Managing High Availability in CPS

Porting All-In-One Policy Builder Configuration to HA

This section describes how to port the Policy Builder configuration from an All-In-One (AIO) environment to a High Availability (HA) environment.

Prerequisites

- All the VMs were created using the deployment process.
- This procedure assumes the datastore that will be used to have the virtual disk has sufficient space to add the virtual disk.
- This procedure assumes the datastore has been mounted to the VMware ESX server, regardless of the backend NAS device (SAN or iSCSI, etc).

Porting the Policy Builder Configuration

Policy Builder configuration can be utilized between environments, however, the configuration for Systems and Policy Enforcement Points is environment-specific and should not be moved from one environment to another.

The following instructions will not overwrite the configuration specific to the environment. Please note that as the Systems tab and Policy Enforcement Points data is not moved, the HA system should have these things configured and running properly (as stated above).

The following steps describe the process to port a configuration from an AIO environment to an HA environment.

Step 1 If the HA environment is currently in use, ensure that SVN backups are up to date.

Step 2 Find the URL that Policy Builder is using to load the configuration that you want to use. You can find this by clicking **Edit** on the initial page in Policy Builder.

The URL is listed in the URL field. For the purpose of these instructions, the following URL will be used for exporting the configuration from the AIO environment and importing the configuration to the HA environment:

http://pcrfclient01/repos/configuration

Figure 1: Repository configuration

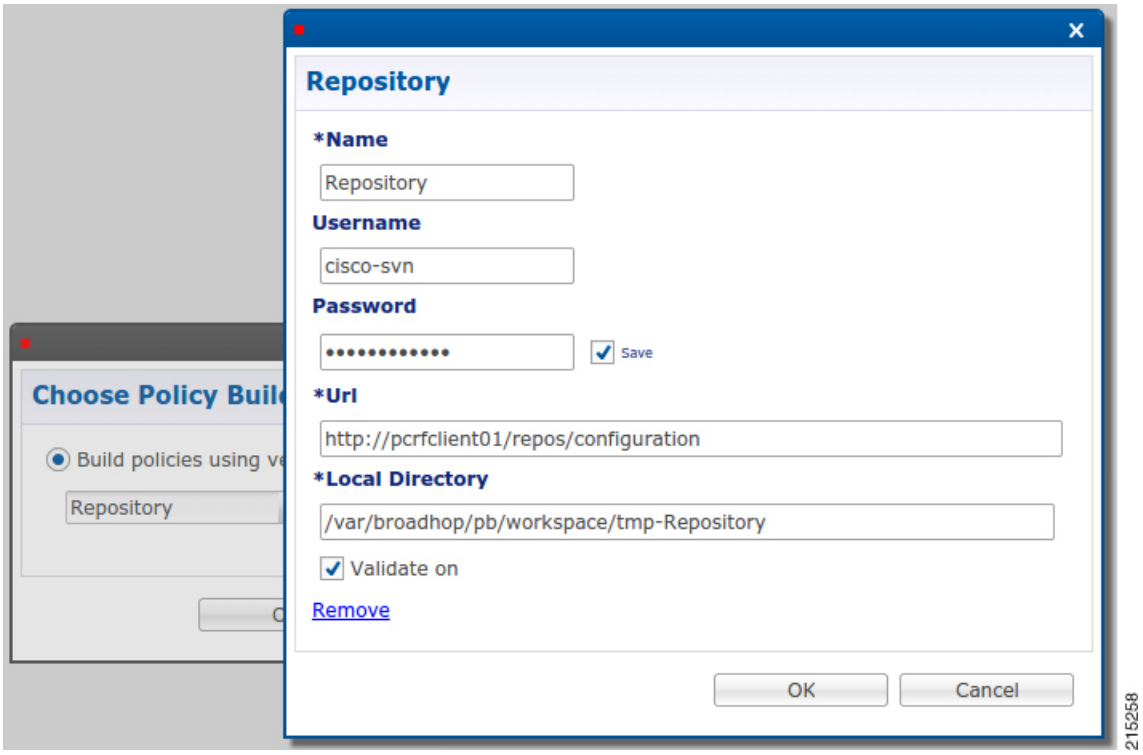

Step 3 On the AIO, export the Policy Builder configuration by entering the following commands:

```
cd /var/tmp
svn export http://pcrfclient01/repos/configuration aio_configuration
```

This creates a directory `/var/tmp/aio_configuration`.

Step 4 Remove the system configuration by entering the following commands:

```
cd aio_configuration
rm -f System* *Configuration* DiameterStack* VoucherSettings* Cluster* Instance*
```

Step 5 Move the `/var/tmp/aio_configuration` directory to `/var/tmp` on your Cluster Manager (using scp, zip and so on).

Step 6 SSH into the pcrfclient01.

The following steps assume you will replace the existing default Policy Builder configuration located at http://pcrfclient01/repos/configuration on your HA environment. If you would like to access your old configuration, copy it to a new location. For example:

```
svn cp http://pcrfclient01/repos/configuration http://pcrfclient01/repos/configuration_date
```

Then set up a new Repository in the HA Policy Builder to access the old configuration.

Step 7 Check out the old configuration (http://pcrfclient01/repos/configuration):

```
svn co http://pcrfclient01/repos/configuration /var/tmp/ha_configuration
```

Step 8 Remove the non-system configuration:

```
svn rm ls | egrep -v '(System|Configuration|DiameterStack|VoucherSettings|Cluster|Instance)'
```

Step 9 Copy in the AIO configuration files:

```
/bin/cp -f /var/tmp/aio_configuration/* .
svn add *
```

Step 10 Commit the configuration:

```
svn ci . -m 'commit configuration moved from AIO'
```

Step 11 If you are already logged into Policy Builder, reload the Policy Builder URL in your browser to access the new configuration.

Step 12 Check for errors in Policy Builder. This often indicates a software mismatch.

Errors are shown with an (x) next to the navigation icons in the left pane of Policy Builder. For example:

Figure 2: Error in Policy Builder

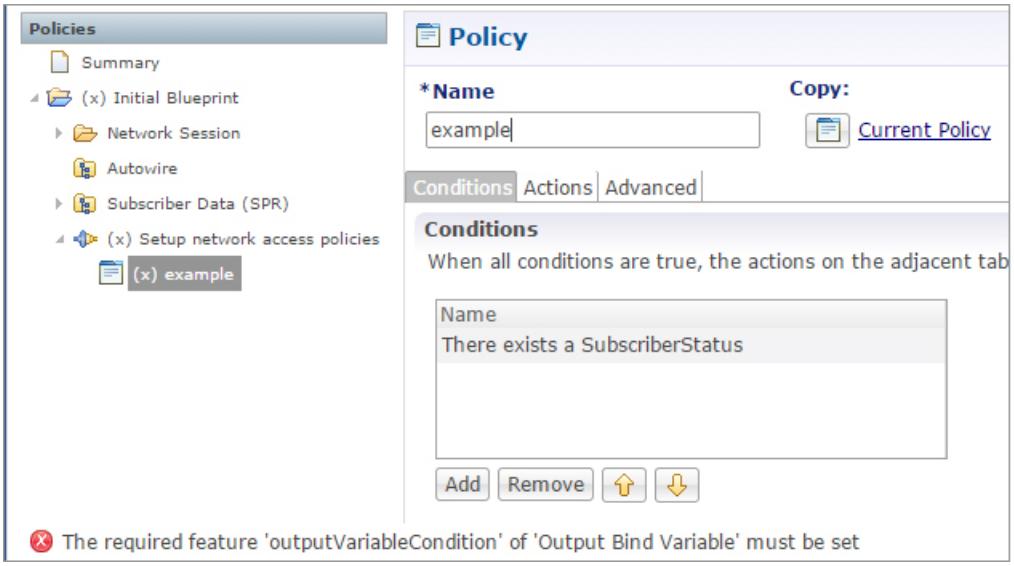

Step 13 Publish the configuration. Refer to the *CPS Mobile Configuration Guide* for detailed steps.

HAProxy

HAProxy is an opensource load balancer used in High Availability (HA) and Geographic Redundancy (GR) CPS deployments. It is used by the CPS Policy Directors (lbs) to forward IP traffic from lb01/lb02 to other CPS nodes. HAProxy runs on the active Policy Director VM.

Documentation for HAProxy is available at http://www.haproxy.org/#docs.

HAProxy Service Operations

Diagnostics

For a general diagnostics check of the HAProxy service, run the following command from any VM in the cluster (except sessionmgr):

diagnostics.sh --ha_proxy

```
QPS Diagnostics Multi-Node Environment
--------------------------
Ping Check for qns01...[PASS]
Ping Check for qns02...[PASS]
Ping Check for qns03...[PASS]
Ping Check for qns04...[PASS]
Ping Check for lb01...[PASS]
Ping Check for lb02...[PASS]
Ping Check for sessionmgr01...[PASS]
Ping Check for sessionmgr02...[PASS]
Ping Check for sessionmgr03...[PASS]
Ping Check for sessionmgr04...[PASS]
Ping Check for pcrfclient01...[PASS]
Ping Check for pcrfclient02...[PASS]
HA Multi-Node Environment
--------------------------
Checking HAProxy status...[PASS]
```

Service Commands

The following commands must be issued from the lb01 or lb02 VM.

To check the status of the HAProxy services, run the following command:

```
monit status haproxy
```

```
[root@host-lb01 ~]# service haproxy status
haproxy (pid 10005) is running...
```

To stop the HAProxy service, run the following command:

```
monit stop haproxy
```

To restart the HAProxy service, run the following command:

```
monit restart haproxy
```

HAProxy Statistics

To view statistics, open a browser and navigate to the following URL:

- **For HAProxy Statistics**: http://<diameterconfig>:5540/haproxy?stats
- **For HAProxy Diameter Statistics**: http://<diameterconfig>:5540/haproxy-diam?stats

Changing HAProxy Log Level

To change HAProxy log level in your CPS deployment, you must make changes to the HAProxy configuration files on the Cluster Manager and then push the changes out to the Policy Director (lb) VMs.

Once deployed, the HAProxy configuration files are stored locally on the Policy Director VMs at `/etc/haproxy/haproxy.cfg.erb` and `/etc/haproxy/haproxy-diameter.erb`.

Note Whenever you upgrade with latest ISO, the log level will be set to default level (err).

Step 1 Log in to the Cluster Manager.

Step 2 Create a backup of the HAProxy configuration file before continuing:

```
cp /var/qps/install/current/puppet/modules/qps/templates/etc/haproxy/haproxy.cfg.erb
/var/qps/install/current/puppet/modules/qps/templates/etc/haproxy/haproxy.cfg.erb-bak-<date>
```

Step 3 Edit the HAProxy files as needed.

By default, the logging level is set as error (err) in `/var/qps/install/current/puppet/modules/qps/templates/etc/haproxy/haproxy-diameter.erb`:

```
log         127.0.0.1       local1 err
```

By default, the logging level in `/var/qps/install/current/puppet/modules/qps/templates/etc/haproxy/haproxy.cfg.erb`:

```
log             127.0.0.1   local3 emerg alert crit err warning
```

The log level can be adjusted to any of the following log levels as needed:

emerg alert crit err warning notice info debug

Step 4 Run `build_all.sh` to rebuild the CPS VM packages.

Step 5 Run `reinit.sh` to trigger all VMs to download the latest software and configuration from the Cluster Manager.

Expanding an HA Deployment

For future installations and network upgrades, this section proposes what hardware and components you should consider as you grow your network. The CPS solution is a robust and scalable software-based solution that can be expanded by adding additional hardware and software components. The following sections explain typical scenarios of when to expand the hardware and software to effect such growth.

Typical Scenarios When Expansion is Necessary

Your network may grow for the following reasons:

- The subscriber base has grown or will grow beyond the initial installation specifications.

  In this case, the number of active or non-active subscribers becomes larger than the initial deployment. This can cause one or more components to reach capacity. New components must be added to accommodate the growth.

- The services or subscriber scenarios have changed, or new services have been introduced, and the transactions per second on a component no longer meet requirements.

  When a new service or scenario occurs, often there is a change in the overall Transactions Per Second (TPS), or in the TPS on a specific component. When this occurs, new components are necessary to handle the new load.

- The operator notices that there are factors outside of the initial design that are causing either the overall system or a specific component to have a high resource load.

  This may cause one or multiple components to reach its capacity for TPS. When this occurs, new components are necessary to handle the new factors.

Hardware Approach to Expanding

Adding a new component may require adding additional hardware. However, the addition of more hardware depends on the physical resources already available, plus what is needed for the new component.

If the number of subscribers exceeds 10 million, then the customer needs to Clone and Repartition sessionmgr Disks. See Manage Disks to Accommodate Increased Subscriber Load.

High Availability Consequences

When adding more hardware, the design must take into consideration the high availability (HA) needs of the system. The HA design for a single-site system is N+1 at the hardware and application level. As a result, adding a new blade incrementally increases the HA capacity of the system.

For example, in a basic installation there are 2 Cisco Policy Server blades handling the traffic. The solution is designed so that if one of the blades fails, the other blade can handle the entire capacity of the system. When adding a third blade for capacity expansion, there are now 2 blades to handle the system load if one of the blades fails. This allows for a more linear scaling approach because each additional blade can be accountable for being able to use its full capacity.

Note When adding new blades to a cluster, the blades in the cluster must be co-located to achieve the proper throughput between other components.

Adding a New Blade

Step 1 Install ESX server to the blade.

Step 2 Open the CPS Deployment Template spreadsheet. This spreadsheet should have been created and maintained during the initial deployment.

Step 3 In the Additional Hosts sheet, add an entry for the new ESX server with IP, Host name and Alias.

Step 4 Save the CSV file and transfer it to the following directory on the Cluster Manager /var/qps/config/deploy/csv

Step 5 Run `/var/qps/install/current/scripts/import/import_deploy.sh` to convert the csv to json.

Component (VM Node) Approach to Expanding

The most common components to be expanded are on the Cisco Policy Servers. As your system begins to scale up, you will need to add more CPS nodes and more SessionMgrs. Expansion for other components can follow the same pattern as described here. The next sections discuss the configurations needed for those specific components to be active in the system.

Adding Additional Component

Step 1 Modify the CPS Deployment Template spreadsheet (this spreadsheet should have been created and maintained during the initial deployment).

Step 2 In the Hosts sheet, add the new VM node with the parameters. See the *CPS Installation Guide for VMware* for details about each column.

Step 3 Save the CSV file and transfer it to the following directory on the Cluster Manager: `/var/qps/config/deploy/csv`.

Step 4 Run `/var/qps/install/current/scripts/import/import_deploy.sh` to convert the csv to json.

Step 5 Deploy the new VM using `/var/qps/install/current/scripts/deployer/deploy.sh` xxx, where xxx is the alias of the new VM to be deployed.

Refer to the *CPS Installation Guide for VMware* for more details about using `deploy.sh`.

Enable SSL

CPS uses encryption on all appropriate communication channels in HA deployments. No additional configuration is required.

Default SSL certificates are provided with CPS but we recommend that you replace these with your own SSL certificates. Refer to Replace SSL Certificates in the *CPS Installation Guide for VMware* for more information.