



Basic Systems Configuration

- [Overview, on page 1](#)
- [Policy Builder Repository Configuration, on page 1](#)
- [System Configuration, on page 13](#)

Overview

The Cisco Policy Suite provides the Policy Builder as an interface for policy management. Policies translate a Service Provider's business rules into actionable, logical processing methods that the Cisco Policy Suite enforces on the network.

The Cisco Policy Suite ships with some standard base policies that serve as a starting point for customization to suit a Service Provider's specific business rules.

Policy Builder Repository Configuration

The Policy Builder uses a Subversion version control repository to store the configuration data created in the UI. The data entered in the UI is translated into XML (Eclipse Modeling Framework xmi files) when saved.

As work is done in the UI, changes are saved to a temporary directory on the perclient01. (The directory is specified in the Repository configuration dialog.) Therefore, you can log out and back in and the latest changes will remain. However, if someone else makes a change and commits, then your local changes are lost.

There are two options for saving configuration changes:

- Publish to Runtime
- Save to Client Repository

When saving to the client repository, the configuration is pushed to Subversion, but it is saved in a client only repository and not copied over to the runtime environment repository. If you 'Publish to Runtime', the configuration is saved to the client repository and also copied to the runtime environment repository. The CPS servers check the runtime environment repository for changes and will update automatically when changes are committed.

Best Practices

Typically, publishing configuration changes to a lab environment to run tests is best. And then when satisfied with the test results, you can publish the new configuration to a production environment.

Revert

As Subversion is a source code tracking repository, each version of a configuration is numbered and stored in the Subversion repository history. Therefore, it is also possible to revert to any version of a configuration. The Policy Builder does not have a way to do this via the GUI, but using the Subversion command line tools, any version of the configuration can be made the current revision. For more information, refer to [Subversion documentation](#) for how to use the command line tools.


Default Repositories

The CPS deployment installs Subversion and creates a default client and runtime repository. The Subversion repositories are synced using Subversion's Master/Slave replication between the pcrfclient01 and pcrfclient02 nodes.

- Client - <http://pcrfclient01/repos/configuration>
- Runtime - <http://pcrfclient01/repos/run>

The Policy Builder start screen shows a dialog that lets you define repositories and choose a repository to check out for editing. A repository definition named “Repository” is installed by default and uses the default client repository (<http://pcrfclient01/repos/configuration>). The default PB user (qns-svn) with the default password is also setup.

Figure 1: Choose Policy Builder Data



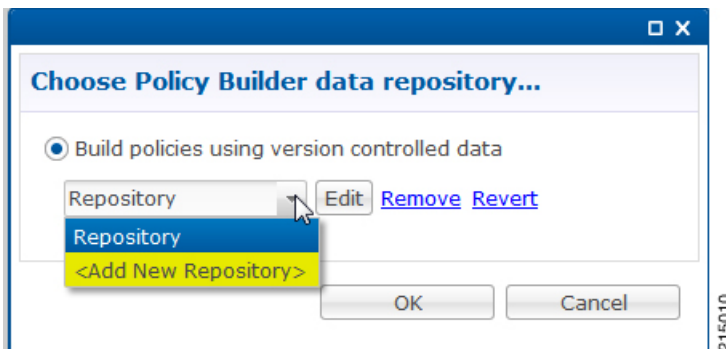
Adding a Client Repository Definition

Procedure

Step 1 Start Cisco Policy Builder.

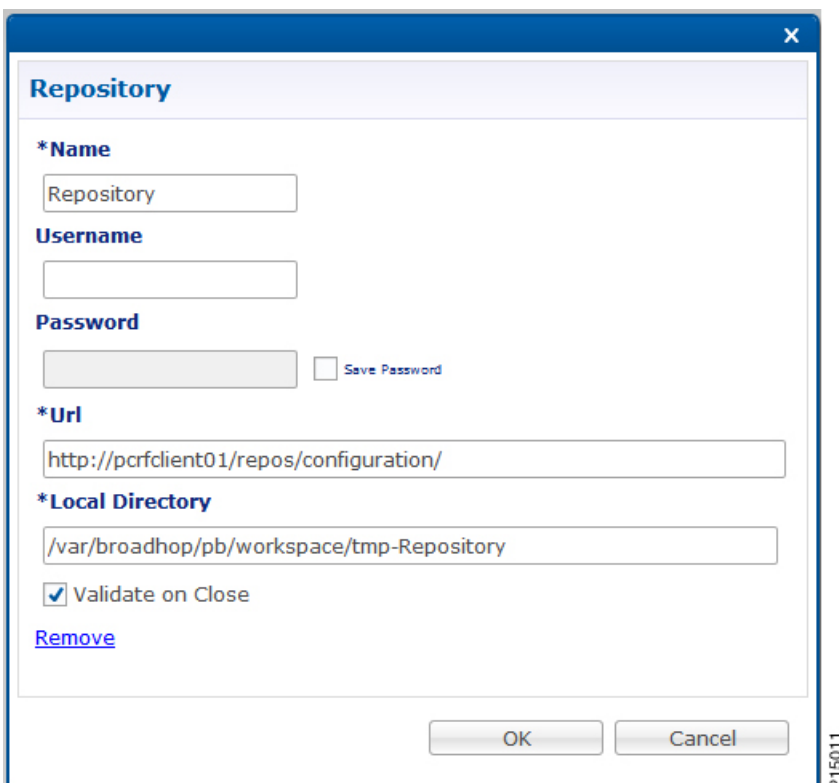
Step 2 In the **Choose Policy Builder data repository ...** dialog box, select **Add New Repository** from the drop-down list.

Figure 2: Adding a New Repository Definition



The **Repository** dialog box appears.

Figure 3: Repository Configuration Fields



The following parameters can be configured under **Repository**:

Configure the parameters according to the network requirements.

Table 1: Repository Parameters

Parameter	Description
Name	<p>This required field uniquely identifies your repository's site with a name.</p> <p>Note We recommend the following format for naming repositories: customername_project_date, where underscores are used to separate customer name, project and date. Date can be entered in the format: MMDDYYYY.</p>
Username and Password	<p>Enter a username that is configured to view Policy Builder data. The password can be saved for faster access, but it is less secure. A password, used with the Username, permits, or denies access to make changes to the repository.</p>
Save Password	<p>Select this check box to save the password on the local hard drive. This password is encrypted and saved as a cookie on the server.</p>
Url	<p>You can have several branches in the version control software to save different versions of configuration data. Create a branch in the version control software before assigning it in this screen.</p> <p>Enter the URL of the branch of the version control software server that is used to check-in this version of the data.</p>
Local Directory	<p>This value do not need to be changed.</p> <p>This is the location on the hard drive where the Policy Builder configuration objects are stored in version control.</p> <p>When you click either Publish or Save to Repository, the data is saved from this directory to the version control application specified in the Url text field.</p> <p>The field supports the following characters:</p> <ul style="list-style-type: none"> • Uppercase: A to Z • Lowercase: a to z • Digits: 1 to 9 • Non-alphanumeric: / <p>Note The user needs to provide the specified characters above.</p>
Validate on Close	<p>Select this check box to see if the values for Username, Password, or the URL are legitimate and unique. If not, the screen displays an error message and provides a chance to correct the errors.</p>

Parameter	Description
Remove	Removes the display of the repository in Cisco Policy Builder. Note The remove link here does not delete any data at that URL. The local directory is deleted.

Step 3 Click **OK** to save your work to the local directory.

Note

When your change screens, Cisco Policy Builder automatically saves your work. Cisco recommends saving your work to the local directory by clicking on the diskette icon on the Policy Builder GUI or **CTRL-S** on the keyboard.

Step 4 If you are ready to commit these changes to the version control software, select **File > Save to Client Repository** on the Policy Builder home screen.

Editing a Client Repository Definition

Use this procedure to change any of the following details of client repository:

- Client repository name
- Username, password, and password save mechanism
- Client repository temporary save URL
- Client repository local directory save file path

Procedure

Step 1 Open a browser and enter the URL of the Cisco Policy Builder.

Step 2 Use the drop-down list in the **Choose Policy Builder data repository...** dialog box to select the desired repository.

Step 3 Click the **Edit** button.

Step 4 In the **Repository** dialog box, make the required changes.

Step 5 Click **OK** to save the changes to the repository definition.

Removing a Client Repository Definition

This procedure removes a repository from Cisco Policy Builder. This procedure does not delete the actual Subversion repository, just the definition for access in the Policy Builder.

Procedure

- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.
 - Step 2** Use the drop-down list in the **Choose Policy Builder data...** dialog box to select the desired repository.
 - Step 3** Click **Remove**. A confirmation dialog box appears.
 - Step 4** Click **OK** to delete the repository.
-

Saving Policy Builder Configuration Data to a Client Repository

Procedure

- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.
 - Step 2** Use the drop-down list on the **Choose Policy Builder data...** dialog box to select the desired repository.
 - Step 3** Click **OK** to open the Policy Builder GUI.
 - Step 4** Make the necessary modifications in the Policy Builder.
 - Step 5** To save the modifications done, select **File > Save to Client Repository**, or click the diskette icon on the Policy Builder GUI or use CTRL-S on the keyboard.
 - Step 6** Enter a commit message for the modifications done.
 - Step 7** Click **OK**. The modified configurations are saved to the client repository for later updating and publishing to the runtime environment.
-

Publishing the Client Repository

To put changes into effect and have the Cisco Policy Builder server recognize the configuration changes made in your client session, use the Publish option and save the changes to the server repository.



Note To save the practice version, publish the client repository to the server. This is the version the server uses for production.

Do not publish to the Policy Builder unless you are completely satisfied with the configuration data in your client repository.

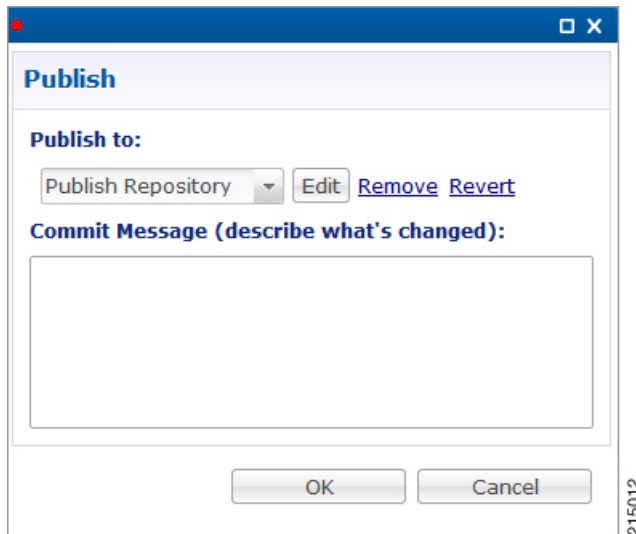
- Use the Policy Builder interface to either commit or set up a commit repository.
- Verify your work either by going to a web browser or by looking at the `config.properties` file.
- Unpublish with an SVN delete and restore.

When you are ready to put your Policy Builder changes into production, you need to publish them to Subversion. This preserves version history.

Procedure

Step 1 To publish in Cisco Policy Builder, select **File > Publish to Runtime Environment**. The **Publish** dialog box appears.

Figure 4: Publishing to the Runtime Environment



Step 2 If you have already set up the repository to publish to, just enter a commit message.

Step 3 If you have not set up the repository, select **Add New Repository** from the **Publish to:** drop-down list and enter the required details for the new repository. For more information, refer to [Adding a Client Repository Definition, on page 3](#).

Step 4 Verify the changes to Production repository:

- All changes are published to Subversion, so they are version-controlled and can be rolled back.
- To verify a publish as part of a troubleshooting process, take the URL seen in the previous screen and put it into a web browser (you may need to substitute the IP). The password is the same as in Cisco Policy Builder.
- If a traditional web browser cannot access the system, you can use a command line browser from the CPS VM's URL.

Adding a Runtime Repository Definition

A repository definition named **Publish Repository** is installed by default and uses the default Runtime repository (<http://pcrfclient01/repos/run>). The default Policy Builder user (*qns-svn*) with the default password is also setup. The Runtime repository matches the value setup in the `/etc/broadhop/qns.conf` file.

The `qns.conf` file is read by all of the active Policy Server and Policy Director nodes and when the policy server process starts up, it checks out the configuration from the Runtime repository.

Figure 5: Runtime Repository Definition



215013

Procedure

-
- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.
 - Step 2** Use the drop-down list on the **Choose Policy Builder data repository...** dialog box to select the desired repository.
 - Step 3** After selecting the required repository, click **OK**.
 - Step 4** Make changes to Policy Configuration data as necessary.
 - Step 5** Select **File > Publish to Runtime Environment...**
 - Step 6** Use the drop-down list to select **<Add New Repository>**.
The **Repository** dialog box appears.
 - Step 7** Enter the necessary values and click **OK** to save your work.
 - Step 8** Enter a commit message and click **OK** to publish to the new repository.
-

Editing a Runtime Repository Definition

Procedure

- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.
 - Step 2** Use the drop-down list on the **Choose Policy Builder data repository...** dialog box to select the desired repository.
 - Step 3** After selecting the required repository, click **OK**.
 - Step 4** Make changes to Policy Configuration data as necessary.
 - Step 5** Select **File > Publish to Runtime Environment...**
 - Step 6** Use the drop-down list to select the desired repository.
 - Step 7** In the **Repository** dialog box, make your changes.
 - Step 8** Click **OK** to save the changes to the repository definition.
 - Step 9** Enter a commit message and click **OK** to publish to the new repository.
-

Removing a Runtime Repository Definition

This procedure removes a runtime repository definition from the Cisco Policy Builder. This procedure does not delete the actual Subversion repository, just the definition for access in the Policy Builder.

Procedure

- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.
 - Step 2** Use the drop-down list on the **Choose Policy Builder data repository...** dialog box to select the desired repository.
 - Step 3** After selecting the required repository, click **OK**.
 - Step 4** Make changes to Policy Configuration data as necessary.
 - Step 5** Select **File > Publish to Runtime Environment...**
 - Step 6** Use the drop-down list to select the desired repository.
 - Step 7** Click **Remove**. A confirmation dialog appears.
 - Step 8** Click **OK** to delete the repository.
 - Step 9** Click **Cancel** to close the dialog box.
-

Saving Policy Builder Configuration Data to a Runtime Repository

Procedure

- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.

- Step 2** Use the drop-down list in the **Choose Policy Builder data repository...** dialog box to select the desired repository.
 - Step 3** After selecting the required repository, click **OK**.
 - Step 4** Make changes to Policy Configuration data as necessary.
 - Step 5** Select **File > Publish to Runtime Environment...**
 - Step 6** Use the drop-down list to select the desired repository.
 - Step 7** Enter a commit message.
 - Step 8** Click **OK**. The data will be saved to the client repository for later updating and publish to the runtime environment.
-

Switching to a Different Client Repository

You may have several variations of your client repository. One may reflect the configuration currently published to the server. Another might be developed for test purposes.

There are two ways to switch to a different repository:

- **File > Switch Repository...**
- **File > Exit/Logout:** You can select the required repository from **Choose Policy Builder data repository...** dialog box.

When you click **Ok** after repository selection, a validation prompt **Are you sure?** is displayed.

- Select **Discard** to discard the uncommitted changes.
- Select **Retain** to retain the uncommitted changes.

If the dialog box is dismissed without clicking escape or close, the uncommitted changes are retained.

Reverting Changes

There are two main SVN repositories (repos) in the system.

- Repository publish which contains ONLY the currently running set of policies.
- Runtime repository which contains a copy of the currently running set of policies along with copies of all previous sets.

To rollback Policy Builder changes, there are two methods:

- Rollback the configuration repository Policy Builder and then perform a publish as described in [Unpublished Changes, on page 10](#).
- Rollback the runtime repository Policy Builder uses and the configuration repository Policy Builder uses.

For more information, refer to [Published Changes, on page 11](#).

Unpublished Changes

If you do not want to save the changes, click the **Revert** link on the Policy Builder start window. All changes that have not been committed to a repository will be removed.

Procedure

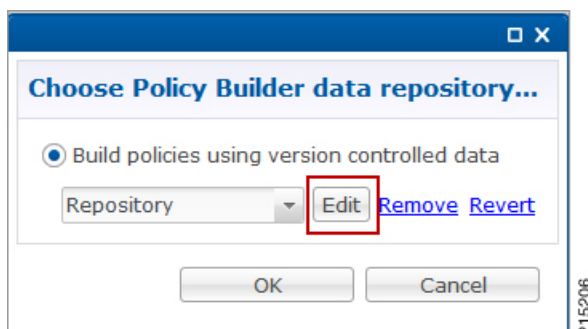
-
- Step 1** Open a browser and enter the URL of the Cisco Policy Builder.
- Step 2** Use the drop-down list in the **Choose Policy Builder data repository...** dialog box to select the desired repository.
- Step 3** Click the **Revert** link. A confirmation dialog appears.
- The **Revert** link is only available if there are uncommitted local changes.
- Step 4** Click **OK** to revert changes to the repository definition.
-

Published Changes

Procedure

-
- Step 1** Check the configuration repository name Policy Builder uses (config_repo). To check the name, use the following steps:
- Open a browser and enter the URL of the Cisco Policy Builder.
 - In the **Choose Policy Builder data repository...** dialog box, click **Edit**.

Figure 6: Choose Policy Builder Data Repository



- In the **Repository** dialog box, look at the contents of the **Url** field to see the repository name used by the Policy Builder. For example, it is **configuration**.

Figure 7: Repository

d) Record the Policy Builder repository name (config_repo). For example, it is **configuration**.

Step 2 To locate the ‘r’ number in the repository used by Policy Builder, execute the following command:

```
svn log http://pcrfclient01/repos/<config_repo> | more
```

The <config_repo> value comes from Step 1.d, on page 12.

The following is an example of the `svn log` command, where <config_repo> is **configuration** as shown in Step 1.d, on page 12.

```
svn log http://pcrfclient01/repos/configuration | more
-----
r367 | qns-svn | 2015-06-18 12:15:34 -0600 (Thu, 18 Jun 2015) | 1 line
second try
-----
r364 | qns-svn | 2015-06-17 15:46:19 -0600 (Wed, 17 Jun 2015) | 1 line
corrected java issue
-----
r361 | qns-svn | 2015-06-16 15:38:28 -0600 (Tue, 16 Jun 2015) | 1 line
Added new Policies
-----
r358 | qns-svn | 2015-06-16 15:06:57 -0600 (Tue, 16 Jun 2015) | 1 line
""
-----
r355 | qns-svn | 2015-06-16 14:58:41 -0600 (Tue, 16 Jun 2015) | 1 line
""
-----
r352 | qns-svn | 2015-06-16 14:52:29 -0600 (Tue, 16 Jun 2015) | 1 line
```

- a) In the example above, the comment we are looking for is in **r361** which is the ‘r’ number we want to rollback to.
- b) Record the **config_repo** ‘r_number’. In this example, it is **r361**.

Step 3 Execute the following command to delete the current version from the configuration repository Policy Builder uses:

```
svn delete http://pcrfclient01/repos/<config_repo> -m 'deleting for rollback'
```

Use the *<config_repo>* value from Step 1.d, on page 12.

The following is an example of the `svn delete` command where *<config_repo>* is **configuration**.

```
svn delete http://pcrfclient01/repos/configuration -m 'deleting for rollback'
```

Step 4 Execute the following command to restore the Policy Builder configuration repository to a previous version.

```
svn cp http://pcrfclient01/repos/<config_repo>@<r_number> http://pcrfclient01/repos/<config_repo> -m 'rolling back to <r_number>'
```

The *<r_number>* value is from Step 2.a, on page 13 and the *<config_repo>* value is from Step 1.d, on page 12. The ‘-m’ option should be used to add a comment indicating what is being done.

The following is an example of the `svn copy` command with the *<r_number>* set to **361** and the *<config_repo>* set to **configuration**:

```
svn cp http://pcrfclient01/repos/configuration@361 http://pcrfclient01/repos/configuration -m 'rolling back to 361'
```

Step 5 Execute the following command to verify if the rollback is successful:

```
svn log http://pcrfclient01/repos/<config_repo> | more
```

The *<config_repo>* value is from Step 1.d, on page 12.

The following is an example of the `svn copy` command:

```
svn log http://pcrfclient01/repos/configuration | more
-----
r367 | qns-svn | 2015-06-18 12:15:34 -0600 (Thu, 18 Jun 2015) | 1 line
rolling back to 361
-----
```

Note

The output should have the ‘-m’ option's text entered in Step 4, on page 13 as the comment.

Step 6 Open Policy Builder and verify the policies to which you have rolled back. Normally the customer should be able to verify the policies in Policy Builder.

Step 7 Perform a publish in Policy Builder and make sure to add a comment indicating that the publish is being done to complete the rollback. For example, “publishing to complete rollback to *<r_number>*”.

System Configuration

The Systems node in the Reference Data tab represents the Cisco Policy Suite runtime environment as it exists in the network environment.

- System: There must always be at least one system defined in the Policy Builder. The system represents the customer deployment. In HA, the system represents a set of PCRF clusters that share the same session

database. System is used to define any common things across the clusters, such as load balancing, and so on.

- **Cluster:** Each system contains one or more clusters - each of which represents a single High-Availability (HA) site environment. A cluster is used for define the configurations related to the blades. A cluster shares the same set of policy directors (that communicates as a group). A customer can take a fully installed PCRF and replicate it to a second cluster.

Each cluster can contain node instances. A node instance corresponds to a physical node in a deployment cluster such as a session manager or Policy Director (load balancer). It is very rare that a deployed system needs to have node instances configured in the Policy Builder. Configurations flow downhill, meaning that if you define a Plugin Configuration for Unified API at the system level, each cluster and subsequently each instance gets that configuration by default.

There are two types of clusters: HA and GR. This document discusses HA clusters. For information related to GR clusters, refer to *CPS Geographic Redundancy Guide* for this release.



Note In an HA environment you should not make any configuration in Cluster node.

Plug-in configuration done at cluster level overrides the same definition at system level. For example, if you configure Custom Reference Data at cluster level, it will override the Custom Reference Data configuration done at system level.

There is a default deployment configuration for mobile. **system-1** is the default system name and **cluster-1** is the default cluster name.

If a customer wants to change the system name, they need to change it in `qns.conf` (`/etc/broadhop/qns.conf`) file also to reflect it in Policy Builder:

```
-Dcom.broadhop.run.systemId=<system name>
```

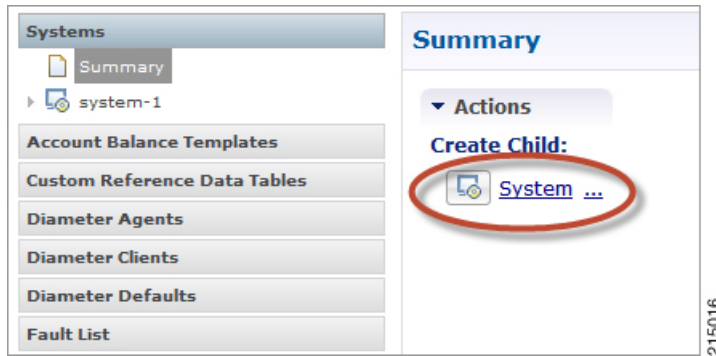
Adding a System

After installation, use this procedure to set up your Cisco Policy Builder by using an example populated with default data. You can change anything that does not apply to your deployment.

Procedure

Step 1 Click the **Reference Data** tab, and then click the **Systems** node to display the **Systems** tree.

Figure 8: Systems Tree



Step 2 Click **System...** under **Create Child:** to open the **System** pane on the right side.

Figure 9: System Pane

Step 3 Fill in the **Name** field, and provide a description of this system. Enter the rest of the parameters based on your network requirements.

Table 2: System Parameters

Parameter	Description
Name	The name of the CPS system.
Description	Description of this entire system.
Session Expiration Hours	<p>If no messages are received in x hours, the session will be removed.</p> <p>Note The maximum value allowed for this parameter is 590 hours. However, the combination of the number of hours specified for this parameter and the number of minutes specified for the Session Expiration Minutes parameter cannot exceed 35,400 minutes.</p> <p>Default value is 8.</p>

Parameter	Description
Session Expiration Minutes	<p>If no messages are received in x minutes, the session will be removed.</p> <p>Note The combined value of Session Expiration Hours multiplied by 60 plus Session Expiration Minutes should not exceed 35,400.</p> <p>Default value is 0.</p>
Timeout for Unknown Session	<p>Time in minutes that CPS keeps a session alive after the subscriber logs off. With this, other network entities involved in the session can let the session close gracefully.</p> <p>Default value is 0.</p>
Timeout For Soft Delete	<p>Determines the time in seconds during which a 'soft delete' session is maintained for a CPS session after session stop.</p> <p>Default value is 30.</p>
Session Limit Overload Protection	<p>This parameter is used to protect the session database from crashing. CPS does not allow session creation when a current system session count exceeds the Session Limit Overload Protection value.</p> <p>The default value is set to 0 which infinitely accepts the Diameter messages and CPS triggers alarms so that you change the value before session count goes beyond the database capacity. This value must be replaced by session capacity that is calculated for each deployment.</p> <p>Default value is 0.</p> <p>Note Session Limit Overload Protection value must be changed than using the default value (0).</p> <p>Note The recommended value for Session Limit Overload Protection has to be derived by Cisco Account representative for each deployment.</p>
Enable Multi Primary Key	<p>Select this check box to allow two primary keys to be utilized by maintaining a map of each separate primary key and storing the 'true' multi-primary key as a UUID related to the two maps. Changing this setting has a negative performance impact and should only be done at the request of the BU. Recommendation is to keep Enable Multi Primary Key unchecked.</p> <p>Default is unchecked.</p>
Enable Number Normalization	<p>Select this check box to enable number normalization under system configuration. The following fields are displayed under Number Normalization List:</p> <ul style="list-style-type: none"> • Number Type - Type of IMSI/MSISDN. • Number Length - Length of the normalized IMSI/MSISDN. • Number Prefix - Prefix to be normalized from the number.

Parameter	Description
Cluster link	Click this link to create a cluster under this system.
Current System Link	Click this link to make a copy of this system, with its clusters and instances.

Step 4 If the created system needs to be used, then after publishing, the following property needs to be updated in the `qns.conf` configuration file:

```
-Dcom.broadhop.run.systemId=<system name>
```

where `<system name>` is the system name defined in the [Step 3, on page 15](#).

Soft Delete Session

A soft delete session is an entry in the session database which maintains session data after session stop with an auto-generated unique primary key, but still maintains needed secondary keys. This allows messages which come after session stop to still be processed while also allowing a session with the same primary key to be immediately created. The CPS code determines when soft delete sessions are required and what secondary keys are needed.

Soft Delete Example (Mobile)

A Gx session with a Gy associated session exists. A Gx CCR-T is received that terminates the CPS session, resulting in a soft-delete session which contains Gy session information and associated Gy secondary keys. A Gy CCR-t is received and the soft-delete session is loaded and updated with the charging information through the end of the session. After the soft delete timeout, the soft delete session is removed.

Adding an HA Cluster

This section describes how to add an HA cluster. A system, a cluster, and an instance are set up at install time. If you have to change the cluster definition, or want to add more clusters, use these steps.

Procedure

Step 1 In the **REFERENCE DATA** tab, under the **Systems** node, click your default system.

Figure 10: System Configuration

The screenshot shows the 'System Configuration' page for a system named 'My System'. The left sidebar contains a navigation menu with items like 'Summary', 'system-1', 'My System', 'Account Balance Templates', 'Charging Rule Retry Profiles', 'Custom Reference Data Tables', 'Diameter Agents', 'Diameter Clients', 'Diameter Defaults', 'Fault List', 'Ldap Server Sets', 'Notifications', 'Policy Reporting', 'Subscriber Data Sources', and 'Tariff Times'. The main content area is titled 'System' and contains the following fields:

- *Name:** My System
- Description:** Built
- *Session Expiration Hours:** 8
- *Session Expiration Minutes:** 0
- *Timeout For Unknown Session:** 0
- *Timeout For Soft Delete:** 30
- Enable Multi Primary Key**
- Actions:**
 - Create Child:**
 - Cluster** (highlighted with a red box)
 - Copy:**
 - Current System

215018

Step 2 Under Actions, click the **Cluster** link to set up your cluster.

Since some data is relevant at the cluster level, you must always have at least one cluster (by default), even if it is a cluster of one instance.

Configure the following cluster parameters.

Table 3: Cluster Parameters

Parameter	Description
Name	The name of the cluster. This name must correspond to the value stated in the <code>com.broadhop.run.clusterId</code> parameter in <code>qns.conf</code> file on the Cisco Policy Server.
Description	A brief description of the cluster.
Db Write Concern	<p>Controls the write behavior of the Session Manager, and specifies for what errors exceptions are raised. Db Write Concern defined in Cluster page applies only to Admin, Trace, and Endpoint databases.</p> <p>Select one of the following options from the drop-down list:</p> <ul style="list-style-type: none"> • <code>OneInstanceSafe</code>: The system waits for writing confirmation in primary member. • <code>TwoInstanceSafe</code>: The system waits for writing confirmation in the primary and one secondary member. <p>Default value is <code>OneInstanceSafe</code>.</p> <p>For more information, see MongoDB documentation.</p>

Parameter	Description
Failover Sla Ms	<p>Specifies the time period, in milliseconds, to wait before starting failover database handling.</p> <p>Default value is 0 milliseconds.</p>
Replication Wait Time	<p>Specifies the time limit, in milliseconds, for the Db Write Concern. This parameter is applicable only if you select TwoInstanceSafe in Db Write Concern.</p> <p>This parameter causes write operations to return with an error after the specified limit, even if the required write concern eventually succeeds. When these write operations return, MongoDB does not undo successful data modifications performed before the write concern exceeded the replication wait time limit.</p> <p>Default value is 100 milliseconds.</p>
Trace Db Size Mb	<p>Determines the size of the policy_trace database capped collection in megabytes. The capped collection never grows higher than the configured size.</p> <p>This parameter is applied to the Trace database.</p> <p>Default value is 512.</p> <p>For more information, see <i>Policy Tracing and Execution Analyzer</i> section in <i>CPS Operations Guide</i>.</p>
SK Evaluation Time In Minutes	<p>This parameter is used to avoid querying same secondary keys.</p> <p>Default vale is 60 minutes.</p>

Parameter	Description
Max Timer T P S	<p>Specifies the maximum number of internally generated transactions per second (TPS) the system produces.</p> <p>This parameter affects the RAR generated by CPS when they are triggered by an internal time event (change of time or quota refresh).</p> <p>For example, if the system needs to generate RAR messages to refresh quotas, the max TPS for the creation of the RAR messages is limited by this value.</p> <p>Default value is 2000.</p> <p>Internal timer TPS is calculated based upon the value specified for Max Timer TPS in Policy Builder, number of session shards, and number of Policy Server (QNS) VMs.</p> <p>Max Timer Expired TPS = (Max Timer TPS / Number of Session Shards) * (Lower of Number of Session Shards or Policy Server (QNS))</p> <p>Here are examples that specify how the internal timer TPS is calculated:</p> <ol style="list-style-type: none"> a. When number of session shards are higher or same as number of QNS VMs: If configured Max Timer TPS in Policy Builder is 2000, number of session shards are 88 and number of QNS are 50. In this case, max TPS per QNS VM is $2000/88$ (no. of session shards) = 22 (value in floored). Actual maximum timer TPS is $22 * 50$ (no. of QNS) = 1100. b. When number of session shards are lower than QNS VMs: If configured Max Timer TPS in Policy Builder is 2000, number of session shards are 12 and number of QNS are 20. In this case, max TPS per QNS VM is $2000/12$ (no. of session shards) = 166 (value is floored). Actual maximum timer TPS is $166 * 12$ (no. of session shards) = 1992. <p>Note In case there is CPS blade expansion, Max Timer T P S parameter must to be tuned to handle internal RAR TPS. For more information, contact your Cisco Account representative.</p>

Parameter	Description
Re-evaluation diffusion buckets	<p>Specifies the number of batches, or “buckets”, into which CPS divides the transactions to be processed when the rate limiting TPS function of CPS is triggered. The rate limiting feature is defined in the Max Timer T P S field above.</p> <p>When a TPS rate limit condition is encountered, CPS divides the total number of re-evaluation transactions to be processed by the number of buckets defined in this field. CPS processes the transactions in each bucket, one by one, while also adding a brief delay between the processing of each bucket. Refer to the Re-evaluation diffusion interval parameter below to configure this delay.</p> <p>This functionality prevents spikes in traffic when a large number of sessions qualify for re-evaluation at the exact same time. This functionality only diffuses the rate at which transactions are sent to the CPS policy engine. It does not wait for end-to-end processing to occur.</p> <p>To disable this functionality, set both Re-evaluation diffusion buckets and Re-evaluation diffusion interval parameters to 0 (zero).</p> <p>If the TPS rate limiting functionality has been disabled, this functionality is also disabled.</p> <p>Default value is 50 (buckets).</p>
Re-evaluation diffusion interval (in milliseconds)	<p>Specifies the delay, in milliseconds, before processing the next bucket. Enter the sum of all the delays between all the buckets.</p> <p>Assuming Re-evaluation diffusion buckets is configured as 50, and the Re-evaluation diffusion interval is configured as 20000 milliseconds, it introduces a delay of 408 milliseconds before proceeding with the next bucket of transactions.</p> <p>Delay between buckets = Diffusion Interval (ms) / (Diffusion Buckets - 1)</p> <p>Delay between buckets = 20000/49</p> <p>Delay between buckets = 408 (ms)</p> <p>Default value is 20 milliseconds.</p> <p>With default configuration, there will be no diffusion. As 20/50 will be 0 (It ignores decimal).</p>
Broadcast Msg Wait Timer Ms	<p>Specifies the time period, in milliseconds, for the Policy Engine to wait between sending each Broadcast Policy Message.</p> <p>Default value is 50.</p>
Max Sessions Per Shard	<p>Specifies the maximum number of sessions that can be stored in each shard that has been created within a Session Database Replica Set.</p> <p>Value ‘0’ in this parameter means that monitoring is not implemented.</p> <p>Note It is recommended to add new shards and rebalance sessions during maintenance window.</p> <p>Default value is 0.</p>

Parameter	Description
Lookaside Key Prefixes	<p>To improve Gx/Rx lookup and caching performance, we can add the lookaside key prefixes. In order to identify the correct shard for subscriber lookup/query, the PCRF needs to know the secondary key (which is internally stored in secondary key cache) for mapping and the exact shard that will be queried for subscriber data. This helps prevent the system from scanning/querying all the available shards in the system to fetch the subscriber record. Reducing the data range for scanning/querying leads to enhanced system performance.</p> <p>The following keys should be added so that the secondary keys for session binding are stored in the secondary key cache.</p> <ul style="list-style-type: none"> • diameter - This key should be added for local session affinity or when Sh interface is in use. • RxTGPPSessionKey - This key should be added only when Rx interface is used. • FramedIpKey - This key should be added if Gx Session lookup is based on IPADDRESS. • framedIpV6PrefixKey - This key should be added if Gx Session lookup is based on IPADDRESS. • USuMSubscriberIdKey – This key should be added only when SPR is used. • MSBMSsubscriberIdKey – This key should be added only when balance is used. • RxClientSessionKey - This key should be added if PAS health check feature is enabled. • imsiApnKey - This key should be added if Gx Session lookup based on IMSI and APN is enabled or Load by IMSI and Load By Called Station ID is enabled. • msisdnApnKey - This key should be added if Gx Session lookup based on MSISDN and APN is enabled. • userIdentityKey - This key should be added when Sh related messages have 'msisdn' AVP or Sh profile has sipUri related information. <p>CPS uses this key to find the session based on the message AVPs present in the request message. For example, when MSISDN is present in request message and sipUri information is present in profile.</p> <p>This would prevent the system from scanning/querying all the available shards in the system to fetch the subscriber record, which eventually leads to enhanced system performance.</p> <p>Note For GR setup, a Lookaside Key Prefix called single must be added under the Cluster session for each Cluster for faster session Lookups based on memcache when SingleSy feature is enabled.</p>

Parameter	Description
Disable Secondary Key Full Scan DB For Selected Keys	<p>Specifies whether full scan for secondary key database lookups is enabled or disabled. By default, it is enabled.</p> <p>Disabling secondary key database lookups helps in reducing PCRf processing latencies.</p> <p>To disable full scan for specific keys:</p> <ul style="list-style-type: none"> • Uncheck Disable Secondary Key Full Scan DB For All Keys • Add the list of key prefixes to be disabled for full scan
Disable Secondary Key Full Scan DB For All Keys	<p>Enable or disable full scan for secondary key database lookups. By default, the secondary key database lookups is enabled (unchecked).</p> <p>Disabling secondary key database lookups helps in reducing PCRf processing latencies.</p> <p>If this is checked, when the memcache look up fails for a secondary key, the full database scan to find the session in database is not performed.</p> <p>Note When the secondary full scan database is disabled, number of 5065 errors thrown increases as memcache works on UDP and looking memcache alone can cause missed sessions count to increase.</p>
Suppress Error Audit Traces To Trace DB	<p>When enabled the application does not write error audit traces logs into <code>error_traces</code> collection.</p> <p>When disabled the application writes error audit traces logs into <code>error_traces</code> collection.</p> <p>Default value is enabled.</p>
Session Tag Padding Configuration	<p>To enable the Session Tag Padding Configuration, select (true) the checkbox. By default the checkbox is not selected (false).</p> <p>Makes sure that the MongoDB subscriber session size is consistent throughout the subscriber session life cycle.</p> <p>Note After performing a fresh installation or an upgrade, you must ensure the following conditions are met:</p> <ul style="list-style-type: none"> • In multi-cluster setup, session tag padding feature related configuration should be same for all the PCRf clusters. • In multi-cluster setup, session tag padding feature related configuration should be same for all the UDC clusters. • Session tag padding feature configurations can be different for PCRf and UDC configuration.

Parameter	Description
Enable Session Query Without Padding	<p>Enables the system to query for secondary key without padding. This parameter is required when there are old subscriber sessions (tag value without padding) and new subscriber sessions (with tag value with padding character) exists in the MongoDB. In case the system does not find the session using the key with padding data, system attempts to find the session using key without padding. This flag enables the system to trigger the query without padding.</p> <p>By default the checkbox is selected (true).</p> <p>Note For fresh installations, set Enable Session Query Without Padding to false and publish the Policy Builder.</p> <p>When upgrading the system, set the Enable Session Query Without Padding to true and publish the Policy Builder.</p>
Initial Number Of Tags	<p>Determines the total number of tags (including reserved tags) in the tag list of subscriber session header part to be created while creating subscriber session.</p> <p>This helps CPS system to have the constant size for subscriber session at runtime.</p> <p>To disable this configuration, set the value to 0. Reserved tags are not created when the value configured as 0.</p> <p>Default value is 10.</p> <p>Note If the number of secondary key/tags exceeds the predefined <i>Initial Number Of Tags</i>, system stores the additional tags.</p> <p>For example, if <i>Initial Number Of Tags</i> is configured as 10 and at runtime list of tags (from subscriber session document) need to store 12 tags, system stores 12 tags without rejecting any request.</p>
Admin Database	<ul style="list-style-type: none"> • Primary Database IP Address: The IP address of the session manager database that holds session information for Cisco Policy Builder and Cisco Policy Server. • Secondary Database IP Address: The IP address of the database that provides fail over support for the primary database. <p>This is the mirror of the database specified in the Primary IP Address field. Use this only for replication or replica pairs architecture. This field is present but deprecated to maintain downward compatibility.</p> <ul style="list-style-type: none"> • Database Port: Port number of the database for session data. <p>Default value is 27717.</p> <p>It is recommended to set port to 27721 or any other port different than 27717 so that the admin database is not in the same replica-set as sessions (it is tmpfs) and data can be lost.</p>

Parameter	Description
Endpoint Database	<p>This database refers to the endpoint collections under diameter and queuing DB. The endpoint database under diameter database has the information about the peers between CPS and other nodes. For example, GW.</p> <p>The endpoint database under queuing database has the information about the connectivity between Policy Director (LB) and Policy Server (QNS) VMs.</p> <ul style="list-style-type: none"> • Primary Database IP Address: The IP address of the session manager database that holds session information for Cisco Policy Builder and Cisco Policy Server. • Secondary Database IP Address: The IP address of the database that provides fail over support for the primary database. <p>This is the mirror of the database specified in the Primary IP Address field. Use this only for replication or replica pairs architecture. This field is present but deprecated to maintain downward compatibility.</p> <ul style="list-style-type: none"> • Database Port: Port number of the database for Session data. <p>Default value is 27717.</p>
Trace Database	<ul style="list-style-type: none"> • Primary Database IP Address: The IP address of the session manager node that holds trace information, which allows for debugging of specific sessions and subscribers based on unique primary keys. • Secondary Database IP Address: The IP address of the database that provides fail over support for the primary database. <p>This is the mirror of the database specified in the Primary IP Address field. Use this only for replication or replica pairs architecture. This field is present but deprecated to maintain downward compatibility.</p> <ul style="list-style-type: none"> • Database Port: Port number of the database for Session data. <p>Default value is 27717.</p> <p>For more information, see <i>Policy Tracing and Execution Analyzer</i> section in <i>CPS Operations Guide</i>.</p>
Data Centre Parameter	Deprecated
Common Time Changes	Deprecated

Step 3 From the Systems tree, open up the cluster that you just added and check the plug-in configurations.

The configurations that you specify here are used only at the cluster level, and cascade down to the instance level if no configuration is set on the instance.

At this point, the plug-ins are available to the cluster but are not configured.

Click any one of them to open the detailed page in the right pane, and check and set your own configuration data. However, there is rarely a need to use the Threading Configuration or the Async Threading Configuration unless instructed to do so.

Step 4 If the configured cluster has to be used, after publishing, the following parameter must be updated in the `qns.conf` configuration file:

```
-Dcom.broadhop.run.clusterId=<cluster name>
```

where, `<cluster name>` is the cluster name defined in Policy Builder.

Adding an Instance

Procedure

Step 1 Begin with a Cluster at the **Systems** node in the **Reference Data** tab.

Step 2 Under **Create Child:**, click the **Instance** link to open the **Instance** pane.

Figure 11: Instance Configuration

Step 3 Type the **Name** and **Description**.

Step 4 From the **Systems** tree, open up the instance node that you just added and check the plug-in configurations.

At this point, plug-ins are available but not configured at the instance level.

Click any one of the plug-ins to open the detailed page in the right pane and check and set your own configuration data.

Any of the configuration data you have here are used at the instance level, overriding any plug-ins set at the system level or the cluster level.

Adding a GR Cluster

At install time, a system, cluster, and instance are set up. If you need to change the cluster definition, or want to add others, use the information mentioned in Adding an HA Cluster and Endpoint Database Configuration.

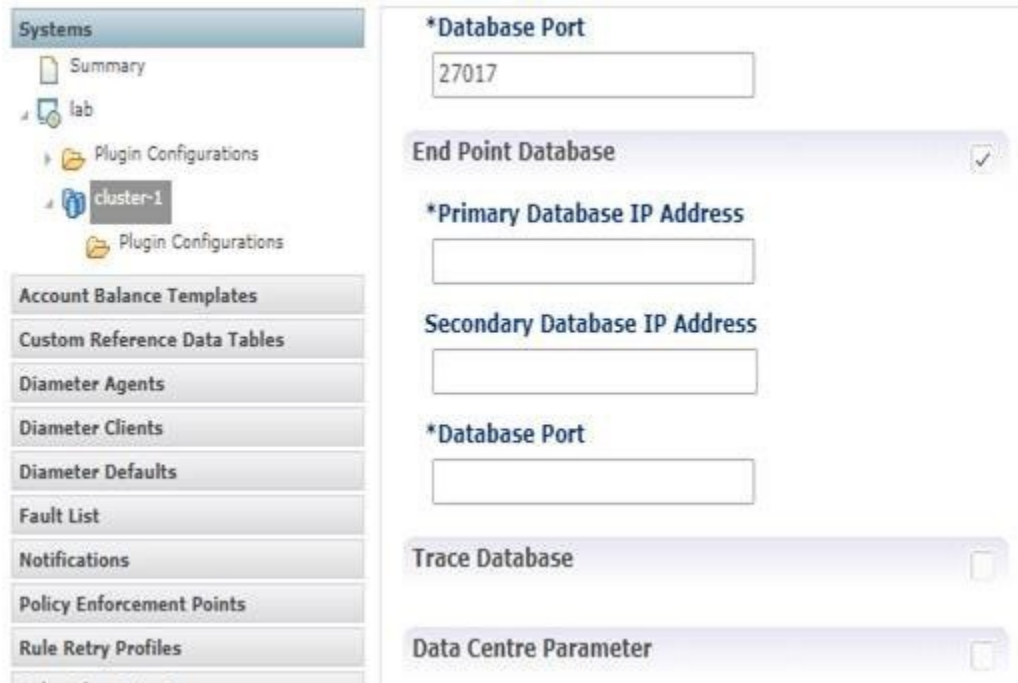
For configuration details, refer to [Adding an HA Cluster, on page 17](#).

For Endpoint Database parameters, refer to [Endpoint Database Configuration, on page 27](#).

Endpoint Database Configuration

For GR setup, Endpoint Database should be configured local to the site under Reporting/Balance Database.

Figure 12: Endpoint Database



The screenshot displays the configuration interface for the Endpoint Database. On the left, a navigation pane shows the following structure:

- Systems
 - Summary
 - lab
 - Plugin Configurations
 - cluster-1
 - Plugin Configurations
- Account Balance Templates
- Custom Reference Data Tables
- Diameter Agents
- Diameter Clients
- Diameter Defaults
- Fault List
- Notifications
- Policy Enforcement Points
- Rule Retry Profiles

The main configuration area on the right includes the following fields:

- *Database Port**: Text input field containing the value 27017.
- End Point Database**: Toggle switch, currently checked.
- *Primary Database IP Address**: Text input field.
- Secondary Database IP Address**: Text input field.
- *Database Port**: Text input field.
- Trace Database**: Toggle switch, currently unchecked.
- Data Centre Parameter**: Toggle switch, currently unchecked.

