



Cisco Policy Suite 7.0.1 Backup and Restore Guide

First Published: January 28, 2015

Last Updated: July 10, 2015

Cisco Systems, Inc.

www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Policy Suite 7.0.1 Backup and Restore Guide
© 2015 Cisco Systems, Inc. All rights reserved.



Preface v

Readers v

Additional Information v

CHAPTER 1

Backup and Restore 1-1

Before You Begin 1-1

Overview 1-1

Backup Strategies 1-2

Backup Schedule 1-2

Cluster Manager 1-2

Cluster Manager Backup 1-2

Cluster Manager Restore 1-4

Mongo Database 1-5

Mongo Database Backup 1-5

General Procedure for Database Backup 1-6

Automatic Database Backup via Cron 1-7

Mongo Database Restore 1-7

General Procedure for Database Restore 1-8

Policy Builder Configuration Data 1-9

Subversion Repository Backup 1-9

Automatic Repository Backup via Cron 1-9

Subversion Repository Restore 1-10

Virtual Machine Backup/Restore 1-10

Virtual Machine Backup 1-10

Virtual Machine Restore 1-12

Full Environment (HA) 1-12

All In One Environment (AIO) 1-12

Validating the Backup 1-13

Next Steps 1-13



Preface

Read about these topics in these sections:

- [Readers, page v](#)
- [Additional Information, page v](#)

Readers

This guide is for:

- Deployment engineers
- Network engineers
- System engineers

You should be familiar with Linux, MySQL, and SVN as well as general backup and restore procedures at an intermediate level.

Additional Information

This document assumes an intermediate level of understanding of network architecture, configuration, and operations.





Backup and Restore

Revised: July 10, 2015

This chapter covers the following sections:

- [Before You Begin, page 1-1](#)
- [Overview, page 1-1](#)
- [Backup Strategies, page 1-2](#)
- [Backup Schedule, page 1-2](#)
- [Cluster Manager, page 1-2](#)
- [Mongo Database, page 1-5](#)
- [Policy Builder Configuration Data, page 1-9](#)
- [Virtual Machine Backup/Restore, page 1-10](#)
- [Validating the Backup, page 1-13](#)
- [Next Steps, page 1-13](#)

Before You Begin

- 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 the backups using either manual or automated methods.

Overview

There are various items from an CPS system that should be backed up. This document focuses on the following three items for backup and restore:

- Cluster Manager
- Databases
- Policy Builder Configuration

Backup Strategies

There are several methods to back up these data areas:

- Use your own company policies and tools.
- Use the instructions provided here.

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 backups on this schedule as a best practice.

Backup this...	...this often
VMs	Monthly
Databases	Daily
Policy Builder Configurations	Weekly or if there are any changes

Cluster Manager

Cluster Manager is the main server that contains CPS cluster software and configurations. In case a CPS VM is corrupted, the VM can be recreated by deploying a new VM from the Cluster Manager. Since Cluster Manager is already staged with the correct software image and configuration, deploying a new VM node is very simple.

This section covers the following topics:

- [Cluster Manager Backup](#)
- [Cluster Manager Restore](#)

Cluster Manager Backup

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



Note

The administrator must schedule a maintenance window to backup the Cluster Manager.


Step 1

Shutdown the Cluster Manager VM.

- Login to the Cluster Manager Linux shell and execute shutdown the server command.

```
shutdown -h now
```

- Administrator can also shutdown the Cluster Manager using vSphere client. Login to the vSphere server that hosts the Cluster Manager using vSphere client.
- Select the Cluster Manager VM, right-click and select **Power > Power Off**.

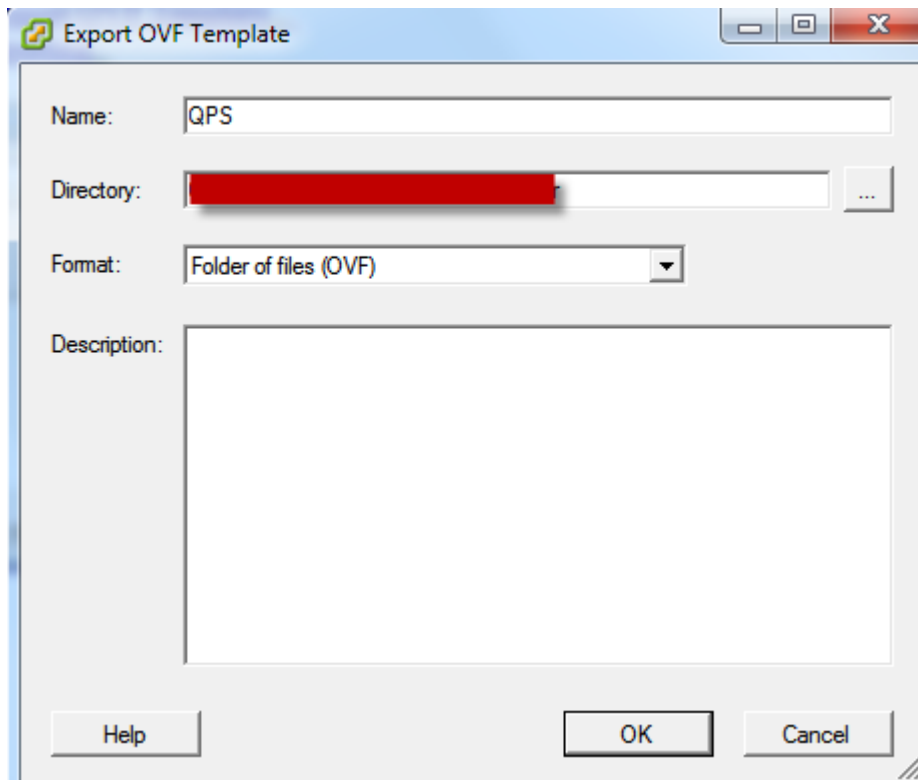


A **Confirm Power Off** message appears. Click **Yes** to confirm the Power Off.

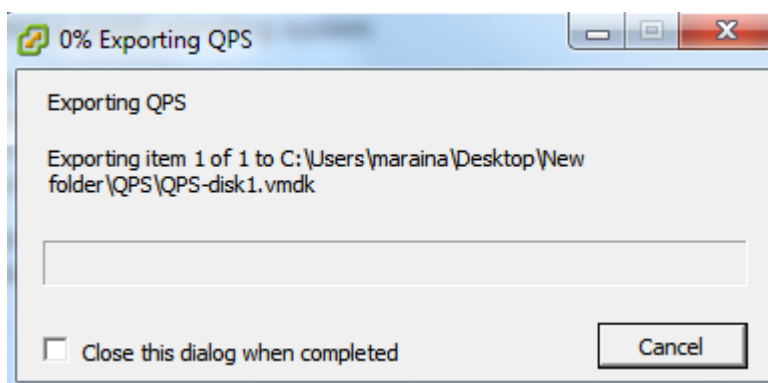
- d. Verify the Cluster VM is powered off from the vSphere Client UI.

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

- a. Select the Cluster Manager from the VM list on the left column.
- b. From the menu, select **File > Export > Export OVF Template**.



- c. Enter the VM name in **Name** field. Also set the directory path where you want to save the backup.
- d. After selecting the required parameters, click **OK**. The backup starts as shown below.



- e. After the export succeeds, save the OVF file.

Cluster Manager Restore

Before restoring the Cluster Manager, configure the ESX server to have enough memory and CPU available. (Refer to *CPS 7.0.1 Installation Guide* for CPU/Memory/Disk requirements). Confirm that the network port group is configured for the internal network.

-
- Step 1** Login to ESX server using vSphere Client.
 - Step 2** Select File > Deploy OVF Template... The Deploy OVF Template wizard appears.
 - Step 3** Click **Browse** to select the required OVF template file and click **Next**.
 - Step 4** Enter the name for the VM in the **Name** field and click **Next**.
 - Step 5** Select the required destination datastore from the **Storage** window where the OVF template will be deployed and click **Next**.
 - Step 6** From **Disk Format** window, select the format in which you want to store the template and click **Next**.
 - Step 7** From **Network Mapping** window, select the network (map the networks used in OVF template to the network in your inventory) and click **Next**.
 - Step 8** Verify the settings from **Ready to Complete** window and click **Finish**.
 - Step 9** After the OVF template is successfully deployed, power on the VM. The Cluster Manager VM is deployed successfully.

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 many other 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.

This section covers the following topics:

- [Mongo Database Backup](#)
- [Mongo Database Restore](#)

Mongo Database Backup

CPS uses Mongo DB for primary system databases. These include:

- Admin
- Audit
- Balance
- Customer Reference Data

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

The Session database (session_cache) runs on 27717 and can be backed up as well, but we strongly discourage backing up the session_cache because it is not useful for production environments. The session database represents transient session data of active network sessions of subscribers on the network and therefore is never the same as what is actually occurring in the network when restoring the data.

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.

Module Name	Database Name	Default Ports
Core	admin	27721
Audit	audit	27725
Balance	balance_mgmt	27718
Customer 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

All-in-One (AIO) Environment:

By default, in AIO environment all databases run on port 27017 except the portal which runs on port 27749 like in the multi-node environment.

General Procedure for Database Backup

Mongo DB provides various tools to assist with database backups. `mongodump` is the recommended tool for the CPS environment.

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/>

```
env_export.sh --mongo /var/tmp/env_export_$(date).tgz
```

where `$date` is the date when the backup file was created.

For example,

```
env_export.sh --mongo /var/tmp/env_export_2014-10-30.tgz
```

Automatic Database 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 database backups on any CPS node. Move them immediately to the Cluster Manager for removal to external storage like a Storage Area Network (SAN).


Note

Because the export script may be used by multiple cron entries, it implements a waiting function. If the script detects that another process is already running, it will wait until the other process has completed and then continue. To avoid conflicts and overwriting data due to multiple cron entries running the script for automated backups, make sure that each cron entry has a unique export file name that includes a timestamp. For example, `/var/tmp/export_mongo_$(date +%Y-%m-%d).tgz`.

The following example creates a backup of the default database set (Admin, Balance, Customer Reference Data, Sharding, and SPR) every night at 10:00 pm:

Step 1 Login to the Cluster Manager (ssh or console login through VMware client) as a root user.

Step 2 To edit the root user's cron tab, execute the command:

```
crontab -e
```

Step 3 Add the following line:

```
22 * * * /var/qps/bin/support/env/env_export.sh --mongo /var/tmp/env_export_mongo_$(date
+%Y-%m-%d).tgz
```


Note

The crontab editor is VI.

Step 4 Save the file and the new cron tab is installed.

Mongo Database Restore

To restore databases in a production environment that use replica sets with or without sharding, a maintenance window is required as the CPS software on all the processing nodes and the sessionmgr nodes must be stopped. The assumption is that the only time the customer does a restore is after an outage or problem with the system and/or its hardware. In that case, service has been impacted and to properly fix the situation, service will need to be impacted again. From a database perspective, the main processing nodes must be stopped so that the system is not processing incoming requests while the databases are stopped and restored. If replica sets are used with or without sharding, then all the database instances must be stopped to properly restore the data and have the replica set synchronize from the primary to the secondary database nodes.

For reference, the following Mongo DB documentation was used to develop the CPS restore procedures.

- <http://docs.mongodb.org/manual/tutorial/restore-sharded-cluster/#restore-sh-cl-dmp>

- <http://docs.mongodb.org/manual/tutorial/restore-replica-set-from-backup/>
- <http://docs.mongodb.org/manual/tutorial/resync-replica-set-member/>

General Procedure for Database Restore

Step 1 Execute the following command to restore the database:

```
/var/qps/bin/support/env/env_import.sh --mongo /var/tmp/env_export_mongo_$date.tgz
```

where *\$date* is the timestamp when the export was made.

For example,

```
env_import.sh --mongo /var/tmp/env_export_2014-10-30.tgz
```

Step 2 Log into the database and verify whether it is running and is accessible:

a. Log into session manager:

```
mongo --host sessionmgr01 --port $port
```

where *\$port* is the port number of the database to check.

For example, 27718 is the default Balance port.

b. Display the database by executing the following command:

```
show dbs
```

c. Switch the mongo shell to the database by executing the following command:

```
use $db
```

where *\$db* is a database name displayed in the previous command. The 'use' command switches the mongo shell to that database.

For example,

```
use balance_mgmt
```

d. To display the collections, execute the following command:

```
show collections
```

e. To display the number of records in the collection, execute the following command:

```
db.$collection.count()
```

For example,

```
db.account.count()
```

The above example will show the number of records in the collection "account" in the Balance database (balance_mgmt).

Policy Builder Configuration Data

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

- [Subversion Repository Backup](#)
- [Subversion Repository Restore](#)

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 Repository 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 the Cluster Manager for removal to external storage like a Storage Area Network (SAN).

**Note**

Because the export script may be used by multiple cron entries, it implements a waiting function. If the script detects that another process is already running, it will wait until the other process has completed and then continue. To avoid conflicts and overwriting data due to multiple cron entries running the script for automated backups, make sure that each cron entry has a unique export file name that includes a timestamp. For example, `/var/tmp/export_mongo_$(date +%Y-%m-%d).tgz`.

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

Step 1 Login to the Cluster Manager (ssh or console login through VMware client) as a root user.

Step 2 To edit the root user's cron tab, execute the command:

```
crontab -e
```

Step 3 Add the following line:

```
22 * * * /var/qps/bin/support/env/env_export.sh --svn /var/tmp/env_export_svn_$(date +%Y-%m-%d).tgz
```

**Note**

The crontab editor is VI.

Save the file and the new cron tab is installed.

Subversion Repository Restore

To restore the Policy Builder Configuration Data from a backup, execute the following command:

```
/var/qps/bin/support/env/env_import.sh --svn /var/tmp/env_export_svn_<math>\$date</math>.tgz
```

where $\$date$ is the date when the cron created the backup file.

Virtual Machine Backup/Restore

This section covers the following topics:

- [Virtual Machine Backup](#)
- [Virtual Machine Restore](#)

Virtual Machine Backup

CPS software is based on the strategy of virtual machines rather than physical machines. Backing up an installer/AIO virtual machine backs up configurations and software applications and we do not recommend backups of individual VMs for full system (HA).

Read this section to backup an installer/AIO virtual machine with vSphere.

To back up installer VM in its entirety, perform the following steps:

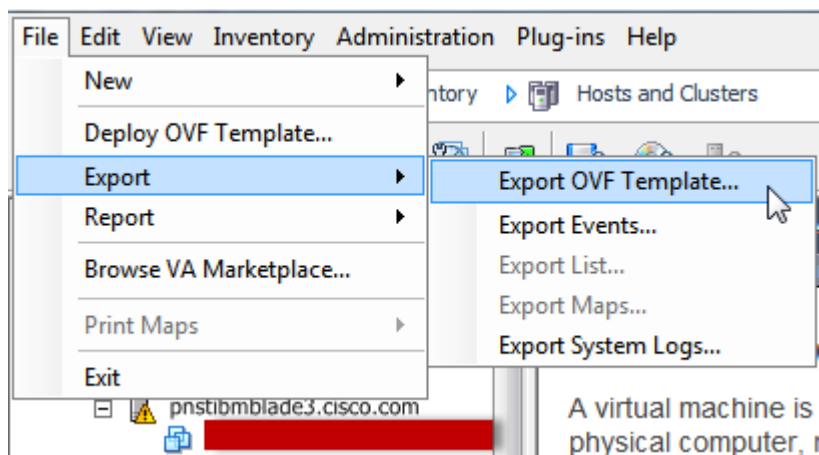
Step 1 Open your vSphere client and log into the ESXi/ESX machine.

Step 2 Power down the virtual machine (VM).




Note In Linux, type **init 0** to shutdown the VM cleanly.


Step 3 When the virtual machine is powered off, select it and select **File > Export > Export OVF Template**.



Step 4 Export OVF Template opens up.



- Step 5** Enter the VM name in **Name** field. Also set the directory path where you want to save the backup.
- Step 6** Select the format from **Format** drop-down list. You can select **Folder of files (OVF)** or **Single file (OVA)**.
- Step 7** Uncheck “Include image files attached to floppy and CD/DVD devices in the OVF package”, By default, it is checked.
- Step 8** After selecting the required parameters, click **OK**. The backup starts as shown below.



Virtual Machine Restore

Installer VM restoring can be done on the same setup from where the backup was taken. Open your vSphere client and log into the ESXi/ESX machine.

Full Environment (HA)

-
- Step 1** Deploy the installer VM backup using vSphere by select File > Deploy OVF Template
- Step 2** Power ON installer VM.
- Step 3** SSH to installer VM.
- Step 4** Re-deploy all the VMs by executing the following commands:
- ```
cd /var/qps/install/current/scripts/deployer/support/
python deploy_all.py
```
- Step 5** Stop all the CPS processes by executing the following command:
- ```
/var/qps/bin/control/stopall.sh
```
- Step 6** Create replica-sets for all sets defined in configuration file, `/etc/broadhop/mongoConfig.cfg`:
- ```
/var/qps/bin/support/mongo/build_set.sh --all -create
```
- Step 7** Transfer all database dumps to installer.
- Step 8** Restore all database dumps by executing the following command:
- ```
/var/qps/bin/support/mongo/import_db.sh <portal/admin/spr.> <.....dump.tar.gz>
```
- Step 9** Verify database status by executing the following command:
- ```
/var/qps/bin/diag/diagnostics.sh --get_replica_status
```
- Step 10** To restore policy information, refer to the section [Subversion Repository Restore](#).
- Step 11** Start all CPS processes by executing the following command:
- ```
/var/qps/bin/control/startall.sh
```
- Step 12** Verify working system by executing the following command:
- ```
/var/qps/bin/diag/diagnostics.sh
```

### All In One Environment (AIO)

- 
- Step 1** Deploy the AIO VM backup using vSphere by selecting File > Deploy OVF Template.
- Step 2** Power ON AIO VM.
- Step 3** SSH to AIO.
- Step 4** Validate setup by executing the following command:

```
/var/qps/bin/diag/diagnostics.sh
```

## 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.

## Next Steps

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.

