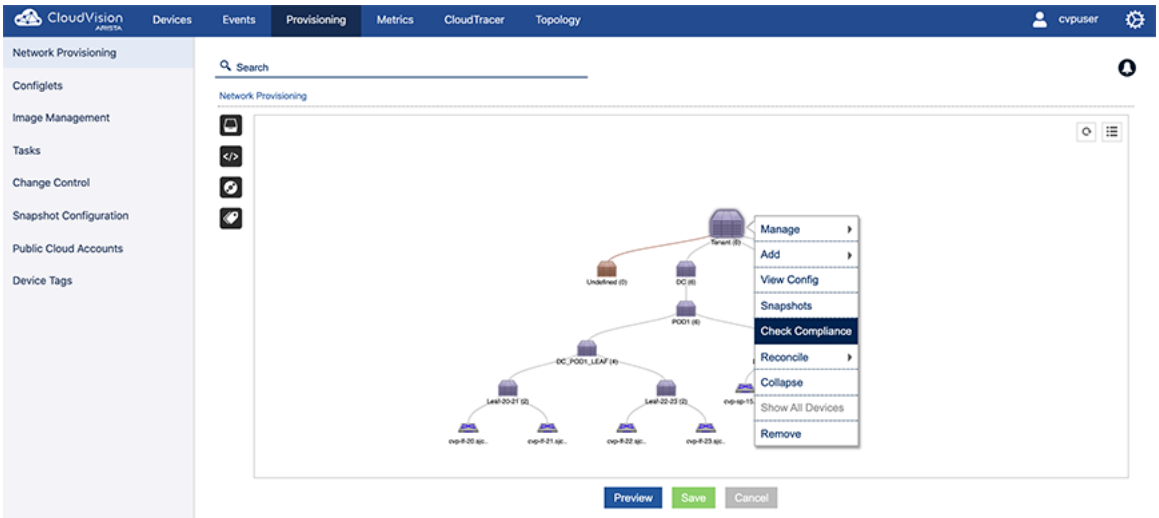

Figure 206: Container-level compliance check

11.4.16.2.2 Running device-level compliance checks

When you run a compliance check on a single device, CloudVision Portal (CVP) compares the running configuration and image on the device against the designed configuration and image for the device. You run the check from the Network Provisioning screen.

Complete these steps to run a device-level compliance check:

1. Make sure the **Network Provisioning** tab is selected.

2. On the **Network Provisioning** screen, locate the device on which you want to run the compliance check.
3. Right-click on the device and choose **Check Compliance**.

   CVP initiates the compliance check. Non-compliant devices are indicated on the screen using device compliance status icons (see Device Compliance Status Indicators).

11.4.16.3 Device Access Alerts

The **Network Provisioning** screen shows device access alerts whenever a device is no longer reachable by CVP. This enables you to easily identify unreachable devices in the screen. Any device that is no longer reachable is represented on the screen using a color coded device icon.

This table lists and describes the color codes used for unreachable devices:

Icon	Description
	**Red** The device is unreachable (CVP cannot connect to the device).

Like device compliance status alerts, CVP also uses color codes for container icons to indicate that a device within the container is unreachable. If a device within a container has an access alert, the container inherits the alert color of the device (red).

This feature enables you to easily see if a device within a container has an alert, even if the device is not visible. It also prevents you from having to open a container to see if a device within it has an alert.

Note: Containers only inherit the alert color of a device if the device is directly underneath the container in the hierarchy. If the device is not directly underneath the container in the hierarchy, the container does not show the alert notification color of the device.

11.4.17 Notifications for Container-level Compliance Checks and Reconciles

CloudVision Portal (CVP) provides notifications for container-level compliance checks and reconciles. When a container-level compliance check or reconcile is completed, CVP automatically generates a notification message, indicating that the action has occurred.

Because container-level compliance check or reconciles are not tracked by tasks, you track them using automated notifications. The notifications can be accessed directly from the **Network Provisioning** screen by clicking the **Notifications** icon. The presentation of the icon indicates whether there are unread notifications.

Figure 207: Read and Unread Notification Icons

The notification list provides the following information:

- Current actions in progress, with a progress bar.
- Unread notifications (shaded in blue).
- Previously viewed notifications (no shading). These are shown at the bottom of the list.

The type of action (Check **Compliance** or **Reconcile**) is indicated for each notification.

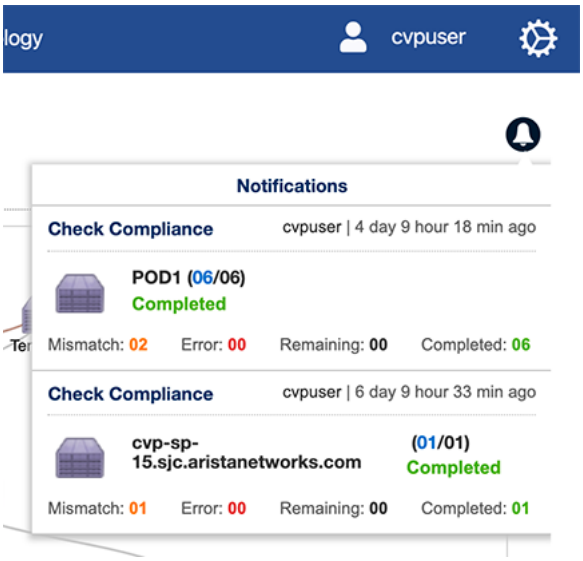

Figure 208: List of notifications

Note: To view notifications for the previous CVP session, click the bell icon and choose **View History**.

For information on container-level compliance checks and reconciles, see:

- Device Compliance Checks
- Reconciling Device Configurations Differences at the Container Level

11.4.18 Global Search

In the **Network Provisioning** module, the user can use the search bar at the top of the module to find a given device or container.

11.4.18.1 Search Behavior in Topology and List View

This search is very different from rest of other search options available in topology. On user starts to type, the list of possible matches will be displayed below as an auto suggestion.

11.4.18.2 Topology Search

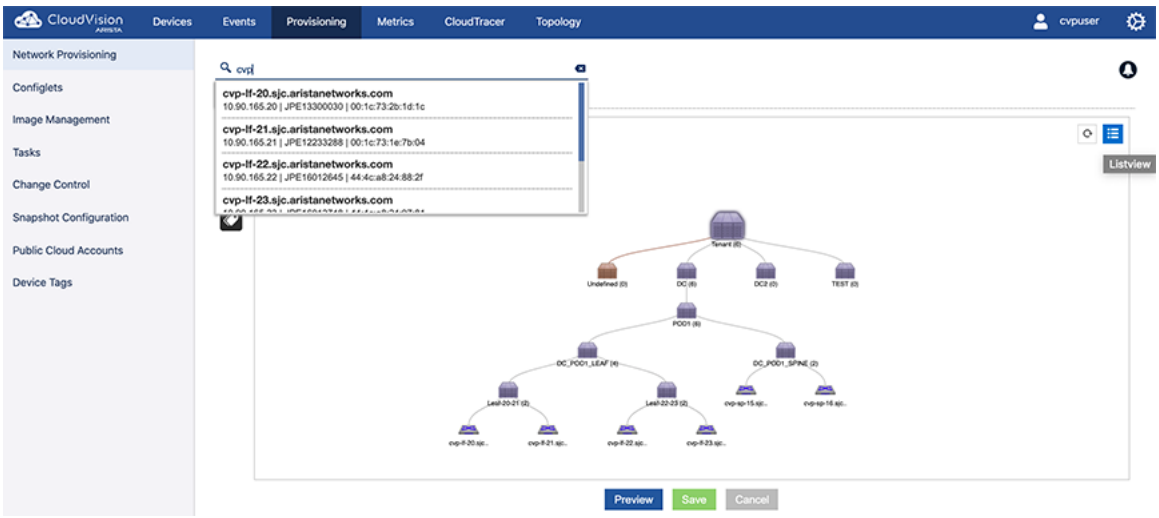

Figure 209: Using search

11.4.18.3 List View Search

The search behaves similar to the topology search.

For a single device search, the selected device will be listed in the grid.

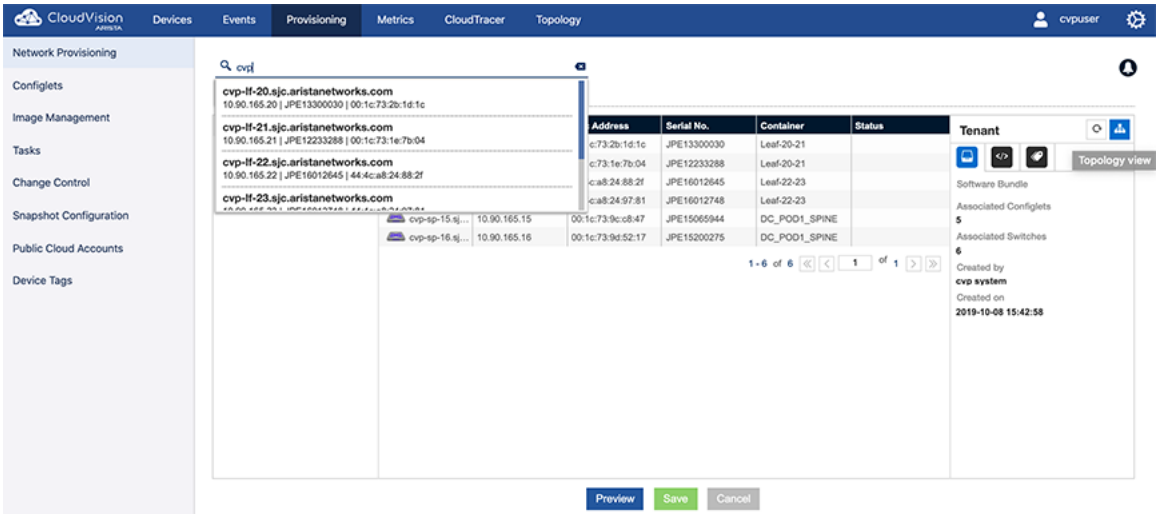

Figure 210: List view search

11.4.18.4 "Search" in Other Grids

During a grid search, the user will not be provided with an auto suggest option. Only the records matching the specified data entered will be filtered and displayed in the grid.

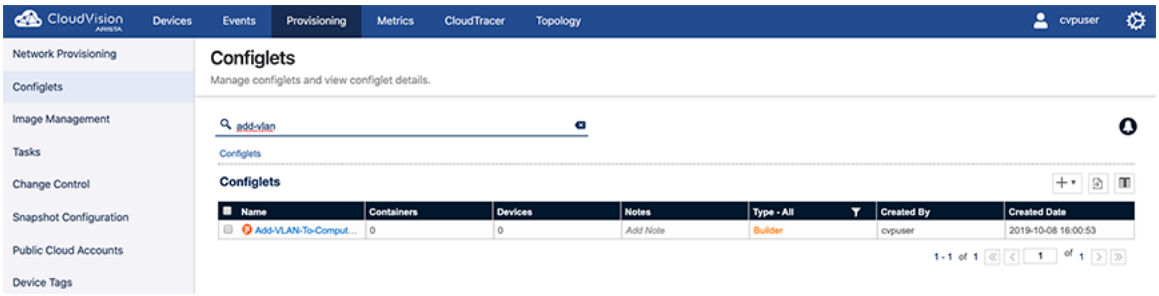

Figure 211: Grid searches

11.4.18.5 Label Search

Use the search bar from the Network Provisioning screen to filter the devices based on labels.

This is a contextual search.

To search a label:

1. Use the keyword Label: followed by the label name.

11.4.18.5.1 AND Operation

Lists all the devices which has both the labels present on it in the hierarchy.

Label: <*Label Name*> AND Label: <*Label Name*>

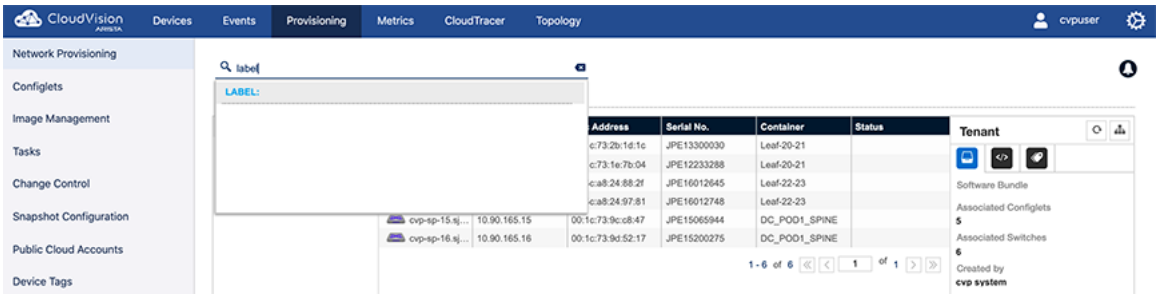

Figure 212: Search AND operation

11.4.18.5.2 OR Operation

Lists all the devices which has either one of the labels present on it in the hierarchy.

Label: <Label Name> OR Label: <Label Name>

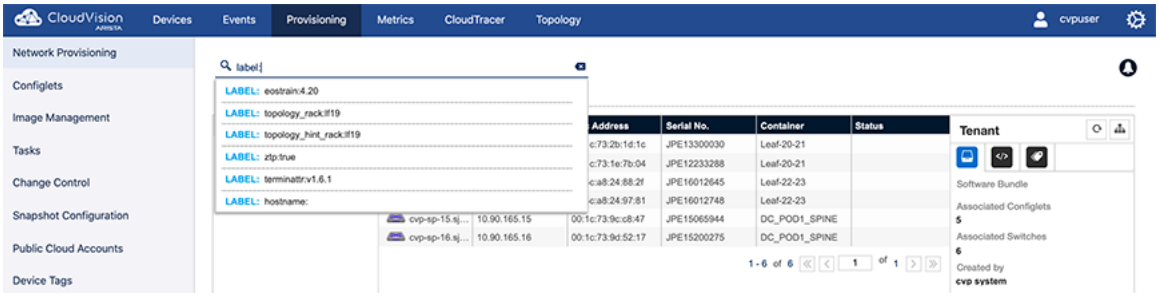

Figure 213: Search OR operation

11.4.18.5.3 NOT Operation

Lists all the devices which has first label one the labels present on it in the hierarchy.

Label: <*Label Name*> AND NOT Label: <*Label Name*>

Figure 214: Search AND NOT operation

11.4.18.6 Preview Option

All the actions performed in **Network Provisioning** module can be previewed before saving the changes.

To access the preview screen:

1. Select the “Preview” button.

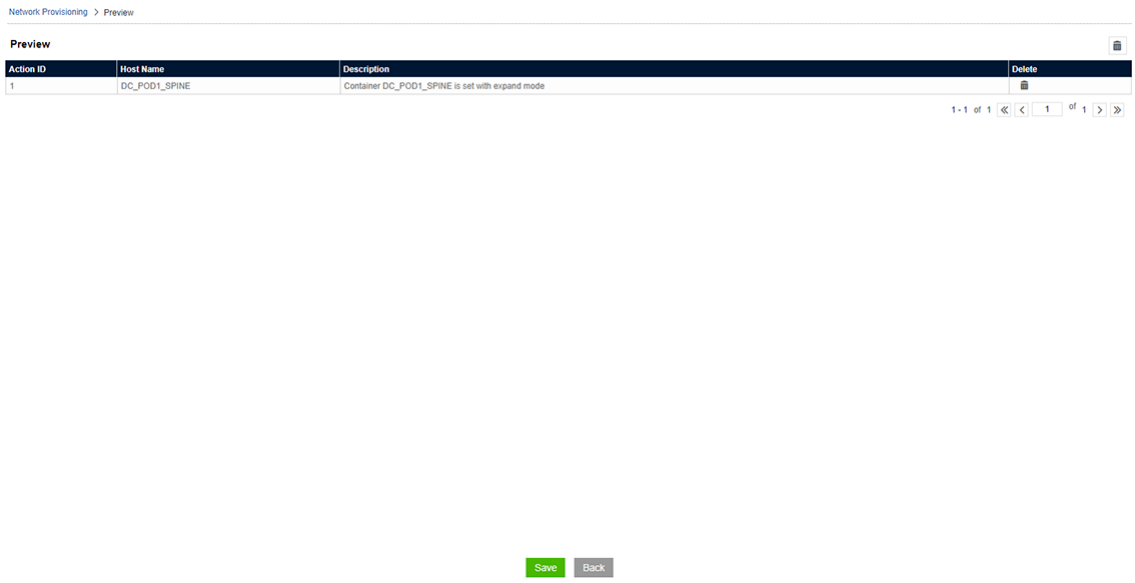

Figure 215: Preview option display

11.4.19 Management IP

The CloudVision Portal tracks the Management IP of each device to use in connecting to it. When this IP address changes, the device becomes unreachable by the portal. You can manually change the IP address used by the portal to communicate with a given device.

11.4.19.1 Changing A Device's Management IP

The management IP address of a device may change for one of the following reasons:

Reason 1:

When a device is provisioned using Zero Touch Provisioning, it may have been assigned a temporary IP address via DHCP. The CloudVision Portal will use this IP address to provision the device. Once the configuration is pushed and the device reboots, this IP address may change.

Reason 2:

1If you change the device IP address directly via the switch console, CloudVision cannot record the change, and the device will become unreachable. **Current management IP** and **proposed management IP** can be used to mitigate this potential issue.

Option 1:

Current Management IP: The IP address used by CloudVision to communicate with a device.

1. Set the proposed IP address before pushing the configlet. This way CloudVision will try to reach the device with this IP address once configuration is pushed.

Option 2:

Proposed Management IP: The IP address that CloudVision uses after pushing the configlet.

1. In the Inventory Management screen and the topology, update the Management IP address. For any unreachable device, set the IP address to bring it back to the network.

11.4.19.2 Setting Proposed Management IP

You can set the Proposed Management IP while adding configlets to the device using the Proposed Management IP menu.

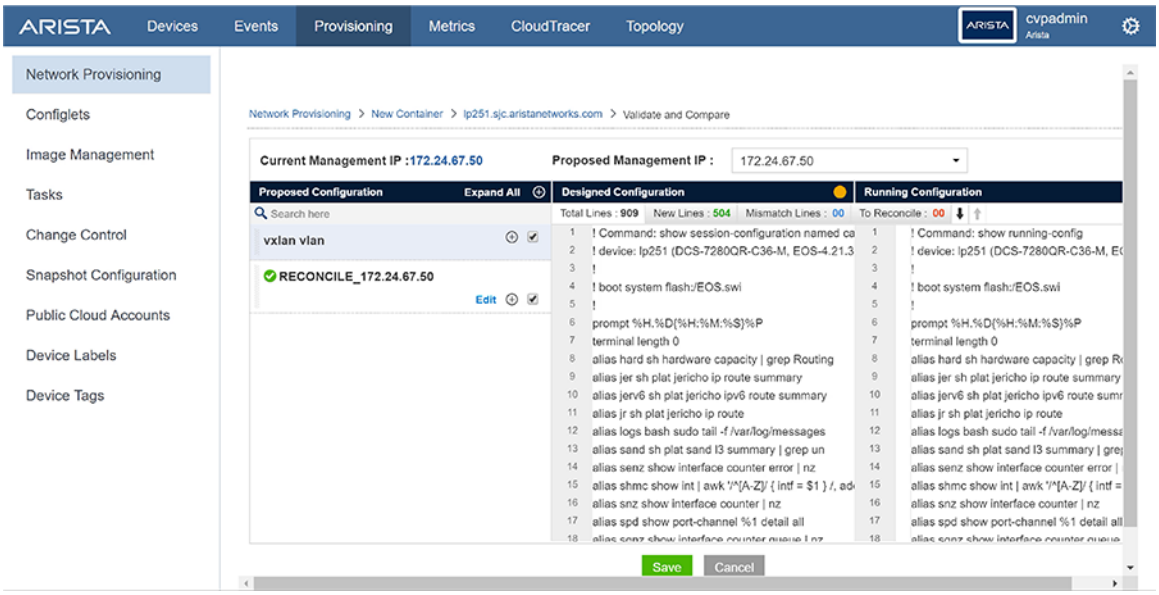

Figure 216: Location of menu for setting Proposed Management IP

If you do not set the Proposed Management IP, you cannot save the configuration as not setting Proposed Management IP.

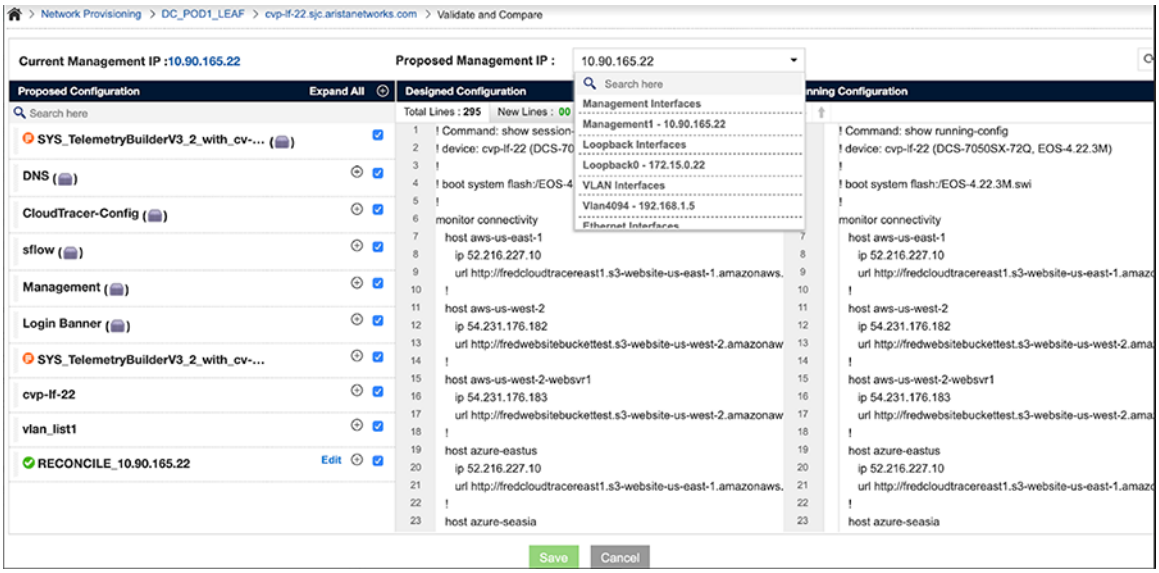

Figure 217: Setting the Proposed Management IP

1. Select the Proposed Management IP using the drop-down menu.

   CloudVision lists the available Management IP, Loop back IP, VLAN IP, and Routed Ethernet IP.
2. Select the desired IP address.
3. Click **Save**.

   A task is spawned to assign the new Proposed Management IP.

11.4.19.3 Changing Current Management IP

1. Go to the **Network Provisioning** screen.
2. Select a device from topology/list view.
3. Right-click the device and choose **Manage > IP Address**

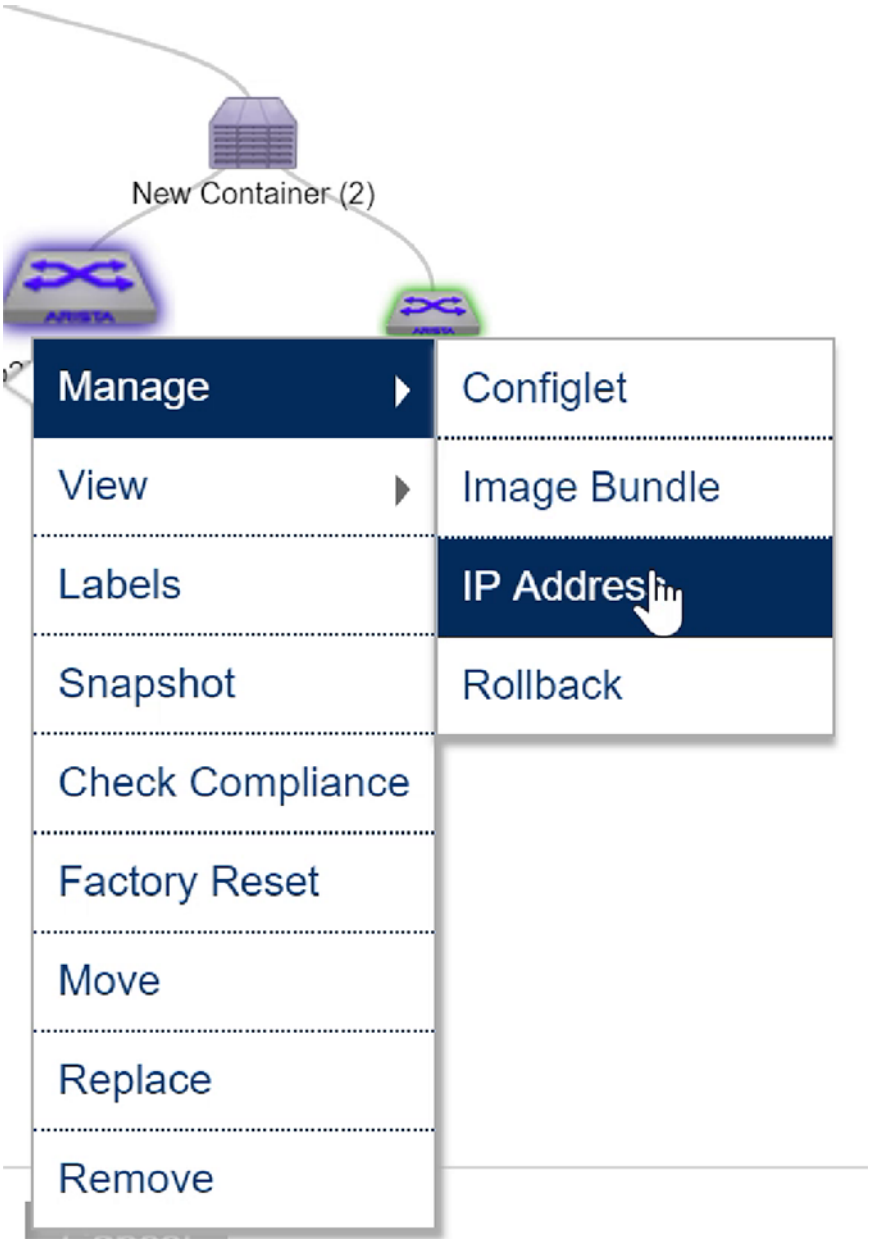

Figure 218: Change Management IP

4. A pop up will appear allowing you to manually add a new IP address.

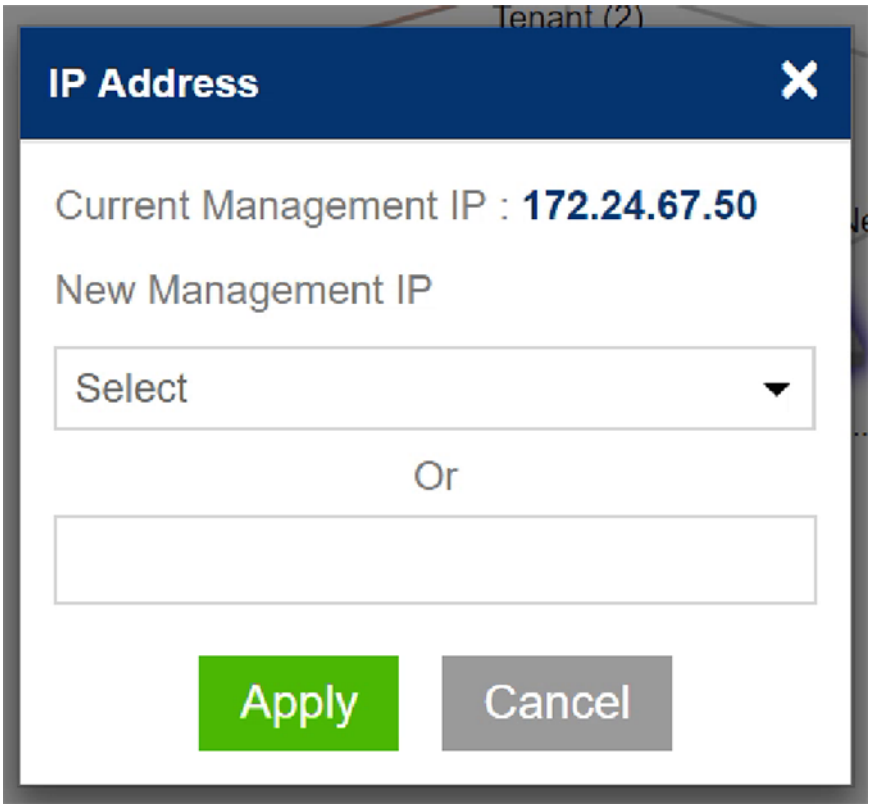

Figure 219: Change IP Address

5. Verify the reachability of new IP address.

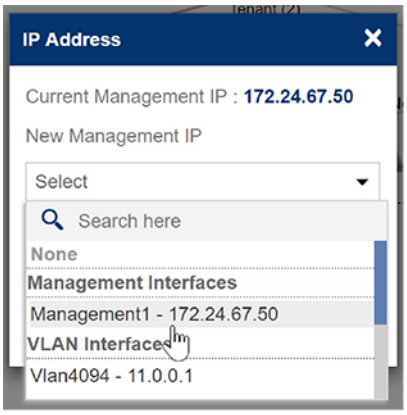

Figure 220: Verify IP Address

Chapter 12

Configlet Management (CVP)

Configlets are portion of configuration that CLOUDVISION user codes and maintains independently under Configlet Management inventory. These Configlets can be later applied to devices or containers in the topology.

Sections in this chapter include:

12.1 Creating Configlets

CloudVision Portal (CVP) enables you to create Configlets using two different methods. You can create Configlets using the CVP Configlet Builder feature, or you can create them manually. You should use the method that is best suited to your intended use of the Configlet.

Note: The Configlet Builder feature is designed to help you create Configlets dynamically based on variables.

For more information, see:

12.1.1 About the Configlet Builder Feature

The Configlet Builder feature enables you to programatically create device configurations (Configlets) for devices that have relatively dynamic configuration requirements. This helps to prevent you from having to manually code Configlets.

The Configlet Builder feature is essentially a set of user interface (UI) widgets and a python script, that when used together, programatically generate Configlets for a device. The python script is embedded into a python interpreter, which is the component that generates Configlets. The UI widgets are essential if you want to use the feature to generate Configlets with user input.

Note: Using UI widgets associated with a Configlet Builder are optional. If the UI widgets are used, the generated Configlets require user input to be created.

The Configlet Builder can be used to create Configlets for both devices or containers, in the same way that static Configlets can be used with devices or containers. Configlets that are created using the Configlet Builder are executed (including the generation of Configlets) at the point when the Configlet Builder is applied to a device or container, or when a device is added to a container that contains a Configlet Builder.

12.1.2 Creating Configlets Using the Configlet Builder

The Configlet Builder enables you to create Configlets (device configurations). The example Configlet Builder shown being created configures the device's management interface based on input you enter through the use of UI widgets.

Complete the following steps to create Configlets using the Configlet Builder:

1. Create a Configlet Builder from the Configlet page.

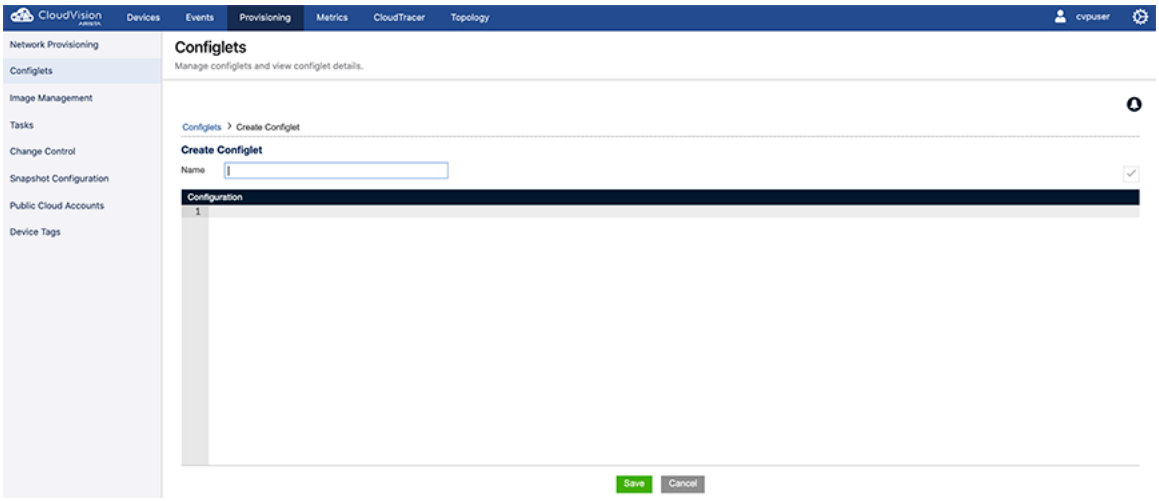

Figure 221: Creating a Configlet Builder

2. (Optional) Define the UI widgets to be associated with the Configlet Builder.

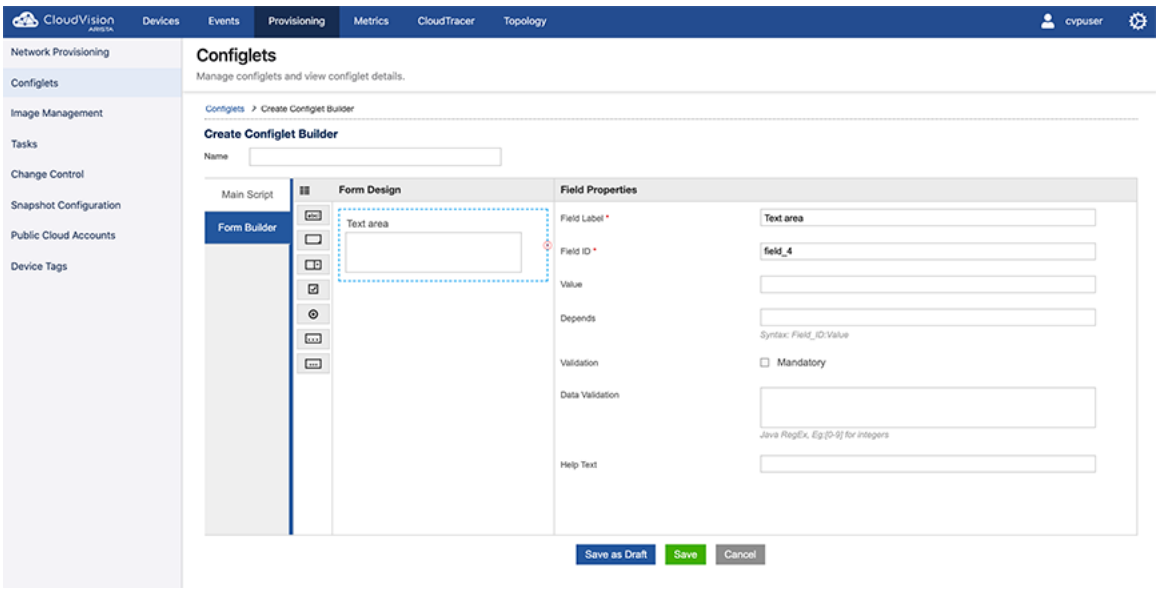

Figure 222: Configlet UI Widgets

The widget types are:

- **Text Box** – Use for single line text entries (for example, descriptions, host name).
- **Text Area** – Use for multiple lines of text (for example, MOTD, or login banner).
- **Drop Down** – Use to select a value from a menu as defined in the Value Field.
- **Tick Box** – Use to select a value from a tick list as defined in the Value Field.
- **Radio Button** – Use to select one option from a set of options as defined in Value Field.
- **IP Address** – Use to specify an IP address (this is a Dotted Decimal Address field).

- **Password** – Use to specify a single line of text (characters are hidden as they are entered).

3. Write a Python script that reads the inputs you entered in the previous step and then generates the Configlet.

Note: The figures listed in this table show examples of the steps involved in writing a script, including an example of use of standard Python syntax to build components of the Configlet.

Figure	Example of	Description
Figure 223: Example (Showing Import of CVP-Specific Internal Libraries)	Importing CVP-specific internal libraries into the script	The CVP-specific internal libraries are used by the script to access form fields and CVP variables.
Figure 224: Example (Showing Specification of Field IDs Defined in the Form Builder)	Specification of field IDs defined in the Form Builder	You must specify the IDs of fields you defined in the Form Builder in **Step 2**. The fields you specify are included in the Configlet content generated by the script.
Figure 225: Example (Showing Use Of Standard Python Script Syntax)	Use of standard Python syntax	The Configlet Builder supports the use of standard Python syntax to build parts of the Configlet. You can also make calls to external files and database.
Figure 226: Example (Showing Print Output)	Print output (Configlet content)	The script automatically produces print output from the CVP internal libraries you imported and the fields you have defined in the script. The print output is the content of the Configlet.

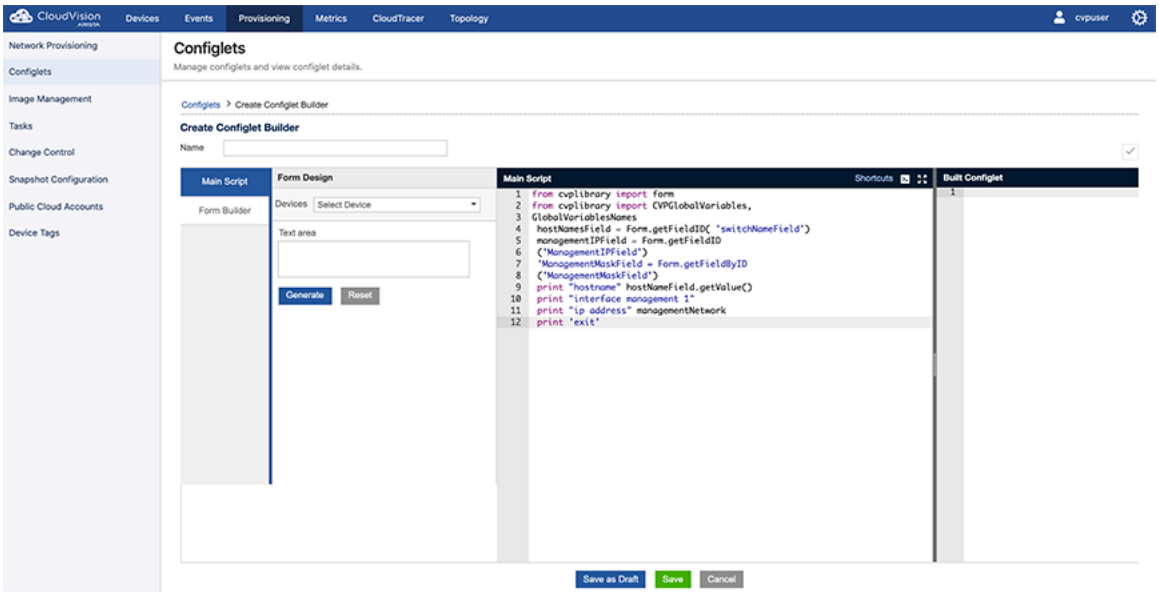

Figure 223: Example (Showing Import of CVP-Specific Internal Libraries)

Main Script Shortcuts

```
from cvplibrary import form
from cvplibrary import CVPGlobalVariables,
GlobalVariablesNames
 hostNamesField = Form.getFieldID( 'switchNameField')
 managementIPField = Form.getFieldID
 ('ManagementIPField')
 'ManagementMaskField = Form.getFieldByID
 ('ManagementMaskField')
 print "hostname" hostNameField.getValue()
 print "interface management 1"
 print "ip address" managementNetwork
 print 'exit'
```

Figure 224: Example (Showing Specification of Field IDs Defined in the Form Builder)

Main Script Shortcuts

```
from cvplibrary import form
from cvplibrary import CVPGlobalVariables,
GlobalVariablesNames
 hostNamesField = Form.getFieldID( 'switchNameField')
 managementIPField = Form.getFieldID
 ('ManagementIPField')
 'ManagementMaskField = Form.getFieldByID
 ('ManagementMaskField')
 print "hostname" hostNameField.getValue()
 print "interface management 1"
 print "ip address" managementNetwork
 print 'exit'
```

Figure 225: Example (Showing Use Of Standard Python Script Syntax)

Main Script Shortcuts

```
from cvplibrary import form
from cvplibrary import CVPGlobalVariables,
GlobalVariablesNames
 hostNamesField = Form.getFieldID( 'switchNameField')
 managementIPField = Form.getFieldID
 ('ManagementIPField')
 'ManagementMaskField = Form.getFieldByID
 ('ManagementMaskField')
 print "hostname" hostNameField.getValue()
 print "interface management 1"
 print "ip address" managementNetwork
 print 'exit'
```

Figure 226: Example (Showing Print Output)

Note: Complete steps 4 and 5 to test the script to make sure it can generate Configlet content.

4. Fill in the Form Design fields.

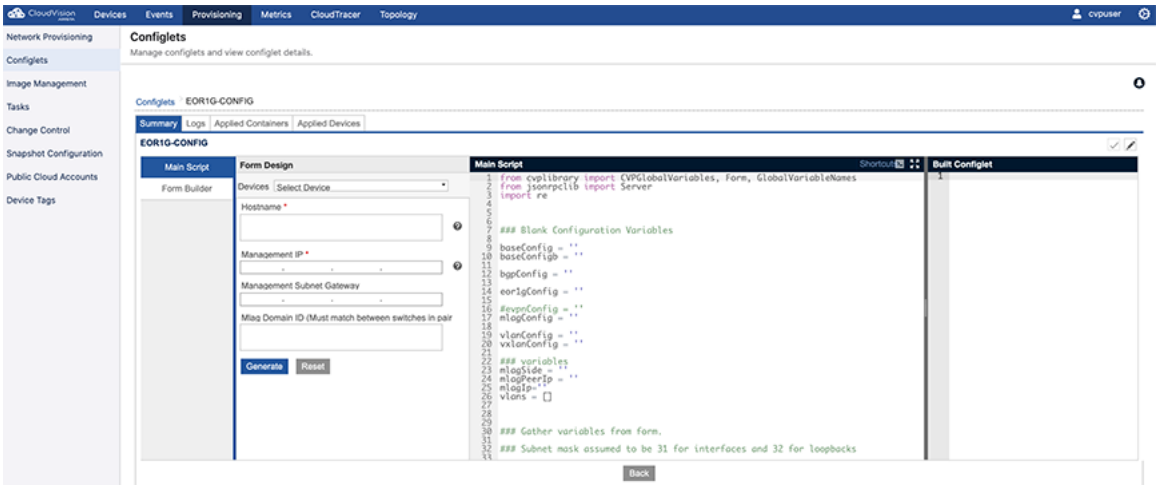

Figure 227: Filling in the Design Fields

5. Click **Generate.**

The Configlet content is generated and shows in the **Built Configlet** pane (see Figure 229: Example (Generating Configlet Content)).

Note: If it is necessary to select a device to generate the Configlet, then select a device from the list of devices under Form Design (see Figure 228: Selecting a Device from the List of Devices Under Form Design).

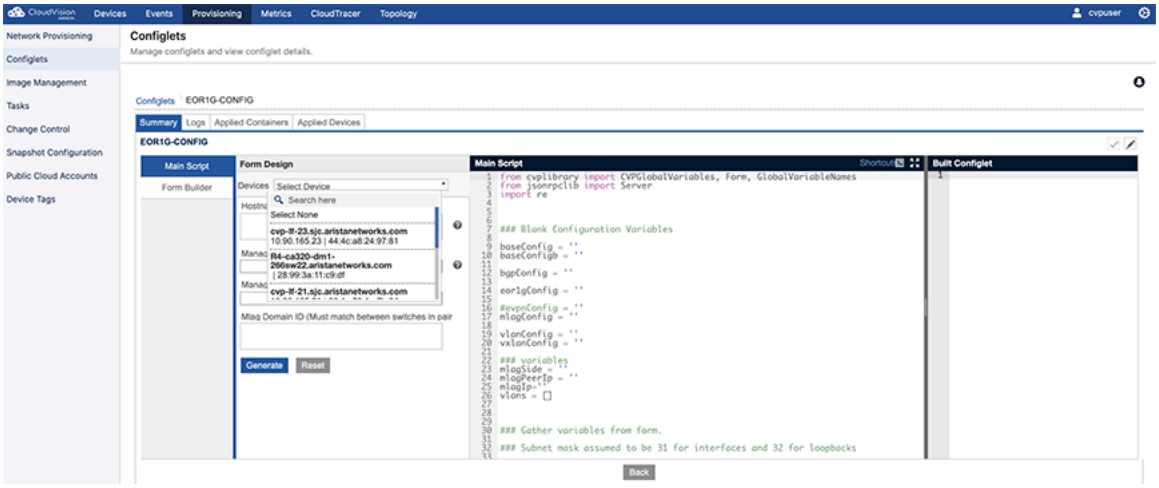

Figure 228: Selecting a Device from the List of Devices Under Form Design

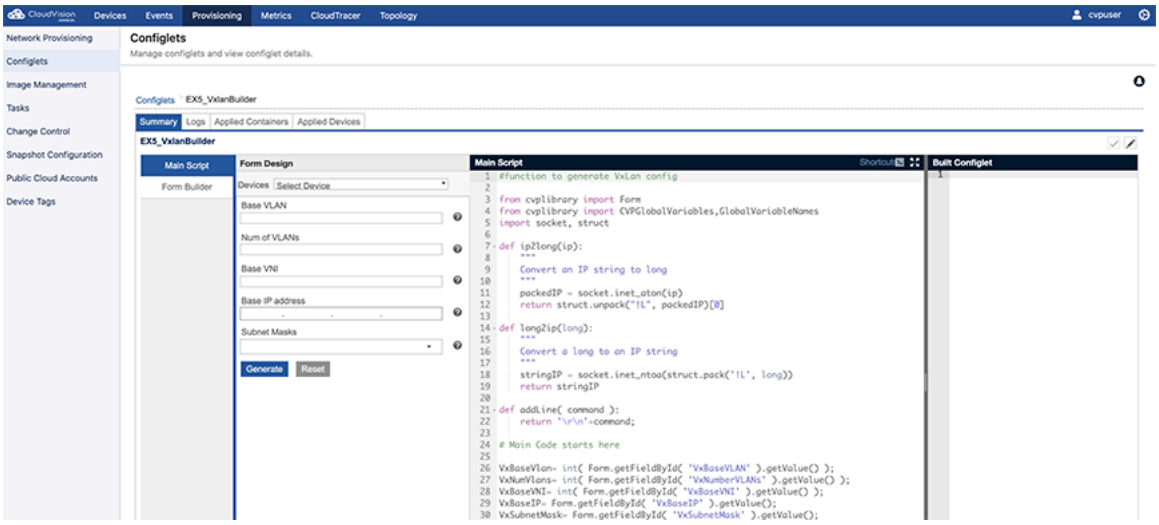

Figure 229: Example (Generating Configlet Content)

6. Validate the generated Configlet on the device by clicking the **Tick** icon at the upper-right of the page.

   The Validate Device dialog appears.

7. In the Validate Device pop-up dialog, click Validate Example Script (Validating Device).

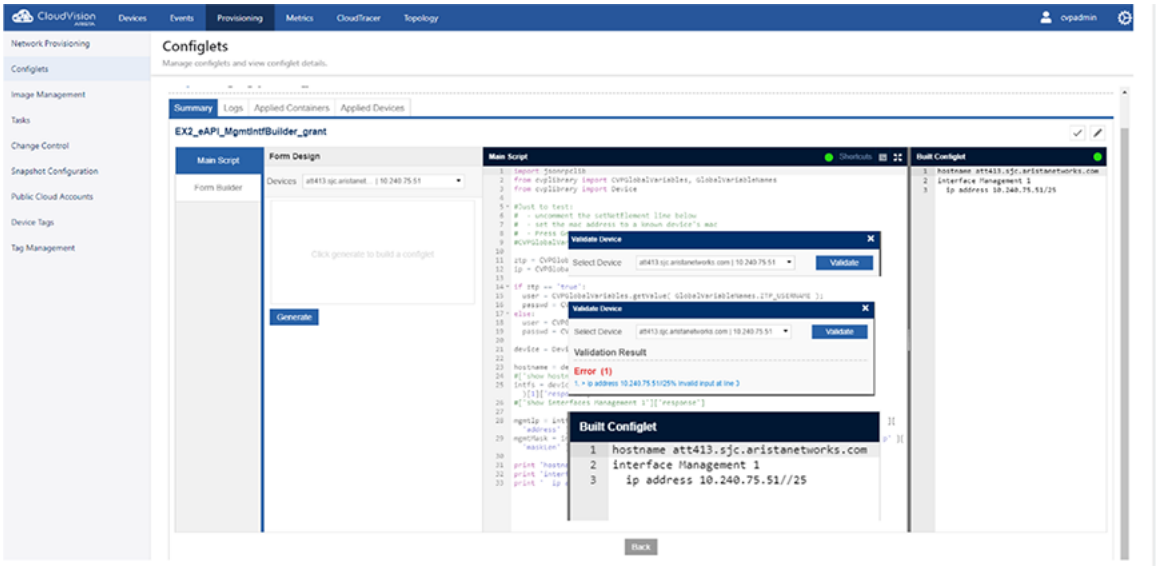

Figure 230: Example Script (Validating Device)

If the device cannot be validated, the error (or errors) are listed in the Validate Device dialog.

8. (If needed) Correct any errors and repeat step 7 to validate the device.

   The Validate Device dialog shows a message to indicate a successful validation.

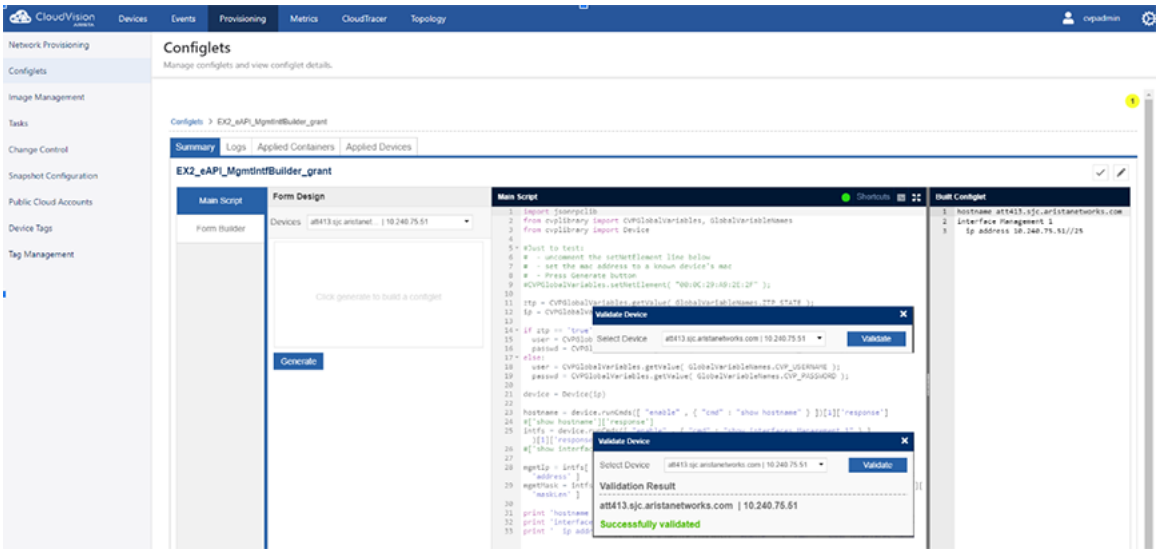

Figure 231: Example Script (Re-Validating Device after Correction)

9. To apply the new Configlet to the container, do the following:
   a. Go the Network Provisioning page.
   b. Right-click the container and choose **Manage > Configlet**.

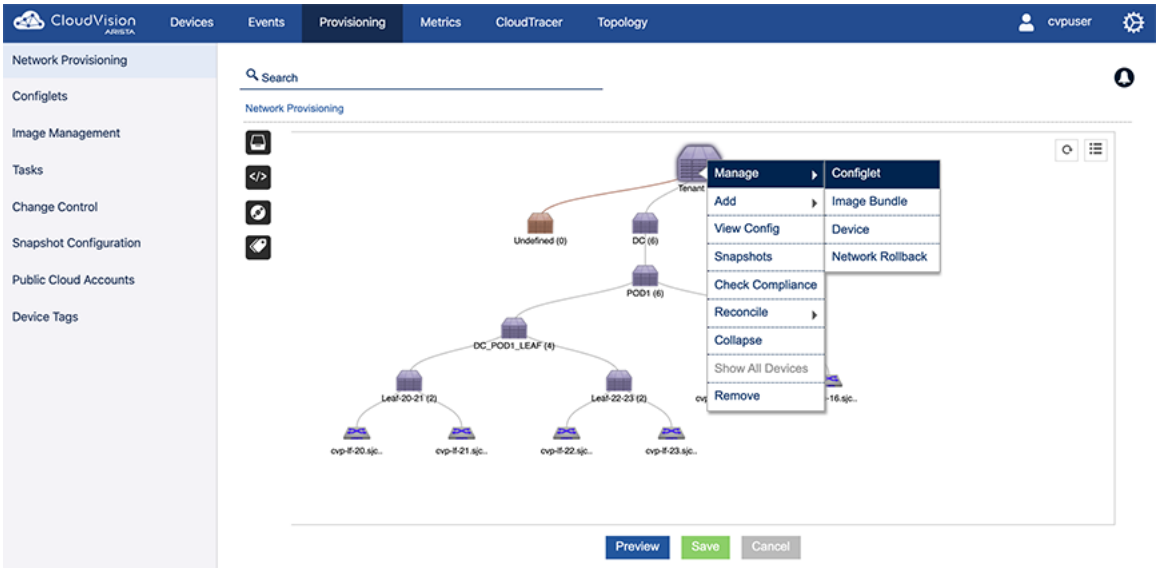

Figure 232: Select the Container to Apply the New Configlet

The list of available Configlets appears on the Configlet page.

10. Select the Configlet to apply to the device by clicking the checkbox next to the name of the Configlet (see .

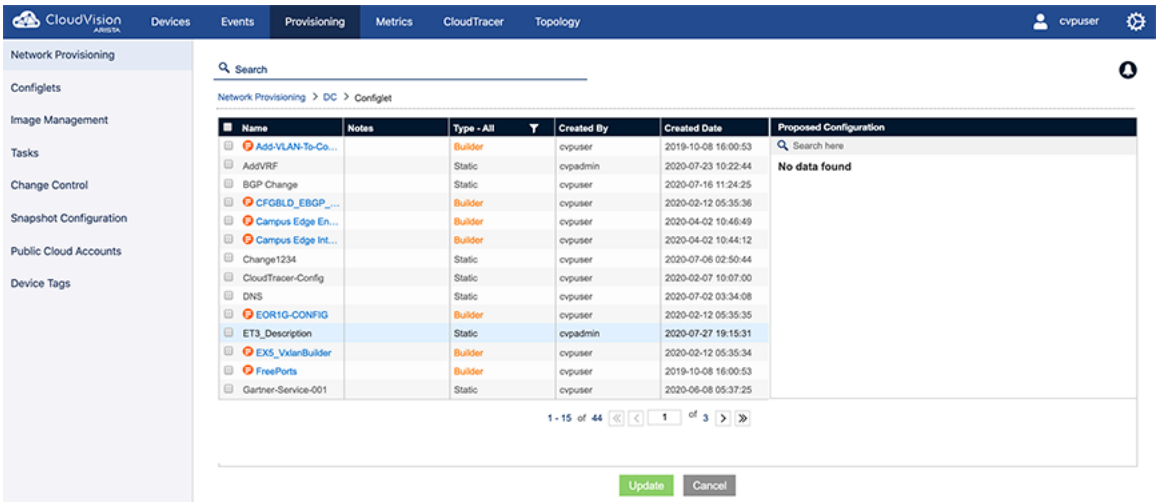

Figure 233: Select Configlet on Configlet Page

11. To add devices to the container, do the following:

    a. Go the **Network Provisioning** page.

    b. Right-click the container and choose **Device > Add**.

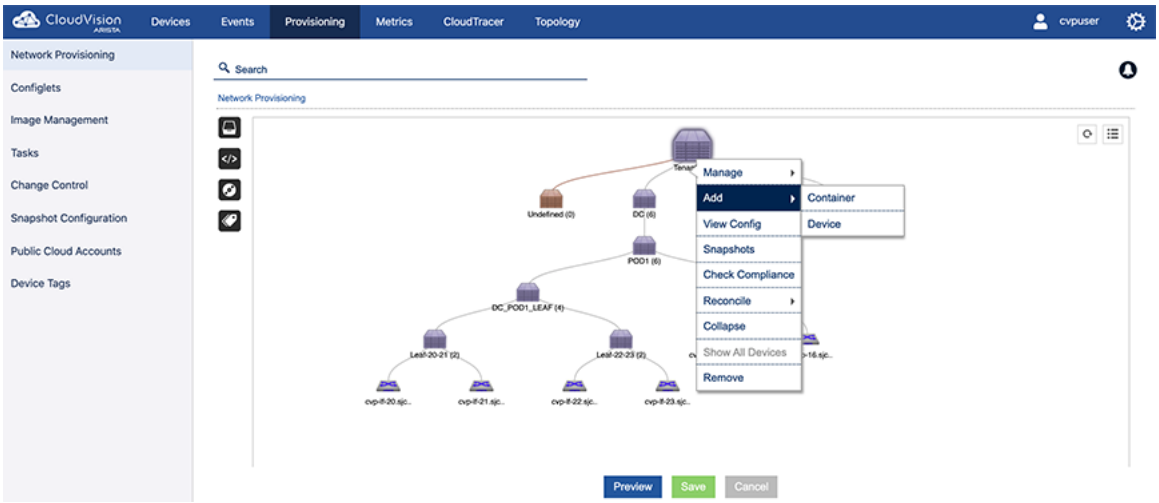

Figure 234: Adding Devices to the Container

12. Do one of the following:

- Click **Yes** to apply the Configlet you selected to all of the devices in the hierarchy.
- Click **No** if you do not want to apply the Configlet you selected to all of the devices in the hierarchy.

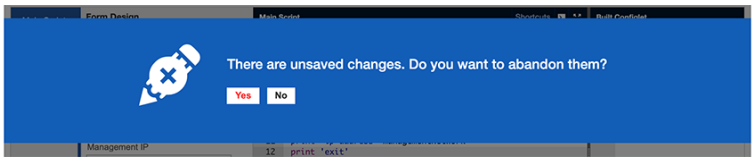

Figure 235: Message Indicating Selection of Hierarchical Container

The Configlet page appears showing the Configlet you selected to apply to the container.

13. To assign the Configlet Builder to the container you selected, select (click) the **Configlet Builder**.

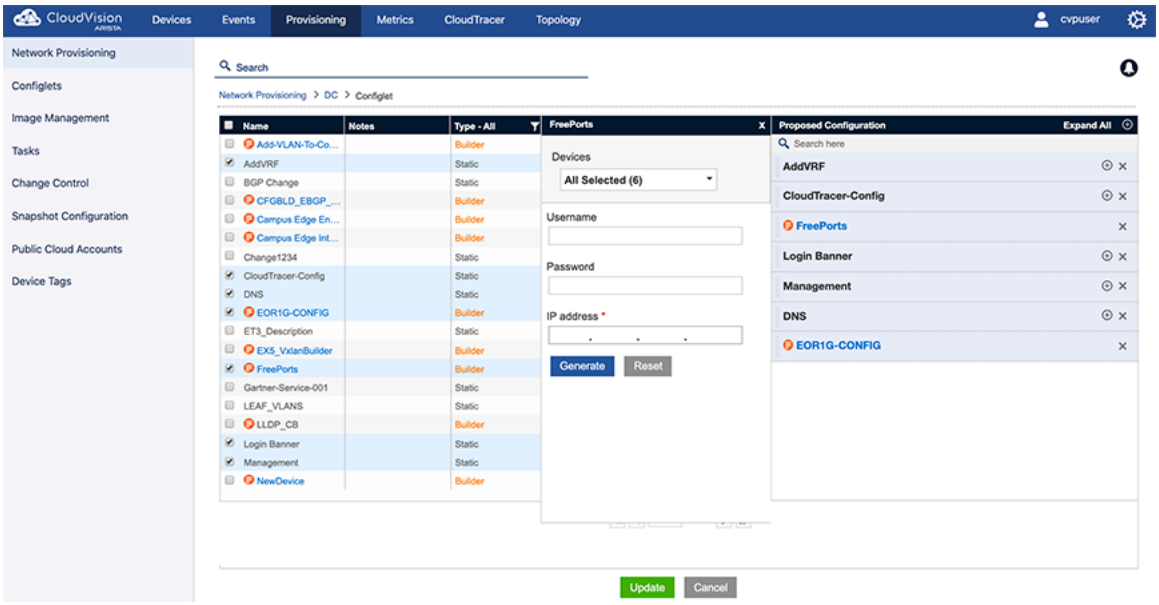

Figure 236: Selecting the Configlet to Assign to the Container

The page loads a form (see Figure 237: Form Loaded on Page after you Select the Configlet Builder).

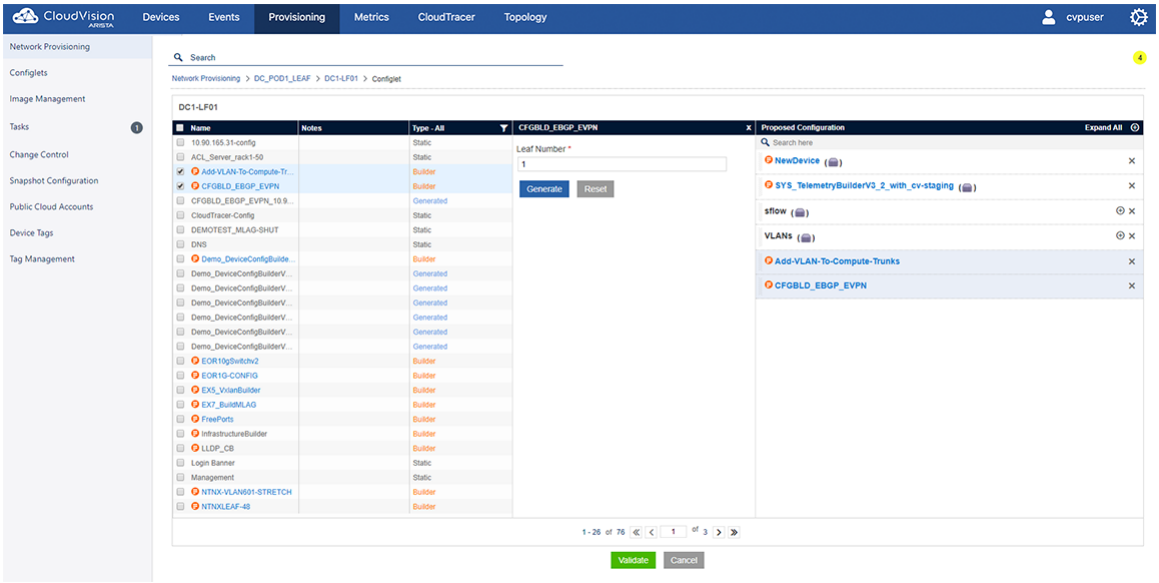

Figure 237: Form Loaded on Page after you Select the Configlet Builder

14. Complete (fill in) the form and then click **Generate**.

    The Configlet Builder creates the new, device-specific Configlet, and the Configlet is shown in the **Built Configlet** pane.

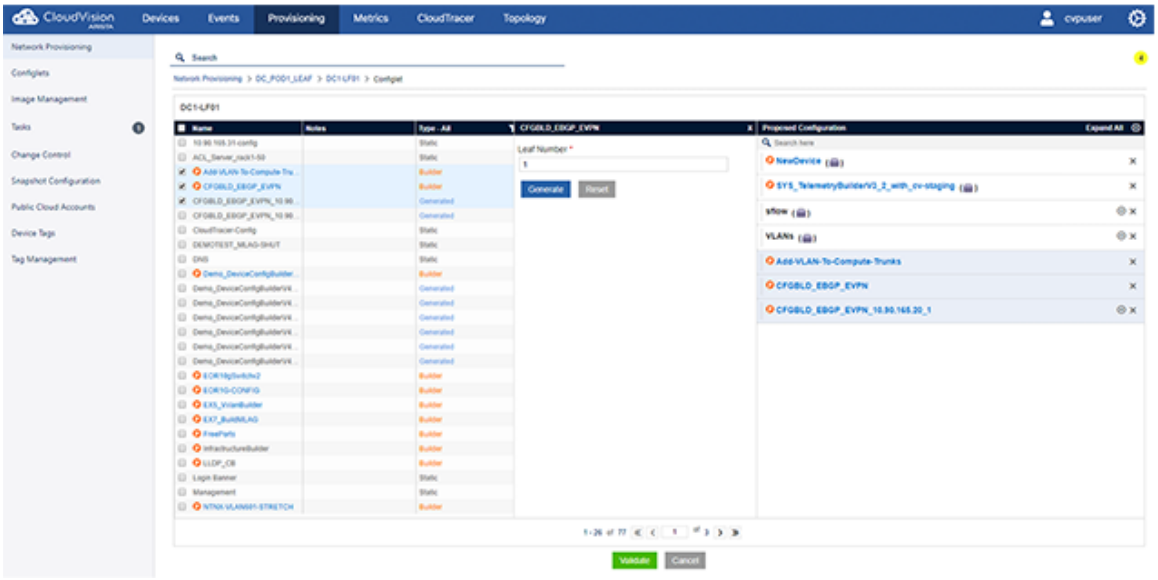

Figure 238: Configlet Page Showing New, Device-Specific Configlet

12.1.3 Using the Provided Configlet Builder Examples

CloudVision Portal (CVP) provides some Configlet Builder examples to help you get started using this feature.

You can load the examples to your CVP instance using the following commands:

- Log into the primary node's Linux shell as root user.
- Change directory to */cvpi/tools* and import the example Configlets using the cvptool.

```
./cvptool.py --host <host> --user <user> --password <pass> --objects
 Configlets --action restore --tarFile examples.tar.
```

The provided examples include:

- Example 1: Form-based management interface Configlet Builder
- Example 2: eAPI-based management interface Configlet Builder
- Example 3: SSH-based management interface Configlet Builder
- #unique_375
- Example 5: Device library based management interface Configlet Builder

12.1.3.1 Example 1: Form-based management interface Configlet Builder

This example uses the form to input the management interface configuration, and generates a new Configlet to preserve the configuration.

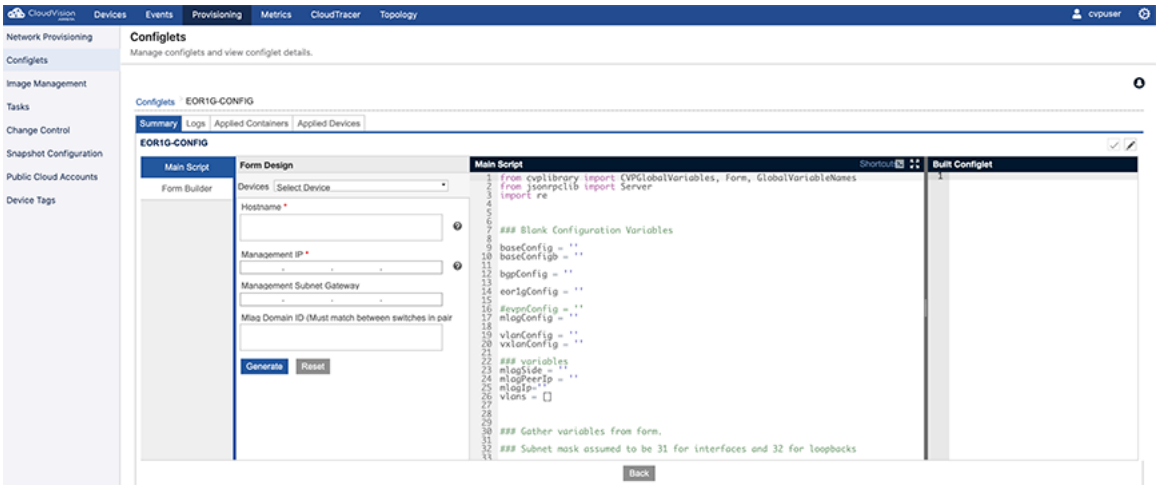

Figure 239: Example 1

12.1.3.2 Example 2: eAPI-based management interface Configlet Builder

This example uses eAPI to read the management interface configuration that the device received from the DHCP server during the ZTP boot, and generates a new Configlet to preserve the configuration.

Note: No UI widgets are associated with the Configlet Builder in this example.

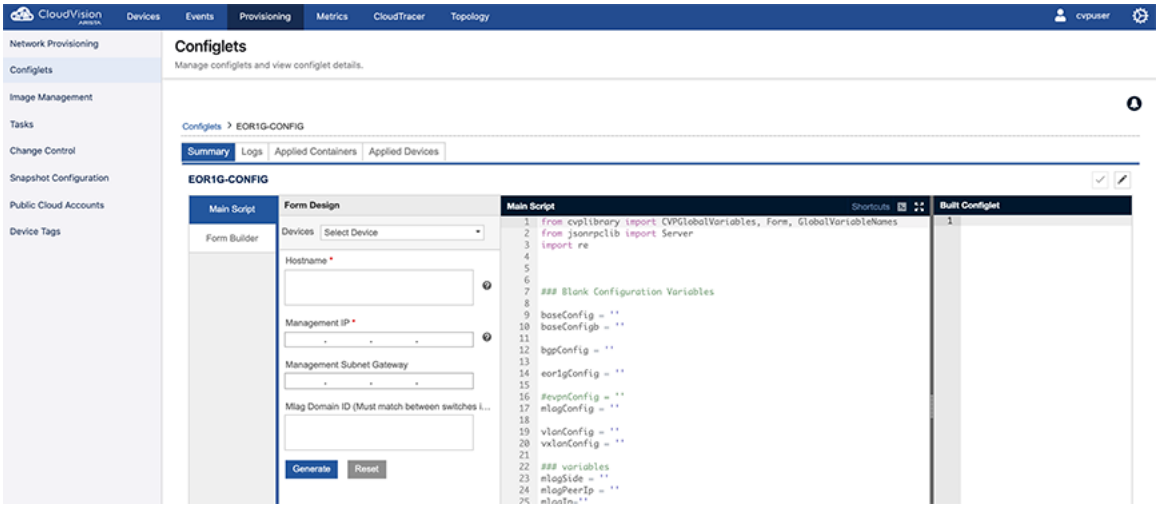

Figure 240: Example 2

12.1.3.3 Example 3: SSH-based management interface Configlet Builder

This example uses SSH to read the management interface configuration that the device received from the DHCP server during the ZTP boot, and generates a new Configlet to preserve the configuration.

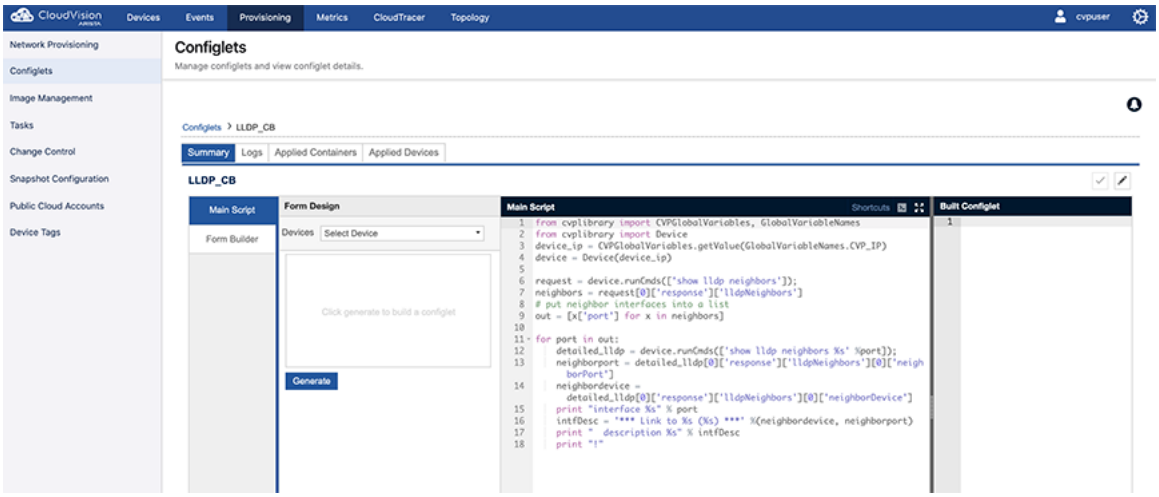

Figure 241: Example 3

12.1.3.5 Example 5: Device library based management interface Configlet Builder

This example uses Device library to read the management interface configuration that the device received from the DHCP server during the ZTP boot, and generates a new Configlet to preserve the configuration.

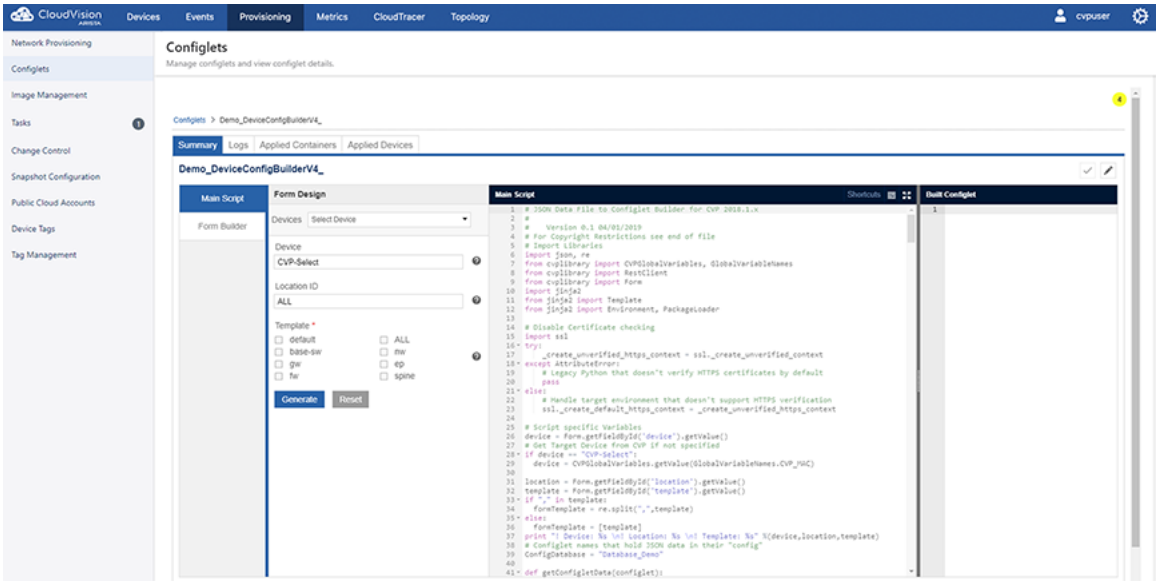

Figure 242: Example 5

12.1.4 Python Execution Environment

The CloudVision Portal (CVP) python execution is supported by several CVP-specific libraries. These libraries provide access to the various CVP services and device state.

12.1.4.2 CVP Global Variables and Supported Methods

This library give access to the current execution context for Configlet Builders (see the provided examples for usage details).

The supplied global variables are:

```
from cvplibrary import CVPGlobalVariables, GlobalVariableNames
CVPGlobalVariables.getValue(GlobalVariableNames.CVP_USERNAME)
```

```
Supported GlobalVariableNames:
   CVP_USERNAME - Username of the current user
   CVP_PASSWORD - Password of the current user
   CVP_IP - IP address of the current device
   CVP_MAC - MAC of the current device
   CVP_SERIAL - Serial number of the current device
   CVP_SESSION_ID - Session id of current cvp user
   ZTP_STATE - ZTP state of the device (true/false)
   ZTP_USERNAME - Default username to login to ztp enabled device
   ZTP_PASSWORD - Password to login to ztp enabled device
   CVP_ALL_LABELS - Labels associated to current device
   CVP_CUSTOM_LABELS - Custom labels associated to current device
   CVP_SYSTEM_LABELS - System/Auto generated labels associated to current
 device
```

12.1.4.3 CVP Rest Client

This library allows a Configlet Builder to access any CVP API endpoint. The following is an example:

```
from cvplibrary import RestClient
url='http://localhost/cvpservice/inventory/devices';
method= 'GET';
client= RestClient(url,method);
if client.connect():
    print client.getResponse()
```

If no certificates are installed on the server, then add the following lines to ignore ssl warnings:

```
import ssl
ssl._create_default_https_context = ssl._create_unverified_contex
```

12.1.5 Creating Configlets Manually

CloudVision Portal (CVP) enables you to create Configlet manually. This method should be used to create Configlets that are relatively static.

Note: If you need to create Configlets that require less user input, you may want to use the Configlet Builder feature.

Complete these steps to manually create Configlets:

1. Select the "**+**" icon in the grid.
2. The **Create Configlet** page appears.

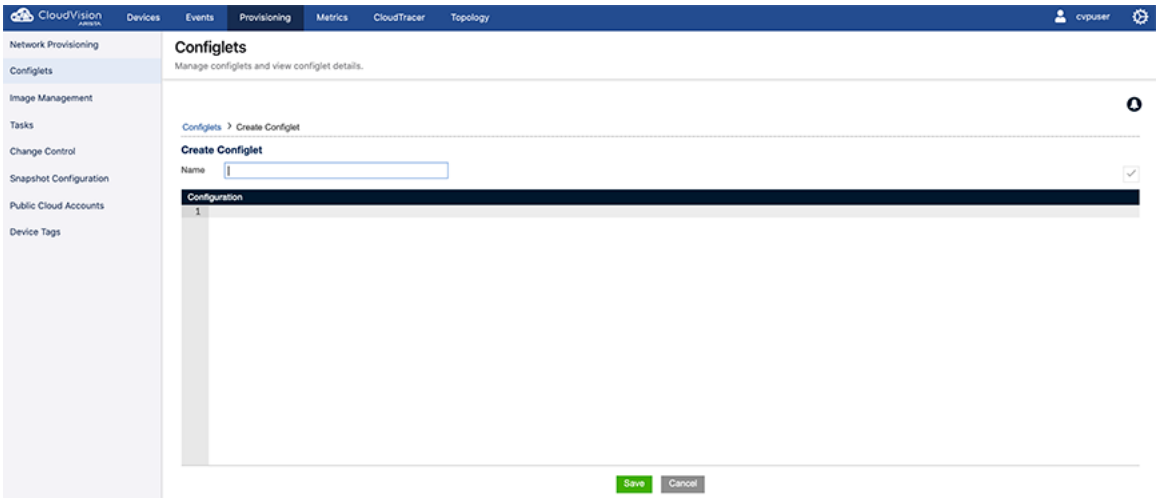

Figure 243: Create Configlet Page

3. Click **Save** to save the Configlet.
4. This will list the Configlet in the Configlet Management grid.

12.1.5.1 Validating a Configlet During Creation

CloudVision provides a facility to enter the Configlet code and validate it before saving the codes.

1. Enter the Configlet codes in the field provided.
2. On the right pane, there is a drop-down menu listing all the switches in CLOUDVISION.
3. Search for the device to be validated.

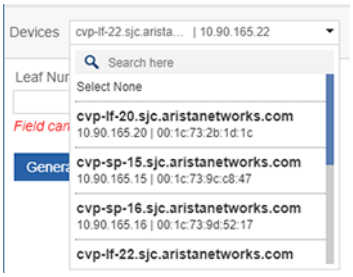

Figure 244: Validate-Search Device

4. Select the switch to validate.

Figure 245: Select Device

5. Select **Validate**.

On successful validation, the message Successfully Validated is displayed.

Figure 246: Validate-Success

When an error occurs, the message error will be displayed.

Figure 247: Validation Error

Related topics:

- Configlet Information Page
- #unique_380
- #unique_381
- Importing and Exporting Configlets

12.2 Configlet Information Page

1. Select the name of the Configlet from the grid to access the Configlet information page.

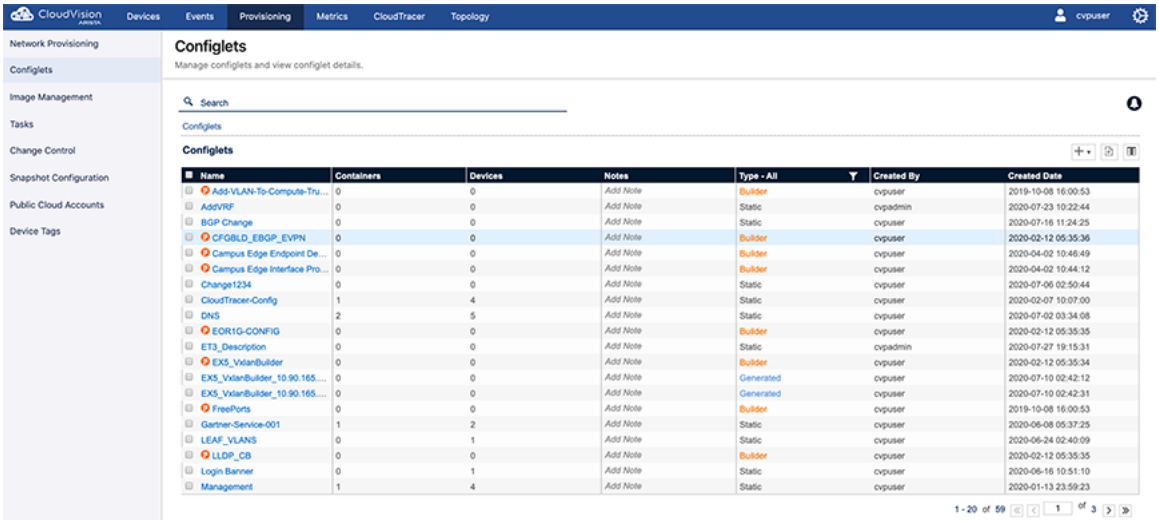

Figure 248: Configlet Information Page

12.2.1 Tabs in Configlet Information Page

The Configlet Information page consists of:

- Summary Tab
- Logs Tab
- Change History Tab
- Applied Containers Tab
- Applied Devices Tab

12.2.1.3 Change History Tab

Any change in the Configlets will be recorded in the **History** tab.

1. Select the **View** option.

   A popup window is opened comparing the last version of the Configlet with the edited version (Figure 249: Configlet History Page).

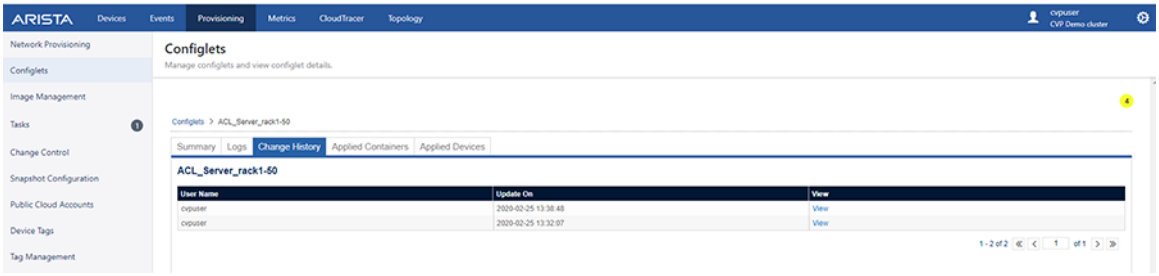

Figure 249: Configlet History Page

12.2.1.4 Applied Containers Tab

This tab gives the details on the containers to which the Configlet is assigned to. This also shows the name of the user who made the assignment (Figure 250: Applied Container Page).

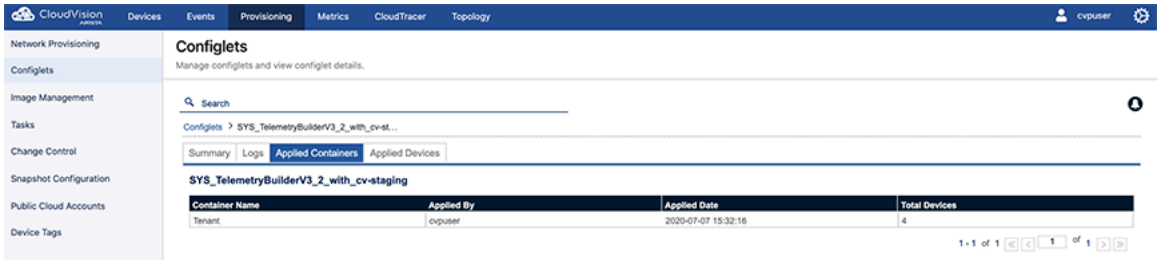

Figure 250: Applied Container Page

12.2.1.5 Applied Devices Tab

The **Applied Devices** tab displays the details on the devices to which the Configlet is associated in addition to other information such as **Parent container**, **Applied by**, and **Applied date**.

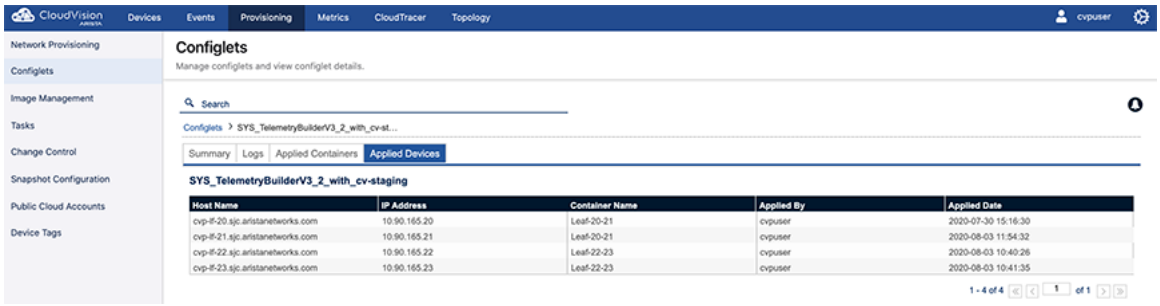

Figure 251: Applied Devices Page

When a Configlet is removed from any device through the Network Provisioning module, the device will be removed from the list.

Related topics:

- #unique_380
- #unique_381
- Importing and Exporting Configlets
- Creating Configlets

12.3 Editing Configlets

You edit Configlets through the Configlet “Summary” page. When you save the edited Configlet, it will update the all the associated tasks and devices in CLOUDVISION.

- Configuration assign tasks which are waiting to be executed in task management that are using the edited Configlet are considered as associated tasks.
- Saving the edited Configlet affects all the associated tasks as follows:

Pending tasks:	Tasks in pending state are auto updated. The spawned configuration points to the updated Configlet.
Failed tasks:	Tasks in a failed state are auto canceled. A new configuration push task is spawned.
Save As:	The edited Configlet can be saved as a new Configlet. Give the new Configlet a unique name.

1. Select the **Edit** (pen) icon in the page.

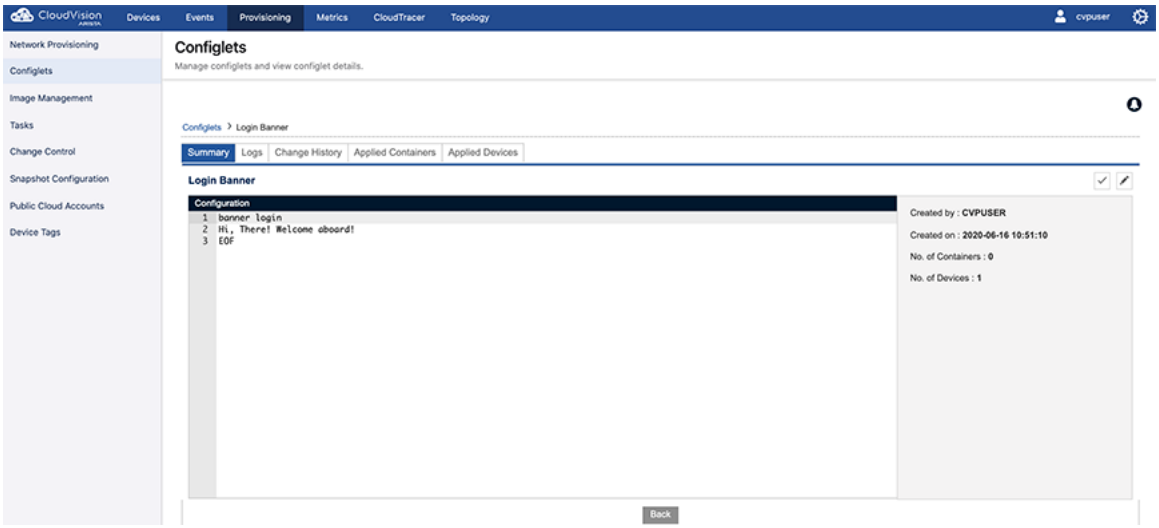

Figure 252: Configlet Summary Page

2. Validate the Configlet with the **Validation** pane.

Figure 253: Edit Configlet Summary

3. Do one of the following:
   - Click **Save** to save the edited configlet.
   - Click **Save As** to save the edited configlet as a new Configlet (the name Configlet).

Related topics:

- Deleting Configlets
- Importing and Exporting Configlets
- Creating Configlets
- Configlet Information Page

12.4 Deleting Configlets

Only unused Configlets can be deleted. If a Configlet is assigned to a device or a container, it cannot be deleted from the inventory. To delete a specific Configlet, its association should be removed from the devices and container.

1. Select a Configlet in the grid. A “trash can” icon will appear.
2. Click the **Trash** icon to delete the Configlet.

Related topics:

- Importing and Exporting Configlets
- Creating Configlets
- Configlet Information Page
- Editing Configlets

12.4.1 Importing and Exporting Configlets

You can import and export Configlets using the CloudVision graphical user interface (GUI). This enables you to easily share Configlets with others and back up specific Configlets.

For Configlets shared with you by another system user, you import Configlets from your desktop. When you share Configlets with another system user, you export Configlets to your desktop. You use the Configlets page to import and export Configlets or Configlet Builders.

Note: Both Configlets and Configlet Builders can be imported and exported using the GUI.

For more information, see:

- Protection from Overwriting Configlets or Configlet Builders
- Importing Configlets or Configlet Builders
- Exporting Configlets or Configlet Builders

12.4.1.1 Protection from Overwriting Configlets or Configlet Builders

CloudVision provides protection from accidentally overwriting exiting Configlets or Configlet Builders when importing a Configlet or Configlet Builder.

If you import a file that contains one or more Configlets or Configlet Builders that are named the same as Configlets or Configlet Builders already in CVP, the system automatically adds a suffix to the names of the items you are importing. The suffix that is added is in the format of “<number>”.

12.4.1.2 Importing Configlets or Configlet Builders

You import Configlets or Configlet Builders into CVP when another system user has shared a Configlet or Configlet Builder with you. Once you import Configlets or Configlet Builders, the imported items are available for use in CVP. You import Configlets or Configlet Builders from your desktop using the Configlets page.

Complete the following steps to import Configlets or Configlet Builders.

1. Open the Configlets page.
2. Click the Import icon, located in the upper right of the page.

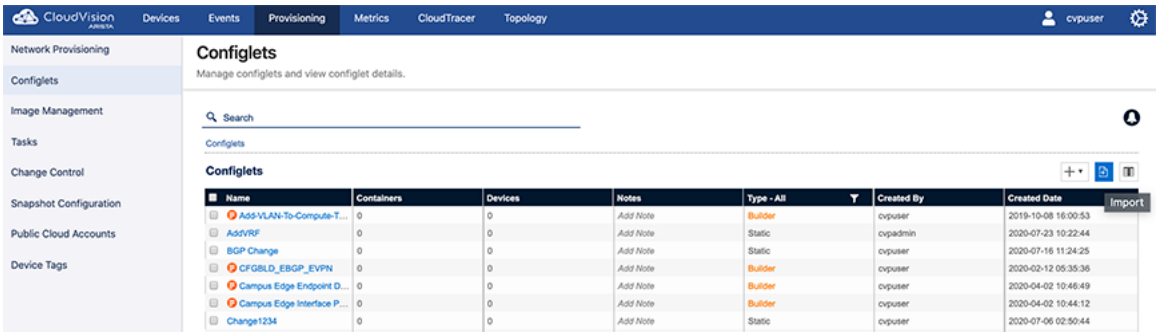

Figure 254: Configlets Page Showing Import Icon

A dialog appears that you use to select the file that contains the Configlets or Configlet Builders you want to import.

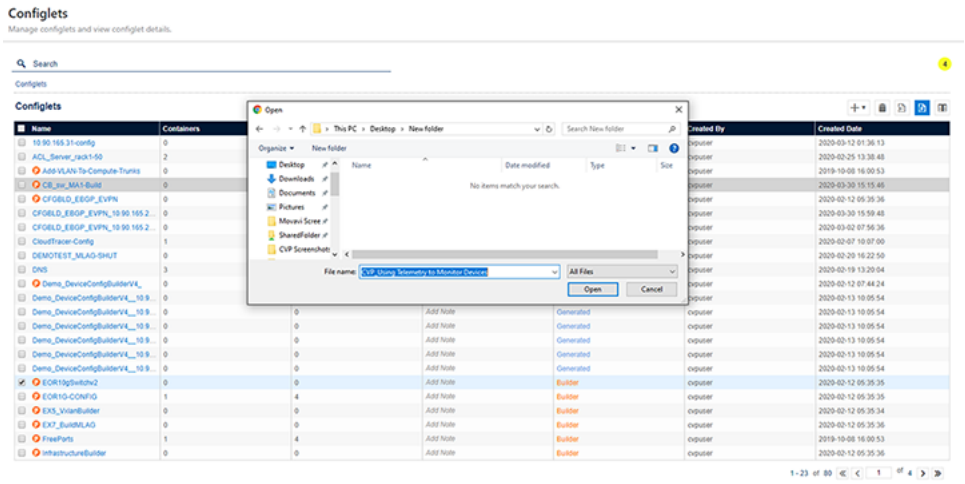

Figure 255: Selecting Configlets or Configlet Builders to be Imported

3. Select the file that contains the items you want to import.
4. Click **Open**.

   The Configlets or Configlet Builders in the file you selected are imported into CVP.

12.4.1.3 Exporting Configlets or Configlet Builders

You export Configlets or Configlet Builders when you want to share them with another system user. Once you export Configlets or Configlet Builders, the exported items are available to be sent to and then imported by the other system user. You export Configlets or Configlet Builders to your desktop using the Configlets page.

Complete the following steps to export Configlets or Configlet Builders.

1. Open the **Configlets** page.
2. Select the checkbox of each Configlet and Configlet Builder you want to export.

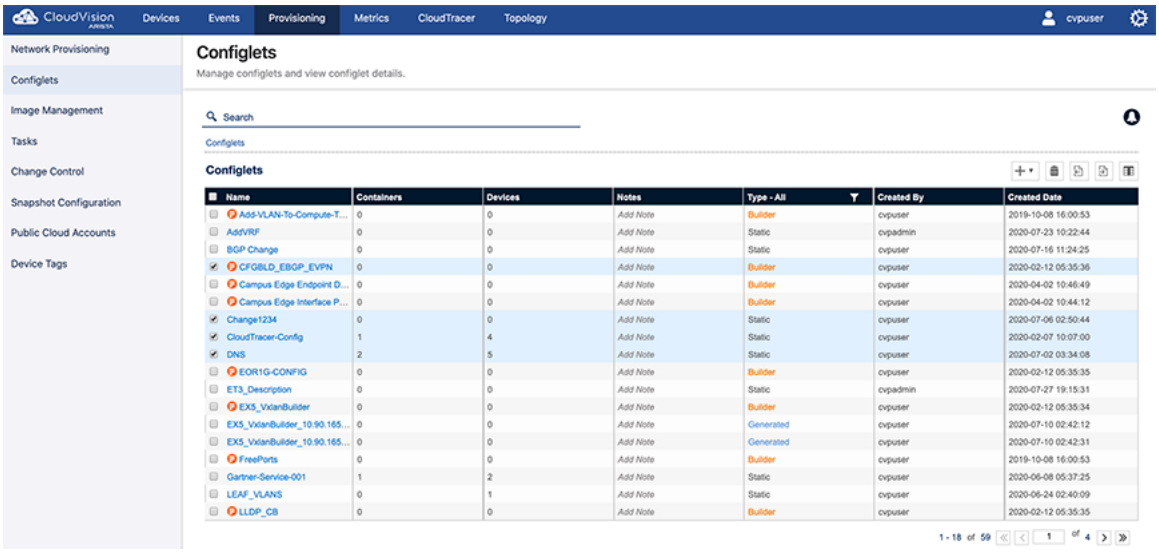

Figure 256: Configlets Page Showing Items Selected to be Exported

3. Click the **Export** icon (located in the upper right of the page).

   A single file (.zip archive) that contains all of the items you selected is automatically downloaded to your desktop.
4. (Optional) You can rename the downloaded file and make a copy of it before sharing it.
5. Share the file with one or more system users.

   Note: The items you share can be imported only on systems that support the import of Configlets and Configlet Builders (the Import icon on the Configlets page indicates support for this feature).

Related topics:

- Creating Configlets
- Configlet Information Page
- #unique_380
- #unique_381

Chapter 13

Image Management (CVP)

The Extended Operating System (EOS) used by the switches are uploaded into CloudVision, and details about them are maintained in the Image Management Inventory.

The main purpose of the Image Management module is to enable you to manage the EOS operating system images across the devices in your current CloudVision environment. It provides you with the functionality required to:

- Validate images
- Upload EOS images to CloudVision
- Maintain the inventory of available EOS images
- Assign images to devices in your CloudVision environment

Sections in this chapter include:

13.1 Image Management Page

The Image Management page shows the current operating system images that are available for upload to CloudVision. Once uploaded, they can be assigned to devices.

You can navigate to the Image Management page through Provisioning > Image Management.

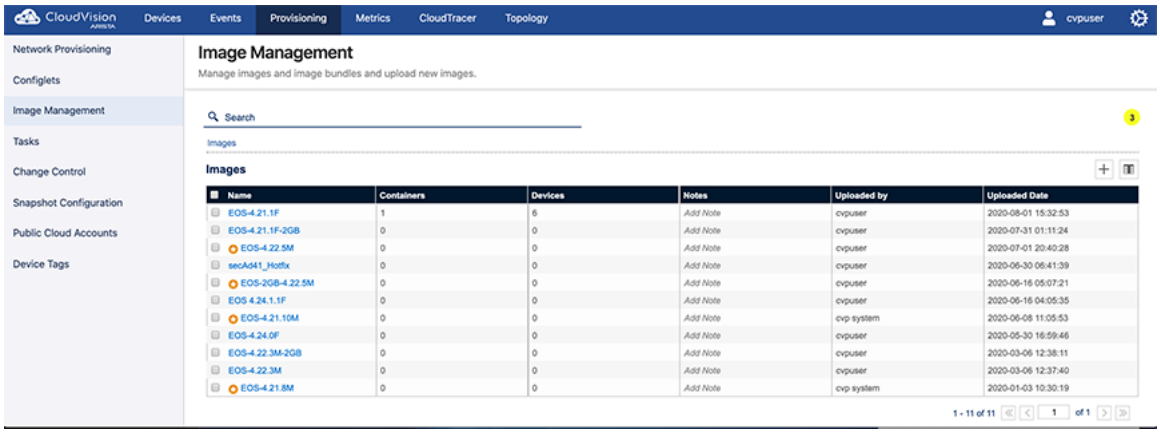

Figure 257: Image Management page

Related topics:

- Validating Images
- Upgrading Extended Operating System (EOS) Images
- Creating Image Bundles

- The Bundle Information Page

13.2 Validating Images

CloudVision Portal (CVP) provides automatic EOS image validation. This automated validation process helps to ensure that all devices in your CVP environment have EOS images that are supported by CVP.

The automatic validation of EOS images takes place whenever you:

- Upload images to CVP or add images to images bundles.
- Add devices to your CVP environment.

The automatic image validation ensures that images that are available to be included in image bundles and assigned to devices are supported by CVP.

Note: EOS images that are not supported cannot be added to an image bundle, or assigned to devices.

13.2.1 Alerts Indicating Unsupported EOS Image Versions

If you attempt to include an unsupported version of an EOS image when creating an image bundle, CVP alerts you with an error to let you know that the upload cannot be done, because the version of the EOS image you are trying to upload is not supported.

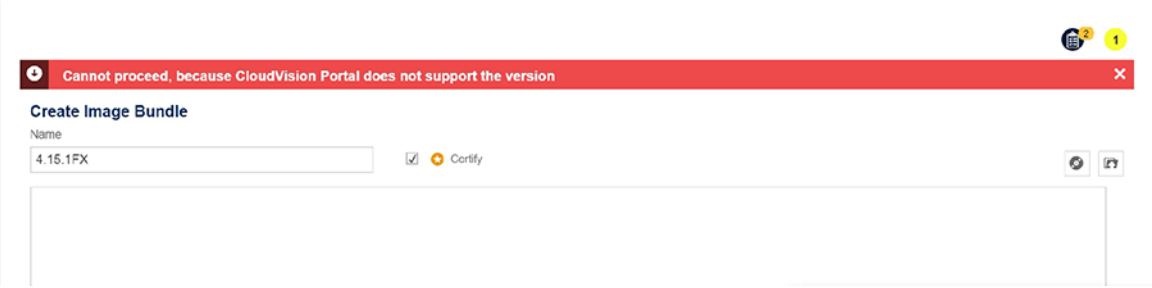

Figure 258: Alerts

If you attempt to add a device to CVP that has an unsupported EOS image, the Status column of the Inventory page indicates that an upgrade is required.

The Network Provisioning page also indicate that the device is running an unsupported image (this alert shows only when placing your cursor over the device icon).

Related topics:

- Upgrading Extended Operating System (EOS) Images
- Creating Image Bundles
- The Bundle Information Page
- Image Management Page

13.3 Upgrading Extended Operating System (EOS) Images

CloudVision Portal (CVP) provides the functionality to upgrade the EOS image on a device. Typically, you upgrade the image on a device to change the version of the image from an unsupported image version to a supported image version.

You upgrade device images by associating an EOS image with a device or a container (the association is referred to as an image association). Image associations follow the same container inheritance

rules as configlet associations. This means that the image you select to be associated is automatically inherited (assigned) to all devices under the level in the hierarchy at which you associate the image.

For more information, see:

- Example of Image Association
- Tip for Handling Multiple Image Association Tasks

13.3.1 Example of Image Association

This example shows the behavior of image associations in a multi-level network hierarchy. The hierarchy in this example contains a tenant container named Demo-Lab. The Demo-Lab container has five child containers named CVX, Host-TOR1, Leaf, Spine, and TOR2.

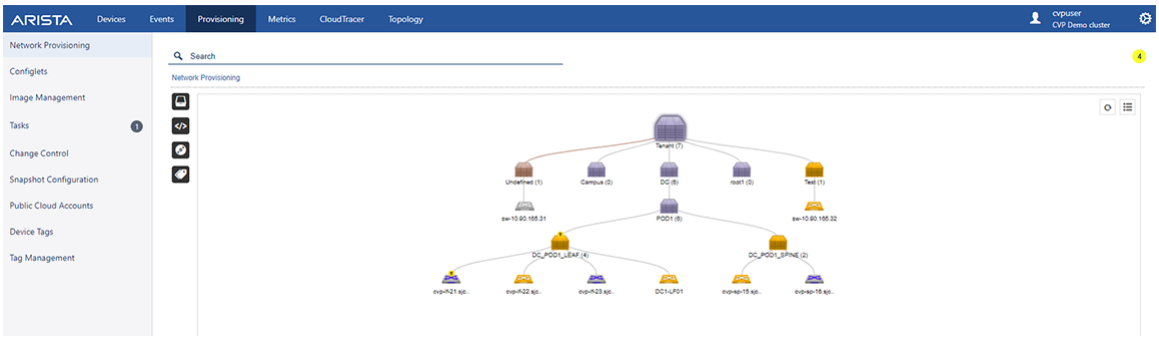

Figure 259: Same Task Scheduled for Every Device in CVX Container

Based on the rules for image association inheritance, the Demo-Lab container could have selected the *4.18.8M* device EOS image.

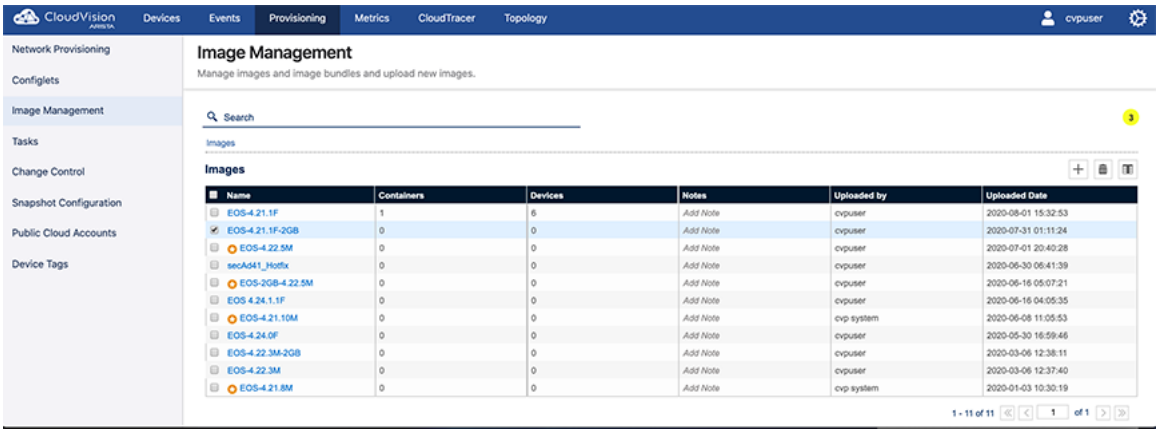

Figure 260: Example of image Association (Example 1)

The CVX container could override that image selection (*4.18.8M* image) for its devices by selecting the *4.20.7M* image. As a result, all of the devices under CVX are assigned the *4.20.7M* image, and the devices under Host-TOR1, Leaf, Spine and TOR2 inherit the *4.18.8M* image from the Demo-Lab container.

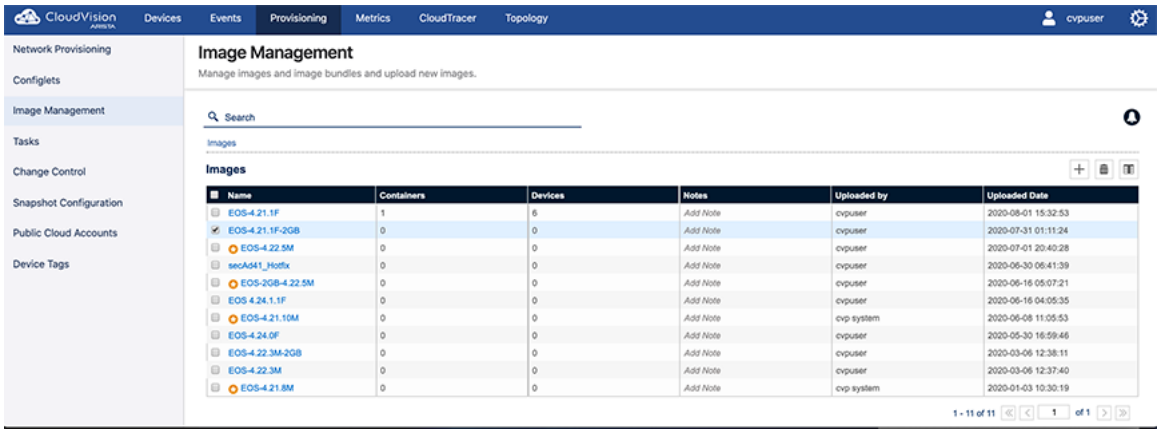

Figure 261: Example of image Association (Example 2)

If an image association is changed at any level, and the change is saved in the **Network Provisioning** page, the following occurs:

- The change impacts all devices under that level.
- A task is automatically created to upgrade the impacted devices.

For example, if the image selection was removed at the CVX level, the following would occur:

- All of the devices under the CVX level would inherit the Demo-Lab image.
- A task would be scheduled for every device in CVX to use the Demo-Lab image.

Related topics:

- Tip for Handling Multiple Image Association Tasks
- Creating Image Bundles
- The Bundle Information Page
- Image Management Page
- Validating Images

13.4 Creating Image Bundles

Creating image bundles is a key image management task. You create image bundles so that you have supported image versions available to be assigned to devices in your CVP environment.

Note: An image bundle must have one `.swi` file. Extensions are optional (not required for image bundles), but you can add one or more extensions to an image bundle.

Pre-requisite: To ensure that you include valid (supported) EOS images in the bundles you create, make sure you validate the images you want to include in the bundle (see Validating Images).

Complete the following steps to create an image bundle:

1. Go to the **Image Management** page.
2. Click the "**+**" icon in the grid.

   This loads the **Create Image Bundle** page.

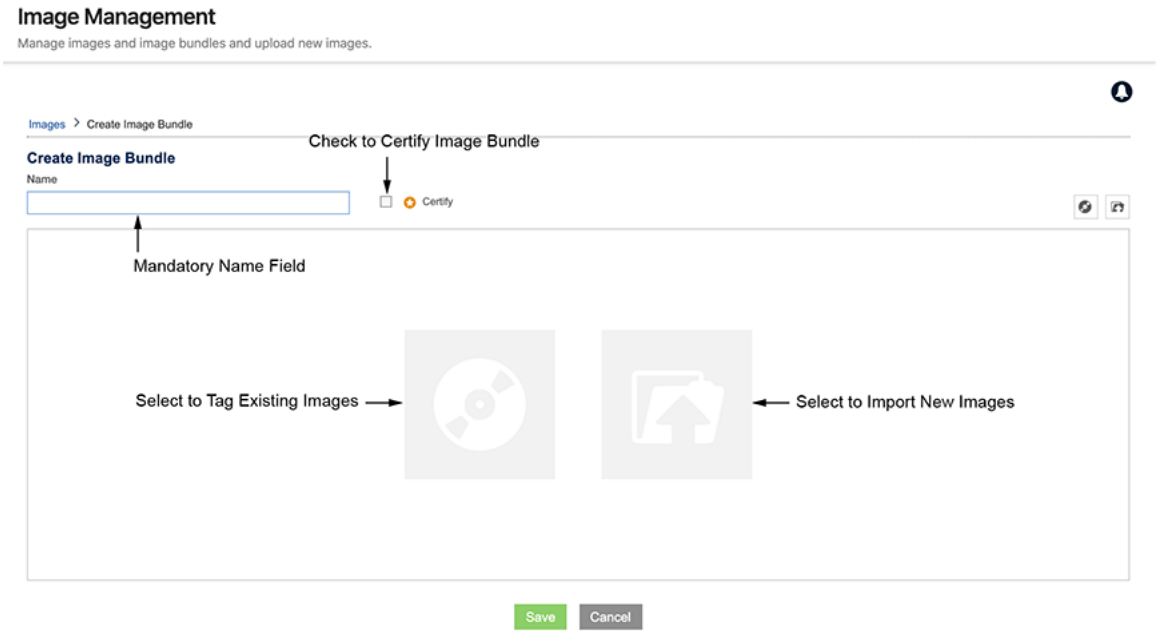

Figure 262: Create Image Bundle page

For more information, see:

- Creating a Bundle by Tagging Existing Image Bundles
- Creating a Bundle by Uploading a New Image
- Adding EOS Extensions to Image Bundles

Related topics:

- Creating a Bundle by Tagging Existing Image Bundles
- Creating a Bundle by Uploading a New Image

13.4.1 Creating a Bundle by Tagging Existing Image Bundles

CloudVision Portal (CVP) enables you to create a new image bundle by tagging existing image bundles. This prevents you from having to import the same image again to create another bundle.

1. Go to the **Image Management** page.
2. Click the "**+**" icon and then the Disk icon.
1. This opens the Images dialog, which lists all of the available images.

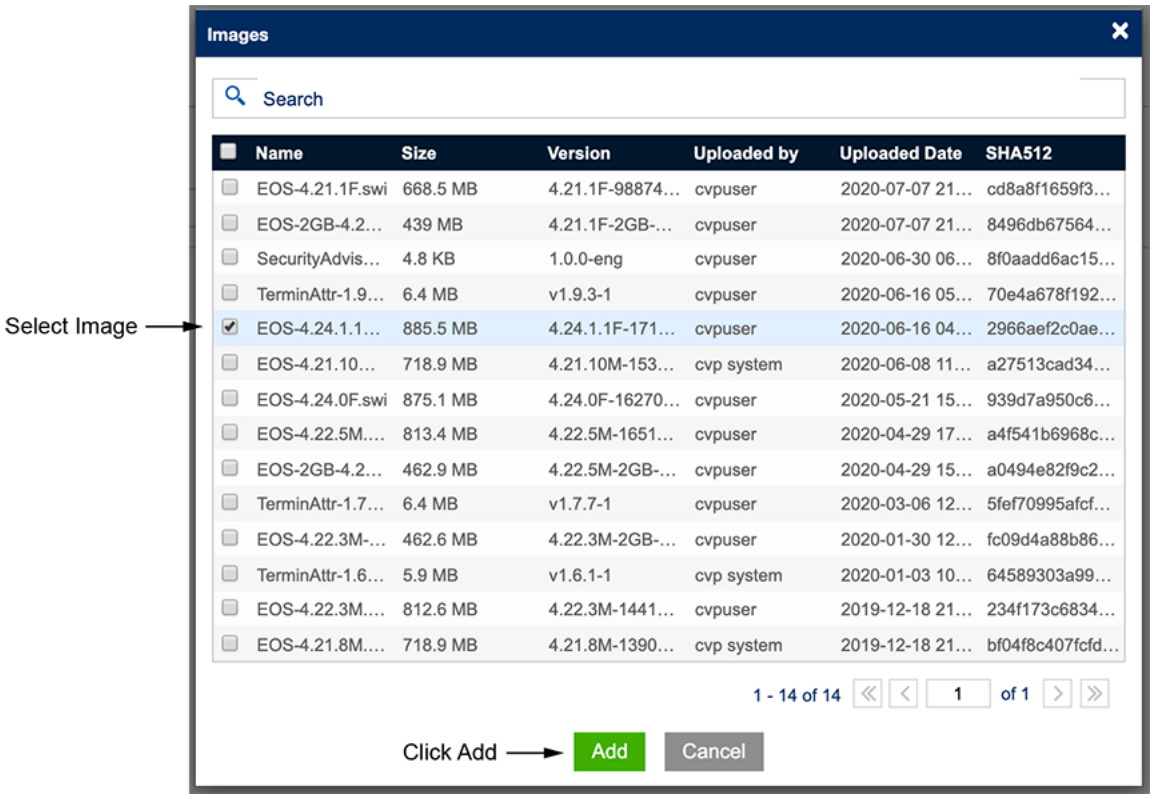

Figure 263: Images dialog

3. Search for the desired image.
4. Select the image and click **Add** to add the image to the bundle.

   The image will be displayed in the grid of the **Create Image Bundle** page.

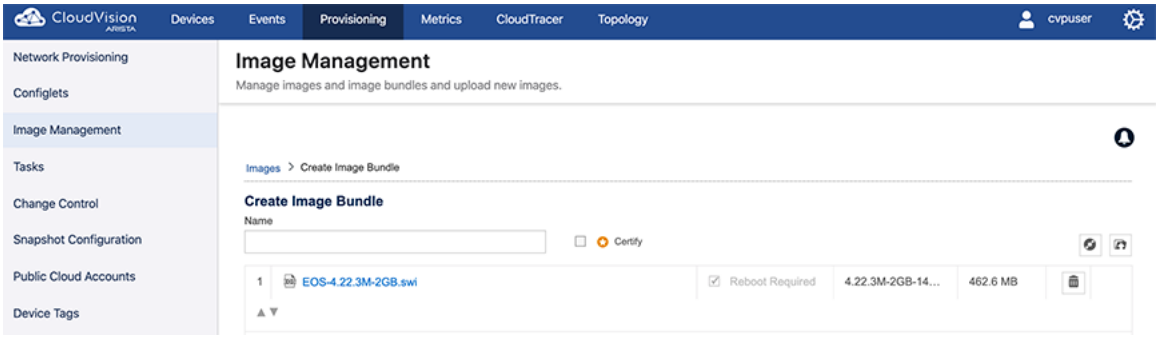

Figure 264: Added image shown in Create Image Bundle page

5. Click **Save** to create the new image bundle.

Related topics:

- Creating a Bundle by Uploading a New Image
- Adding EOS Extensions to Image Bundles

13.4.2 Creating a Bundle by Uploading a New Image

CloudVision Portal (CVP) enables you to create new image bundles by uploading new images to CVP.

1. Go to the **Create Image Bundle** page.
2. Click the upload from local icon available next to disk icon.

   This opens a dialog to search and upload .swi files from system.
3. Navigate to the desired .swi file and upload it to CVP.

The upload bar on the page shows the progress of the upload.

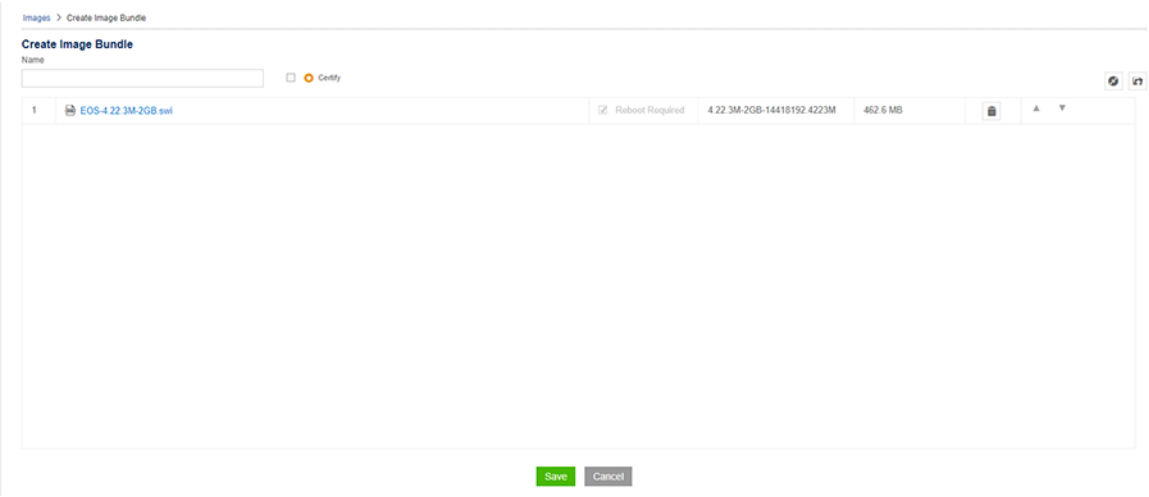

Figure 265: Uploading .swi files to CVP (upload in progress)

4. Click **Save** to create the new image bundle.

Related topics:

- Adding EOS Extensions to Image Bundles
- Creating a Bundle by Tagging Existing Image Bundles

13.4.3 Adding EOS Extensions to Image Bundles

CloudVision Portal (CVP) enables you to add EOS extensions to image bundles along with .swi images. Extensions are either .rpm files or .swix files. You upload .rpm or .swix files using the Images page. Extensions are optional for image bundles

Note: To verify that all the extensions you selected are installed and running on the device, run a compliance check on the device after you install the image bundle on the device.

Complete these steps to add EOS extensions to an image bundle:

1. Go to the **Create Image Bundle** page.
2. Click the upload from local icon.

   This opens a dialog to search and upload EOS extensions (.rpm or .swix files) from the system
3. Navigate to the desired .rpm or .swix files and upload them.

   The upload bar on the page shows the progress of the upload. The extensions you uploaded are shown in the Create Image Bundle page

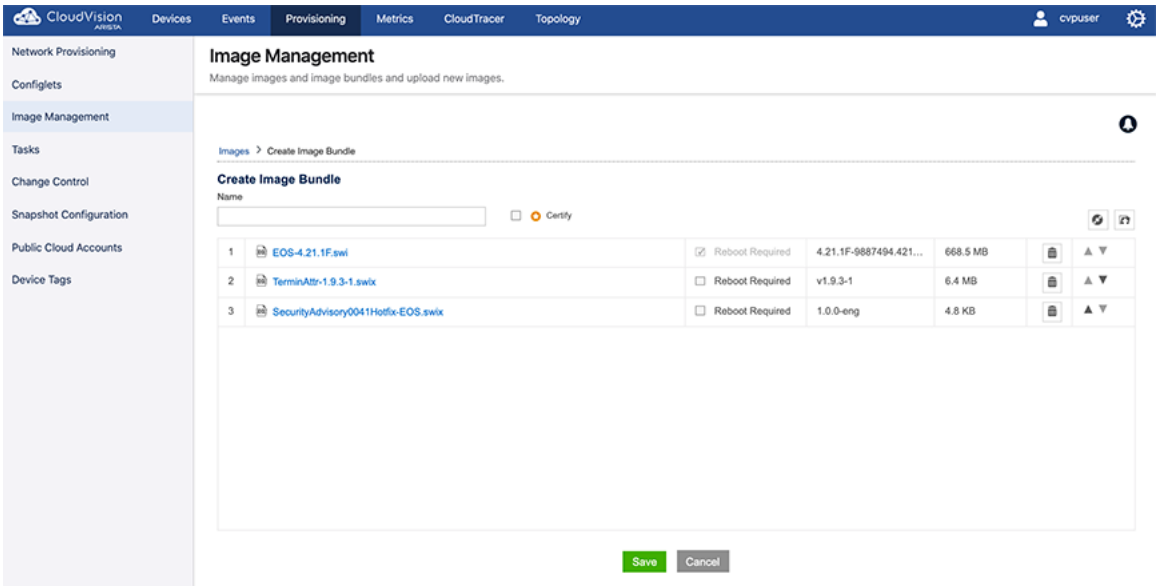

Figure 266: Create Image Bundle showing uploaded extensions

4. Select **Reboot Required** check-boxes for all extensions that require a reboot. (All uploaded extensions in this example require a reboot.)
5. Click **Save**. The extensions are added to the image bundle.

   Once the image bundle is assigned to a device, a reboot task will be generated. The newly added extensions are installed on the device when the reboot task is executed. Any extensions that were previously installed but are not part of the current bundle are removed from the device.

13.5 The Bundle Information Page

The Image Management page provides high-level information about an image bundle (for example, the number of containers to which an image bundle is associated, and the number of devices to which an image bundle is assigned).

To view more detailed information about image bundles, use the Bundle Information page, which you can open from the Image Management page.

Complete these steps to open the **Bundle Information** page.

1. Go to the **Image Management** page.
2. Click the name of image bundle for which you want to view information.

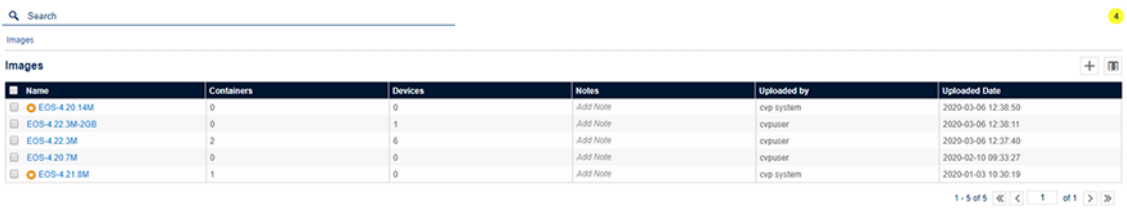

Figure 267: Opening the Bundle Information page

The **Bundle Information** page appears, showing information for the selected image bundle. Use the following tabs to view specific information about the selected image bundle.

- Summary Tab
- Logs Tab

- Applied Containers Tab
- Applied Devices Tab

13.5.1 Summary Tab

The Summary tab provides basic information about the Image Bundle. It also provides options to go back to the **Image Management** page, to open the dialog used to update image bundles, and to delete corresponding image bundle and its extensions.

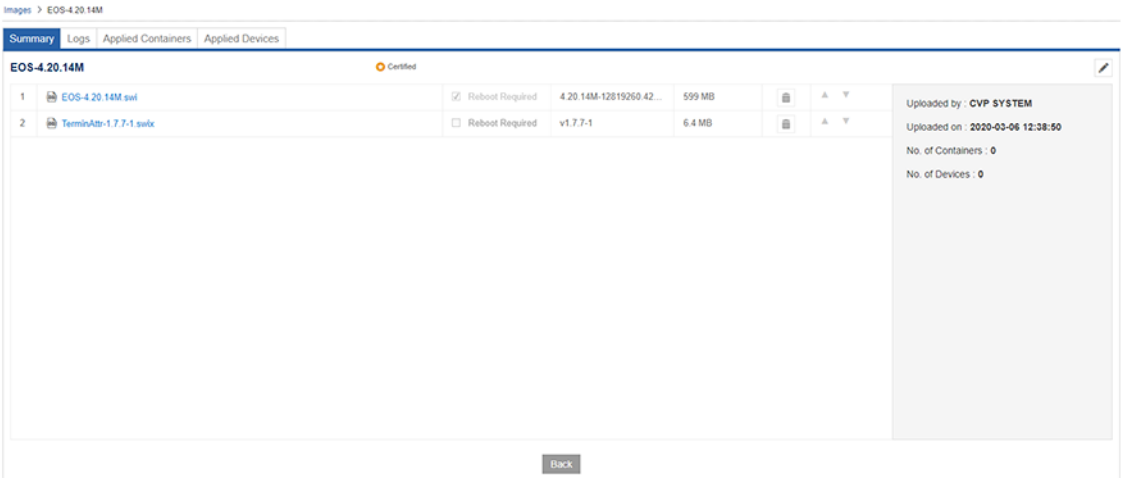

Figure 268: Summary tab

For details on the steps used to edit image bundles and delete image bundles, see:

- Updating Bundles
- Deleting Bundles

13.5.2 Logs Tab

The Logs tab provides complete information on the image assignment to devices and execution details. It also provides the option to go back to the **Image Management** page.

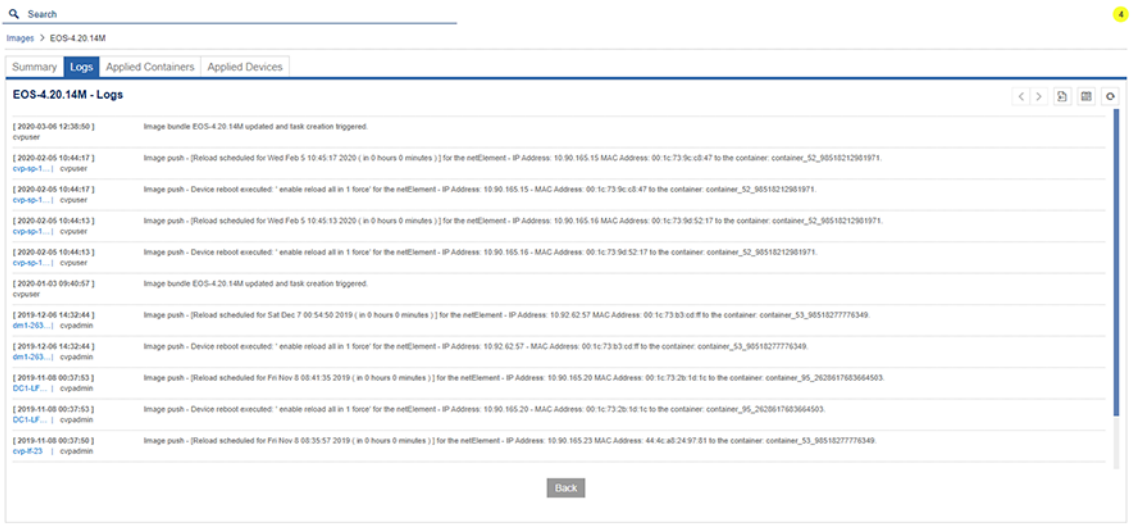

Figure 269: Logs tab

13.5.3 Applied Containers Tab

The Applied Containers tab displays the details on the containers to which the bundle has been applied. It also displays the name of the user that applied the bundle and the date it was applied.

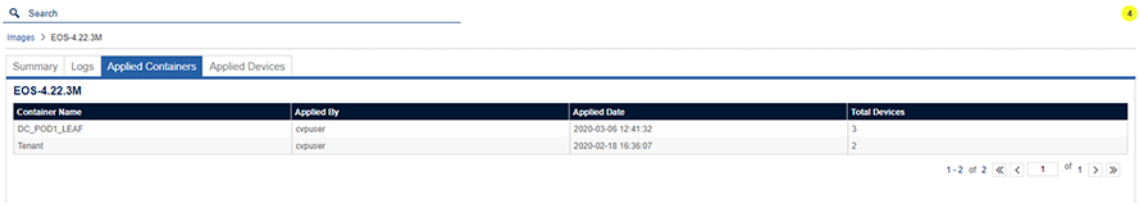

Figure 270: Applied Container tab

13.5.4 Applied Devices Tab

The **Applied Devices** tab displays the details on the devices to which the bundle is assigned, along with other information such as the parent container for the device, and the name of the user that applied the bundle and the date it was applied.

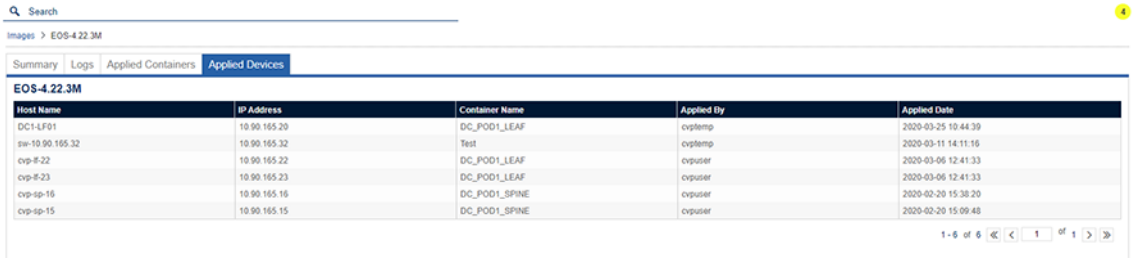

Figure 271: Applied Devices tab

Related topics:

- Summary Tab
- Logs Tab
- Applied Containers Tab

13.5.5 Updating Bundles

Perform the following steps to update a bundle:

1. Go to the **Image Management** page.
2. Click the name of image bundle that you want to update.

   The system displays the **Summary** tab.

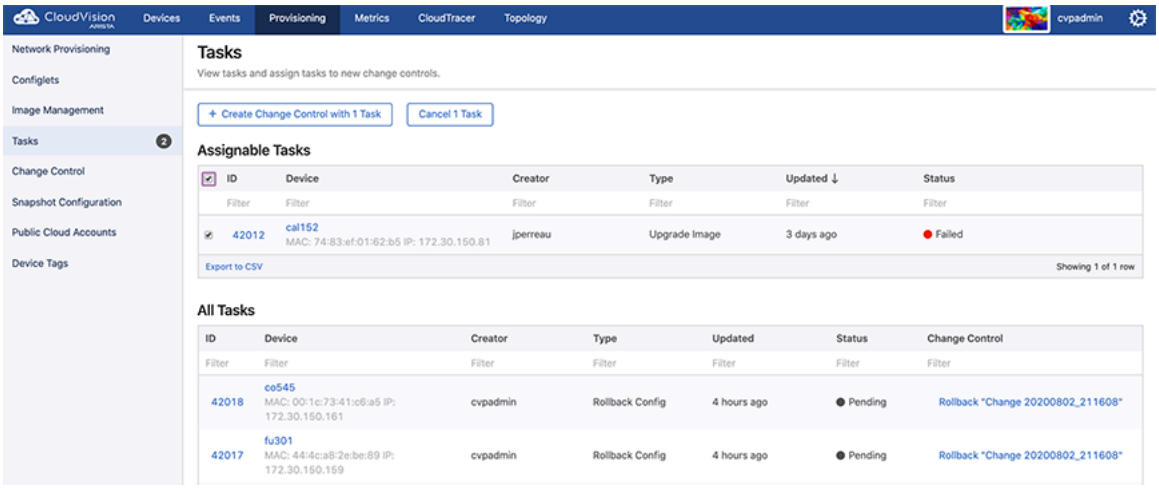

Figure 272: Summary page showing bundle selected for edit

3. Click the edit icon at the upper right corner of the Summary section.
4. Edit the bundle as needed.
5. Click **Save**.

Related topics:

- Deleting Bundles

13.5.6 Deleting Bundles

Only unused bundles can be deleted. If a bundle is assigned to a device or a container, it cannot be deleted from the inventory.

Perform the following steps to delete a bundle:

1. Go to the **Image Management** page.
2. Click the name of image bundle that you want to delete.

   The system displays the **Summary** tab.
3. Click the edit icon at the upper right corner of the **Summary** section.

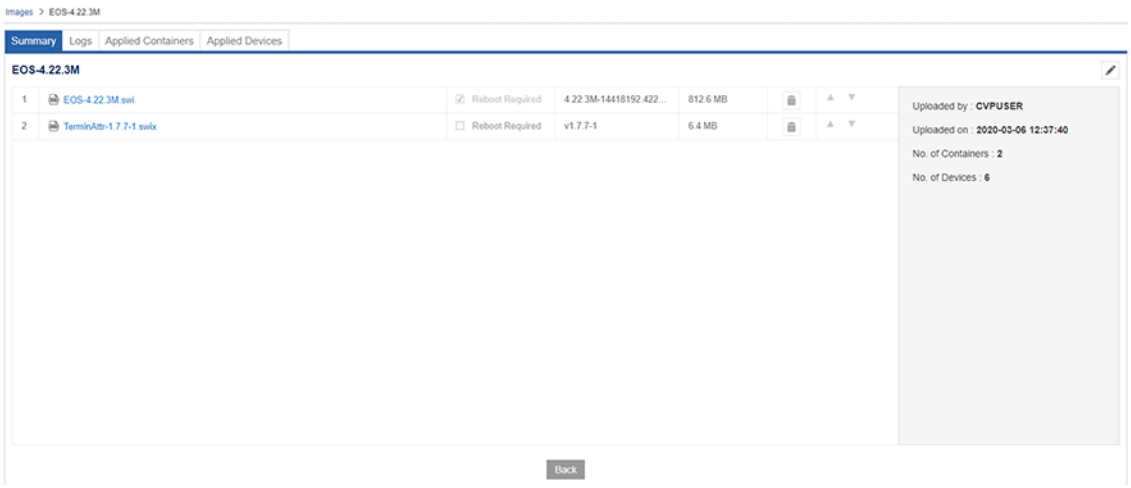

Figure 273: Summary page showing bundle selected for deletion

4. Click the trash icon to delete the selected bundle from the inventory.

   The system prompts to confirm the deletion.
5. Click **Yes** to confirm deletion.

6. Click **Save**.

Note: The association can be removed only if a new bundle is assigned to device or container.

Note: When an image bundle is assigned to a container, no task will be spawned to the subordinate devices.

Related topics:

- Updating Bundles

Chapter 14

Change Control

Task Management is an inventory of all the tasks generated in CloudVision. You can create a Change Control or cancel a task in task management.

Sections in this chapter include:

14.1 Basic Options for Handling Tasks

CloudVision provides two basic ways to handle tasks. You can handle tasks individually (task by task), or by groups of tasks.

To view and cancel tasks individually, use the Task Management module, which you can access by navigating to **Provisioning > Tasks** from the CloudVision Portal. For detailed information on the Tasks module, see Using the Tasks Module.

To execute grouped tasks (multiple tasks in the same group), use the Change Control module from either Tasks or Change Control screens. To access the Change Control screen, navigate to **Provisioning > Change Control** from the CloudVision Portal. For detailed information on the Change Control module, see Using the Tasks Module.

14.1.1 Creating Tasks

The following actions that affect the performance of devices are automatically generated as tasks:

- Assigning Configuration (assigning a configuration to a device or container)
- Adding Devices (adding a device from the undefined container to a defined container)
- Managing Devices (Moving or removing devices from a container)

14.1.1.1 Assigning Configuration

1. Go to the Network Provisioning screen.
2. Select a device or container.
3. Assign configuration.
4. Save the topology to generate the task.

Note: Editing a configlet also generates a task.

14.1.1.2 Adding Devices

1. Go to the Network provisioning screen.
2. Select a container.
3. Add devices to the container.
4. Save the topology to generate the task.

Note: If the hierarchy of the container has images or configlets, the created task will also include image push and configuration push tasks.

14.1.1.3 Managing Devices

1. Go to the Network provisioning screen.
2. Select a container.
3. Move or remove devices from the container.
4. Save the topology to generate the task.

14.2 Using the Tasks Module

This module covers the following sections:

- Accessing the Tasks Summary Screen
- Creating Change Controls from the Change Controls Summary Screen
- Accessing the Tasks Details Screen
- Task Status

14.2.1 Accessing the Tasks Summary Screen

Use the **Tasks Summary** screen to create Change Controls, cancel tasks, view assignable and assigned tasks, navigate to the appropriate task details screen, and navigate to the device overview screen. See **Task Screen** below.

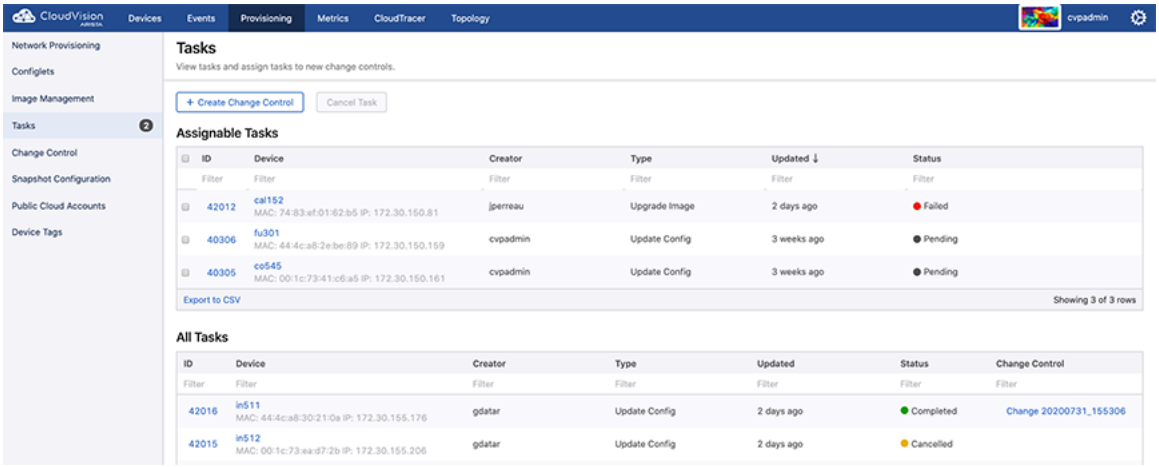

Figure 274: Tasks Screen

To access the **Tasks Summary** screen, go to the **Provisioning** screen and click **Tasks** in the left menu.

The **Tasks Summary** screen consists of the following entities:

- **+ Create Change Control button** - Click this button to create a Change Control
- **Cancel Task(s) button** - Click this button to cancel selected assignable tasks
- **Assignable Tasks Table** - Lists assignable tasks with the following information:
  - **Task ID** - Displays the task ID.

    Click the **Task ID** go to the appropriate task details screen.
  - **Device** - Displays the device name on which this task is performed.

    Click the device name to open the appropriate **Device Overview** screen.

- **Created By** - Displays who created the task.
- **Type** - Displays the task type.
- **Last Updated** - Displays when the task was last updated.
- **Status** - Displays the task status.

- **Assigned Tasks Table** - Lists assigned tasks with the following information:
  - **Task ID** - Displays the task ID.

    Click the task ID go to the appropriate task details screen.
  - **Device** - Displays the device name on which this task is performed.

    Click the device name to open the appropriate **Device Overview** screen.
  - **Created By** - Displays who created the task.
  - **Type** - Displays the task type.
  - **Last Updated** - Displays when the task was last updated.
  - **Status** - Displays the task status.
  - **Change Control** - Displays the Change Control name.

    Click the Change Control name to go to the appropriate **Change Control Details** screen.

14.2.2 Creating Change Controls from the Tasks Summary Screen

The Change Control module selects and executes a group of tasks that you want to process simultaneously. While creating a Change Control, you add tasks with pending or failed status to the Change Control.

Complete the following steps to create a Change Control from the tasks summary screen:

1. On the CloudVision Portal, click **Provisioning > Tasks.**

   The system displays the tasks summary screen.
2. Under the Assignable Tasks table, select tasks you want to include in the Change Control by selecting appropriate checkboxes.

   Note: If you do not select any tasks, the system creates a Change Control without tasks.
3. Click **+ Create Change Control** with *n tasks* where n is the count of selected tasks.

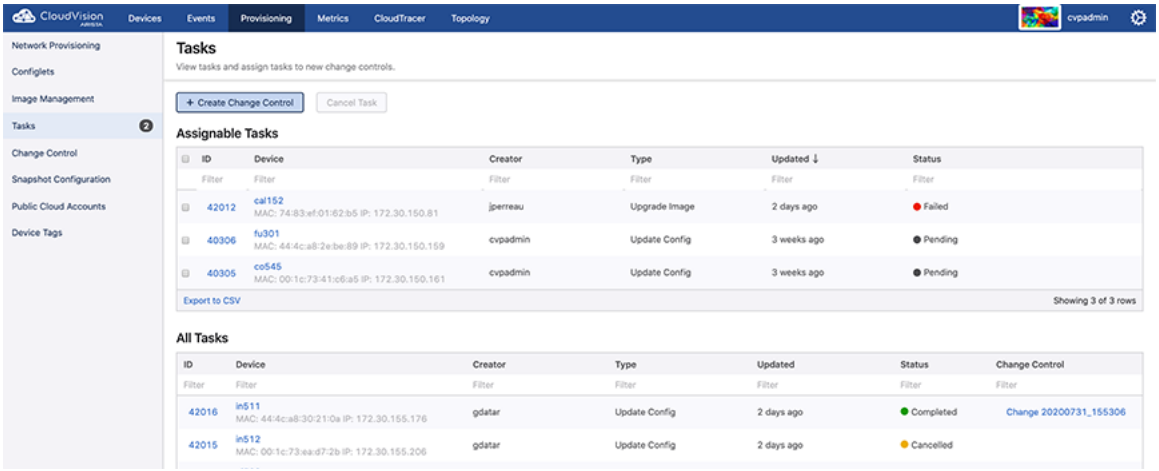

Figure 275: Create Change Control Button

The system displays the appropriate Change Control details screen.

14.2.3 Accessing the Tasks Details Screen

The **Tasks details** screen provides detailed information for any given task. To access the Tasks details screen, click the task ID under the **Task ID** column in the **Tasks summary** screen.

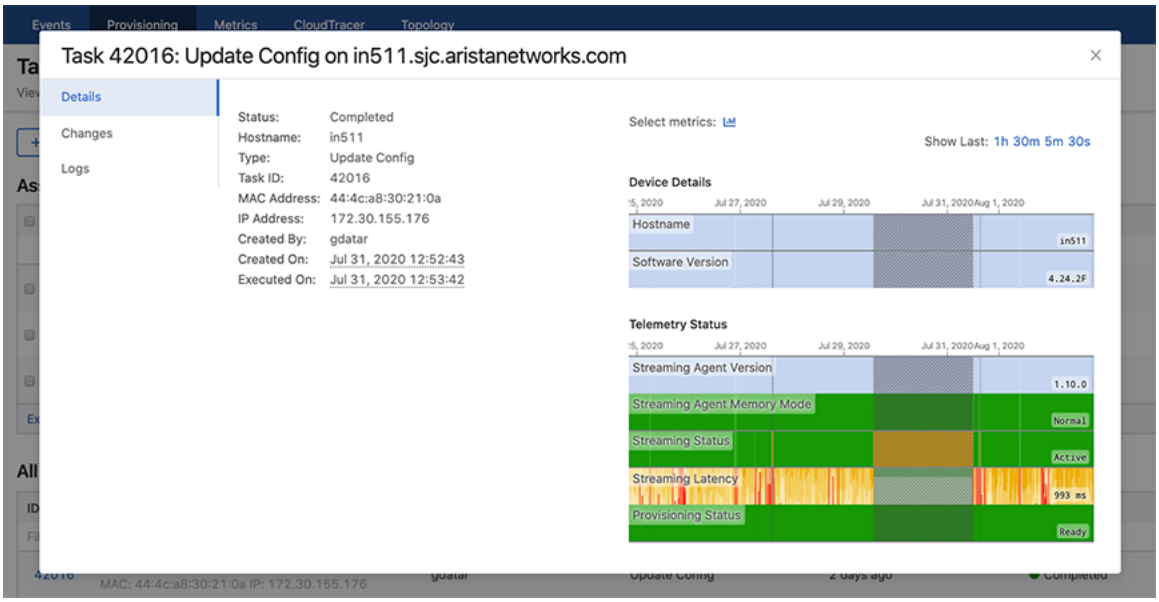

Figure 276: Task Details Screen

The **Tasks Details** screen provides the specified information in following tabs:

- **Pending tasks** icon - Displays the count of pending tasks
- **Notifications** - Displays the count of unread notifications.
- **Logs** tab - Displays logs of the appropriate task.

  Note: This tab is displayed only for completed tasks.

- **View Image** tab - Provides detailed information on image changes.

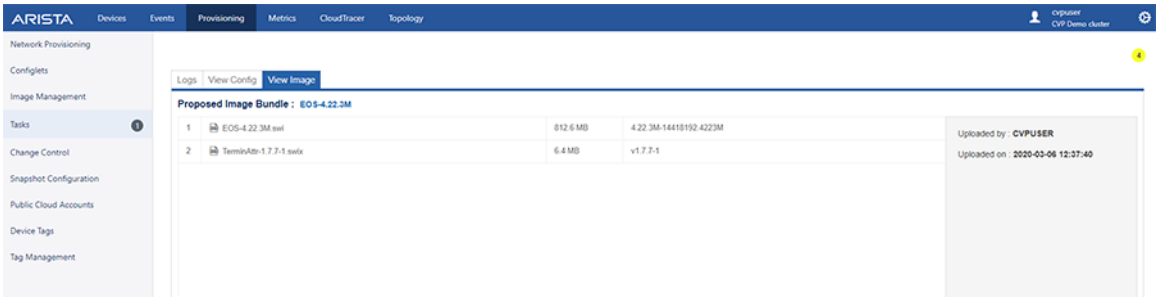

Figure 277: View Image Tab

- **View Config** tab - Displays provisioned, designed, and running configuration changes.

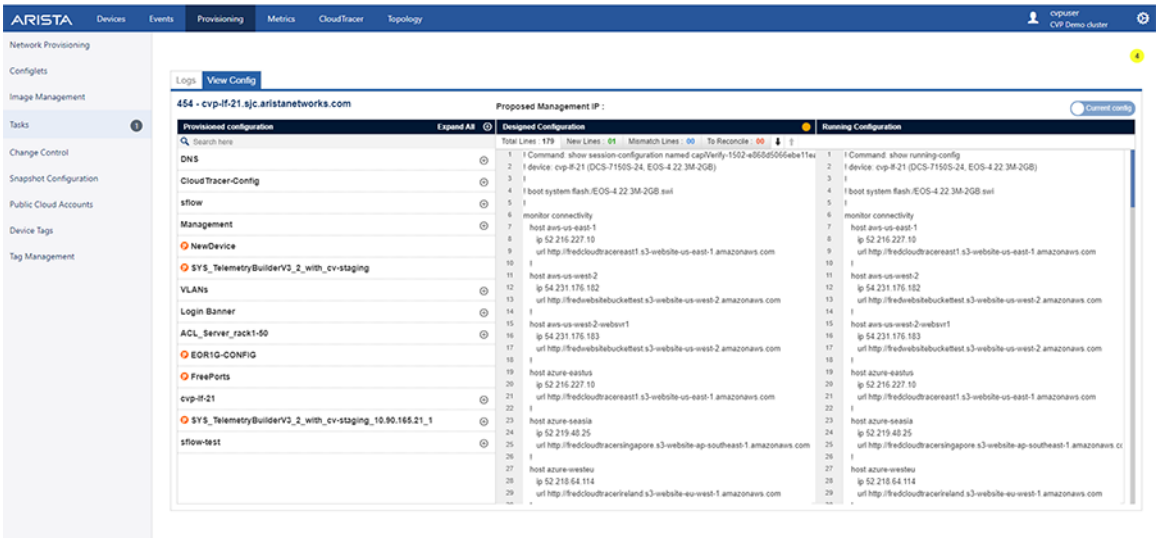

Figure 278: View Config Tab

14.2.4 Task Status

All CloudVision Portal (CVP) tasks are automatically assigned a specific status by the system. The system automatically updates tasks status to indicate the current status of a task.

The task statuses are:

- Pending
- In-Progress
- Completed
- Failed
- Canceled

14.2.4.1 Pending

Any new task is generated with a 'Pending' status. This means that the task has been generated but not executed. You can execute a pending task at any time. Once the task is successfully executed (completed without failure), the status of the task changes to Completed.

14.2.4.2 In-Progress

A task being executed moves to “In-progress” state.

- Config assign, pushes the configuration on the device.
- Image assign, copies the image from CLOUDVISION to the device.
- In-Progress tasks can be canceled.

Various statuses during the Change Control execution are:

- Execution In Progress
- Device Reboot In Progress
- Task Update In Progress
- Configlet Push In Progress
- Image Push In Progress
- Rollback Config Push In Progress
- Rollback Image Push In Progress
- Cancel In Progress
- ZTR Replacement In Progress

14.2.4.3 Completed

A task that has been completed. Upon completion, the status changes to Completed. Tasks with Completed status can't be executed or canceled.

14.2.4.4 Failed

A task moves to failed state due to multiple reasons such as:

- Device not reachable
- Wrong configuration
- Application problem

14.2.4.5 Canceled

A task that is removed from the queue of pending tasks. Tasks with the status of Completed or tasks that have already been canceled, cannot be canceled. Tasks with any status other than Canceled or Completed can be selected and canceled.

14.3 Using the Change Control Module

The **Change Control** module selects and executes a group of tasks that you want to process simultaneously. Selecting tasks and creating Change Controls function similarly in **Change Control** and **Task Management** modules.

Change Controls provides the following benefits:

- Sequencing tasks
- Adding unlimited snapshots to every device impacted by the Change Control execution
- Adding custom actions
- Pushing images via Multi-Chassis Link Aggregation (MLAG) In-Service Software Upgrade (ISSU) or Border Gateway Protocol (BGP) maintenance mode
- Reviewing the entire set of changes to approve Change Controls

Note: Snapshots display the state of impacted devices before and after the execution.

For more information about Change Controls, see:

- Accessing the Change Control Summary Screen
- Creating Change Controls from the Tasks Summary Screen
- Accessing the Open Change Control Details Screen

14.3.1 Accessing the Change Control Summary Screen

The Change Control summary screen is used to manage Change Controls.

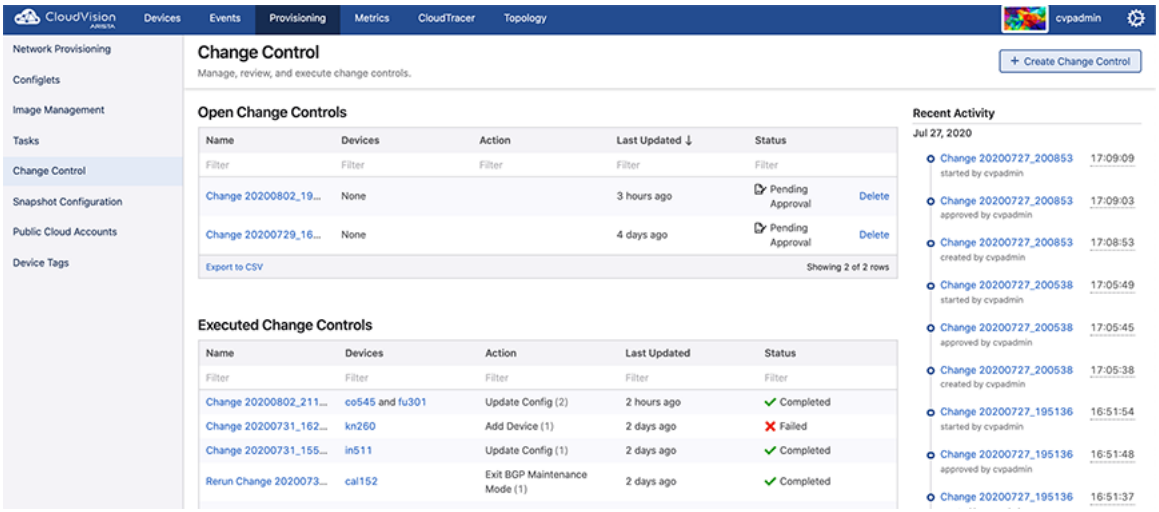

Figure 279: Change Control Summary Screen

To access the Change Control screen, go to the Provisioning screen, and click Change Control in the left menu.

The Change Control screen consists of the following entities:

- **Open Change Controls** and **Executed Change Controls** tables - Lists corresponding Change Controls with the following information:
  - **Name** - Displays the Change Control name

    Click the Change Control name to go to the appropriate Change Control details screen.
  - **Devices** - Displays devices used in the Change Control

    Click the device name to go to the appropriate Device Overview screen.
  - **Action** - Displays types of actions to be executed by the Change Control
  - **Last Updated** - Displays when the Change Control was last updated
  - **Status** - Displays the Change Control status

    Note:

    - Under the **Status** column of the **Open Change Controls** table, a pending Change Controls is represented with a doc-edit icon and an approved Change Controls is represented with a user-check icon.
    - Under the **Status** column of the **Open Change Controls** table, a failed Change Control is represented with a cross mark and a completed Change Control is represented with a tick mark.
    - Hover the cursor on the status icon in **Open Change Controls** table to view how long ago the current approval status was updated. When you hover the cursor on the status icon in **Executed Change Controls** table, it also displays the approver's name.
- In the **Open Change Controls** table, click **Delete** to delete the appropriate Change Control.

  Note: After you delete an open Change Control, the system returns any tasks used by the deleted Change Control to the assignable tasks pool for reallocation.
- **Recent Activity** pane - Lists most recent activities like updated, executed, and deleted Change Controls.

  Note: Click on the Change Control name to go to the appropriate Change Control details screen.
- **+ Create Change Control** - Click this button to create a Change Control
- **Export to CSV** - Exports the summary data to a CSV file.

14.3.2 Creating Change Controls from the Change Controls Summary Screen

The first step involved in using the **Change Control** module to manage tasks is to create a Change Control. While creating a Change Control, you add tasks with pending or failed status to the Change Control. By default, all tasks in the same Change Control are added in parallel. If you want to change the execution order, you can drag and drop the action cards on the **Change Control Details** screen. You can execute grouped tasks after a Change Control is created, reviewed, and approved.

Note: If you do not add any tasks, the system creates a Change Control without tasks.

Complete the following steps to create a Change Control from the **Change Control Summary** screen:

1. On the CloudVision Portal, click **Provisioning > Change Control**.

   The system displays the **Change Control Summary** screen.

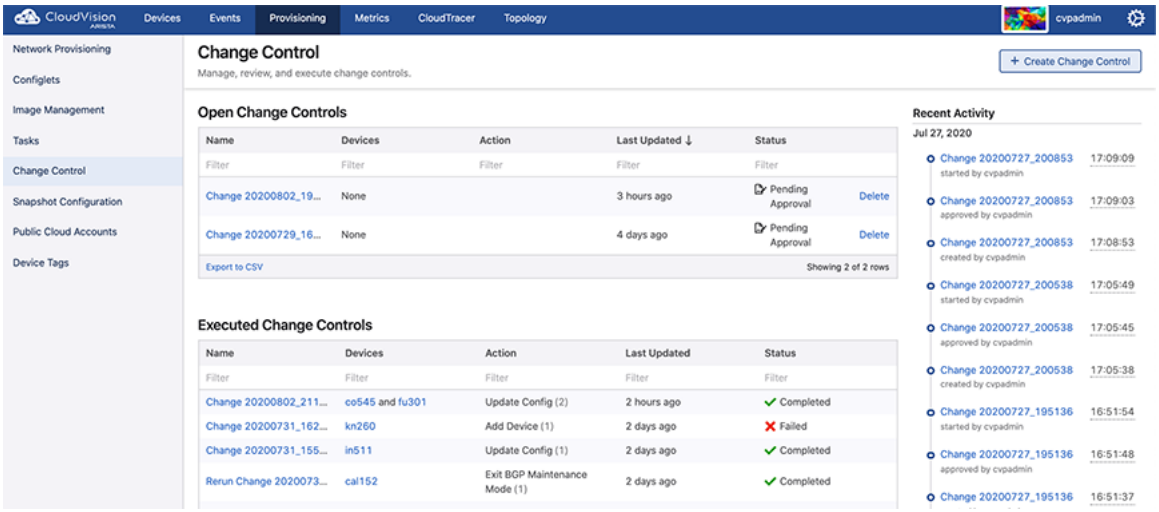

Figure 280: Change Control Summary Screen

2. Click **+ Create Change Control** button at the upper right corner.

   The system displays the **Assignable Tasks** dialog box.

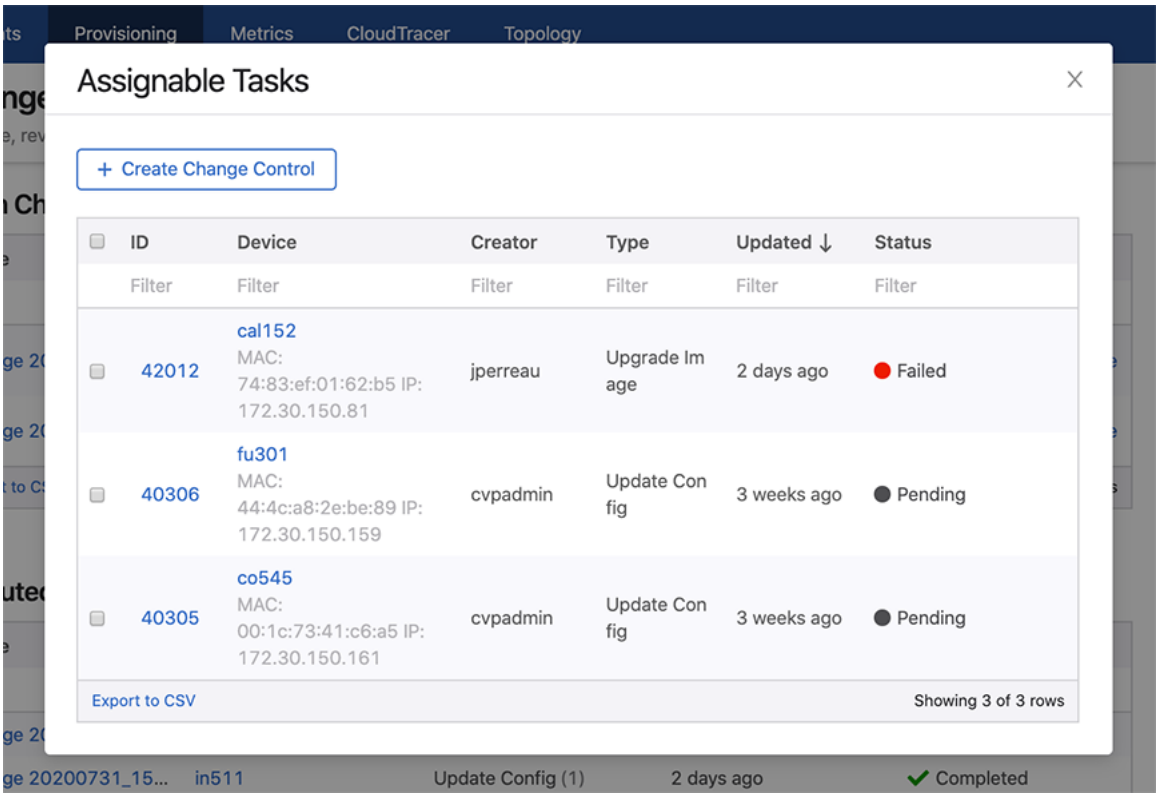

Figure 281: Assignable Tasks Dialog Box with No Tasks Selected

3. Select tasks you want to include in the Change Control by selecting appropriate checkboxes.

   Note: If you do not select any tasks, the system creates a Change Control without tasks.

4. Click **+ Create Change Control** with n tasks where n is the count of selected tasks.

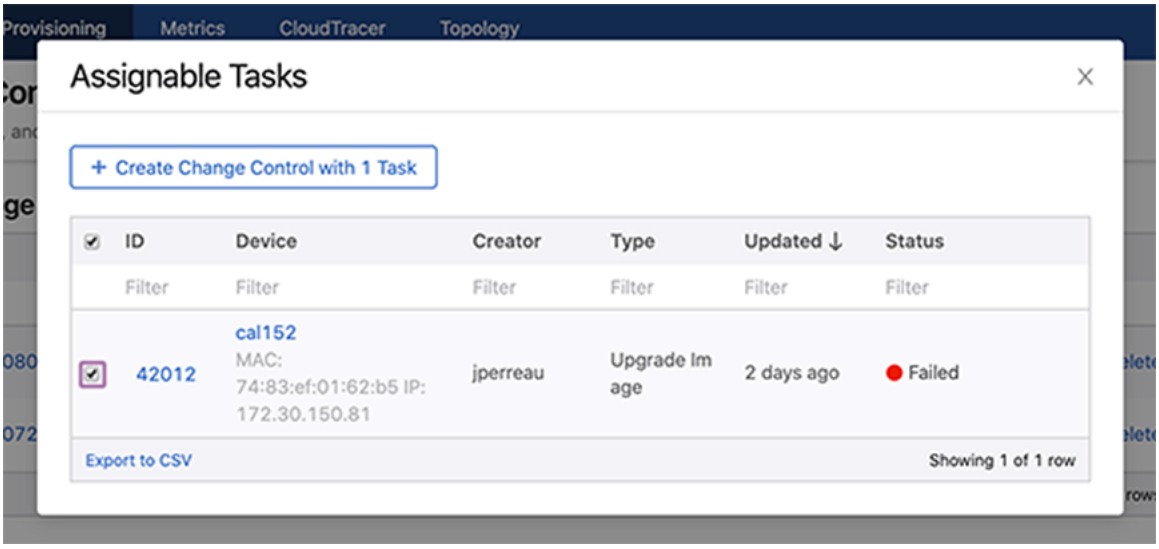

Figure 282: Assignable Tasks Dialog Box with Tasks Selected

The system displays the appropriate **Change Control Details** screen.

14.3.3 Accessing the Open Change Control Details Screen

The open Change Control details screen performs the following functions:

- Displays Change Control information
- Adds actions to Change Control
- Adds, edits, and deletes child stages
- Reviews and approves Change Control

Perform the following steps to access the Change Control details screen:

1. On the CloudVision Portal, click **Provisioning** > **Change Control**.

   The system displays the Change Control summary screen.
2. Under the **Open Change Controls** table, click one of the listed Change Controls.

   The system displays the Change Control details screen.

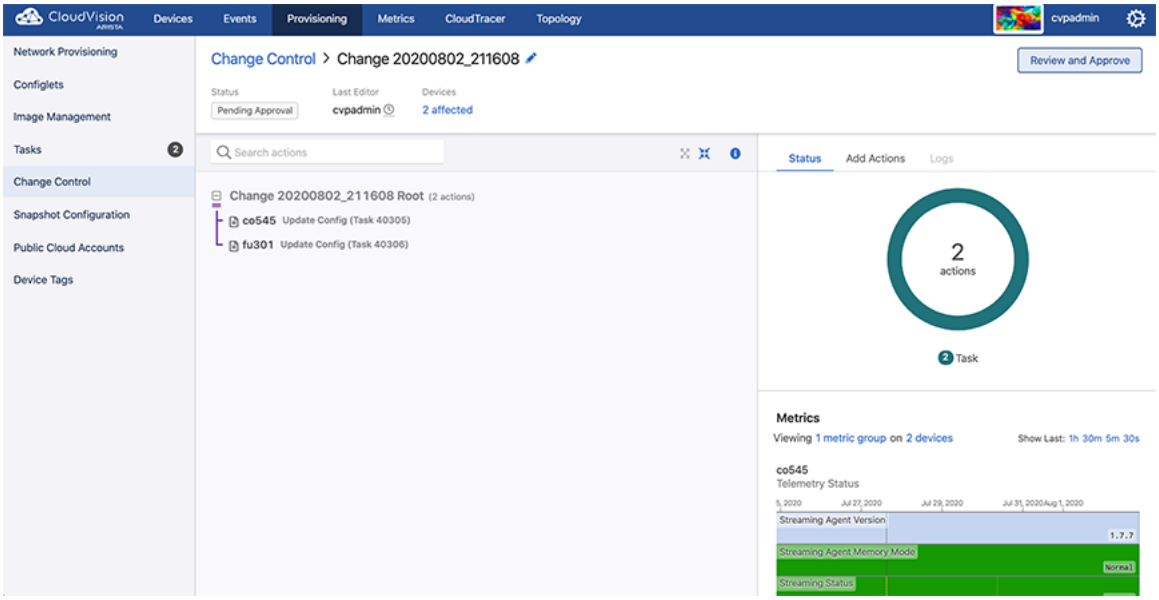

Figure 283: Change Control Details Screen

The Change Control details screen consists of the following panels:

- Header Panel
- Main Panel
- Edit Panel

Header Panel

This primary panel provides the following basic information on the Change Control:

- Edit icon to update the Change Control name
- Change Control information -
  - The open Change Control details screen displays the status, last editor, and count of affected devices.

    Note:

    - Hover the mouse cursor over the clock icon to view last time of action.
    - Hover the cursor on the count of affected devices to view their list. Clicking on an affected device opens the corresponding Device Overview screen.
  - The executed Change Control details screen displays the status, approver, time of start, last editor, and count of affected devices.

    Note:

- Click **Review** next to the status for details on review and approve process.

- **Review and Approve** - Click **Review and Approve** in open Change Controls for assessing Change Control updates. These updates include configuration differences, image bundle changes when appropriate, and commands that run as part of a CLI snapshot.

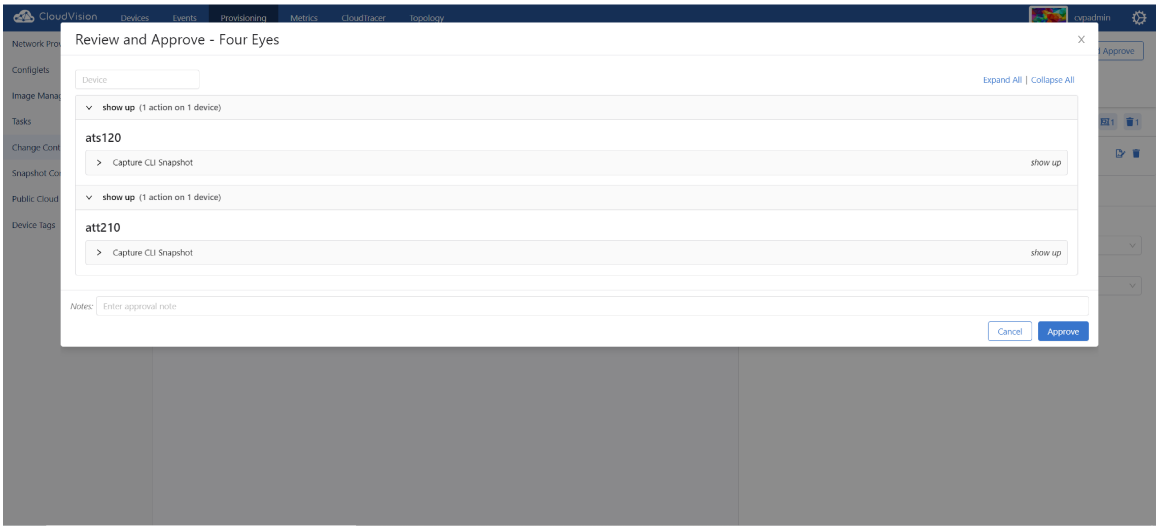

Figure 284: Review and Approve Pop-Up Window

Click **Approve** to accept Change Control updates.

Note: (Optional) Approver can leave comments in the ***Notes:*** field.

- On the approved Change Control details screen, click **Unapprove** to revert the approval status and **Execute Change Control** to run approved Change Controls.

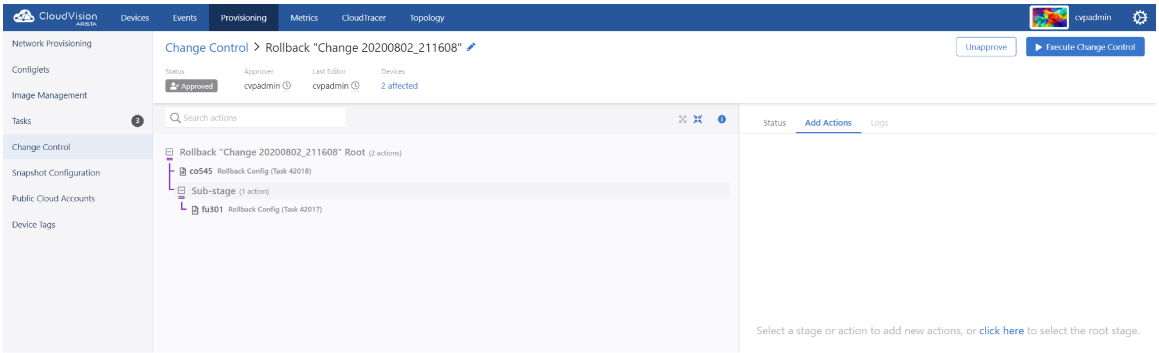

Figure 285: Approved Change Control

Note: CVP executes Change Controls in the following ways:

- Runs approved Change Controls immediately if sufficient privileges are set for the **Change Control Management** permission.
- Stops the change automatically if an action fails.
- Runs actions in progress until complete.

- On the failed Change Control details screen, click **Rerun** to repeat the execution of a completed but failed Change Control. This creates a new Change Control that must be approved again.

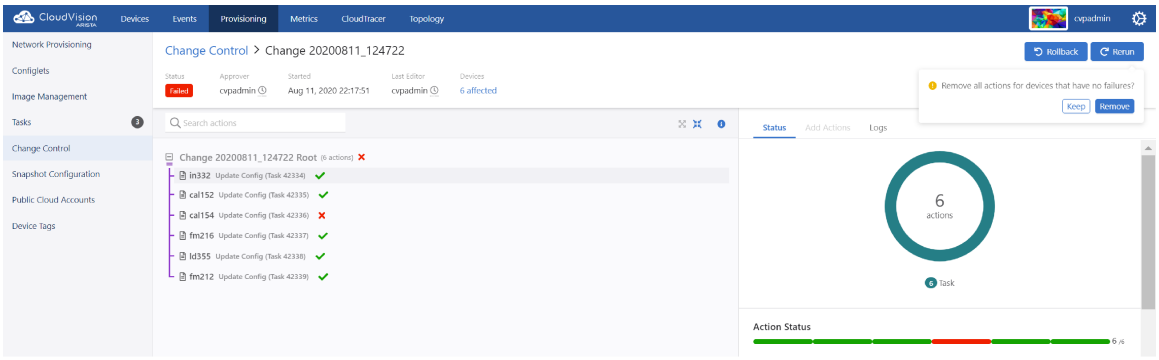

Figure 286: Rerun Change Control

Note: Click **Remove** when CVP prompts you with **Remove all actions for devices that have no failures?** for skipping the rerun of completed actions.

- Click **Rollback** in executed Change Controls to open the Rollback *Change Control* pop-up window. To create a rollback after evaluating the executed Change Control, select tasks to rollback from the table and click **Create Rollback Change Control**.

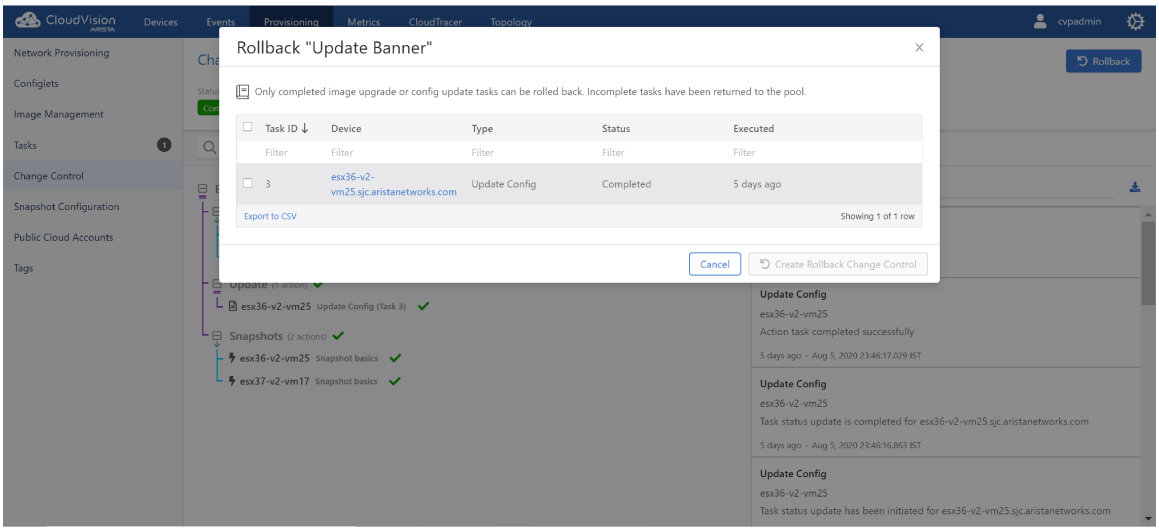

Figure 287: Rollback Pop-Up Window

Note: CVP rolls back only completed configuration updates and image upgrade tasks.

Main Panel

This main panel consists of the following entities:

- Search bar - Enter a string to perform a search in the Change Control tree.
- Expand icon - Click to expand all stages.
- Collapse icon - Click to collapse all stages.
- Information icon - Click to get help on Change Control.
- Change Control tree - Change Controls are composed of actions and stages. Action types include tasks, CLI snapshots, health checks, custom scripts, enter BGP maintenance mode, and exit BGP maintenance mode.

Note: Different icons represent various task types like adding a new device, updating configuration on a device, and updating software image bundle on a device. Actions are represented with a bolt symbol.

Actions are grouped and nested within stages via drag and drop. Each stage executes its children in series (represented with a down arrow) or parallel (represented with an equal sign).

Note:

- Tasks being executed in parallel do not block subsequent actions in that branch.
- In a series execution, the Change Control execution starts from the first item and works its way from top to bottom. The next action starts only when the previous action completed successfully.
- You can toggle the option by clicking the stage type dropdown menu in the edit panel.

Edit Panel

This panel edits stages and actions.

- Edit a stage - Click the required stage in the main panel. The edit panel provides the following options:
  - Show details icon - Click to view associated configuration differences, image bundle changes, and action details.
  - Remove icon - Click to delete the stage.

    Note: Select multiple tasks to view details and delete multiple tasks simultaneously. Use **command**-click or **Ctrl**-click to select multiple items. To select a range of items, click the first item and then **Shift**-click the last item.
  - Group icon - Select multiple tasks to group them into sub-stages.
  - Edit icon - Click to edit the stage name.
  - Change Control stage type dropdown menu - Click to select the Change Control stage type.

    Note: By default, all tasks and actions execute in parallel.
  - Plus icon - Click to add a child stage.
  - Status - Displays telemetry of each device in the stage.

    Note:

    - Hover the cursor on ***n* metric group** to view selected metric groups.

      Note: *n* represents the count of selected metric groups.
    - Hover the cursor on ***n* device(s)** to view selected metric groups.

      Note: *n* represents the count of selected devices.
  - Add actions - Adds actions to open Change Control. Select the required action and placement from corresponding dropdown menus; and click **Add to change control** to update selected changes.

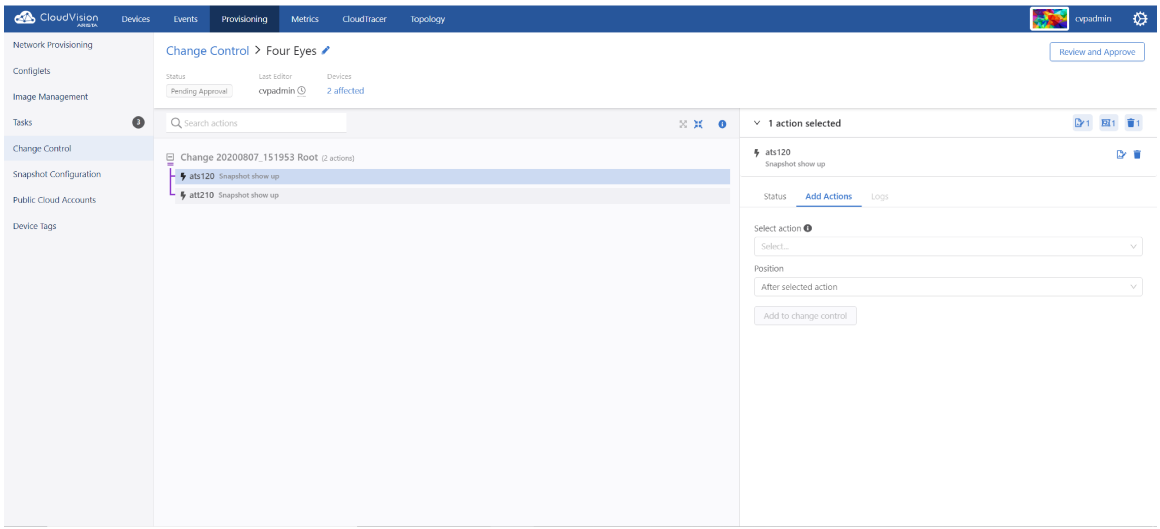

Figure 288: Add Actions to Change Control

- **Logs** - Displays logs of each update in the executed Change Control process.

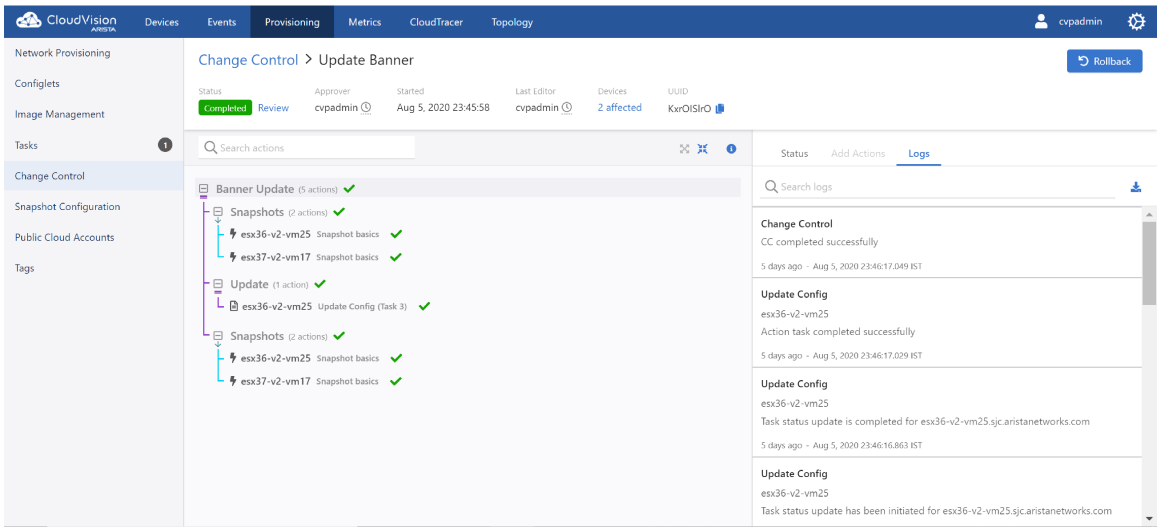

Figure 289: Change Control Logs

Note:

- Use the search logs bar for filtering logs based on a string.
- Click the download icon to download logs to your local drive.

14.3.3.1 Change Control Drop-Down Menu

Click the Change Control drop-down menu to select another Change Control.

14.3.3.2 Change Control Edit Drawer

The system provides collapsed and expanded views of the edit Change Control drawer.

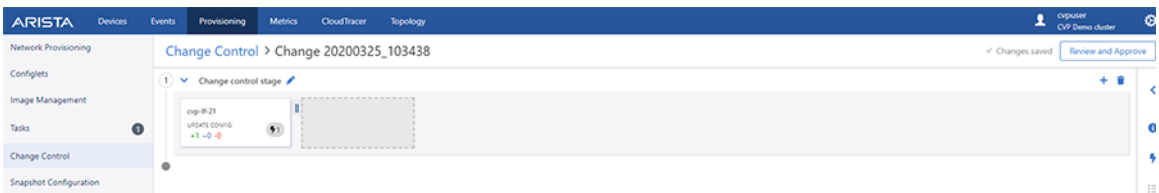

Figure 290: Collapsed View of the Edit Change Control Drawer

Each icon in the collapsed view corresponds to the appropriate drawer section. The chevron button expands the drawer, displaying the most recently used section. Click any of the active icons in the collapsed view to expand the Change Control drawer with the selected section.

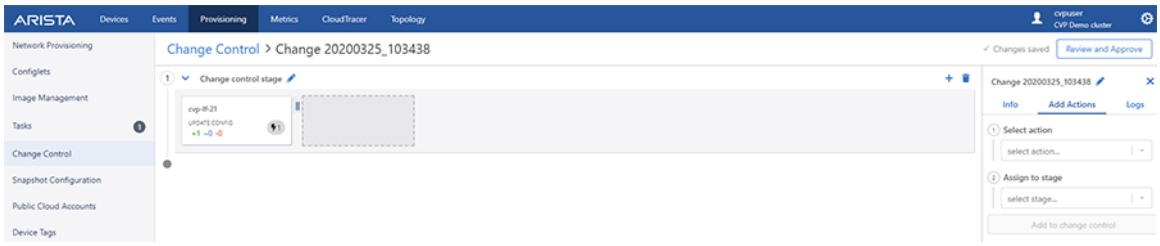

Figure 291: Expanded View of the Edit Change Control Drawer

The Change Control edit drawer consists of the following entities:

- Edit Change Control name - Click the Change Control name to edit the name.

  Note: Alternatively, click the edit icon next to Change Control name to edit the name.

- Info tab - Provides information of the current Change Control and displays the list of affected devices. Hover the mouse on any of the affected devices to view appropriate device details. .

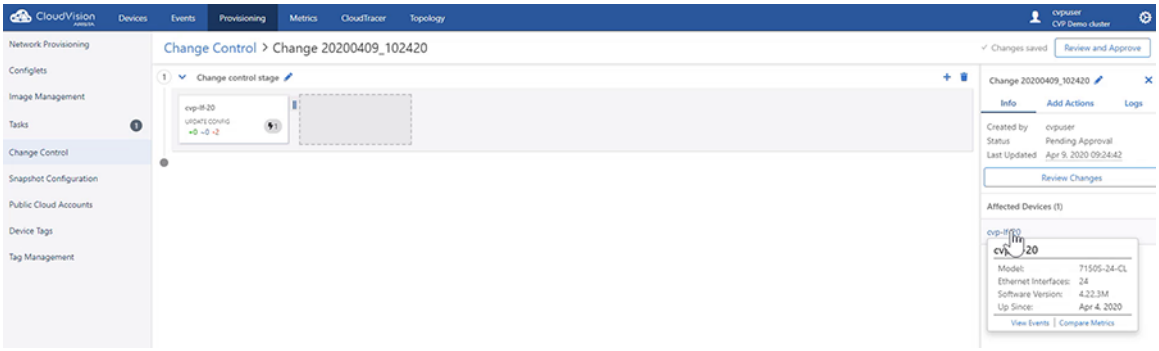

Figure 292: Affected Devices Popup in Info Tab

Click **View Events** to view events of the appropriate device. Click **Compare Metrics** to view metrics of the appropriate device. Click on any of the affected devices to view the appropriate device overview screen.

- **Add Actions** tab - Adds actions, assigns to a stage, and adds them to assigned stage.

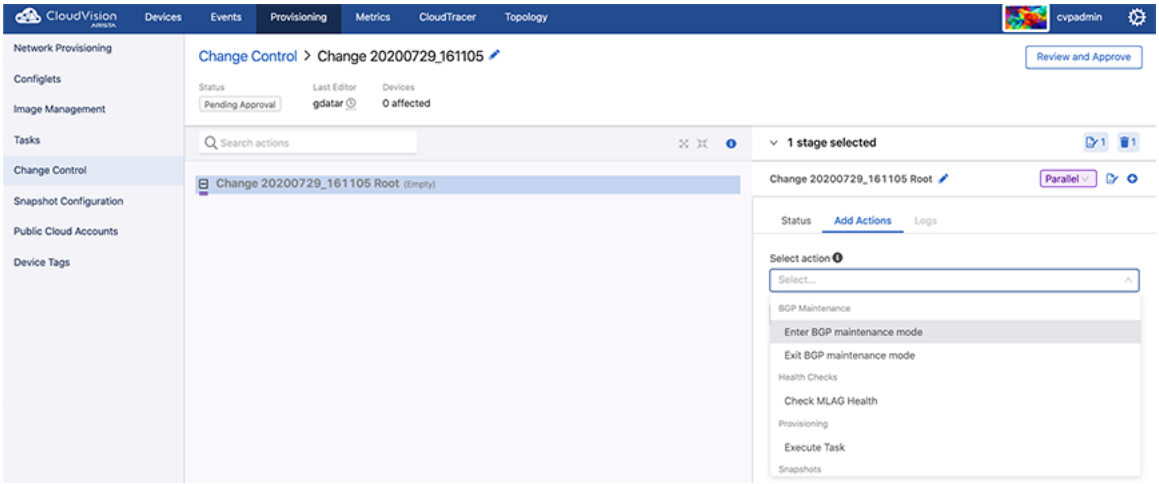

Figure 293: Add Actions Tab in Edit Change Control Pane

- **Logs** tab - Displays logs only when the Change Control in either running or has been executed.

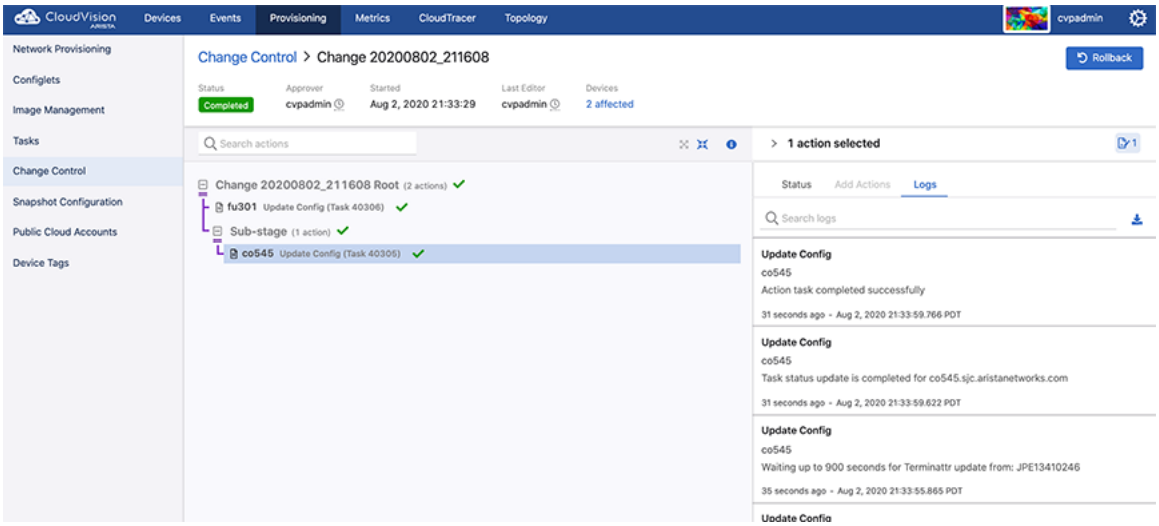

Figure 294: Logs Tab in Edit Change Control Pane

Note: This tab is available only for completed Change Controls.

14.3.3.3 Change Control Stages

These panes consists of the following entities:

- Change Control stage name - Click either the Change Control name or the corresponding edit icon to update the name.
- Add a stage icon - Click the plus icon at the upper right corner of the stage to add a stage.
- Delete a stage icon - Click the appropriate trash icon at the upper right corner of the stage to delete the corresponding stage.
- Edit actions icon - Click the thunder icon within a card to edit or view the appropriate leaf.
- For open Change Controls, the system displays the actions window to edit the appropriate leaf.

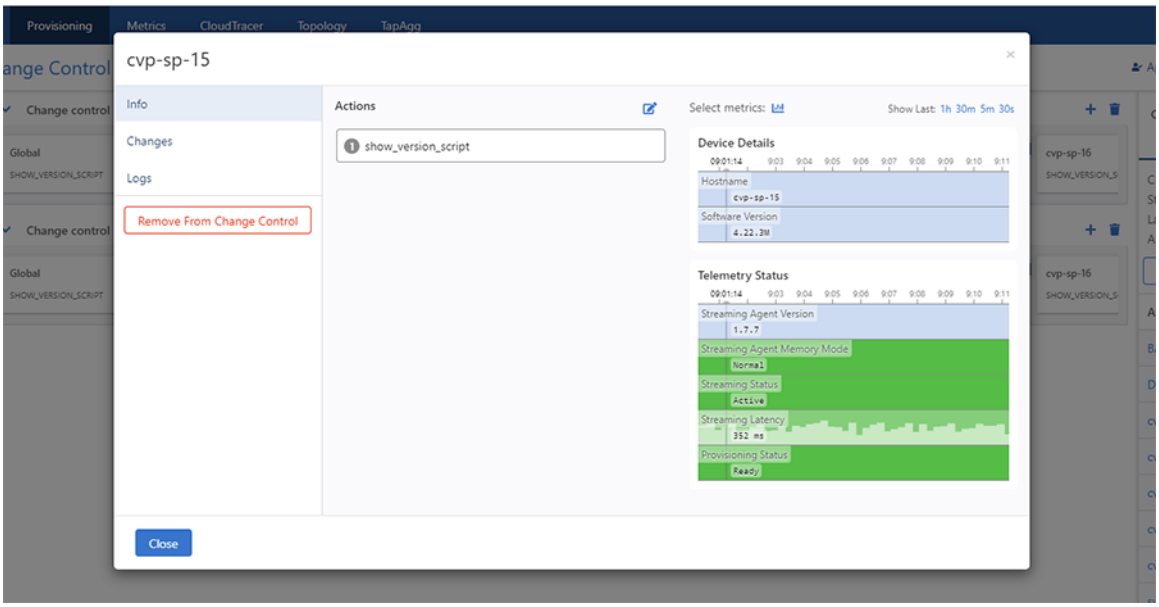

Figure 295: Info Tab in Edit Actions

Note: For completed Change Controls, the system displays the actions window to view the appropriate leaf.

This window consists of the following entities:

- **Info** tab - This tab lists the actions to be run, edits actions, and displays action details.

  Click the edit icon to reorder and edit actions.

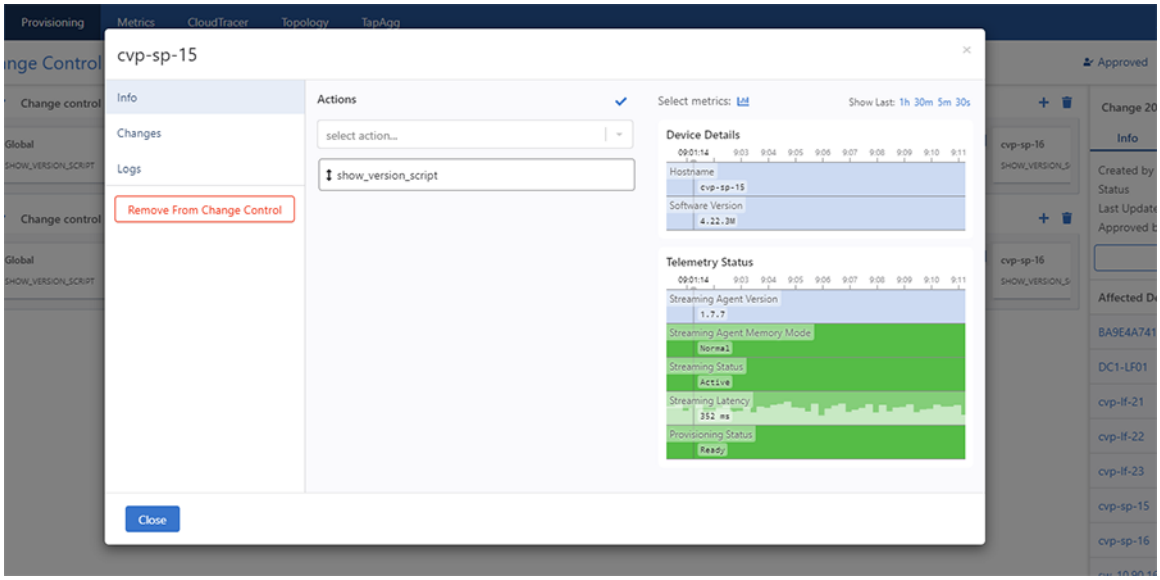

Figure 296: Reorder and Edit Actions Screen

  - Click the select action drop-down menu and select the required action.

    Note: The system displays selected actions beneath the select action drop-down menu.

  - Click **Clear** at the end of a field to delete the appropriate action.

    Note: This option is available only for a card with multiple actions. The main action in a card is not available to clear.

  - Click the check-mark to save changes.

    Note: Here, actions comprise of provisioning, Border Gateway Protocol (BGP) maintenance, health checks, and snapshots.

- **Configuration Changes** tab - For tasks, this tab displays any configuration or image differences that will be applied as part of the task.

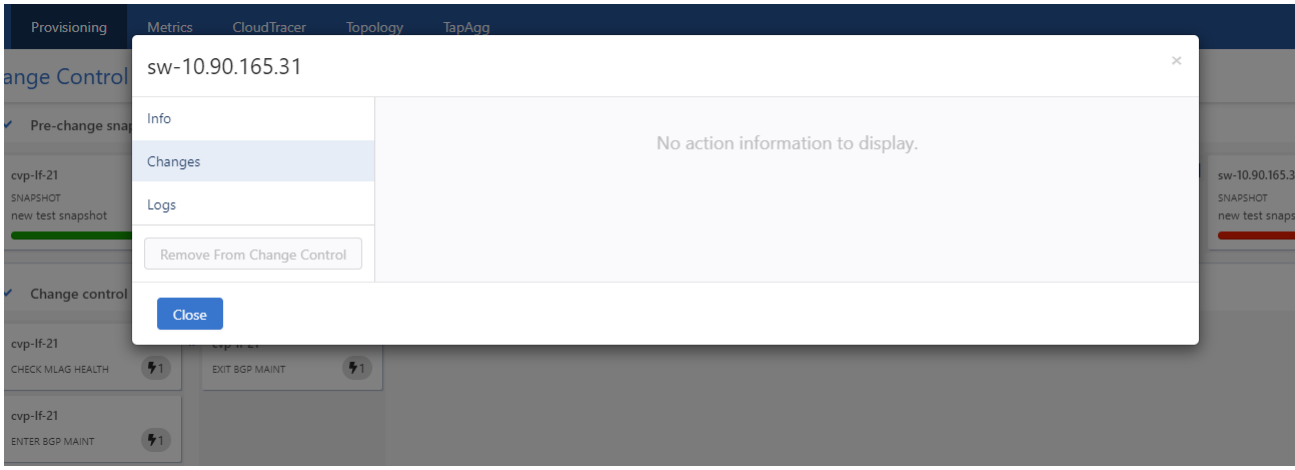

Figure 297: Configuration Changes Tab in Edit Actions

- **Logs** tab - This tab displays log information of completed Change Controls.

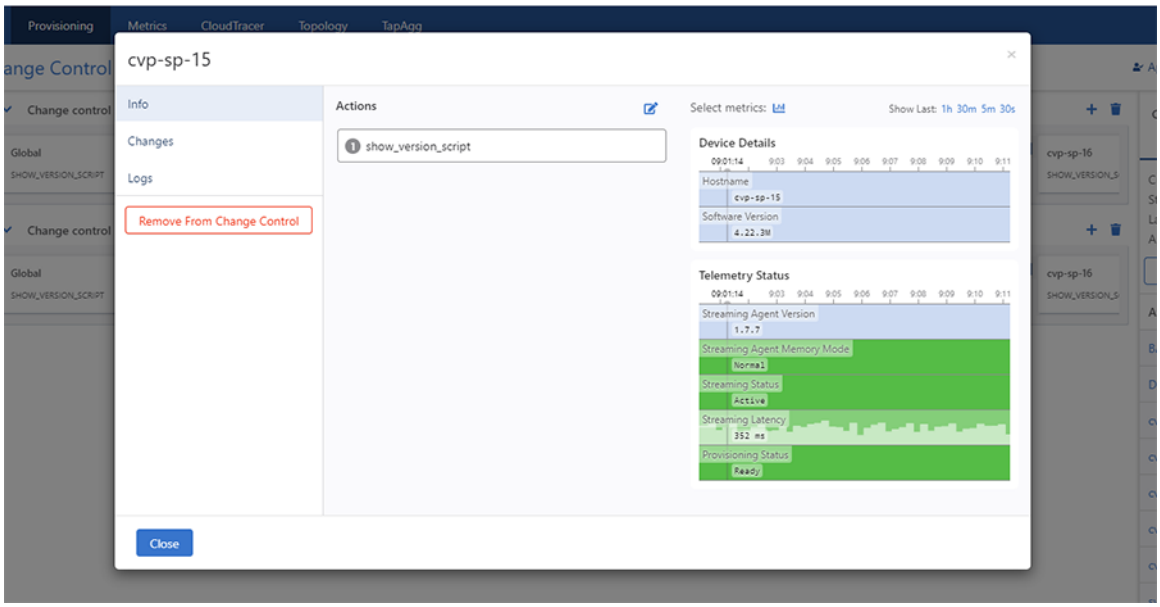

Figure 298: Logs Tab in Edit Actions

- **Remove from Change Control** button - Click Remove from Change Control to remove this task from the stage.

  Note: Click **Remove** on the **Confirm** pop-up dialog box to confirm the deletion.

  - **Done** button - Click **Done** to save changes.
- Trashbin icon - Click the trashbin icon at the upper right corner of the pane to delete the stage.

14.3.3.4 Review and Approve

Click the **Review** and **Approve** button at the upper right corner of the Change Control screen to review and approve the Change Control. This button displays the **Review and Approve** dialog box for the selected Change Control.

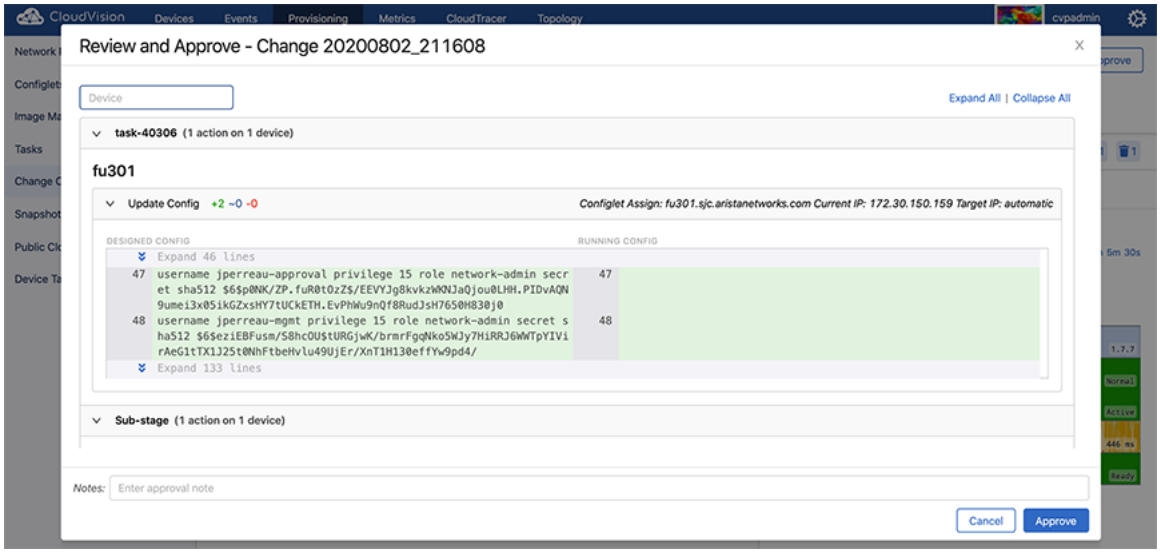

Figure 299: Review and Approve Dialog Box

This window consists of a device search field and a list of changes by Change Control stages.

Type the device name in the search field and if available, the system displays the list of changes for the specified device.

The expanded Change Control stage list displays details of the actions to be executed in each stage, grouped by a device.

If you are happy with configuration changes, click the **Approve** button at the lower right corner of the dialog box to approve the Change Control.

14.3.3.5 Execute Change Control

After approval, the **Review and Approve** button is replaced with the Execute Change Control button.

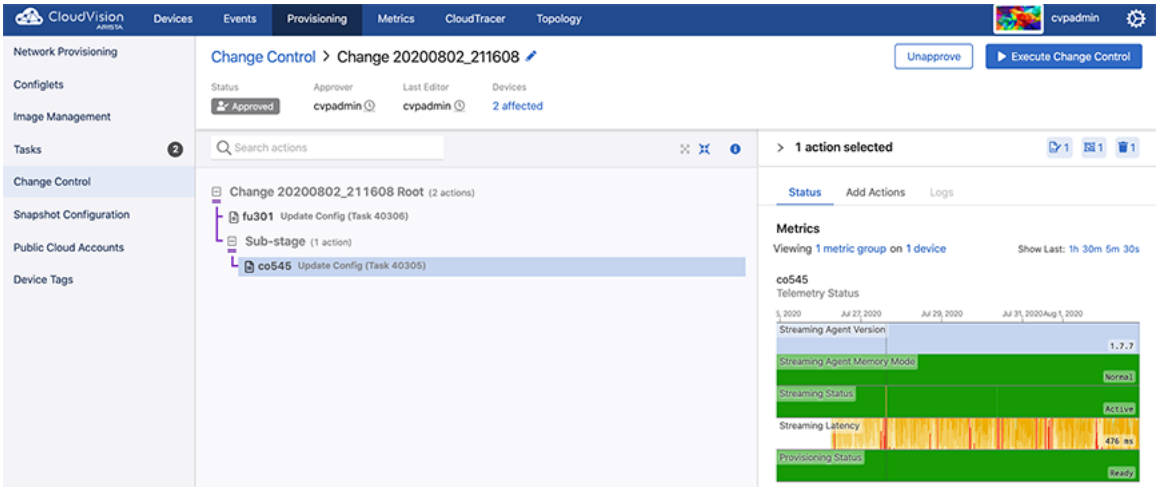

Figure 300: Execute Change Control Button

Click the **Execute Change Control** button to execute the Change Control.

Note: A Change Control is executed until all actions are either completed or there is a failure in one or more of the actions.

14.3.3.6 Stop Change Control

While the system is executing changes specified in Change Control, it replaces the **Execute Change Control** button with the **Stop Change Control** button.

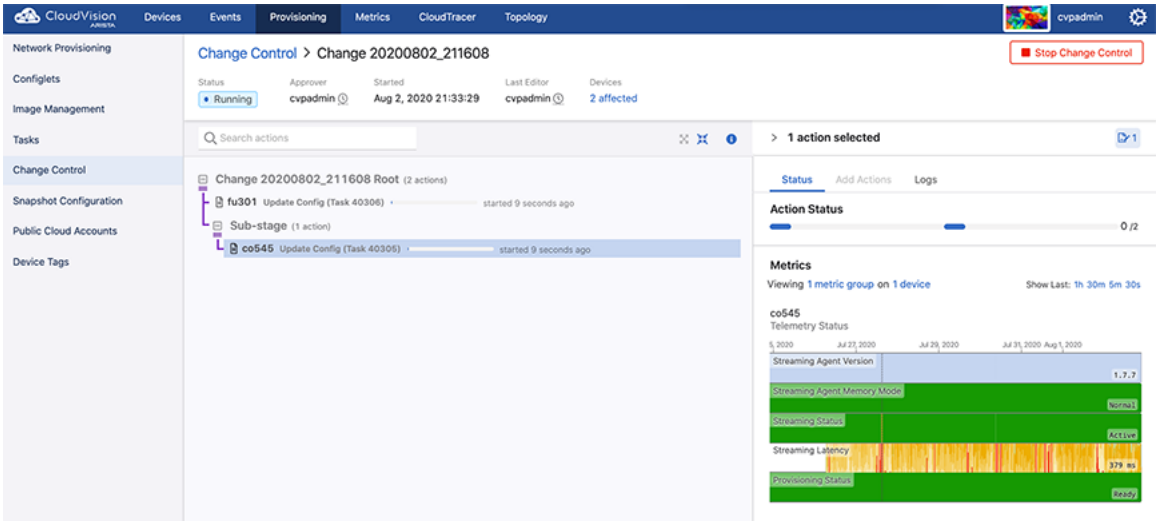

Figure 301: Stop Change Control Button

Click the **Stop Change Control** button to stop the execution of Change Control.

Note: Clicking the **Stop Change Control** button returns failed and incomplete tasks to the assignable tasks pool for reallocation.

If a Change Control has revertible actions, the system replaces the Stop Change Control button with the **Rollback Change** button after the execution of all actions.

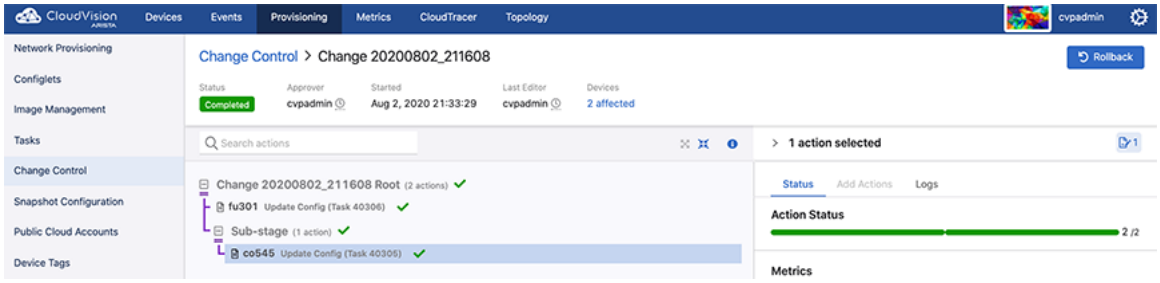

Figure 302: Rollback Change Button

Click the **Rollback Change** button to rollback the execution of Change Control.

Chapter 15

Authentication & Authorization (CVP)

Authentication determines if the provided user credentials (username/password) are correct. If authentication succeeds, the user is logged in.

Authorization determines what operations the user can perform after login. Authorization can be for no access, read access, or read and write access.

In the Access Control page, the type of Authentication and Authorization can be defined. AAA servers are defined in this page.

This module guides account management administrators to manage AAA servers, user accounts, and user roles. It provides the functionality required to manage all aspects of user accounts.

Note: Only account management administrators have the permissions to manage accounts.

Sections in this chapter include:

15.1 Access Requirements for Image Bundle Upgrades

If AAA is configured (enabled) on the switch, you must have certain access rights before you can perform image bundle upgrades on the switch.

The specific access rights required to perform image bundle upgrades when AAA is configured are:

- Config session
- Bash

The access rights to execute bash commands is required because the following bash command must be executed to upgrade image bundles:

```
bash timeout 10 sudo rm -f /mnt/flash/boot-extensions && echo -e '' > /
mnt/flash/boot-extensions
```

Note: If AAA is enabled and you attempt to perform image bundle upgrades without having these required access rights, the upgrade will fail and the following error occurs:

```
Jul 11 11:36:45 cd342 Aaa: %AAA-4-CMD_AUTHZ_FAILED: User cvpadmin
 failed authorization to execute command 'bash timeout 10 sudo rm -f
```

```
/mnt/flash/boot-extensions && echo -e '' > /mnt/flash/boot-
extensions
```

Related topics:

- Access to the Access Control Page
- Modifying AAA Servers

15.2 Managing AAA Servers

The system uses the following functionalities to manage AAA servers:

- Adding AAA Servers
- Modifying AAA Servers
- Removing AAA Servers

15.2.1 Adding AAA Servers

1. Navigate to the **Access Control** Page.
2. Click the Authentication source drop-down menu and select either RADIUS or TACACS.

   The Access Control page lists all current servers. See Access to the Access Control Page.
3. Click **+ New Server** at the upper right corner of the **Servers** section.

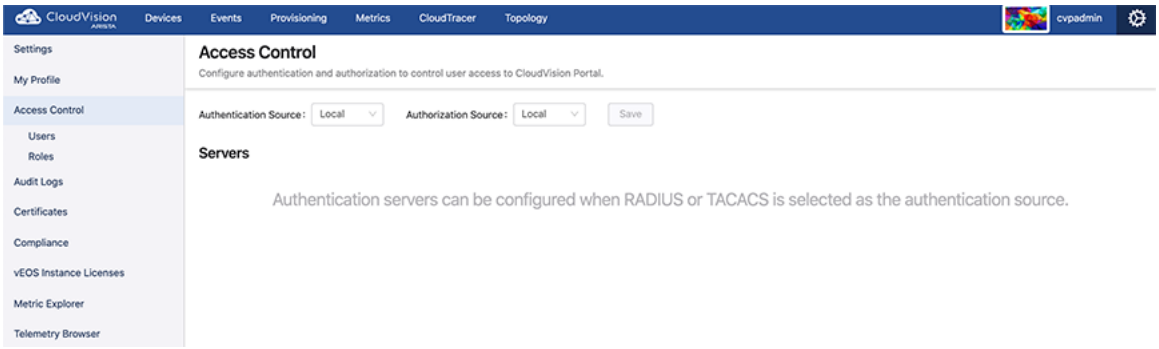

Figure 303: + New Server in Access Control Page

The system pops-up the New Server window.

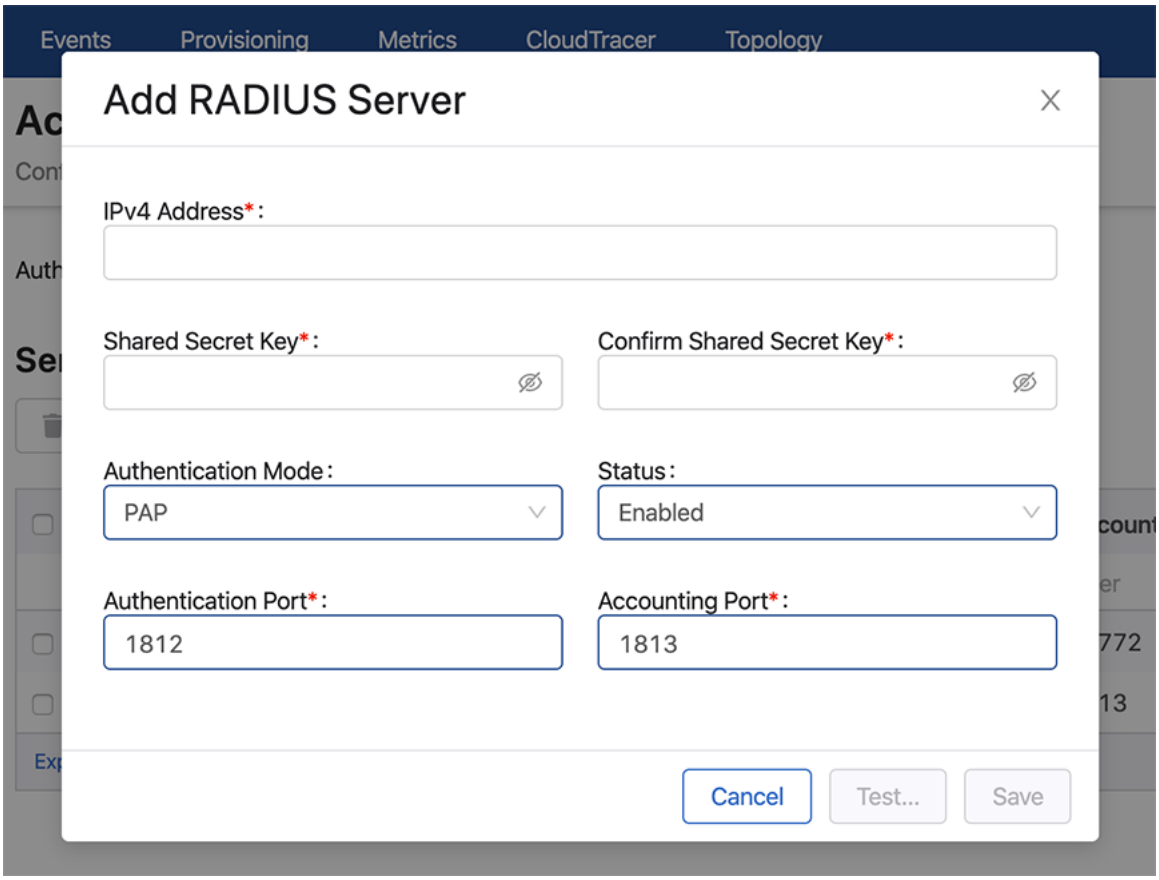

Figure 304: New Server Pop-Up Window

4. Provide the required Information in corresponding fields.
5. If required, click **Test** for testing the new configuration. Else, skip to step 8.
6. Enter your credentials when the **Test Server** pop-up prompts for it.

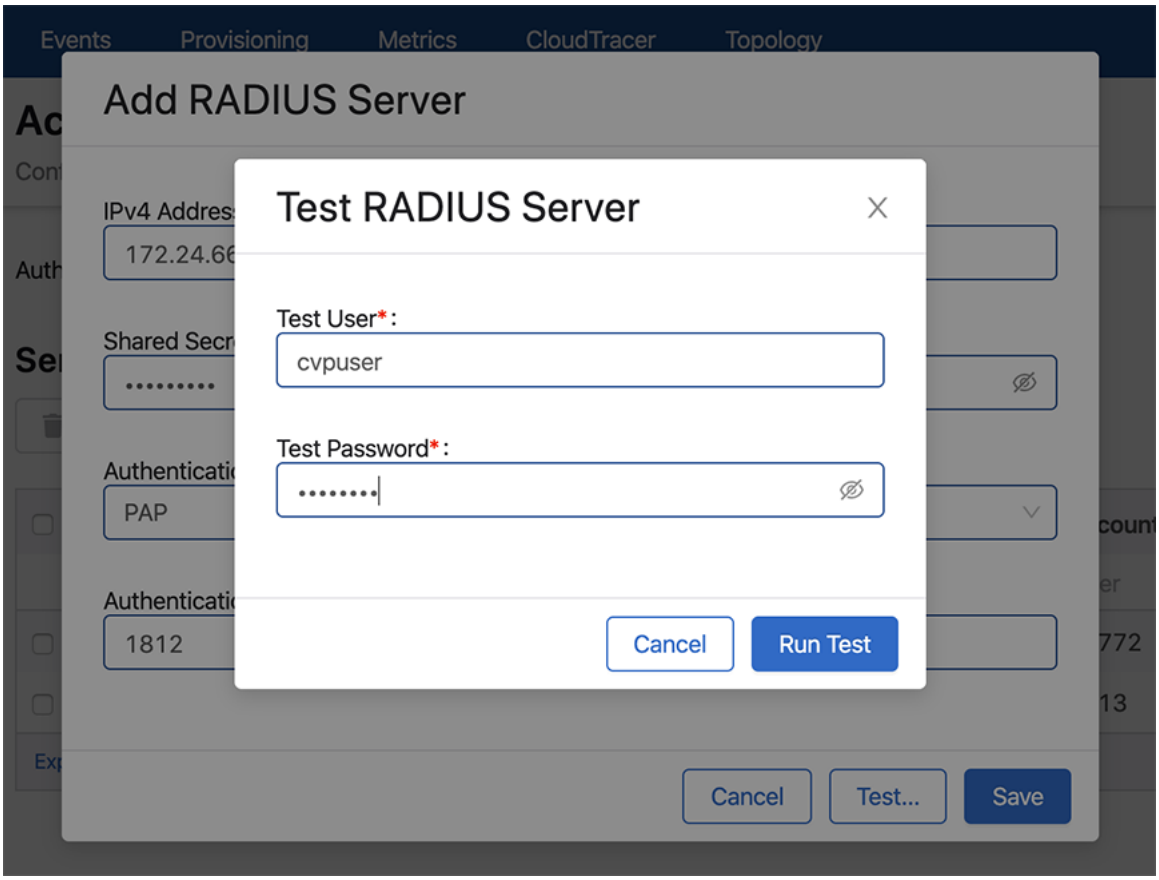

Figure 305: Test Server Pop-Up Window

7. Click **Run Test**.

   The system displays test results. If required, modify the configuration based on the test result.
8. Click **Save**.

   The server is added to the list of servers in the AAA grid.

   Related topics:

   - Access to the Access Control Page
   - Modifying AAA Servers
   - Removing AAA Servers

15.2.2 Modifying AAA Servers

1. Navigate to the **Access Control** Page.
2. Select desired modes from **Authentication source** and **Authorization source** drop-down menus

   The system lists all registered servers of the selected AAA server type. See Figure 320: AAA Access Control Page.
3. Click the edit icon available next to IP address of the corresponding server.

   The system pops-up the Edit Server window.

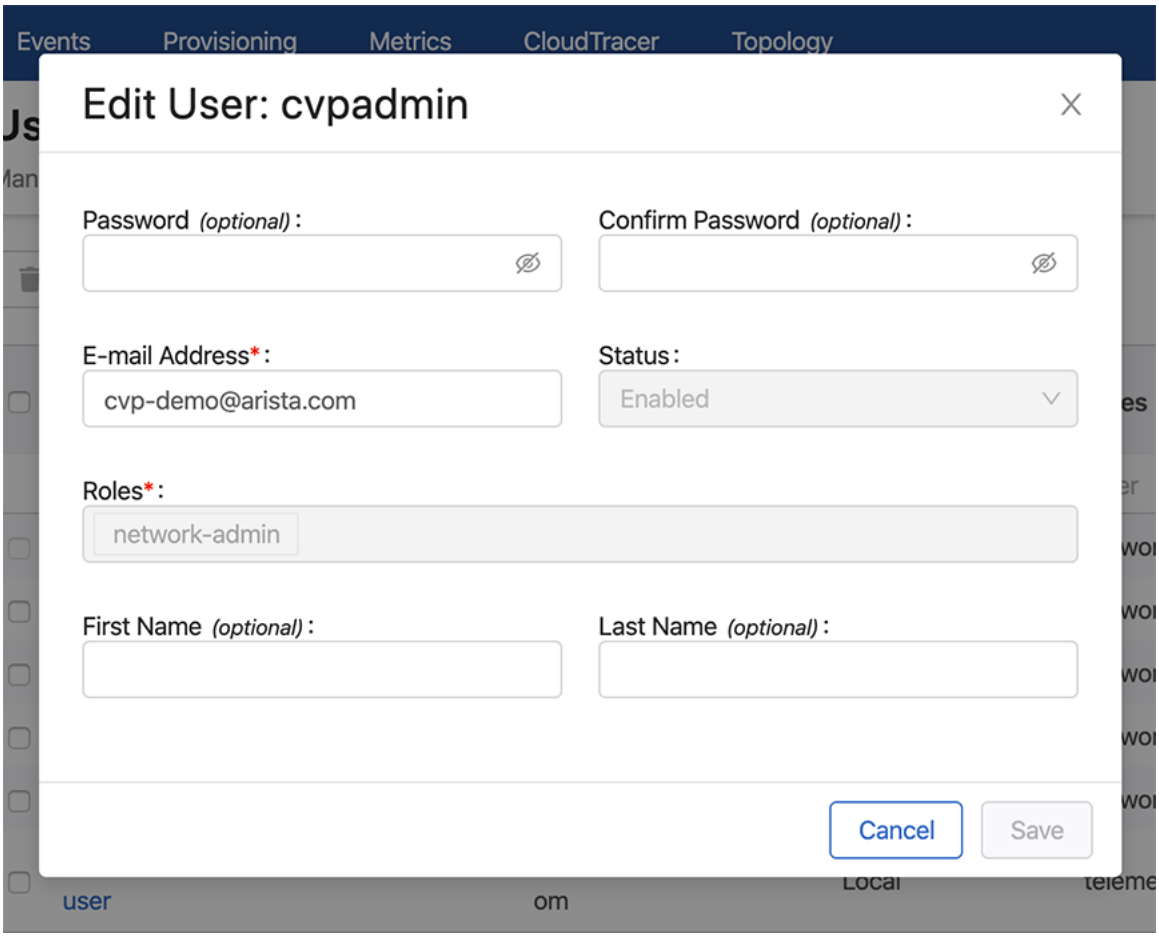

Figure 306: Edit Server Pop-Up Window

4. Modify the required information.
5. If required, click **Test** to verify latest changes.
6. Click **Save**.

Note: To apply external authentication, there should be at least one enabled server listed in the page.

15.2.2.1 Adding Vendor Specific Codes to AAA Servers

You can add vendor specific codes to AAA servers for the following:

- RADIUS
- TACACS+
- CISCO ACS

15.2.2.1.1 RADIUS

Arista Vendor Specific Code: add it to the RADIUS dictionary.

```
VENDOR Arista 30065
BEGIN-VENDOR Arista
ATTRIBUTE Arista-AVPair 1 string
END-VENDOR Arista
```

To specify role for a user

```
"bob"     Cleartext-Password := "Pa$sW04d"
```

```
                    Arista-AVPair = "shell:cvp-roles=network-admin",
                    Service-Type = NAS-Prompt-User
```

15.2.2.1.2 TACACS+

For TACACS+ there is no vendor specific code, just different strings.

Note: CloudVision support for TACACS+ servers can be affected with the setting of the “service” parameter. Some TACACS servers may require "service = shell" instead of "service = exec" in the TACACS+ configuration (*tacacs.conf*).

This example configures user “bob” in the admin group and specifies certain attributes. It specifies a "cvp-roles" attribute for the CloudVision role name (it can also be a list of roles).

```
A. tacacs.conf
group = admingroup {
   default service = deny
   service = exec {
      default attribute = permit
      priv-lvl = 15
      cvp-roles = network-admin
   }
 enable = nopassword

}
user = bob {
   login = cleartext "secret"
 member = admingroup
}
B. CVP AAA settings
C. Switch AAA configlet
```

15.2.2.1.3 CISCO ACS

To ensure that authentication and authorization work properly, complete the following procedures.

- Creating Identity Groups and Users
- Creating a Shell Profile using ACS
- Creating and Mofiying Access Policy

Creating Identity Groups and Users

1. Select **Users and Identity Stores**, and then select **Identity Groups**.
2. Make sure a group named <*user-group*> exists. If this group does not exist, add it.
3. Add new users under the group named <*user-group*>.

15.2.2.1.4 Supported TACACS Types

CloudVision Portal (CVP) supports different types of TACACS. Table **Supported TACACS Types** lists the supported types of TACACS, including the following information for each TACACS type:

- Supported version
- Service shell (whether it is supported for each type)
- Service exec (only the following attributes are supported):
  - acl
  - default
  - double-quote-values
  - message
  - optional
  - protocol

- return
- script
- set

Table 14: Supported TACACS Types

TACACS Type	Supported Version	Service Shell	Service Exec
tac_plus (Shruberry)	F4.0.4.26	Not Applicable	Supported
tac_plus (Probono)	201706241310 201503290942/DES	Supported	Supported
CISCO ACS	4.4.0.46 5.3.0.40	Supported	Not Applicable

Related topics:

- Access to the Access Control Page
- Adding AAA Servers
- Removing AAA Servers

15.2.3 Removing AAA Servers

Complete these steps to remove AAA servers:

1. Navigate to the **Access Control** page.
2. Select required options from **Authentication source** and **Authorization source** drop-down menus.

   The systems lists all current servers.
3. Select required servers for removal.
4. Click **Remove Server(s)** at the upper right corner of the **Servers** section.

   The systems lists all current servers.

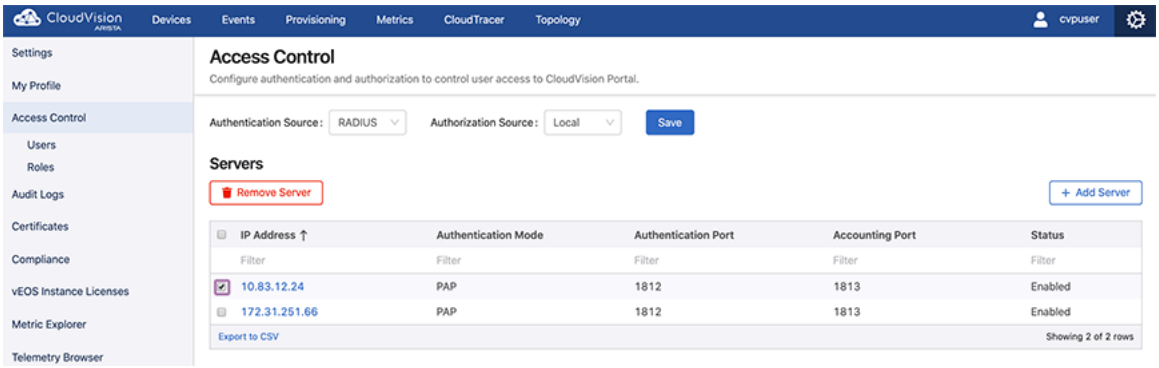

Figure 307: Remove AAA Servers

5. Click **Delete**.

   The system deletes selected AAA servers.

Related Topics:

- Access to the Access Control Page
- Adding AAA Servers
- Modifying AAA Servers

15.3 About Users and Roles

Account management is based on users and roles. In the CloudVision Portal, users and roles have specific meaning.

Users	A user is a person who uses the CVP application and is authenticated by the system through the use of account credentials (username and password). which is maintained by CVP or external enterprise servers. Only the users with account management module credentials (Account management administrator) can create and manage users. The account management administrator specifies the authentication credentials, name and contact information, status, and CVP permissions when creating user accounts for new users. Account management administrators control which CVP modules users are authorized to use by assigning roles to users (the role assignments can be changed as needed at any time). **Note:** Activity of CVP users is logged and can be viewed in the Audit Logs page.
Roles	A role is a set of read and write module permissions that defines user authorization to modules in CloudVision Portal. The account management administrator specifies the read and write permissions of each module when they create roles. Only account management administrators can create and manage roles. Roles enable account management administrators to efficiently manage user permissions by assigning roles to users, and by changing the role assigned to users. CloudVision Portal provides two default roles, one for the system administrator (network-admin) and one for a basic operator (network-operator).

15.4 Managing User Accounts

The system uses the following functionalities to manage user accounts:

- Adding New User Accounts
- Modifying User Accounts
- Removing User Accounts

15.4.1 Adding New User Accounts

When you create a new user account, you specify the login information (authentication credentials) of a person that needs to use one or more CVP modules. Personal information for the new user account is optional and can be specified when you create the new user or at a later time.

By default, new user accounts are enabled. The new user is able to use the CVP modules they are permitted to use, based on the role assigned to them. If you do not want the new user to use CVP at this time, select the Disable option (a Status option). You can enable the user account at a later time.

Note: As an alternative to creating user accounts in CVP, you can point CVP to an external AAA server that automatically creates users and maps them to roles during first login.

Complete these steps to create a new user:

1. Navigate to the **Access Control** page.
2. Under **Access Control** in the left menu, click **Users**.

   The Users page lists all current users.

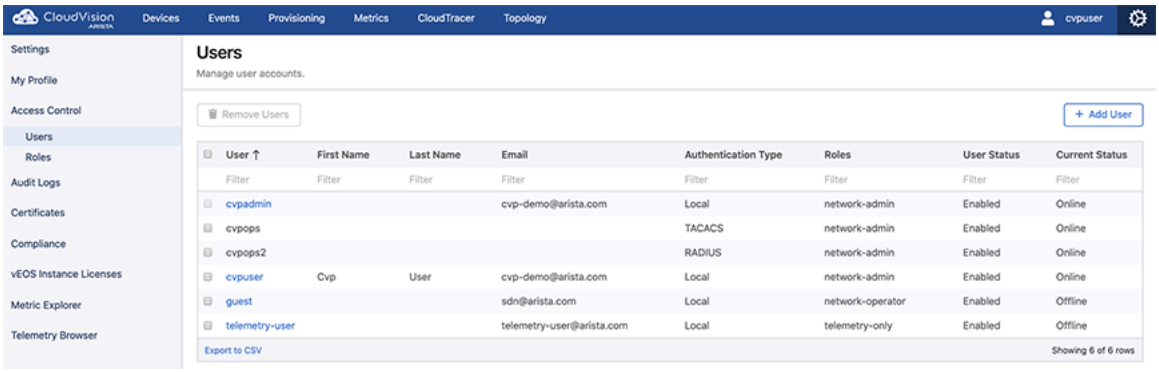

Figure 308: Users Page

3. Click **+ New User** at the upper right corner of the Users page.

   The system pops-up the **New User** window.

   Note: The **New User** pop-up window creates users only with the 'Local' authentication type.

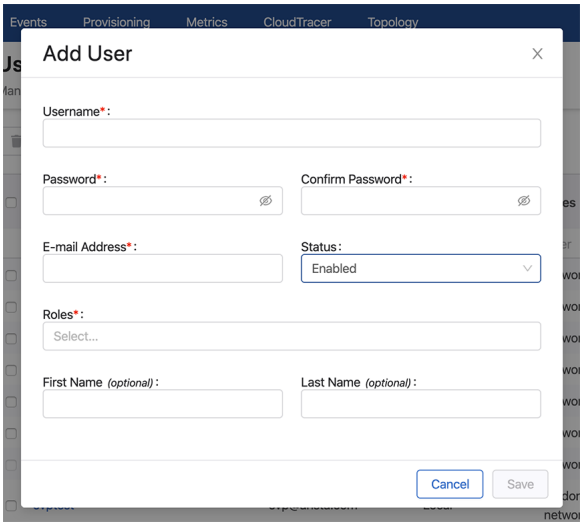

Figure 309: New User Pop-Up Window

4. Provide the required information in corresponding fields.
5. Click **Save**.

   The new user account is created.

Note: If the specified role is unavailable in the local CVP, then the network-operator role is automatically assigned to either the RADIUS or TACACS user. Unless you set the account status to disabled, the new user is active using CVP modules based on the role assigned to the user. If user roles conflict when multiple roles are assigned to a user account, the user role with higher privileges is applied to the user account.

Related topics:

- Modifying User Accounts
- Removing User Accounts
- Viewing Activity Logs

15.4.2 Modifying User Accounts

Modifying user accounts enables you to change the following aspects of existing user accounts:

- Login information (password)
- Contact information (email address)
- Status (enabled or disabled)
- Role(s) (the CVP role(s) assigned to the user)
- Personal information (first and last names)

Note: Once changes are saved, they are implemented immediately.

Complete these steps to modify a user account.

1. Navigate to the Access Control page.
2. Under **Access Control**, click **Users**.
3. In the **Users** page, click the edit icon available next to the corresponding user name.

   The system pops-up the **Edit User** window displaying all information related to the corresponding user.

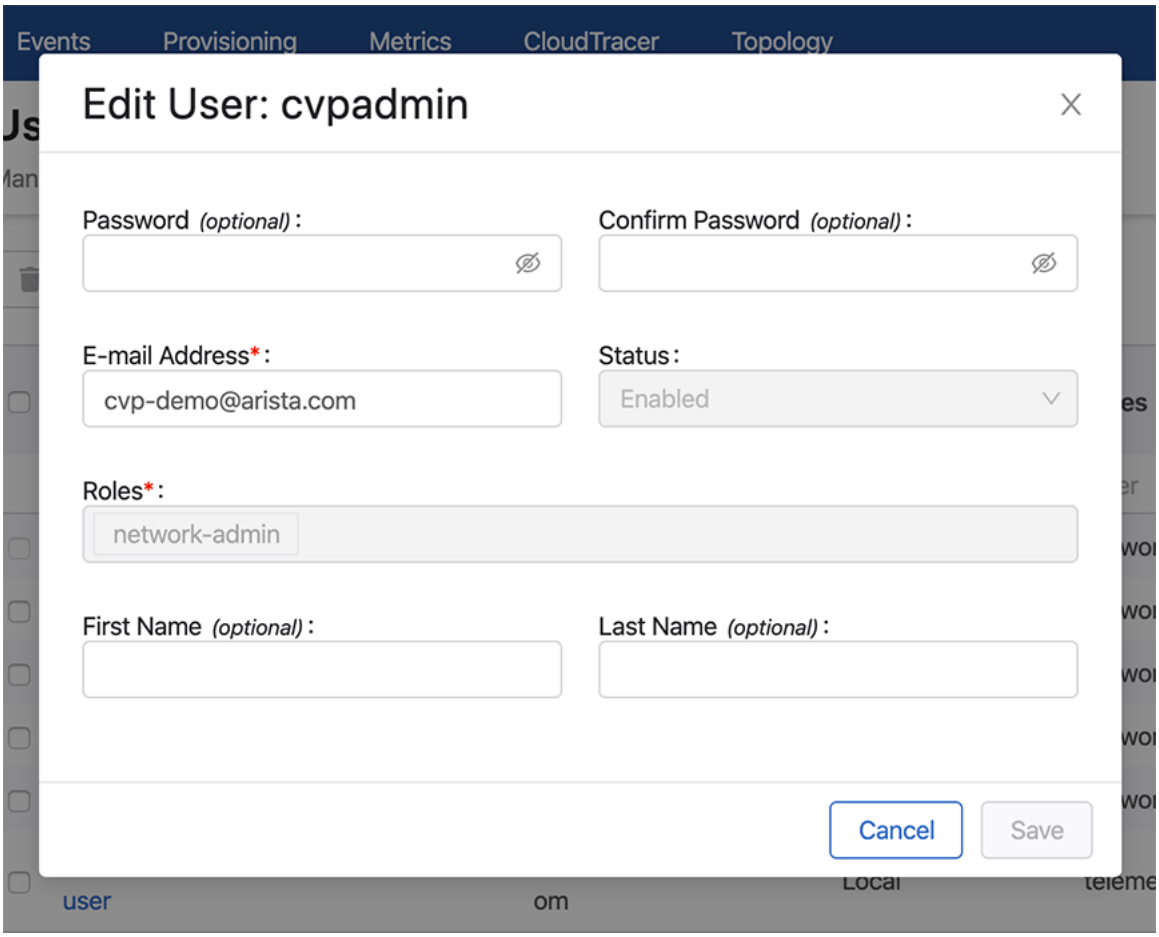

Figure 310: Edit User Pop-Up Window

4. Modify the required information.
5. Click **Save**.

Related Topics:

- Adding New User Accounts
- Removing User Accounts
- Viewing Activity Logs

15.4.3 Removing User Accounts

Complete these steps to remove a user account:

1. Navigate to the **Access Control** page.
2. Under **Access Control** in the left, click **Users**.

   The **Users** page appears displays all current user accounts.
3. Select the users for removal.
4. Click **Remove User/Remove Users** at the upper right corner of the Users page.

   The system prompts to confirm deletion.

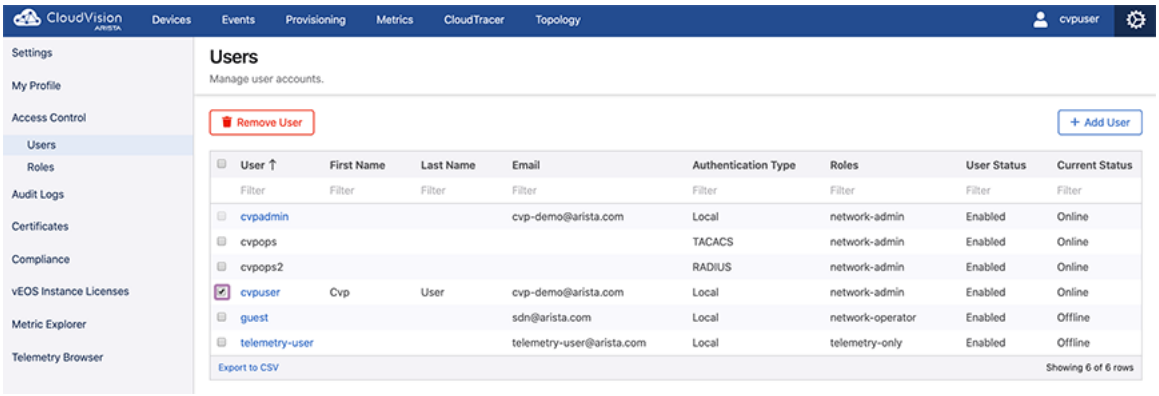

Figure 311: Remove User Account

5. Click **Delete.**

   The system deletes selected user accounts.

   Related Topics:

   - Adding New User Accounts
   - Modifying User Accounts
   - Viewing Activity Logs

15.5 Managing User Roles

The system uses the following functionalities to manage user roles:

- Adding New User Roles
- Modifying User Roles
- Removing User Roles

15.5.1 Adding New User Roles

CloudVision Portal enables you to create new roles as needed to ensure that you are able to efficiently manage CVP user permissions. When you create a new role, you specify the read and write permissions for each CVP module.

Once a role has been created, it is automatically added to the list of Available roles, and you can assign it to users that should have the permissions defined in the role. When you assign the role to a user, they inherit the read and write permissions defined in the role.

Complete the following steps to create new roles:

1. Navigate to the **Access Control** page.
2. Under **Access Control** in the left menu, click **Roles**.

   The Roles page lists all current roles.

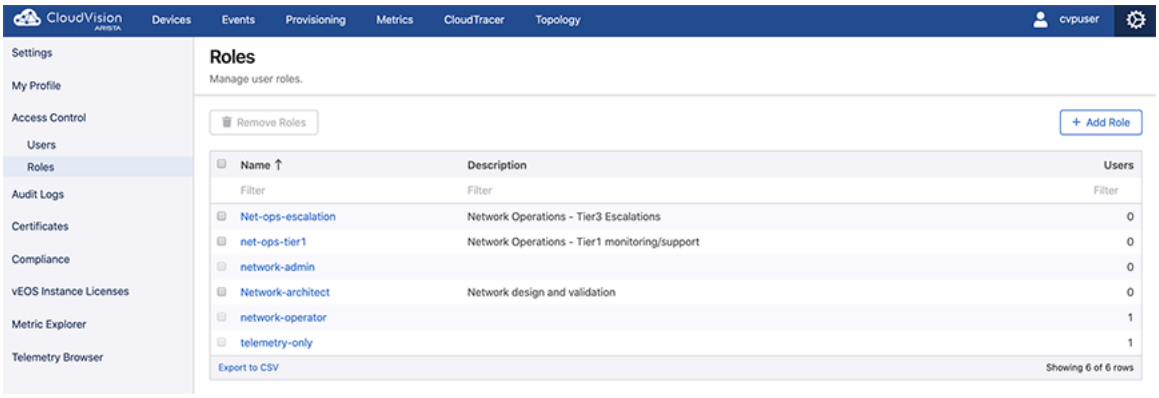

Figure 312: Roles Page

3. Click **+ New Role** at the upper right corner of the Roles page.

   The system pops-up the New Role window.

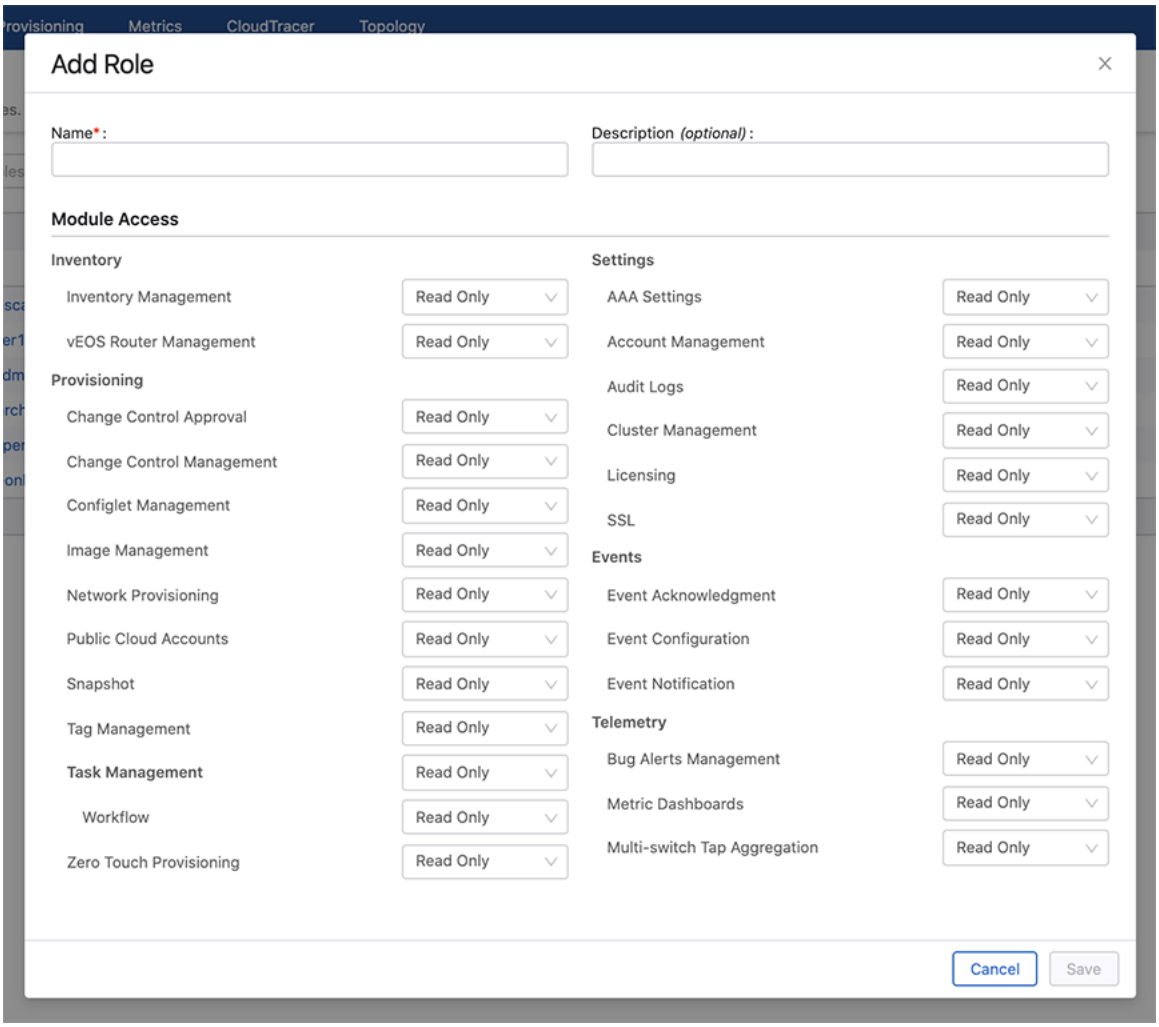

Figure 313: New Role Pop-Up Window

4. Provide the required information in corresponding fields.
5. Click **Save**.

   The new role is saved to the CVP database and is available to be assigned to users.

Note: The roles created can be assigned to locally created users or by the external AAA server to its known users.

Related topics:

- Adding New User Roles
- Modifying User Roles
- Viewing Activity Logs

15.5.2 Modifying User Roles

CloudVision Portal provides the functionality required to change the permissions of an existing role. This enables you to efficiently change the permissions of all users that are assigned the role. After you modify the role, all users assigned the role inherit the read and write permissions defined in the new version of the role.

Complete the following steps to modify an existing role:

1. Navigate to the **Access Control** page.
2. Under in the left menu, click **Roles**.
3. In the **Roles** page, click the edit icon available next to the corresponding role name.

   The system pops-up the **Edit Role** window displaying all information related to the corresponding role.

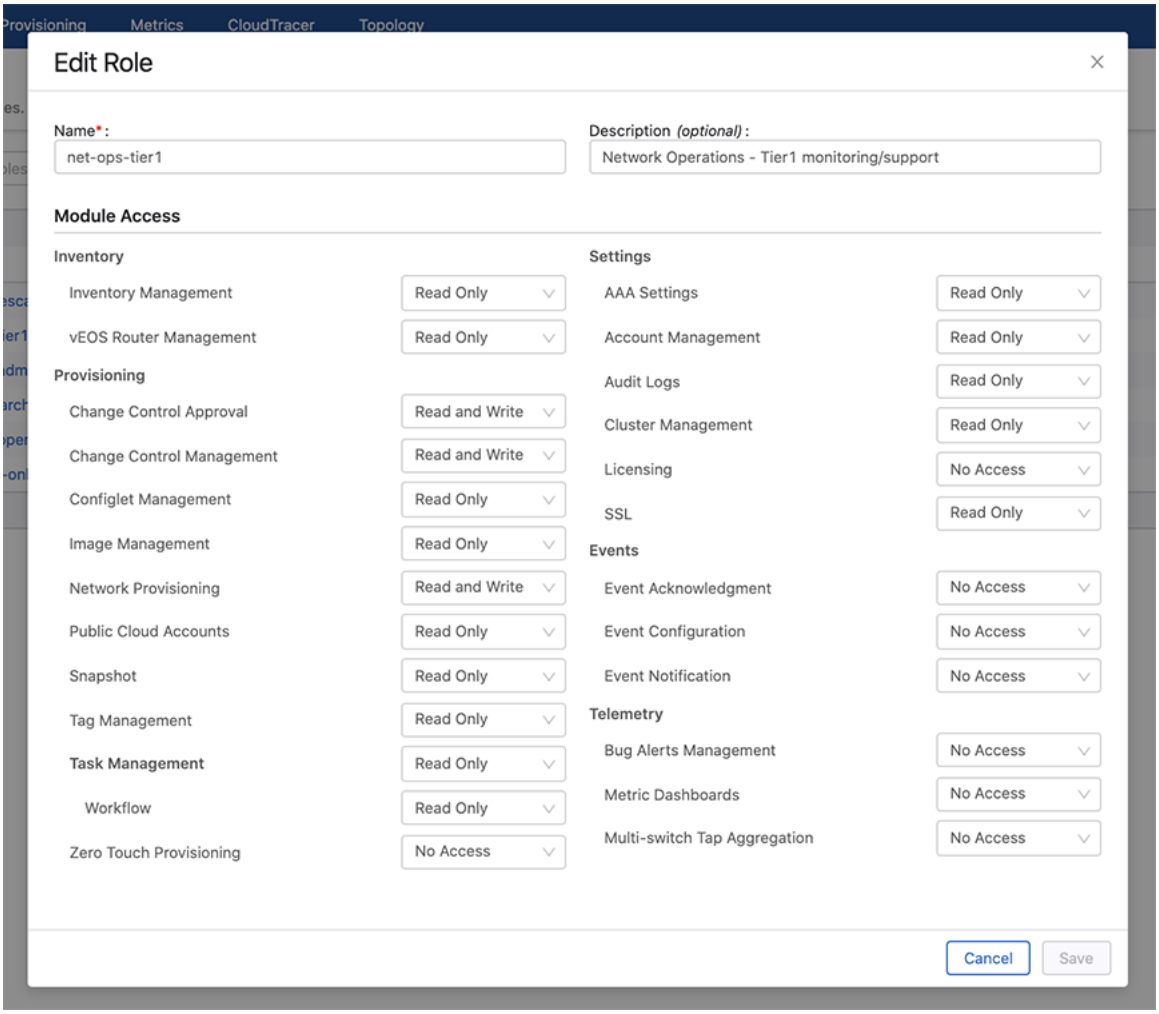

Figure 314: Edit Role Pop-Up Window

4. Modify the required Information.
5. Click **Save**.

   The new version of the role is saved to the CVP database.

   Note: All users assigned the role inherit the read and write permissions defined in the new version of the role.

   Related topics:

   - Adding New User Roles
   - Removing User Roles
   - Viewing Activity Logs

15.5.3 Removing User Roles

Complete these steps to remove a user role:

1. Navigate to the **Access Control** page.
2. Under **Access Control** in the left menu, click **Roles**.

   The Roles page lists all current user roles.
3. Select the required user roles for removal.
4. Click **Remove Role/Remove Roles** at the upper right corner of the **Roles** page.

   The system prompts to confirm removal.

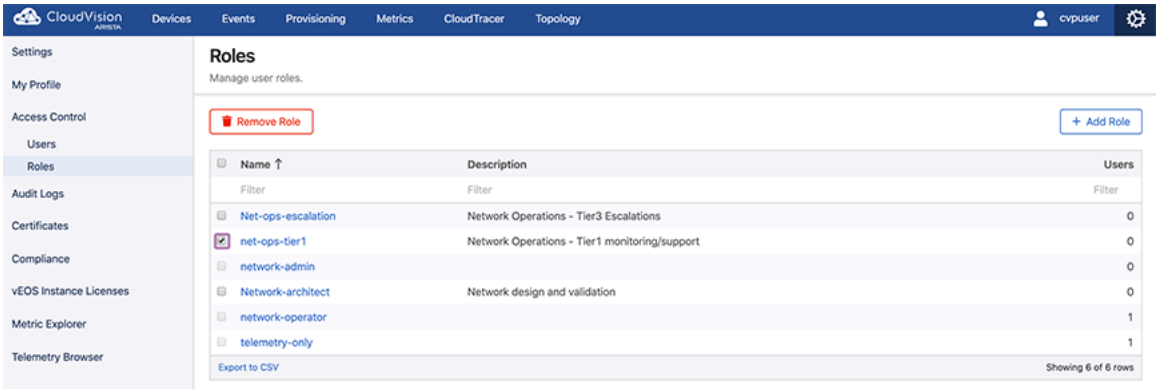

Figure 315: Remove User Role

5. Click **Delete**.

   The system deletes selected user roles.

   Note: A role assigned to user(s) cannot be deleted.

   Related topics:

   - Adding New User Roles
   - Modifying User Roles
   - Viewing Activity Logs

15.6 Viewing Activity Logs

The **Audit Logs** page displays activity logs of user accounts and user roles.

Complete these steps to view activity logs:

1. Click the gear icon at the upper right corner of the CVP page.

2. Click **Audit Logs** on the left menu.

   The system displays the Audit Logs page.
3. Select desired options from **View** logs for drop-down menus.

   The system displays corresponding logs.

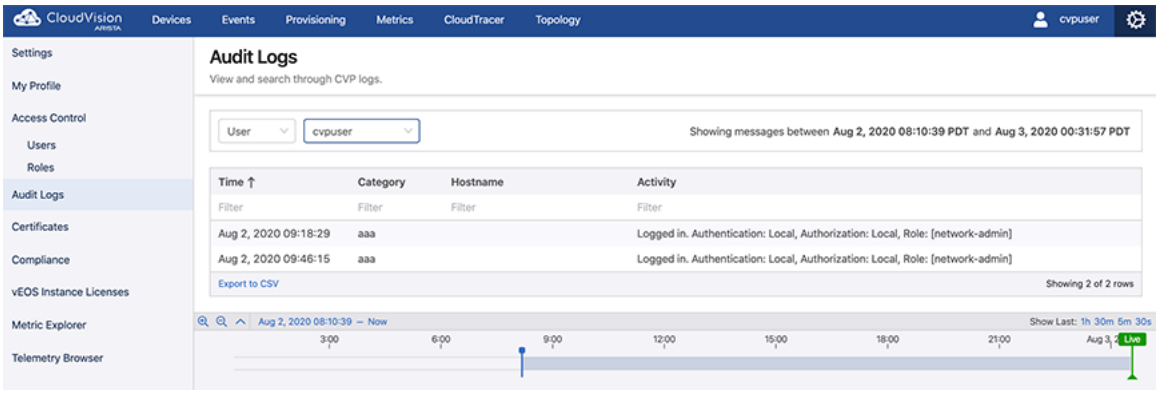

Figure 316: Audit Logs Page

15.7 Advanced Login Options

Multi-Factor Authentication (MFA) and One-Time Passwords authenticate all CVP managed devices when you authenticate with CVP. CVP runs CLIs on managed devices by sending eAPI requests over the gRPC connection established by TerminAttr.

Note:

- Under **Cluster Management** on the settings screen, enable **Advanced login options for device provisioning** to use MFA and one-time passwords.
- CVP needs TACACS to perform command authorization and accounting as per EOS configuration.
- Use the new Device class to make eAPI requests for using this mechanism in Configlet Builder python scripts.

Pre-requisities to install this feature are:

- Devices must run CVP 2018.2.3 or later releases
- Managed devices must have TerminAttr version 1.5.0 or later versions

  Note: TerminAttr is included with EOS, but may be a version earlier than v1.5.0. Newer versions are available as an extension (swix)

  Refer to CVP and TerminAttr release notes available at https://www.arista.com/en/support/software-download for detailed information on compatible TerminAttr versions with CVP and EOS.
- Ensure that the eAPI unix domain socket is enabled with `management api http-commands` and `protocol unix-socket` configurations in devices running EOS releases prior to 4.20

To enable MFA and One-Time Passwords authentication, enable **Advanced login options for device provisioning** using the toggle button under **Cluster Management** on the Settings page. See the figure below.

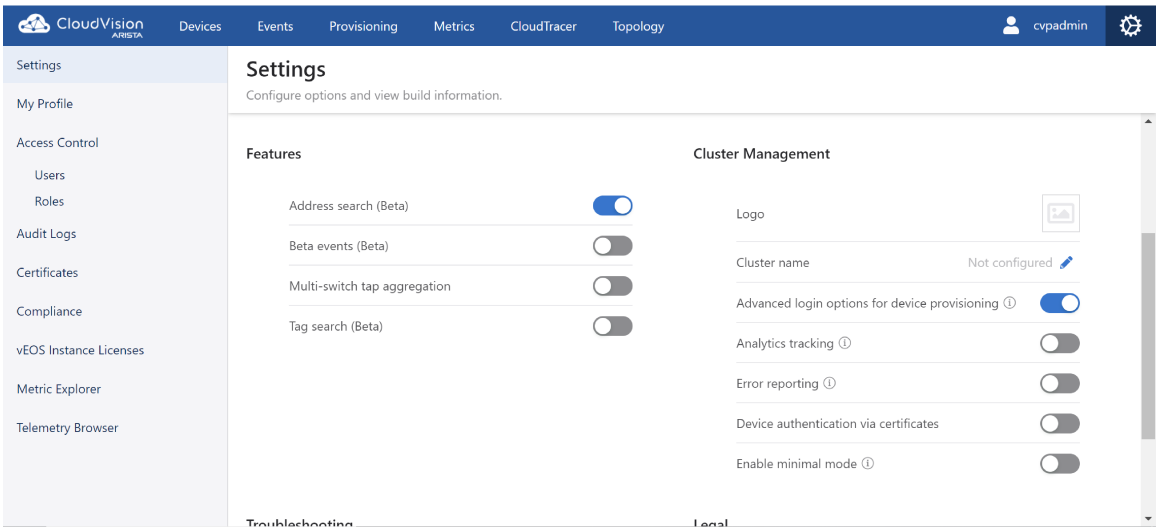

Figure 317: Advanced Login Options for Device Provisioning Toggle Button

15.8 Access to the Access Control Page

To gain access to the Access Control Page, complete the following:

1. Click the gear icon on the home page.

Figure 318: Gear Icon

2. Click Access Control in the left menu.

   The system displays the Initial Access Control Page.

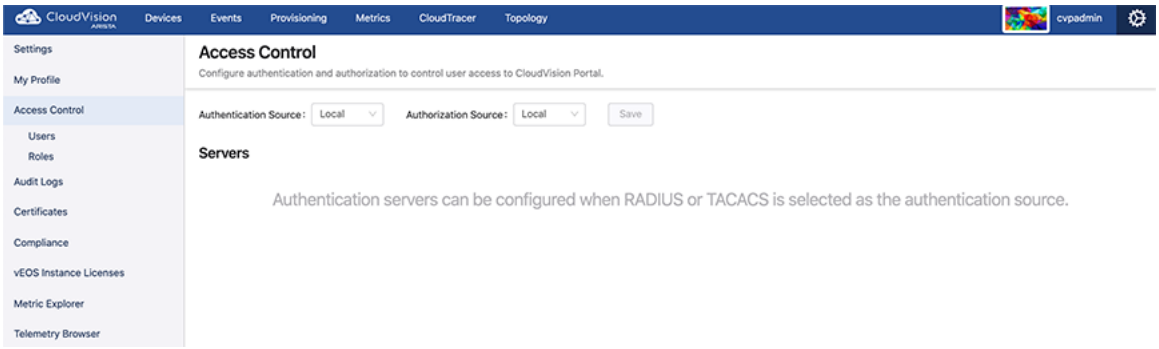

Figure 319: Initial Access Control Page

The system displays the Servers section when either RADIUS or TACACS is selected as Authentication source.

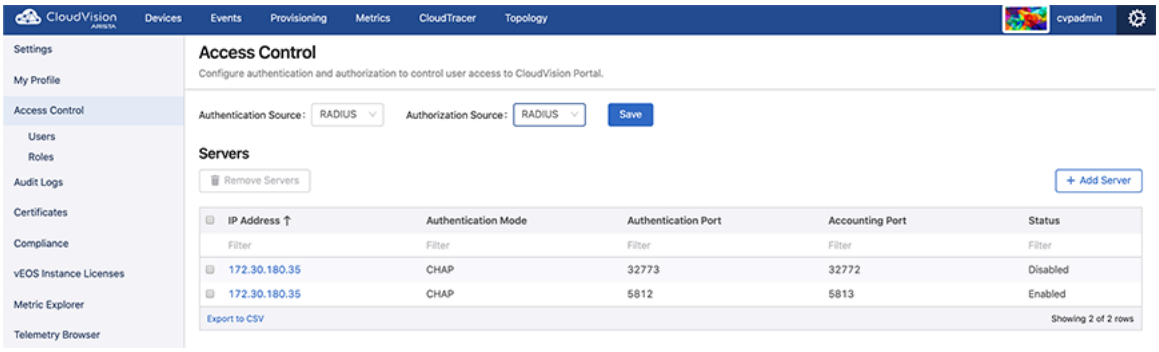

Figure 320: AAA Access Control Page

- If the authentication is local, the authorization must be done locally.
- If the authentication is done externally, the authorization can be done locally or externally.

Table 15: Server Authentication and Authorization

Authentication	Authorization
Local	Local
RADIUS	Local RADIUS
TACACS	Local TACACS

Note: External servers supported by CloudVision are RADIUS and TACACS.

Related topics:

- Managing AAA Servers
- Managing User Accounts
- Managing User Roles
- Access Requirements for Image Bundle Upgrades

Chapter 16

CloudTracer

Cloud Tracer tracks connectivity to monitor metrics streamed from EOS devices. The section in this chapter includes:

- Accessing the CloudTracer Screen
- CloudTracer Latency Anomaly Events

16.1 Accessing the CloudTracer Screen

To view data metrics, open to the CloudTracer screen by clicking **CloudTracer** on the CloudVision Portal (CVP).

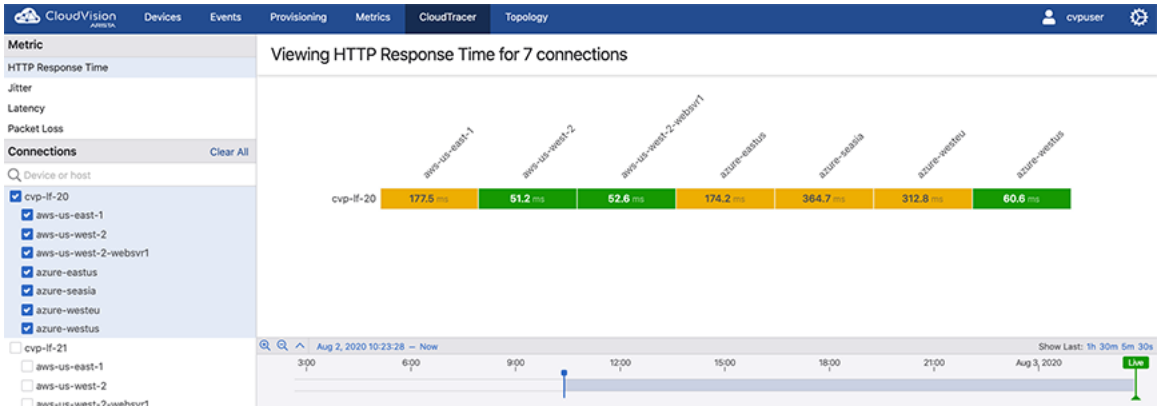

Figure 321: CloudTracer Screen

This screen is divided into the following two panels:

- Left Panel of the CloudTracer Screen
- Right Panel of the CloudTracer Screen

16.1.1 Left Panel of the CloudTracer Screen

This panel provides the following metric options:

- **Metric** pane - Click any of the following entities to view the corresponding current metric for n connections where n is the count of selected devices and hosts:
  - HTTP Response Time
  - Jitter
  - Latency
  - Packet Loss
- **Connections** pane
  - Device or host search string - Type the device or host name for a quick search
  - Configured devices - Select the required devices and hosts to view corresponding metrics
- **Clear All** - Click to clear all selections

16.1.2 Right Panel of the CloudTracer Screen

This panel displays metrics of selected options in the following ways:

- Current information of the selected metric type from selected devices and hosts

  Note: Metrics are streamed whenever data is gathered on EOS switches. The default interval to query metrics data is five seconds.

- Click on a metric to view detailed information.

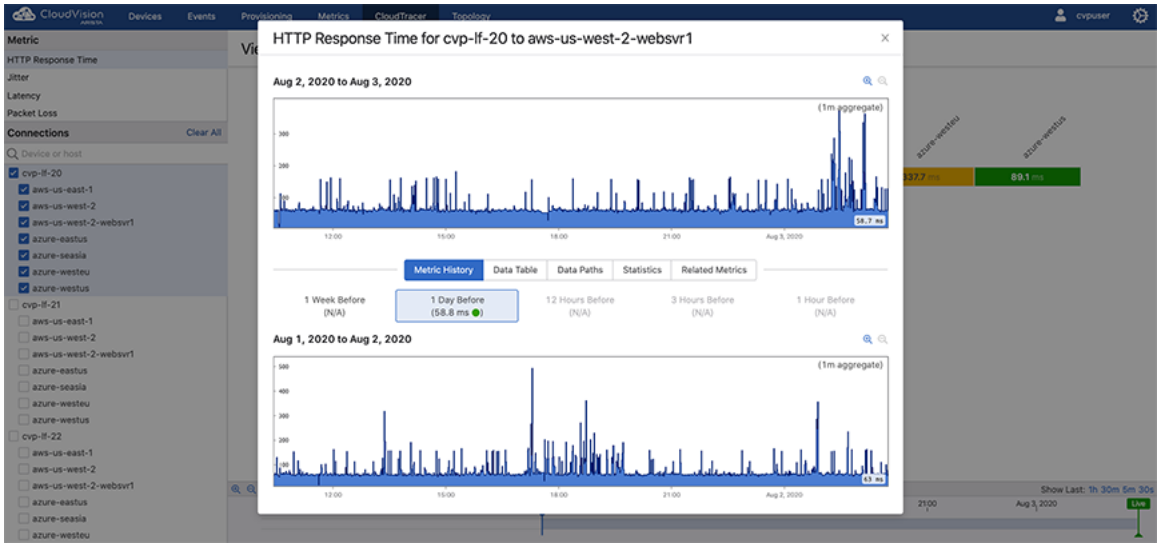

Figure 322: Detailed Metrics

The upper panel of this screen provides graphical presentation of the metric. The lower panel of this screen displays the metric through following categories:

- **Metric History** tab - Displays the metric history ranging from the last hour to the last week.

  Click the required timeline to view corresponding metrics.

  Note: Click Zoom In and Zoom Out options to view metrics ranging from every 15 minutes to every minute.

- **Raw Data** tab - Displays indexes, timestamps, and values of raw data.

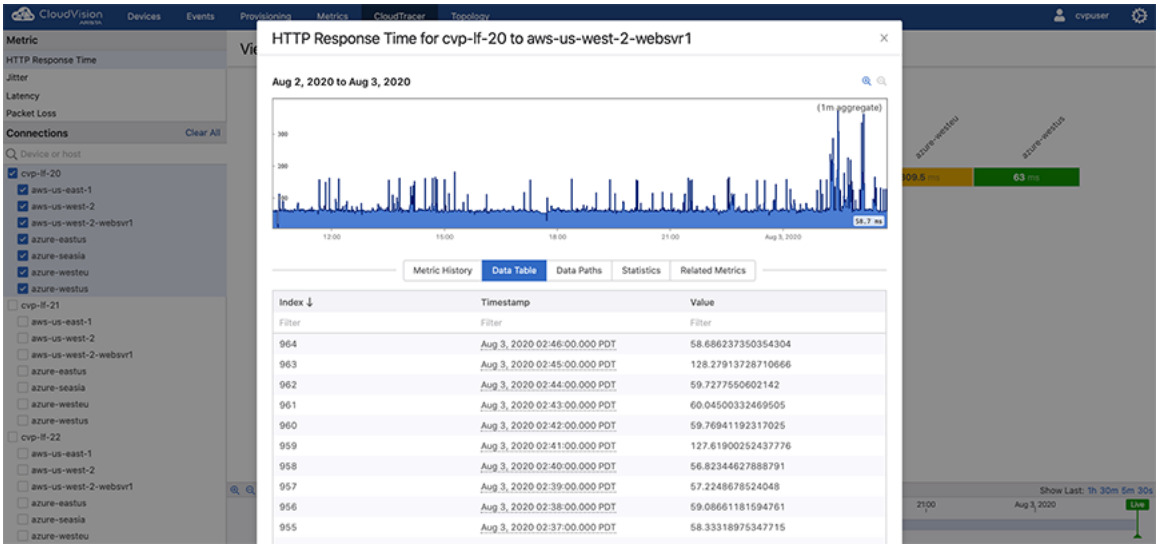

Figure 323: Raw Data Tab

- **Data Paths** tab - Displays keys and data paths used to compute the data for this metric.

Figure 324: Data Paths Tab

Note: Clicking required link navigates to the corresponding path in the telemetry browser.

- **Statistics** tab - Displays statistics of the selected device.

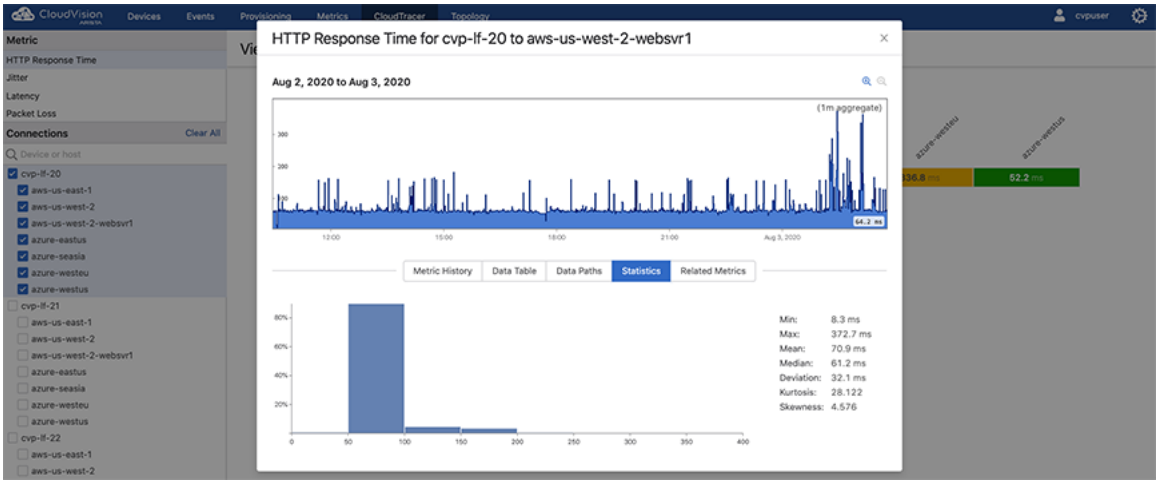

Figure 325: Statistics Tab

- Hover the cursor on metric to view metrics from all metric types.

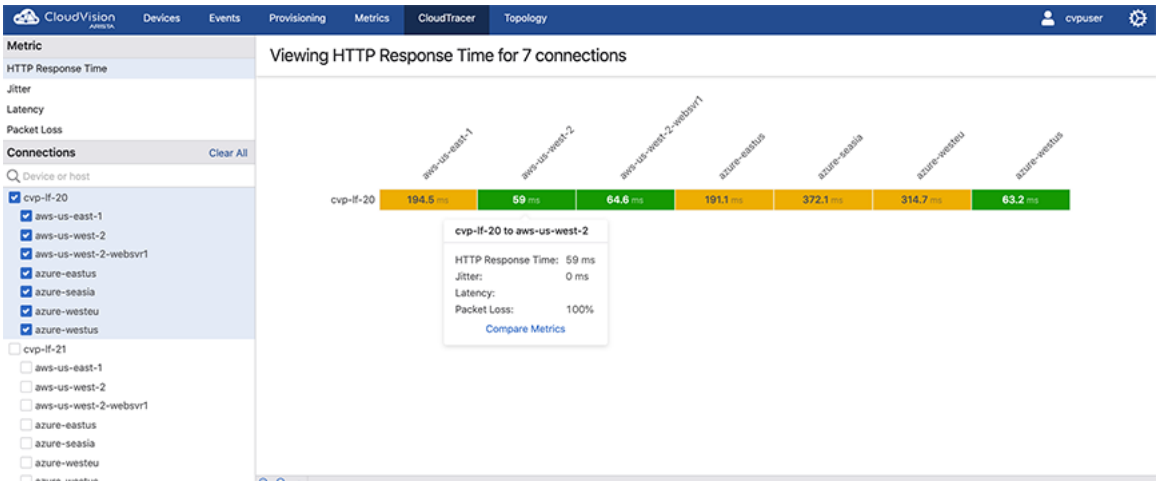

Figure 326: Metrics from All Metric Types

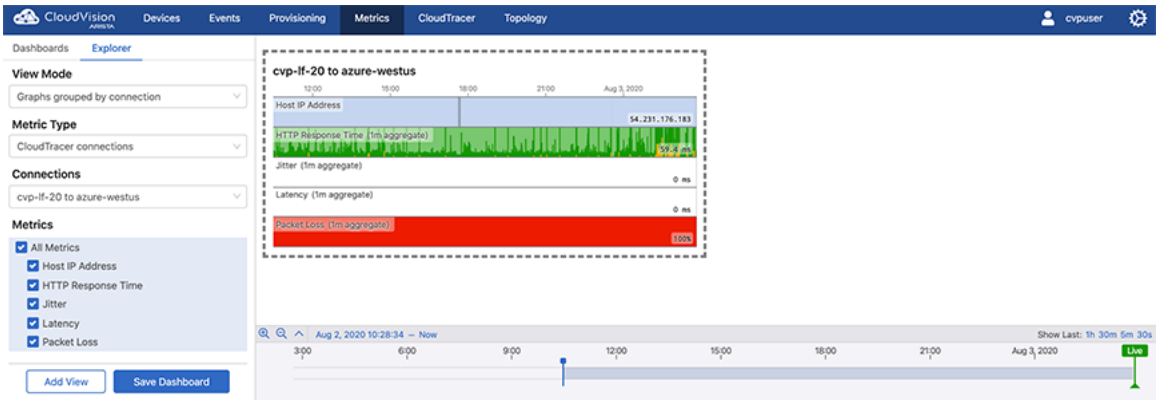

Figure 327: Metrics History of Selected Device

16.2 CloudTracer Latency Anomaly Events

The cloudtracer latency anomaly event monitors the latency metric between devices and configured hosts. The events are designed to alert the user when the latency between a device and a configured host is outside of recent historical bounds.

Figure 328: Anomaly Event View is a sample event view for one of these events between the device with hostname `Oslo` and the cloudtracer host endpoint `www.bbc.co.uk`.

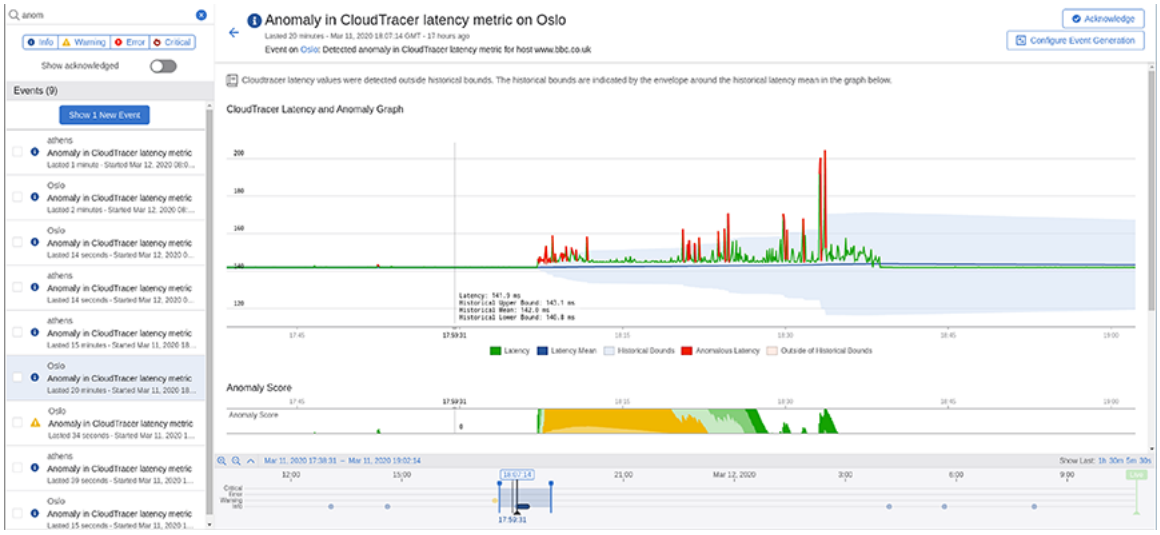

Figure 328: Anomaly Event View

Figure 329: Anomaly Event View Overlay explains various stages of this event.

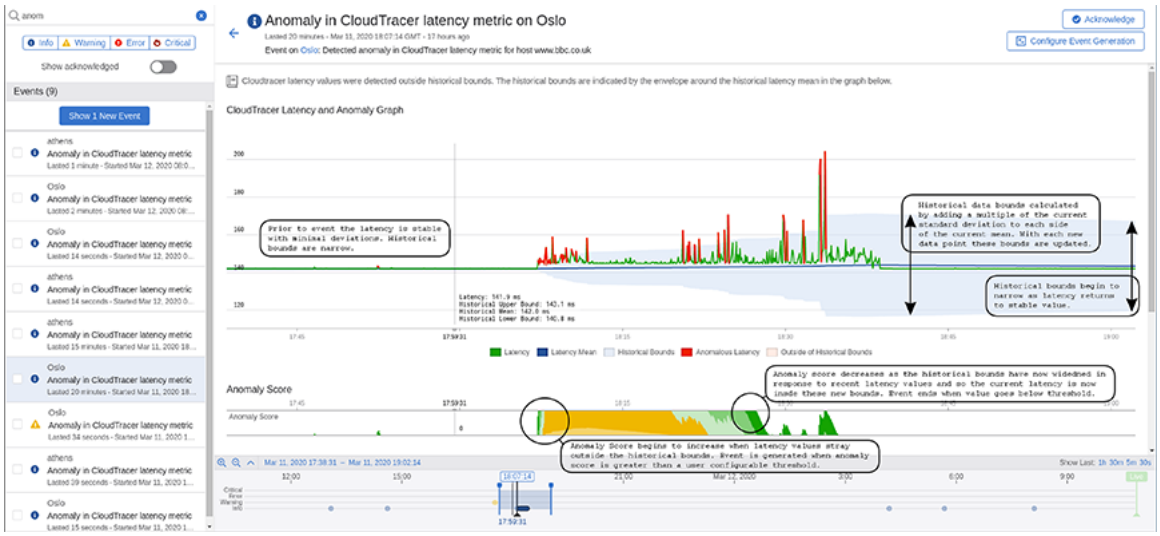

Figure 329: Anomaly Event View Overlay

Prior to this event in Figure 329: Anomaly Event View Overlay, the latency metric (green line in upper graph) is stable with minimal deviations. The historical bounds (blue shaded region) that determine when the metric is in a normal state has a small range with both the upper and lower bounds near the historical mean (dark blue line). The historical bounds are computed by adding and subtracting a fixed multiple of the current latency standard deviation to the current mean.

The anomaly score starts to increase from zero when the latency value strays outside of the historical bounds. The latency values that are outside the bounds are highlighted in red. The anomaly score is the total number of standard deviations outside the historical bounds. The anomaly score is the positive cumulative sum of the number of standard deviations outside of the historical bounds. For example, if the bounds are set as 3 standard deviations outside of the mean and we get a value of the latency that is 5 times the standard deviation away from the mean, the anomaly score will increase by 2. If the next latency value was 1.5 times the standard deviation outside of then mean then we would subtract 1.5 from the anomaly score. The anomaly score therefore keeps track of the cumulative deviation of the latency outside of the historical bounds. It is bounded below by zero.

Figure 330: Anomaly Score Computation provides a detailed explanation on computing the anomaly score.

Figure 330: Anomaly Score Computation

The event is generated when the anomaly score exceeds a threshold for a set period of time.

Note: You can configure the threshold and time duration in the event configuration rules.

The anomaly score starts to decrease when the latency values are inside the historical bounds. The historical bounds have increased based on recent deviations in latency which makes the system less sensitive than prior to the event. The event ends when the anomaly score is below the threshold for a set period of time.

Figure 331: Decreasing of Anomaly Score provides a detailed explanation of the anomaly score decreasing when an event ends.

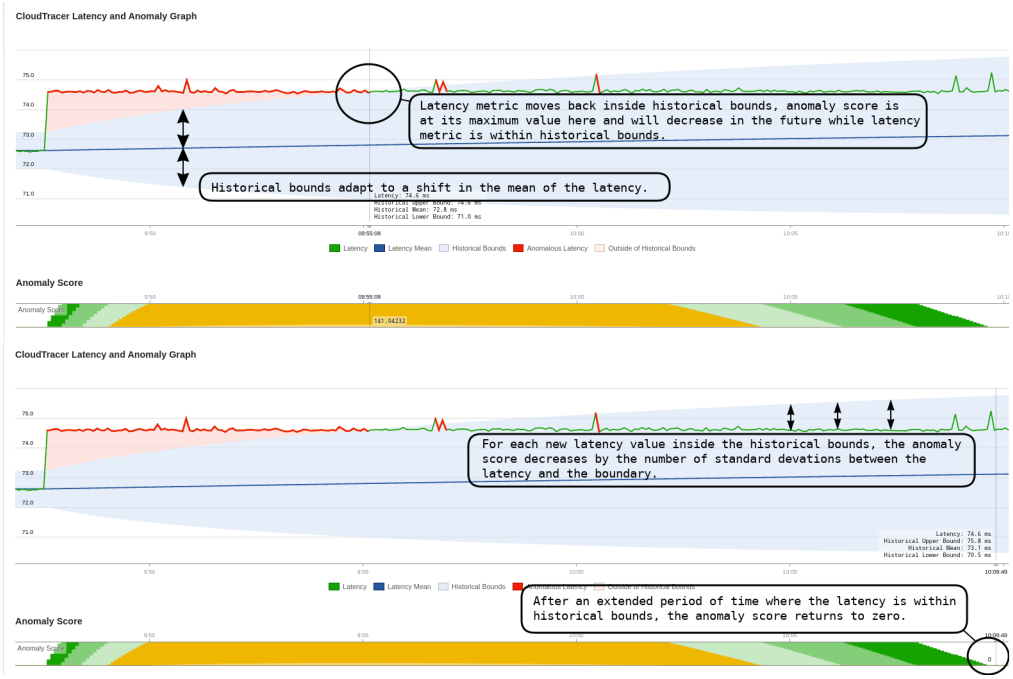

Figure 331: Decreasing of Anomaly Score

At the end of the time range, historical bounds are narrowing as the latency has now returned to a stable value with minimum deviations. The history needs approximately six hours to have negligible impact on the statistics and bounds.

This screen also provides the following additional metrics of this event (see Figure 332: CloudTracer Event Additional View):

- The other CloudTracer metrics are displayed for this device and host pair
- The latency metric between other devices and this host
- The latency metric between this device and other hosts

Figure 332: CloudTracer Event Additional View

Chapter 17

CloudVision Topology

The CloudVision Topology screen provides an explicit visual representation of the connectivity of your network, allowing you to understand your network's structure and performance more easily. It provides the following benefits:

- Easily understand parts of your network by collapsing or filtering out irrelevant parts
- Explore the historical state and performance of your network or watch it update live
- Support for both datacenter and campus style network connectivity

CloudVision topology provides Virtual Extensible LAN (VXLAN), Internet Protocol Security (IPsec), Distributed Path Selection (DPS), and Link Layer Discovery Protocol (LLDP) network links between endpoints.

Note:

- Information and Statistics for each member link is accessed from the side panel. See Topology Overview.
- If this screen does not display any devices, refer to the CVP release notes at https://www.arista.com/en/support/software-download for compatibility issues.

To view the Topology screen, click the **Topology** tab on the CloudVision Portal.

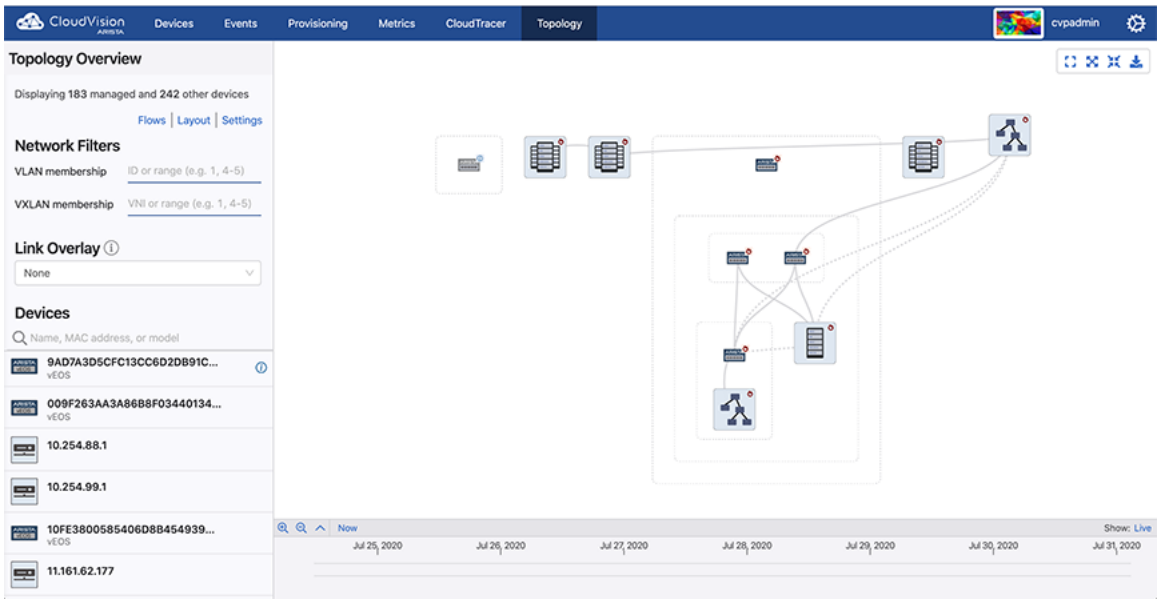

Figure 333: Topology Screen

This screen is divided into main and side panels. The main panel displays the main topology visualization. Devices are drawn with paths to connect them if they share at least one network connection. They are grouped into containers that can be expanded or collapsed to control which portions of the network are displayed in detail. See Main Panel of the Topology Screen.

The side panel provides the following panes to perform the specified functionalities:

- To customize the network view:

- Topology Overview
- Topology Layout Pane
- Topology Options Pane

- To view the component information:
  - Container Details Pane
  - Device Details Pane
  - Link Details Panel
  - Flow Visibility

17.1 Main Panel of the Topology Screen

The main panel displays the network topology where devices are grouped into containers according to their connectivity or assigned role in the network.

The icons in the following table represents specified containers:

Table 16: Icons Used in Network Topology

Cloud 	Datacenter 	Campus 
Building 	Floor 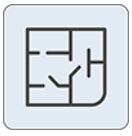	Pod 
Rack 	Spine 	

The icons in the following table represents specified devices:

Table 17: Device Icons

Switch 	Wireless Access Point   **Note:** Blue WAP represents managed devices. Gray WAP represents unmanaged devices.	Management Device Badge   **Note:** This badge next to a device icon represents a management device.

Computer	Third Party Device	Telephone
		

This panel provides the following options for a detailed view:

- Zoom to fit icon - Click to fit the topology on the screen.
- Expand containers icon - Click to expand all containers in the topology.
- Collapse containers icon - Click to collapse all containers in the topology.
- Alternatively, right-click on the main panel to get **Expand Network**, **Expand All**, and **Collapse All** options.

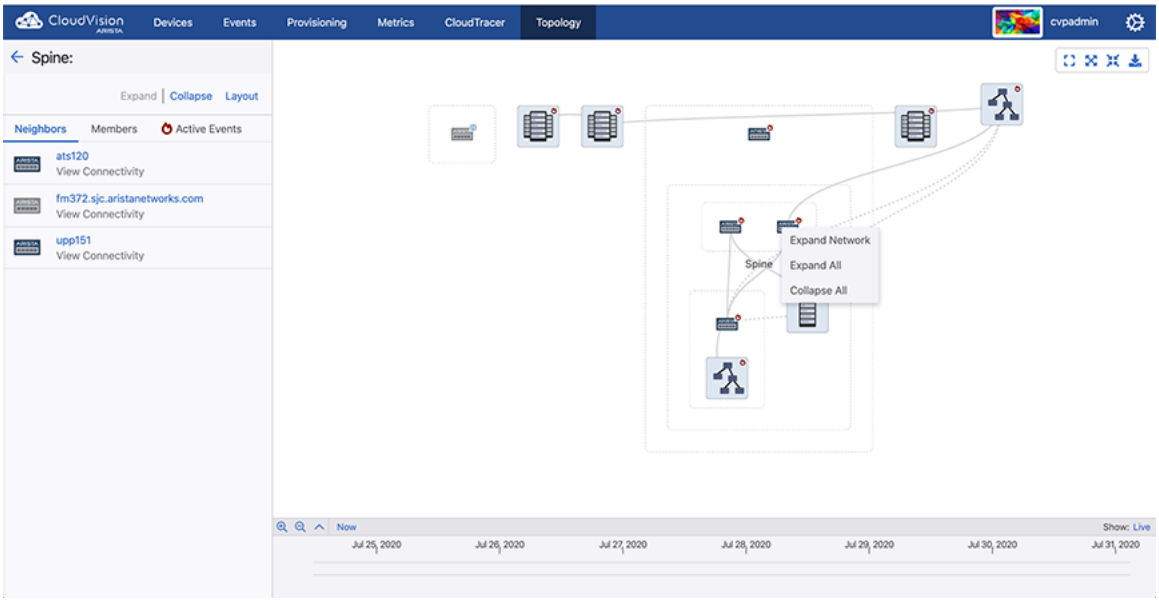

Figure 334: Right-Click on a Device

Note: Right-click on a cluster to get cluster specific context menu options.

- Download icon - Click to open the Export Preview pop-up window. Click **Export** for downloading the current topology image to your local drive in either PNG or SVG formats with selected image resolution.

Note: We recommend to select higher resolutions for readable device labels in bigger topologies.

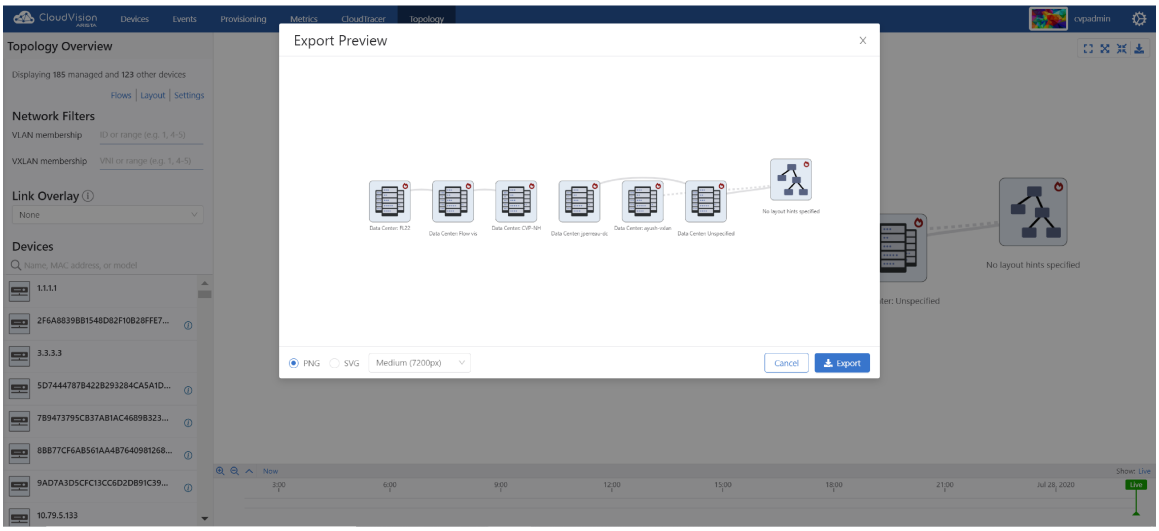

Figure 335: Export Preview Pop-Up Window

- Double-click on a container to expand it.
- To collapse a container, hover the cursor on a dotted rectangular box and click on the displayed hyphen symbol.

Figure 336: Collapse a Container

- Click container component(s) to view corresponding information on the left panel.
- Selected components are highlighted with dashed frame.

> **Note:** Press and hold the shift key while selecting multiple devices. Press and hold the shift key while dragging to select a region.

- Hover the cursor on a topology component to view the count of corresponding events.

> **Note:** You must enable the option to view events.

17.2 Topology Overview

The Topology Overview pane provides the following options:

- **Layout** - Click to view the **Topology Layout** pane. See Topology Layout Pane.

- **Options** - Click to view the **Topology Options** pane. See Topology Layout Pane.
- **Network Filters** - Provides the following options to filter networks:
  - **Management network** - Display or hide management networks using the toggle button.
  - **VLAN membership** - To view desired VLAN(s), type either a VLAN ID or a range of VLANs.

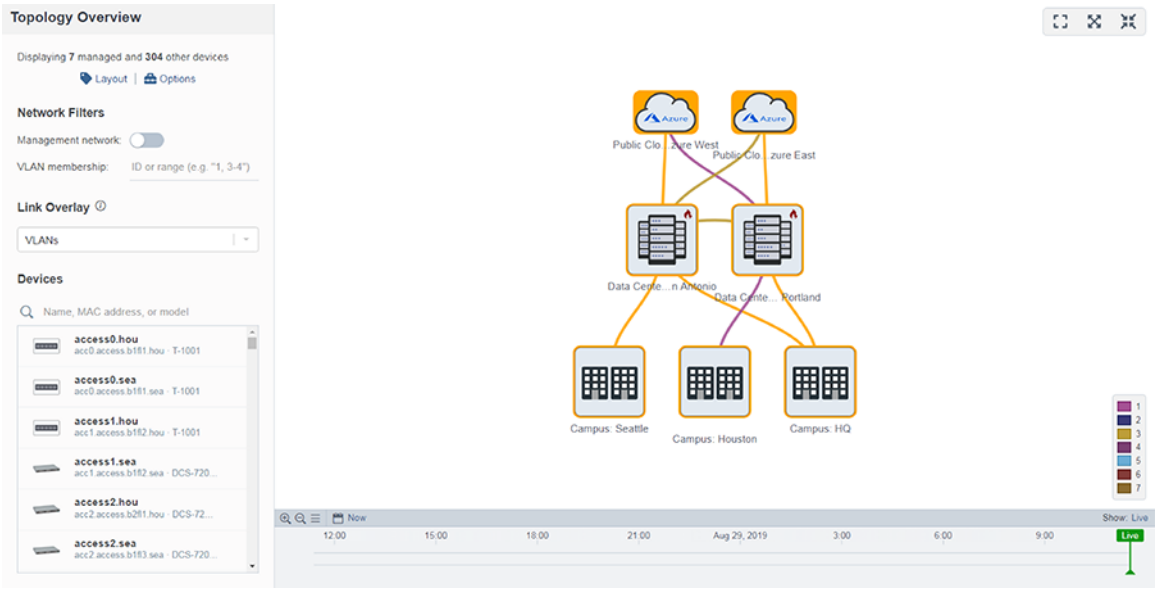

Figure 337: VLANs in Topology

Note: The right panel displays selected VLAN(s) distinguished with various colors.

- **Link Overlay** drop-down menu - Select an overlay to color each link based on selected metric type. Options include:
  - Active Events
  - Bandwidth Utilization
  - Discard Rate
  - Error Rate
  - Traffic Throughput
  - VLANs
  - None
- Devices
  - Search field - Type the device name, MAC address, or model to perform a quick search.
  - List of devices - Click on a device to view the detailed information of corresponding device. See Device Details Pane.

17.3 Topology Layout Pane

On the Topology Overview pane, click **Layout** and select a container component from the topology on the right panel to edit layout hints of multiple device(s) in the **Topology Layout** pane.

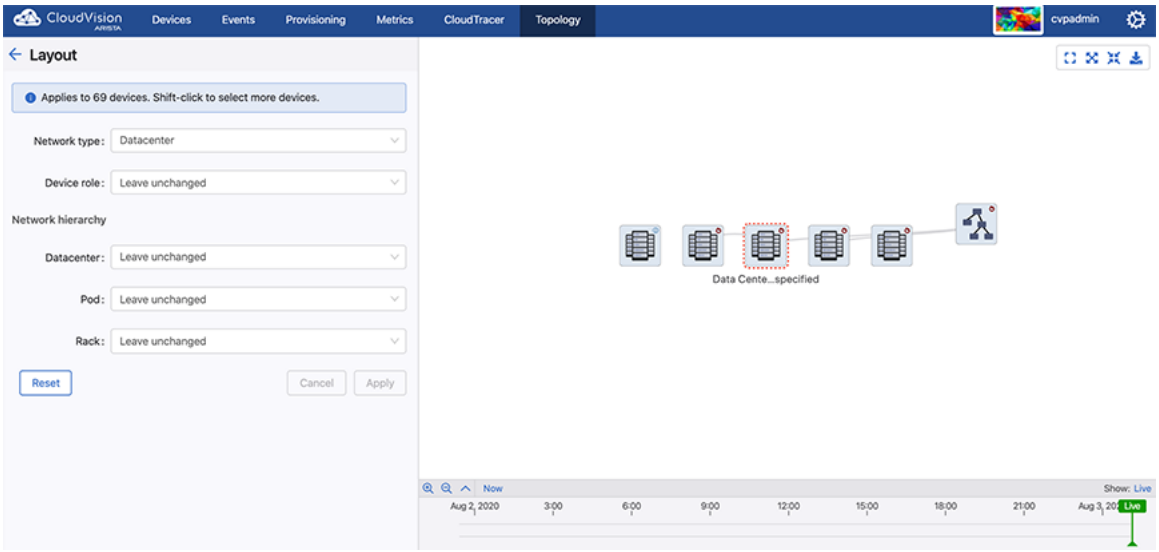

Figure 338: Topology Layout Pane

Topology automatically tries to guess a layout with specified containers and roles for your devices based on their connectivity and advertised LLDP capabilities. However, you might sometimes find that the automatic categorization is incorrect, or you simply want a custom layout different from what was originally envisioned. The **Layout** pane lets you override the automatic categorizations and control the layout more directly.

The layout works on the basis of hints that describe the role of a device, whether it exists within a datacenter or campus network, and where it should go in that network. Devices with similar roles and positions in the hierarchy are grouped together. Parallel hierarchies like network pods or racks are created if different names are used.

Examples

- A device named *athens* is a datacenter leaf switch, but it has no rack server connections yet and is miscategorized as an edge switch. You can click on athens and then select **Node type** as **leaf** to force it to take on a leaf role. It moves into the leaf position inside its datacenter hierarchy.
- To partition your network into New York and San Francisco datacenters, multi-select the devices or containers that must go in the New York datacenter, type **New York** in the **Datacenter** field, and confirm it. Repeat the same process for San Francisco. Now, your network is divided between these two datacenters, and you can expand or collapse New York and San Francisco datacenters independently to view only one datacenter at a time.

This pane provides the following selections:

- **Network type** drop-down menu - Select the network type that most closely matches your network arrangement. It provides the following options:
  - **Campus** - Devices are manually arranged in containers for different buildings and floors. It provides the following options:
    - **Node type** drop-down menu - Select the preferred device type or roles.
    - **Building** drop-down menu - Select the building name that the selected device preferred to be placed into.
    - **Floor** drop-down menu - Select the preferred floor number in the selected building.

- **Devices** drop-down menu (Optional) - Set a name to be used to group devices in the selected floor.
- **Datacenter** - Aspine-and-leaf type layout is used and devices are arranged into pods and racks. It provides the following options:
  - **Node Type** drop-down menu - Select the preferred device type or roles.
  - **Pod** drop-down menu - Select the pod name that the selected device preferred to be placed into.

    Note: Devices in different pods of the same datacenter appear in different pod containers that can be expanded and collapsed independently.
  - **Rack** drop-down menu - Select the name of a rack similar to pod.

- **Show Advanced** - Click to view the **Skip Auto-Generated Classifications** drop-down menu.

  Note: Click **Hide Advanced** to hide the **Skip Auto-Generated Classifications** drop-down menu. If the **Skip Auto-Generated Classifications** option is enabled, CVP does not automatically identifies the device(s). Only manually-provided layout hints affect the layout of the selected device(s).
- **Set all to Auto** - Use the automatic layout classification exclusively; all manually-specified layout hints are removed from selected devices.
- **Save** button - Click to save latest changes.

17.4 Topology Options Pane

On the **Topology Overview** pane, click **Options** to edit display settings of topology.

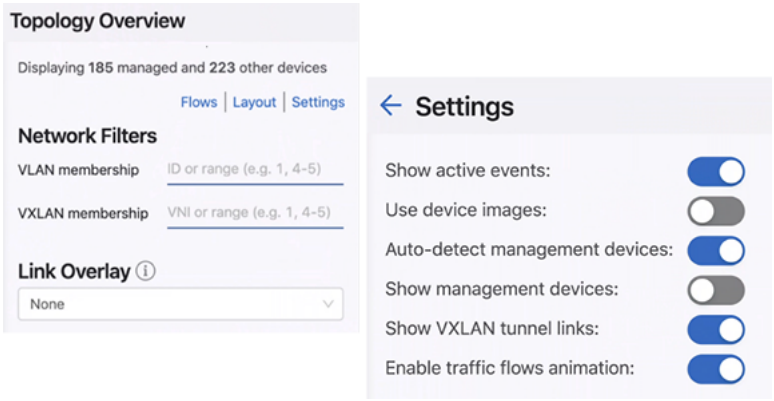

Figure 339: Topology Options Screen

This pane provides the following selections:

- **Show active events:** toggle button - If this option is enabled, active events are shown as badges on devices. These are the same events that are displayed on the Events page. If the same device has multiple events, the badge type of the highest severity event is displayed. Containers also show badges if they contain any devices with active events. This allows you to quickly find active events anywhere in a large network.

  Note: This option is enabled by default.
- **Use device images:** toggle button - Enable this option to view photorealistic device images for identified devices. If this option is disabled, icons are used instead. See Figure 340: Network Hierarchy Tree with Images.

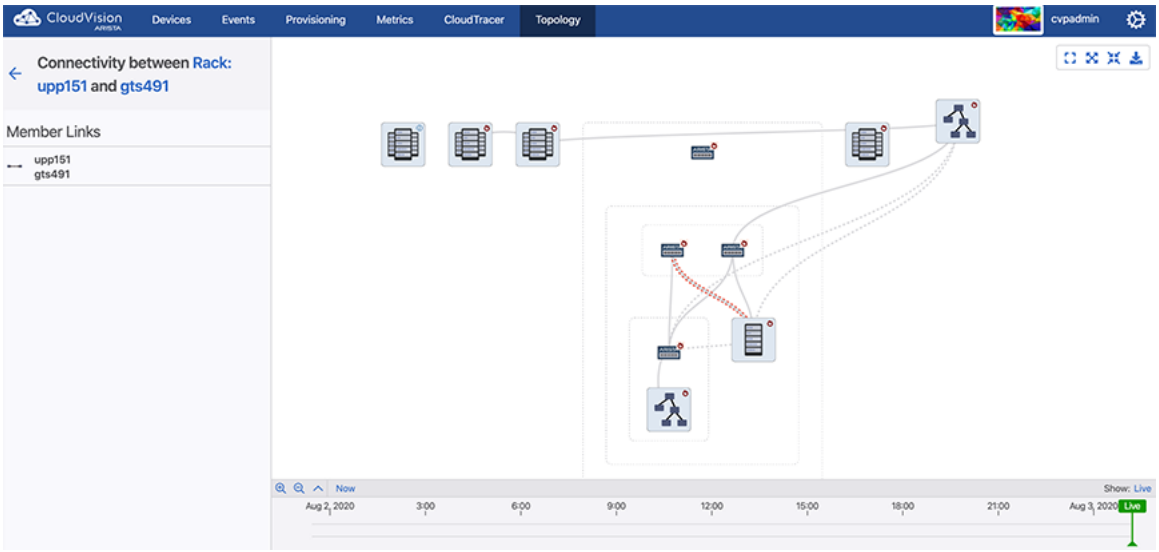

Figure 340: Network Hierarchy Tree with Images

- **Auto-detect management devices:** - If this option is disabled, CVP will not attempt to automatically identify management devices. Devices are considered management devices if they are known to have a relatively high number of connections over a management interface.
- **Auto tagger hints** pane - Influences the way devices are arranged. If a device's hostname matches the provided text string or regular expression, it will automatically be tagged with the given role. Options include:
  - **Spine Hint**: - Type a text string that is used to identify matching spine devices.
  - **Leaf Hint**: - Type a text string that is used to identify matching leaf devices.
- **Save** button - Click to save latest changes.

17.5 Container Details Pane

To view more information about a device or the devices in a container, click the corresponding device or container on the right panel.

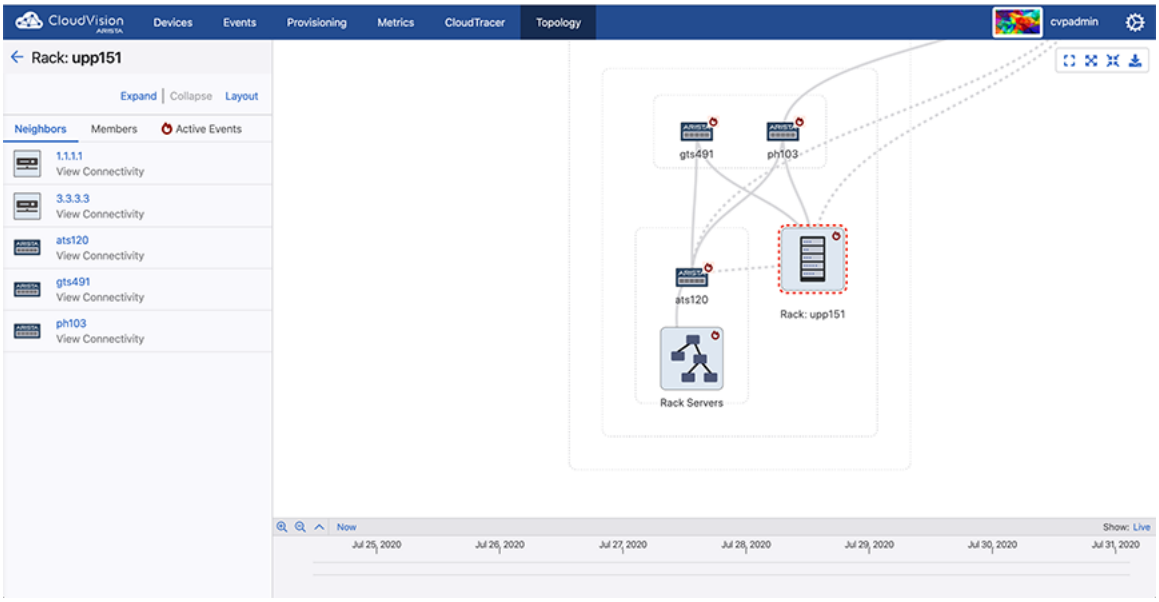

Figure 341: Container Pane

This screen provides the following functionalities:

- **Expand** - Expands the selected container.
- **Collapse** - Collapses the selected container.
- **Layout** - Edits layout hints of the selected container. See Topology Layout Pane.
- **Neighbors** - Displays the list of connected devices from neighboring container.

  Note: Click on any neighboring device name to view the corresponding device pane. See Device Details Pane.
- **Members** - Displays the list of container members. Each entry provides the following options:
  - **Device name** - Click to view the corresponding device pane. See Figure 342: Device Details Pane.
  - **View Connectivity** - Click to view the connectivity between selected device and neighboring device. See Link Details Panel.
- **Active Events** (Optional) - Displays events of the selected container. Click on an event link to view the corresponding event details screen.

  Note: This option is available only when the **Show active events** option is enabled in the Topology Options pane. See Topology Options Pane.

17.6 Device Details Pane

To get a device pane, click on a device (switch, wireless access point, server, or telephone) on the right panel. See Figure 342: Device Details Pane.

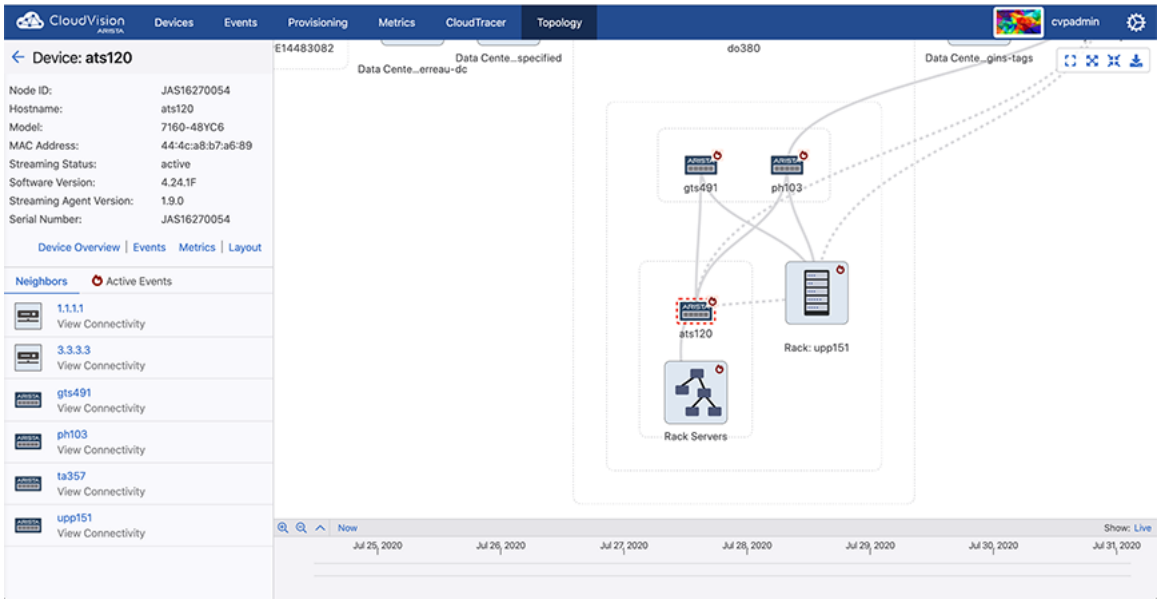

Figure 342: Device Details Pane

This screen provides the following functionalities:

- Additional information on the device.
- **Device Overview** - Click to view the Interface Overview screen. Device Overview
- **Events** - Click to view the Events summary screen. See Events Summary Screen.
- **Metrics** - Click to view the Explorer screen. See Explorer Tab.
- **Layout** - Click to edit layout hints of the selected device. See Topology Layout Pane.
- **Neighbors** - Displays the neighbors list of selected device. Each entry provides the following options:

- ***Device name*** - Click to view the corresponding device pane.
- **View Connectivity** - Click to view the connectivity between selected device and neighboring device. See Link Details Panel.

- **Active Events** (Optional) - Displays events of the selected device. Click on an event link to view the corresponding **Event Details** screen.

> **Note:** This option is available only when the **Show active events** option is enabled in the **Topology Options** pane. See Topology Options Pane.

17.7 Link Details Panel

To view the links panel, click on a connectivity link between two components on the right panel.

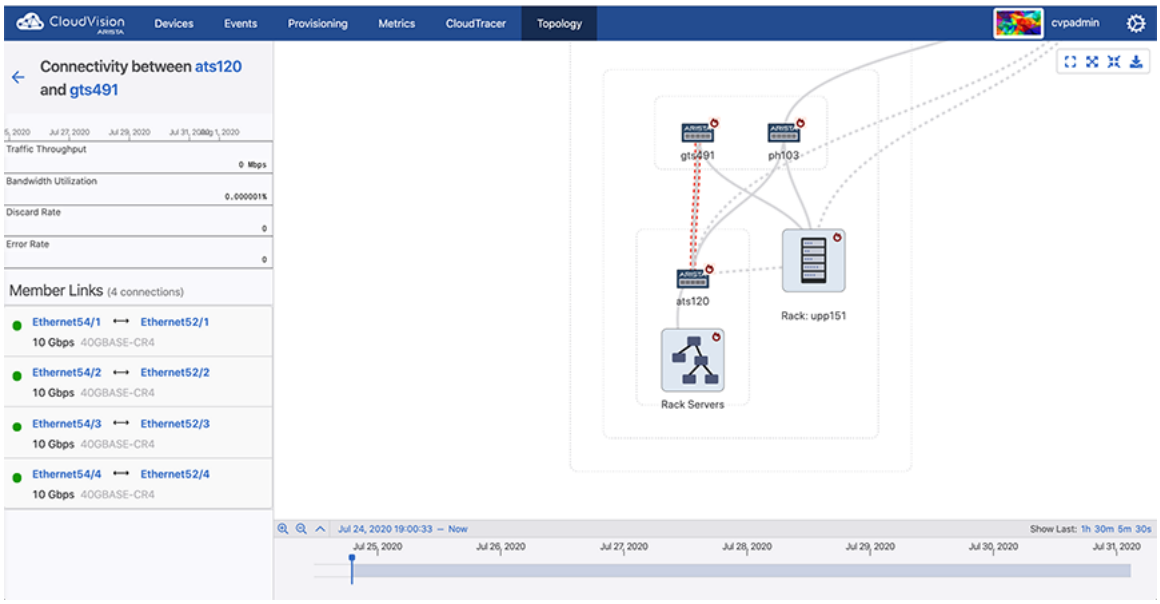

Figure 343: Links Panel

Links represent connections between devices or clusters of devices. If two devices or clusters have at least one network connection, a link is drawn to connect them. If they have many network connections, they still have a single link in the topology view and information provided for the link is aggregated over those connections. Expanding and collapsing containers expand and collapse links; you may sometimes want to expand containers to see links in greater detail.

This screen provides the following information of the selected connectivity link:

- Click on a device name to view the corresponding device panel.
- Metrics - Displays statistics of traffic throughput, bandwidth utilization, discard rate, and error rate.

  > **Note:** Hover the cursor on the metrics to view metrics at the corresponding time.

- **Member Links** - Displays the list of connected ports.

  > **Note:** Click on any connected port link to view the corresponding **Interface Overview** screen.

- **Events** - Displays events of the selected connectivity link. Click on an event link to view the corresponding **Event Details** screen.

  > **Note:** This option is available only when the **Show active events** option is enabled in the **Topology Options** panel. See Topology Options Pane.

17.8 Flow Visibility

On the Topology Overview pane, click **Flows** to open the **Topology Flows** panel. This screen displays traffic flows detected by EOS devices on the network.

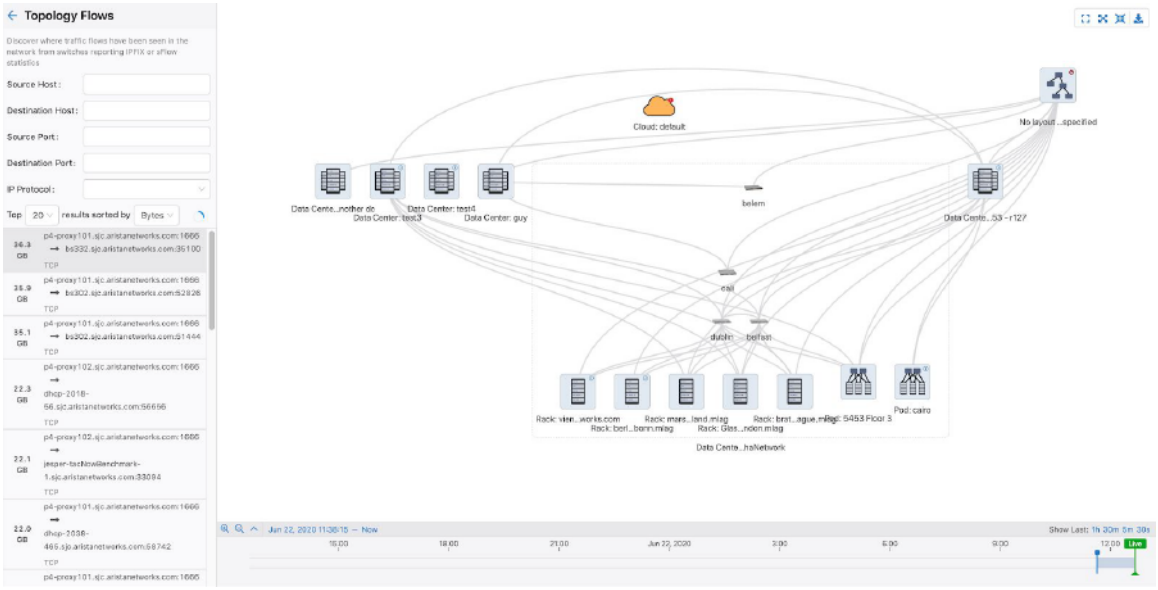

Figure 344: Topology Flow Search

Note:

- CVP displays traffic flows only when SFLOW of IPFIX are configured on EOS devices.
- For complete flow visibility, flow collectors are required on all devices along the traffic flow path.

The **Topology Flows** panel searches for traffic flows via specified IP address, hostname, ports or IP protocol and lists the flow results that match the given search parameters.

You can limit the count of displayed flows via the options available in the **Top** dropdown. Traffic flows sorted by the selected metric (**Bytes**, **Packets**, and **Newest**) from the **results sorted by** dropdown menu are displayed on the top of the list.

The listed traffic flows in the side panel displays the five-tuple information. The arrow indicates the direction of traffic flow.

Figure 345: Topology Host showing Flows

In this example, TCP protocol is used in the traffic flowing from p4-proxy101.sjc.aristanetworks.com via 1666 port to bs332.sjc.aristanetworks.com via 37150 port. 36.6GB of data is flown over the given time window.

Flows are displayed based on the timeline selected at the bottom of the Window. To search previous flows, select an earlier time by either using the timeline's time selector, or by dragging the displayed time window to a different position.

Note: Live view updates the data every 60 seconds.

Flow Highlight

Clicking on a listed traffic flow result highlights the nodes and edges in the graph where the flow has been seen. Animated dots indicate the direction of the traffic flow.

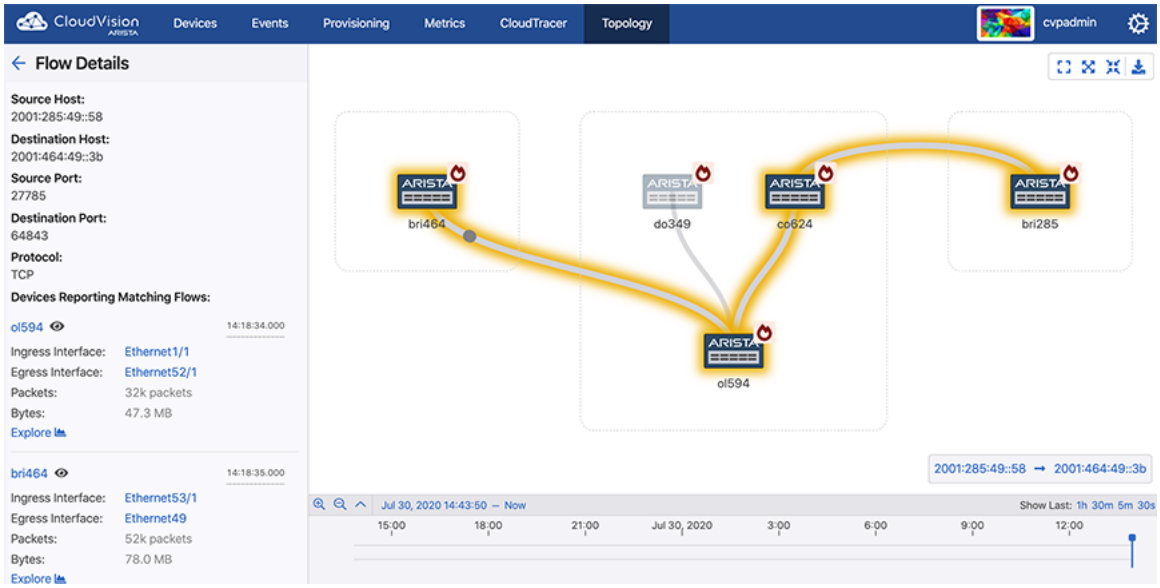

Figure 346: Highlighted Traffic Flow

Note:

- In environments that capture flow data through sFlow, devices may not capture short-lived or small flows, especially if the selected time window is small.
- This highlight does not guarantee to capture the exact path; it just displays all the devices and links where that flow was seen in the given time window.

The **Devices Reporting Matching Flows** section displays the five-tuple information and lists devices that reported the flow. Each device entry includes the ingress and egress interface, packets, bytes and the timestamp when this flow was seen given the time window.

Click on the following entities to view the corresponding specified information:

- Eye icon to magnify the device on the main panel
- *Device hostname* to view the Device Overview page
- *Interface* to view the Interface Overview page
- **Explore** button to the Traffic Flows

Flow Animation

Few browsers consume high amounts of CPU to render traffic flow animations. To avoid the traffic flow animation, click **Settings** on the **Toplogy Overview** panel and disable it using the **Enable traffic flows animation** toggle button.

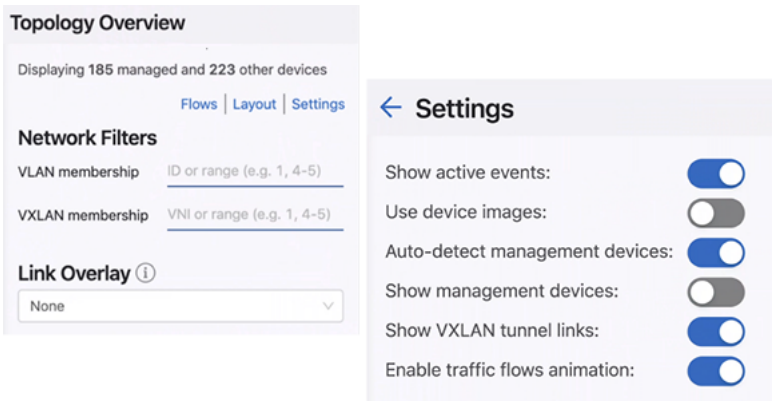

Figure 347: Enabling Traffic Flow Animation in Settings

Animated dots are replaced with static arrows indicating the direction of flow.

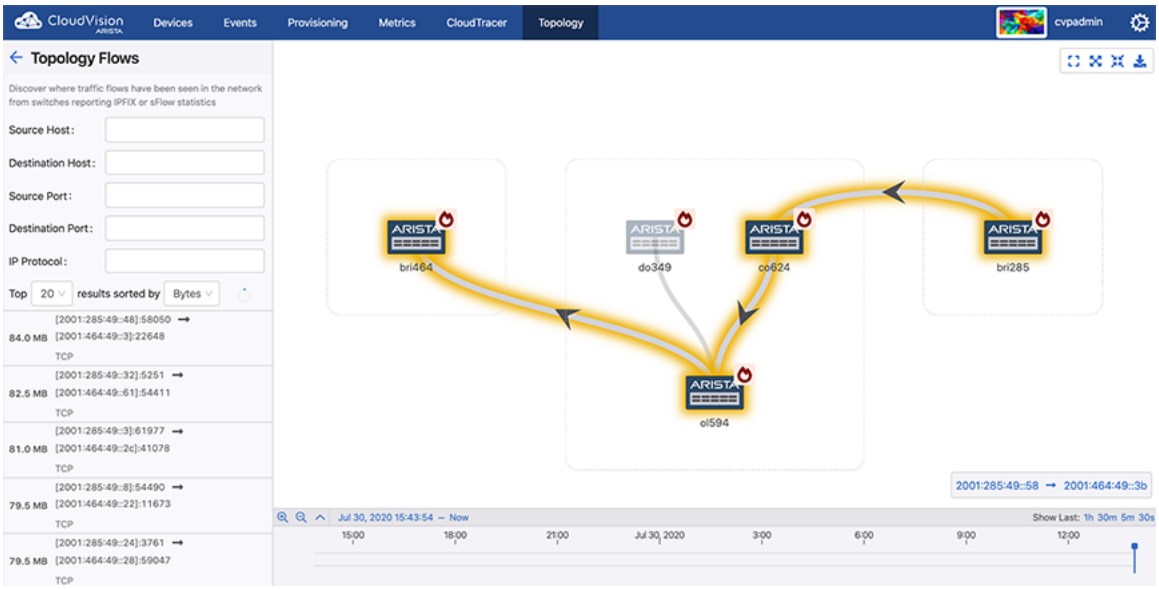

Figure 348: Topology with Disabled Traffic Flow Animation

Chapter 18

Tap Aggregation (CVP)

Arista EOS provides unprecedented visibility for rapidly identifying and troubleshooting application and performance problems with tracers such as VM Tracer and MapReduce Tracer. EOS integrates with Apache Hadoop systems to track big data workloads, aggregates and monitors business critical applications across thousands of devices, and provides deep visibility and integration with virtualization platforms such as VMware vSphere.

Arista EOS also simplifies tap aggregation with the Arista Data Analyzer (DANZ) feature set. For organizations with compliance requirements to aggregate and capture traffic, Arista EOS enables traffic collection at high data volumes with minimal infrastructure investment and without impacting network performance.

The Arista EOS CloudVision platform further enhances network visibility through a network-wide database approach. By consolidating the network state to a central database, the network operator can visualize the environment.

Sections in this chapter include:

18.1 Integration with CloudVision

In CloudVision's multi-switch tap aggregation, a datacenter network feeds taps into a layer of switches. These switches forward their traffic to an aggregation layer which subsequently sends traffic to tool ports. Thereby in CloudVision Portal (CVP), you can monitor and manage clusters of switches working in concert.

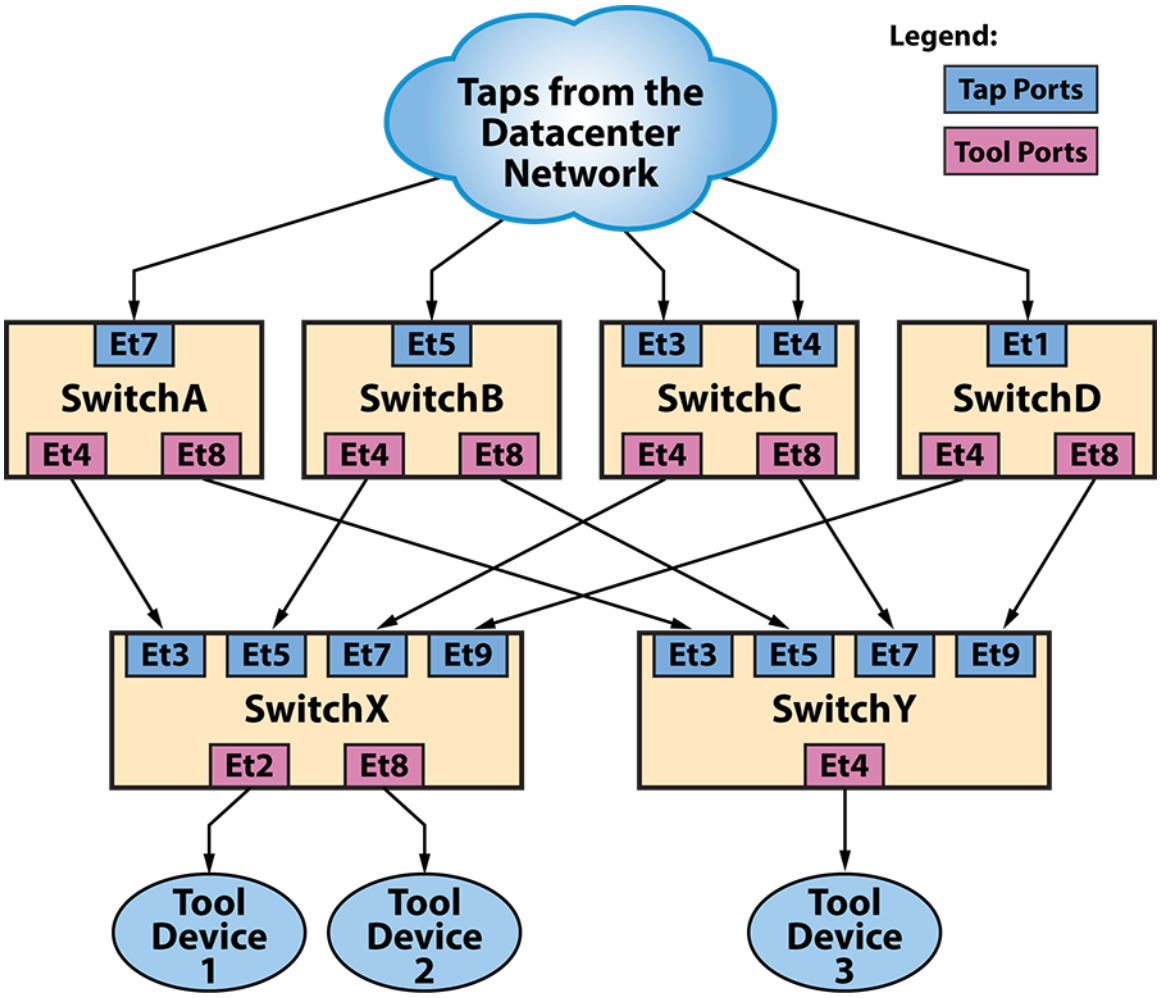

Figure 349: Multi-Switch Tap Aggregation Topology

CloudVision assigns a unique VLAN ID to each external tap port. It tags the traffic arriving on each external tap port with the appropriate VLAN ID and forwards it to each tool-facing device. The traffic arrived on the tool-facing switch passes through a large policy map that matches the VLAN ID of the packet and then sent to the default groups configured on the original tap port. Tool ports that are configured as members of that group receives the packet and forwards it to the external tool device.

You can access the tap aggregation screen for each switch. The CVP multi-switch tap aggregation provides the following functionalities:

- Configures an interface's switchport mode as either tap port or tool port.
- Configures default groups on an external tap port
- Configures the group membership on an external tool port
- Automatically manages policy-maps to correctly steer packets from external tap ports to external tool ports
- Provides built-in verification and reconciliation tools to ensure consistent and valid configuration in devices
- Instinctively monitors details of traffic throughput, interface status, and tap aggregation
- Integrates with CloudVision's other telemetry features including events, notifications, device and interface detail views, and metric comparisons.

18.1.1 Initial Setup for Multi-Switch Tap Aggregation

Initial setup for multi-switch tap aggregation includes the following tasks:

1. Prerequisites
2. Creating a Tap Aggregation Cluster
3. Setting Up Tap and Tool Devices
4. Configuring Internal Fabric

18.1.1.1 Prerequisites

The prerequisites to create a multi-switch tap aggregation cluster are provided below:

- CVP version 2019.1.0 and above
- Ensure that devices are:
  - In tap aggregation mode

    See the *Tap Aggregation Configuration* section in the *EOS Configuration Guide*.
  - Streaming via TerminAttr agent to a CVP node or cluster
  - Provisioned
  - Physically connected
  - Have Port-Channels configured (if they are being used)
- **Advanced login options for device provisioning and Multi-switch tap aggregation** options are enabled in CVP. See Enabling Multi-Switch Tap Aggregation.

Note: When prerequisite conditions are met, CVP displays the list the configured tap aggregation devices on the Tap Aggregation screen. See Figure 367: Initial Tap Aggregation Screen.

18.1.1.2 Creating a Tap Aggregation Cluster

Perform the following steps to create a tap aggregation cluster:

1. On CVP, click **Provisioning** > **Device Tags**.

   The system displays the Device Tags screen.

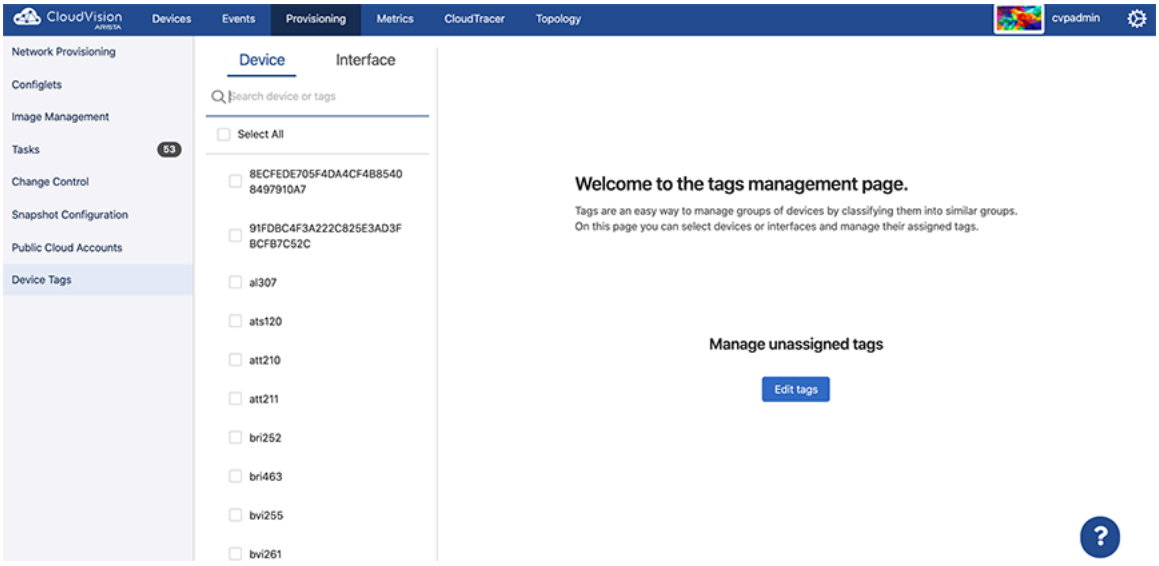

Figure 350: Device Tags Screen

Note: To assign tags to interfaces, click the **Interface** tab.

2. On the main panel, select device(s) of your tap aggregation cluster that you want to create a tag for.

   The system displays the **Assigned tags** panel.

Note:

- In general, tags should be of the form <*label*>: <*value*>.
- (Optional) Use the search bar for searching required devices.

3. Under **User Tags** > **Add or create tags**, type tapAggCluster: <*clusterName*> in the text box.

Note:

- To create and assign tap and tool tags, add tags of tapAggType: tap or tapAggType: tool to appropriate devices.
- The **System Tags** panel displays tags automatically created by CVP.

4. Click **Create and Assign**.

The new tag is displayed under **Manage assigned tags**.

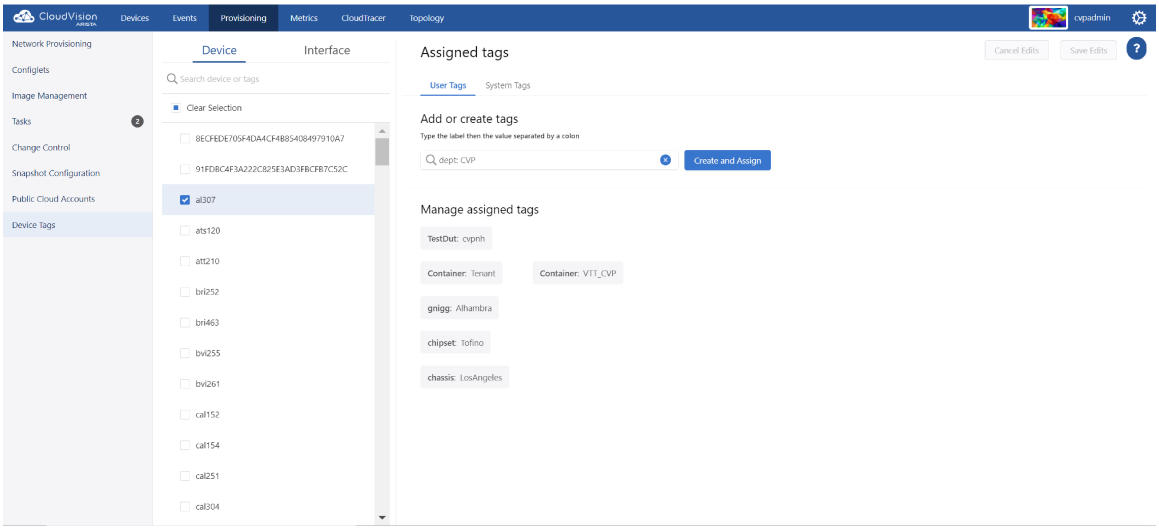

Figure 351: Create and Assign

Note: To delete a tag, click on the inessential tag > the minus sign > **Save edits**.

18.1.1.3 Setting Up Tap and Tool Devices

Devices are classified as either tap devices or tool devices by using tags with the tapAggType type. Perform the following steps to classify devices with ports:

1. On the CloudVision Portal, click **Provisioning** > **Device Tags**.
2. Click **Interface** to open the interface tags panel.
3. Select desired tap interfaces.

The system displays the **Assigned tags** panel.

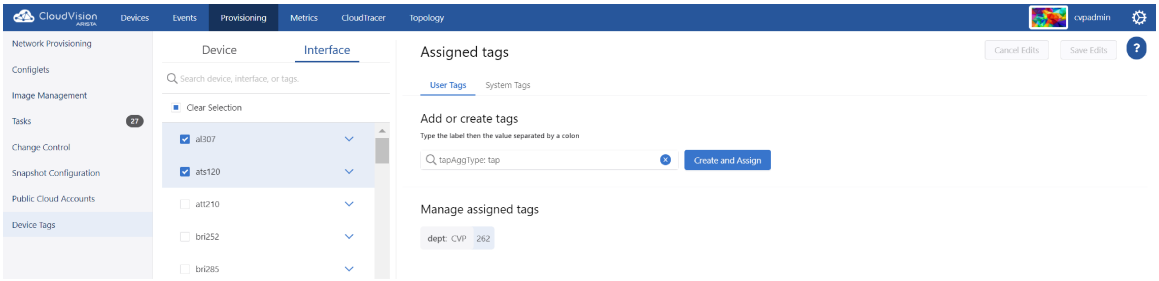

Figure 352: Assigned Tags Panel

4. Under **User Tags** > **Add or create tags**, type tapAggType: tap in the text box.
5. Click **Create and Assign**.
6. Select desired tool interfaces.

The system displays the **Assigned tags** panel.

7. Under **User Tags** > **Add or create tags**, type tapAggType: tool in the text box.
8. Click **Create and Assign**.

18.1.1.4 Configuring Internal Fabric

We must manually specify all connections between the devices in our tap aggregation cluster's internal fabric so that CVP can determine the cluster's topology, which will later be used for generating the cluster policy.

Perform the following steps to configure internal fabric:

1. On the CloudVision Portal, click **TapAgg**.

   The system displays the tap aggregation screen.

   Note: If you are configuring internal fabric for the first time, CVP displays the 'You do not have exactly one connection between each of your cluster devices. Update internal connections warning.

2. Select the desired cluster from the **Cluster** drop-down menu at the upper left corner.
3. Select Internal Fabric from the **Table** drop-down menu.

   The system displays the internal fabric screen.

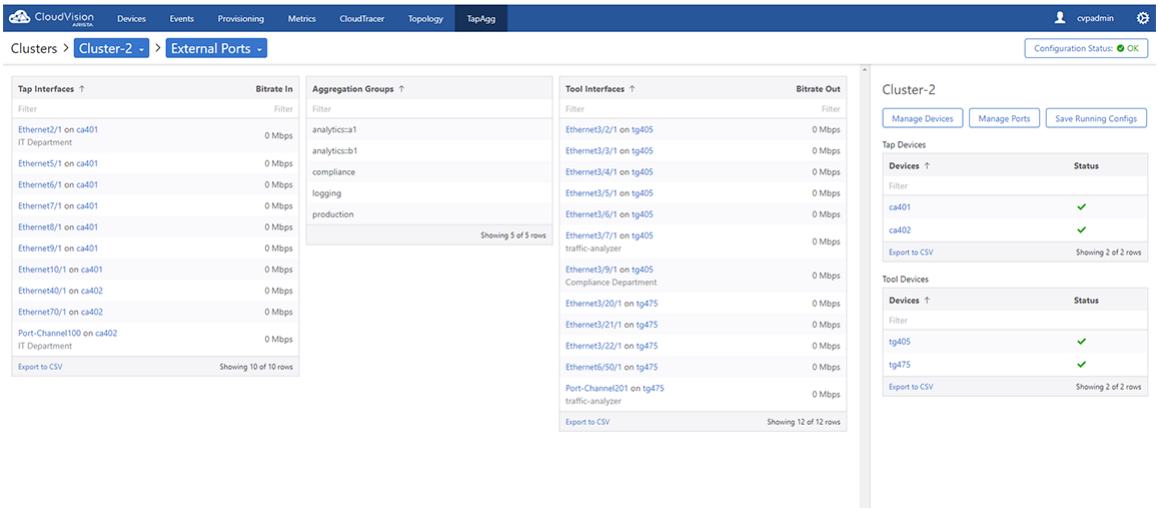

Figure 353: Internal Fabric Screen

4. Provide the following information in corresponding fields to add a connection:
   - Source Device
   - Source Interface
   - Destination Device
   - Destination Interface
5. Click **Add Connection**.

   The system automatically configures the source and destination interface as tool and tap ports respectively.

18.1.2 Accessing the Tap Aggregation Screen

The tap aggregation screen configures internal fabric and provides a summary of all ports and groups configured in tap aggregation clusters.

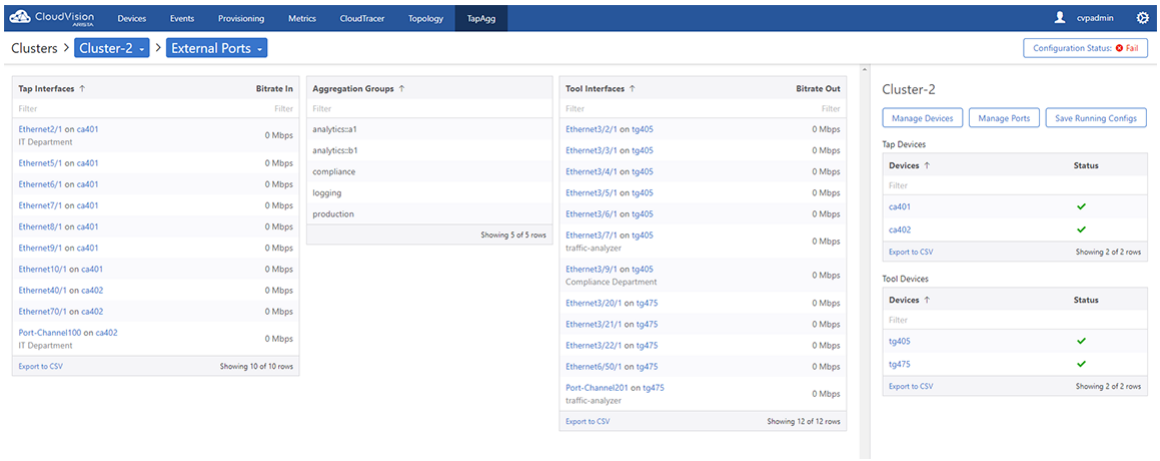

Figure 354: Tap Aggregation Screen

This screen provides the following information:

- **Cluster** drop-down menu - Select the desired cluster to switch among various tap aggregation clusters.
- **Table** menu - Select the desired table. Available options are:
  - External Ports - Manages external ports. See **External Ports Table Type**.
  - Group Table - Displays an overview of all groups created in the tap aggregation cluster.

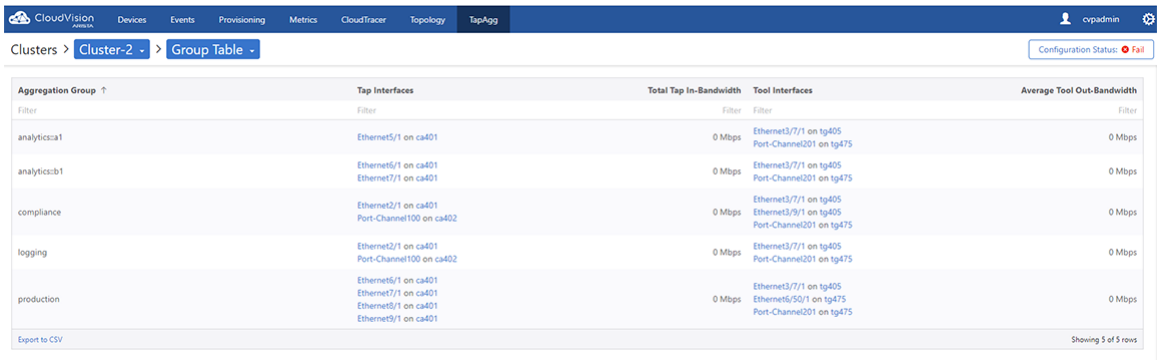

Figure 355: Groups Overview

  - Internal Fabric - Configures internal fabric. See Configuring Internal Fabric.
- **Tap Interfaces** column - Lists all configured tap ports.

  Note: Clicking on the interface link displays the **Interface Overview** screen. Clicking on the device link displays the **Device Overview** screen.
- **Bitrate In** column - The bitrate of incoming packets.
- **Aggregation Groups** column - Lists all aggregation groups.
- **Tool Interfaces** column - Lists all configured tap ports.

  Note: Clicking on the interface link displays the **Interface Overview** screen. Clicking on the device link displays the **Device Overview** screen.
- **Bitrate Out** column - The bitrate of outgoing packets.
- **Export to CSV** - Click to download the appropriate table contents to your local drive.

18.1.2.1 External Ports Table Type

Select External Ports from the **Table** drop-down menu to access the following functionalities:

- Cluster Management
- ACLs and Tap Ports Management

- Tool Ports Management
- Groups Management

18.1.2.1.1 Cluster Management

Cluster management includes the following functionalities:

- Adding and Removing Devices
- Managing Tap and Tool Ports
- Saving Running-Configuration
- Verifying Running-Configuration

Adding and Removing Devices

Click the **Manage Devices** button to open the Device Tags screen where you can add or remove devices from a cluster. See Assigning devices to a Tap Aggregation Cluster.

Managing Tap and Tool Ports

Click the **Manage Ports** button to open the Manage Ports pop-up window.

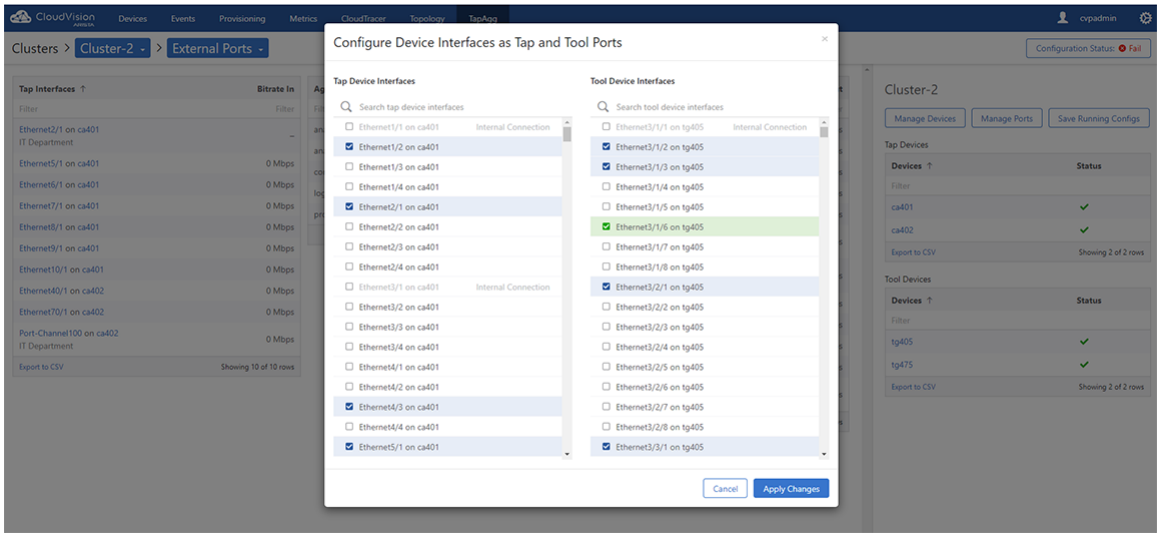

Figure 356: Manage Ports Pop-Up Window

This screen provides the following functionalities:

- View all current tap and tool ports
- Add or remove multiple tap and tool ports

Note: Click **Apply Changes** to save configuration changes.

Saving Running-Configuration

Click the **Save Running Configs** button to save the running-configuration of all devices in the cluster as startup configuration.

The system displays the Save Running Configs pop-up window.

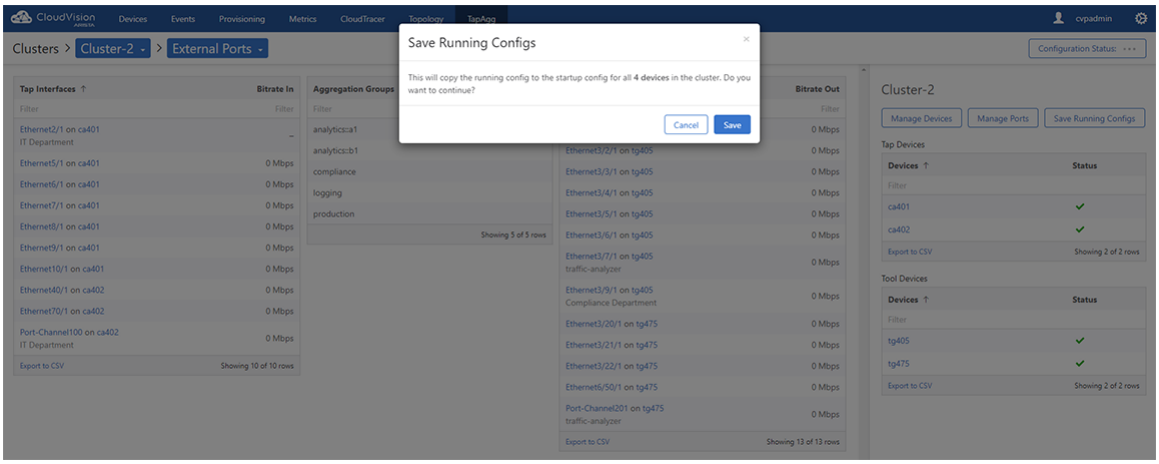

Figure 357: Save Running Configs Pop-Up Window

Note: Click **Save** to confirm running-configuration changes.

Verifying Running-Configuration

Click the **Configuration Status** button to verify that all devices in the cluster are configured correctly.

The system displays the Verify Running Configs screen which lists verification results of each rule that the application checks for. Click **Verify Configuration** to verify all current configurations.

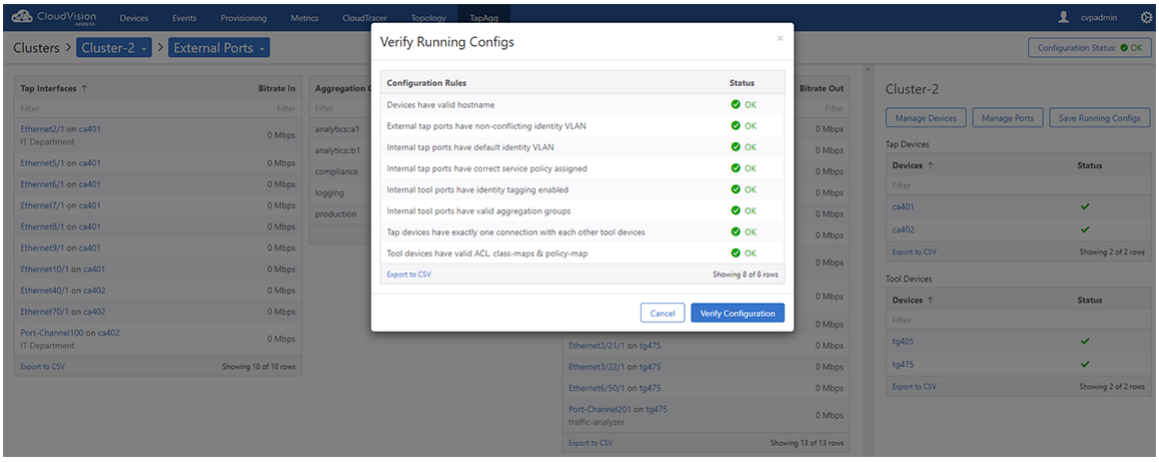

Figure 358: Verify Running Configs Pop-Up Window

In case of an error, click **Fix Configuration** to resolve the configuration error(s).

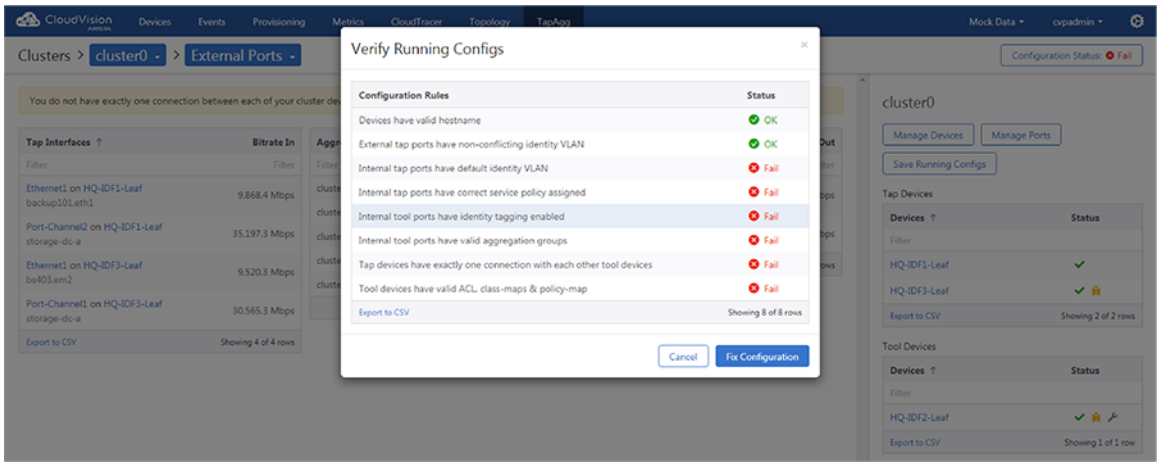

Figure 359: Running configuration errors

The system computes all commands required to fix the current configuration and applies the correct configuration on devices in the Tap Aggregation cluster.

Note: Click **Export to CSV** to download the table in csv format to your local drive.

18.1.2.1.2 ACLs and Tap Ports Management

Perform the following steps to manage ACLs and Tap Ports:

1. Select a tap port by clicking on a row in the Tap Interfaces table.
   The system displays the appropriate tap port's configuration and metrics in the right panel.

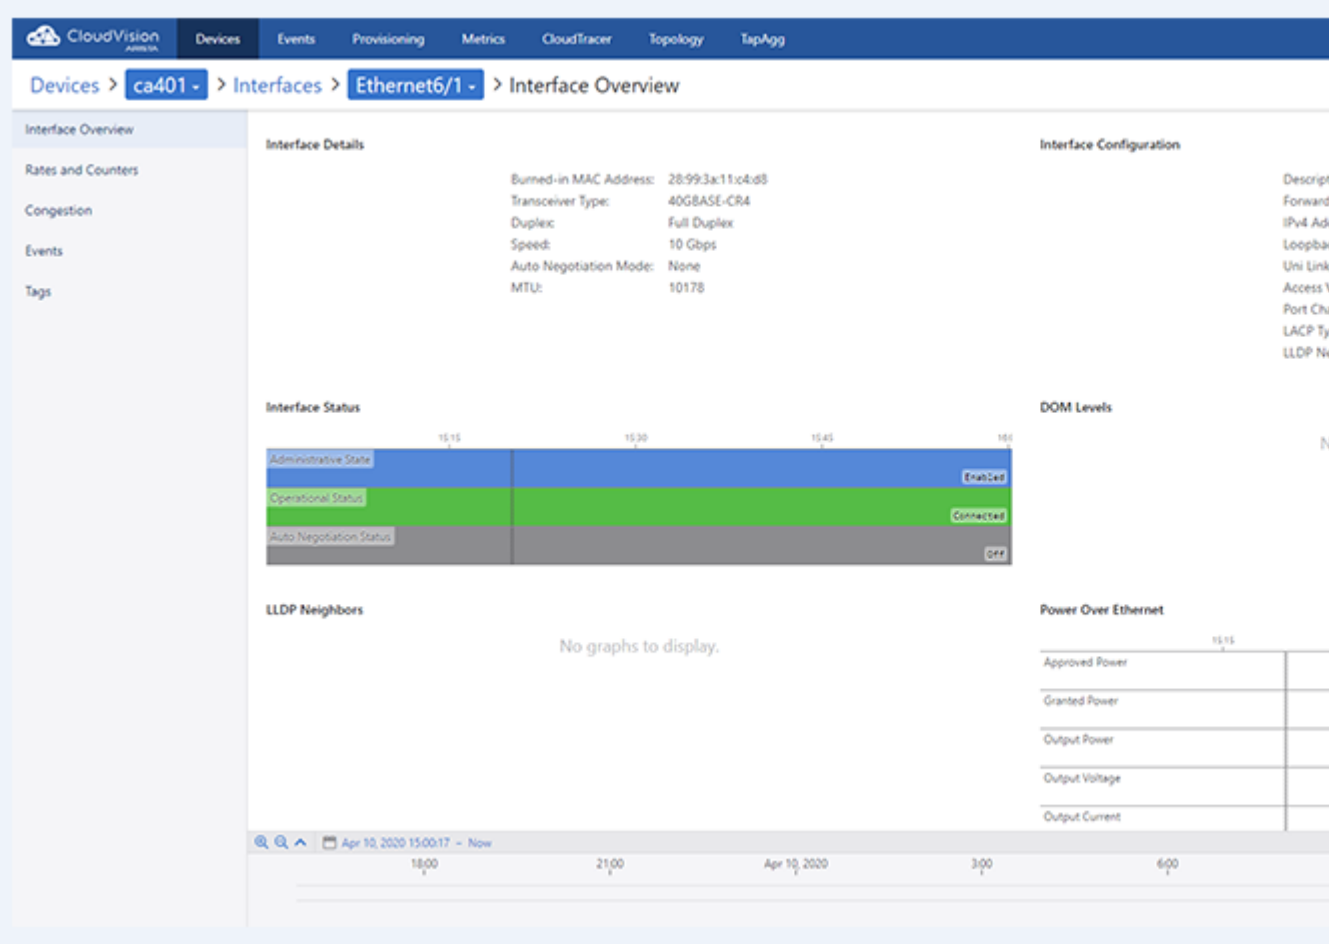

Figure 360: Tap Port's Configuration and Metrics Panel

2. On the right panel, perform the following steps to execute specified functionalities:
   - Creating an ACL
   - Modifying an ACL
   - Modifying Traffic Steering
   - Modifying Default Groups

Creating an ACL

1. Click the **+ Add Match Statement** button.
   The system displays a **Match Statement Card #1** pane.
2. Select Create ACL from the **Match ACL** drop-down menu.
   The system displays the Create ACL pop-up window.

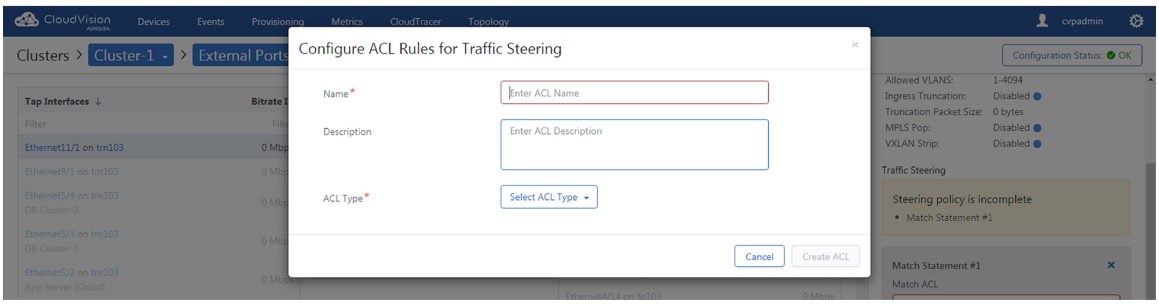

Figure 361: Create ACL Pop-Up Window

3. Provide the required information in the corresponding entities:

- Name
- Description
- ACL Type

4. Click **Create ACL**.
   The system confirms when configuration changes are applied successfully.

Modifying an ACL

1. Click the **Add Match Statement** button.
   The system displays a **Match Statement Card #1** pane.
2. Select the edit icon next to the required ACL from the **Match ACL** drop-down menu.
   The system displays the Manage ACL pop-up window.

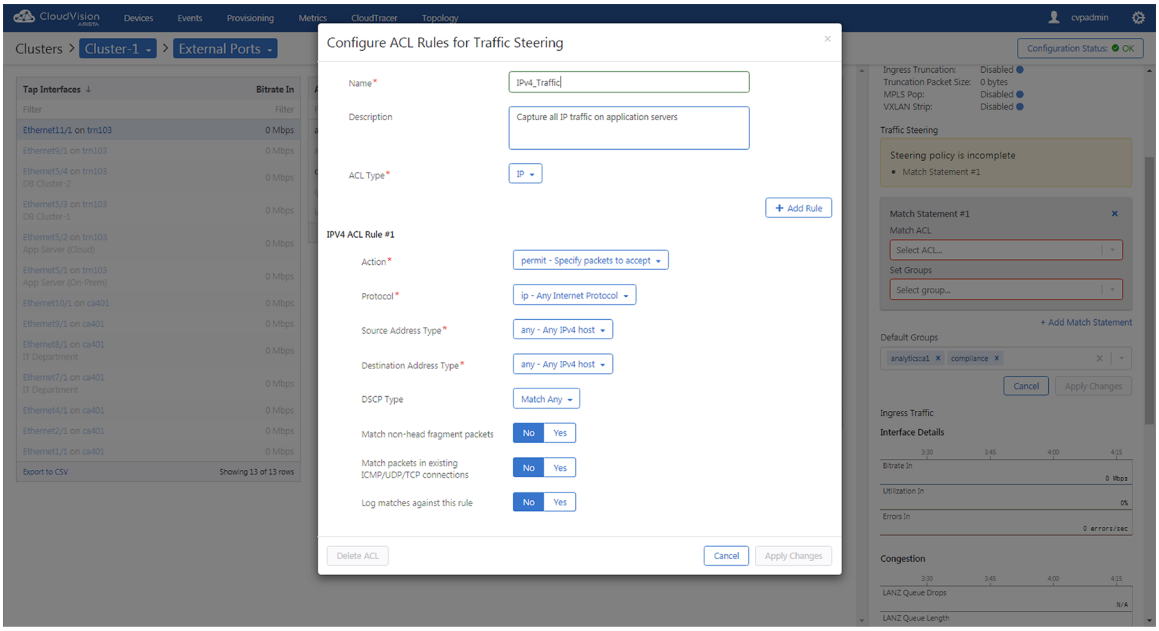

Figure 362: Manage ACL Pop-Up Window

3. Update required changes.
4. Click **Apply Changes** to confirm updated changes.

Note: Click **Delete ACL** to delete the appropriate ACL.

Modifying Traffic Steering

1. Click the **Add Match Statement** button.
   The system displays a **Match Statement Card #1** pane.
2. Select the required options from **Match ACL and Set Groups** drop-down menu.
3. Click **Apply Changes**.

Modifying Default Groups

Select required group(s) from the multi-purpose **Default Groups** widget.

Tool Port Management

Perform the following steps to add or remove groups from the tool port:

1. Select a tool port by clicking on a row in the **Tool Interfaces** table.
   The system displays the appropriate tool port s configuration and metrics in the right panel.

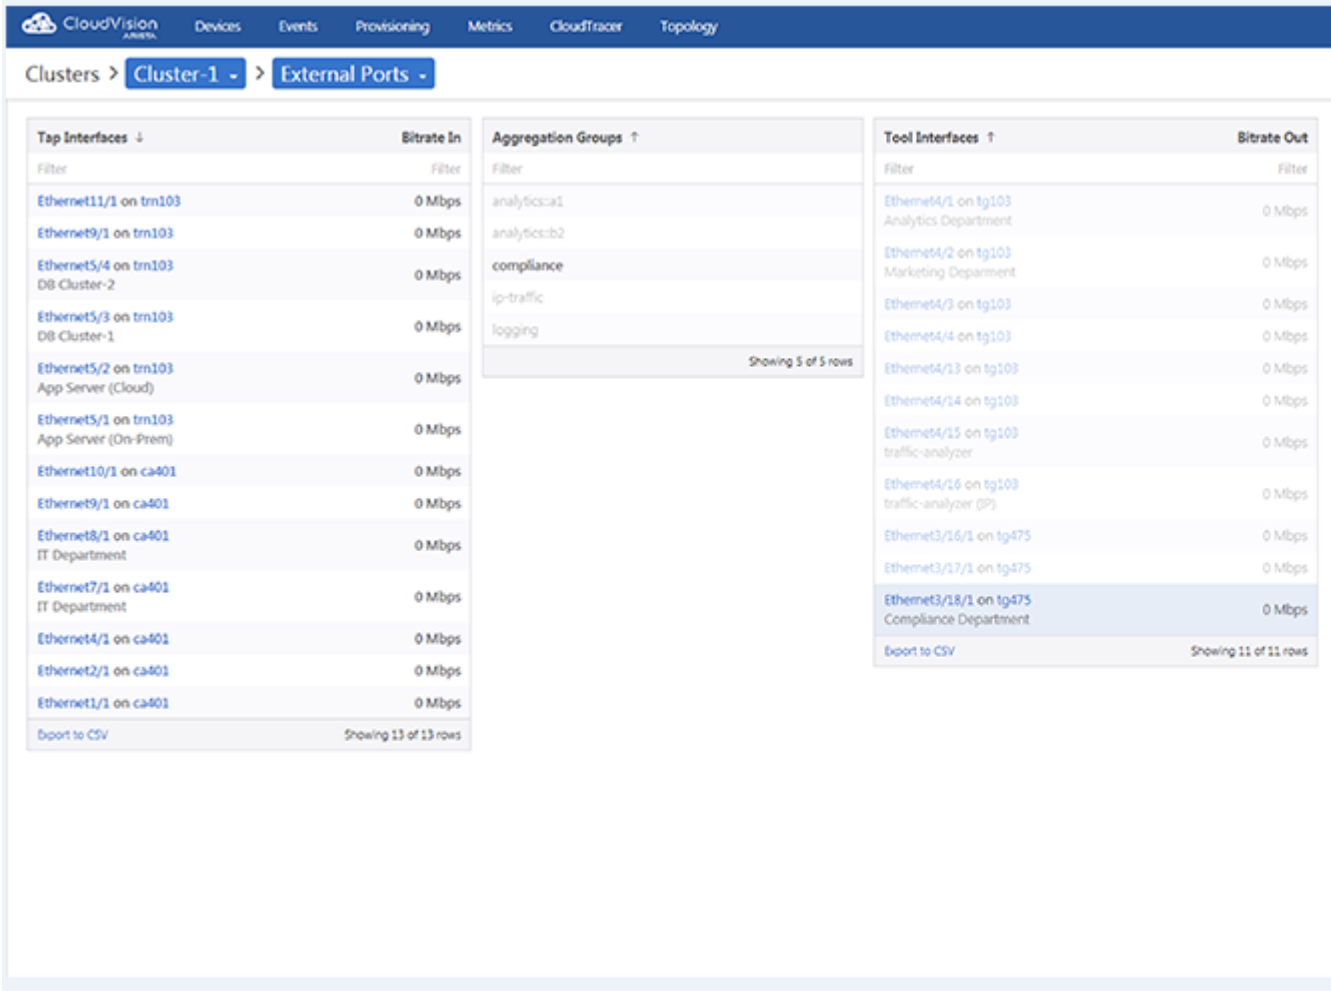

Figure 363: Tool Port's Configuration and Metrics Panel

2. Select required group(s) from the multi-select **Group Membership** drop-down menu.
3. Click **Apply Changes**.

Groups Management

Select the required port from either Tap Interfaces or Tool Interfaces pane to initiate the following functionalities in the right panel:

- Creating Group Membership
- Groups Management

Creating Group Membership

Perform the following steps to create group membership:

1. Type the new aggregation group name in the multi-purpose Default Groups widget.
   The system displays the Create group “group_name” option. See Figure 364: Create Group Option.

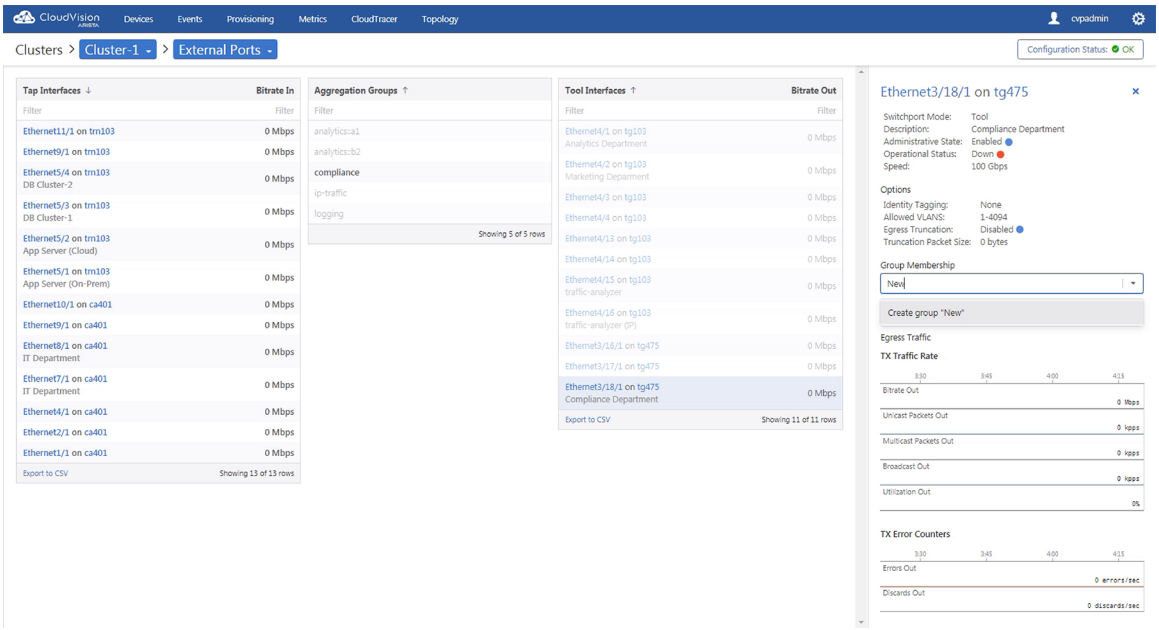

Figure 364: Create Group Option

2. Click the **Create group "*group_name*"** option.
3. Click **Apply Changes**.

   The system creates the aggregation group and applies it on the selected port.

Modifying Group Membership

Perform the following steps to modify group membership:

1. Select the required group from the Aggregation Groups pane.

   The system displays the appropriate group's configuration and metrics in the right panel.
2. Click the Modify Membership button.

   The system displays the Manage Group Membership pop-up window.

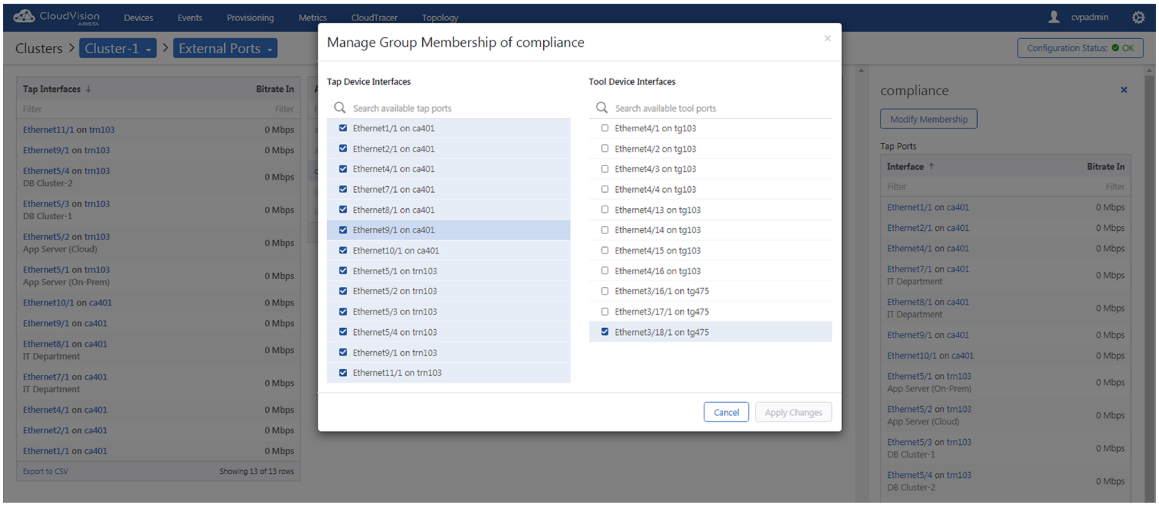

Figure 365: Manage Group Membership Pop-Up Window

3. Choose required ports.
4. Click Apply Changes.

Note: The system configures selected ports and deconfigures unselected ports that were previously selected.

18.2 Enabling Multi-Switch Tap Aggregation

Perform the following steps if you do not find the **TapAgg** tab on the CVP screen:

1. Click the gear icon at the upper right corner of the screen.
   The browser displays the Settings screen.
2. Under the Beta Features pane, enable **Multi-switch tap aggregation** using the toggle button. See Enabling Multi-Switch Tap Aggregation.

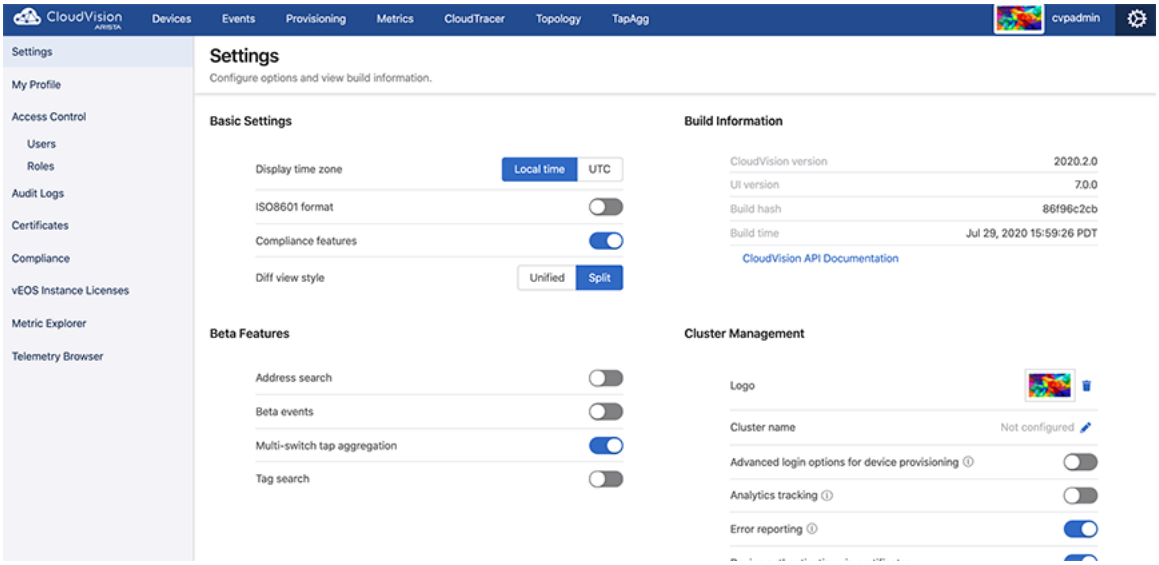

Figure 366: Enable Multi-Switch Tap Aggregation

Note: We recommend to enable **Advanced login options for device provisioning** under the **Cluster Management** pane. This performs configuration changes over the connection between CVP and the device's TerminAttr agent.

18.3 Configuring Tap Aggregation Devices

CVP enables you to select and configure devices for tap aggregation. When you configure a device, you specify the tap aggregation interfaces, aggregation groups, and tool interfaces. You can also view the running configuration on the device and the differences between the designed configuration and running configuration.

You use the tap aggregation screen to select the device for configuration, and the **Tap Aggregation Manager** to configure the device.

Complete these steps to configure a device:

1. Go to the tap aggregation screen.

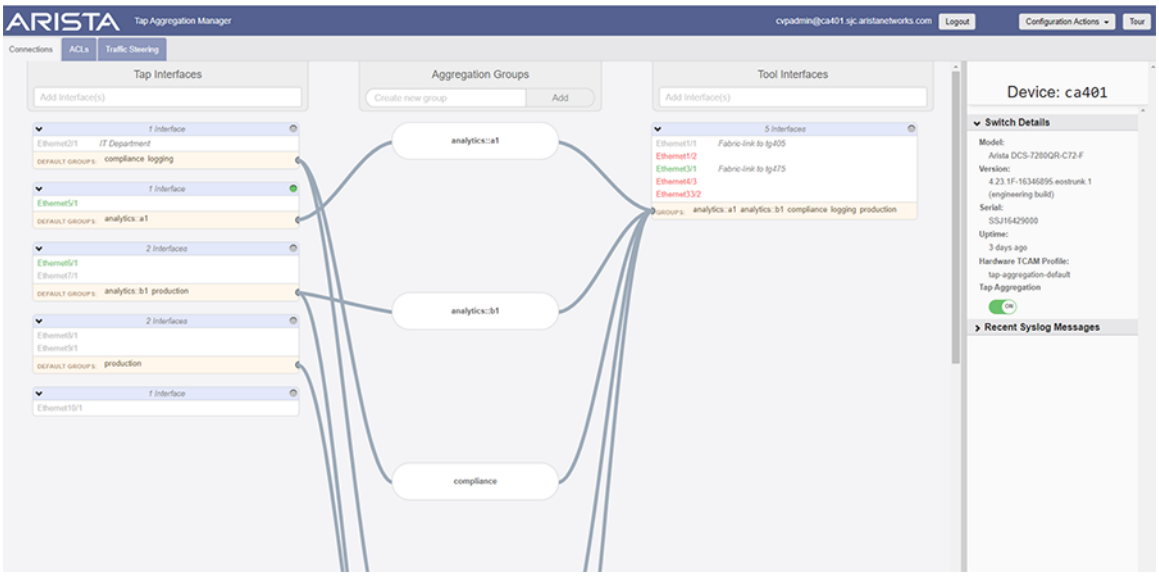

Figure 367: Initial Tap Aggregation Screen

2. Click the pop-out icon of device you want to configure.

Note: In case of a huge list, search for the device using the **Filter** search box.

The Tap Aggregation Manager appears for the device you selected.

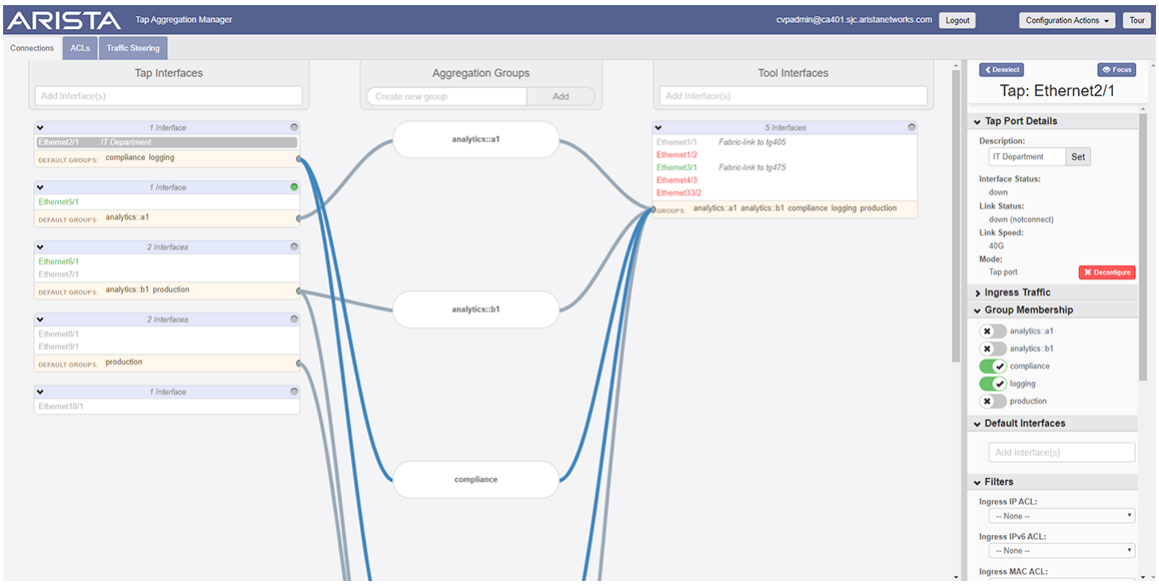

Figure 368: Tap Aggregation Manager for Selected Device

3. Specify the tap aggregation interfaces, aggregation groups, and tool interfaces as needed.
4. (Optional) To view the running configuration for the device, click the **Running Config** button.
5. Click **Save** to save the configuration for the device.

Chapter 19

Using Snapshots to Monitor Devices

CloudVision enables you to monitor changes in the state of the devices in your network over time through the use of snapshots.

Note: Starting from *2018.2.0* release, snapshots UI is available as part of the **Device View** in **Telemetry**.

Sections in this chapter include:

19.1 About Snapshots

In CloudVision, the snapshot service runs as a scheduler to capture device snapshots periodically.

The information recorded in snapshots provides you with insights on the configuration, EOS image, and other aspects of the device. Snapshots are captured for individual devices (single switches) only.

19.2 Standard Information in Snapshots

The information recorded in the snapshot reflects the state of the device at the time snapshot was captured. A snapshot only contains outputs of custom commands that are part of a snapshot template. (You must select a snapshot template when you capture a snapshot.) See Defining Custom Snapshot Templates and Editing Custom Snapshot Templates for information on using snapshot templates.

When upgrading to the *2018.2* train, only snapshot templates are migrated but not previous snapshots. CloudVision stores migrated templates without any device list associated with them. Hence, they are marked as unscheduled. However, these templates can be used to capture snapshots before and after change controls.

19.3 How to Use Snapshots

In CloudVision, snapshot service schedules and periodically captures the outputs of commands that are specified in the template. The frequency of capturing command outputs is based on the scheduling frequency mentioned in the snapshot template. The information recorded in snapshots can provide you with insights on the configuration, EOS image, and other aspects of the device. Snapshots are captured for individual devices (single switches) only.

The main uses of snapshots are:

- Viewing snapshots to understand the state of a device at a given time, or over time.
- Comparing snapshots to see the change in state of a device between two points in time.
- Comparing snapshots to see the state of a device before and after a change control.

19.4 Accessing Snapshots

Snapshots are stored under the CVP dataset, which you can access any time for detailed analysis. The Snapshots page displays all valid snapshots created over time. Each valid snapshot provides the following additional information:

- **Name** - The name of the template (you assign the name when you create the template).
- **Capture Time** - The date and time when the snapshot was last captured.
- **Last Executed By** - The user that captured the snapshot.

It also allows navigating to snapshots of the corresponding snapshot template.

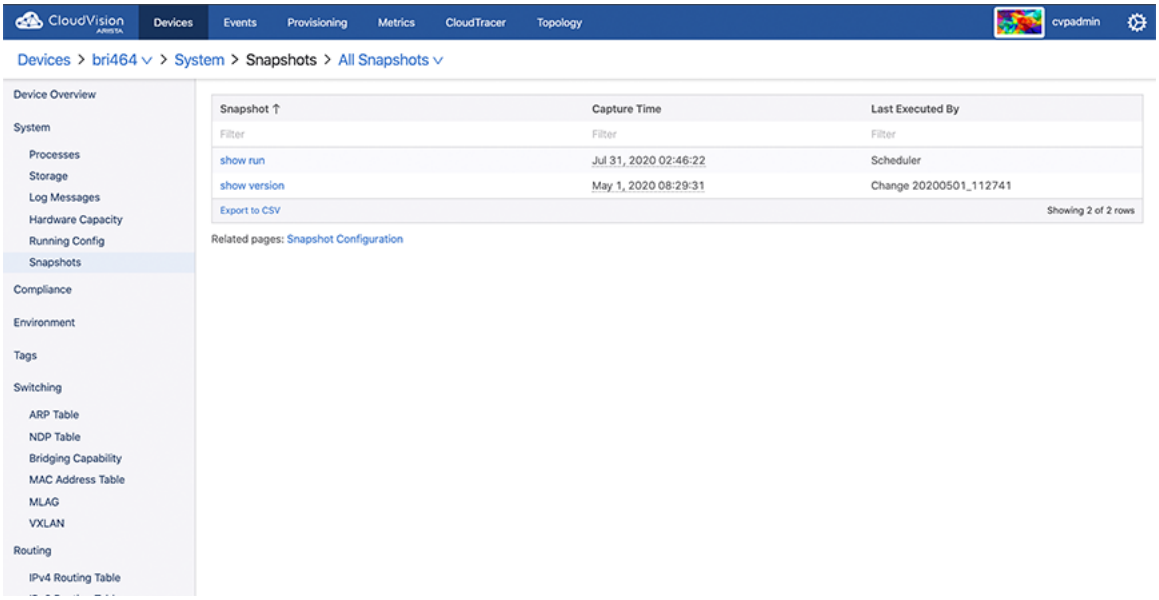

Figure 369: Snapshots Page

You can navigate to the Snapshots page through one of the following paths:

- **Inventory > Device_ID > Snapshots**
- **Network Provisioning** > Right-click on the required device > **Snapshot**.

19.5 Accessing Snapshot Configurations

The Snapshot Configuration page displays all snapshot templates created over time. It further allows you to edit current snapshot configuration, navigate to the Snapshots page, view the status of each snapshot configuration, and create a new custom snapshot configuration.

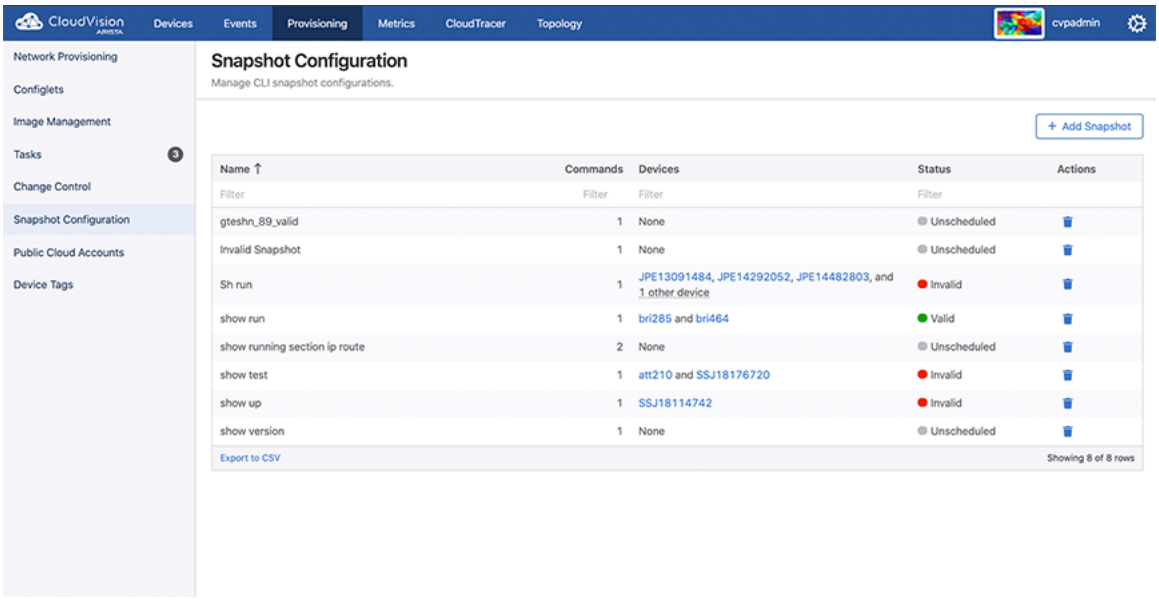

Figure 370: Snapshot Configuration Page

You can navigate to the Snapshot Configuration page through one of the following paths:

- **Inventory > *Device_ID* > Snapshots > Snapshot Configuration**
- **Network Provisioning** > Right-click on the required device > **Snapshot** > **Snapshot Configuration**.

19.6 Defining Custom Snapshot Templates

To ensure that snapshots contain the information you need for effectively monitoring changes in the state of devices over a certain period of time, CloudVision allows you to define custom snapshot templates.

A snapshot template defines commands, outputs of which need to be captured as part of the snapshot using that template. When you create a snapshot template, associate a list of devices, and set an execution frequency with it, the snapshot service starts capturing and storing snapshots for that template based on the scheduled frequency.

Complete the following steps to define a new custom snapshot template:

1. Navigate to **Inventory > *Device_ID* > Snapshots > Snapshot Configuration**.

   The Snapshot Configuration page displays currently available snapshot templates.
2. Click the **(or create a new configuration)** hyperlink at the lower right side of the page.

   The **Snapshot Configuration** page displays the **Add Snapshot Configuration** section.

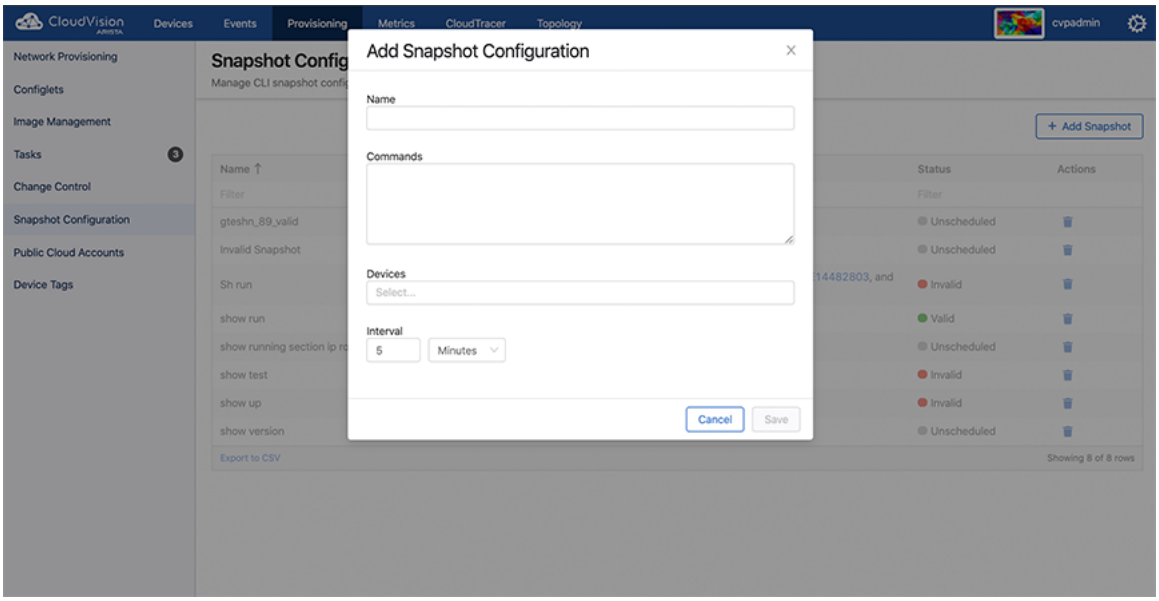

Figure 371: Add Snapshot Configuration Section

3. In the **Name** field, type the name of the custom snapshot template.
4. In the **Commands** field, enter the EOS CLI commands to be executed by the snapshot.
5. If necessary, click the **Devices** drop-down and select required devices.
6. Under **Interval**, Specify the frequency for capturing snapshots in either minutes, hours, or days.
7. Click **Save**.

   The Snapshot Configuration page immediately displays the latest configuration along with the list of current configurations.

> **Note:** A snapshot configuration that is created without a device is saved and marked as unscheduled. Snapshot templates with bash commands are marked as invalid. However, these unscheduled and invalid templates can still be selected while creating a Change Control to capture pre and post change control snapshots.

19.7 Editing Custom Snapshot Templates

Complete the following steps to go to defined templates:

1. Navigate to **Inventory > *Device_ID* > Snapshots > Snapshot Configuration**.

   The Snapshot Configuration page displays currently available snapshot templates.
2. Click the snapshot name for editing the corresponding snapshot template..

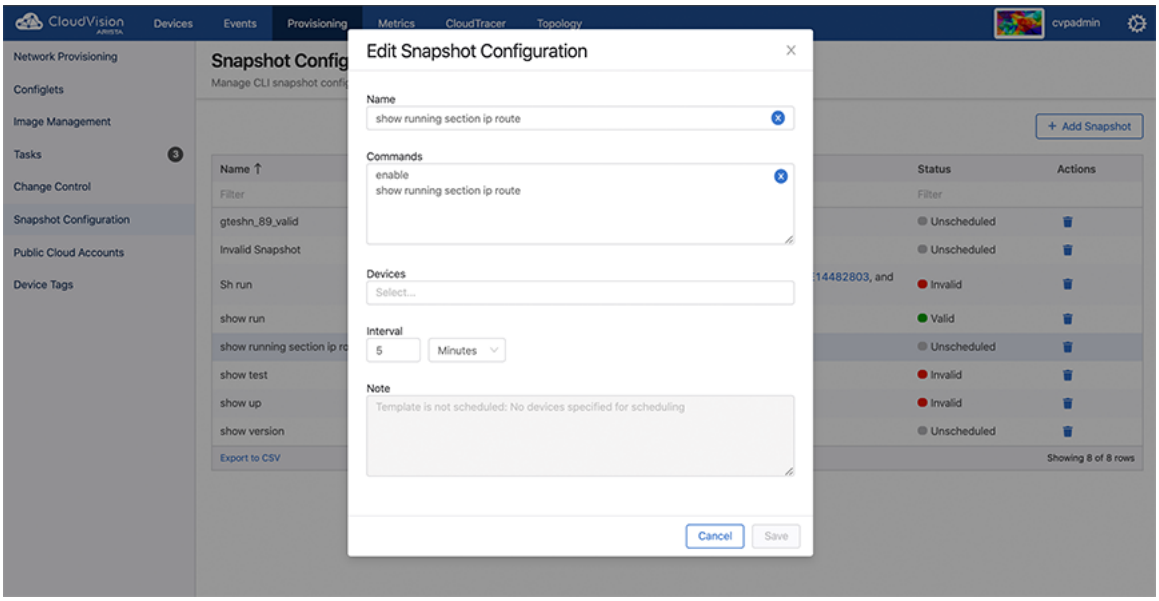

Figure 372: Edit Snapshot Configuration Section

3. Modify the required information in corresponding fields.
4. Click **Save**.

19.8 Viewing Snapshots Differences

You can take snapshots of single devices only. The exact set of information and presentation of the information in the snapshot is determined by the snapshot template you choose when capturing the snapshot.

Complete the following steps to view snapshots of a device:

1. Go to the **Network Provisioning** page.
2. Locate the device for which you want to view snapshots.
3. Right-click on the device icon, then click **Snapshot**.

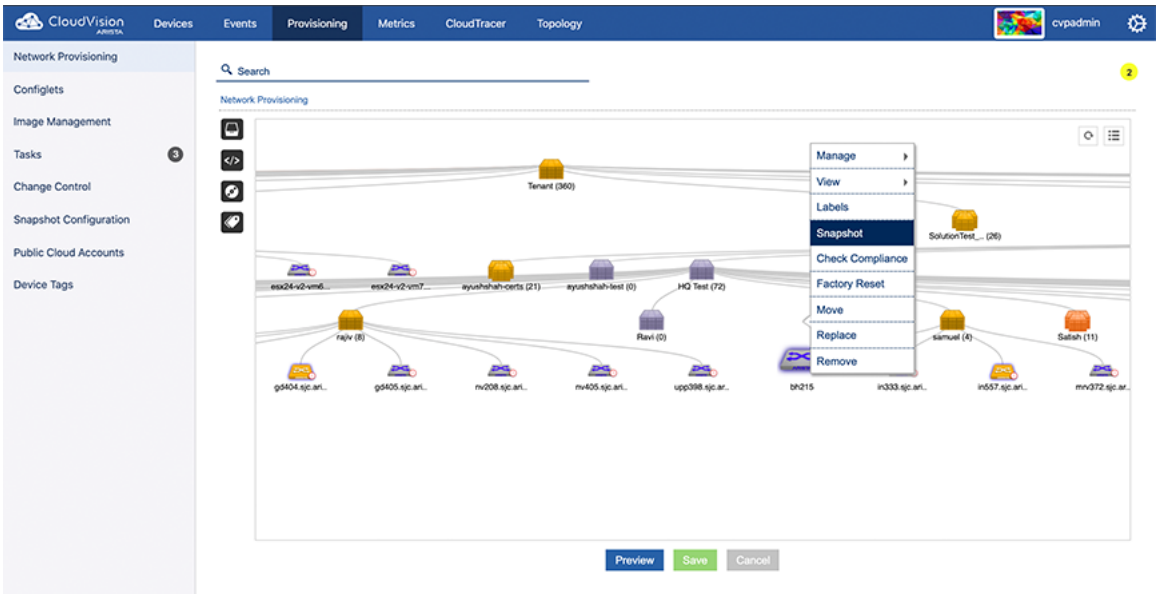

Figure 373: Initiate Viewing Snapshot

The **All Snapshots** page displays all valid snapshots.

Note:

You can also navigate to the **All Snapshots** page through **Telemetry > Devices > *Device_ID* > Snapshots**.

4. Click on the snapshot template name for viewing the corresponding snapshot.

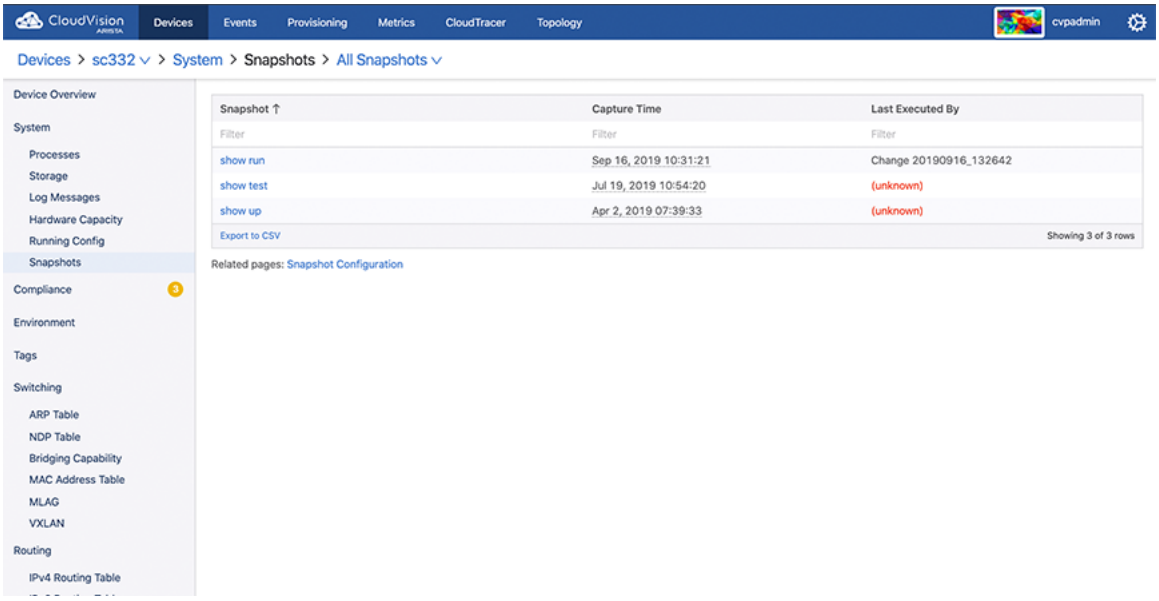

Figure 374: All Snapshots Page

5. Click the date and time breadcrumb for viewing all snapshots of the corresponding template.

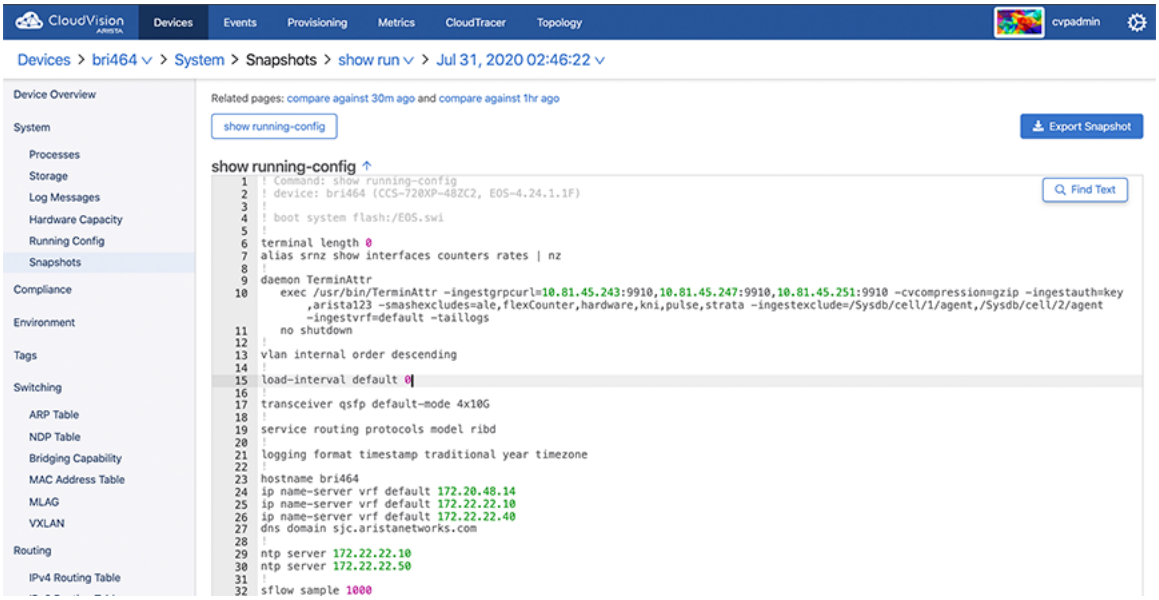

Figure 375: View All Snapshots

6. Click the required snapshot to view the corresponding output.

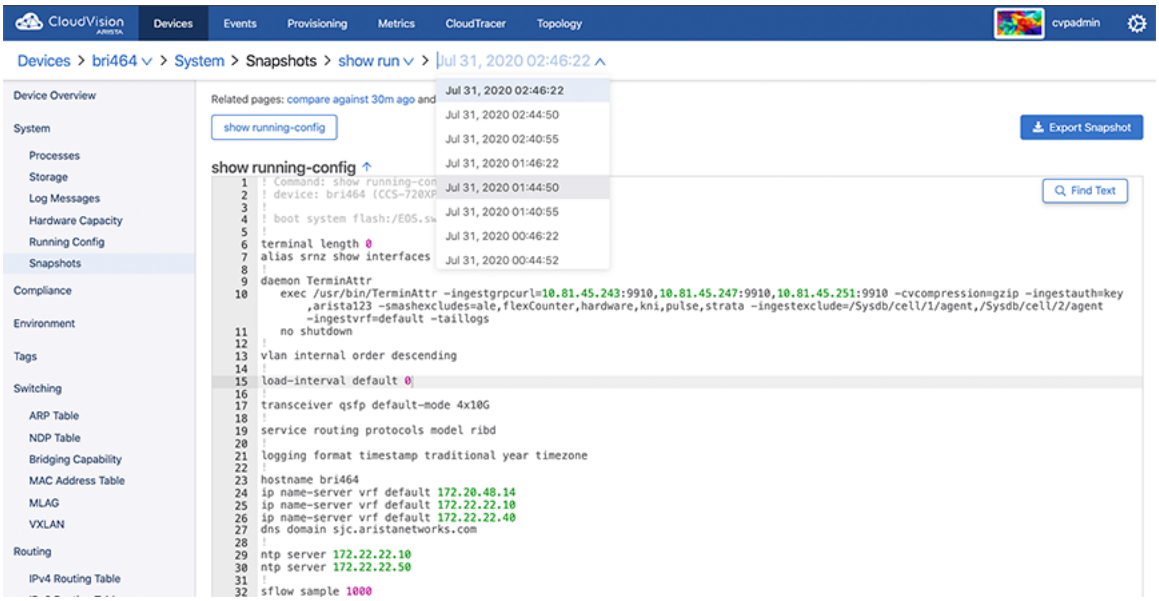

Figure 376: Select Snapshot

7. Click Compare against a previous time for viewing corresponding snapshot differences.
8. The page displays corresponding snapshot differences.

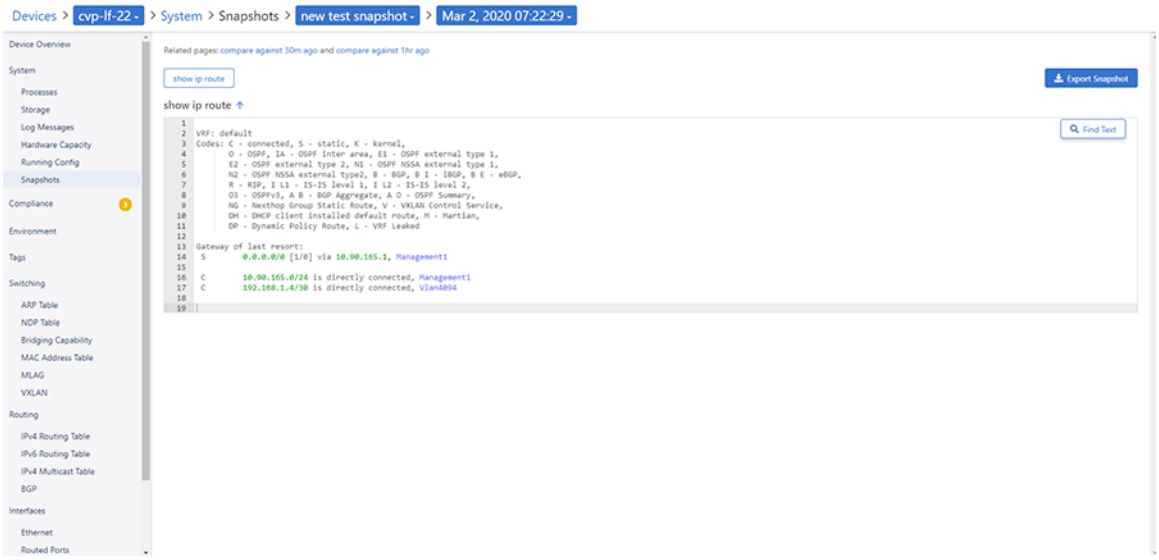

Figure 377: Compare Snapshots

Note: Snapshot differences are displayed in color codes to quickly identify significant changes in the state of the device over time. Click the Split tab for viewing snapshot differences in different windows.

Chapter 20

Backup & Restore, Upgrades, DNS NTP Server Migration

This document provides details on how to perform backup and restore operations and upgrading CloudVision Portal (CVP).

20.1 Backup and Restore

CloudVision Portal (CVP) enables you to backup and restore the complete CVP provisioning dataset, including containers, devices, configlets, images, and configlet / image assignments. You can use commands to backup and restore CVP data.

Arista provides a simple script at /cvpi/tools/backup.py which is scheduled by default to run daily to backup CVP data, and retain the last 5 backups in /data/cvpbackup/. Backing up and restoring data saves information about the CVP instance to a tgz file, and then restores the information from the tgz file to a new CVP instance. The CVP commands provide all of the functionality required to complete backup and restore operations.

Note: It is a good practice to regularly create and export backups to ensure that you have an adequate supply of backup files available to you that you can use to restore CVP data.

Note: There is no backup or restore of the Telemetry analytics dataset.

The current CVP release does not support restoring backups taken from previous CVP releases. If you would like to restore a backup from a previous CVP release, install the previous release, restore the backup, and then upgrade to the current release. After you have successfully upgraded to the current release, take another backup so that you can directly restore that into current main release in the future.

For more information, see:

20.1.1 Requirements for Multi-node Installations

The basic requirements for backup and restore operations are the same for single-node installations and multi-node installations.

20.1.2 Using CVPI Commands to Backup and Restore CVW Data

Arista recommends to back up wifimanager regularly and especially before performing any upgrades.

20.1.2.1 Restore CVW Data

You can restore wifimanager from a backup using the `cvpi restore wifimanager </path/to/backup/file>` command.

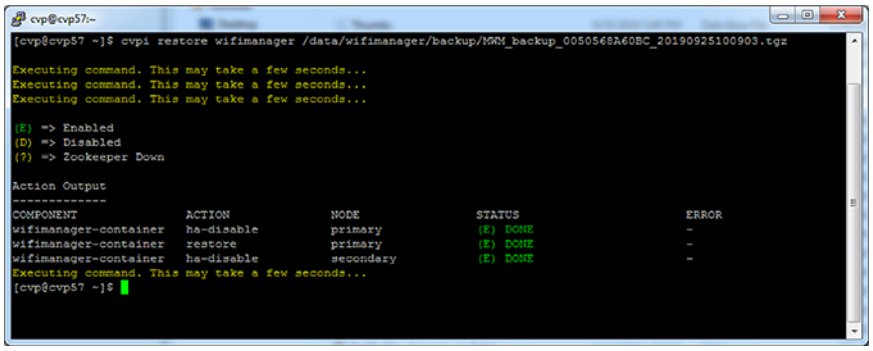

Figure 378: Restore CVW Data

Note: For a CV cluster, you can run this command only on the primary node. If no backup was carried out before the upgrade, you can use a scheduled backup under the `/data/wifimanager/data/data/backup` directory to restore wifimanager.

20.1.2.2 RMA

For RMA or recovery issues, contact support-wifi@arista.com.

Note: Back up wifimanager on any node before submitting it for an RMA. When the node is re-deployed post-RMA, you can restore earlier wifimanager data from a backup that you have stored elsewhere.

20.1.3 Using CVPI Commands to Backup and Restore CVP Provisioning Data

Backup and restore are CVPI functionalities of CVPI components.

Note:

The default directory to save and restore backup data files is `/data/cvpbackup`.

The default directory for backup/restore log files is `/cvpi/logs/cvpbackup`.

The default directory for temporary files during backup/restore is `/data/tmp/cvpbackup`.

The following commands are used to backup and then restore the containers, devices, configlets, images, and configlet or image assignments that are defined in CVP.

Note: When restoring devices, use the username and password that can access the devices being registered.

20.1.3.1 Backup CVP Provisioning Data

Use the `cvpi backup` command for saving a copy of CVP data as backup.

```
cvpi backup cvp
```

Note: To check the progress of the backup, read the latest `backup_cvp.*.log file in /cvpi/logs/cvpbackup`.

This command creates the backup files for the CVP component.

```
[cvp@cvp108 bin]$ cvpi backup cvp
```

20.1.3.2 Restore CVP Provisioning Data

Use the `cvpi restore` command to restore backup files for the CVP component.

```
cvpi restore cvp cvp.timestamp.tgz eosimages.timestamp.tgz
```

The `cvp.<timestamp>.tgz` parameter contains provisioning data from the DataBase (DB) of the CVP application. The `cvp.eosimages.<timestamp>.tgz` parameter contains EOS images and extensions stored in the DataBase (DB) of the CVP application.

Note: To check the progress of the restore, read the latest `restore_cvp.*.log file in /cvpi/logs/cvpbackup`.

This command restores the backup files of the CVP component.

```
[cvp@cvp108 bin]$ cvpi restore cvp cvp.2019.1.0.tgz cvp.eosimages
.2019.1.0.tgz
```

Note:

To check the progress of the backup, tail `-f/cvpi/logs/cvpbackup/backup_cvp.20190606020011.log`.

CVP backup creates two backup files in the /data/cvpbackup directory for restoration. The `eosimages.tgz` is generated only when it differs from the currently available copy of the `eosimages.tgz`, and is an optional parameter for restore if the CVP system already contains the same EOS image.

The `cvpi backup` command can be run anytime and does not disrupt the cvp application. However, the `cvpi restore` command will stop the cvp application and disrupt the service for the duration of the restore. If the restore is from a backup on a different CVP system to a new CVP system, it may also be required to on-board the EOS devices or restart the Terminattr daemons on the EOS devices after the restore.

20.1.3.2.1 Troubleshooting CVP Restore Failure of Provisioning Data

If the cvpbackup directory does not exist in /data when copying the restore files to a newly built VM, you must create it and assign the ownership to the cvp user and group in either of the following two ways:

- Login as cvp user and create the cvpbackup directory

  Use the `su cvp` command to login as cvp user and the `mkdir -p /data/cvpbackup` command to create the cvpbackup directory.
- Create the folder as root and change the ownership

  Use the `mkdir -p /data/cvpbackup` command to create the folder as root and the `chown -R cvp:cvp /data/cvpbackup/` command to change the ownership of cvpbackup directory and its files to cvp user and group.

Verifying the Ownership of cvpbackup Directory

Use one of the following commands to verify the ownership of cvpbackup directory:

- **ls**

  This example verifies the ownership of cvpbackup directory using the `ls` command.

```
[root@cvp-2019 data]# ls -l /data/ | grep cvpbackup
drwxrwxr-x. 2 cvp cvp 236 Mar 16 02:01 cvpbackup
```

- **stat**

  This example verifies the ownership of cvpbackup directory using the `stat` command.

```
[root@cvp-2019 data]# stat /data/cvpbackup/ | grep Access
Access: (0775/drwxrwxr-x) Uid: (10010/ cvp) Gid: (10010/ cvp)
```

Verifying the Ownership of Files Inside the cvpbackup Directory

The following example verifies the ownership of files inside the cvpbackup directory using the `ls` command:

```
[root@cvp-2019 data]# ls -l /data/cvpbackup
total 18863972
-rw-rw-r-- 1 cvp cvp 6650171 Mar 14 02:01 cvp.20200314020004.tgz
-rw-rw-r-- 1 cvp cvp 9642441292 Mar 14 02:08 cvp.eosimages.202003140200
02.tgz
```

Correcting the Ownership of cvpbackup Directory Files

Use the chown command to correct the ownership of cvpbackup directory files.

```
chown cvp:cvp cvp.<timestamp>.tgz cvp.eosimages.<timestamp>.tgz
```

The `cvp.<timestamp>.tgz` parameter contains provisioning data from the DataBase (DB) of the CVP application. The `cvp.eosimages.<timestamp>.tgz` parameter contains EOS images and extensions stored in the DataBase (DB) of the CVP application.

This example changes the ownership of all cvpbackup directory files.

```
[root@cvp-2019 data]# chown cvp:cvp cvp.20200319020002.tgz cvp.eosimages
.20200314020002.tgz
```

20.2 Upgrading CloudVision Portal (CVP)

Similar to Arista EOS, CVP is packaged and released in trains.

Note: While upgrading CVP, refer to the latest release notes available at Arista Software Download page; and upgrade procedures.

Devices under management must:

- be running supported EOS version
- have supported TerminAttr version installed
- have the TerminAttr agent enabled and successfully streaming telemetry to CVP.

The following steps can be taken at any point on an existing cluster as part of preparing for an upgrade to the current version:

1. Upgrade existing CVP clusters to the latest CVP release
2. Upgrade all EOS devices under management to the supported release train.

3. For devices running EOS releases prior to *4.20*, ensure that the eAPI unix domain socket is enabled with the following configuration:

```
management api http-commands
   protocol unix-socket
```

4. Install supported TerminAttr on all EOS devices under management.
5. Enable state streaming from all EOS devices under management by applying the **SYS_StreamingTelemetry** configlet and pushing the required configuration to all devices.
6. Ensure that all devices are successfully streaming to the CVP cluster.
7. Ensure that all devices are in image and config compliance.
8. Complete regular backups. Complete a final backup prior to upgrade.
9. Ensure that all tasks are in a terminal state (Success, Failed, or Canceled).
10. Ensure that all Change Controls are in a terminal state.

Note: After the cluster is upgraded to the latest CVP release, systems running unsupported TerminAttr versions fail to connect to the CVP cluster. These devices will have to be first upgraded to a supported TerminAttr version by re-onboarding them from the CloudVision UI. You cannot rollback a device to a time before it was running the supported TerminAttr version.

The upgrade from the previous CVP release to the current CVP release trains include data migrations that can take several hours on larger scale systems.

- Upgrades
- CVP Node RMA
- CVP / EOS Dependencies
- Upgrade CVW As Part of a CV Upgrade

20.2.1 Upgrades

Upgrades do not require that the VMs be redeployed, and do not result in the loss of logs. .

The CVP cluster must be functional and running to successfully complete an upgrade. As a precaution against the loss of CVP data, it is recommended that you back-up the CVP data before performing an upgrade

To upgrade CVP to the current release, you must first upgrade CVP to the supported release that supports an upgrade to the current release. For more information, refer the CVP release notes at Arista Software Download page.

20.2.1.1 Verifying the Health of CVP before Performing Upgrades

Upgrades should only be performed on healthy and fully functional CVP systems. Before performing the upgrade, make sure that you verify that the CVP system is healthy.

Complete the following steps to verify the health of CVP.

1. Enter into the Linux shell of the primary node as **cvp user**.
2. Execute the `cvpi status all` command on your CVP:

   This shows the status of all CVP components.
3. Confirm that all CVP components are running.
4. Log into the CVP system to check functionality.

Once you have verified the health of your CVP installation, you can begin the upgrade process.

- Upgrading CloudVision Portal (CVP)

20.2.2 CVP Node RMA

Use this procedure to replace any node of a multi-node cluster. Replacing nodes of multi-node cluster involves removing the node you want to replace, waiting for the remaining cluster nodes to recover, powering on the replacement node, and applying the cluster configuration to the new node.

When you replace cluster nodes, you must replace only **one node at a time**. In case, you plan to replace more than one node of a cluster, you must complete the entire procedure for each node to be replaced.

Note: It is recommended that you save the cvp cluster configuration to a temporary file, or write down the configuration on a worksheet. The configuration can be found in `/cvpi/cvp-config.yaml`.

1. Power off the node you want to replace (primary, secondary, or tertiary).
2. Remove the node to be replaced.
3. Allow all components of the remaining nodes to recover.
4. Use the `cvpi status all` command to ensure that remaining nodes are healthy.

```
[cvp@cvp73 root]$ cvpi status all

Current Running Command: None
Executing command. This may take a few seconds...
primary         78/78 components running
secondary       89/89 components running
tertiary        NODE DOWN
```

5. Power on the replacement node.
6. Log in as *cvpadmin*.
7. Enter the cvp cluster configuration.

```
CentOS Linux 7 (Core)
Kernel 3.10.0-957.1.3.el7.x86_64 on an x86_64

localhost login: cvpadmin
Last login: Fri Mar 15 12:24:45 on ttyS0
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
Enter a command
[q]uit [p]rint [s]inglenode [m]ultinode [r]eplace [u]pgrade
>r
Please enter minimum configuration to connect to the other peers
*Ethernet interface for the cluster network: eth0
*IP address of eth0: 172.31.0.216
*Netmask of eth0: 255.255.0.0
*Default route: 172.31.0.1
*IP address of one of the two active cluster nodes: 172.31.0.161
 Root password of 172.31.0.161:
```

8. Wait for the RMA process to complete. No action is required.

```
Root password of 172.31.0.161:
External interfaces, ['eth1'], are discovered under /etc/sysconfig/
network-scripts
These interfaces are not managed by CVP.
Please ensure that the configurations for these interfaces are correct.
Otherwise, actions from the CVP shell may fail.
Running : /bin/sudo /sbin/service network restart
[  334.001886] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
```

```
[  334.004577] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[  334.006315] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[  334.267535] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[  348.252323] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[  348.254925] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[  348.256504] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[  348.258035] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
Fetching version information
Run cmd: sudo -u cvp -- ssh 172.31.0.156 cat /cvpi/property/version.txt
 0.18
Fetching version information
Run cmd: sudo -u cvp -- ssh 172.31.0.216 cat /cvpi/property/version.txt
 10.19
Fetching version information
Run cmd: sudo -u cvp -- ssh 172.31.0.161 cat /cvpi/property/version.txt
 0.16
Running : cvpConfig.py tool...
[  392.941983] vmxnet3 0000:0b:00.0 eth0: intr type 3, mode 0, 9
 vectors allocated
[  392.944739] vmxnet3 0000:0b:00.0 eth0: NIC Link is Up 10000 Mbps
[  392.946388] IPv6: ADDRCONF(NETDEV_UP): eth0: link is not ready
[  393.169460] IPv6: ADDRCONF(NETDEV_CHANGE): eth0: link becomes ready
[  407.229180] vmxnet3 0000:13:00.0 eth1: intr type 3, mode 0, 9
 vectors allocated
[  407.232306] vmxnet3 0000:13:00.0 eth1: NIC Link is Up 10000 Mbps
[  407.233940] IPv6: ADDRCONF(NETDEV_UP): eth1: link is not ready
[  407.235728] IPv6: ADDRCONF(NETDEV_CHANGE): eth1: link becomes ready
[  408.447642] Ebtables v2.0 unregistered
[  408.935626] ip_tables: (C) 2000-2006 Netfilter Core Team
[  408.956578] ip6_tables: (C) 2000-2006 Netfilter Core Team
[  408.982927] Ebtables v2.0 registered
[  409.029603] nf_conntrack version 0.5.0 (65536 buckets, 262144 max)
Stopping: ntpd
Running : /bin/sudo /sbin/service ntpd stop
Running : /bin/sudo /bin/systemctl is-active ntpd
Starting: ntpd
Running : /bin/sudo /bin/systemctl start ntpd.service
Waiting for all components to start. This may take few minutes.
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status zookeeper' 0.45
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status zookeeper' 0.33
Checking if third party applications exist
Run cmd: su - cvp -c '/cvpi/zookeeper/bin/zkCli.sh ls /apps | tail -1'
 0.72
Running : cvpConfig.py tool...
Stopping: cvpi-check
Running : /bin/sudo /sbin/service cvpi-check stop
Running : /bin/sudo /bin/systemctl is-active cvpi-check
Starting: cvpi-check
Running : /bin/sudo /bin/systemctl start cvpi-check.service
```

9. Continue waiting for the RMA process to complete. No action is required.

```
[Fri Mar 15 20:26:28 UTC 2019] :
Executing command. This may take a few seconds...

(E) => Enabled
(D) => Disabled
(?) => Zookeeper Down

Action Output
-------------
COMPONENT           ACTION              NODE                STATUS
         ERROR
```

```
hadoop              cluster             tertiary            (E) DONE

hbase               cluster             tertiary            (E) DONE

Executing command. This may take a few seconds...

(E) => Enabled
(D) => Disabled
(?) => Zookeeper Down

Action Output
-------------
COMPONENT           ACTION              NODE                STATUS
        ERROR
aerisdiskmonitor    config              primary             (E) DONE

aerisdiskmonitor    config              secondary           (E) DONE

aerisdiskmonitor    config              tertiary            (E) DONE

apiserver           config              primary             (E) DONE

apiserver           config              secondary           (E) DONE

apiserver           config              tertiary            (E) DONE

cvp-backend         config              primary             (E) DONE

cvp-backend         config              secondary           (E) DONE

cvp-backend         config              tertiary            (E) DONE

cvp-frontend        config              primary             (E) DONE

cvp-frontend        config              secondary           (E) DONE

cvp-frontend        config              tertiary            (E) DONE

geiger              config              primary             (E) DONE

geiger              config              secondary           (E) DONE

geiger              config              tertiary            (E) DONE

hadoop              config              primary             (E) DONE

hadoop              config              secondary           (E) DONE

hadoop              config              tertiary            (E) DONE

hbase               config              primary             (E) DONE

hbase               config              secondary           (E) DONE

hbase               config              tertiary            (E) DONE

kafka               config              primary             (E) DONE

kafka               config              secondary           (E) DONE

kafka               config              tertiary            (E) DONE

zookeeper           config              primary             (E) DONE
```

```
zookeeper           config              secondary           (E) DONE

zookeeper           config              tertiary            (E) DONE

Executing command. This may take a few seconds...
secondary       89/89 components running
primary         78/78 components running
Executing command. This may take a few seconds...
COMPONENT           ACTION              NODE                STATUS
         ERROR
Including: /cvpi/tls/certs/cvp.crt
Including: /cvpi/tls/certs/cvp.key
Including: /etc/cvpi/cvpi.key
Including: /cvpi/tls/certs/kube-cert.pem
Including: /data/journalnode/mycluster/current/VERSION
Including: /data/journalnode/mycluster/current/last-writer-epoch
Including: /data/journalnode/mycluster/current/last-promised-epoch
Including: /data/journalnode/mycluster/current/paxos
Including: /cvpi/tls/certs/ca.crt
Including: /cvpi/tls/certs/ca.key
Including: /cvpi/tls/certs/server.crt
Including: /cvpi/tls/certs/server.key
mkdir -p /cvpi/tls/certs
mkdir -p /data/journalnode/mycluster/current
mkdir -p /cvpi/tls/certs
mkdir -p /etc/cvpi
mkdir -p /cvpi/tls/certs
mkdir -p /cvpi/tls/certs
mkdir -p /cvpi/tls/certs
mkdir -p /data/journalnode/mycluster/current
mkdir -p /cvpi/tls/certs
mkdir -p /data/journalnode/mycluster/current
mkdir -p /data/journalnode/mycluster/current
mkdir -p /cvpi/tls/certs
Copying: /etc/cvpi/cvpi.key from secondary
rsync -rtvp 172.31.0.161:/etc/cvpi/cvpi.key /etc/cvpi
Copying: /cvpi/tls/certs/cvp.crt from secondary
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/cvp.crt /cvpi/tls/certs
Copying: /cvpi/tls/certs/server.key from secondary
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/server.key /cvpi/tls/certs
Copying: /cvpi/tls/certs/ca.crt from secondary
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/ca.crt /cvpi/tls/certs
Copying: /cvpi/tls/certs/cvp.key from secondary
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/cvp.key /cvpi/tls/certs
Copying: /cvpi/tls/certs/ca.key from secondary
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/ca.key /cvpi/tls/certs
Copying: /data/journalnode/mycluster/current/last-writer-epoch from
 secondary
rsync -rtvp 172.31.0.161:/data/journalnode/mycluster/current/last-
writer-epoch /data/journalnode/mycluster/current
Copying: /cvpi/tls/certs/kube-cert.pem from secondary
Copying: /cvpi/tls/certs/server.crt from secondary
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/server.crt /cvpi/tls/certs
Copying: /data/journalnode/mycluster/current/VERSION from secondary
rsync -rtvp 172.31.0.161:/data/journalnode/mycluster/current/VERSION /
data/journalnode/mycluster/current
Copying: /data/journalnode/mycluster/current/paxos from secondary
rsync -rtvp 172.31.0.161:/data/journalnode/mycluster/current/paxos /
data/journalnode/mycluster/current
Copying: /data/journalnode/mycluster/current/last-promised-epoch from
 secondary
rsync -rtvp 172.31.0.161:/data/journalnode/mycluster/current/last-
promised-epoch /data/journalnode/mycluster/current
rsync -rtvp 172.31.0.161:/cvpi/tls/certs/kube-cert.pem /cvpi/tls/certs
```

```
Starting: cvpi-config
Running : /bin/sudo /bin/systemctl start cvpi-config.service
Starting: cvpi
Running : /bin/sudo /bin/systemctl start cvpi.service
Running : /bin/sudo /bin/systemctl start cvpi-watchdog.timer
Running : /bin/sudo /bin/systemctl enable docker
Running : /bin/sudo /bin/systemctl start docker
Running : /bin/sudo /bin/systemctl enable kube-cluster.path
```

10. Enter "q" to quit the process after the **RMA process is complete!** message is displayed.

```
Waiting for all components to start. This may take few minutes.
[  560.918749] FS-Cache: Loaded
[  560.978183] FS-Cache: Netfs 'nfs' registered for caching
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 48.20
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.73
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 7.77
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.55
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.23
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.64
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.59
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.07
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.70
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.51
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.57
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.40
Run cmd: su - cvp -c '/cvpi/bin/cvpi status all --cluster' 2.24
Waiting for all components to start. This may take few minutes.
Run cmd: su - cvp -c '/cvpi/bin/cvpi -v=3 status all' 9.68
RMA process is complete!
[q]uit [p]rint [e]dit [v]erify [s]ave [a]pply [h]elp ve[r]bose
>q
```

11. Use the `cvpi status all` command to ensure that the cluster is healthy.

```
[cvp@cvp87 ~]$ cvpi status all

Current Running Command: None
Executing command. This may take a few seconds...
primary         78/78 components running
secondary       89/89 components running
tertiary        45/45 components running
```

Related topics:

- CVP / EOS Dependencies
- Upgrades

20.2.3 CVP / EOS Dependencies

To ensure that CVP can provide a base level of management, all EOS devices must be running at least EOS versions *4.17.3F* or later. To ensure device compatibility supported EOS version advice should be sought from the Arista account team.

CVP should not require any additional EOS upgrades to support the standard features and functions in later versions of the appliance. Newer features and enhancements to CVP may not be available for devices on older code versions.

Refer to the latest Release Notes for additional upgrade/downgrade guidance.

Related topics:

- Upgrades

- CVP Node RMA

20.2.4 Upgrade CVW As Part of a CV Upgrade

In case of a CV upgrade, services go through the following steps:

1. Services or service containers (such as CVW) are stopped.
2. Existing container images are deleted.
3. New component RPMs are installed.
4. The server is rebooted and all services are started again.

   A service on CV is upgraded only if its version is different from the pre-upgrade version (CV stores its pre-upgrade state to decide this). The wifimanager component follows a similar process. When CV boots up after an upgrade, wifimanager starts and upgrades only if the CV upgrade has resulted in a new wifimanager version. The following actions precede every wifimanager **start** operation:

   a. `load`: Loads the wifimanager container image into docker when CV boots up for the first time after an upgrade.
   b. `init`: Initializes wifimanager before the start. The wifimanager `init` is versioned*init-8.8.0-01*, for example. The `init-<version>` handler initiates a wifimanager upgrade if needed. Thus, if the wifimanager version has not changed after the CV upgrade, the wifimanager upgrade is not invoked. If the wifimanager version has changed, then a wifimanager upgrade is called before its start.

   Note: Load and init are internal actions to the wifimanager start operation; they are not run separately. The CVW service might take longer to start than other CV services.

20.3 DNS / NTP Server Migration

You can migrate your DNS / NTP server after you have completed your initial deployment of CloudVision. Migrating the DNS / NTP server is typically done if you want to or need to change the DNS / NTP server that CloudVision currently uses.

For example, if the current CloudVision DNS / NTP server was intentionally isolated during the initial CloudVision installation, you need to migrate the server to make it accessible by external resources.

- Migrating the DNS and NTP Server

20.3.1 Migrating the DNS and NTP Server

The process for migrating the DNS / NTP server after the completion of the initial CloudVision installation involves updating the DNS and NTP server entries on each cluster node and modifying the `/cvpi/cvp-config.yaml` file (on each node) to reflect the updates to the server entries.

Pre-requisites

Before you begin the migration process, make sure that:

- The IP addresses and hostnames (fqdn) of the nodes must not change.
- For each node, make sure that:
  - At least one DNS server entry is present in the /cvpi/cvp-config.yaml file.
  - The DNS server that corresponds to the DNS server entry in the `/cvpi/cvp-config.yaml` file can be accessed by the cluster throughout the migration process. (The reason for this is that any changes made to resolv.conf take effect immediately upon saving the file.)
- The time difference between the old NTP server and new NTP server should be negligible.
- The old NTP server and new NTP server should be in same time zone.

Complete these steps to migrate the DNS / NTP server.

1. On each node, **add** the new server to `/etc/resolv.conf`, by adding a new nameserver line at the top of the file. For example, `nameserver 172.22.22.40`.
2. On each node, **remove** the old server from `/etc/resolv.conf`, by removing the old nameserver line.
3. On each node, do the following to update the NTP server:
   a. Run the `ntpstat` command to make note of the current NTP server.
   b. In `/etc/ntp.conf`, add the new NTP server entry and **comment out** the entry for the old NTP server.
   c. Run the `service ntpd restart` command.
   d. Run the `ntpstat` command to verify that the NTP server has been changed on all nodes.
4. On each node, edit the `/cvpi/cvp-config.yaml` file to reflect the changes to the DNS and NTP server entries you made in the previous steps.

Related topics:

- Backup and Restore
- Backup and Restore

Chapter 21

Supplementary Services

This document provides configurations steps and examples for supplementary setup procedures for CloudVision Portal (CVP).

21.1 HTTPS Certificates Setup

CVP uses nginx to front and terminate all HTTPS connections. To support HTTPS, the server must be configured with a certificate. A selfsigned certificate is generated at first bootup.

- To install your own certificate and key, copy the certificate to the following location on CVP:

```
/etc/nginx/cvp.crt
/etc/nginx/cvp.key
```

- To generate a new selfsigned certificate, follow the example below:

```
openssl req -new -nodes -x509 -days 365 -out
/etc/nginx/cvp.crt -keyout /etc/nginx/cvp.key -subj /CN=self.signed
```

- You should restart the nginx service if you install a new certificate/key:

```
service nginx reload
```

Related topics:

- DHCP Service for Zero Touch Provisioning (ZTP) Setup
- RADIUS or TACACS Authentication Setup
- Background Tasks
- Resetting cvpadmin Password

21.2 Customizing TLS and SSH Ciphers

CVP uses nginx to front and terminate all HTTPS connections. To support HTTPS, the server must be configured with a certificate. A selfsigned certificate is generated at first bootup.

- Configuring Custom TLS Ciphers
- Configuring Custom SSH Ciphers

21.2.1 Configuring Custom TLS Ciphers

Complete these steps to configure custom TLS ciphers.

1. Create a file named */etc/nginx/conf.d/locations/cvp-ciphers.https.conf* that contains all of the SSL ciphers you need. Any open ssl cipher string can be used.
2. Run the following command to make sure the configuration does not contain any errors:

```
/usr/sbin/nginx -t -c /etc/nginx/conf.d/cvpi-server.conf
```

3. Run the following command to reload nginx with the updated configuration.

```
systemctl reload nginx
```

21.2.2 Configuring Custom SSH Cipher

Complete these steps to configure custom SSH ciphers.

Note: Upgrading CVP removes custom SSH ciphers. You must reconfigure SSH ciphers after the upgrade.

1. Edit the `/etc/cvpi/sshd_config` to include custom ciphers and MAC definitions.
2. Run the following command to make sure the configuration does not contain any errors:

```
sshd -t -f /etc/cvpi/sshd_config
```

3. Run the following command to reload sshd with the updated configuration.

```
systemctl reload sshd
```

21.3 DHCP Service for Zero Touch Provisioning (ZTP) Setup

The ZTP process relies on a DHCP server to get devices registered with CVP. The DHCP server can be on the CVP, but is more commonly an external DHCP server.

1. Ensure the DHCP server is installed (it is installed by default in CVP).

```
rpm -qa | grep dhcp
dhcp-common-4.1.1-43.P1.el6.x86_64
dhcp-4.1.1-43.P1.el6.x86_64
```

2. Edit the `/etc/dhcp/dhcpd.conf` file to include the option bootfile-name, which provides the location of the script that starts the ZTP process between CVP and the device.

   In this example, DHCP is serving the *172.31.0.0/16* subnet.

   Note: The *172.31.5.60* is the IP address of a CVP node, and that you must use the HTTP (and not HTTPS) URL to the bootstrap file. This ensures that the specified devices, after they ZTP, will show up under the undefined container of the specified CVP.

```
[root@cvp1-dhcp dhcp]# cat dhcpd.conf
#
# DHCP Server Configuration file.
#   see /usr/share/doc/dhcp*/dhcpd.conf.sample
#   see 'man 5 dhcpd.conf'
#
subnet 172.31.0.0 netmask 255.255.0.0 {
  range 172.31.3.212 172.31.5.214;
  option domain-name "sjc.aristanetworks.com";
```

```
}
host esx21-vm20 {
  option dhcp-client-identifier 00:0c:29:f9:21:99;
  fixed-address 172.31.3.211;
  option bootfile-name "http://172.31.5.60/ztp/bootstrap";
}
host esx21-vm22 {
  option dhcp-client-identifier 00:0c:29:d1:64:e1;
  fixed-address 172.31.3.213;
  option bootfile-name "http://172.31.5.60/ztp/bootstrap";
}
```

3. Restart the DHCP service after any configuration changes with the `service dhcpd restart` command.
4. Configure dhcpd to start on system boot with the `chkconfig dhcpd on` command.

Related topics:

- RADIUS or TACACS Authentication Setup
- Background Tasks
- Resetting cvpadmin Password
- HTTPS Certificates Setup

21.4 RADIUS or TACACS Authentication Setup

1. Edit the client file `/etc/raddb/clients.conf` by adding the following:

```
# CVP
client 172.31.0.0/16 {
        secret = cvpsecret
```

2. To add more, enter the following.

```
# Arista Networks
client 172.17.0.0/16 {
        secret = cvpsecret
}
client 172.18.0.0/16 {
        secret = cvpsecret
}
client 172.20.0.0/16 {
        secret = cvpsecret
}
client 172.22.0.0/16 {
        secret = cvpsecret
}
```

The default `clients.conf` file will have a section for local host. The user should either delete the whole section or comment it out. If CVP will be connecting to RADIUS on local host. You have to add a client entry for *127.0.0.0/16* (same as above).

1. Edit the users file `/etc/raddb/users` by adding the following:

```
# CVP
cvpuser Cleartext-Password := "cvpuser"
        Service-Type = NAS-Prompt-User

start radiusd:  sudo service radiusd start
enable radiusd on boot: sudo chkconfig radiusd on
```

2. If RADIUS is not working, run the server in debug mode.

```
# service radiusd stop
# /usr/sbin/radiusd -X -f
```

RADIUS will now run on the terminal with verbose output. This will let you know if RADIUS is receiving auth requests and what failure is being hit for the request. After you are done debugging, Control-C the process and start radiusd as a service.

Note: You may have to either disable iptables or firewall.serviced depending on the OS version. You could also configure it to allow traffic on ports 1812 and 1813 on the Radius server.

Related topics:

- Background Tasks
- Resetting cvpadmin Password
- HTTPS Certificates Setup
- DHCP Service for Zero Touch Provisioning (ZTP) Setup

21.5 Background Tasks

CloudVision provides command-line tools that can be executed from the linux shell or scheduled as cronjobs either on a CVP node or on an external server, for the following tasks:

- Compliance checks
- Snapshots
- Backups

The tools are available by default on the CVP nodes in the `/cvpi/tools/` directory. The tools can be used on an external linux server by downloading the `cvp-tools-<version> .tgz` from https://www.arista.com to the external linux server.

Detailed help on the tool is available by using the –h option with the tool:

```
cvpi/tools/compliance.py -h
cvpi/tools/backup.py -h
```

21.5.1 Scheduling and Viewing Cronjobs To schedule cronjobs to perform periodic compliance checks or snapshots, insert commands into the crontab using the following command:

```
crontab -e
```

Note: Note When inserting commands to schedule cronjobs, you only need to do this on one node of the cluster.

Example

To schedule a periodic compliance check and snapshot to be performed hourly on the tenant container, and a backup to be performed daily at 2:00 am, insert the following lines into the crontab file on the primary node if not already present. In this example, the user is named “**me**” and the password is “**pwd**”.

```
0 * * * * /cvpi/tools/compliance.py --user me --password pwd --containers
 tenant
0 2 * * * /cvpi/tools/backup.py --limit 5
```

To see the active cronjobs, use the following command:

```
crontab -l
```

To view the console outputs of the cronjobs tail, view (open) the following log file:

```
tail -f /var/log/cron
```

Related topics:

- Resetting cvpadmin Password
- HTTPS Certificates Setup
- DHCP Service for Zero Touch Provisioning (ZTP) Setup
- RADIUS or TACACS Authentication Setup

21.5.1 Scheduling and Viewing Cronjobs

To schedule cronjobs to perform periodic compliance checks or snapshots, insert commands into the crontab using the following command:

```
crontab -e
```

Note: When inserting commands to schedule cronjobs, you only need to do this on one node of the cluster.

Example

To schedule a periodic compliance check and snapshot to be performed hourly on the tenant container, and a backup to be performed daily at 2:00 am, insert the following lines into the crontab file on the primary node if not already present. In this example, the user is named "**me**" and the password is "**pwd**".

```
0 * * * * /cvpi/tools/compliance.py --user me --password pwd --containers
 tenant
0 2 * * * /cvpi/tools/backup.py --limit 5
```

To see the active cronjobs, use the following command:

```
crontab -l
```

To view the console outputs of the cronjobs tail, view (open) the following log file:

```
tail -f /var/log/cron
```

Related topics:

- Resetting cvpadmin Password
- HTTPS Certificates Setup
- DHCP Service for Zero Touch Provisioning (ZTP) Setup
- RADIUS or TACACS Authentication Setup

21.6 Resetting cvpadmin Password

If the *cvpadmin* password is lost or forgotten, you can reset it from any of the CVP nodes using the following steps.

1. Log into a CVP node Linux shell as root user.
2. Navigate to `cd /cvpi/lib`

3. Execute the following command:

```
/cvpi/tools/update-mgmt-password -password <new password>
```

Note: Do not set the new password to the string "*cvpadmin*".

Related topics:

- HTTPS Certificates Setup
- DHCP Service for Zero Touch Provisioning (ZTP) Setup
- RADIUS or TACACS Authentication Setup
- Background Tasks

Chapter 22

Troubleshooting and Health Checks

If you encounter an issue when using CloudVision appliance, check to see if there are troubleshooting steps for the issue.

- System Recovery
- Health Checks
- Resource Checks

22.1 System Recovery

System recovery should be used only when the CVP cluster has become unusable and other steps, such as performing a `cvpi watchdog off, cvpi stop all, and then, cvpi start all, cvpi watchdog` on have failed. For example, situations in which, regardless of restarts, a `cvpi status all` continues to show some components as having a status of UNHEALTHY or NOT RUNNING.

There are two ways to completely recover a CVP cluster:

- VM Redeployment
- CVP Re-Install without VM Redeployment

Note: A good backup is required to proceed with either of these system recoveries.

22.1.1 VM Redeployment

Complete the foloowing steps:

1. Delete all the CVP VMs.
2. Redeploy the VMs using the procedures in.
3. Issue a `cvpi status all` command to ensure all components are running.
4. Login to the CVP GUI as `cvpadmin/cvpadmin` to set the cvpadmin password.
5. From the **Backup & Restore** tab on the **Setting** page, restore from the backup.

22.2 Health Checks

The following table lists the different types of CVP health checks you can run, including the steps to use to run each check and the expected result for each check.

Component	Steps to Use	Expected Result
Network connectivity	ping -f across all nodes	No packet loss, network is healthy.

Component	Steps to Use	Expected Result
HBase	echo `list | /cvpi/hbase/bin/hbase shell | grep -A 2 row\(`	Prints an array of tables in Hbase created by CVP Hbase, and the underlying infrastructure works.
All daemons running on all nodes, bypass `cvpi status all`	On all nodes: `su - cvp -c "/cvpi/jdk/bin/jps"`	On primary and secondary nodes, 9 processes including jps. • 3149 HMaster • 2931 NameNode • 2797 QuorumPeerMain • 12113 Bootstrap • 3040 DFSZKFailoverController • 2828 JournalNode • 11840 HRegionServer • 12332 Jps • 2824 DataNode On tertiary 6 processes: • 2434 JournalNode • 4256 HRegionServer • 2396 QuorumPeerMain • 2432 DataNode • 4546 Jps • 8243 Bootstrap
Check time is in sync between nodes	On all nodes run `date +%s`	UTC time should be within a few seconds of each other (typically less than one second). Up to 10 seconds is allowable.
I/O slowness issues	The disk I/O throughput is at an unhealthy level (too low).	Use the `cvpi resources` command to find out whether the disk I/O throughput is at a **healthy level** or **unhealthy level**. The disk I/O throughput reported in the command output is measured by the Virtual Machine. See Running Health Checks for an example of the output of the `cvpi resources` command.

- Running Health Checks

22.2.1 Running Health Checks

Run the `cvpi resources` command to execute a health check on disk bandwidth. The output of the command indicates whether the disk bandwidth is at a healthy level or unhealthy level. The threshold for healthy disk bandwith is 20MBS.

The possible health statuses are:

- **Healthy** - Disk bandwidth above 20MBs
- **Unhealthy** - Disk bandwidth at or below 20MBs

The output is color coded to make it easy to interpret the output. Green indicates a healthy level, and red indicates an unhealthy level (see the example below).

This example shows output of the `cvpi resources` command. In this example, the disk bandwidth status is healthy (above the 20MBs threshold).

```
[root@varuns-cvpfoster ~]# su cvp
[cvp@varuns-cvpfoster root]$ cvpi status all

Current Running Command: None
Executing command. This may take a few seconds...
primary         128/128 components running
[cvp@varuns-cvpfoster root]$ cvpi resources

+---------------------------+----------------------+
|            NODE           |       PRIMARY        |
+---------------------------+----------------------+
| N/w bandwidth to all nodes | 14.60 MB/s           |
| CPU Count                 | 8                    |
| Disk Throughput for /data | 172.437 MB/s         |
| Total Memory              | 21.4G                |
| N/w latency to all nodes  | 0.05 ms              |
| NTP Status                | synchronized         |
| Size of /data             | 1023.6G (941.2G)     |
| System Time               | 2019-03-14T02:40:42Z |
+---------------------------+----------------------+
[cvp@varuns-cvpfoster root]$ cvpi status cvp

Current Running Command: None
Executing command. This may take a few seconds...
primary         17/17 components running
[cvp@varuns-cvpfoster root]$
```

Figure 379: Example output of cvpi resources command

Related topics

- Resource Checks

22.3 Resource Checks

CloudVision Portal (CVP) enables you to run resource checks on CVP node VMs. You can run checks to determine the current data disk size of VMs that you have upgraded to CVP version 2017.2.0, and to determine the current memory allocation for each CVP node VM.

Performing these resource checks is important to ensure that the CVP node VMs in your deployment have the recommended data disk size and memory allocation for using the Telemetry feature. If the

resource checks show that the CVP node VM data disk size or memory allocation (RAM) are below the recommended levels, you can increase the data disk size and memory allocation.

These procedures provide detailed instructions on how to perform the resource checks and if needed, how to increase the CVP node VM data disk size and CVP node VM memory allocation.

- Running CVP node VM Resource Checks
- Increasing Disk Size of VMs Upgraded to CVP Version 2017.2.0
- Increasing CVP Node VM Memory Allocation

22.3.1 Running CVP node VM Resource Checks

CloudVision Portal (CVP) enables you to quickly and easily check the current resources of the primary, secondary, and tertiary nodes of a cluster by running a single command. The command you use is the `cvpi resources` command.

Use this command to check the following CVP node VM resources:

- Memory allocation
- Data disk size (storage capacity)
- Disk throughput (in MB per second)
- Number of CPUs

Complete the following steps to run the CVP node VM resource check.

1. Login to one of the CVP nodes as **root**.
2. Execute the `cvpi resources` command.

   The output shows the current resources for each CVP node VM

   - If the total size of sdb1 (or vdb1) is approximately 120G or less, you can increase the disk size to 1TB (seeIncreasing Disk Size of VMs Upgraded to CVP Version 2017.2.0).
   - If the memory allocation is the default of 16GB, you can increase the RAM memory allocation (see Increasing CVP Node VM Memory Allocation).

```
[cvp@cvp56 root]$ cvpi resources

+--------------------------+----------------------+----------------------+----------------------+
|           NODE           |       PRIMARY        |       SECONDARY      |       TERTIARY       |
+--------------------------+----------------------+----------------------+----------------------+
| N/w bandwidth to all nodes | 14.98/13.52/10.57 MB/s | 11.87/19.32/13.76 MB/s | 10.96/12.06/10.78 MB/s |
| CPU Count                | 8                    | 8                    | 8                    |
| Disk Throughput for /data | 103.575 MB/s         | 179.037 MB/s         | 99.010 MB/s          |
| Total Memory             | 15.5G                | 15.5G                | 15.5G                |
| N/w latency to all nodes | 0.04/0.23/0.23 ms    | 0.20/0.03/0.77 ms    | 0.35/0.18/0.05 ms    |
| NTP Status               | synchronized         | synchronized         | synchronized         |
| Size of /data            | 1023.6G (970.1G)     | 1023.6G (970.1G)     | 1023.6G (970.1G)     |
| System Time              | 2019-03-18T06:27:40Z | 2019-03-18T06:27:40Z | 2019-03-18T06:27:40Z |
+--------------------------+----------------------+----------------------+----------------------+
[cvp@cvp56 root]$
```

Figure 380: Using the cvpi resource command to run CVP node VM resource checks

22.3.2 Increasing Disk Size of VMs Upgraded to CVP Version 2017.2.0

If you already upgraded any CVP node VMs running an older version of CVP to version 2017.2.0, you may need to increase the size of the data disk of the VMs so that the data disks have the 1TB disk image that is used on current CVP node VMs

CVP node VM data disks that you upgraded to version 2017.2.0 may still have the original disk image (120GB data image), because the standard upgrade procedure did not upgrade the data disk image. The standard upgrade procedure updated only the root disk, which contains the Centos image along with rpms for CVPI, CVP, and Telemetry.

Note: It is recommended that each CVP node have 1TB of disk space reserved for enabling CVP Telemetry. If the CVP nodes in your current environment do not have the recommended

reserved disk space of 1TB, complete the procedure below for increasing the disk size of CVP node VMs.

Pre-requisites

Before you begin the procedure, make sure that you:

- Have upgraded to version 2017.2.0. You cannot increase the data disk size until you have completed the upgrade to version 2017.2.0 (see Migrating the DNS and NTP Server).
- Have performed the resource check to verify that the CVP node VMs have the data disk size image of previous CVP versions (approximately 120GB or less). See Running CVP node VM Resource Checks.

Procedure

Complete the following steps to increase the data disk size.

1. Turn off cvpi service by executing the `systemctl stop cvpi` command on all nodes in the cluster. (For a single-node installation, run this command on the node.)
2. Run the `cvpi -v=3 stop all` on the primary node.
3. Perform a **graceful power-off** of all VMs.

   Note: You do not need to unregister and re-register VMs from vSphere Client or undefine and redefine VMs from kvm hypervisor.
4. Do the following to increase the size of the data disk to 1TB using the hypervisor:
   - **ESX**: Using vSphere client, do the following:
     a. Select the **Virtual Hardware** tab, and then select **hard disk 2**.
     b. Change the setting from 120GB to **1TB**.
     c. Click **OK**.
   - **KVM**: Use the `qemu-img resize` command to resize the data disk from 120GB to 1TB. Be sure to select **disk2.qcow2.**

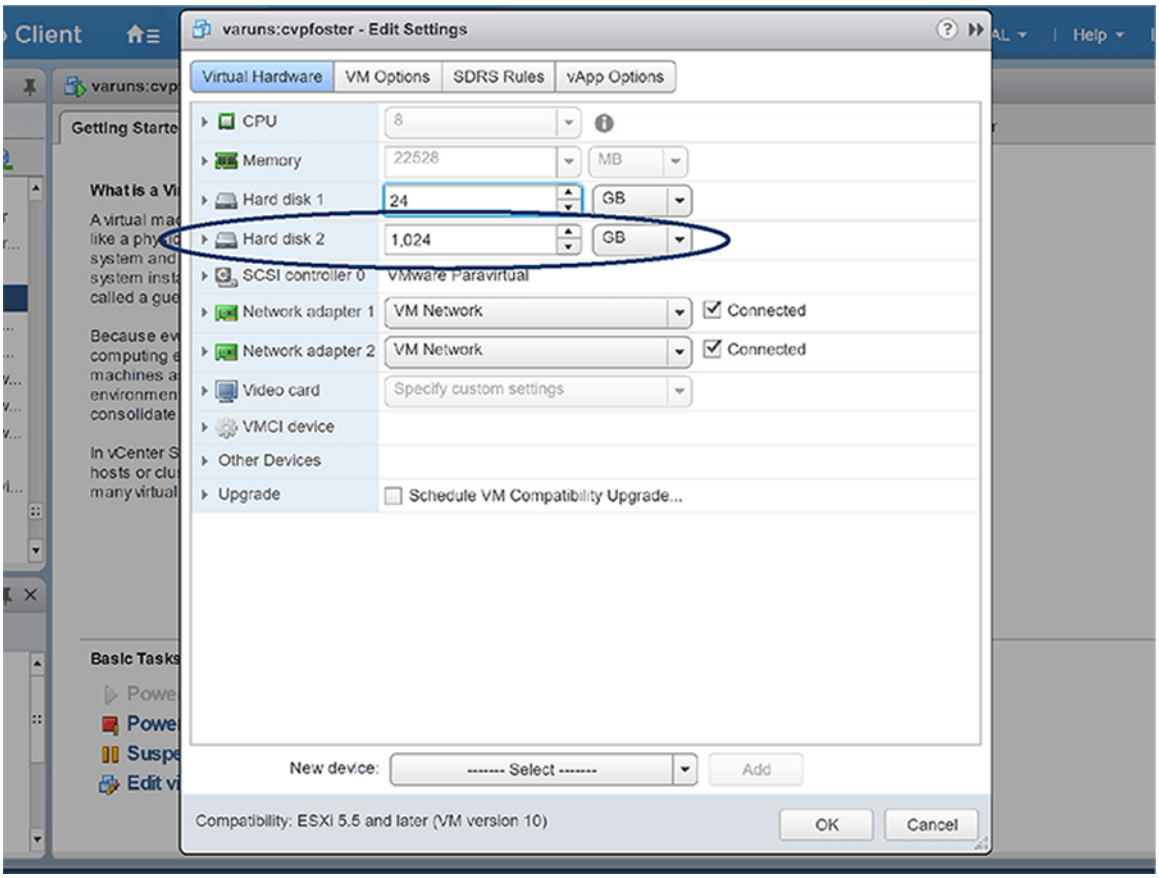

Figure 381: Using vSphere to increase data disk size

5. Power on all CVP node VMs, and wait for all services to start.
6. Use the `cvpi status all` command to verify that all the cvpi services are running.
7. Run the `/cvpi/tools/diskResize.py` command on the primary node. (Do not run this command on the secondary and tertiary nodes.)
8. Run the `df -h /data` command on all nodes to verify that the /data is increased to approximately 1TB.
9. Wait for all services to start.
10. Use the `cvpi -v=3 status all` command to verify the status of services.
11. Use the `systemctl status cvpi` to ensure that cvpi service is running.

22.3.3 Increasing CVP Node VM Memory Allocation

If the CVP Open Virtual Appliance (OVA) template currently specifies the default of 16GB of memory allocated for the CVP node VMs in the CVP cluster, you need to increase the RAM to ensure that the CVP node VMs have adequate memory allocated for using the Telemetry feature.

Note: It is recommended that CVP node VMs have 32GB of RAM allocated for deployments in which Telemetry is enabled.

You can perform a rolling modification to increase the RAM allocation of every node in the cluster. If you want to keep the service up and available while you are performing the rolling modification, make sure that you perform the procedure on only one CVP node VM at a time.

Once you have completed the procedure on a node, you repeat the procedure on another node in the cluster. You must complete the procedure once for every node in the cluster.

Pre-requisites

Before you begin the procedure, make sure that you:

- Have performed the resource check to verify that the CVP node VMs have the default RAM memory allocation of 16GB (see Running CVP node VM Resource Checks).
- Make sure that you perform a GUI-based backup of the CVP system and copy the backup to a safe location (a location off of the CVP node VMs). The CVP GUI enables you to create a backup you can use to restore CVP data.

Procedure

Complete the following steps to increase the RAM memory allocation of the CVP node VMs.

1. Login to a CVP node of the cluster as **cvp user**.
2. Using the `cvpi status cvp shell` command, make sure that all nodes in the cluster are operational.

```
[cvp@cvp56 root]$ cvpi status cvp

Current Running Command: None
Executing command. This may take a few seconds...
primary          17/17 components running
secondary        17/17 components running
tertiary         17/17 components running
[cvp@cvp56 root]$
```

Figure 382: cvpi status cvp shell command

3. Using vSphere client, shutdown one CVP node VM by selecting the node in the left pane, and then click the **Power off the virtual machine** option.

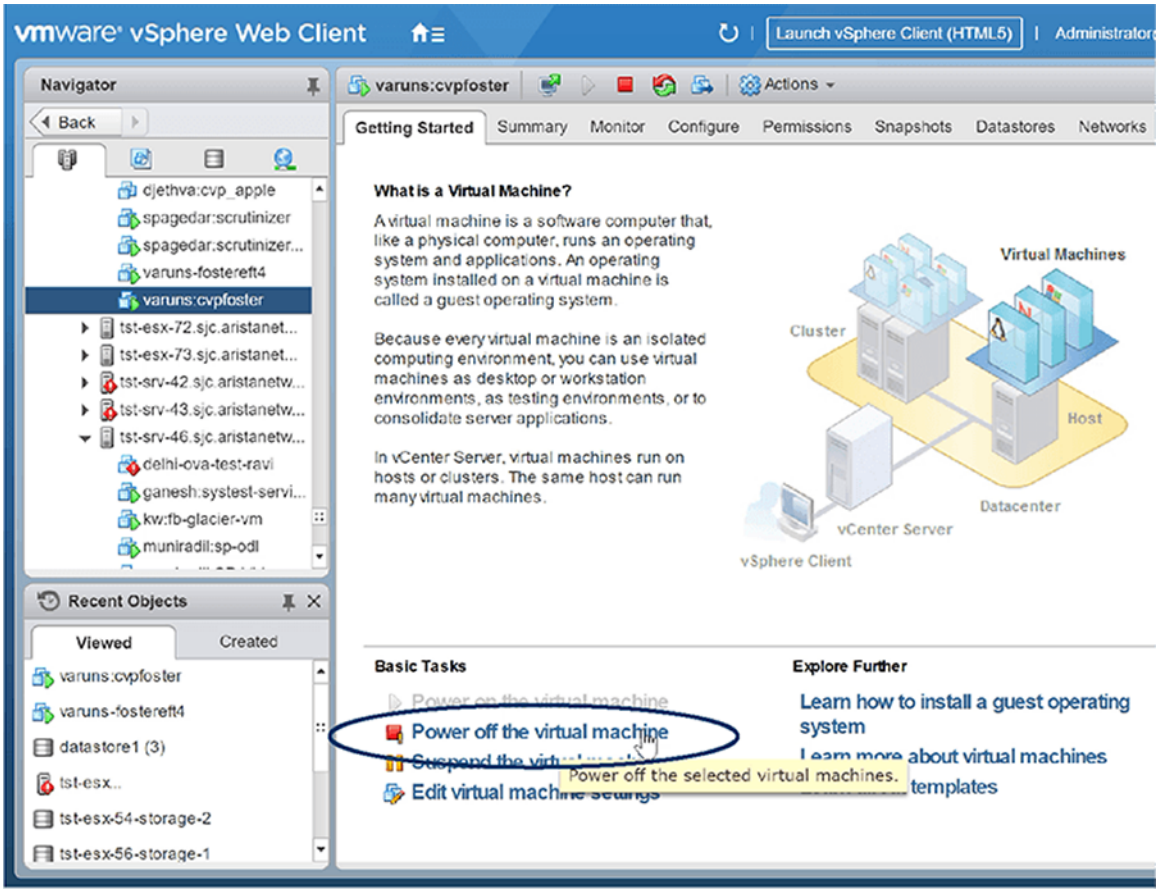

Figure 383: Power off the virtual machine

4. Click to confirm powering off the virtual machine.

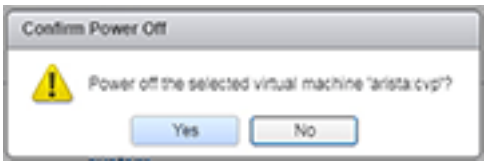

Figure 384: Powering off confirmation

5. On the CVP node VM, increase the memory allocation to 32GB by right-clicking the node icon, and then choose **Edit Settings**.

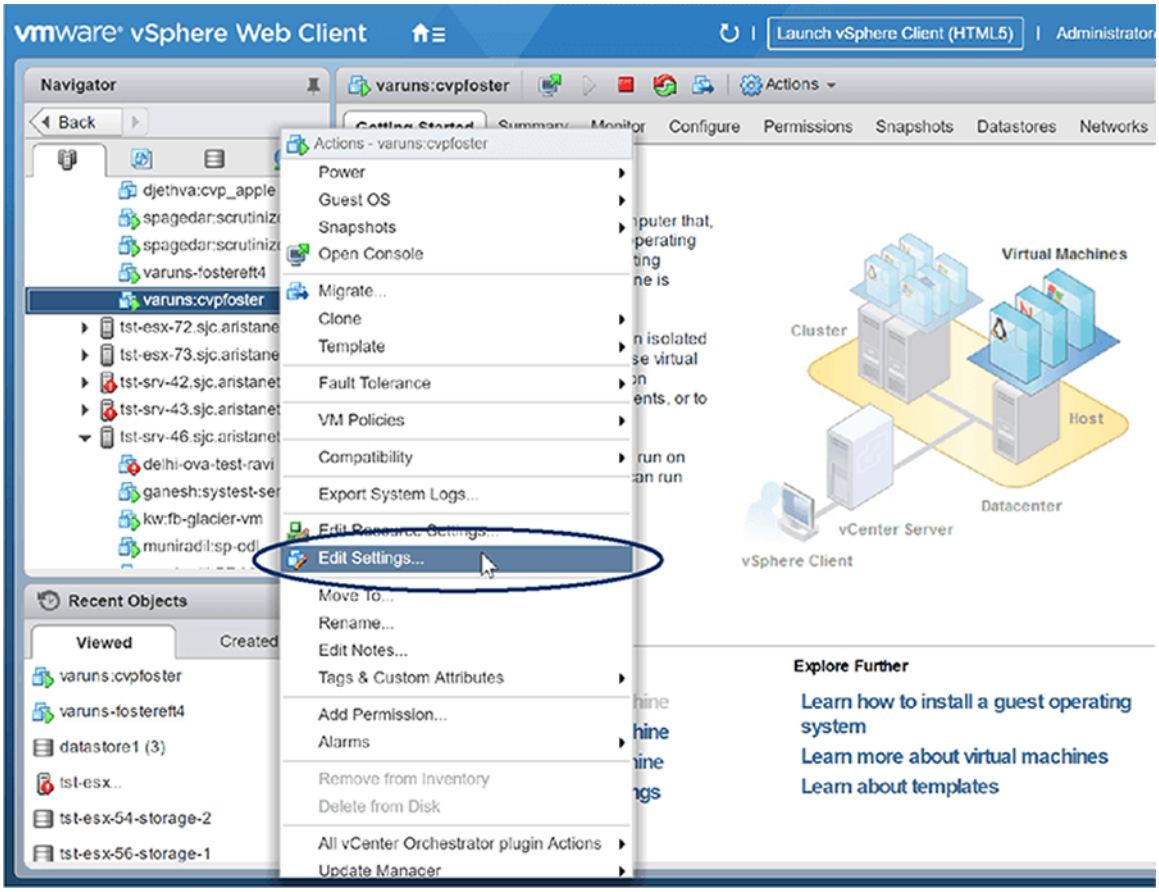

Figure 385: Edit Settings

The **Edit Resource Settings** dialog appears.

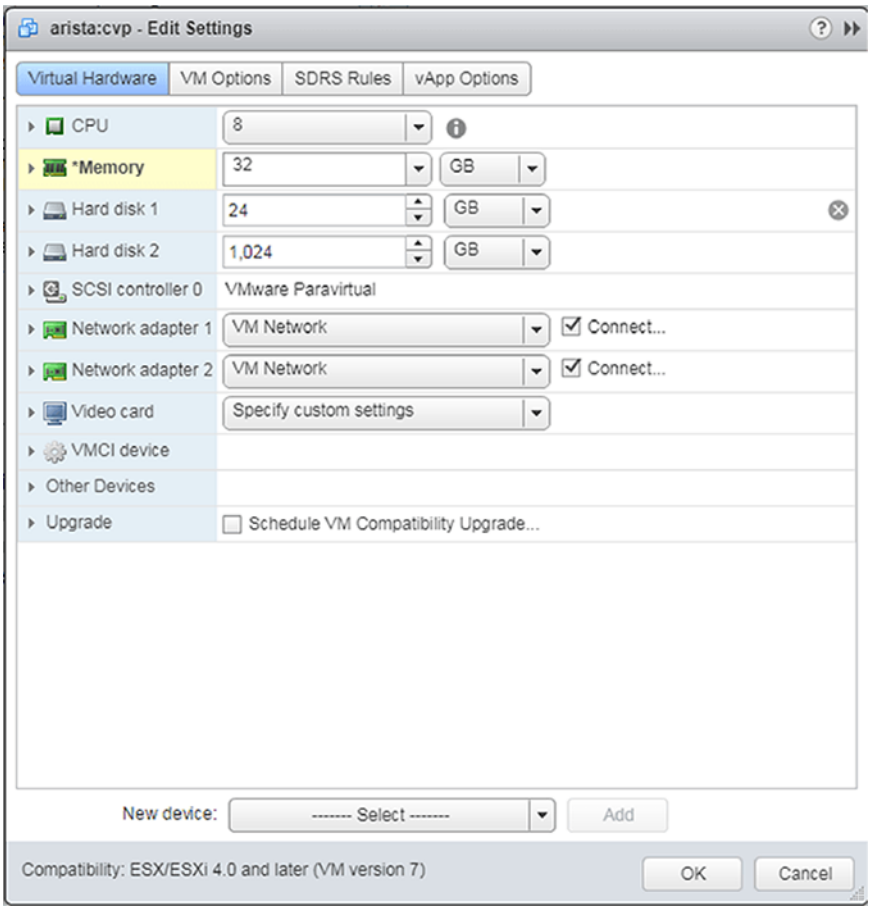

Figure 386: Edit Resources Settings

6. Do the following to increase the memory allocation for the CVP node VM:
   - Using the **Memory** option, click the up arrow to increase the size to **32GB**.
   - Click the **OK** button.

   The memory allocation for the CVP node VM is changed to 32GB. The page refreshes, showing options to power on the VM or continue making edits to the VM properties.
7. Click the **Power on the virtual machine** option.

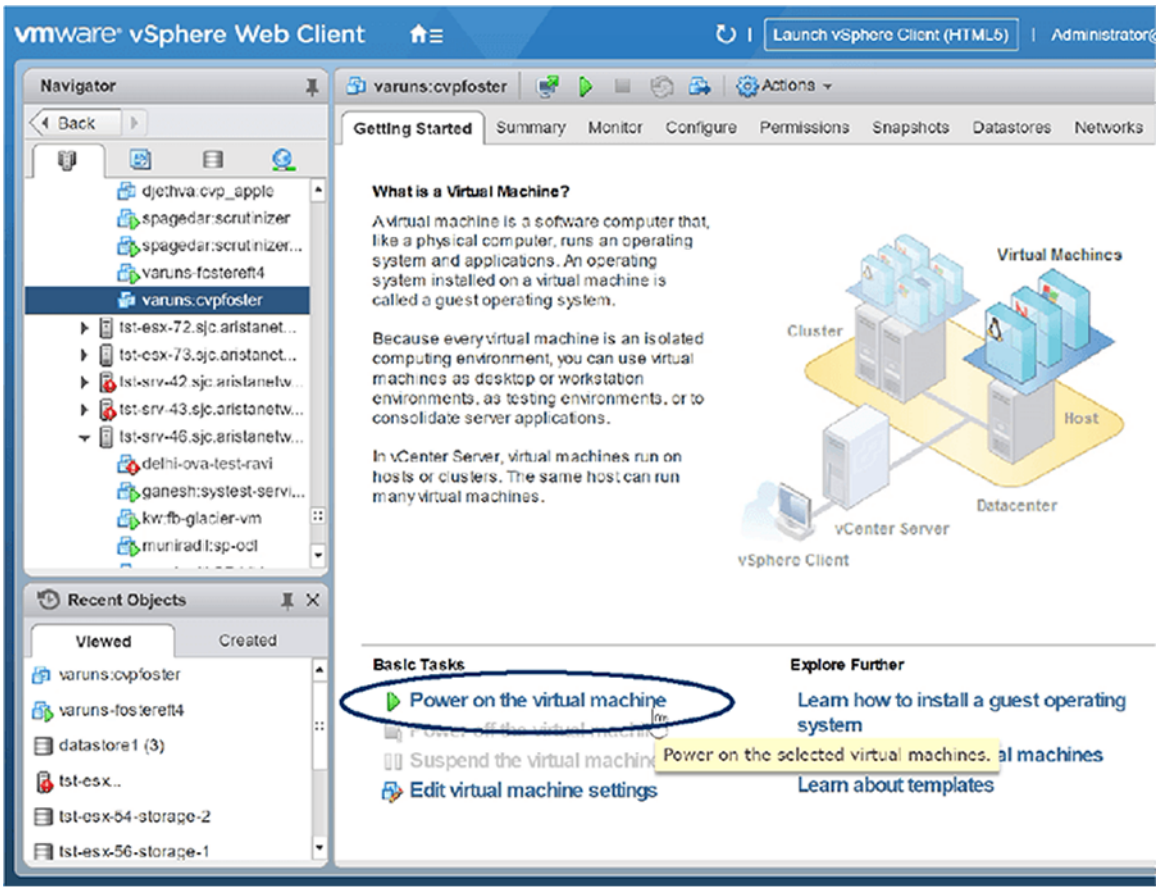

Figure 387: Power on the virtual machine

8. Wait for the cluster to reform.
9. Once the cluster is reformed, repeat **step 1 through step 7** one node at a time on each of the remaining CVP node VMs in the cluster.

Related topics:

- System Recovery
- Health Checks