



Back up CPS

- [Overview](#), on page 1
- [Back Up the Cluster Manager VM](#), on page 2
- [Back up CPS VMs](#), on page 4
- [Mongo Database](#), on page 5
- [Policy Builder Configuration Data](#), on page 7
- [Validating the Backup](#), on page 8
- [Back up Grafana Dashboard](#), on page 8

Overview

Various items in a Cisco Policy Suite (CPS) cluster require periodic backups.

This document describes the procedure to use the `config_br.py` script to back up the configuration and the data from a CPS cluster, and the steps to restore the backups, in case of any failures.

For general information about the `config_br.py` script and all the options, see the [About the Backup and Restore Script](#).

Before You Begin

Before you begin any backup and restore procedures, ensure you have performed the following tasks:

- Install CPS and have it running successfully. Backups are stored on customer-provided hardware, preferably in a location apart from where CPS is currently running.
- Initiate backups from the cluster manager VM. Ensure that there is adequate storage space on the cluster manager VM before taking a backup.
- `perfcient01/OAM01` VM is up and running for SVN repository backup.
- `lbvip01` VIP is available on Policy Director (LB) VMs.



Note Cisco recommends that the backup of the cluster manager be performed by taking a snapshot of the cluster manager VM. This operation requires administrator access to OpenStack or VMware (depending on the environment that CPS is deployed in).

Backup Schedule

Your first backup operation should occur after a successful installation and configuration. This provides a baseline and tests your backup procedures with respect to hardware, software, and protocols.

Then, do backup on this schedule as a best practice.

Table 1: Backup Schedule

Backup this...	...this often
Cluster Manager VM	Monthly and after any configuration changes, patch updates, or upgrades
Databases	Daily
Policy Builder Configurations	Weekly or after any changes

Back Up the Cluster Manager VM

The backup and restore procedures for the Cluster Manager do not require a maintenance window. The CPS cluster can continue to operate successfully without an operational Cluster Manager. Any CPS administrative scripts which use the Cluster Manager would not be usable while the Cluster Manager is offline.

The following sections describe two options for backing up the Cluster Manager VM:

- [Back Up the Cluster Manager VM in OpenStack](#)
- [Back Up the Cluster Manager VM in VMware](#)

It is not recommended to perform backups of the other CPS VMs. Instead, these VMs can be redeployed at any time. For more information, refer to [Restore a CPS VM](#).

Back Up the Cluster Manager VM in OpenStack

To back up the cluster manager VM in OpenStack, you must first create a snapshot of the VM.

Back up Cluster Manager VM

Step 1 Use the following command to view the nova instances and note the name of the cluster manager VM instance:

```
nova list
```

Step 2 Create a nova snapshot image as shown in the following command:

```
nova image-create --poll <cluman_instance_name> <cluman_snapshot_name>
```

Note Ensure that you have enough disk space for the snapshot.

Important In case if VM becomes unreachable after snapshot creation, check status of VM using `nova list` command. If it is in "SHUTOFF" state, you need to start the VM manually.

Step 3 View the image list with the following command:

```
nova image-list
```

Figure 1: Example Output

ID	Name	Status	Server
146719e8-d8a0-4d5a-9b15-2a669cfab81f	CPS_10.9.9_20160803_100301_112.iso	ACTIVE	
1955d56e-4ecf-4269-b53d-b30e73ad57f0	base_vm	ACTIVE	
2bbfb51c-cd05-4b7c-ad77-8362d76578db	cluman_snapshot	ACTIVE	4842ae5a-83a3-48fd-915b-6ca6361adb2c

Step 4 When a snapshot is created, the snapshot image is stored in OpenStack Glance. To store the snapshot in a remote data store, download the snapshot and transfer the file. To download the image, use the following command in OpenStack:

```
glance image-download --file <snapshot_name_on_filesystem> <snapshot_id>
```

For example:

```
glance image-download --file snapshot.raw 2bbfb51c-cd05-4b7c-ad77-8362d76578db
```

Step 5 List the downloaded images as shown in the following command:

```
ls -ltr *snapshot*
```

Example output:

```
-rw-r--r--. 1 root root 10429595648 Aug 16 02:39 snapshot.raw
```

Step 6 Store the snapshot of the Cluster Manager VM to restore in the future.

Back Up the Cluster Manager VM in VMware

The following section describe how to back up the entire Cluster Manager VM to a VMware OVF template. Backing up a Cluster Manager VM backs up all configurations and software applications.

Back up Cluster Manager Using OVF Template

To take the backup of the Cluster Manager, perform the following steps:

1. Shutdown the Cluster Manager VM using either of the following methods:

From the Cluster Manager VM: Log in to the Cluster Manager VM and run the following command to shutdown the VM.

```
shutdown -h now
```

From the vSphere Web Client:

- a. Log in to the vSphere server that hosts the Cluster Manager using vSphere Web Client.
- b. Right-click the Cluster Manager VM and select **Power > Power Off**.
A Confirm Power Off message appears. Click **Yes** to confirm the Power Off.
- c. Verify the Cluster VM is powered off from the vSphere Web Client UI.

2. Export the OVF template of the Cluster Manager VM.

- a. Select the Cluster Manager from the VM list on the left column.
- b. Select **Edit Settings...**
- c. Uncheck **Connected** near CD/DVD drive where you have linked the data store ISO file.
- d. Right-click the Cluster Manager and select **Template > Export OVF Template**.
The **Export OVF Template** dialog opens.
- e. In *Name* field, type the name of the template.
- f. (Optional) In the *Annotation* field, type a description.
- g. Uncheck “Include image files attached to floppy and CD/DVD devices in the OVF package”. By default, it is checked.
- h. Select the **Enable advanced options** checkbox if you want to include additional information or configurations in the exported template. The advanced settings include information about the BIOS UUID, MAC addresses, boot order, PCI Slot numbers, and configuration settings used by other applications.
- i. After selecting the required parameters, click **OK**. The backup starts.
- j. After the export succeeds, you are prompted to save each file associated with the template (.ovf, .vmdk, .iso). Save the files to the desired location.

Back up CPS VMs

To back up CPS VMs, see the following table for a description of the options that you use with the `config_br.py` script:

Table 2: `config_br.py` Script Options

VM	Options to include with script
Policy Director (LB)	<code>--network --haproxy --users</code>
OAM (pcrfclient)	<code>--etc-oam --svn --stats --grafanadb --auth-htpasswd --users</code>
QNS	<code>--users</code>
Session Manager	<code>--mongo --mongo-all --users</code>
All VMs	<code>--all</code>



Note For more information about each of the command options, see [Script Options](#).

The following examples illustrate how to use the command from Cluster Manager and options for various VMs:

- **Session Manager VM:** `config_br.py -a export --mongo-all --users /mnt/backup/sessionmgr_backup_27102016.tar.gz`
- **OAM VM:** `config_br.py -a export --etc-oam --svn --stats --grafanadb --auth-htpasswd --users /mnt/backup/oam_backup_27102016.tar.gz`
- **Policy Director (LB) VM:** `config_br.py -a export --network --haproxy --users /mnt/backup/lb_backup_27102016.tar.gz`
- **QNS VM:** `config_br.py -a export --users /mnt/backup/qns_backup_27102016.tar.gz`
- **All VMs:** `config_br.py -a export --all /mnt/backup/cps_backup_27102016.tar.gz`

To restore data from CPS VMs, see [Restore a CPS VM](#).

Mongo Database

In a production environment, databases need to use replication to help guarantee data integrity. Mongo DB calls its replication configuration replica sets as opposed to Master/Slave terminology used for Relational Database Management System (RDBMS).

Replica sets create a group of database nodes that work together to provide the data backup. There is a primary (the master) and 1..n secondaries (the slaves). Additionally, each replica set requires another node called the Arbiter. The Arbiter is used as a non-data-processing node that helps decide which node becomes the primary in the case of failure. For example, if there are four nodes: primary, secondary1, secondary2 and the arbiter, and if the primary fails, the remaining nodes “vote” for which of the secondary nodes becomes the primary. Since there are only two secondaries, there would be a tie and failover would not occur. The arbiter solves that problem and “votes” for one node breaking the tie.

Mongo DB has another concept called Sharding that helps redundancy and speed for a cluster. Shards separate the database into indexed sets which allow for much greater speed for writes which improves overall database performance. Sharded databases are often setup so that each shard is a replica set. Replica Sets and Sharding both require some special handling for backup. Mongo DB recommends that for each replica set being backed up, one secondary is shut down and that node is used for the backup. After backup, that node is brought back up and integrated back into the replica set.

Mongo Database Backup



Note RADIUS-based policy control is no longer supported in CPS 14.0.0 and later releases as 3GPP Gx Diameter interface has become the industry-standard policy control interface.

CPS uses MongoDB for primary system databases. These include:

- Admin
- Audit
- Balance
- Custom Reference Data

- Policy Reporting
- Portal
- Radius
- Sharding
- SPR
- Vouchers

The session database (session_cache) represents the transient session data of the active network sessions of subscribers on the network. Due to the transient nature of this data, it does not make sense to backup or restore this database. The script does not provide an option to back up the Session Cache database (configured by default on port 27717).

Full Environment:

The following table lists the Module Name, the database name, and the default ports when using Replica Sets via the `/etc/broadhop/mongoConfig.cfg` file.

Table 3: Using Replica-set

Module Name	Database Name	Default Ports
Core	admin	27721
Audit	audit	27725
Balance	balance_mgmt	27718
Custom Reference Data	cust_ref_data	27717
Policy Intel	policy_trace	27719
Portal	portal	27749
Radius	radius	27717
Core	sharding	27717
SPR	spr	27720
Voucher	vouchers	27717

General Procedure for Database Backup

Use the following command to generate a backup of the CPS databases.

```
config_br.py -a export --mongo-all /mnt/backup/backup_28092016.tar.gz
```

For reference, the following Mongo DB documentation was used to develop the CPS backup procedures:

<http://docs.mongodb.org/manual/tutorial/backup-sharded-cluster-with-database-dumps/>

Automatic Backup via Cron

Using a cron job, it is possible to automate backups. It is best to schedule automated backups when least amount of traffic is running through the CPS system.



Note Do not store repository backups on any CPS node. Move them immediately to an external storage like a Storage Area Network (SAN).

The following example procedure creates a backup of the policy configuration subversion repository every night at 10:00pm:

1. Login to the Cluster Manager as the root user.
2. To edit the root user's cron tab, execute the command:

```
crontab -e
```

3. Add the following line:

```
22 * * * config_br.py -a export --svn /mnt/backup/backup_$(date +%Y-%m-%d).tar.gz
```



Note The crontab editor is VI.

Save the file and the new cron tab is installed.

Policy Builder Configuration Data

The Policy Builder uses a Subversion (SVN) repository to store the policy configurations. The following sections outline the backup and restore procedures for the Subversion repository.

Subversion Repository Backup

The Subversion repository is setup like a master/slave with the master repository in pcrfclient01 and the slave repository in pcrfclient02. All commits go to the master and are replicated to the slave using the Subversion hooks process. Hooks are scripts that get executed by the SVN binary automatically. Typically in deployments, policy configuration does not change very often once the system is live, so automated weekly backups of the repository are usually sufficient.

Automatic Backup via Cron

Using a cron job, it is possible to automate backups. It is best to schedule automated backups when least amount of traffic is running through the CPS system.



Note Do not store repository backups on any CPS node. Move them immediately to an external storage like a Storage Area Network (SAN).

The following example procedure creates a backup of the policy configuration subversion repository every night at 10:00pm:

1. Login to the Cluster Manager as the root user.
2. To edit the root user's cron tab, execute the command:

```
crontab -e
```

3. Add the following line:

```
22 * * * config_br.py -a export --svn /mnt/backup/backup_$(date +%Y-%m-%d).tar.gz
```



Note The crontab editor is VI.

Save the file and the new cron tab is installed.

Validating the Backup

After you make a backup of any database, you can check these things to make sure the backup is valid:

- Observe and correct any errors or warnings during the backup. For example, the backup may be aborted if there is not enough file space available or if the media is corrupt.
- Make sure that the file size of the backup is the same as the original, and that it is not zero.

Open the backup database with an appropriate third-party tool.

With these instructions, your backup routines should be adequate and timely. If in doubt, try to restore backups to a test environment and gauge your success. Please contact your Cisco technical representative at any time with questions or concerns.

Back up Grafana Dashboard

You can back up Grafana dashboard using following command:

```
config_br.py -a export --grafanadb /mnt/backup/backup_$(date +%Y-%m-%d).tar.gz
```