CPS Basic Operations

Starting and Stopping CPS

This section describes how to start and stop Cisco Policy Server nodes, VMs, and services.

Starting VMs Using VMware GUI

Step 1 Start a VMware vSphere session.

Step 2 Right-click the VM and select **Power** > **Power On**.

Important If the Policy Server (QNS) VM was previously powered off, it must be powered on only during Maintenance Window or low traffic time. If the VM is powered on during high traffic, then when the qns java process comes up and it immediately starts taking up load. As a result there can be timeouts and high CPU until around 60 seconds from the Policy Server (QNS) VM during the JVM hotspot warmup time. Once the JVM warmup phase is completed, the VM must be able to handle traffic smoothly.

Step 3 After the VM has started, log into the VM from Cluster Manager and verify that the processes are running.

Shutting Down the Cisco Policy Server Nodes

The following sections describe the commands to shut down the Cisco Policy Server nodes:

Policy Director (LB) or Policy Server (QNS) Nodes

Step 1 SSH to the lbxx or qnsxx node from Cluster Manager:
`ssh lbxx` or `ssh qnsxx`

Step 2 Stop all CPS processes on the node:
`/usr/bin/monit stop all`

Step 3 Check the status of all the processes. Verify that all processes are stopped before proceeding.
`/usr/bin/monit summary`

Step 4 Stop the monit process:
`service monit stop`

Step 5 Shut down lbxx/qnsxx:
`shutdown -h now`

OAM (pcrfclient) Nodes

Step 1 SSH to the pcrfclient*xx* node from Cluster Manager:
`ssh pcrfclient`*xx*

Step 2 Stop all CPS processes on the node:
`/usr/bin/monit stop all`

Step 3 Check the status of all the processes. Verify that all processes are stopped before proceeding:
`/usr/bin/monit summary`

Step 4 Stop the monit process:
`service monit stop`

Step 5 Stop the licenses process:
`service lmgrd stop`

Step 6 Shut down pcrfclient*xx*:
`shutdown -h now`

sessionmgr Nodes

Step 1 SSH to the sessionmgr*xx* node from Cluster Manager:
```
ssh sessionmgrxx
```

Step 2 Stop all CPS processes on the node:
```
/usr/bin/monit stop all
```

Step 3 Check the status of all the processes. Verify that all processes are stopped before proceeding:
```
/usr/bin/monit summary
```

Step 4 Stop the monit process:
```
service monit stop
```

Step 5 For CPS nodes, such as sessionMgrs, there are mongo processes running that require special steps to stop. First, determine which processes are running by executing:
```
ls /etc/init.d/sessionmgr*
```

Step 6 Make sure the mongo replica set is in secondary:
```
/usr/bin/mongo --port $PORT --eval "rs.stepDown(10)"
```
where, PORT is the port number found in the previous step, such as 27717.

Step 7 Stop the MongoDB processes.
For example:
```
service sessionmgr-27717 stop
```

Step 8 Shut down sessionmgr*xx*:
```
shutdown -h now
```

Restarting the Cisco Policy Server

CPS is composed of a cluster of nodes and services. This section describes how to restart the different services running on various CPS nodes.

Restarting Database Services

Each database port and configuration is defined in the `/etc/broadhop/mongoConfig.cfg` file.

The scripts that start/stop the database services can be found in the `/etc/init.d/` directory on the CPS nodes.

To stop and start a database, log into each Session Manager VM and execute the commands as shown below. For example, to restart the sessionmgr 27717 database, execute:

```
service sessionmgr-27717 stop
```

```
service sessionmgr-27717 start
```

or:

```
service sessionmgr-27717 restart
```

Note It is important not to stop and start all of the databases in the same replica set at the same time. As a best practice, stop and start databases one at a time to avoid service interruption.

Restarting Policy Server Services

If the Policy Server (QNS) VM was previously powered off, it must be powered on only during Maintenance Window or low traffic time. If the VM is powered on during high traffic, then when the qns java process comes up and it immediately starts taking up load. As a result there can be timeouts and high CPU until around 60 seconds from the Policy Server (QNS) VM during the JVM hotspot warmup time. Once the JVM warmup phase is completed, the VM must be able to handle traffic smoothly.

Restarting All Policy Server Services

To restart all Policy Server (QNS) services on all VMs, execute the following from the Cluster Manager:

```
/var/qps/bin/control/restartall.sh
```

Note This script only restarts the Policy Server (QNS) services. It does not restart any other services.

Use summaryall.sh or statusall.sh to see details about these services.

Restarting All Policy Server Services on a Specific VM

To restart all Policy Server (QNS) services on a single CPS VM, execute the following from the Cluster Manager:

```
/var/qps/bin/control/restartqns.sh <hostname>
```

where <*hostname*> is the CPS node name of the VM (qns01, qns02, lb01, pcrfclient01, and so on).

Restarting Individual Policy Server Services on a Specific VM

Step 1 Log into the specific VM.

Step 2 To determine what Policy Server (QNS) services are currently running on the VM, execute:

```
monit summary
```

Output similar to the following appears:

```
The Monit daemon 5.5 uptime: 1d 17h 18m
```

```
Process 'qns-4' Running
Process 'qns-3' Running
Process 'qns-2' Running
Process 'qns-1' Running
```

Step 3 Execute the following commands to stop and start the individual Policy Server (QNS) process:

```
monit stop qns-<instance id>
```

```
monit start qns-<instance id>
```

Restarting Services Managed by Monit

The Monit service manages many of the services on each CPS VM.

To see a list of services managed by `monit` on a VM, log in to the specific VM and execute:

```
monit summary
```

To stop and start all services managed by `monit`, log in to the specific VM and execute the following commands:

```
monit stop all
```

```
monit start all
```

To stop and start a specific service managed by Monit, log in to the specific VM and execute the following commands:

```
monit stop <service_name>
```

```
monit start <service_name>
```

where <*service_name*> is the name as shown in the output of the `monit summary` command.

Restarting Other Services

Restarting Subversion

To restart Subversion (SVN) on OAM (pcrfclient) nodes, execute:

```
service httpd restart
```

Restarting Policy Builder

To restart Policy Builder on OAM (pcrfclient) nodes (pcrfclient01/pcrfclient02), execute:

```
monit stop qns-2
```

```
monit start qns-2
```

Restarting Control Center

To restart Control Center on OAM (pcrfclient) nodes (pcrfclient01/pcrfclient02), execute:

```
monit stop qns-1

monit start qns-1
```

Restarting Services on Policy Director (lb01 and lb02)

The following commands are used to restart the services on the Policy Director (lb) nodes only (lb01 and lb02).

Step 1 Login to lb01/lb02.

Step 2 To restart the service that controls the virtual IPs (lbvip01 and lbvip02 are virtual IP addresses shared between lb01 and lb02 for High Availability), execute the following command:

```
monit restart corosync
```

Step 3 To restart the service that balances and forwards IP traffic (port forwarding service) from lb01/lb02 to other CPS nodes, execute:

```
monit restart haproxy
```

Recovering After a Power Outage

If there is a controlled or uncontrolled power outage, the following power on procedures should be followed to bring the system up properly.

Step 1 Power ON the Cluster Manager.

Step 2 Power ON pcrfclient01.

Step 3 Power ON all Session Manager nodes (sessionmgr0x).

Step 4 Validate that the databases are all online by running:

```
diagnostics.sh --get_replica_status
```

Step 5 Power ON Policy Director node 2 (lb02).

Step 6 Power ON Policy Director node 1 (lb01).

Step 7 Power ON all Policy Server (QNS) nodes.

Step 8 Power ON pcrfclient02.

Step 9 On pcrfclient01, run the following commands to reinitialize the services:

```
monit stop all

monit start all
```

Step 10 Run `diagnostics.sh` to validate system is functioning properly.

Recovery Control

Due to the operational inter-dependencies within the CPS, it is necessary for some CPS services and components to become active before others.

CPS can monitor the state of the cluster through the various stages of startup. It also includes functionality to allow the system to gracefully recover from unexpected outages.

Cluster State Monitoring

CPS can monitor the state of the services and components of the cluster from the OAM (pcrfclient) VMs. By default, this functionality is disabled.

This functionality can be enabled by setting the cluster_state_monitor option to true in the CPS Deployment Template (Excel spreadsheet).

To update an existing deployment to support this functionality, modify this setting in your CPS Deployment Template and redeploy the csv files as described in the *CPS Installation Guide for VMware*.

This monitoring system reports the state of the system as an integer value as described in the following table:

Table 1: Cluster State Monitoring

Cluster State	Description	Values
0	unknown state/pre-inspection state	The system will report '0' until both conditions have been met under '1': lbvip02 is UP AND databases are accessible. Various systems can be coming online while a '0' state is being reported and does not automatically indicate an error. Even if the system cannot proceed to '1' state, Policy Builder and Control Center UIs should be available in order to manage or troubleshoot the system.
1	lbvip02 is alive and all databases in `/etc/broadhop/mongoConfig.cfg` have an accessible primary	All backend databases must be available and the lbvip02 interface must be UP for the system to report this state.

Cluster State	Description	Values
2	lbvip02 port 61616 is accepting TCP connections	Backend Policy Server (QNS) processes access lbvip02 on this port. When this port is activated, it indicates that Policy Server (QNS) processes can proceed to start.
3	at least 50% of backend Policy Server (QNS) processes are alive	Once sufficient capacity is available from the backend processes, the Diameter protocol endpoint processes are allowed to start.

The current cluster state is reported in the following file on the OAM (pcrfclient): `/var/run/broadhop.cluster_state`

The `determine_cluster_state` command logs output of the cluster state monitoring process into `/var/log/broadhop/determine_cluster_state.log`.

Controlled Startup

In addition to the monitoring functionality, CPS can also use the cluster state to regulate the startup of some of the CPS services pending the appropriate state of the cluster.

By default this functionality is disabled. It can be enabled for the entire CPS cluster, or for troubleshooting purposes can be enabled or disabled on a per-VM basis.

Note Cluster State Monitoring must be enabled for Controlled Startup to function.

Enable/Disable For All VMs in Cluster

The Controlled Startup functionality is enabled by the presence of the `/etc/broadhop/cluster_state` file.

To enable this feature on all CPS VMs in the cluster, execute the following commands on the Cluster Manager VM to create this file and to use the syncconfig.sh script to push those changes out to the other VMs.

```
touch /etc/broadhop/cluster_state
```

```
syncconfig.sh
```

To disable this feature on all VMs in the cluster, remove the cluster_state file on the Cluster Manager VM and sync the configuration:

```
rm /etc/broadhop/cluster_state
```

```
syncconfig.sh
```

Enable/Disable For Specific VM

To enable this feature on a specific VM, create a `/etc/broadhop/cluster_state` file on the VM:

touch /etc/broadhop/cluster_state

To disable this feature again on a specific VM, delete the `/etc/broadhop/cluster_state` file on the VM:

rm /etc/broadhop/cluster_state

Note This is temporary measure and should only be used for diagnostic purposes. Local modifications to a VM can be overwritten under various circumstances, such as running **syncconfig.sh**.

Switching Active and Standby Policy Directors

In CPS, the active and standby strategy applies only to the Policy Directors (lb). The following are the two Policy Directors in the system:

- lb01
- lb02

Determining the Active Policy Director

Step 1 Log in to the pcrfclient01 VM.

Step 2 Run the following command to SSH to the active Policy Director (typically lb01):

```
ssh lbvip01
```

Step 3 You can also confirm an active Policy Director by ensuring it has the virtual IP (VIP) associated with it by running the following command:

```
ifconfig -a
```

If you see the eth0:0 or eth1:0 interfaces present in the list and marked as “UP” then that is the active Policy Director.

For example:

```
eth0:0   Link encap:Ethernet  HWaddr 00:0C:29:CD:7E:4C
         inet addr:172.26.241.240  Bcast:172.26.241.255  Mask:255.255.254.0
     --> UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1 The passive or standby load balancer
will not have active VIPs
    shown in the
    ifconfig -a output (no eth0:0 and eth1:0).
```

Switching Standby and Active Policy Directors

Step 1 Log in to the active Policy Director (lb) VM. See Determining the Active Policy Director, on page 9 for details to determine which Policy Director is active.

Step 2 Restart the Heartbeat service using the following command:

```
monit restart corosync
```

This command will force the failover of the VIP from the active Policy Director to the standby Policy Director.

Step 3 To confirm the switchover, SSH to the other Policy Director VM and run the following command to determine if the VIP is now associated with this VM:

```
ifconfig -a
```

If you see the eth0:0 or eth1:0 interfaces in the list and marked as "UP" then that is the active Policy Director.

Backing Up and Restoring

As a part of routine operations, it is important to make backups so that if there are any failures, the system can be restored. Do not store backups on system nodes.

For detailed information about backup and restore procedures, see the *CPS Backup and Restore Guide*.

Adding or Replacing Hardware

Hardware replacement is usually performed by the hardware vendor with whom your company holds a support contract.

Hardware support is not provided by Cisco. The contact persons and scheduling for replacing hardware is made by your company.

Before replacing hardware, always make a backup. See the *CPS Backup and Restore Guide*.

Unless you have a readily available backup solution, use VMware Data Recovery. This solution, provided by VMware under a separate license, is easily integrated into your CPS environment.

The templates you download from the Cisco repository are partially pre-configured but require further configuration. Your Cisco technical representative can provide you with detailed instructions.

Note You can download the VMware software and documentation from the following location:

http://www.vmware.com/

Export and Import Service Configurations

You can export and import service configurations for the migration and replication of data. You can use the export/import functions to back up both configuration and environmental data or system-specific information from the configuration for lab-to-production migration.

You can import the binary in the following two ways:

- Import the binary produced by export - All configuration exported will be removed (If environment is included, only environment will be removed. If environment is excluded, environment will not be removed). The file passed is created from the export API.
- Additive Import - Import the package created manually by adding configuration. The new configurations get added into the server without impacting the existing configurations. The import is allowed only if the CPS running version is greater than or equal to the imported package version specified in the configuration.

Step 1 In a browser, navigate to the export/import page, available at the following URLs:
HA/GR: https://<<*lbvip01*>:7443/doc/import.html

All-In-One (AIO): http://<*ip*>:7070/doc/import.html

Step 2 Enter the API credentials.

Step 3 Select the file to be imported/exported.
The following table describes the export/import options:

Table 2: Export and Import Options

Option	Description
Export	
All data	Exports service configuration with environment data, which acts as a complete backup of both service configurations and environmental data.
Exclude environment	Exports without environment data, which allows exporting configuration from a lab and into another environment without destroying the new system's environment-specific data.
Only environment	Exports only environment data, which provides a way to back up the system-specific environmental information.
Export URL	Found in Policy Builder or viewed directly in Subversion.
Export File Prefix	Provide a name (prefix) for the export file. **Note:** The exported filename automatically includes the date and time when the export was performed, for example: *prefix*_2016-01-12_11-03-56_3882276668.cps **Note:** The file extension **.cps** is used so that the file is not opened or modified by mistake by another application. The file should be used for export/import purposes only.

Option	Description
Import	
Import URL	URL is updated/created. We recommend importing to a new URL and use Policy Builder to verify/publish.
Commit Message	Message recorded with the import. Provide details that are useful to record.

After you select the file, the file's information is displayed.

Step 4 Select **Import** or **Export**.
CPS displays response messages that indicate the status of the export/import.