GR Installation - VMware

GR Installation Process

In this chapter, Active/Standby Geographic Redundancy model has been used to describe the database and configuration changes required to modify the current installed HA system into Geo-HA.

If you want to deploy historical Active/Active model, just deploy additional flipped pair of this active/standby model.

Overview

An overview of active/standby model has been provided in this section.

1. Active/Standby solution has only one CPS cluster at each site, CA-PRI (referenced as ClusterA Primary henceforth) at S1 site (referenced as Geo-site-1/site-1 henceforth) and CA-SEC (referenced as ClusterA secondary henceforth) at S2 site (referenced as Geo-site-2/site-2 henceforth).

Figure 1: Geographical Sites

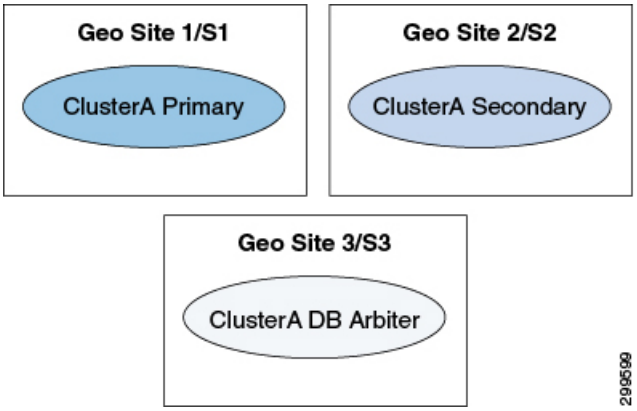

In the above figure, you have primary cluster (Geo Site 1/S1), secondary cluster (Geo Site 2/S2) and arbiter (Geo Site 3/S3).

- Geo site 1/S1 could be any site (for example, Mumbai)
- Geo site 2/S2 could be any site (for example, Chennai)
- Geo site 3/S3 could be any site (for example, Kolkata)

2. For Site1 PCEF, there are two CPS clusters. One is primary, CA-PRI on S1 and other is secondary, CA-SEC on S2. They are geographically redundant.
3. Upon failure of primary CPS Cluster, secondary CPS cluster would seamlessly serve the subscriber's sessions. For that, session replication is enabled between primary and secondary clusters and for session replication high bandwidth is expected. For more information, contact your Cisco technical representative.
4. We recommend to use Replication interface for Database replication between Geo sites (that is, S1 and S2) to segregate Network traffic with Database replication traffic. For example, setting up separate VLAN's for segregating Network and Database traffic.
5. The secondary CPS cluster is not configured as passive.
6. We recommend to place the arbiter on site-3.
7. We recommend the SPR and balance databases to be on SSD and session database to be on tmpfs for optimized performance.

Prerequisites

- Base install (CPS-HA) has been completed on both sites and verified basic validation on both sites.
- Call model has been validated on both HA sites as per your TPS/traffic.
- CPS VMs should have Replication IP address.
- Familiarity with *CPS Installation Guide for VMware*.

  Familiarity with *CPS Release Notes*.

- For third site, Arbiter must be deployed and running the same build (The ISO used to prepare the Geo-Redundant Setup).
- The database configuration is planned.

Reference for CPS VM and Host Name Nomenclature

Note This section is for reference only. You need to follow the nomenclature based on your network requirements. As a prerequisite, HA must be already deployed.

For better usability of the system, install the HA system according to the following nomenclature:

1. In order to know the exact geo site details, we recommend to have the following entries in VMSpecification sheet of `CPS_deployment_config_template.xlsm` or `VMSpecification.csv`.

   Host Name Prefix field value as Sx:

   Table 1: Host Name Prefix Example

Cluster Name	Recommended Value
CA-PRI	S1-
CA-SEC	S2-

2. In order to know the exact cluster name and role (primary/secondary) details, we recommend to have the following entries in Hosts sheet of CPS_deployment_config_template.xlsm or Hosts.csv:
   - Guest Name field value as:

     CA-PRI-XXX for primary cluster (like CA-PRI-lb01, CA-PRI-qns01, and so on.) and CA-SEC-XXX for secondary cluster (like CA-SEC-qsn01, CA-SEC-lb01, and so on.)

3. We recommend to distribute session manager VMs equally between primary and secondary clusters, example:

   sessionmgr01, sessionmgr02, sessionmgr03, sessionmgr04 on CA-PRI and

   sessionmgr01, sessionmgr02, sessionmgr03, sessionmgr04 on CA-SEC

4. The following convention must be used while creating cross site replica-set for the session database:

   You must create the session database replica-set members on same VM and same port on both sites. For example, among four replica-set members (except arbiter), if `sessionmgr01:27717` and `sessionmgr02:27717` are two members of replica-set from SITE1 then choose `sessionmgr01:27717` and `sessionmgr02:27717` of SITE2 as other two replica-set members as shown in following example:

   ```
   [SESSION-SET]
            SETNAME=set01
            OPLOG_SIZE=5120
            ARBITER1=SITE-ARB-sessionmgr05:27717
            ARBITER_DATA_PATH=/var/data/sessions.1/set1
            PRIMARY-MEMBERS
            MEMBER1=SITE1-sessionmgr01:27717
   ```

```
MEMBER2=SITE1-sessionmgr02:27717
SECONDARY-MEMBERS
MEMBER1=SITE2-sessionmgr01:27717
MEMBER2=SITE2-sessionmgr02:27717
DATA_PATH=/var/data/sessions.1/set1
[SESSION-SET-END]
```

5. pcrfclient01 and pcrfclient02 of each site require Management/Public IP
6. Site1 HA Blade naming conventions of VMs looks like (This information is for reference only):

Table 2: Naming Convention

Blade	Virtual Machines
CC Blade 1	S1-CA-PRI-cm S1-CA-PRI-lb01 S1-CA-PRI-pcrfclient01
CC Blade 2	S1-CA-PRI-lb02 S1-CA-PRI-pcrfclient02
CPS Blade 1	S1-CA-PRI-qns01 S1-CA-PRI-sessionmgr01
CPS Blade 2	S1-CA-PRI-qns02 S1-CA-PRI-sessionmgr02
CPS Blade 3	S1-CA-PRI-qns03 S1-CA-PRI-sessionmgr03
CPS Blade 4	S1-CA-PRI-qns04 S1-CA-PRI-sessionmgr04

7. Site2 HA configuration looks like (This information is for reference only):

Table 3: Naming Convention

Blade	Virtual Machines
CC Blade 1	S1-CA-SEC-cm S1-CA-SEC-lb01 S1-CA-SEC-pcrfclient01
CC Blade 2	S1-CA-SEC-lb02 S1-CA-SEC-pcrfclient02
CPS Blade 1	S1-CA-SEC-qns01 S1-CA-SEC-sessionmgr01

Blade	Virtual Machines
CPS Blade 2	S1-CA-SEC-qns02 S1-CA-SEC-sessionmgr02
CPS Blade 3	S1-CA-SEC-qns03 S1-CA-SEC-sessionmgr03
CPS Blade 4	S1-CA-SEC-qns04 S1-CA-SEC-sessionmgr04

Arbiter Installation

Note If you want to add the MongoDB authentication on Arbiter, refer to *General Configuration* section in *CPS Installation Guide for VMware*. You need to mention password for all the sites separately using CSV file and that should be same for all the sites.

On Third Site

Important Currently, SNMP and statistics are not supported on third site arbiter.

Do not install Arbiter if third site is not there or Arbiter is already installed on primary site.

Additionally, if third site blades are accessible from one of the GR sites, you can spawn the Arbiter VM from one of the sites, say, Site1, and installer will sit on third site blades. In that case also, this section is not applicable. Just have appropriate configurations done (On Primary Site, on page 7) so that destination VM is on a third site's blade.

The automatic GR site failover happens only when arbiters are placed on third site thus we recommend the MongoDB arbiter to be on third site that is, S3.

Note Arbiter VM name should be sessionmgr*xx*.

Site3 HA configuration looks like (This information is for reference only):

Table 4: Naming Convention

Blade	Virtual Machines	vCPU	Memory (GB)
CPS Blade 1	S3-ARB-cm S3-CA-ARB-sessionmgr01	1 4	8 8

For more information about deploying VMs, refer to *CPS Installation Guide for VMware*.

Step 1 Configure system parameters for deployment for new Arbiter VM. We need the following CSV files to deploy and configure arbiter VM:

They are:

- VLANs.csv
- Configuration.csv
- VMSpecification.csv
- AdditionalHosts.csv
- Hosts.csv

1. VLAN.csv: Here configurations need to be as per targeted deployment/availability. An example configuration is shown:

Table 5: VLAN.csv

VLAN Name	Network Target Name	Netmask	Gateway	VIP Alias
Internal	VM Network	x.x.x.x	x.x.x.x	-
Management	VM Network-1	x.x.x.x	x.x.x.x	-

2. Configuration.csv: Here configurations need to be as per targeted deployment/availability.
3. VMSpecification.csv: Here configurations need to be as per targeted deployment/availability.
4. AdditionalHosts.csv: Here configurations need to be as per targeted deployment/availability. An example configuration is shown where we need to provide site1 and site2 session managers details:

Table 6: AdditionalHosts.csv

Host	Alias	IP Address
ntp-primary	ntp	x.x.x.x
ntp-secondary	btp	x.x.x.x
CA-PRI-sessionmgr01	-	x.x.x.x
CA-PRI-sessionmgr02	-	x.x.x.x
CA-SEC-sessionmgr01	-	x.x.x.x
CA-SEC-sessionmgr02	-	x.x.x.x

5. Hosts.csv: Take the template file `/var/qps/install/current/scripts/deployer/templates` from Cluster Manager VM and make changes.

   An example configuration is shown:

Figure 2: Hosts.csv

Hypervisor Name	Guest Name	Role	Alias	Datastore	Networks -->	Internal	Management Gx
x.x.x.x	CA-ARB-sessionmgr01	smarb	sessionmgr01	x.x.x.x		x.x.x.x	

299596

Note
- Datastore name should be as per deployment/availability.
- Arbiter VM alias should be sessionmgr*XX* only. In the above example, it is sessionmgr01.

Step 2 Convert the Excel spreadsheet into a CSV file and upload the file to the Cluster Manager VM in `/var/qps/config/deploy/csv` directory.

a) Execute the following commands to import CSV files and conversion to JSON data:

```
/var/qps/install/current/scripts/import/import_deploy.sh
```

b) Execute the following command to validate the imported data:

```
cd /var/qps/install/current/scripts/deployer/support/
python jvalidate.py
```

The above script validates the parameters in the Excel/csv file against the ESX servers to make sure ESX server can support the configuration and deploy VMs.

Step 3 For each host that is defined in the Hosts sheet of the deployment spreadsheet, perform the manual deployment (Refer to the *Manual Deployment* section in the *CPS Installation Guide for VMware*).

Example:

An example is shown below:

```
/var/qps/install/current/scripts/deployer
./deploy.sh sessionmgr01
```

Step 4 If you want to enable mongo authentication on Arbiter, create the file `/etc/facter/facts.d/mongo_auth.txt` using the following data. The password should be same with all other sites.

```
db_authentication_enabled=TRUE
db_authentication_admin_passwd=XXXXX
db_authentication_readonly_passwd=YYYYY
```

where, *XXXXX* and *YYYYY* are encrypted passwords.

On Primary Site

Note Optional: Do not perform the following steps if Arbiter is installed on third site.

If third site is not available then deploy arbiter VM on Primary Cluster that is, CA-PRI.

Note Arbiter VM name should be sessionmgr*XX* only. *XX* should be replaced with a digit higher than the last used digit of the session manager. For example, if there are a total of six sessionmgrs (sessionmgr01-sessionmgr06) then, the Arbiter session manager must be sessionmgr07.

To deploy arbiter on primary site, perform the following steps:

Step 1 Configure system parameters for deployment.

Add the following arbiter entry in **Hosts** sheet of deployment template sheet or **Hosts.csv** file. An example entry is shown below m:

Figure 3: Arbiter Entries

Hypervisor Name	Guest Name	Role	Alias	Datastore	Networks -->	Internal	Management Gx	Replication
x.x.x.x	CA-PRI-sessionmgr07	smarb	sessionmgr07	datastore1		x.x.x.x		x.x.x.x

299597

Step 2 Import modified CSV files using the following command:

```
/var/qps/install/current/scripts/import/import_deploy.sh
```

Step 3 Execute the following command to validate the imported data:

```
cd /var/qps/install/current/scripts/deployer/support/
```

```
python jvalidate.py
```

Note The above script validates the parameters in the Excel/csv file against the ESX servers to make sure ESX server can support the configuration and deploy VMs.

Step 4 For each host that is defined in the Hosts sheet of the excel document perform the manual deployment (Refer to the *Manual Deployment* section in the *CPS Installation Guide for VMware*).

An example is shown below:

```
/var/qps/install/current/scripts/deployer
```

```
./deploy.sh sessionmgr07
```

Standalone Arbiter Deployment On VMware

Note If you want to add the MongoDB authentication on Arbiter, refer to *General Configuration* section in *CPS Installation Guide for VMware*. You need to mention password for all the sites separately using CSV file and that must be same for all the sites.

To install Arbiter on VM, perform the following steps:

Step 1 Convert the Cluster Manager VM to an Arbiter (VMware).

Note Here you are converting the Cluster Manager deployed at Site3 to an Arbiter. For more information on how to deploy Cluster Manager VM, refer to *Deploy the Cluster Manager VM* section in the *CPS Installation Guide for VMware*.

Step 2 Run `install.sh` from ISO directory.

```
cd /mnt/iso
./install.sh
Please enter install type [mobile|mog|pats|arbiter|andsf|escef]: arbiter ----> Select arbiter for
this option
Would you like to initialize the environment... [y|n]: y     ----> Enter y to continue
```

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.

Note Currently, eSCEF is an EFT product and is for Lab Use Only. This means it is not supported by Cisco TAC and cannot be used in a production network. The features in the EFT are subject to change at the sole discretion of Cisco.

Step 3 When prompted for `Please pick an option for this setup:`.

Select **1** for new Arbiter deployment.

Step 4 To enable the firewall, it is required to add the following configuration in `/etc/facter/facts.d/qps_firewall.txt` file:

```
firewall_disabled=0
internal_address=XX.XX.XX.XX ---> update XX.XX.XX.XX to your internal IP address
internal_device=0            ---> update 0 to your device ID
internal_guest_nic=eth0      ---> update eth0 to other port if it is not using default NIC for
internal address
```

Step 5 When `install.sh` finishes its run, execute the `reinit.sh` script to apply the appropriate configurations to the system:

```
/var/qps/install/current/scripts/upgrade/reinit.sh
```

Step 6 If you want to enable mongo authentication on Arbiter, create the file `/etc/facter/facts.d/mongo_auth.txt` using the following data. The password should be same with all other sites.

```
db_authentication_enabled=TRUE
db_authentication_admin_passwd=XXXXX
db_authentication_readonly_passwd=YYYYY
```

where, *XXXXX* and *YYYYY* are encrypted passwords. For encrypted passwords, you need to SSH to a Cluster Manager and execute the following command:

```
/var/qps/bin/support/mongo/encrypt_passwd.sh <Password>
```

Step 7 Edit `/etc/hosts/` and add the information related to all the replica members entries as per your requirement (replica members in `mongoConfig.cfg` file).

Example:

```
cat /etc/hosts

192.168.1.1 arbiter-site3
     192.168.1.2 sessionmgr01-site1
     192.168.1.3 sessionmgr02-site1
     192.168.1.4 sessionmgr01-site2
```

```
192.168.1.5 sessionmgr02-site2
```

Step 8 After performing the upgrade/new installation, unmount the ISO image. This prevents any "device is busy" errors when a subsequent upgrade/new installation is performed.

```
cd /root
umount /mnt/iso
```

Step 9 (Optional) After unmounting the ISO, delete the ISO image to free the system space.

`rm` *xxxx.iso*

where, *xxxx.iso* is the name of the ISO image used.

Step 10 (Optional) Change the host name of the Arbiter.

a) Run `hostname` *xxx*, where *xxx* is the new host name for the Arbiter.
b) Edit `/etc/hostname` to add the new host name for the Arbiter.

Multiple Arbiter Installation - VMware

Step 1 Update the arbiter member information in `/etc/broadhop/mongoConfig.cfg` file.

Example:

```
[SESSION-SET1]
SETNAME=set01a
OPLOG_SIZE=5120
HEARTBEAT_TIMEOUT=3
ARBITER1=arbitervip:27717
ARBITER2=arbiterscale02:27717
ARBITER3=arbiterscale03:27717
ARBITER_DATA_PATH=/var/data/sessions.1/WSP1/set01a
MEMBER1=WSP1SM01:27717
MEMBER2=WSP2SM01:27717
MEMBER3=SFP1SM01:27717
MEMBER4=SFP2SM01:27717
DATA_PATH=/var/data/sessions.1/WSP1/set01a
[SESSION-SET2-END]
```

Note With the above configuration, CPS can handle sessions even in case of three cluster failure. But if one site and arbiter goes down, then primary selection may fail. To avoid this, Cisco recommends to increase the number of arbiters to three or more.

Step 2 Rebuild `etc` directory on cluster with the updated `mongoConfig.cfg` file.

```
/var/qps/install/current/scripts/build/build_etc.sh
```

Step 3 Copy `mongoConfig.cfg` file to all the nodes using `copytoall.sh` from Cluster Manager.

```
copytoall.sh /etc/broadhop/mongoConfig.cfg /etc/broadhop/mongoConfig.cfg
```

Configure Remote/Peer Site VM

Session Manager VM Information on Local Site

Note The following steps need to be performed on other sites as well.

Note In this section, to configure remote/peer site VM in local site, sessionmgr has been taken as an example. You can use this section to add peer policy server (qns) and peer policy directors (lbs) entries in `AdditionalHosts` file.

Step 1 Add the following entry in `AdditionalHosts` sheet of CPS deployment template or `AdditionalHosts.csv` on CA-PRI-cm: Objective of this section is for primary cluster to add other cluster (that is, secondary cluster) session manager's details.

a) Add sessionmgr VM information of secondary cluster (that is, Name, Replication Interface IP addresses).
b) In Alias column add `psessionmgr`*xx* (that is, peer sessionmgr).
c) Add arbiter VM entry, also in Alias column add the same host name.

- If it is on third site, then add IP address of arbiter VM which is reachable from all other sessionmgrs from both sites.
- Else add internal interface IP address of arbiter VM.

Example:

Example of `/var/qps/config/deploy/csv/AdditionalHosts.csv` (on CA-PRI-cm):

```
Host,Alias,IP Address
-----
CA-SEC-sessionmgr01,psessionmgr01,xx.xx.xx.xx
CA-SEC-sessionmgr02,psessionmgr02, xx.xx.xx.xx
CA-ARB-sessionmgr01,CA-ARB-sessionmgr01,xx.xx.xx.xx
-----
```

Step 2 Import modified CSV files by executing the following command:

```
/var/qps/install/current/scripts/import/import_deploy.sh
```

Step 3 Execute the following command to validate the imported data:

```
cd /var/qps/install/current/scripts/deployer/support/

python jvalidate.py
```

Note The above script validates the parameters in the Excel/csv file against the ESX servers to make sure ESX server can support the configuration and deploy VMs.

Step 4 Execute the following command in Cluster Manager to copy updated `/etc/hosts` file to all deployed VMs:

```
SSHUSER_PREFERROOT=true copytoall.sh /etc/hosts /etc/hosts
```

Step 5 Validate setup using `diagnostics.sh` script.

Policy Director (lb) VM Information on Local Site

Before you begin

Redis must be enabled as IPC. For more information on how to enable REDIS, refer to *CPS Installation Guide for VMware*.

Step 1 Add the following entry in `AdditionalHosts` sheet of CPS deployment template or `AdditionalHosts.csv` on CA-Site1-cm: Objective of this section is for local site to add other site (that is, remote clusters) policy director (lb) VM details.

a) Add policy director (lb) VM information of secondary cluster (that is, Name, Policy Director (LB) External Interface Name).

b) In Alias column `add plbxx` (that is, peer policy director (lb). For example, plb01, plb02 and so on).

Add IP address of remote policy director (lb) VM which is reachable from all policy director (lb) VMs of primary cluster.

Example:

Example of `/var/qps/config/deploy/csv/AdditionalHosts.csv` (on CA- Site1-cm):

```
Host,Alias,IP Address
-----
CA- Site2-lb01,plb01,xx.xx.xx.xx
CA- Site2-lb02,plb02, xx.xx.xx.xx
-----
```

Step 2 Add the number of remote redis instances in `Configuration.csv` with key as *remote_redis_server_count* and value as the number of redis instances running on remote site:

Example:

If the remote site contains three redis instances per policy director (lb), add the following:

```
remote_redis_server_count,3
```

For more information in *remote_redis_server_count*, refer to *CPS Installation Guide for VMware*.

Step 3 Import modified CSV files by executing the following command:

```
/var/qps/install/current/scripts/import/import_deploy.sh
```

Step 4 Execute the following commands in Cluster Manager to copy updated `/etc/hosts` file to all deployed VMs:

```
SSHUSER_PREFERROOT=true copytoall.sh /etc/hosts /etc/hosts

copytoall.sh /etc/broadhop/redisTopology.ini /etc/broadhop/redisTopology.ini
```

Step 5 Verify that the `/etc/broadhop/redisTopology.ini` contains the remote policy director (lb) redis instances entry as follows:

```
cat /etc/broadhop/redisTopology.ini
```

```
policy.redis.qserver.1=lb01:6379
policy.redis.qserver.2=lb01:6380
policy.redis.qserver.3=lb01:6381
policy.redis.qserver.4=lb02:6379
policy.redis.qserver.5=lb02:6380
policy.redis.qserver.6=lb02:6381
remote.policy.redis.qserver.1=plb01:6379
remote.policy.redis.qserver.2=plb01:6380
remote.policy.redis.qserver.3=plb01:6381
remote.policy.redis.qserver.4=plb02:6379
remote.policy.redis.qserver.5=plb02:6380
remote.policy.redis.qserver.6=plb02:6381
```

If the number of redis instances/lb instances are to be increased/decreased on the remote cluster(s), the same should first be updated in all other clusters in the CSV files as mentioned in Step 1, on page 12 and Step 2, on page 12.

Repeat the steps from Step 3, on page 12 to Step 5, on page 12 after changing the CSV files so as to update the `redisTopology` file on all the VMs.

Database Configuration

Note While configuring mongo ports in a GR environment, there should be a difference of 100 ports between two respective sites. For example, consider there are two sites: Site1 and Site2. For Site1, if the port number used is 27717, then you can configure 27817 as the port number for Site2. This is helpful to identify a mongo member's site. By looking at first three digits, one can decide where the mongo member belongs to. However, this is just a guideline. You should avoid having mongo ports of two different sites close to each other (for example, 27717 on Site-1 and 27718 on Site2).

Reason: The reason is that the `build_set.sh` script fails when you create shards on the site (for example, Site1). This is because the script calculates the highest port number in the `mongoConfig` on the site where you are creating shards. This creates clash between the replica-sets on both sites. Since the port number which it allocates might overlap with the port number of `mongoConfig` on other site (for example, Site2). This is the reason why there should be some gap in the port numbers allocated between both the sites.

Step 1 Log in to Cluster Manager as a root user.

Step 2 Modify the `/etc/broadhop/gr_cluster.conf` file. For example, if `/etc/broadhop/qns.conf` file has the following entries for Site ID:.

```
-DSiteId=clusterA
```

```
-DRemoteSiteId=clusterB
```

a) Add cluster name that is, clusterA followed by pcrfclient01/pcrfclient02 management interface public IP address of the Primary ClusterA.

For example, `clusterA:`*`a.b.c.d,w.x.y.z`*

where,

a.b.c.d is the pcrfclient01 management interface public IP address of the Primary ClusterA.

a.b.c.d is the pcrfclient02 management interface public IP address of the Primary ClusterA.

b) On next line add remote site name that is, clusterB followed by pcrfclient01/pcrfclient02 Management-interface public IP address of the Secondary ClusterB (these public IP addresses should be pingable from Site1).

For example, `clusterA:e.f.g.h,p.q.r.s`

where,

e.f.g.h is the pcrfclient01 management interface public IP address of the Secondary ClusterB.

p.q.r.s is the pcrfclient02 management interface public IP address of the Secondary ClusterB.

These entries need to match with site name entries given in `qns.conf` file.

File contents look like:

```
cat /etc/broadhop/gr_cluster.conf
#<site name>:<pcrfclient01 IP address>:<pcrfclient02 IP address>
#Primary sites
clusterA:a.b.c.d,w.x.y.z
#Secondary sites
clusterB:e.f.g.h,p.q.r.s
```

Step 3 **Verify MongoConfig**: Do not miss to add `#SITEx_START` and `#SITEx_END` tags to the block of replica set entries in `/etc/broadhop/mongoConfig.cfg` file, where *x* is the site number. To add these tags at proper location, refer to sample configuration file (`geo_mongoconfig_template`) present in `/etc/broadhop` directory. The SiteIDs must be obtained from `/etc/broadhop/qns.conf` file from the field `-DSiteID`.

Example:

For example, if Site1 (clusterA_PRI) and Site2 (clusterA_SBY) are the two sites, then you need to add Site1 and Site2 entries in `mongoconfig.cfg` file as per the sample configuration file (`geo_mongoconfig_template`) present in `/etc/broadhop` directory.

Example:

To monitor the Arbiter VM alarms, the `mongoConfig.cfg` file must identify the local and remote replica-sets separately. The sets should be marked by SiteID of the site. This SiteID must be obtained from `/etc/broadhop/qns.conf` file from the field `-DSiteID`. For marking/separating the two sites replica-sets, identify the start and end of the replica-sets.

Note Make sure case sensitivity for SiteID as well as START and END tags.

For example, one site is having SiteId as Site1 (clusterA_PRI) and other has Site2 (clusterA_SBY). So, the `mongoconfig.cfg` looks like:

```
#Site1_START
#clusterA_PRI_START
.
.
<All replica sets for Site1.>
.
.
#clusterA_PRI_END
#Site1_END

#Site2_START
#clusterA_SBY_START
.
.
<All Replica sets for Site2.>
.
.
```

```
#clusterA_SBY_END
#Site2_END
```

Step 4 To install the database and synchronize the `mongoConfig.cfg` file across the cluster execute the following commands:

```
/var/qps/install/current/scripts/build/build_etc.sh

/var/qps/bin/update/syncconfig.sh
```

Step 5 Set priority using `set_priority.sh` command. The following are example commands:

```
cd /var/qps/bin/support/mongo/; ./set_priority.sh --db session
cd /var/qps/bin/support/mongo/; ./set_priority.sh --db spr
cd /var/qps/bin/support/mongo/; ./set_priority.sh --db admin
cd /var/qps/bin/support/mongo/; ./set_priority.sh --db balance
```

The primary member of individual replica-sets are on respective sites.

Step 6 Verify replica set status and priority is set correctly using the following command:

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

Note If a member is shown in an unknown state, it is likely that the member is not accessible from one of other members, mostly an arbiter. In that case, you must go to that member and check its connectivity with other members.

Also, you can login to mongo on that member and check its actual status.

Step 7 When the Session replication is configured then the Host collection of the Cluster database should have all the “admin replica-set” members, entries with Internal and Replication VLAN IP's.

By default, `db.hosts` file gets populated if you configure `/etc/broadhop/gr_cluster.conf` file. If the entries are not present then use the following commands to add these entries (XX is internal/replication IP address of “admin replica-set” and YY is siteId (defined in `qns.conf`):

```
mongo --host <admin DB primary host> --port <admin DB port> clusters
> db.hosts.update({"ip" : "XX"},{"siteName" : "YY", "ip" : "XX"}, true)
```

Example:

```
mongo CA-PRI-sessionmgr02:27721/clusters
MongoDB shell version: 2.6.3
connecting to: CA-PRI-sessionmgr02:27721/clusters
set05:PRIMARY> db.hosts.find()
{ "_id" : ObjectId("545e0596f4ce7b3cc119027d"), "siteName" : "clusterA_PRI", "ip" :
"192.168.109.127" }
{ "_id" : ObjectId("545e0596f4ce7b3cc119027e"), "siteName" : "clusterA_PRI", "ip" :
"192.168.109.128" }
{ "_id" : ObjectId("545e0596f4ce7b3cc1190281"), "siteName" : "clusterA_SBY", "ip" :
"192.168.109.227" }
{ "_id" : ObjectId("545e0596f4ce7b3cc1190282"), "siteName" : "clusterA_SBY", "ip" :
"192.168.109.228" }
{ "_id" : ObjectId("545e0596f4ce7b3cc119027d"), "siteName" : "clusterA_PRI", "ip" : "11.11.11.127"
}
{ "_id" : ObjectId("545e0596f4ce7b3cc119027e"), "siteName" : "clusterA_PRI", "ip" : "11.11.11.128"
}
{ "_id" : ObjectId("545e0596f4ce7b3cc1190281"), "siteName" : "clusterA_SBY", "ip" : "11.11.11.227"
}
{ "_id" : ObjectId("545e0596f4ce7b3cc1190282"), "siteName" : "clusterA_SBY", "ip" : "11.11.11.228"
}
```

1. (Optional) By default, db.hosts gets populated if there is a difference between IP addresses of sessionmgr* VMs in `/etc/hosts` file on both sites.

Example:

For sessionmgr01 SITE-A, in `/etc/hosts` file, if the entry is: `10.10.10.1 sessionmgr01 sessionmgr01-SITE-A`

and for SITE-B on sessionmgr01, in `/etc/hosts`, if the entry is: `172.20.20.1 sessionmgr01-SITE-A`

As, IP addresses of sessionmgr VMs are different in this case, user needs to run the following scripts on both SITEs.

```
cd /var/qps/bin/support/mongo/; ./set_clusterinfo_in_admindb.sh
```

Step 8 From this Cluster Manager, copy `mongoConfig.cfg` and `gr_cluster.conf` files to peer Cluster Managers (CM).

Balance Backup Database Configuration

CPS provides extra high availability for balance database during failover. During failover or switchover, or when primary is not available due to network reachability, balance writes happen in the backup database. After primary database is available, the records in backup database are reconciled with the primary.

The balance backup database configuration in mongoConfig appears like any other balance database with two members of replica-set on a given site. The replica-set must be created in the same manner in which regular balance database replica-sets are created.

Step 1 In Policy Builder, click **Reference Data** > **Systems** > ***name of your primary system*** > **Plugin Configurations** and select **Balance Configuration** from right side. In **Balance Configuration**, configure the primary database information. For parameter description, refer to *CPS Mobile Configuration Guide*.

An example configuration is shown:

Figure 4: Balance Backup Database Configuration - 1

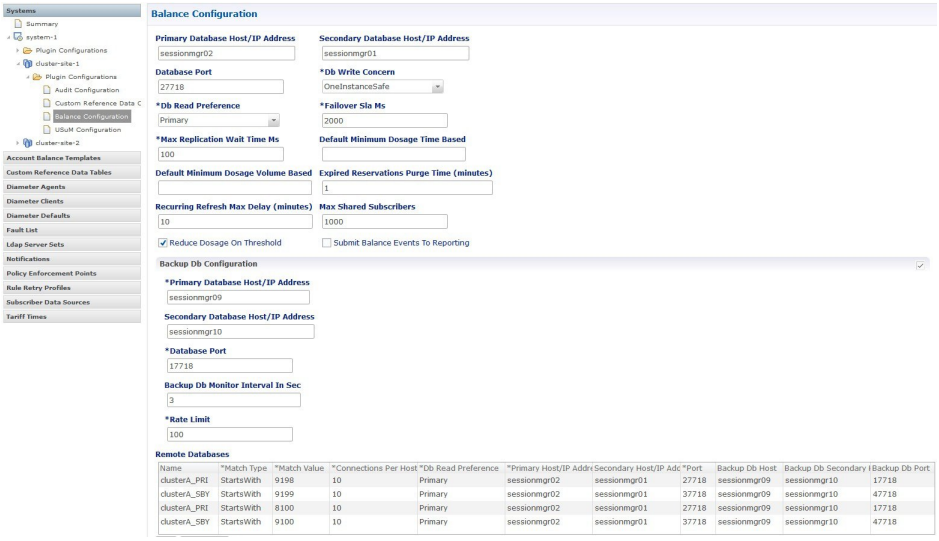

Step 2 In Policy Builder, click **Reference Data** > **Systems** > ***name of your backup system*** > **Plugin Configurations** and select **Balance Configuration** from right side. In **Balance Configuration**, configure the backup database information. For parameter description, refer to *CPS Mobile Configuration Guide*.

An example configuration is shown:

Figure 5: Balance Backup Database Configuration - 2

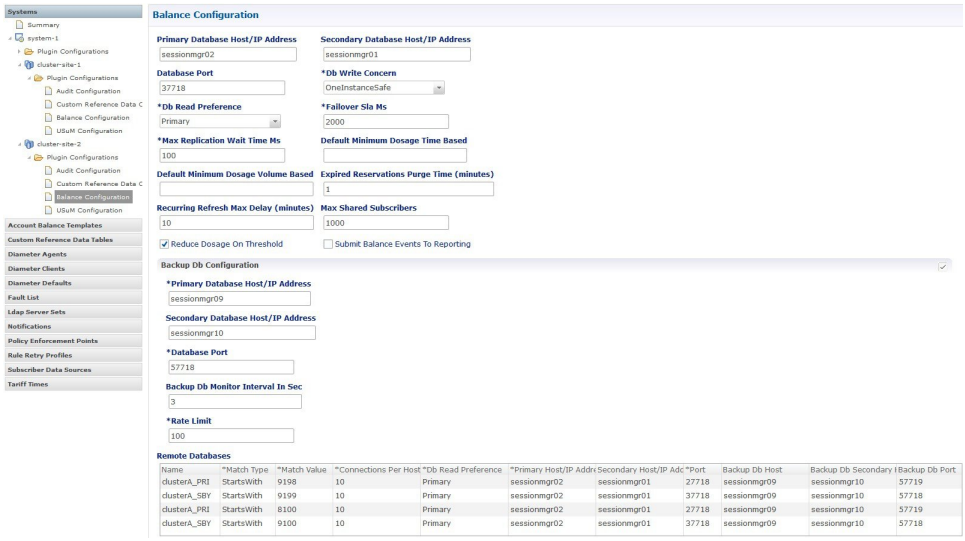

Step 3 The following is an example output for balance backup database:

```
diagnostics.sh --get_re
```

The balance database replica-sets for Site1 and Site2 are displayed in the example output.

```
CPS Diagnostics GR Multi-Node Environment
--------------------------
Checking replica sets...
|-----------------------------------------------------------------------------------------------------------------|
| Mongo:x.x.x                        MONGODB REPLICA-SETS STATUS INFORMATION OF SITE1           Date
: 2016-09-26 10:59:52 |
|-----------------------------------------------------------------------------------------------------------------|
|   SET NAME  -  PORT  : IP ADDRESS   - REPLICA STATE -           HOST NAME         -  HEALTH  -
LAST SYNC   - PRIORITY |
|-----------------------------------------------------------------------------------------------------------------|
| ADMIN:set08
                        |
|  Member-1  -  27721 : 172.20.18.54  -  ARBITER      -  L2-CA-ARB-sessionmgr15    - ON-LINE  -
--------    -    0      |
|  Member-2  -  27721 : 172.20.17.83  -  PRIMARY      -  L2-CA-PRI-sessionmgr09    - ON-LINE  -
--------    -    4      |
|  Member-3  -  27721 : 172.20.17.87  -  SECONDARY    -  L2-CA-PRI-sessionmgr10    - ON-LINE  -
1 sec       -    3      |
|  Member-4  -  27721 : 172.20.19.53  -  SECONDARY    -  L2-CA-SEC-sessionmgr09    - ON-LINE  -
1 sec       -    2      |
|  Member-5  -  27721 : 172.20.19.57  -  SECONDARY    -  L2-CA-SEC-sessionmgr10    - ON-LINE  -
1 sec       -    1      |
|-----------------------------------------------------------------------------------------------------------------|
| BALANCE:set05
                        |
|  Member-1  -  27718 : 172.20.18.54  -  ARBITER      -  L2-CA-ARB-sessionmgr15    - ON-LINE  -
--------    -    0      |
|  Member-2  -  27718 : 172.20.17.40  -  PRIMARY      -  L2-CA-PRI-sessionmgr02    - ON-LINE  -
--------    -    4      |
|  Member-3  -  27718 : 172.20.17.38  -  SECONDARY    -  L2-CA-PRI-sessionmgr01    - ON-LINE  -
0 sec       -    3      |
|  Member-4  -  27718 : 172.20.19.29  -  SECONDARY    -  L2-CA-SEC-sessionmgr02    - ON-LINE  -
0 sec       -    2      |
```

```
|  Member-5  -  27718 : 172.20.19.27  -  SECONDARY    -  L2-CA-SEC-sessionmgr01    - ON-LINE  -
0 sec      -    1     |
|-------------------------------------------------------------------------------------------------------------|
| BALANCE:set09a
                      |
|  Member-1  -  47718 : 172.20.18.54  -  ARBITER      -  L2-CA-ARB-sessionmgr15    - ON-LINE  -
--------    -     0     |
|  Member-2  -  47718 : 172.20.17.87  -  PRIMARY      -  L2-CA-PRI-sessionmgr10    - ON-LINE  -
--------    -     2     |
|  Member-3  -  47718 : 172.20.17.83  -  SECONDARY    -  L2-CA-PRI-sessionmgr09    - ON-LINE  -
0 sec      -    1     |
|-------------------------------------------------------------------------------------------------------------|
| BALANCE:set10a
                      |
|  Member-1  -  17718 : 172.20.18.54  -  ARBITER      -  L2-CA-ARB-sessionmgr15    - ON-LINE  -
--------    -     0     |
|  Member-2  -  17718 : 172.20.17.83  -  PRIMARY      -  L2-CA-PRI-sessionmgr09    - ON-LINE  -
--------    -     2     |
|  Member-3  -  17718 : 172.20.17.87  -  SECONDARY    -  L2-CA-PRI-sessionmgr10    - ON-LINE  -
0 sec      -    1     |
|-------------------------------------------------------------------------------------------------------------|
| REPORTING:set07
                      |
|  Member-1  -  27719 : 172.20.18.54  -  ARBITER      -  L2-CA-ARB-sessionmgr15    - ON-LINE  -
--------    -     0     |
|  Member-2  -  27719 : 172.20.17.44  -  PRIMARY      -  L2-CA-PRI-sessionmgr04    - ON-LINE  -
--------    -     2     |
|  Member-3  -  27719 : 172.20.17.42  -  SECONDARY    -  L2-CA-PRI-sessionmgr03    - ON-LINE  -
0 sec      -    1     |
|-------------------------------------------------------------------------------------------------------------|
| SESSION:set01
                      |
|  Member-1  -  27717 : 172.20.18.54  -  ARBITER      -  L2-CA-ARB-sessionmgr15    - ON-LINE  -
--------    -     0     |
|  Member-2  -  27717 : 172.20.17.38  -  PRIMARY      -  L2-CA-PRI-sessionmgr01    - ON-LINE  -
--------    -     4     |
|  Member-3  -  27717 : 172.20.17.40  -  SECONDARY    -  L2-CA-PRI-sessionmgr02    - ON-LINE  -
0 sec      -    3     |
|  Member-4  -  27717 : 172.20.19.27  -  SECONDARY    -  L2-CA-SEC-sessionmgr01    - ON-LINE  -
0 sec      -    2     |
|  Member-5  -  27717 : 172.20.19.29  -  SECONDARY    -  L2-CA-SEC-sessionmgr02    - ON-LINE  -
0 sec      -    1     |
|-------------------------------------------------------------------------------------------------------------|
|-------------------------------------------------------------------------------------------------------------|
| Mongo:x.x.x                      GODB REPLICA-SETS STATUS INFORMATION OF SITE2             Date
: 2016-09-26 11:00:06 |
|-------------------------------------------------------------------------------------------------------------|
|  SET NAME  -  PORT  : IP ADDRESS       - REPLICA STATE -           HOST NAME      -  HEALTH  -
LAST SYNC   - PRIORITY |
|-------------------------------------------------------------------------------------------------------------|
| BALANCE:set25
                      |
|  Member-1  -  37718 : 172.20.18.54     -  ARBITER      -  L2-CA-ARB-sessionmgr15  - ON-LINE  -
--------    -     0     |
|  Member-2  -  37718 : 172.20.19.29     -  PRIMARY      -  L2-CA-SEC-sessionmgr02  - ON-LINE  -
--------    -     4     |
|  Member-3  -  37718 : 172.20.19.27     -  SECONDARY    -  L2-CA-SEC-sessionmgr01  - ON-LINE  -  0
 sec      -     3     |
|  Member-4  -  37718 : 172.20.17.38     -  SECONDARY    -  L2-CA-PRI-sessionmgr01  - ON-LINE  -  0
 sec      -     2     |
|  Member-5  -  37718 : 172.20.17.40     -  SECONDARY    -  L2-CA-PRI-sessionmgr02  - ON-LINE  -  0
 sec      -     1     |
|-------------------------------------------------------------------------------------------------------------|
| BALANCE:set26
```

```
                    |
|  Member-1  -  57719 : 172.20.18.54     -  ARBITER      -  L2-CA-ARB-sessionmgr15  - ON-LINE  -
--------     -    0     |
|  Member-2  -  57719 : 172.20.19.57     -  PRIMARY      -  L2-CA-SEC-sessionmgr10  - ON-LINE  -
--------     -    2     |
|  Member-3  -  57719 : 172.20.19.53     -  SECONDARY    -  L2-CA-SEC-sessionmgr09  - ON-LINE  -  0
 sec       -    1     |
|-------------------------------------------------------------------------------------------------------------------------|
| BALANCE:set27
                    |
|  Member-1  -  57718 : 172.20.18.54     -  ARBITER      -  L2-CA-ARB-sessionmgr15  - ON-LINE  -
--------     -    0     |
|  Member-2  -  57718 : 172.20.19.53     -  PRIMARY      -  L2-CA-SEC-sessionmgr09  - ON-LINE  -
--------     -    2     |
|  Member-3  -  57718 : 172.20.19.57     -  SECONDARY    -  L2-CA-SEC-sessionmgr10  - ON-LINE  -  0
 sec       -    1     |
|-------------------------------------------------------------------------------------------------------------------------|
```

Session Cache Hot Standby

Important Cisco recommends to configure standby session for GR.

CPS runs a distributed database called MongoDB. MongoDB uses a replication concept for high availability called replica-sets. A replica-set is made up of independent MongoDB instances that run in one of the following three modes:

- Primary: A primary database is the only database available that can accept writes.
- Secondary: A secondary database is a database that is read only and is actively synchronizing to a primary database by replaying the primary's oplog (operations log) on the local node.
- Recovering: A secondary database that is currently synchronizing to the primary and has not caught up to the primary.

Session data is highly concurrent, the application always reads and writes from the primary database. The secondary database(s) provide HA for the primary in the event of VM shutdown or process shutdown. Hot standby session cache replica set is configured to take over the load while primary database is failing over to secondary session cache database. In this fail-over process, it minimizes the call failures and provides high system availability.

Prerequisites

- Hot standby replica-set must be created on different blades (for maximum protection).
- Admin database and session cache databases must be separate replica-sets.
- Hot standby replica-set should be added to shard configuration as backup database true.

Configuration

Step 1 The hotstandby database must be configured just like any other session cache database in mongo config and a replica-set needs to be created.

The following is an example backup database configuraton in mongoDB:

```
[SESSION-SET1] SETNAME=set01
ARBITER1=pcrfclient01-prim-site-1:37718
ARBITER_DATA_PATH=/data/sessions.3
MEMBER1=sessionmgr01-site1:27718
MEMBER2=sessionmgr02-site1:27718
DATA_PATH=/data/sessions.3
[SESSION-SET1-END]
```

Note Hotstandby replica sets must be on different ports.

This needs to be created for VMware. For more information, refer to *CPS Installation Guide for VMware*.

For OpenStack, user `/api/system/config/replica-sets`. For more information, refer to *CPS Installation Guide for OpenStack*.

Step 2 Verify CPS application is running on both the sites (pcrfclient01 and pcrfclient02) and without any application errors.

Example:

By executing `diagnostics.sh` script you can get the diagnostics of application. The `diagnostics.sh` output should not contain any application errors.

Step 3 Verify whether shard command is available in OSGi console or not. From pcrfclient01, login as root user into the OSGi console and run the help.
You can find the following shard command:

```
telnet qns01 9091
osgi> help
---QNS Commands---
     reload - Reload reference data
     genpassword <db password>
---Sharing Commands---
addshard seed1[,seed2] port db-index [backup]
rebalance
migrate
---Controlling the Console---
more - More prompt for console output
disconnect - Disconnects from telnet session
help <commmand> - Display help for the specified command.
```

Step 4 To configure hot standby session management, execute the following commands:

```
telnet qns01 9091
addshard sessionmgr03,sessionmgr04 27717 1 Site1 backup
addshard sessionmgr03,sessionmgr04 27717 2 Site1 backup
addshard sessionmgr03,sessionmgr04 27717 3 Site1 backup
addshard sessionmgr03,sessionmgr04 27717 4 Site1 backup
rebalance
migrate
disconnect
y
```

Step 5 To verify the configuration:

1. Login to primary administration database using port#<admin DB port> and verify the collection shards in shading database.

```
mongo sessionmgr01:27721
MongoDB shell version: 2.6.3
connecting to: sessionmgr01:27721/test
set05:PRIMARY> use sharding
switched to db sharding
set05:PRIMARY> db.shards.find()
{ "_id" : 1, "seed_1" : "sessionmgr03", "seed_2" : "sessionmgr04", "port" : 27717, "db" :
"session_cache_2", "online" :
true,
"count" : NumberLong(0), "backup_db" : true }
{ "_id" : 2, "seed_1" : "sessionmgr03", "seed_2" : "sessionmgr04", "port" : 27717, "db" :
"session_cache_3", "online" :
true,
"count" : NumberLong(0), "backup_db" : true }
{ "_id" : 3, "seed_1" : "sessionmgr04", "seed_2" : "sessionmgr04", "port" : 27717, "db" :
"session_cache_4", "online" :
true,
"count" : NumberLong(0), "backup_db" : true }
{ "_id" : 4, "seed_1" : "sessionmgr03", "seed_2" : "sessionmgr04", "port" : 27717, "db" :
"session_cache", "online" :
true,
"count" : NumberLong(0), "backup_db" : true }
set05:PRIMARY
```

Failover Detection

There are three possible ways a MongoDB node can fail and trigger a fail over to another node:

1. Replica set step down: This scenario is the cleanest method since it disconnects all client sockets and immediately initiates a new election for a master node.
2. Process abort: This scenario occurs when a node aborts due to an internal failure. Since this is unplanned, the other replica nodes will not request a new election for a master node until a majority of the nodes have detected the master as down.
3. VM power off: This scenario occurs when a node is powered off without a proper shutdown. In this case sockets are usually hung on all clients and the other replica nodes will not request a new election for a master node until a majority of the nodes have detected the master as down.

The Cisco Policy Server detects client failure by:

1. Utilizing the pre-packaged MongoDB Java driver software to detect failure.
2. Detecting server power off via server pings to rapidly detect power off situations.

Limitation

You can configure only one backup database. Thus, in GR Active/Standby configuration, if you configure backup database on the standby site, during local primary to local secondary database fail over on active site, the sessions would be saved on the backup database which is on secondary site. This might increase cross-site traffic temporarily.

Policy Builder Configuration

Step 1 Configure and publish Policy Builder changes from each site. Use **about.sh** command to find out Policy Builder URL.

Cisco recommends to configure and publish Policy Builder data separately from each site. But if the user wants to publish Policy Builder data from single site to all other sites then it is difficult to access Policy Builder data from other sites when the primary site goes down.

To access Policy Builder data from other sites when primary site is down, refer Access Policy Builder from Standby Site when Primary Site is Down, on page 25.

Step 2 Set appropriate Primary Database IP address, Secondary Database IP address and Port numbers for the following plug-ins:

- USuM Configuration
- Balance Configuration
- Custom Reference Data Configuration
- Voucher Configuration
- Audit Configuration

Step 3 Set **Balance Configuration** > **Db Read Preference** as **SecondaryPreferred** for all databases except balance database.

Figure 6: Db Read Preference

An example **Cluster** configuration is given:

Figure 7: Policy Builder Screen

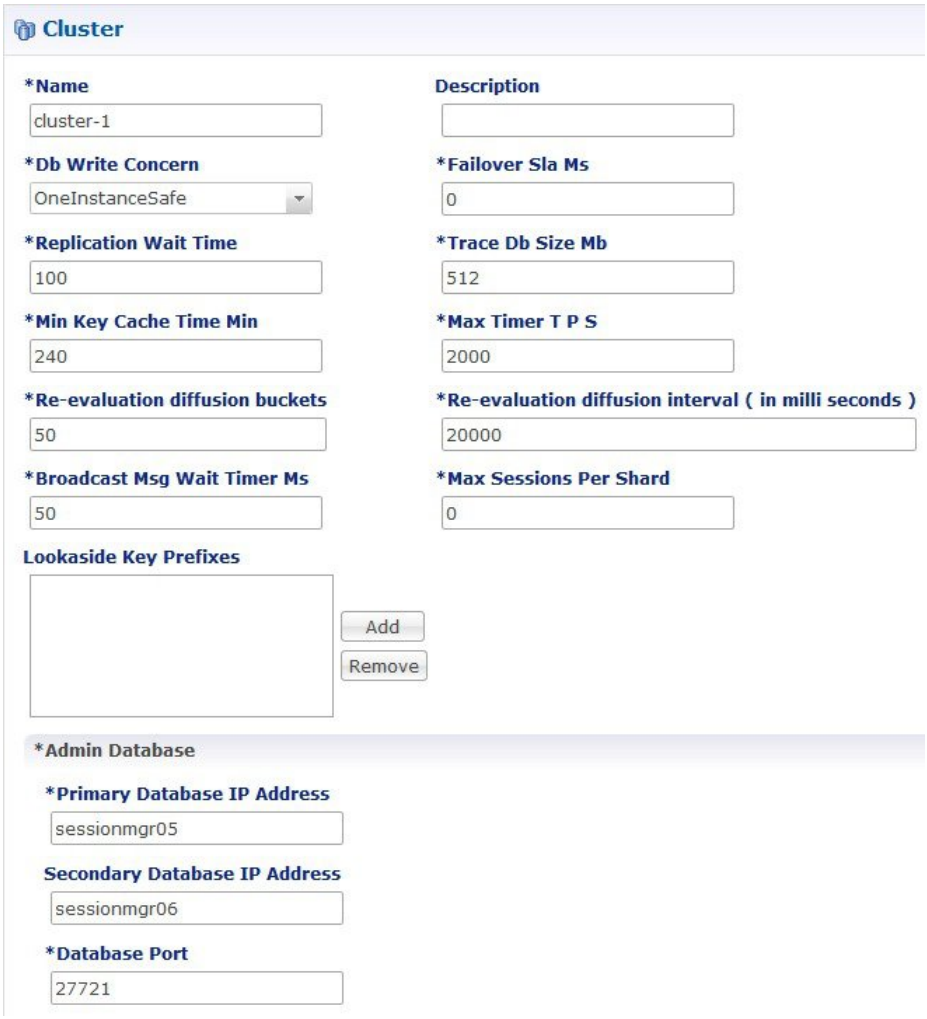

Also update **Lookaside Key Prefixes** and **Admin Database** sections. For more information, refer to *CPS Mobile Configuration Guide*.

Step 4 It is recommended to publish Policy Builder changes from each site. If a user is using primary site to publish Policy Builder changes then publishing into all the following cluster repositories is not recommended:

Table 7: Publishing

Cluster	Publish URL
CA-PRI	http://<*Management interface public IP address of CA-PRI-pcrfclient01*>/repos/run
CA-SEC	http://< *Management interface public IP address of CA-SEC-pcrfclient01*>/repos/run

Step 5 Add all the above repositories. Repository and Publish screen looks like:

Figure 8: Repository Screen

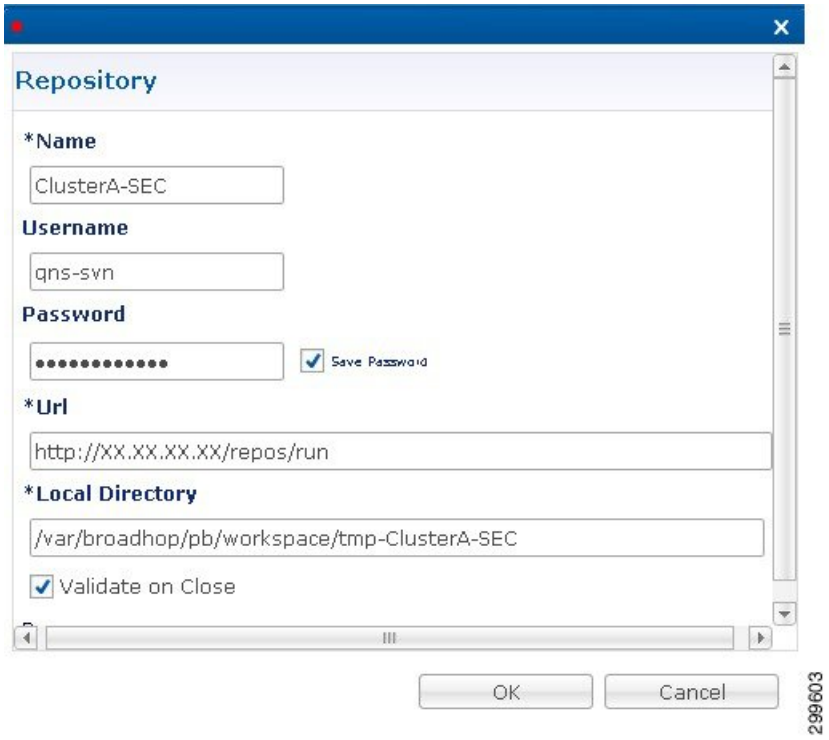

Figure 9: Publish Screen

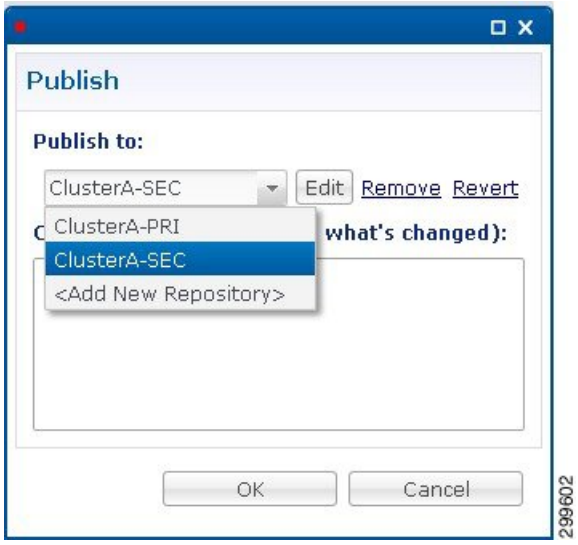

Step 6 Validate both setups using **diagnostics.sh**, after publishing the repository (wait for five minutes) OR follow the Validate VM Deployment section in the *Cisco Policy Suite Installation Guide* for this release.

Access Policy Builder from Standby Site when Primary Site is Down

Step 1 This is recommended only when primary site is down and secondary is only for reading/viewing purpose. It is applicable only where user publishes policy builder data from single site that is, primary site to all other sites.

Step 2 Open Policy Builder from secondary site (use **about.sh** command to find out PB URL).

Step 3 Create new data repository SEC-RUN-RO using URL 'http://< Management interface public IP address of secondary pcrfclient01>/repos/run', screen looks like:

Figure 10: New Data Repository

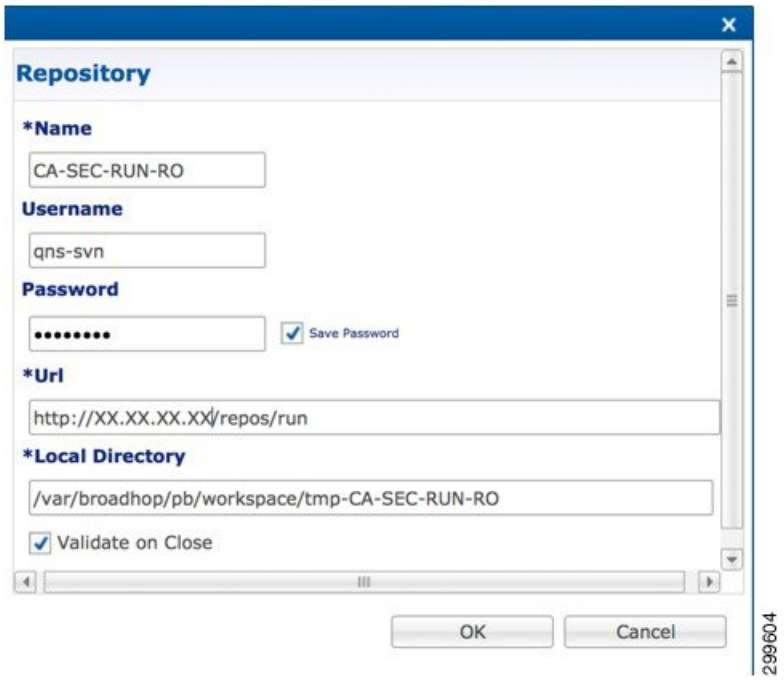

Step 4 Access Policy Builder from secondary site using newly created repository.

qns.conf Configuration Changes for Session Replication

The following changes are required in `qns.conf` file when session replication is required for active/active or active/standby GR deployments.

For active/active GR deployment, Geo HA feature needs to be enabled. For more information, refer to Active/Active Geo HA - Multi-Session Cache Port Support.

Step 1 Add the following GR related parameters in `/etc/broadhop/qns.conf` file of Cluster A Primary cluster manager VM that is, CA-PRI-cm:

```
-DGeoSiteName=clusterA_PRI
-DSiteId=clusterA_PRI
-DRemoteSiteId=clusterA_SBY

-DheartBeatMonitorThreadSleepMS=500
-Dcom.mongodb.updaterConnectTimeoutMS=1000
-Dcom.mongodb.updaterSocketTimeoutMS=1000
-DdbConnectTimeout=1200
-Dmongo.client.thread.maxWaitTime=1200
-DdbSocketTimeout=600
-DclusterFailureDetectionMS=2000
```

Step 2 Add the following GR related parameters in `/etc/broadhop/qns.conf` file of Cluster A Secondary cluster manager VM that is, CA-SEC-cm:

```
-DGeoSiteName=clusterA_SBY
-DSiteId=clusterA_SBY
-DRemoteSiteId=clusterA_PRI

-DheartBeatMonitorThreadSleepMS=500
-Dcom.mongodb.updaterConnectTimeoutMS=1000
-Dcom.mongodb.updaterSocketTimeoutMS=1000
-DdbConnectTimeout=1200
-Dmongo.client.thread.maxWaitTime=1200
-DdbSocketTimeout=600
-DclusterFailureDetectionMS=2000
```

Step 3 For multi-cluster, the following setting should be present only for GR (multi-cluster) CPS deployments:

```
-DclusterFailureDetectionMS=1000
```

Note In an HA or GR deployment with local chassis redundancy, the following setting should be set to true. By default, this is set to false.

```
-Dremote.locking.off
```

Step 4 Create `etc` directory on each cluster using `/var/qps/install/current/scripts/build/build_etc.sh` script.

Step 5 Copy the changes in `qns.conf` to other VMs:

```
copytoall.sh /etc/broadhop/qns.conf /etc/broadhop/qns.conf
```

Step 6 Restart all software components on the target VMs:

```
restartall.sh
```

Step 7 Validate setup using `diagnostics.sh` or follow *Validate VM Deployment* section in the *CPS Installation Guide for VMware* for this release.

Configurations to Handle Database Failover when Switching Traffic to Standby Site Due to Load Balancer Fail/Down

Note To understand traffic switch over, refer to Load Balancer VIP Outage.

Step 1 Add the list of databases that needs to be migrated to primary on other site after traffic switch over in `mon_db_for_lb_failover.conf` file (`/etc/broadhop/mon_db_for_lb_failover.conf`) in Cluster Manager.

Note Contact your Cisco Technical Representative for more details.

Add the following content in the configuration file (`mon_db_for_lb_failover.conf`):

The following is an example and needs to be changed based on your requirement.

```
#this file contains set names that are available in mongoConfig.cfg. Add set names one below other.
#Refer to README in the scripts folder.
SESSION-SET1
SESSION-SET2
BALANCE-SET1
SPR-SET1
```

Step 2 Rebuild etc directory on cluster by executing the following command:

```
/var/qps/install/current/scripts/build/build_etc.sh
```