Cisco Policy Suite 7.0 Wi-Fi Configuration Guide

First Published: September 26, 2014

Last Updated: July 10, 2015

Cisco Systems, Inc.
www.cisco.com

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

Text Part Number:

Preface

Welcome to the Cisco Systems, Inc. *Cisco Policy Builder Wi-Fi Configuration Guide*.

This document describes configuration tasks for the Cisco Cisco Policy Builder 7.0. The document assists installers, network operators, and network engineers to tune and configure the Cisco Policy Server 7.0 using the Cisco Policy Builder client interface.

This preface covers the following sections:

Audience

This guide is best used by these readers:

- Network administrators
- Network engineers
- Network operators
- System administrators

This document assumes a general understanding of network architecture, configuration, and operations.

For further documentation and support:

- Contact your Cisco Systems, Inc. technical representative.
- Write to Cisco Systems, Inc. at support@cisco.com
- Refer to your other documents.

Terms and Definitions

This document uses certain terms and definitions specific to the Cisco Policy Builder software application. Please refer to the common Cisco Glossary.

Version Control Software

Cisco Policy Builder uses version control software to manage its various data repositories. The default installed version control software is Subversion, which is provided in your installation package.

Logging In

To log in, you need to have the URL of the Cisco Policy Builder interface, and you may need a username, and a password. Cisco Policy Builder can be configured to either require or waive the use of a password.

The URL for your site is similar to this one:

http://<IPAddress:7070/pb>

where, IPAddress:7070 is the IP address specific to your network.

When you log in successfully, you see the top Cisco Policy Builder screen. The Reference Data tab is selected and displays the Reference Data tree on the left.

Multi-user Logins

Cisco Policy Builder also supports multiple user logins and uses a password to manage this.

When first logging in, and you want to configure Cisco Policy Builder to permit concurrent user logins, click Edit on the initial dialog or publish dialog.

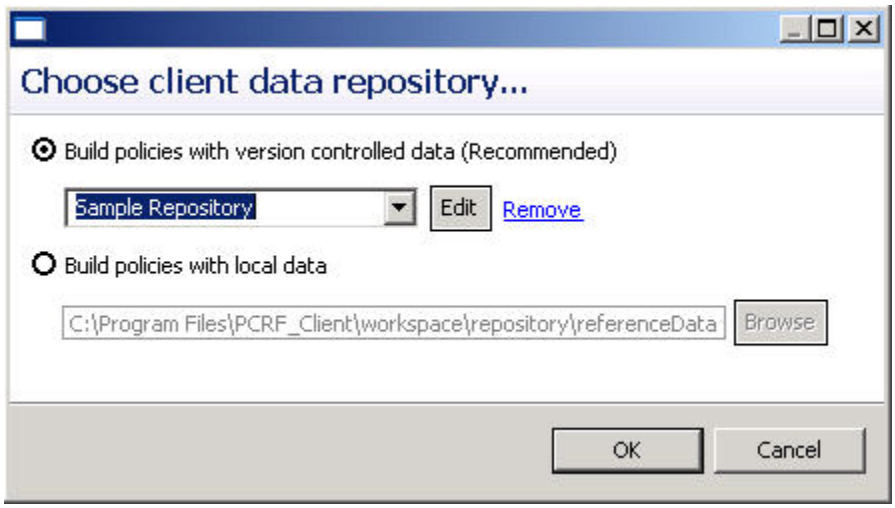

Note To redisplay the above screen after you have been working in Cisco Policy Builder, click Files > Switch Repositories.

Step 1 On the Repository screen, delete any values in the username and password fields.

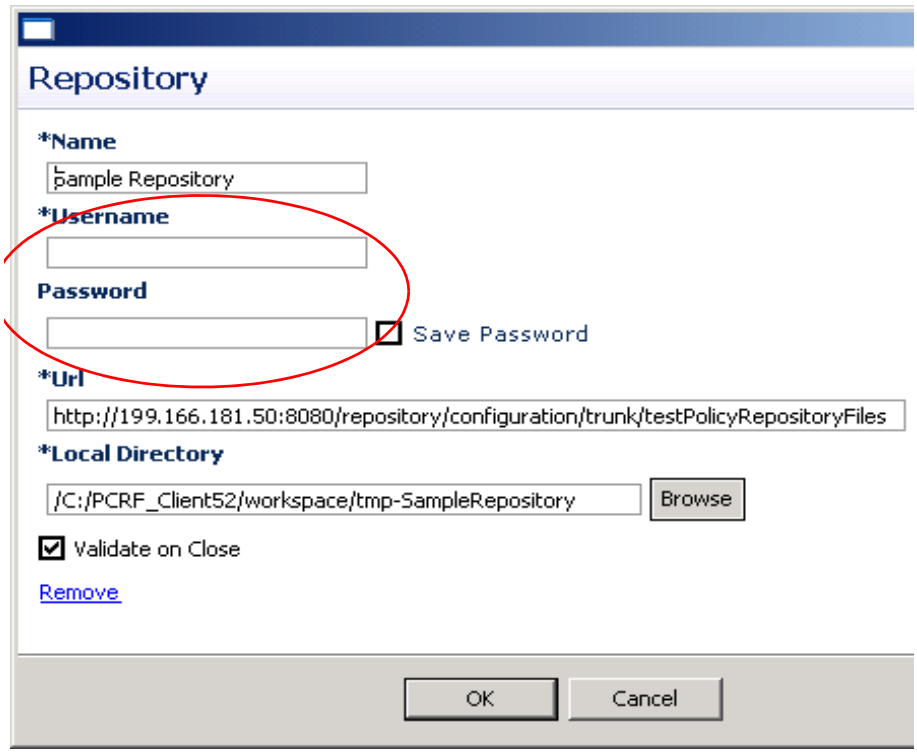

You are then prompted to enter a unique username and password.

This username/password is stored is cached for the duration of your session until a web browser refresh or when you switch repositories.

Understanding Services

Revised: July 10, 2015

Read this chapter before you begin configuring your service options and services so that you can quickly complete those steps.

Note The URL of the web-based Cisco Policy Builder is *https://lbvip01:7443/pb*.

This chapter covers the following sections:

Overview of Services

Configuration tasks are based on a hierarchy of concepts and software screens.

- Services screens and domain screens are closely related to the Cisco Subscriber Services Portal interface.
- Customer Reference Data Tables are closely related to the Cisco Control Center interface.
- Services are a collection of one or more Service Options and are offered to subscribers.
- Service Options are copies of use case templates and define specific aspects of a particular use case.
- Copy a use case template and either use the copy as is or change the values to match your needs.

  The use cases themselves, which you use as templates, are described on Base Use Case Templates Description.
- Cisco provides many use case templates, the basic construct. You can create others as well.

- You create the service options and services.

To examine these concepts in a hierarchical view, see the next section A Top-down View.

A Top-down View

Another way to understand the concepts used to configure services is to examine some example services and look at their contributing parts. As an example, the hierarchy of the service presented looks like this:

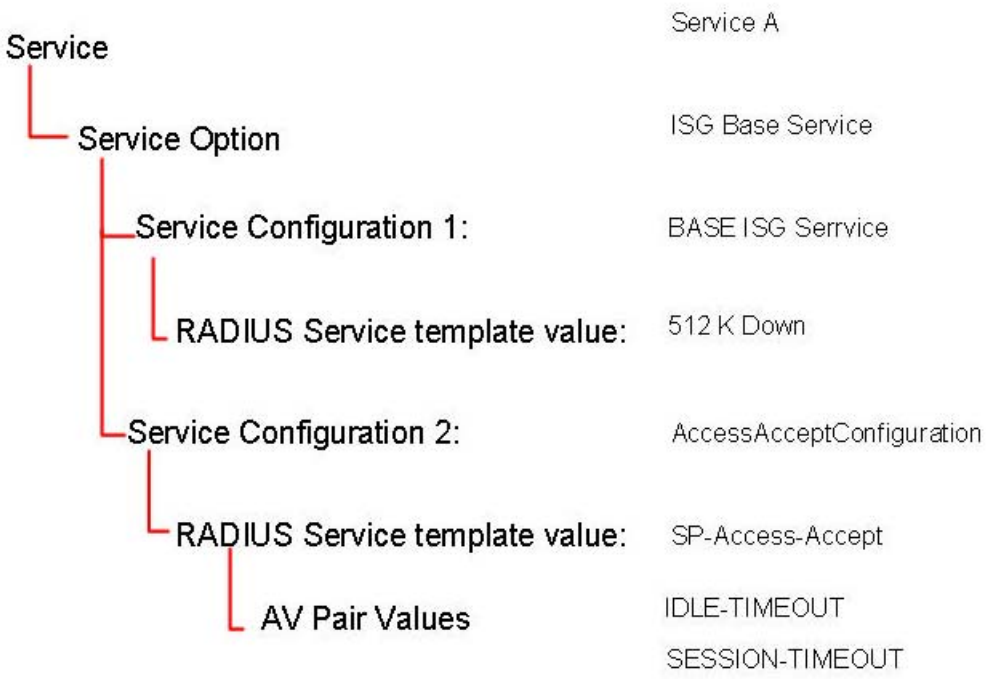

Next, view this hierarchy via the screens used to configure it.

Step 1 In a browser window, open a Cisco Policy Builder interface, or examine the figures provided here.

Step 2 Click Service tab > Services node > Services folder > Service A.

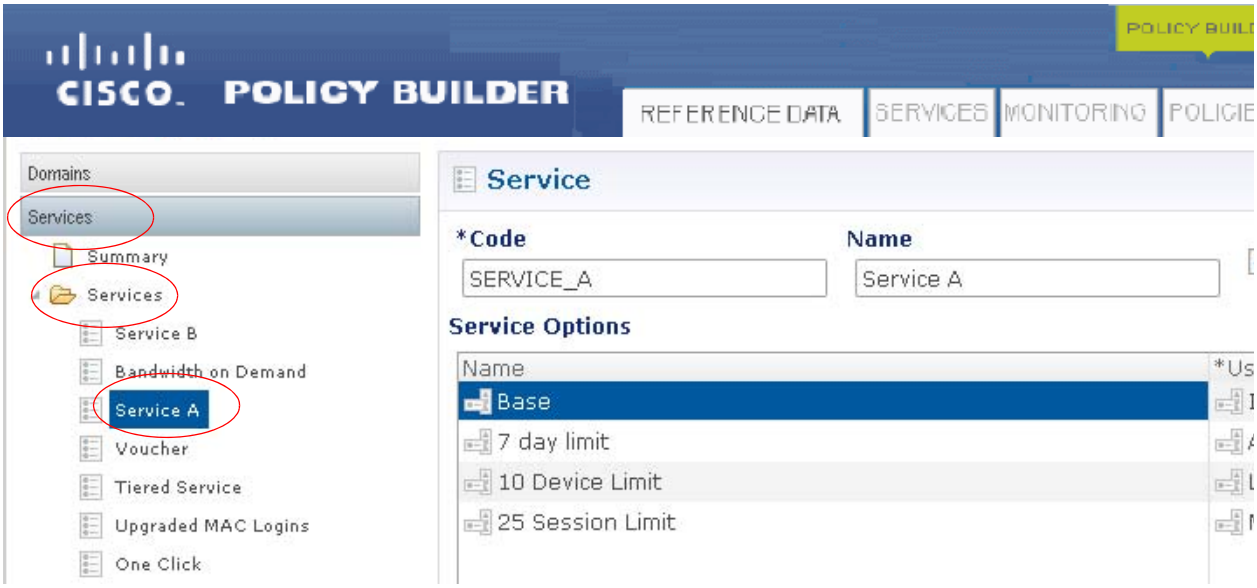

Service A is comprised of these service options:

- Base
- 7 day limit
- 10 Device Limit
- 25 Session Limit

This Service Option...	is based on this use case template:
Base	ISG Base Service
7 day limit	Auto Register MAC Credential
10 Device Limit	Limit Max MAC Registrations
25 Session limit	Max Concurrent Sessions

Step 3 Click the Base service option in the tree and look at its content.

The Base service option uses the Base ISG Service service configuration and uses the 512K DOWN service template for a transmission speed.

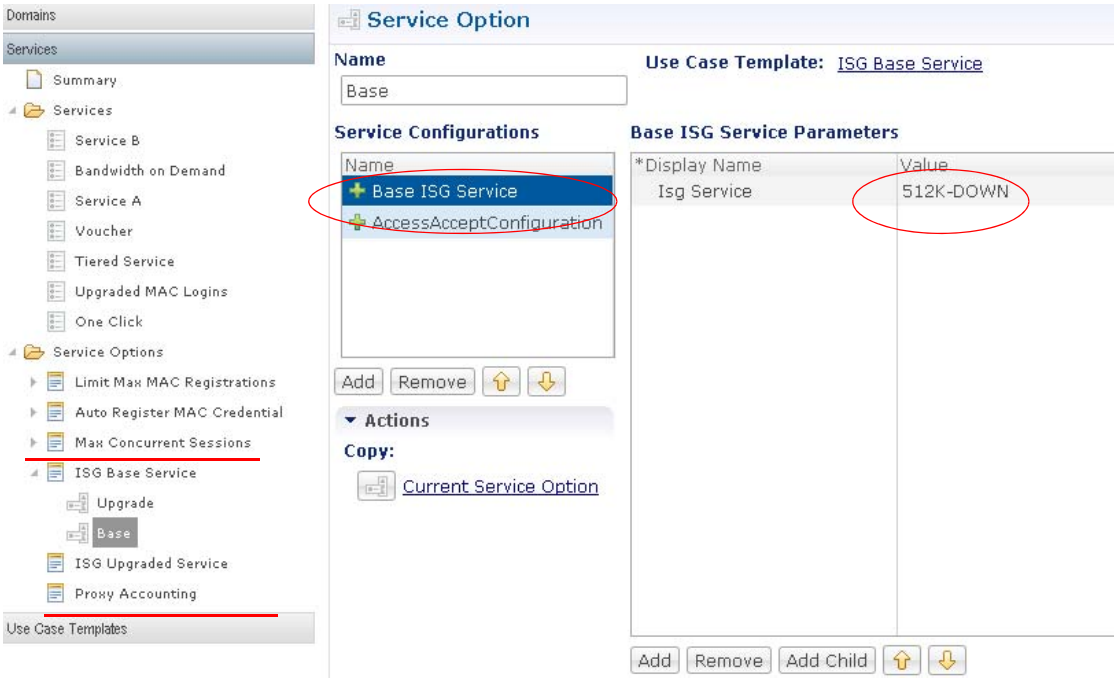

Note Check the RADIUS Service Template node under the Reference Data tab to see where this originated.

Step 4 Also notice the Service Option screen for ISG Base Service.

The Base Service Option also uses the Access Accept Configuration service configuration, which uses the ISG_ACCESS_ACCEPT RADIUS Service template.

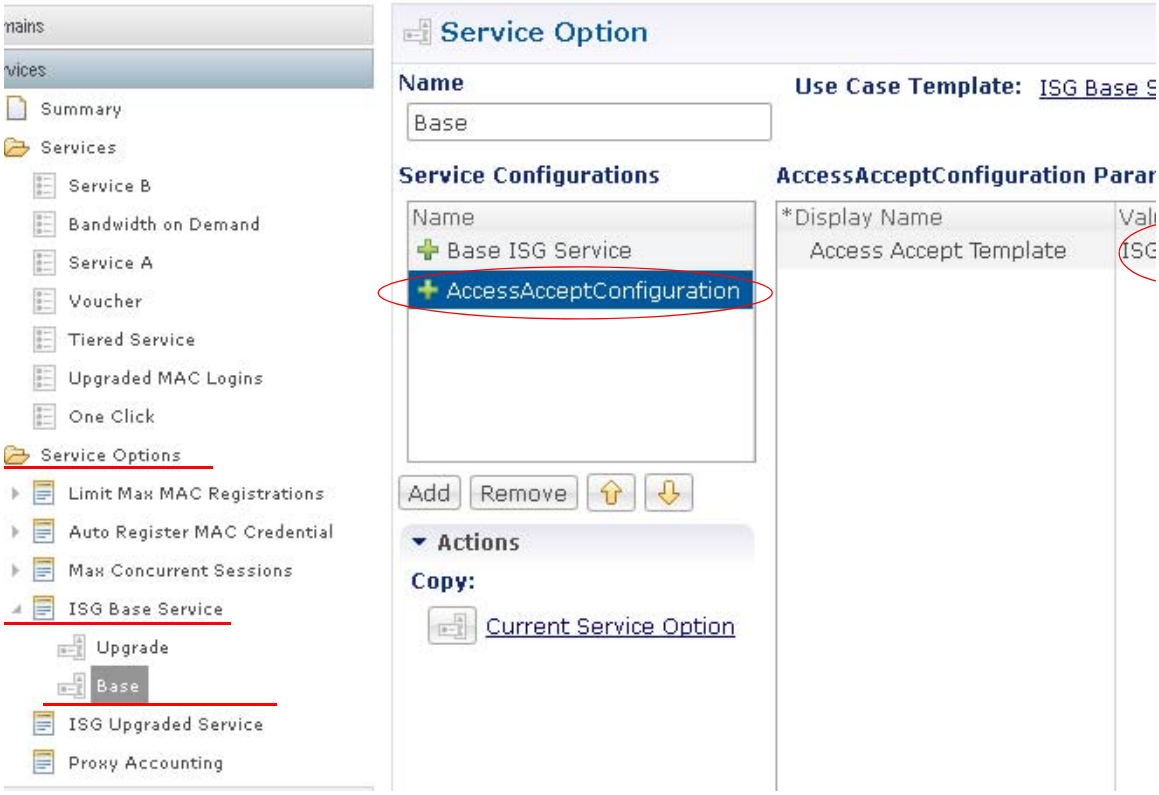

Note The ISG_ACCESS_ACCEPT is defined by the RADIUS Service Template for Service Provider Specific Template, in the AV pairs table. This template lets you specify idle time outs and session time outs, for example.

Only the Base service option is shown in this example, but the same relationships exist for the other service options that make up Service_A.

Configure Subscriber Services

Use the procedures in the following chapters to configure your subscriber services:

- Required Service Configuration shows the detailed steps for creating:
  - PEPs
  - RADIUS Service Templates
- Domain Configurations shows the detailed steps for creating:
  - Domains
- Elective Service Configurations shows the detailed steps for creating:
  - Service Options
  - Services

Base Use Case Templates Description

Use case templates are the building blocks of the Cisco Policy Builder service manager architecture. They are the basic constructs upon which you create service options.

Cisco Cisco Policy Builder provides the following uses cases for you to adapt or use as is. Recall that you do not use the provided use case templates. Rather, you copy them, calling them service options, and then make any changes you need in the parameter values.

For example, the use case template for ISG Base Service shows that the minimum number of authorized MACs is 5. For your service, you may want a subscriber's number of authorized MACs to be 10 to accommodate several devices for several different family members.

The use case templates provided by Cisco are these:

Limit Max MAC Registration	This use case restricts the number of devices a subscriber can register at the same time.
Auto Register MAC Credential	This use case causes the system to auto-register a MAC address as subscriber credential. Optionally, you can pick a period of time to keep the credential valid.
Max Concurrent Sessions	This use case limits the number of concurrent sessions a subscriber may have. The default is 1, but you may choose to allow more because subscribers may have more than one device.
ASR9K	The ASR9K is a device which can do the following: • Enforce QOS (Quality of Service or Bandwidth speed) to the end subscriber. • Receive accounting messages. • CPS can send down any RADIUS AVP to set QoS or other QoS policies on the ASR9K.
ISG Upgraded Service	This use case provides upgraded service above the normal base level service to a subscriber. For example, a turbo boost for a specified period of time is an upgraded service and the subscriber would revert to the lower level base service when the turbo boost has expired.
Proxy Accounting	This use case is preferred when RADIUS accounting messages need to be forwarded, or proxied, to another system.
ISG Base Service	This use case assigns the base or default service for a subscriber. This service has the lowest priority and the ISG Upgraded Service use case, when active, replaces the Base service for a subscriber.
Bandwidth Monitor no UM	This application includes the information needed to install a monitoring key, the different thresholds that is used to flag the subscriber. It installs any additional monitoring keys needed while reusing any monitoring keys that are already installed. In order to reuse a monitoring key the same monitoring key name is used in both BandwidthMonitor and UsageMonitoringKey objects.

Auto-Provision Quota	This use case is selected in situations when end customers should get a certain amount of quota when they come on the network, but without the need for additional provisioning. Typically, auto-provisioning quota done for cases like Fair Use, where a user would be automatically given a certain amount of quota for a month, and above which you could downgrade their service.
Prepaid Data Service	This use case is used for vouchers, and sets a valid time limit for the voucher holder. It uses USuM as a subscriber profile repository (SPR).

You are not required to use all of these use cases in service options, but they are available so you can expand your service offerings.

Basic Configuration Steps

Fundamentally, the steps for configuring your services are:

- Configure the required use cases one time.
- Set up at least one domain.
- Copy a use case template and create a service option.
- Adapt that service option to meet your business rules.
- Combine one or more service options and create a service.

Using the Screens

To perform your configuration tasks, you use two main screens:

- Service Options Screen
- Services Screen

Service Options Screen

Step 1 Open up a browser window and display the Cisco Policy Builder interface.

Step 2 Click Service tab > Services node > Service Options folder.

In the tree on the left, two service options use the Limit Max MAC Registrations service option.

One service option limits the subscriber to 10 devices, and is used in a basic level service.

Another version of the service option sets a 20 device limit that can be used to create a higher level service for a premium subscriber.

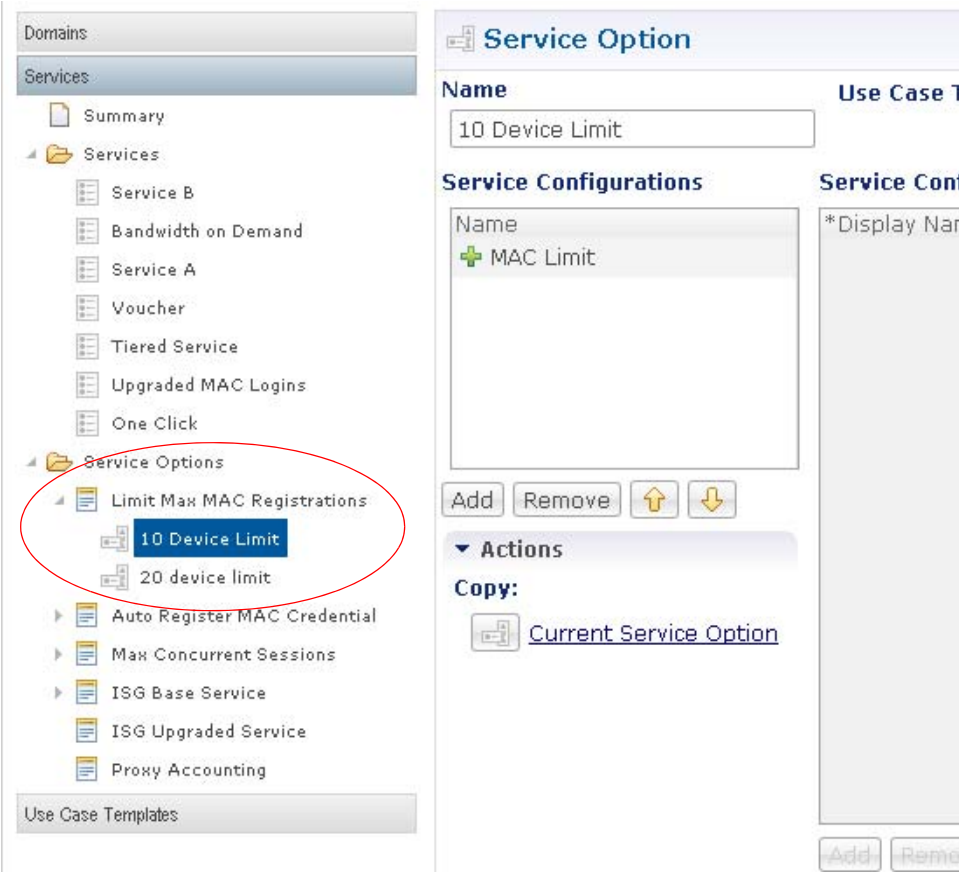

Step 3 In the Service option main screen, adapt the Value field to create the parameters you need for a subsequent service.

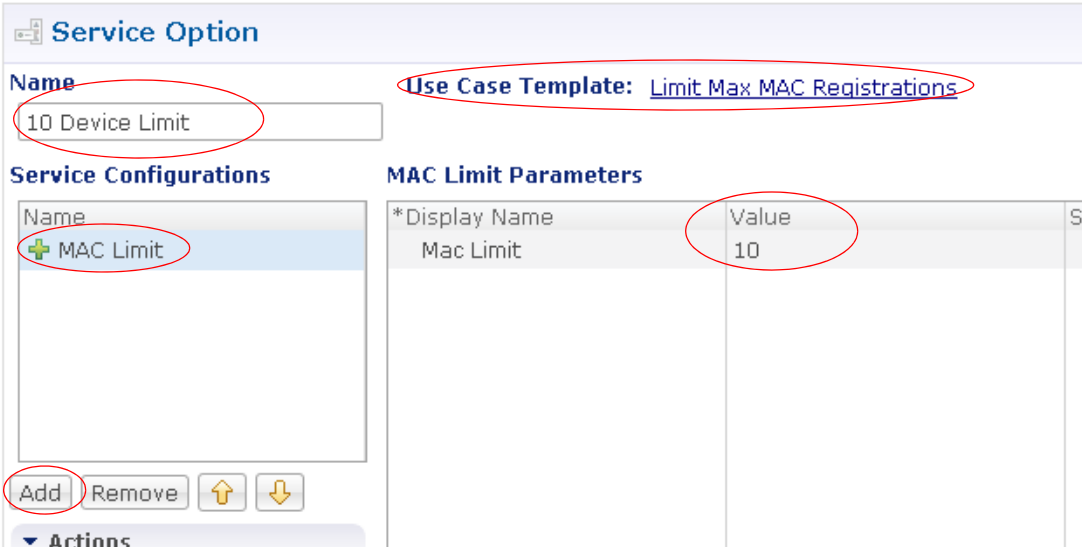

Our example shows a maximum MAC device limit of 10, which is a change from the default value of 5 in the original use case template.

Creating a New Service Option

The basic steps used to create a new service option, based on an available use case are these.

Step 1 Click Services tab > Services node > Service options folder > base service option.

Step 2 Select the Service Option link on the right.

Step 3 Click OK to display an empty Service Option screen.

Step 4 Select the Service Configuration on the left to display the parameters.

Step 5 Click in the table on the right and change the parameters.

If you do not see the service option you want in the tree, notify your Cisco technical agent. They can construct a use case template for your needs.

Services Screen

In the example below, the tree on the left shows a variety of services available to offer subscribers.

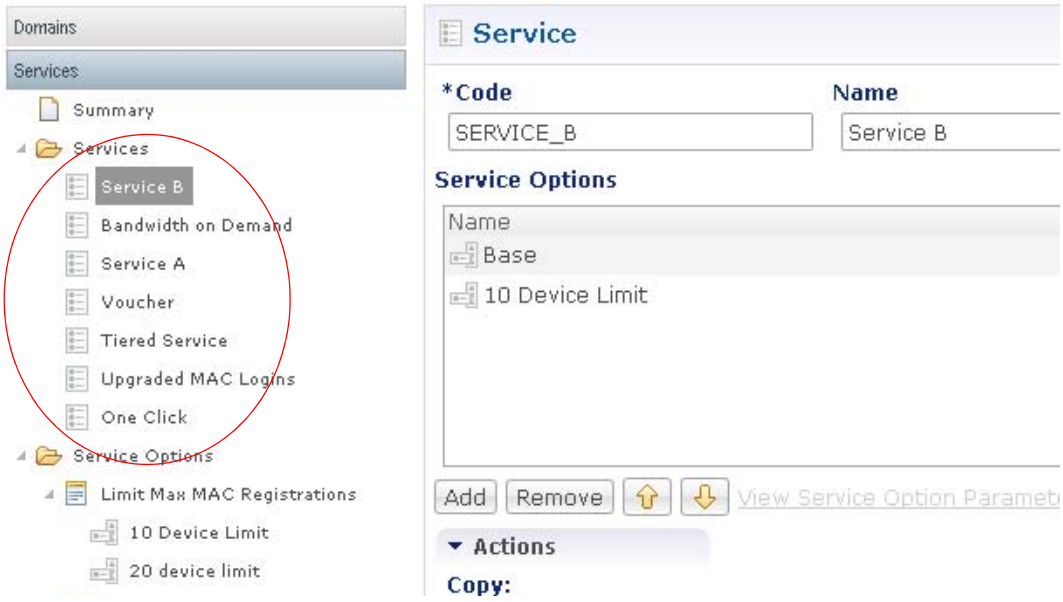

Service B, uses these service options:

- Base
- 10 Device Limit

The 10 Device Limit service option is used in Service B, but may also be used in other services as well. The same is true for the Base service option.

To augment or restrict Service B, use the Add and Remove buttons to add or remove the service options that comprise it. For example, if you want to put a limit on the number of concurrent sessions, add the service option 25 Session Limit.

Creating a New Service

To create a new service, select one or several previously created service options.

Step 1 Click Services tab > Services node > Service folder > Service link on the right.

Step 2 Fill in the Service screen, with a Code and Name.

Step 3 Use the Add and Remove buttons to add and remove service options.

If you do not see the service option you want for your service, create one with the Service Options folder.

Services and Service Options in Detail

Now that you understand the basic concepts for configuring your services,

- Perform the required configuration tasks described in Required Service Configuration.
- Create a Default domain as described in Domain Configurations.
- Create any elective services as described in Elective Service Configurations.

Required Service Configuration

Revised: July 10, 2015,

When you have completed the required configurations:

- Create a default domain as described in Domain Configurations.
- Examine Elective Service Configurations on and configure your own service options and services.
- Test your configuration. See Test the Configuration.

This chapter covers the following sections:

Before You Begin

Before you perform the required configurations tasks discussed here, be sure that you have these tasks completed:

- Install Cisco Policy Suite and have it process test traffic.
- Have the proper licenses enabled.
- Have access to the Cisco Policy Builder web page.
- Have access to the Control Center interface.
- Have your NAS or PEP installed and configured.
- Have the ISG or SCE installed and configured. ACLs and accounting lists must be set up on the ISG.

- Begin all procedures from the Cisco Policy Builder interface.
- Know what your basic service consists of.
- Plan what your additional services look like, perhaps prepaid or not.
- Decide on the names and levels of tiered service.
- Be familiar with any AV pairs you use.
- Know what speeds define your uploads and downloads.
- Specify what domains permit and provision.
- Know the numeric specifics of IP addresses, MAC maximums, time limits.
- Understand what limitations and provisions each service have.

Register Policy Enforcement Points

This configuration task defines a Policy Enforcement Point (PEP), either an ISG or a RADIUS pool or a Diameter.

You may have another type of PEP, but can easily follow the example provided here.

When a subscriber tries to access your network or server, the PEP describes the subscriber's attributes to other entities on the system. The PEP assigns the Policy Decision Point (PDP) the task of deciding whether or not to authorize the subscriber based on the description of the subscriber's attributes. From the architectural perspective, any network access device such as a PDSN, PCEF, SCE or ISG may act as the PEP. However, PEPs are not limited to such devices.

These configuration steps have you enter shared secrets, time out values, accounting list names, and IP addresses for your system.

This configuration task is required. Work through this task only one time per PEP device.

Steps

Note This procedure sets up either an ISG pool or a RADIUS pool. Your PEP have many of the same fields. The RADIUS pool uses a subset of the fields that the ISG has and instructions are not duplicated. If you use any other devices, see Other PEP Devices.

Step 1 Open a browser and log on to the Cisco Policy Builder interface.

Step 2 Click Reference Data tab > Policy Enforcement Points node.

Depending on the plug-ins you install, you'll have different policy enforcement points in the tree than shown here.

Note In the examples here, you use an ISG Pool folder, used to configure Cisco ISG enforcement points. The RADIUS Device Pools folder configures RADIUS for devices other than an ISG, if that device is not shown in the list.

Step 3 Click ISG Pools folder > ISG Pool link on the right.

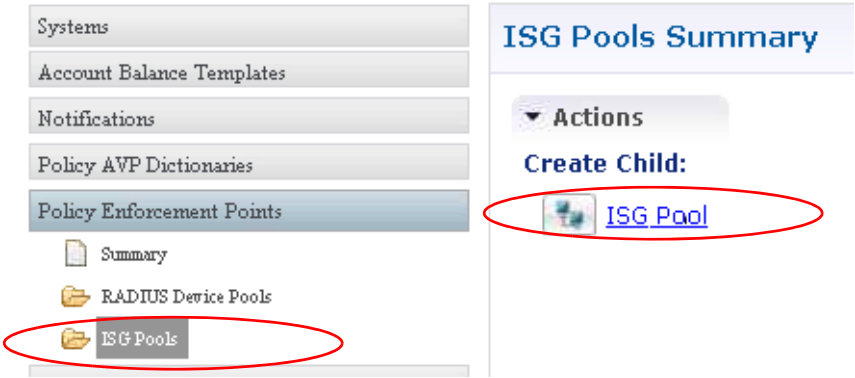

Note Most other network devices use a subset of the ISG Pool screen fields, the ones **not** outlined in red in the figure below. Other network devices do not use the fields Port Bundle Key Length, Change Service Rule, and Accounting List. Individual device screens are not discussed separately. For other network device pools see Other PEP Devices.

Step 4 Fill in the ISG Pool screen.

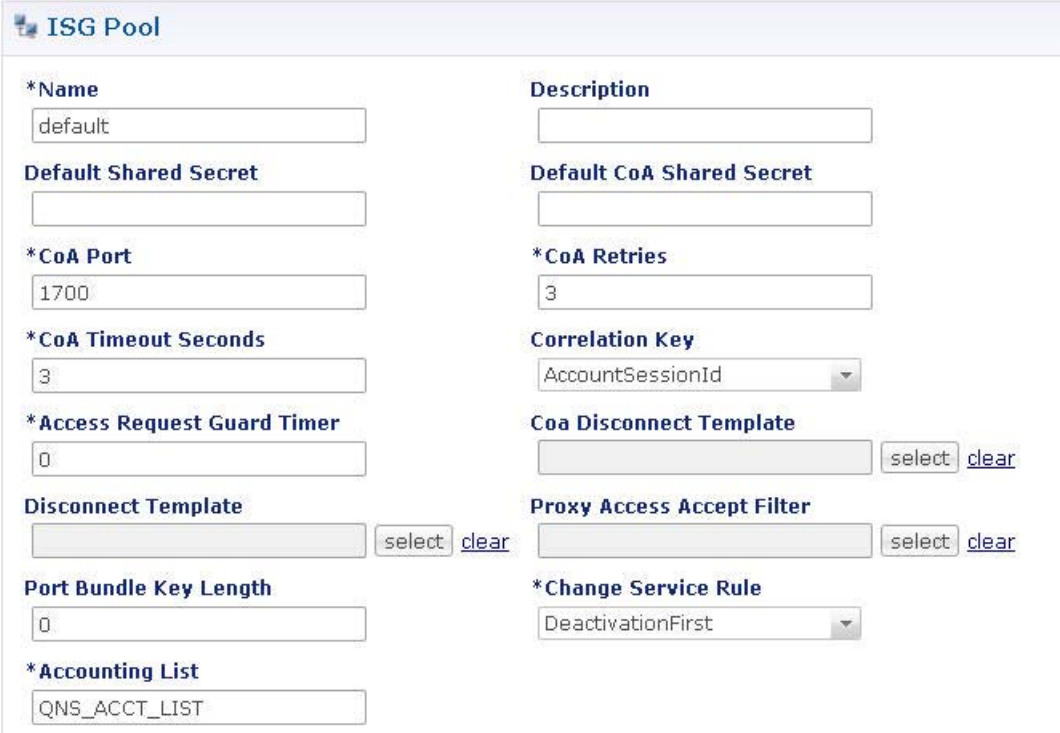

Name	A friendly name for the ISG pool. All the devices in the pool all share the information in the top of this screen. The rows in the Devices table at the bottom let you create exceptions to this general sharing.
Description	Helpful information about the device pool.
Default Shared Secret	The shared password or phraseword between Cisco Policy Builder and the ISG device.
Default CoA Shared Secret	The shared secret used between Cisco Policy Builder and the ISG devices unless a different one is specified in the Devices table below.
CoA Port	The hardware port on the ISG that listens for authentication tries. The default CoA port is 1700.
CoA Retries	The number of times that Cisco Policy Builder tries to authenticate with an ISG device in the list below.
CoA Timeout Seconds	The number of seconds that Cisco Policy Suite tries to authenticate with an ISG device in the list below.

Correlation Key	The key to correlate between the subscriber authentication request and the rest of the requests. Possible correlation keys are: • AccountSessionId • callingStationId • Tgpp2CorrelationId • UserId
Port Bundle Key Length	For PBHK (port bundle host key), the port bundle number: Includes a range of sequential port numbers, starting with a base port number. Is approximated by range of sequential port numbers=2 port bundle length. The port bundle length is an integer between 1 and 16. The default value on the ISG is 4. A RADIUS pool does not have this field.
Change Service Rule	This drop-down menu describes what happens when a subscriber changes services when logged in to the subscriber portal. • ActivationFirst The new services are activated and then the old services are deactivated. This ensures that there is no interruption of service. This method requires that you have unique ACLs set up on the ISG, and the subscriber is activating different services, not the same ones. • DeactivationFirst, the default. The old service is deactivated, and then the new service is started. This method ensures that stopping and starting services by means of their ACLs is cleanly done and without error. • SameCoA Starting and stopping of services occurs in the same message. Older versions of ISGs may not be able to accommodate this method. Make sure your ISG supports starting and stopping in the same message before making this selection. A RADIUS pool does not have this field.

Accounting List	The accounting list that CPS sends to the ISG and the ISG sends back. The accounting list must be defined on the ISG. The list defined on the ISG must list CPS as a server in order for CPS to manage sessions. A RADIUS pool does not have this field.
Devices Table	The Devices table specifically defines the IP addresses for the ISGs in the ISG pool and any loopback addresses. **Note** Loopback addresses must be set here. You cannot use the management address of the ISG. If loopback address are not set properly here, the system does not function.
IP Address	The IP address of the ISG device or a loopback address you are using.
Shared Secret	The shared password or phrase word between Cisco Policy Builder and the ISG device. This secret overrides the default set above. If no shared secret is specified here, the value in the Default CoA Shared Secret field above is used.
CoA Shared Secret	The shared password of phraseword between Cisco Policy Builder and the ISG device for purposes of authentication. This secret overrides the default set above. If no secret is specified here, the value in the Default CoA Shared Secret field above is used.
Loopback Addresses	Loopback addresses are set here. You cannot use the management address of the PEP. If loopback addresses are not set properly here, the system does not function.

Step 5 Go to Configure an Access Accept Template for the Subscriber.

Configure an Access Accept Template for the Subscriber

With this configuration task, you'll create a collection of services from your own service provider-specific templates.

These configuration steps have you specify the RADIUS AV pair template and select RADIUS AV pairs, which you can customize further.

This configuration task enables:

- RADIUS and Diameter authentication protocols and can be used with both of these.

- Wireless Internet Service Provider roaming (WISPr) protocol 1.0 and 2.0.
- Transparent Automatic Logon (TAL), which enables subscribers to maintain an always-on connection without the need to authenticate on each connect.

This configuration task is required. Work through this task only one time.

Steps

Step 1 Open a browser and log on to the Cisco Policy Builder interface.

Step 2 Click Reference Data tab > RADIUS Service Templates node > Summary > RADIUS Service Template Group link.

Step 3 Select the Service Provider Specific Templates node from the tree and open it.

Step 4 In the tree, open the node, Service Provider Specific Templates.

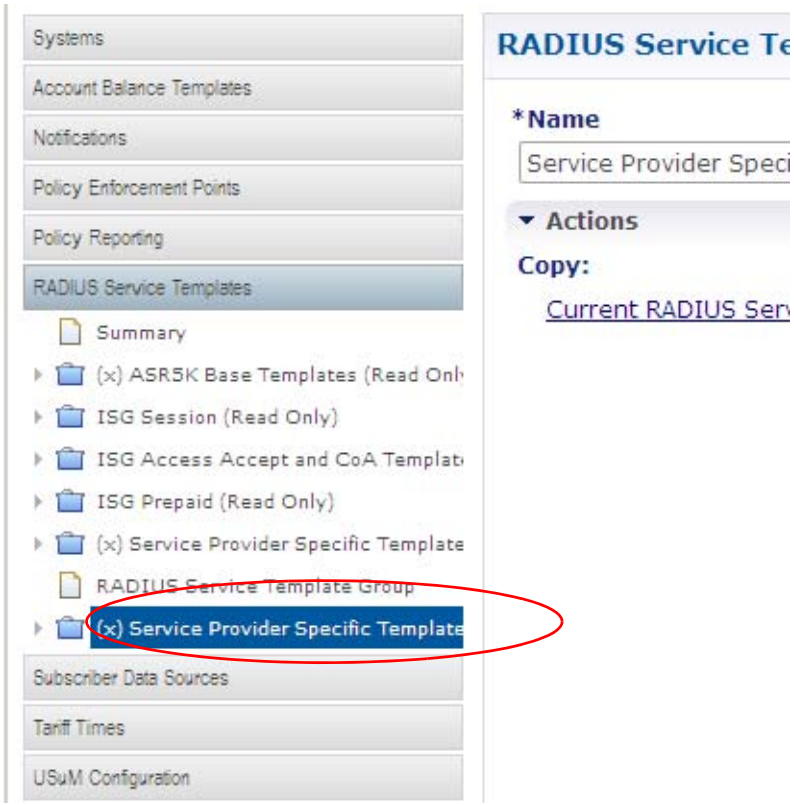

Step 5 Note that the screen displays all of the RADIUS templates specific to your enterprise, as developed by you or your department.

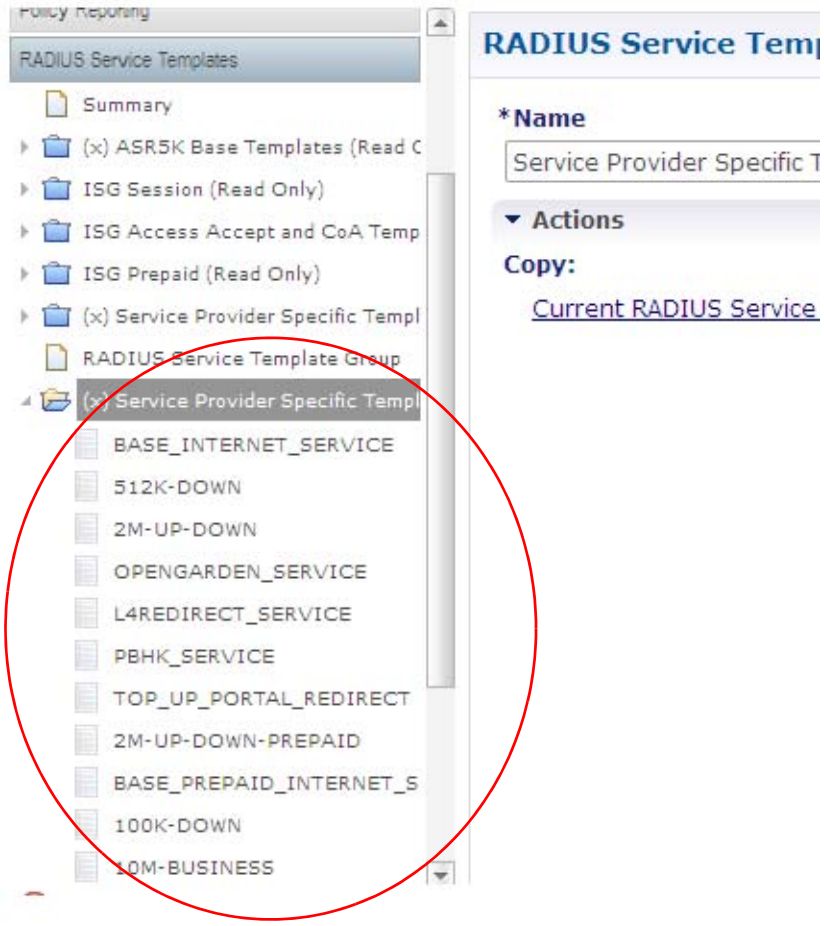

Step 6 Select the SP-ACCESS-ACCEPT node.

Step 7 Review the screen, making sure of it's content and making any needed changes.

You may need to change or add AV pairs to the RADIUS Service Template to customize your service group.

Adding AV pairs lets you fully customize the base template you uses later to define all services via a service option.

a. Click the link Show Available AV Pair Attributes To Add. (It changes to say Hide Available AV Pairs...)

From this list of AV pairs, a common item to add is Idle Timeout. The next step lets you customize the base idle timeout for your company.

b. Select <Radius> in the Vendor column and IDLE-TIMEOUT in the Attributes column. Doing so immediately populates the table above the Vendors / Attributes List.

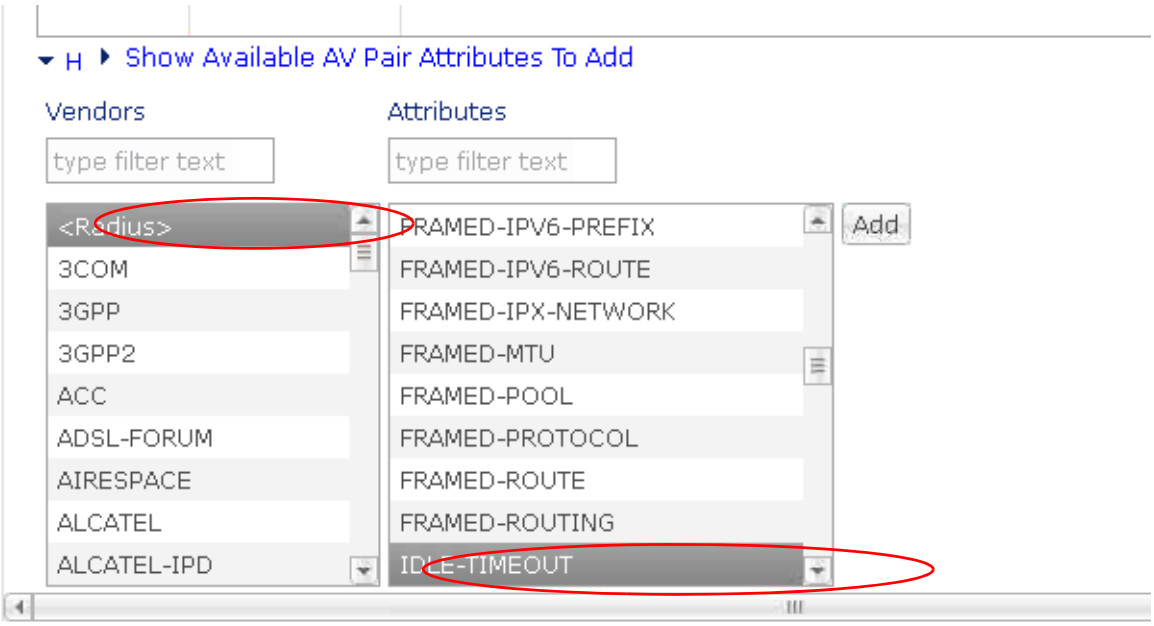

c. Click the Add button to the right of the Attributes list.

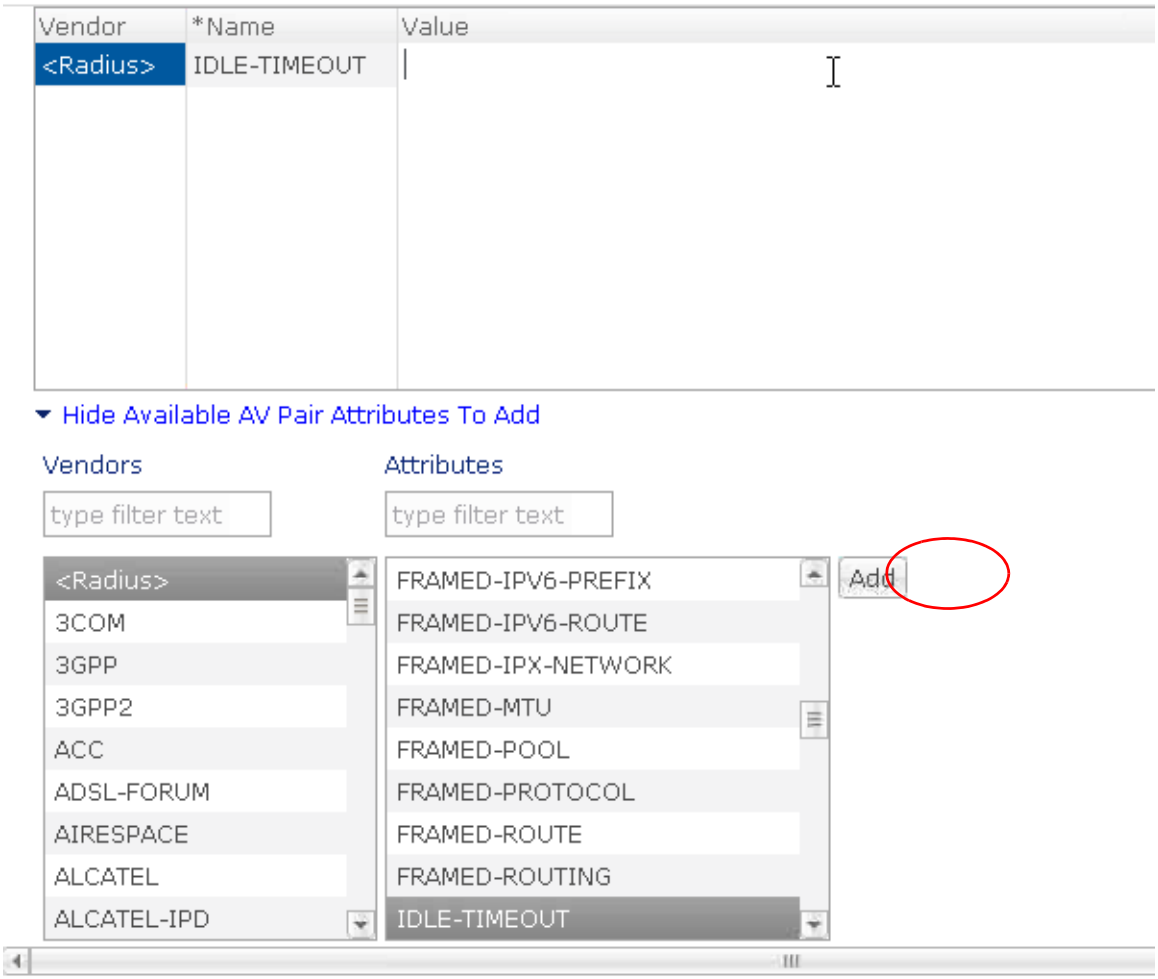

d. In the table above Vendors/Attributes list, click in the Value column and set your idle timeout as measured in seconds, to 600 (10 minutes).

e. Customize the session time out.

f. Click Show Available AV Pair Attributes again.

g. Select <Radius> in the Vendors column and SESSION-TIMEOUT in the Attributes column.

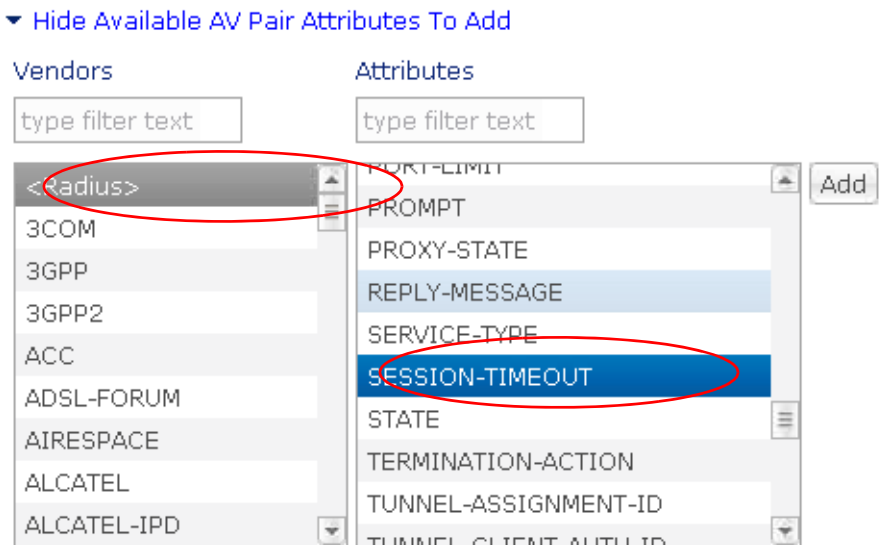

Note that the table above is populated immediately.

h. Click the Add button to the right.

i. Click in the Value column in the table above, and enter 3600 seconds (one hour).

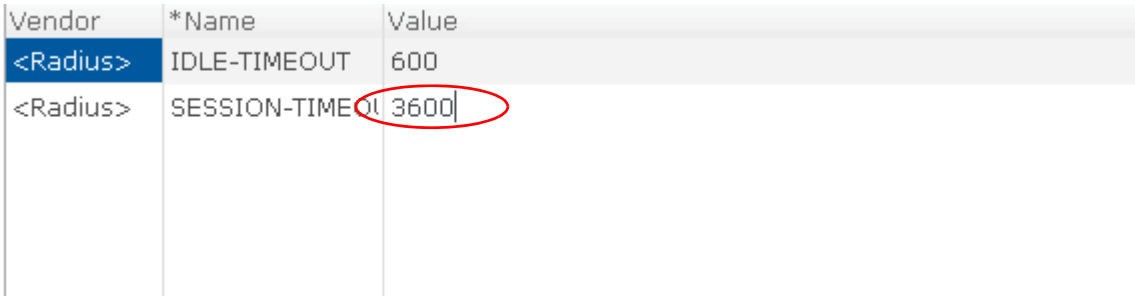

Step 8 Go to Configure Subscriber Authentication and Authorization.

Configure Subscriber Authentication and Authorization

Several services are provided for you. Take time to examine the 512K-DOWN and the 2M-UP-DOWN services for example. If you would like to develop your own services, work through this procedure.

Configuring the method for authenticating and authorizing subscribers is provided in this procedure, Defining a Base Prepaid Service.

If you want to more fully understand how you can provide service templates with other attributes, review the section Power Understanding — Refining the BASE_PREPAID_INTERNET_SERVICE Configuration Task.

Defining a Base Prepaid Service

This procedure shows how to define a base service and is a required configuration task. Start out by using a standard, predefined service from the example.

These configuration steps have you specify a RADIUS AV pair template and select RADIUS AV pairs which you can customize further.

Note AVP value substitution is only supported for String values and if the value substitution is IP address, it does not support AVP substitution.

The RADIUS Service Template named BASE_INTERNET_SERVICE is base template from which you builds all Internet services.

Steps

Step 1 Open a browser and log on to the Cisco Policy Builder interface.

Step 2 Select Reference Data tab > RADIUS service template node > Service Provider Specific Templates > RADIUS service template > BASE_PREPAID_SERVICE.

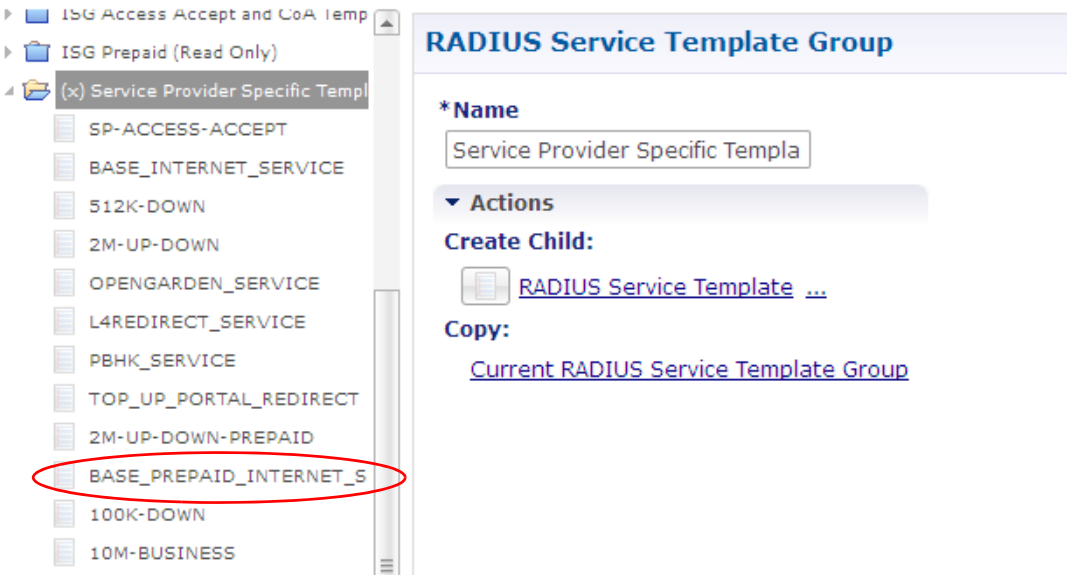

Step 3 Click the Select button next to the Base Template field.

Step 4 Open the item ISG Prepaid.

Step 5 Select the PREPAID_ACCESS template object.

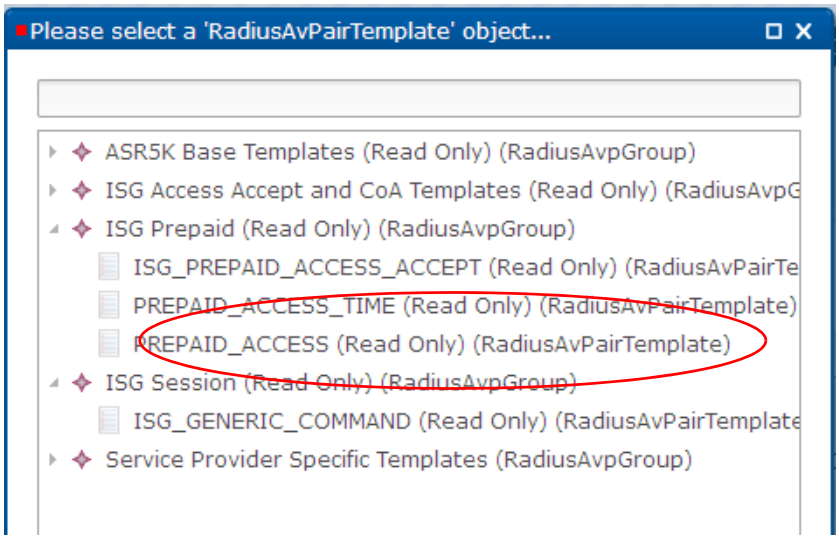

This is what is used for your enterprise's standard Internet service.

Step 6 Click OK.

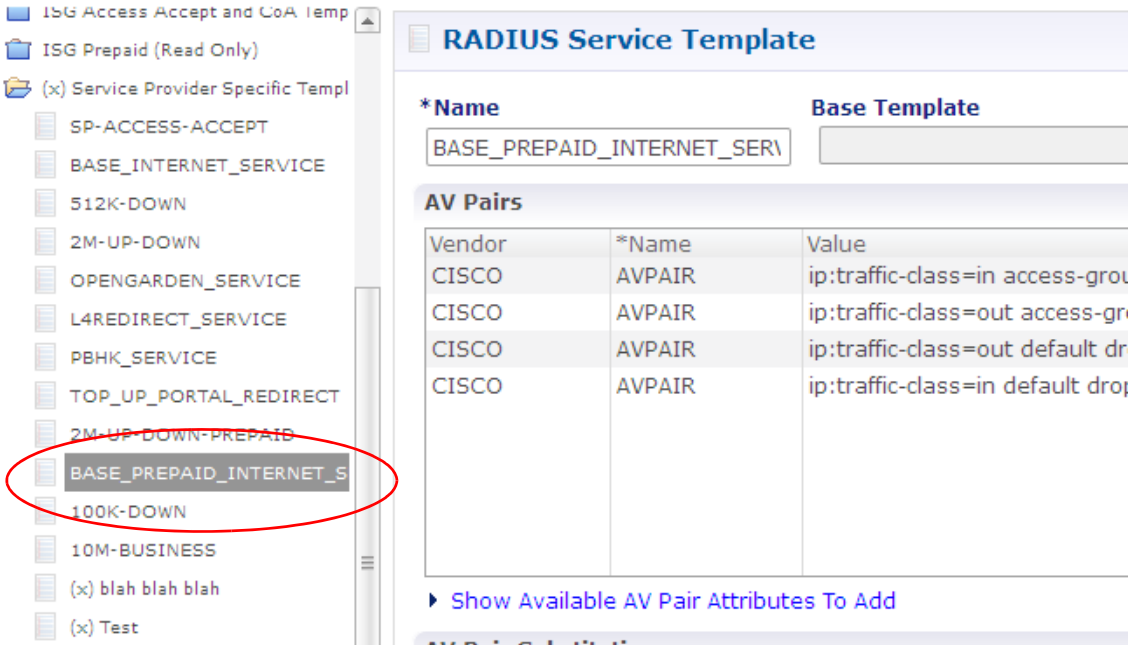

The proper AV Pairs are presented for you.

Step 7 Make any changes based on the advice of your technical representative.

Step 8 Go to Power Understanding — Refining the BASE_PREPAID_INTERNET_SERVICE Configuration Task.

Power Understanding — Refining the BASE_PREPAID_INTERNET_SERVICE Configuration Task

This section is not required to configure your system, but working through it provides additional understanding.

To compare and contrast how different service templates appear, this procedure has you define a second service template that specifies a 512K download speed.

Steps

Step 1 Click the Reference Data tab > RADIUS Service Template > Service Provider Specific Templates.

Step 2 Select the object Service Provider Specific Templates and open it.

Step 3 In the tree, click the 512K-DOWN item and make sure that screen reflects the upload and download speeds for your 512K up down service.

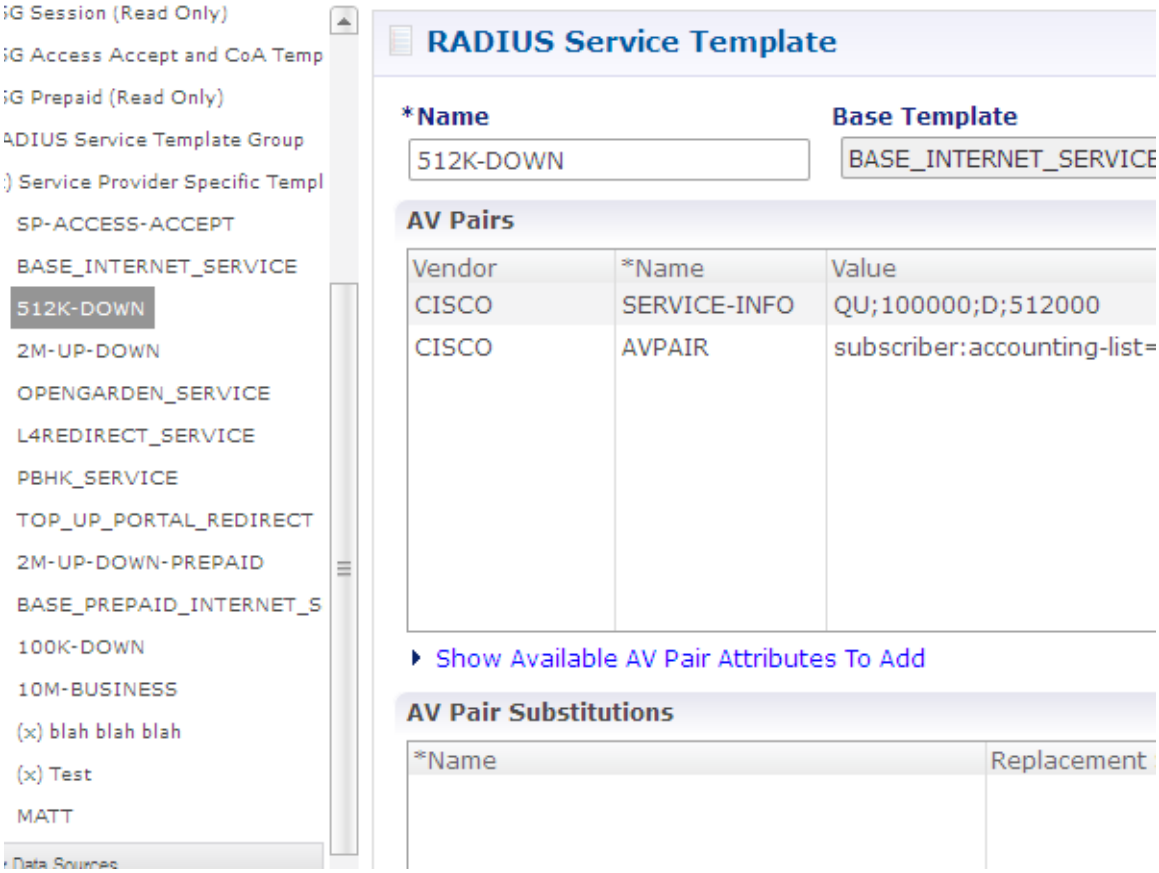

If you would like to see this configuration completed with service options and services, continue to work through the sections in Power Understanding— Service Options and Services and use the service templates you just defined.

If you are done working on the required configuration tasks, go to Domain Configurations and then Elective Service Configurations.

Power Understanding— Service Options and Services

This section is not required to configure your system, but working through it provides additional understanding of how service templates are used.

- Configuring a Service Option
- Configuring a Service

Configuring a Service Option

This example creates a service option that uses the RADIUS templates you just created. You probably uses this service option in creating your own services.

Recall that a service option is based on a use case template previously defined and provided for you by Cisco.

You can mix and match many service options to create a variety of subscriber services.

Step 1 Click Services tab > Services node > Service Options node > Service Options folder > ISG Base Service > Service Option link.

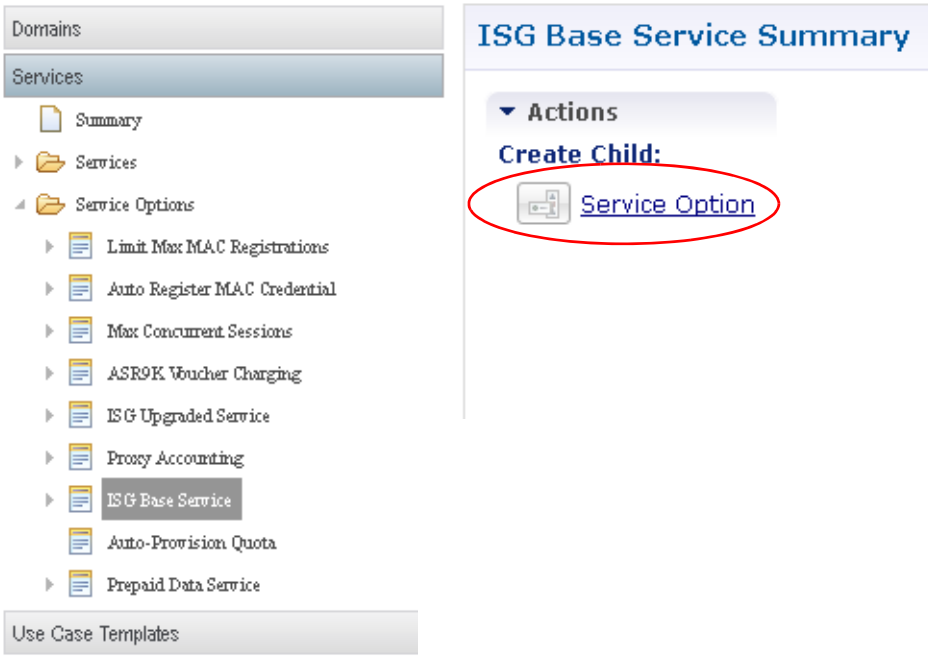

Note Note that current service options are listed in the Service Configurations list on the left. You can select one of them and examine the display name and the value if you like.

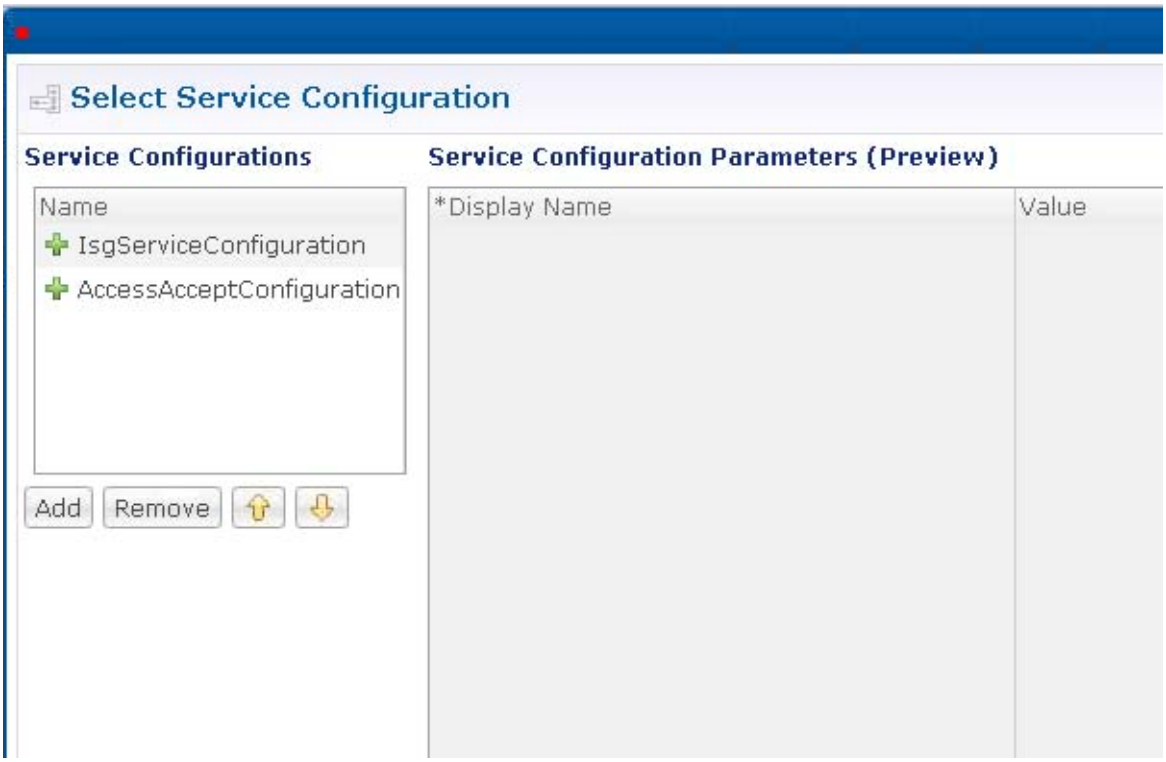

Step 2 Click OK to display the Service Option page.

Step 3 In the Name column, double click and enter the name Basic, which is used as your basic service.

Step 4 Select ISGServiceConfiguration from the Service Configuration list.

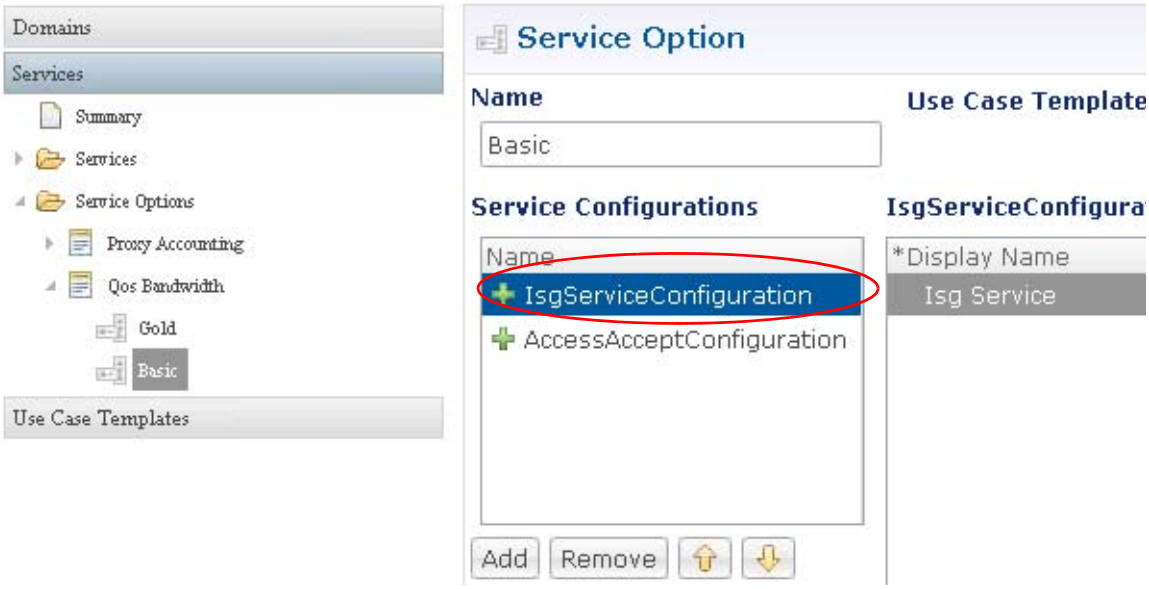

Step 5 In the Value column, click the example drop-down button .

The RADIUS AV Pair templates list appears.

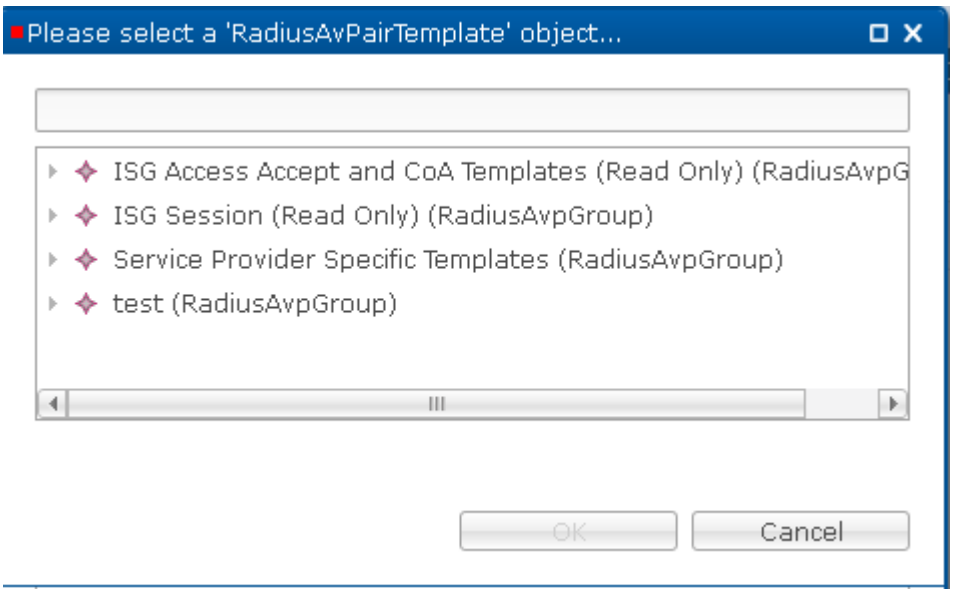

Step 6 In this object list, open the Service Provide Specific Templates you configured earlier and select the 512K-DOWN object. (See Power Understanding — Refining the BASE_PREPAID_INTERNET_SERVICE Configuration Task.

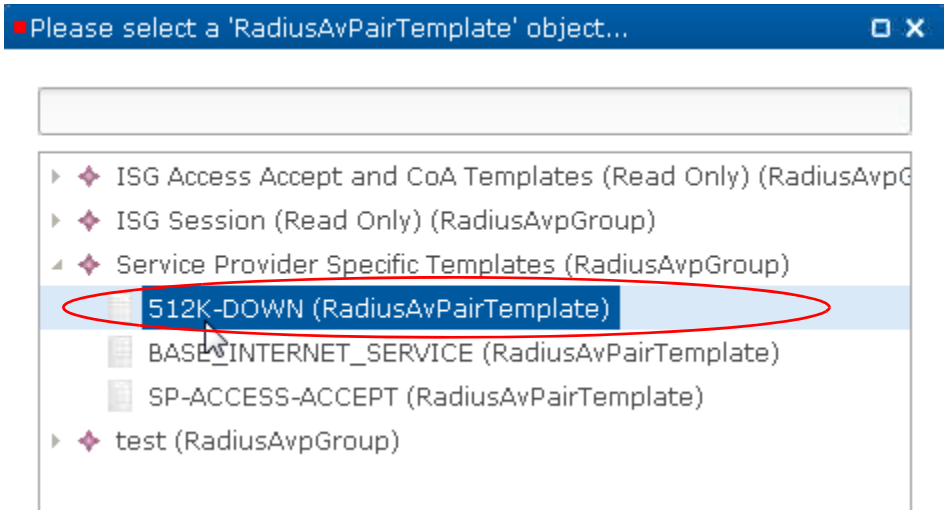

Step 7 On the Service Option screen, select the AccessAcceptConfiguration to define it.

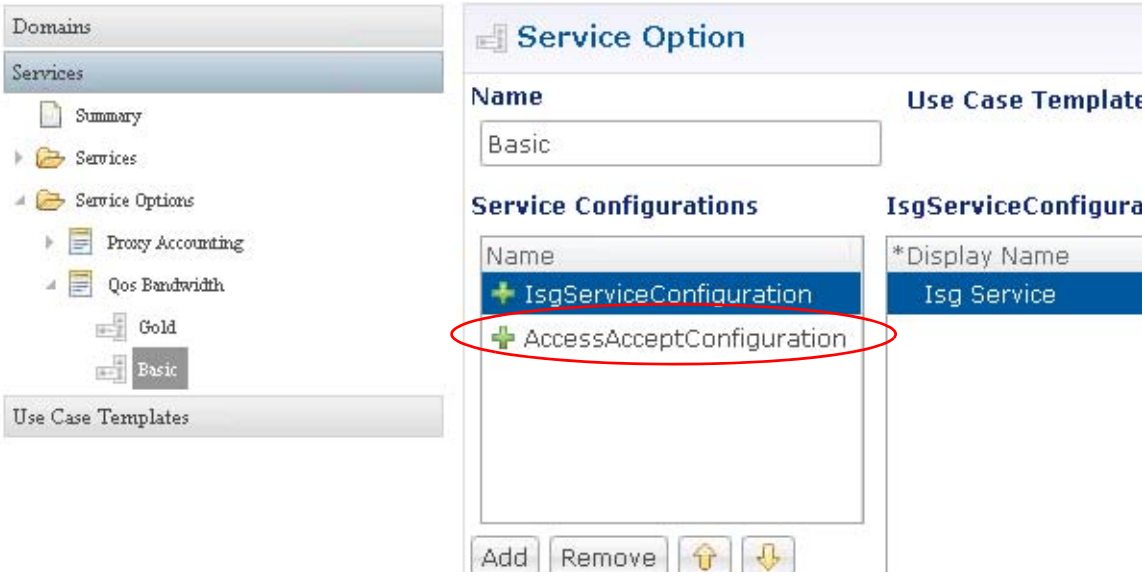

Step 8 In the Value column, click the example drop-down button. 

Step 9 In the RADIUS AV Pair template list, open the Service Provider Specific Template and choose the SP-ACCESS-ACCEPT template you defined previously in Configure an Access Accept Template for the Subscriber.

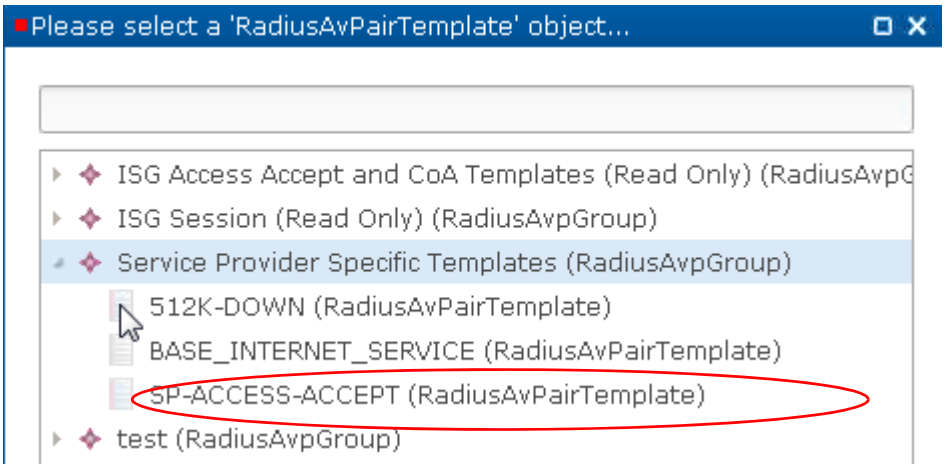

Step 10 Go to Configuring a Service and use the components you have just defined.

Configuring a Service

These steps have you create a new service that subscribers can subscribe to. This service uses the service option created in Configuring a Service Option.

Step 1 Click Service tab > Services node > Services folder > Service link to display an empty Services screen.

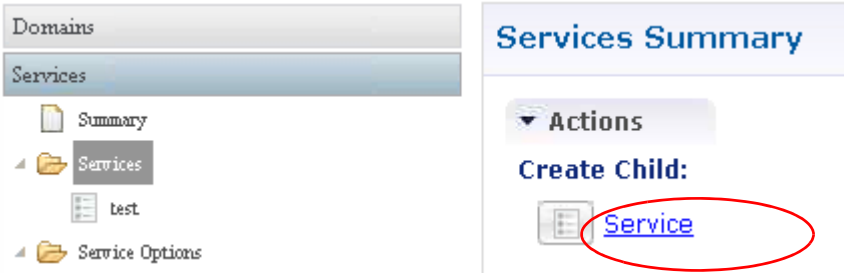

Step 2 Define the Code field as Service_A.

Step 3 Define the Name field as Service_A.

Step 4 Click the Add button to add service options to the service.

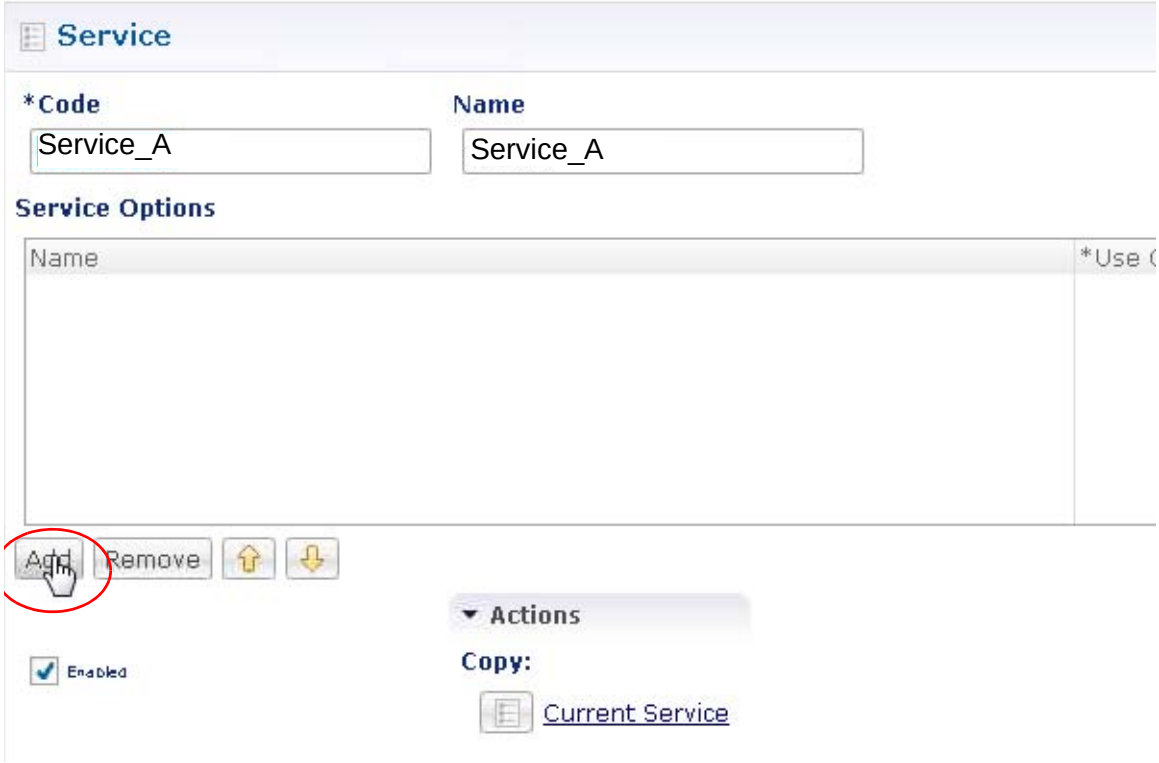

Step 5 Click Base to assign the basic service option to the service SERVICE_A.

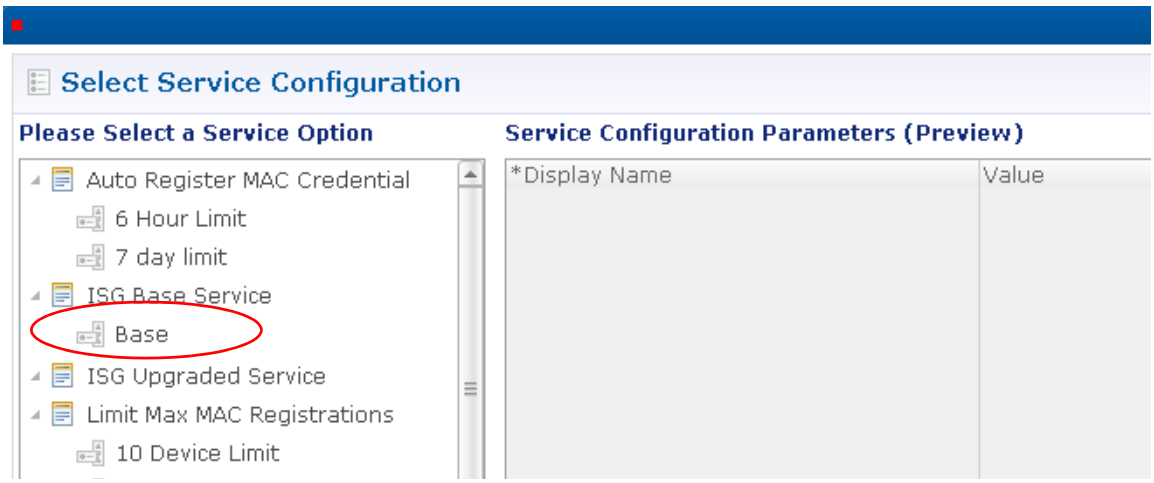

Now the service SERVICE_A has the base bandwidth on it, that is the basic service level for your enterprise. Subscribers are able to subscribe to Service_A.

Your subscribers must access their service by logging in to a domain.

Step 6 Next, create a domain with the procedures at Domain Configurations.

Subnet based RADIUS Client

For network architectures with multiple RADIUS clients, it is difficult to configure same shared secret for multiple client devices.To overcome this difficulty, these devices are configured with a subnet rather than just a single IP. The same shared secret would be applicable for all IP Addresses lying inside the Subnet. Cisco Policy Suite (CPS) defines a RADIUS device with a single IP and shared secret.

CPS provides the capability to enter the Radius Client IP Address in CIDR (Classless Inter Domain Routing) notation instead of a single IP address. The same shared secret is used for all devices with IP Addresses lying within the IP range specified by the subnet defined. All Policy Enforcement Points such as WLC, ISG, ASR5K, ASR9K, MAG, IWAG, etc. are provided with the ability to define Subnet based RADIUS clients sharing the same secret.

When a RADIUS request is received, it is checked if the NAS-IP Address lies within the range of the IP pool defined in every Policy Enforcement Point. If it lies within an IP range, the Policy Enforcement Point is selected and verified if the secret matches with the RADIUS secret key in the incoming request. If the IP overrides with multiple IP ranges, suitable error messages are displayed.

To implement the Subnet based RADIUS client mechanism, the Policy Enforcement Point is configured in the Policy Builder.

Configuring Subnet based RADIUS Client

To configure the Policy Enforcement Point, perform the following steps:

Step 1 Login to Policy Builder.

Step 2 Click **Reference Data > Policy Enforcement Point > ISG Pools.**

Step 3 In the ISG Pools Summary Window, click **ISG Pool** to create a new ISG pool.

Step 4 Enter the values for the required fields according to your requirement.For example, see the figure given below:

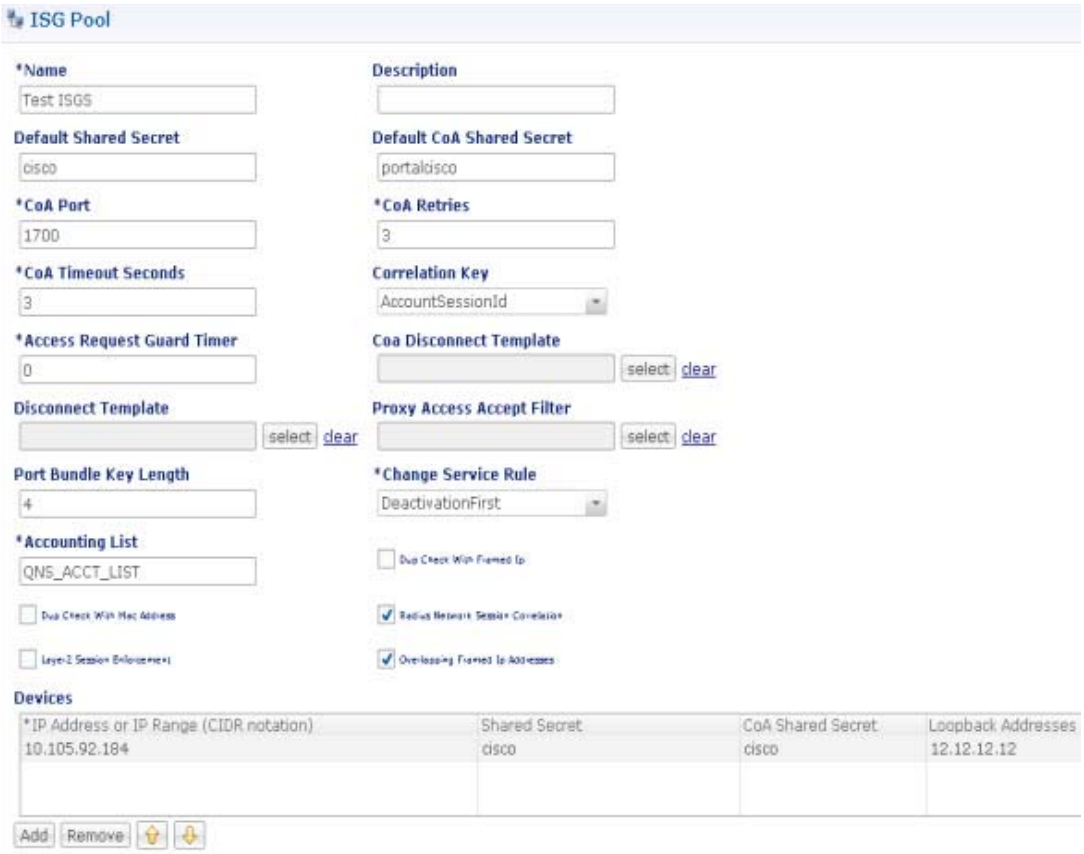

Refer to Register Policy Enforcement Points for additional details.

Step 5 In the Devices section, enter the Subnet or IP Range (CIDR notation). To add an IP Range, click **Add**. By default the IP Range is 0.0.0.0. Edit the IP Range according to your requirement in the CIDR notation by clicking on the default value as shown in the figure given below:

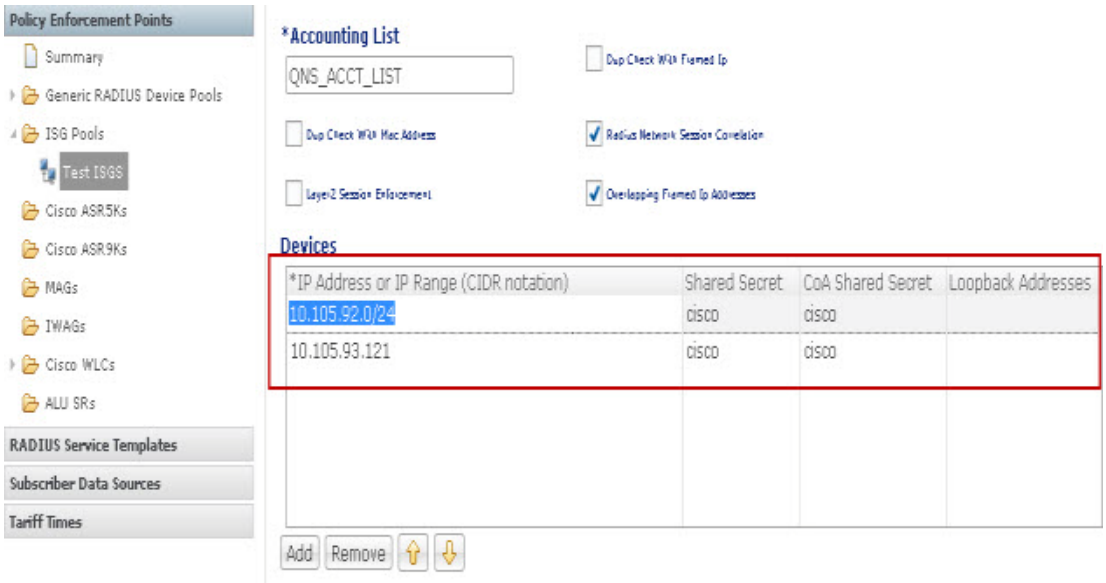

Step 6 Enter the value for Shared Secret and CoA Shared Secret by selecting the blank row of the column respectively. For example, refer to the figure shown above.

If the IP Range in one device definition overrides with any other IP Range or any IP Address in the same or other device definitions, the Policy Builder performs a validation check and displays suitable error messages against the Policy Enforcement Point, which has an overlapping IP range. Refer to the figure given below showing error messages due to IP Range overlap.

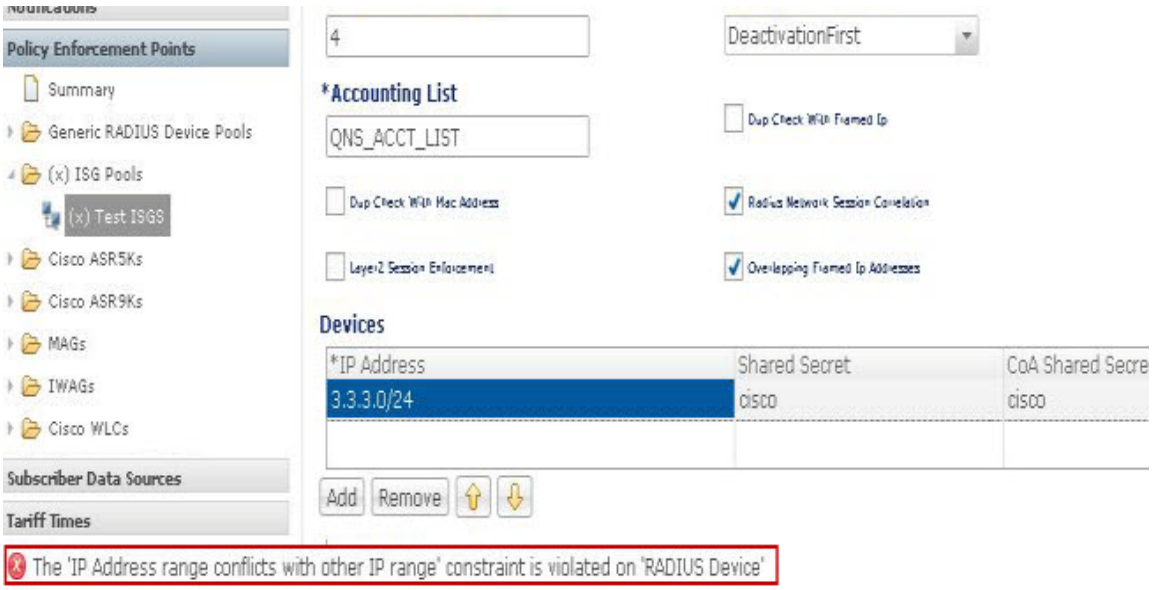

Configuration and Restrictions

- Configuration of Loopback Address in CIDR notation is not supported.

- If a Loopback Address is configured, the corresponding IP Address column should have a single IP Address and not a range of IP Address. This leads to an incorrect configuration.

SPR Cleanup for Inactive Subscribers

When a subscriber is found to be idle for a period of time due to expiration of services or insufficient account balances and so on, CPS marks the subscriber as inactive and removes it from the database.

Cisco Policy Suite (CPS) provides an automated mechanism to cleanup inactive subscribers from the database eliminating the manual process.

CPS provides the **SubscriberInactivity** event to handle the mechanism.

The possible event states for the **SubscriberInactivity** are given in the following table:

Event State	Description
READY	When the scheduler finds the subscriber to be inactive, it puts the Subscriber Inactivity event into the READY state for processing by the events processor.
SUB_INACTIVE	The events processor takes the event in the READY state and passes the inactivity event to the policy engine for marking the subscriber as inactive. The state of the event changes to SUB_INACTIVE in the events collection.
SUB_ACTIVE_CHECK	When there is an update to the SPR profile (service/balance/credentials) change through the unified API, the unified API notifies the scheduled events NDM by pushing a message in the policy Engine. The EventsNDM processes this message and checks if there is an event present in the events database for the subscriber, if the event is present and is either in READY or SUB_INACTIVE, the event state is changed to SUB_ACTIVE_CHECK. If there is no event present the EventsNDM discards the message.
SUB_ACTIVE	If the event is currently in the SUB_ACTIVE_CHECK, the event processor checks if the subscriber is active again. If the subscriber is not active, it marks the event in SUB_INACTIVE state. If the subscriber is in the active state the even processor marks the event to SUB_ACTIVE state. The SUB_ACTIVE state is reached only when the subscriber satisfies the condition criteria.
SUB_DELETED	When the event in the events database stays in the SUB_INACTIVE state for a period of time greater than the number of days configured against the event, the event is marked to SUB_DELETED state and the timeToLive field is updated to current time.

Subscriber Inactivity Event Configuration

The Scheduled Events Plugin Configuration is configured in the Policy Builder to implement the SubscriberInactivityEvent. New fields **Service** and **Max Number of Days** are added to support the subscriber inactivity event.

Prerequisite

To enable the scheduled events framework this feature has to be enabled in the feature set of Policy Server and Policy Builder. The following packages when added to the respective servers, deploys the functionality of scheduledEvents during a session:

- In the Policy Builder - /etc/broadhop/pb/features, the **com.broadhop.client.feature.scheduledevents** package is added
- In the Policy Server - /etc/broadhop/pcrf/features, the **com.broadhop.scheduledevents.service.feature** package is added.

Steps

To implement the Subscriber Inactivity event, perform the following steps:

Step 1 Login to Policy Builder.

Step 2 Click **Reference Data > System > Plugin Configuration > Scheduled Events Plugin Configuration**.

Step 3 In the **Scheduled Event Configuration** page, enter appropriate values for the fields provided. For example, see the figure given below:

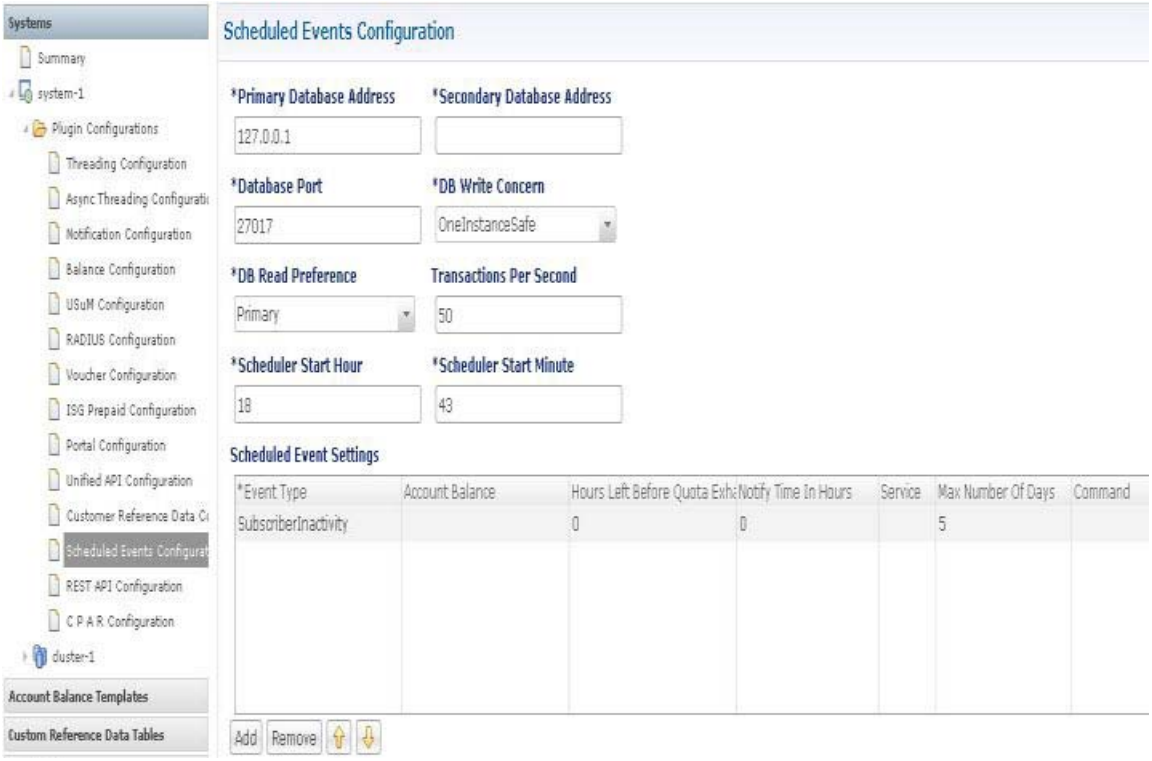

The Schedule Start Hour and Scheduled Start Minute are new fields that are added which defines the time at which the SubscriberInactivity event is triggered.

- The value for **Scheduled Start Hour** should be in the range of 0 to 23 (24 hours format).
- The value for **Scheduled Start Minute** should be in the range 0 to 59.

Step 4 In the **Scheduled Event Settings** section enter the values for the columns shown in the figure.

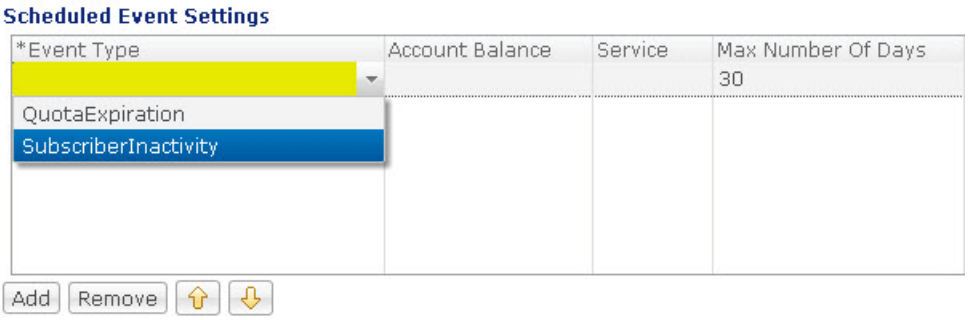

Step 5 Assign the Event Type to **SubscriberInactivity** by selecting from the drop-down list.

Parameter	Description
Event Type	Specifies the event type to be triggered. SubscriberInactivity should be configured to look for inactive subscribers
Account Balance	Specifies only those subscribers whose account balance is specified in the configuration, other subscribers are ignored for processing.
Service	Specifies only those subscribers who has the configured service associated, other subscribers are ignored for processing. Therefore the Account Balance and Service act to filter out subscribers having the configured balance and service.If these columns are not specified, the event looks up for all the subscribers.
Max Number of Days	Specifies the duration in days to retain the subscriber in the inactive state. If the status of a subscriber remains inactive for longer than the configured maximum number of days, the subscriber is automatically deleted from the database.

The figure shows the **Inactive** status of the subscriber in the Control Center.

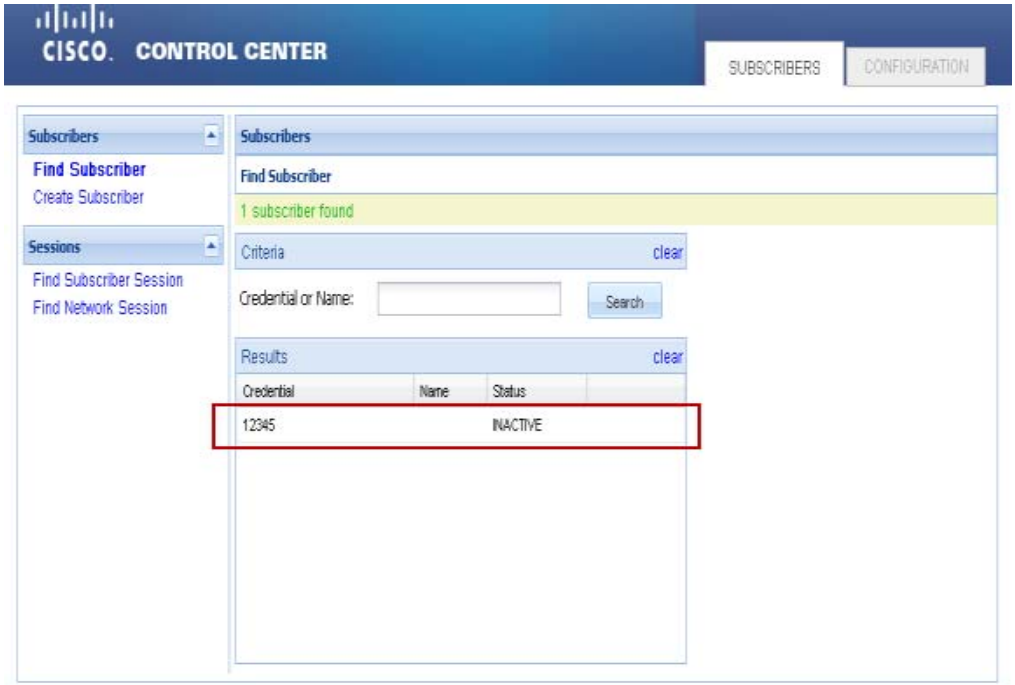

REST Technology between CPAR and CPS with JSON Interface

Cisco Policy Suite (CPS) provides support to handle multiple Cisco Prime Access Registrar (CPAR) sessions over the REST interface. The Representational state transfer (REST) interface provides the endpoints for both the subscriber and the session having the capability to perform create, read, update and delete operations. CPS exposes the REST endpoints to perform CRUD operations on the session and the subscriber database as requested by CPAR server.

The session and subscriber databases are configurable in the Policy Builder.

We use the HTTP methods to distinguish whether the request is for CREATE, READ, UPDATE or DELETE. The following table shows a mapping of the HTTP methods to the type of request and the operation received.

HTTP Method	Operation
POST	Create
GET	Read or Get
PUT	Update
DELETE	Delete

The key to the session database is **session_id** and the key for subscriber database is **sub_id**. These values are the necessary input parameter in the REST request. Also, if the request contains the JSON payload, it should contain the key attribute (**session_id/sub_id**) with its value.

For example,

- HTTP GET for <IP>:8080/qps/rest/cpar/subscriber?sub_id=9811122222 will fetch the details of the subscriber 9811122222 from the database.
- HTTP POST for <IP>:8080/qps/rest/cpar/subscriber?sub_id=9811122222 with a JSON payload will insert the details to the subscriber with the subscriber_id 9811122222.
- HTTP DELETE for <IP>:8080/qps/rest/cpar/subscriber?sub_id=9811122222 will delete the subscriber with the subscriber_id 9811122222 from the database.
- HTTP PUT for <IP>:8080/qps/rest/cpar/subscriber?sub_id=9811122222 with a JSON payload will update the subscriber with subscriber_id 9811122222 in DB with the new details provided.

If the operation succeeds a success response is sent back in the JSON format. To ensure duplicate keys in the database validation is performed and suitable response is displayed.

The Error Response messages that might be displayed during a session are shown in the following table:

Error Type	Error Message
GENERIC_ERROR	Request failed
DUPLICATE_KEY_ERROR	Duplicate value for Unique Data Constraint: sub_id
DATA_REQUIRED_ERROR	Required Data: sub_id
DELETE_ERROR	The object sub_id: <sub_id> was not found
SUCCESS_RESPONSE	Request completed successfully

CPAR and REST API Configuration

To perform the CPAR configuration and REST API configuration, perform the following steps:

Step 1 Login to Policy Builder

Step 2 Click **Reference Data > System > Plugin Configuration > CPAR Configuration**

Step 3 Enter the Subscriber and Session details as shown in the figure given below.

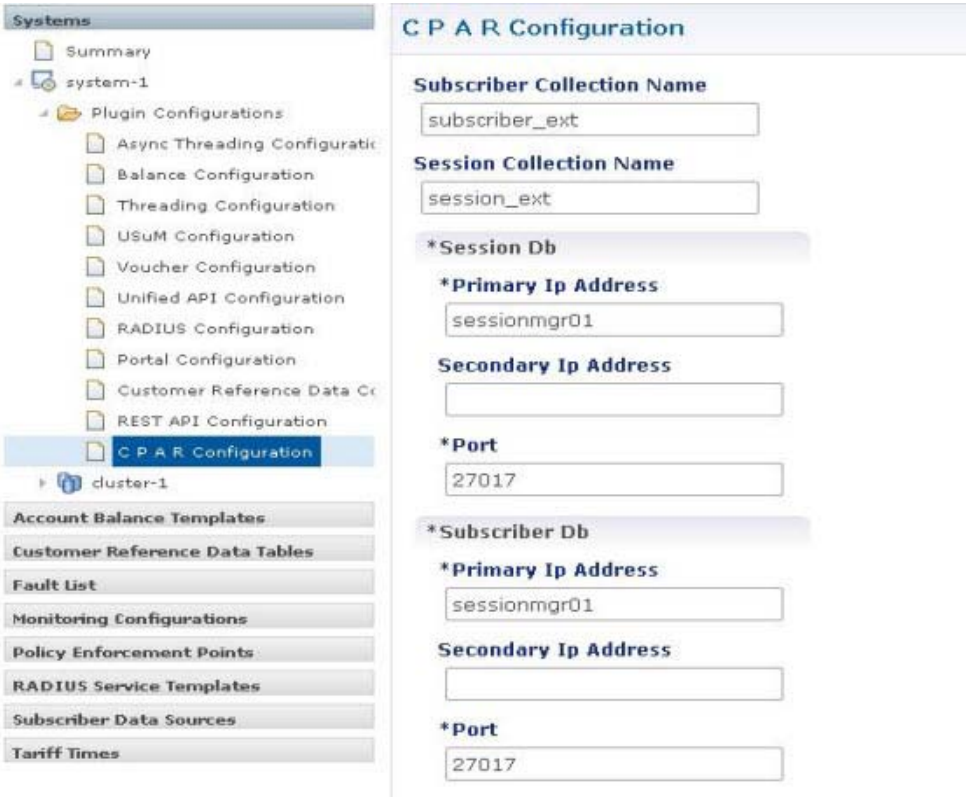

Step 4 To configure REST API configuration, click **System > Plugin Configuration > REST API Configuration** and check the **Json Payload** checkbox to enable JSON payload.as shown in the figure given below:

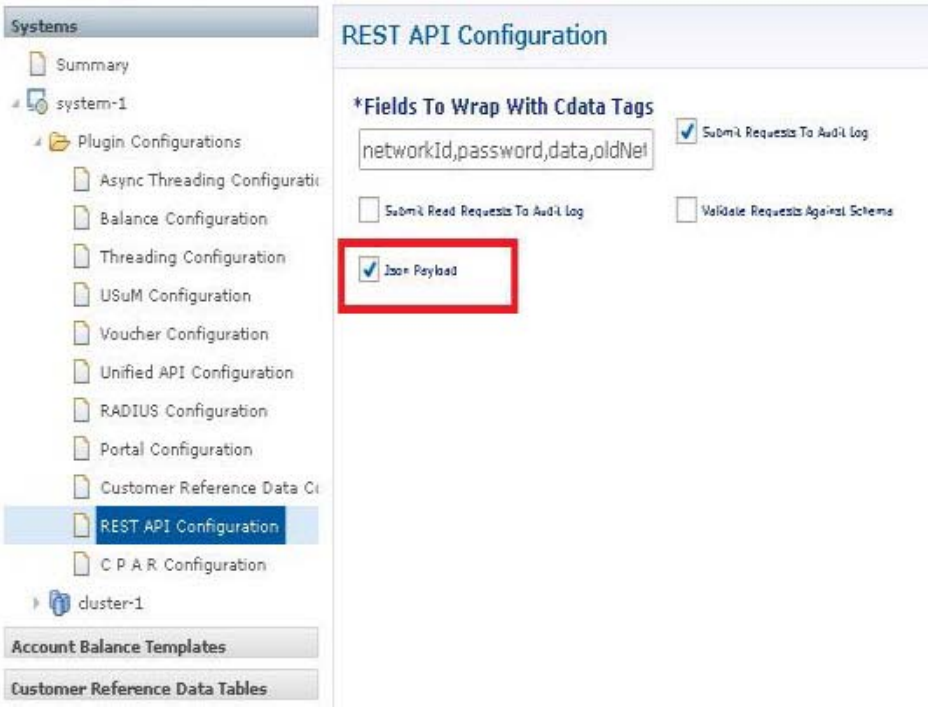

Other PEP Devices

These devices are also available and have many of the fields previously discussed. Although they appear on the interface, please contact your Cisco technical representative before using them.

- MAGs Device
- Cisco WLC
- Cisco ASR5000
- Cisco ASR9000
- ALU SRs
- SCE Device Pool

Domain Configurations

Revised: July 10, 2015

This chapter shows how to configure domains with Cisco Policy Builder.

Open a browser and log on to the Cisco Policy Builder interface to perform domain configuration tasks.

A domain is a way to partition subscribers based on an attribute of the triggering session creation message with a set of authorization, auto provisioning, secondary profile load, locations, and various advanced domain options.

This list provides you with a detailed explanation of these concepts:

- Authorization section of the message—defines the method that can be used for authorizing session creation. The following methods are appropriate for mobile use cases:
  - USuM Authorization—performs a lookup in the Cisco Unified SuM SPR for the subscriber's profile. This method requires a username (network identifier) and optionally a password. However, a password is not common in mobile scenarios. Common username fields are MSISDN or IMSI.
  - Allow All Users—performs no validation and always allows session creation. This method is commonly used if a secondary subscriber data source is the sole source of subscriber data or if auto provisioning is defined.
- Provisioning section—defines whether auto provisioning of subscribers within the SPR should occur. This method is generally used in scenarios where the system is configured to "auto-learn" subscribers and assign a default service profile. For example, the Authorization section would be configured with "Allow All Users". The "Provisioning" section would be configured to provision users with a key of the MSISDN.
- Additional Profile Data section—defines whether an LDAP search or another search method (Sh) should be utilized to retrieve profile data that is not stored in the local SPR. This is often used to integrate an external profile repository.
- Locations section—defines the rules used to guide the requests to a non-default domain. The supported rules relevant to mobile sessions include these:
  - Framed IP Mask, in CIDR format
  - Realm Mask, defined as realm/abc.madeup.com
- Advanced Rules—determines if unknown subscribers can come into the system and defines the unknown service. This is often used if subscribers self-provision and so are initially unknown. An example of this scenario is:
  - Cisco Unified SuM authorization, IMSI Authorization

- No Auto Provisioning
- Unknown Subscribers Allowed—Unknown Package - Portal Self-service

This chapter covers the following sections:

Creating a Domain

This section contains steps for configuring these types of domains.

- Defining a Default Domain
- Defining a Domain to Limit Framed IPs
- Defining a Registered Users Domain

Defining a Default Domain

This procedure shows you how to create a default domain for all subscribers in the system. At any time, there must be one domain defined in the system and that domain is assigned to a session if the location rules do not resolve to another domain.

This domain specifies that when a request is received, the Cisco Unified SuM SPR profile is loaded using the IMSI. No provisioning is triggered, and no additional profile data is retrieved. All advanced options are set to default.

Steps

Step 1 Click Services tab > Domains node > Domain link.

General Subtab

Step 1 Display the Domain screen > General subtab.

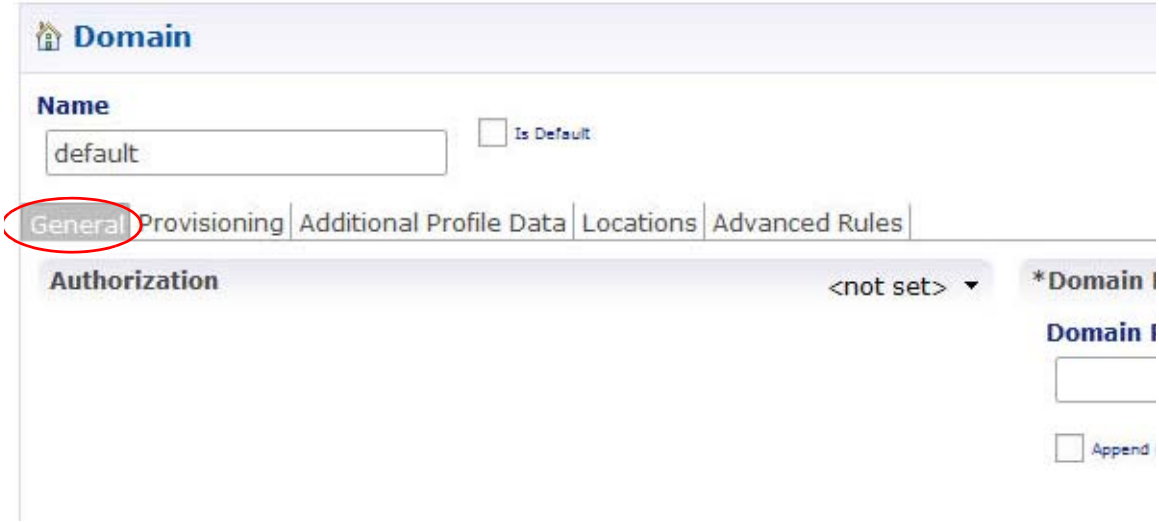

Step 2 Set the Name field to Default.

Step 3 For the Default domain, make sure the **Is Default** check box is selected. Doing so means that any subscriber unknown to the system accesses the network through the Default domain, and use its authorization, provisions, and location criteria.

Step 4 Set the Authorization mode.

a. Open the drop-down menu in the Authorization area.

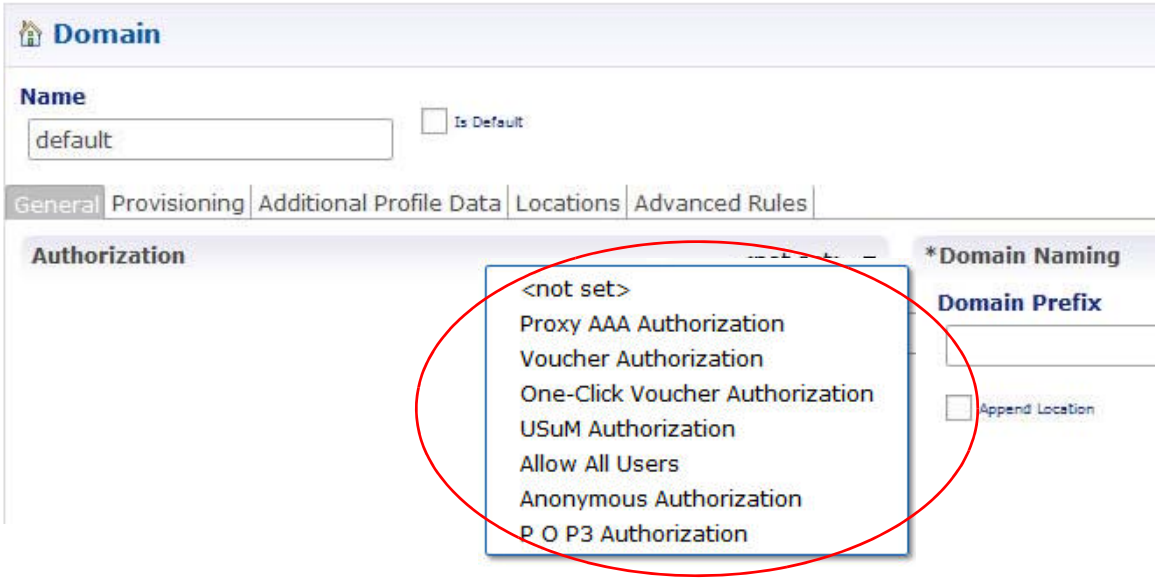

---Proxy AAA Authorization uses an AAA server to authenticate a subscriber.

---USuM Authentication sets a domain that authorizes the subscriber against Cisco Unified SuM.

---Allow All Users sets a domain that has no criteria for authorizing subscribers.

---Anonymous Authorization uses previously stored subscriber information.

b. Because this is a default domain, select USuM Authorization. This restricts the authorization to only those subscribers pre-registered in the system.

Step 5 Set the Domain Prefix to DEF.

The Domain Prefix concisely prevents similar subnet IP addresses or access attempts from mingling with each other.

Provisioning Subtab

No configuration is needed in the Provisioning subtab. No special provisioning are provided to a default user.

Additional Profile Data Subtab

No configuration is needed in the Additional Profile Data subtab. A Default domain is not intended to use special profile such as Sh, generic LDAP, LDAP bind, or SPR profiles.

Locations Subtab

No configuration is needed in the Locations subtab because any location is acceptable to the Default domain.

Advanced Rules Subtab

No configuration is needed in the Advanced Rules subtab. This screen is for TAL:, EAP, and unknown services.

Step 6 Click the Provisioning subtab > the Provisioning drop-down icon > Primary Credential Select button.

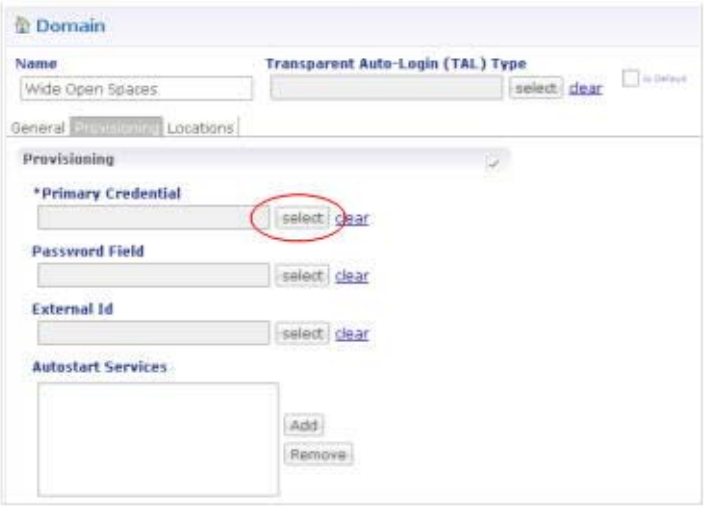

This displays an object list to pick from.

Step 7 Click RADIUS MAC Address in the object list.

This finds the MAC address that comes from the login attempt.

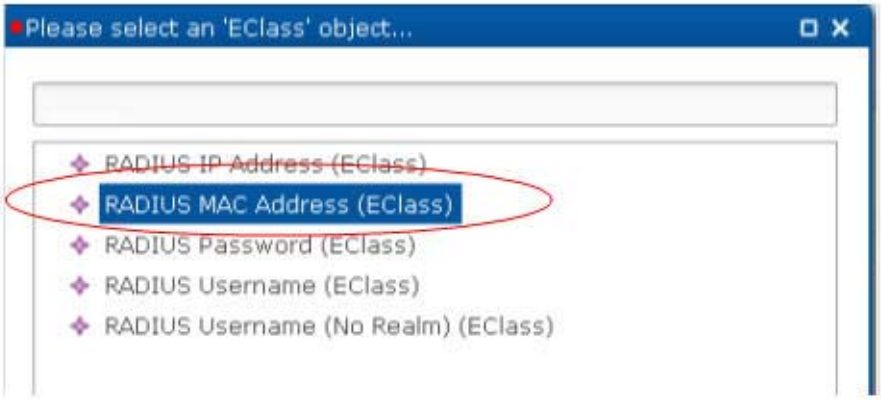

Step 8 Click the Add Button next to Autostart Services.

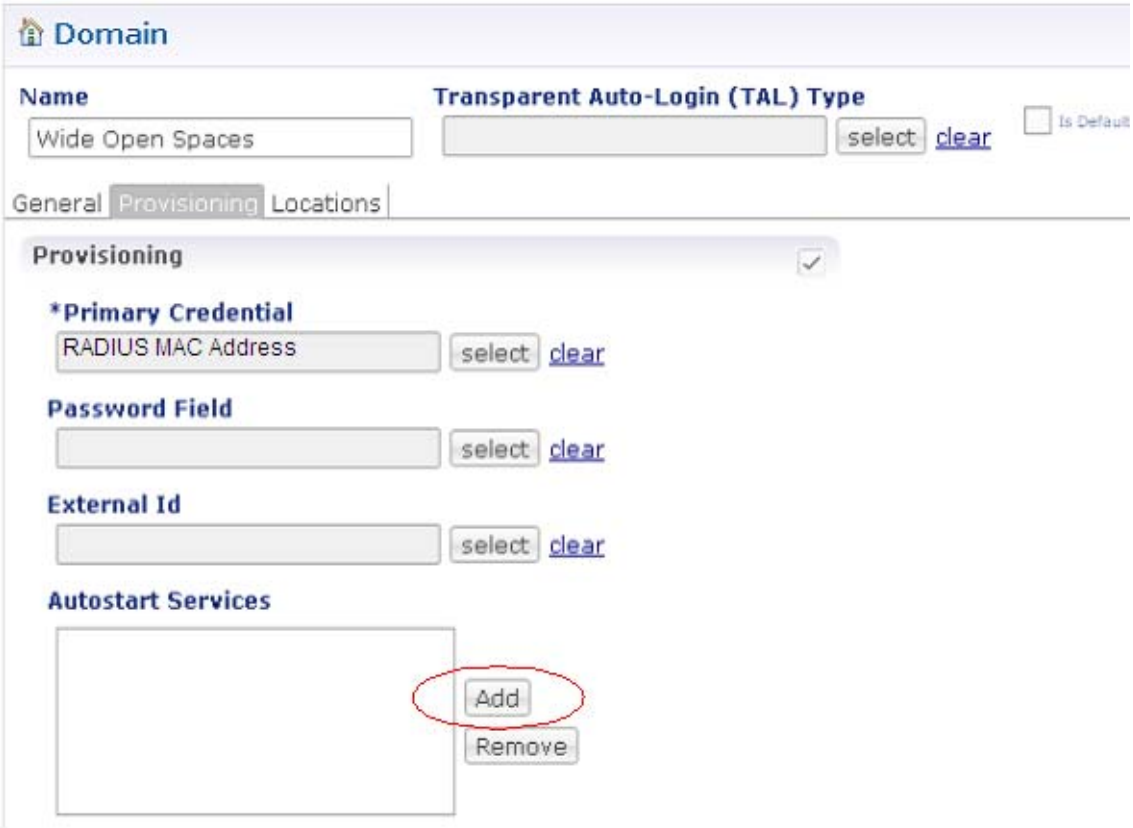

Step 9 Select Service A in the window.

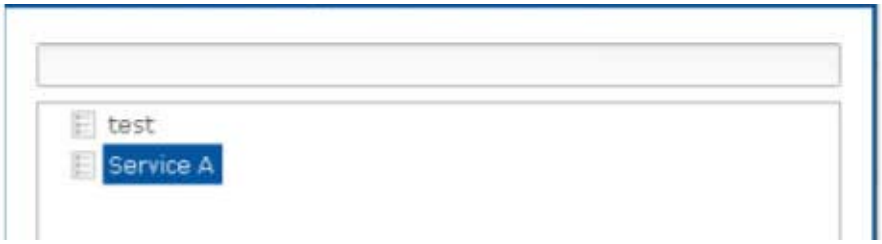

This provisions the subscriber with Service A when he attempts to log in.

Step 10 Make sure your domain looks like this under the Provisioning subtab.

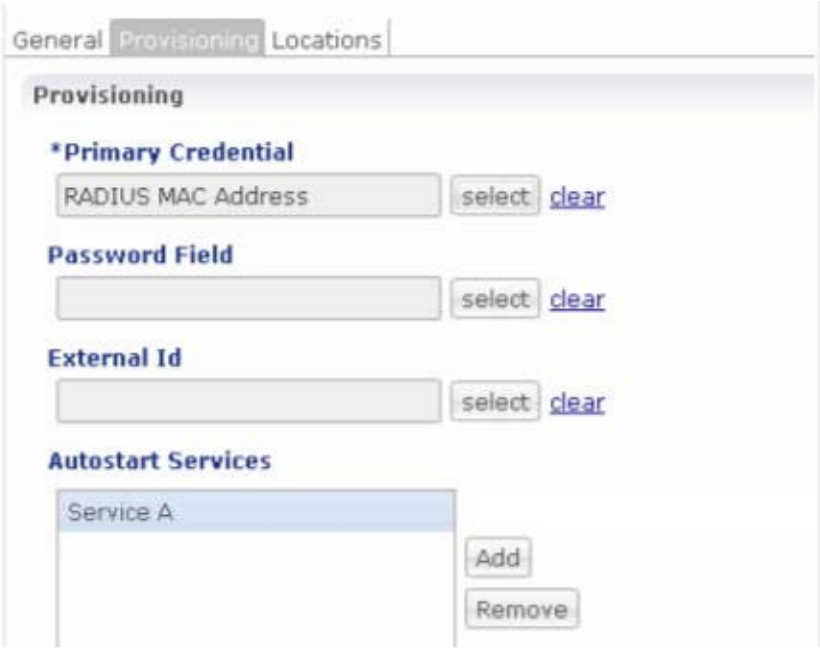

Defining a Domain to Limit Framed IPs

The domain created in these steps checks for specific framed IPs from which access to the network is provided.

Step 1 Click the Services tab > Domains node > Domain link.

General Subtab

Step 2 Select the General subtab.

Step 3 Enter the name for the new domain as Wide Open Spaces for example.

Step 4 De-select the check box for Is Default. This example is not a default domain, it is a domain for a specific purpose.

Step 5 Select the Authorization drop-down menu and select Allow All Users.

Step 6 Change the Domain Prefix to WOS.

The Domain Prefix concisely prevents similar subnet IP address or access attempts from mingling with each other. In this example, the Framed IP Location Type concretely defines the subscriber as being a subscriber in the WOS domain.

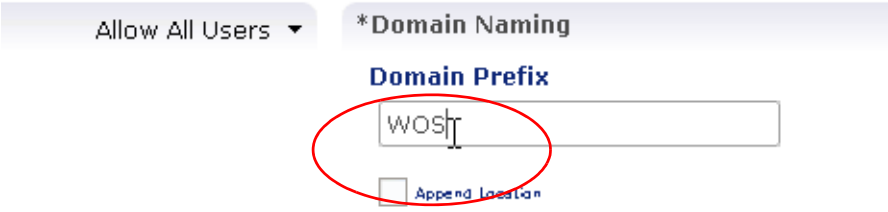

Provisioning Subtab

Step 7 Select the Provisioning subtab.

Step 8 Click the drop-down menu on the right.

Step 9 Click the Select button next to Primary Credential and choose USuM registration.

Step 10 Select the Add button next to Autostart Services and choose Service A.

Service_A is provided for the subscriber of this domain at log on.

Additional Profile Data Subtab

No work is necessary under this subtab.

Locations Subtab

Step 11 Click the Locations subtab.

Step 12 Next to the Location Matching Type field, click the Select button.

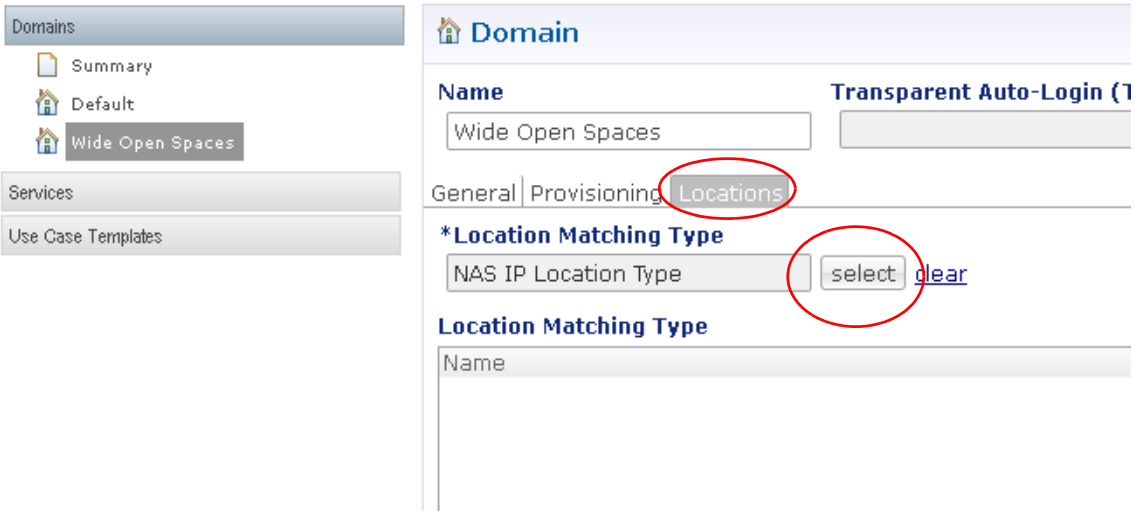

Note When using the Subscriber Services Portal GUI, there is no need to configure "Location Matching Type".

Step 13 Select the object Framed IP Location Type from the object list.

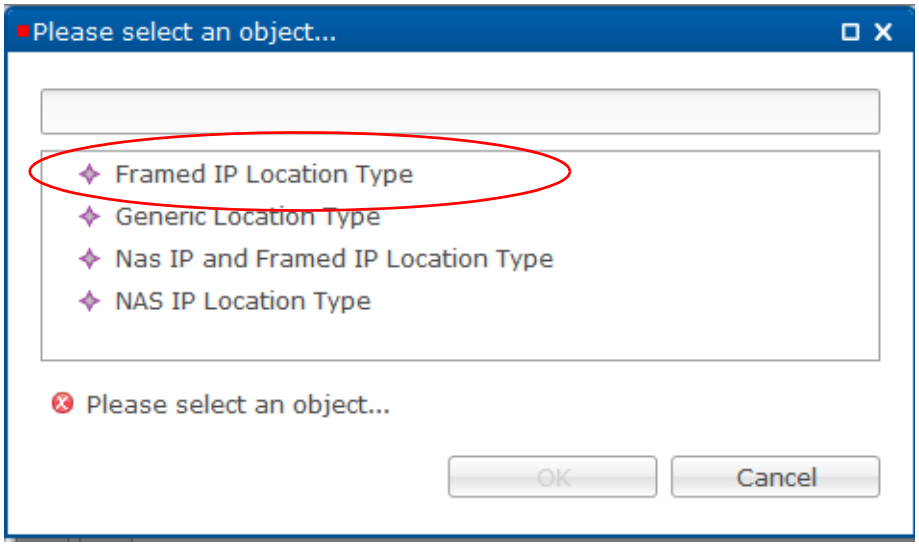

This checks the IP address before assigning the subscriber to the domain.

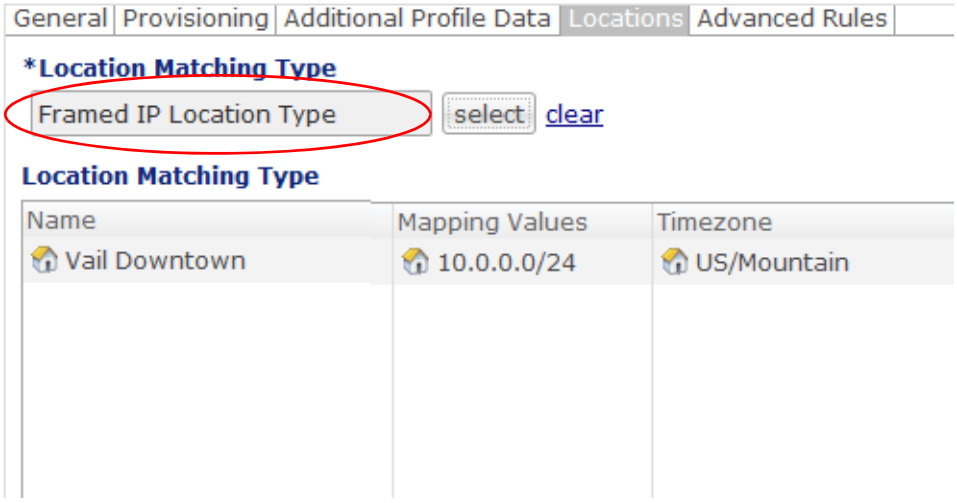

Step 14 Click the Add button to add a row to the Locations table.

Step 15 In this row, for Name enter Vail Downtown, as in the example.

Step 16 For Mapping value, click in that column, then click the drop-down menu icon on the right.

Step 17 Enter the IP addresses of the Vail downtown subnet, 10.0.0.0/24 perhaps.

Step 18 Click Add button and then OK.

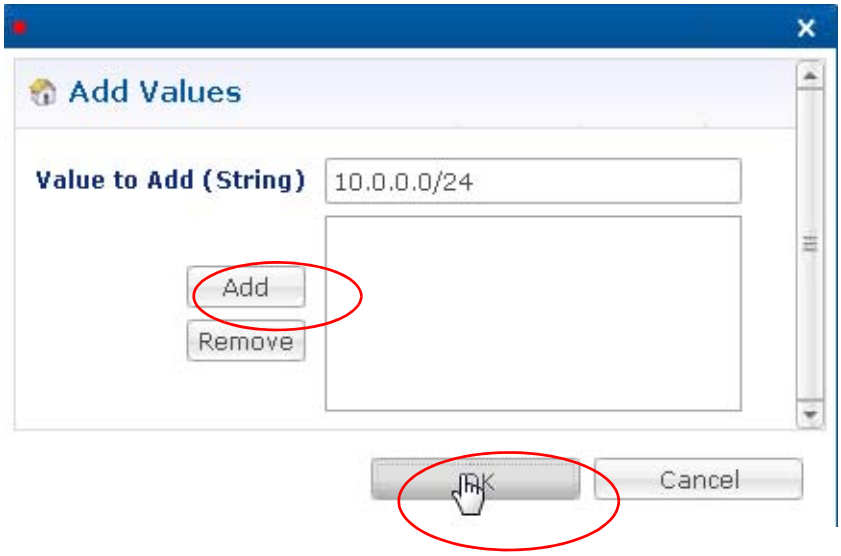

This determines the subnet IP addresses that limit this domain.

Step 19 Make sure your Domain screen Locations tab looks similar to this one.

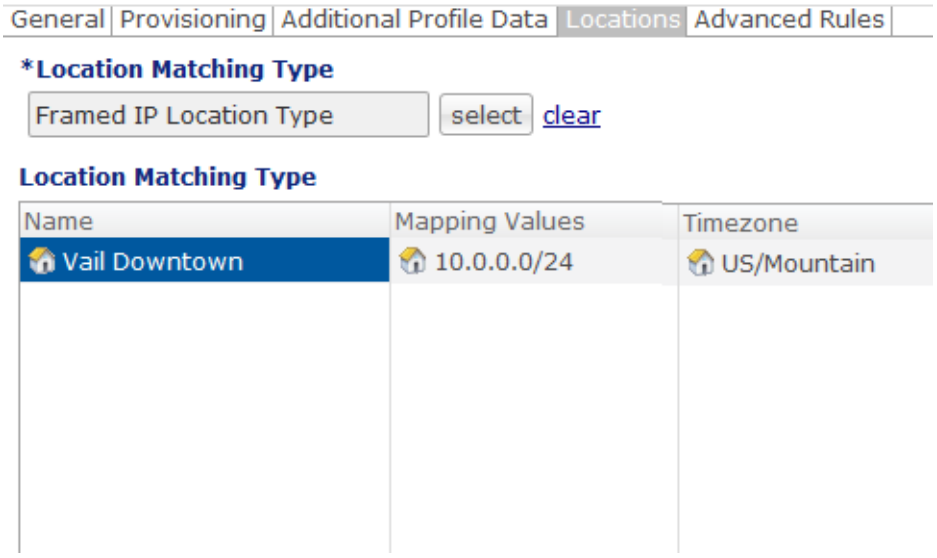

Advanced Rules Subtab

No configuration is necessary under the Advanced Rules subtab.

Step 20 Go to Step 7.

Defining a Registered Users Domain

A Registered Users domain type is frequently used, and with it you can create a more likely domain that looks for a subscriber name/password case, using a known subscriber.

Often, a Registered Users Domain is the default domain, with authorization set to Cisco Unified SuM. These subscribers access the network with a known Username and a Password at log on.

General Subtab

Step 1 Set up a domain called Registered Users Domain so that it looks like the figure below under the General subtab.

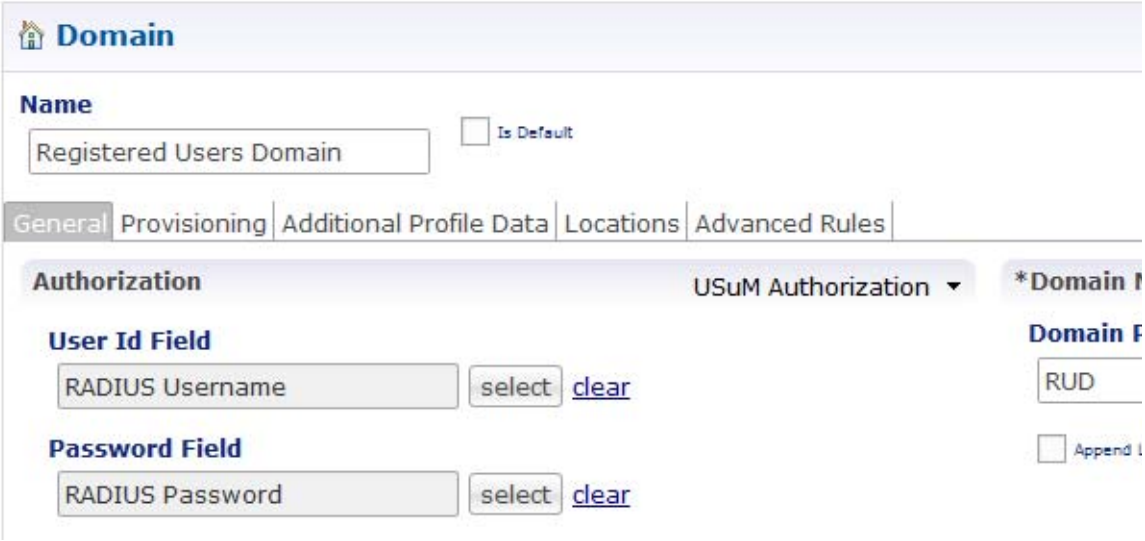

For this domain,

- No configuration is necessary under the Provisioning subtab because subscribers are already provisioned under Cisco Unified SuM.
- No configuration is necessary under the Additional Profile Data subtab.
- No locations or subnets need to be set under the Locations subtab.
- No configuration is necessary under the Advanced Rules subtab because subscribers are authorized and authenticated with RADIUS usernames and passwords.

Go to Step 7.

Enhanced Location Query

In previous versions of Cisco Policy Builder, the only location information collected from the WLC was this:

- IP address
- MAC address

The WLC controller has been updated so that when a RADIUS message is sent, the message can be configured to include output for the SSID and AP_MAC, and AP_Group as well.

Location queries in CPS can include all of these location data:

- IP address, as before
- MAC address, as before
- SSID
- AP_MAC - access point machine access code
- AP_Group - access point group ID

Location Query Overview

Configuring location query has two tasks:

- Configuring the WLC
- Configure Cisco Policy Builder

The order you use to perform these tasks does not matter, but they are both required.

If you have your location queries configured to use IP address or MAC address, this configuration is not affected by updating to CPS 7.0.

If location information cannot be determined, a default location is used and no error is generated.

Configuring the WLC

Configure the WLC so that it sends the SSID and AP_MAC in the RADIUS request. This topic is not addressed in this document. Instructions for configuring the WLC are found in the configuration guide named *Cisco Wireless LAN Controller Configuration Guide*, Release 7.0, June 2010, OL-21524-01, Chapter 6, section Configuring RADIUS.

Here is an example of the screen you must use to configure the WLC.

Configure Cisco Policy Builder

Step 1 Log in to Cisco Policy Builder and click Services tab > Domains > domain name > Location sub tab.

You can use any domain that relies on a WLC. For example, our example uses the USuM authorization domain.

Step 2 Click the Locations subtab. A typical Location usage is shown below.

Step 3 Configure the Location Matching Type field as NAS IP Location Type.

Step 4 Set the other subtabs as you require.

Step 5 In the Location Matching Table, configure a row that shows a friendly name and a subnet range. You may use a time zone if you want to further restrict your location matching.

Step 6 To set Mapping values, add subnets with the ellipses icon on the right.

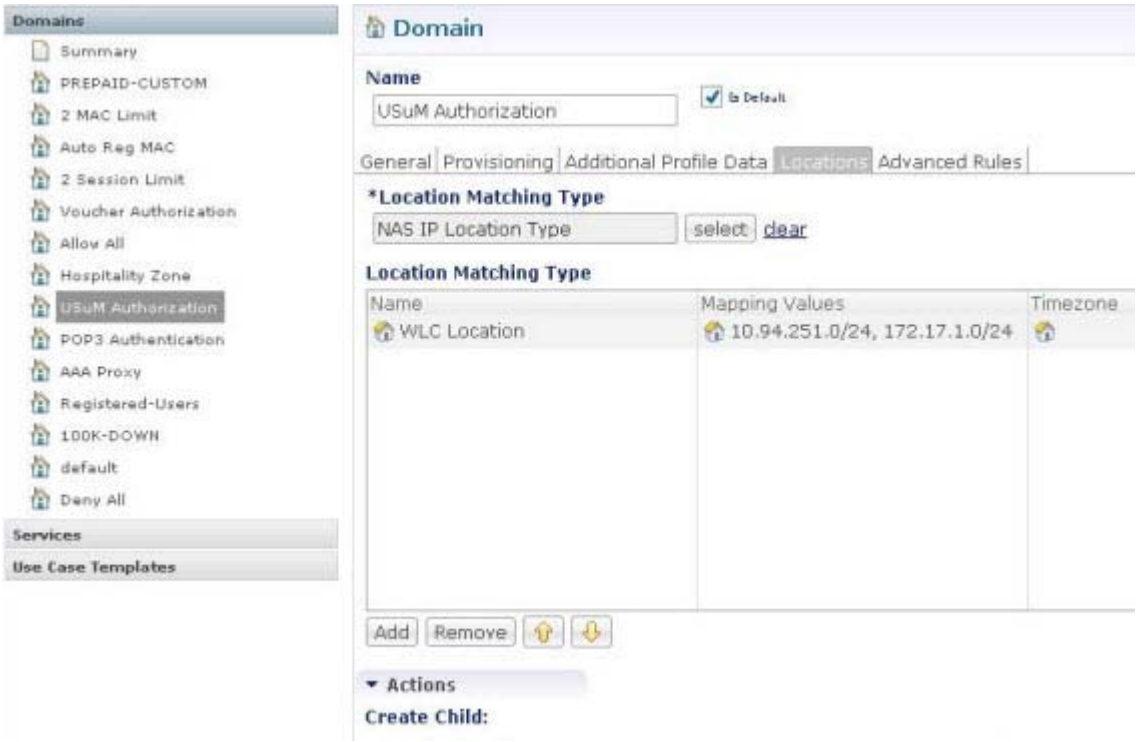

Name	The name of the domain.
Is Default	De-select the check box for Is Default. This example is not a default domain, it is a domain for a specific purpose.
Location Matching Type	Select NAS IP Location Type for this field. The subscriber will be matched on the location type in the table below.
Name	The name for the WLC domain.
Mapping Values	This is the range of IP addresses the domain will cover. Add and remove IP address with the ellipses icon on the right. ...
Timezone	Select a time zone if you want to have more specificity about your subnet and location.

Step 7 \Test Your Work

Test your configuration with these steps:

Step 1 To check for correct access, log on using a test subscriber logon.

Step 2 Verify correct portal redirection,

Step 3 On the portal, validate the subscriber and that they are registered in Cisco Unified SuM.

To configure the use cases needed for your company's deployment, refer Elective Service Configurations.

Elective Service Configurations

Revised: July 10, 2015

This chapter shows how to configure your software for your specific service configurations.

Elective service configurations are ones that you choose to implement and are shown in the list below. You can use one, many, or all of the elective service configurations to serve your subscribers.

Steps for testing your work are included as part of each procedure.

Open a browser window and log on to the Cisco Policy Builder interface for all tasks.

Before you configure the CPS, make sure that you have these tasks completed:

- Complete the required use cases as described in Required Service Configuration. You must have all of the required service configurations complete before you can implement any of the elective service configurations.
- Choose the use cases you want from this chapter and configure them
- Set up at least one domain as described at Domain Configurations.
- When you are all done, perform the tasks described in Elective Service Configurations.

Elective Use Cases

Configure one, many, or all of these use cases to enable the services you want to provide to your subscribers.

Note Before you work through any of these use cases, you *must* have completed all of the required use cases shown in Required Service Configuration.

- Example Service Plan
- Configuring a MAC-based TAL Use Case
- WISPr Use Case
- Tiered Services Use Case
- Voucher-based Services
  - Configuring a Voucher-based Service
  - Configuring a Time-based One-click Voucher Service
- Concurrent Logons Service Option

- Bandwidth on Demand Use Case

Example Service Plan

A service (that is, a service offering or service plan) is combination of service options. This example shows a service that uses several service options.

Although you may use a scheme where one service option is used as one service, it is more typical to bundle service options into a single plan.

Step 1 In a browser window, open a Cisco Policy Builder interface, or examine the figures provided here.

Step 2 Click Service tab > Services node > Services folder > Service A.

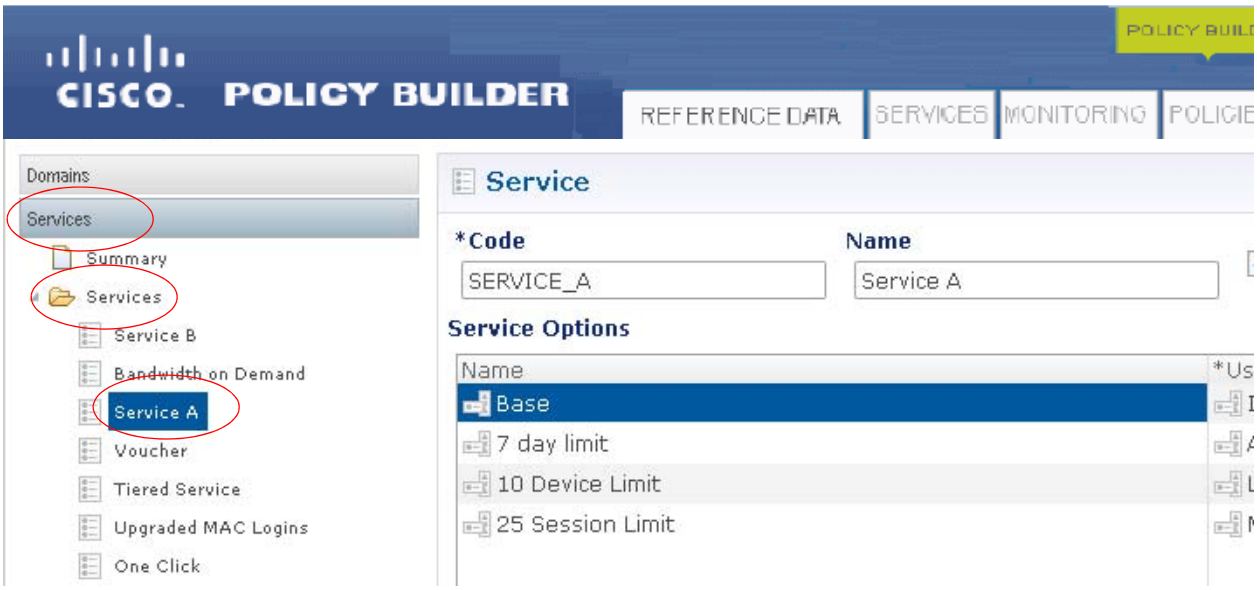

Service A is comprised of these service options:

- Base
- 7 day limit
- 10 Device Limit
- 25 Session Limit

This Service Option...	is based on this use case template:
Base	ISG Base Service
7 day limit	Auto Register MAC Credential
10 Device Limit	Limit Max MAC Registrations
25 Session limit	Max Concurrent Sessions

Step 3 Click the Base service option in the tree and look at its content.

The Base service option uses the Base ISG Service service configuration and uses the 512K DOWN service template for a transmission speed.

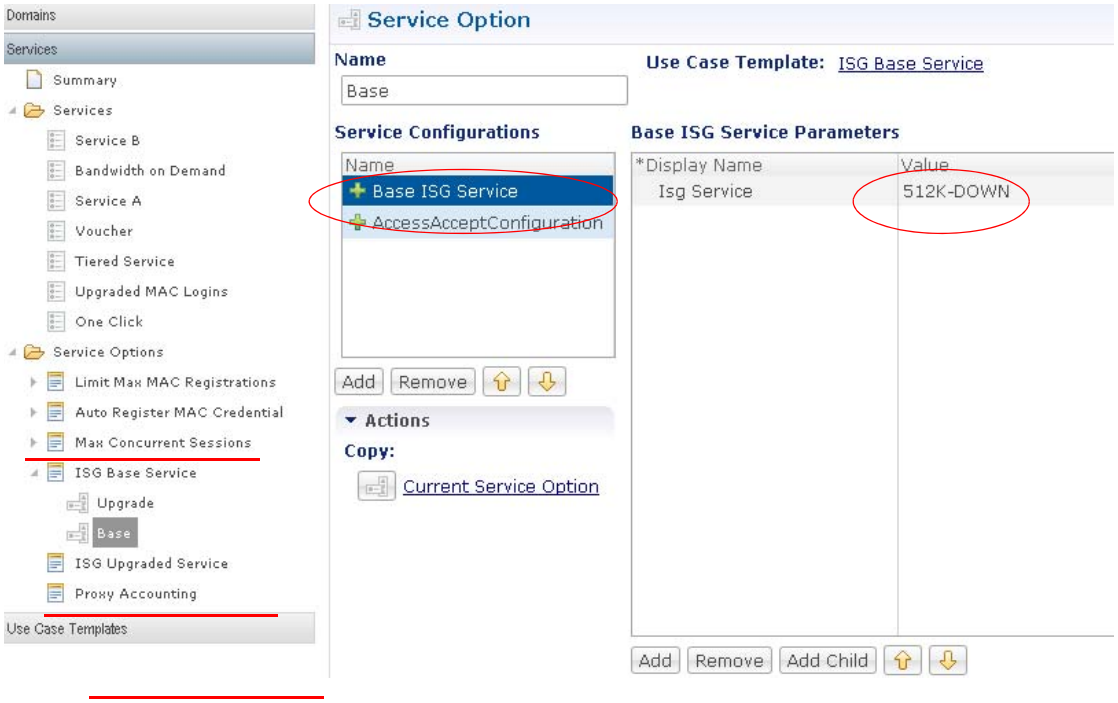

Note Check the RADIUS Service Template node under the Reference Data tab to see where this originated.

Step 4 Also notice the Service Option screen for ISG Base Service.

The Base Service Option also uses the Access Accept Configuration service configuration, which uses the ISG_ACCESS_ACCEPT RADIUS Service template.

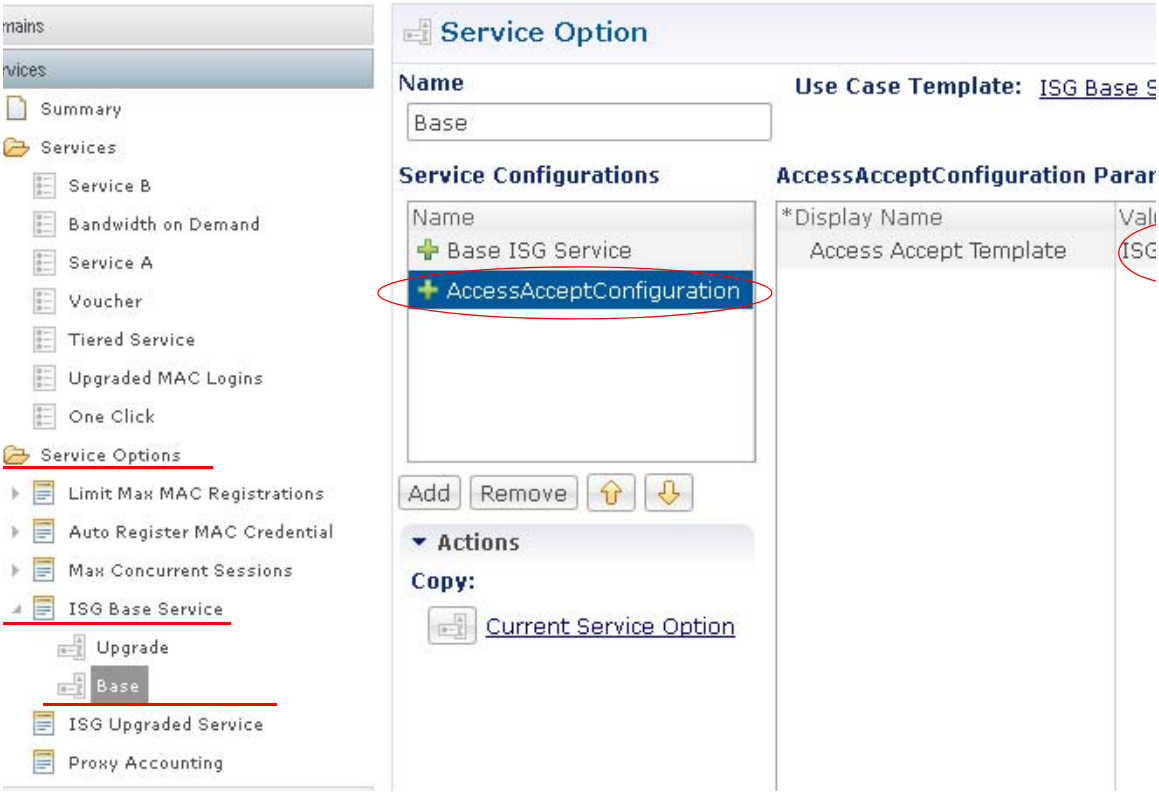

Note The ISG_ACCESS_ACCEPT is defined by the RADIUS Service Template for Service Provider Specific Template, in the AV pairs table. This template lets you specify idle time outs and session time outs, for example.

Only the Base service option is shown in this example, but the same relationships exist for the other service options that make up Service_A.

Configuring a MAC-based TAL Use Case

This procedure shows how to configure MAC-based TAL log ons. This use case stores the subscriber for a specific number of days, letting the subscriber log on, and then automatically re-log on based on the hardware MAC address and the days elapsed since the last log on attempt.

Enable this use case to ensure that any subsequent use cases are available to the domain, which enforces the number of log ons.

For example, a subscriber would log on with their usual credentials, and once the normal credentials are authorized, CPS stores those credentials for a designated period of time. The next time the subscriber tries to access the network, CPS logs on the subscriber by recognizing their MAC address.

Consider a restaurant or coffee shop that provides Wi-Fi service. The subscriber views the initial log on screen and then logs in. CPS learns the MAC address and stores it so the next time the subscriber comes in, CPS recognizes the subscriber.

The predefined use case template for this use case is Auto Register MAC Credential.

Steps

This use case presents these examples:

- 1—Steps for Limiting the Number of Days
- 2—Steps for Limiting the Number of Devices
- 3—Steps for Auto-provisioning the MAC Address

1—Steps for Limiting the Number of Days

These steps show how to set up CPS to remember a Wi-Fi subscriber for a set number of days after initial logon, based on their MAC address.

Step 1 Click Service tab > Services node.

Step 2 Click Service Options node in the tree and open it to show its content.

Step 3 Select the Auto Register MAC Credential item in the list.

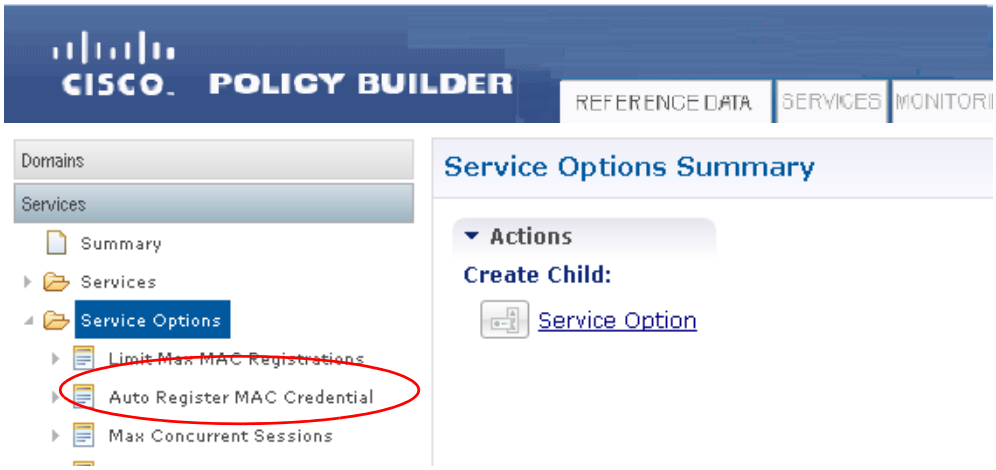

Step 4 Click the Service Option link on the main window.

Step 5 Click the Registration Limit service configuration to review its parameters.

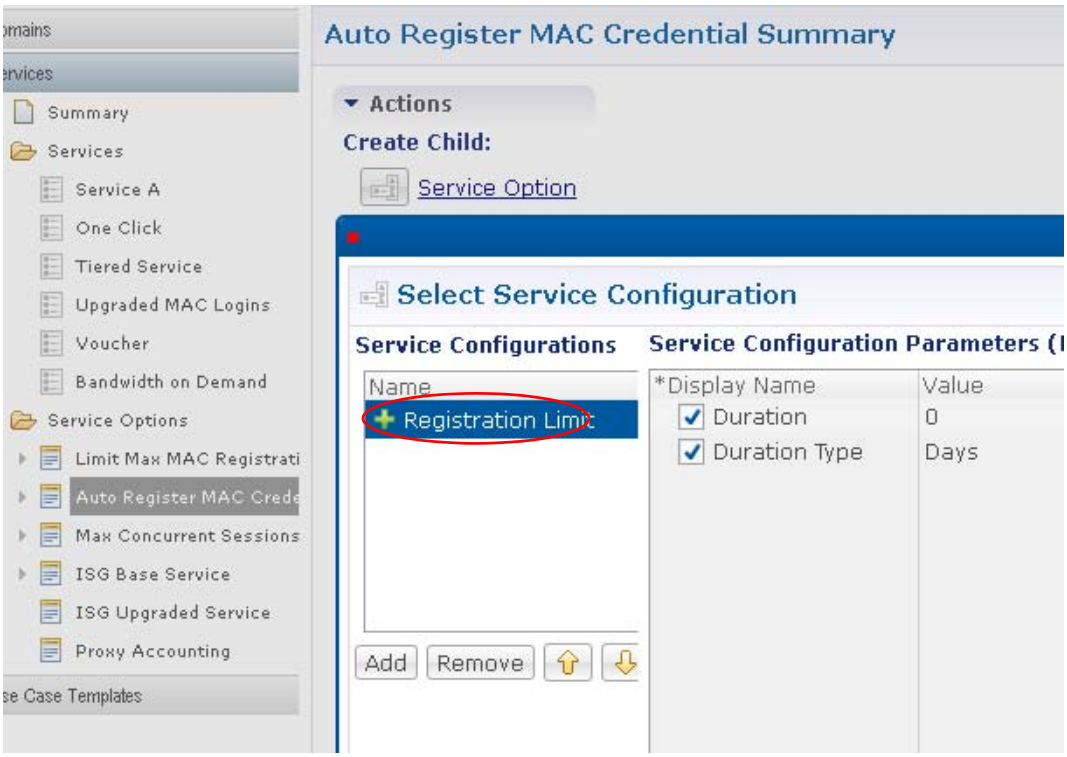

Step 6 Click OK to accept.

This displays the Service Option screen. This is where you set the MAC-based TAL.

On the Service Option screen that appears, notice that it shows the use case template being used as a base is the Auto Register MAC Credential.

Step 7 In the Name field, call this service option **7 day limit**.

This indicates that CPS remembers the patron for seven days after the initial logon.

Step 8 Double click the Value column and set the Duration to 7 and Duration Type to Days.

Now you have a service option that limits the number of days a device is remembered. This service option can be used in any service or service plan.

Step 9 Test your service options with Test MAC-based TAL.

2—Steps for Limiting the Number of Devices

This example also uses the MAC-based TAL use case, but limitation is based on the number of MAC address instead of on the amount of time.

For example, you may want to permit a subscriber to register only 5 MAC addresses to reduce ID sharing, but still leave enough MAC addresses so that a subscriber could log on with any of several of their own devices. Consider a cable account holder. If you set the limit of devices to 10, you permit up to 10 devices within this account to share the Wi-Fi privilege.

This example uses the Limit Max MAC Registrations use case.

Step 1 Click Services tab > Service Options node > Limit Max MAC Registrations > Service Options link.

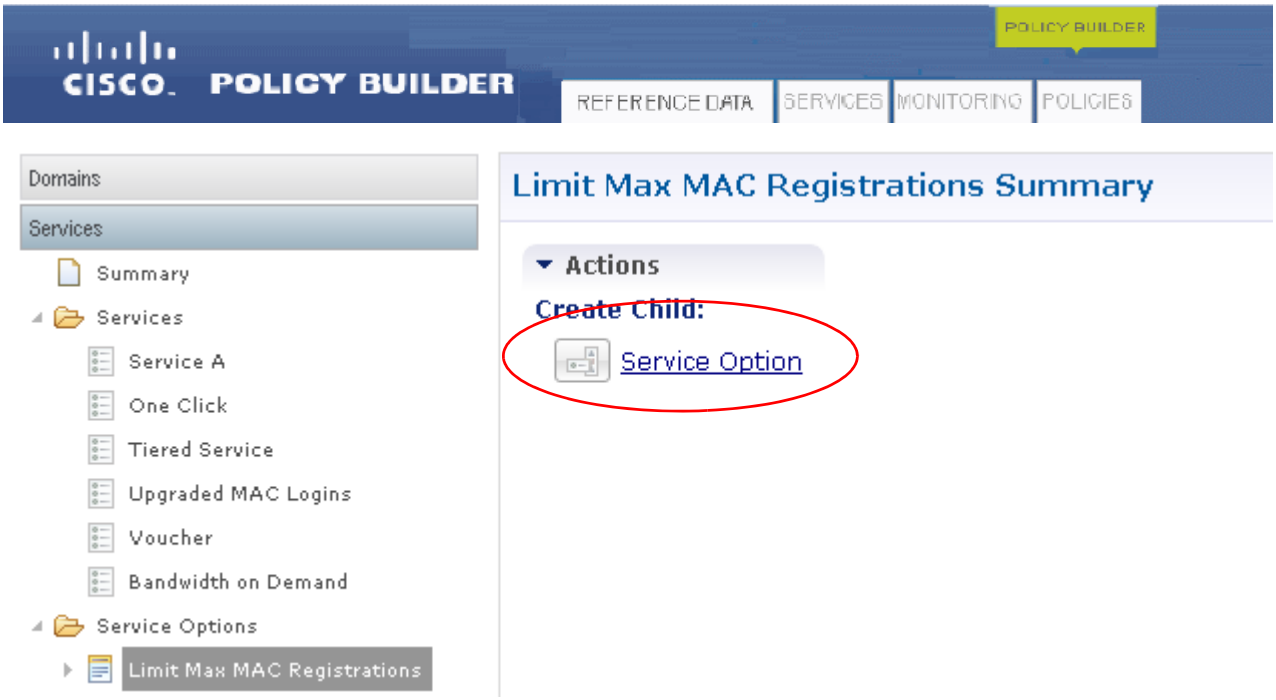

Step 2 Select the MAC Limit service configuration on the left and review its parameters.

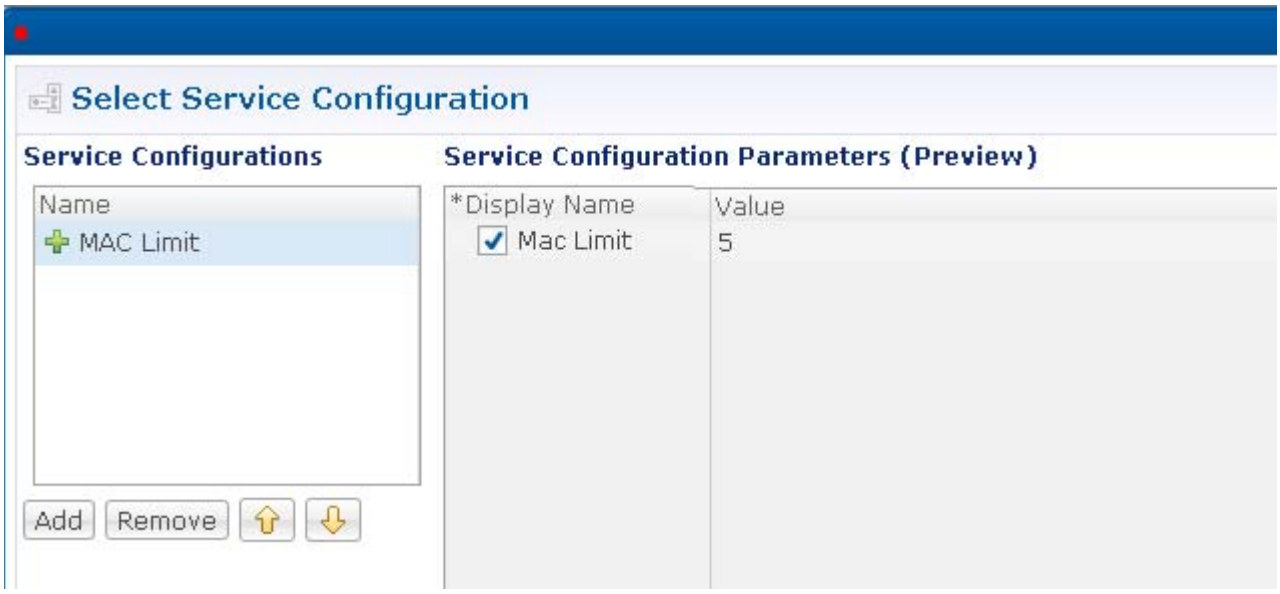

Step 3 Click OK to accept.

On the Service Option screen that displays, notice that it shows the use case template being used is the Limit Max Mac Registrations.

Step 4 Name this Service option 10 Device Limit.

Step 5 Click the MAC Limit item in the Service Configurations list to display the current values.

Step 6 Double click in the Value column and change the 5 to a 10.

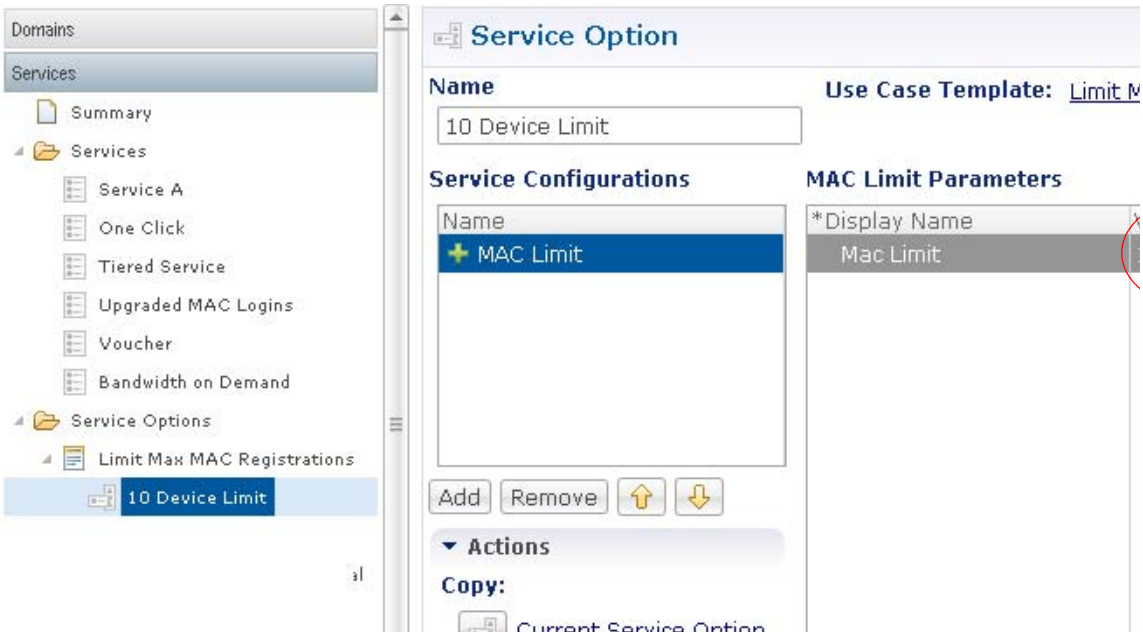

Now you have a service option that limits the number of devices logged in with MAC-based TAL. This service option can be used in any service or service plan.

Step 7 Test your service options with Test MAC-based TAL.

3—Steps for Auto-provisioning the MAC Address

This example uses the one-click scenario. You provide the subscriber with a one-click logon page on your portal so that they can log on quickly if they do so within a 24-hour period.

- Create a one-click service option
- Create a service that uses the one-click service option
- Create a domain that automatically provides the one-click service when the subscriber logs in to that domain.

3a—Create a One-Click Service Option

Step 1 Click Services tab > Service Option node > Auto Register MAC Credential > Service Option link.

Step 2 Click the Registration Limit item in the Service Configurations list and review its parameters.

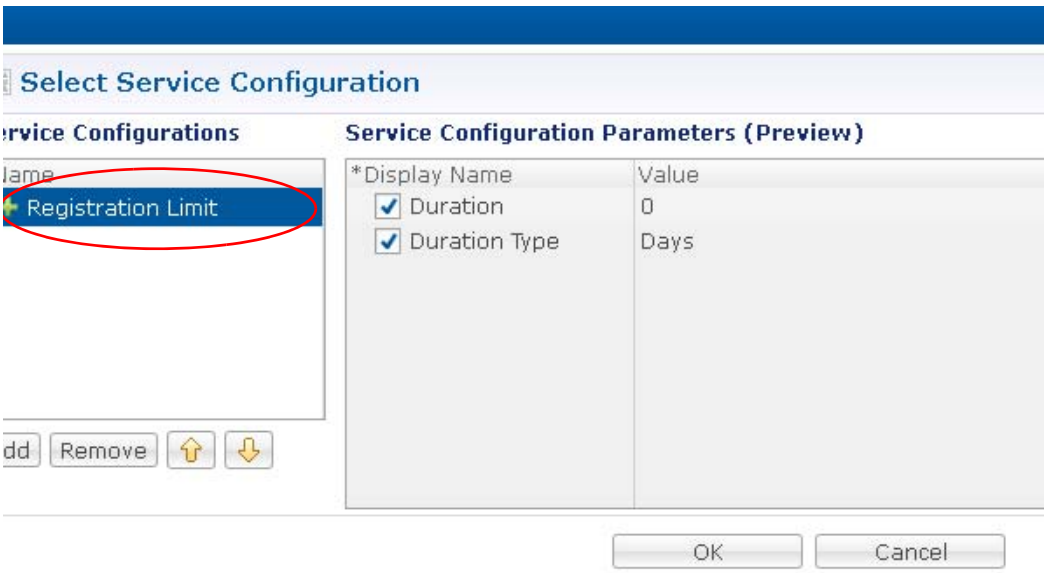

Step 3 Click OK. Your new service option is based on these parameters.

Step 4 In the Name field, provide the name 6 Hour Limit.

Notice that this Service Option is based on the use case template Auto Register MAC Credential.

Step 5 In the Service Option screen that displays, click the Registration Limit item in the Service Configurations list.

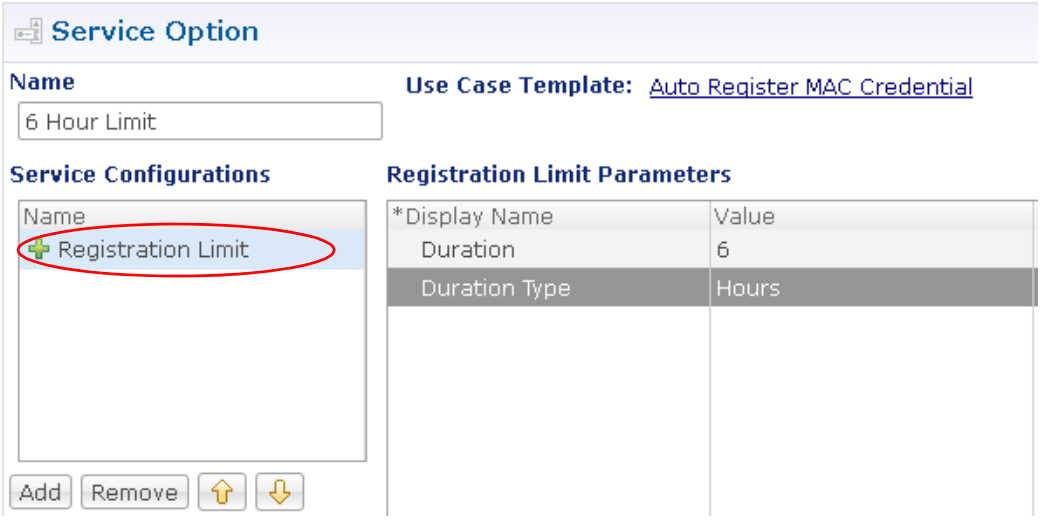

Step 6 Double click in the Value column and change the Duration value to 6 and the Duration Type to hours.

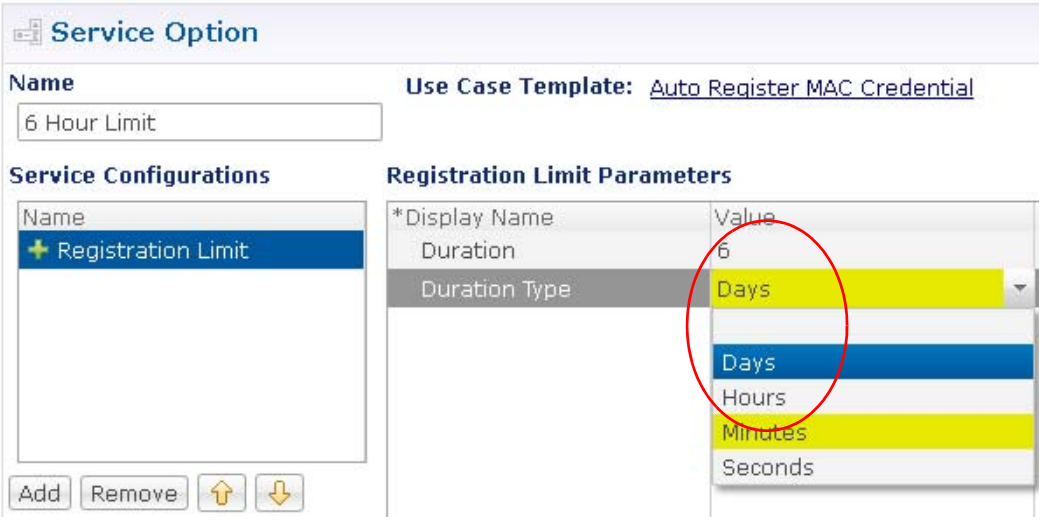

Now you have a service option to use that requires the subscriber to click and log on at least every 6 hours.

3b—Create a Service that Uses the One-Click Service Option

Step 1 Click Services tab > Services node > Services folder > Service link.

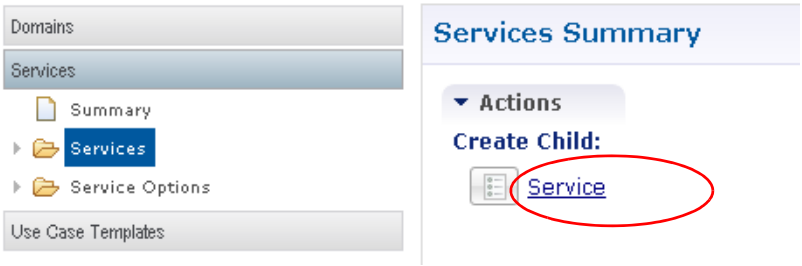

Step 2 Fill in the Services screen.

- For Code, enter ONE-CLICK.
- For Name, enter One-Click.
- Select the Enabled check box (service is displayed in Control Center if this check box is selected).
- Click the Balance Service check box.

Step 3 Click the Add button to add service options.

Step 4 Click on Base service to add that as a service option to the One-click service.

Step 5 Click OK.

Step 6 Click Add again and click on **Add 6 Hour Limit** to add that service option.

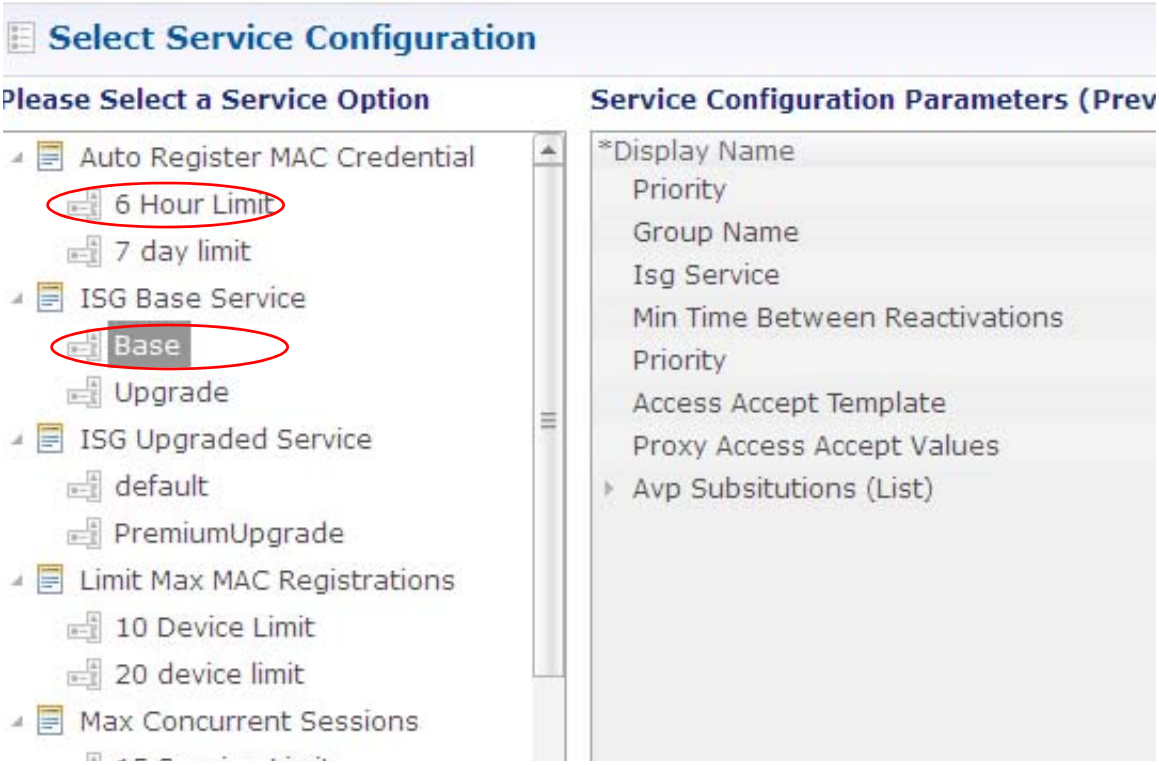

Step 7 Click OK.

Now you have a service that provides a basic service level *and* requires the subscriber to re-authenticate every six hours.

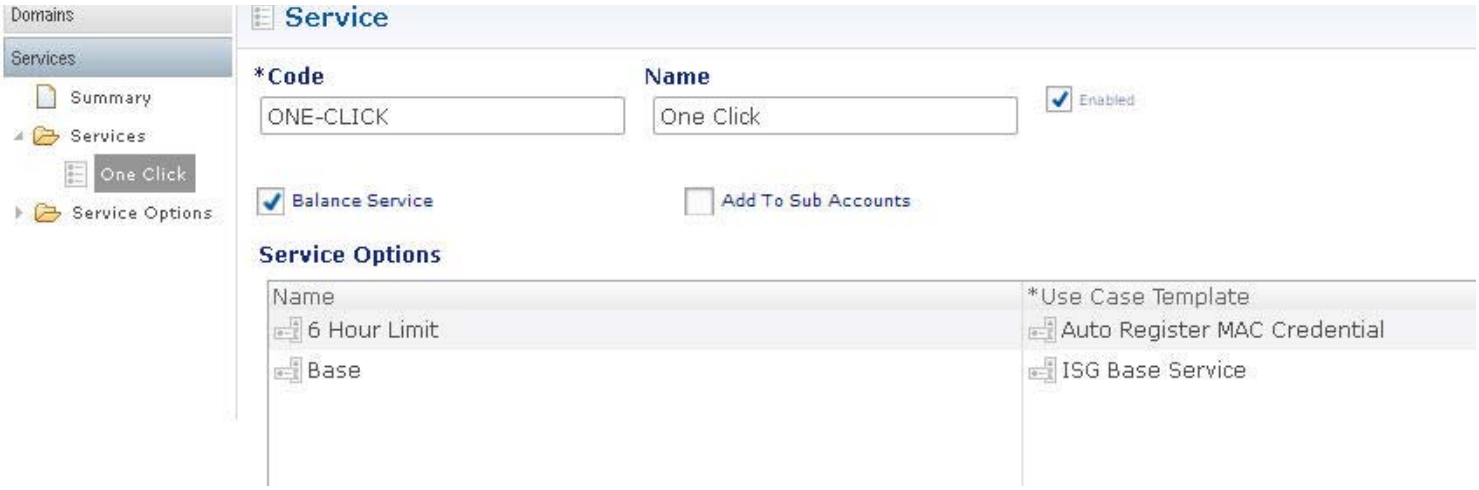

3c—Create a Domain to Provide the One-click Service

To create a new domain that supports a one-click portal page use these steps.

Step 1 Click Service tab > Domains node > Domain link.

General Subtab

Step 2 Configure data in the General subtab.

a. Select the General subtab.

b. Enter the name as One-click.

c. Click the drop-down arrow for Authorization and select Anonymous Authorization.

Anonymous Authorization uses previously stored information.

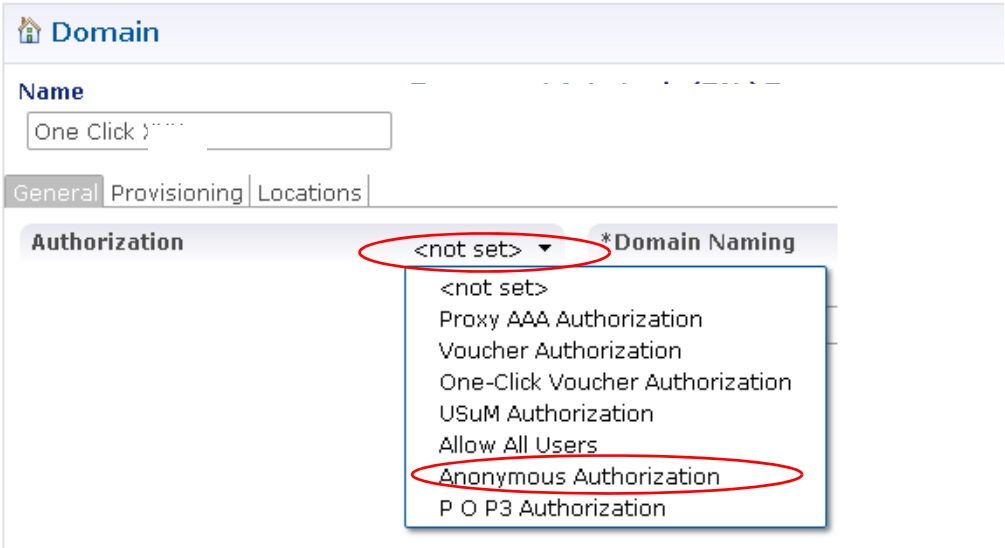

d. Click the select button next to the User Id Field.

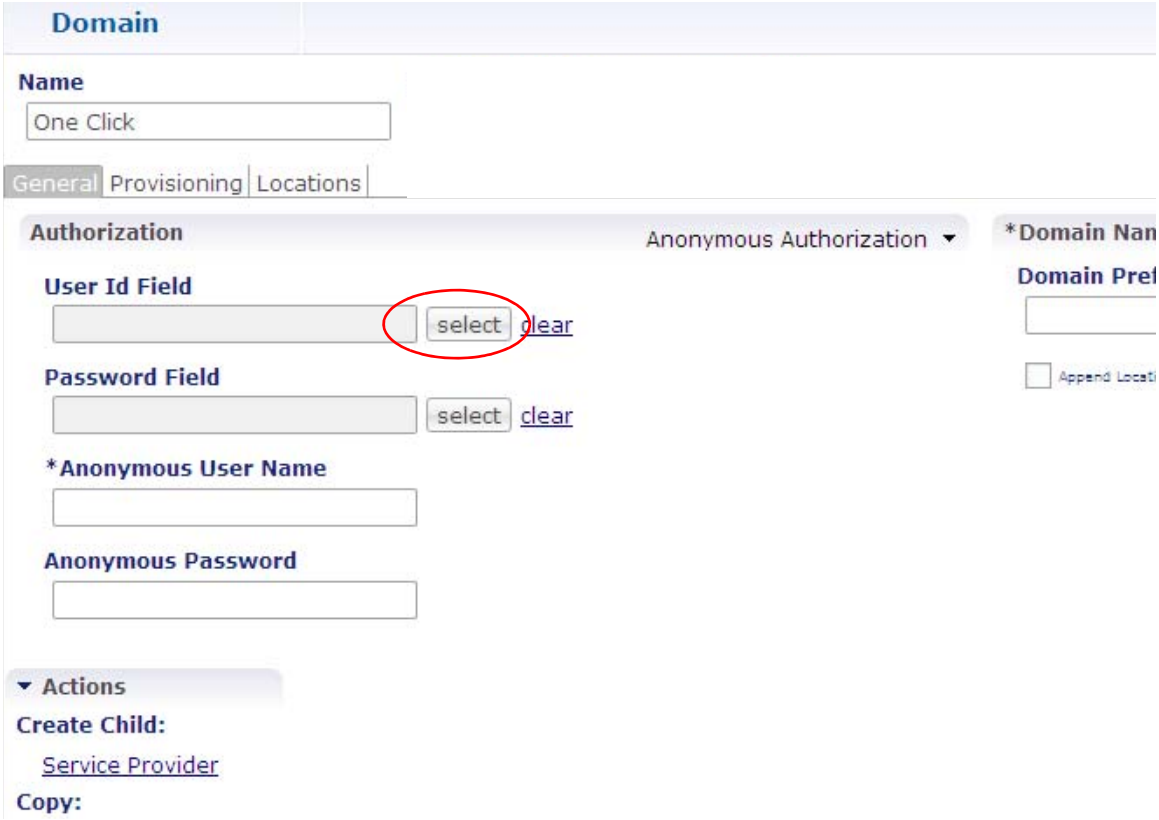

e. Select RADIUS Username from the list.

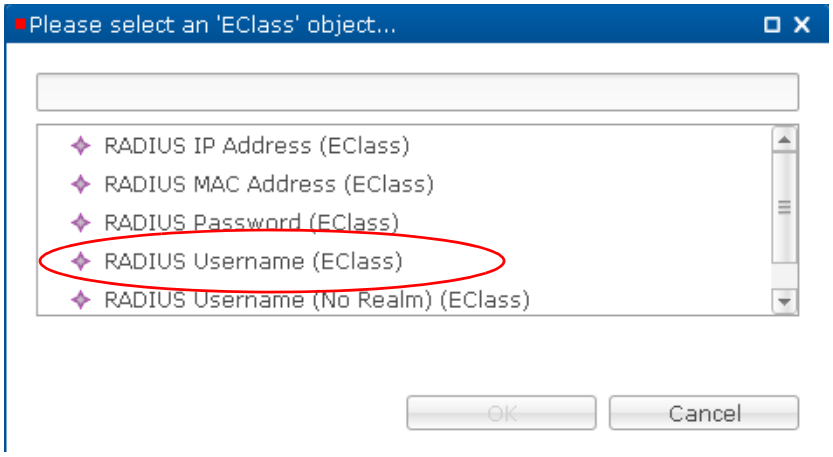

f. Click the select button next to the Password Field and select RADIUS Password from the list.

g. For Anonymous User Name, enter one-click.

This automatically informs the one-click portal page of the user name.

Note Your one-click portal page must match what you enter here.

h. For Anonymous Password, enter p@ssword.

This automatically informs the one-click portal page of the user password.

i. For Domain Prefix, enter one-click, a unique domain to keep track of exactly the subscribers that log on from the one-click portal page.

Your Domain screen General subtab should look like this:

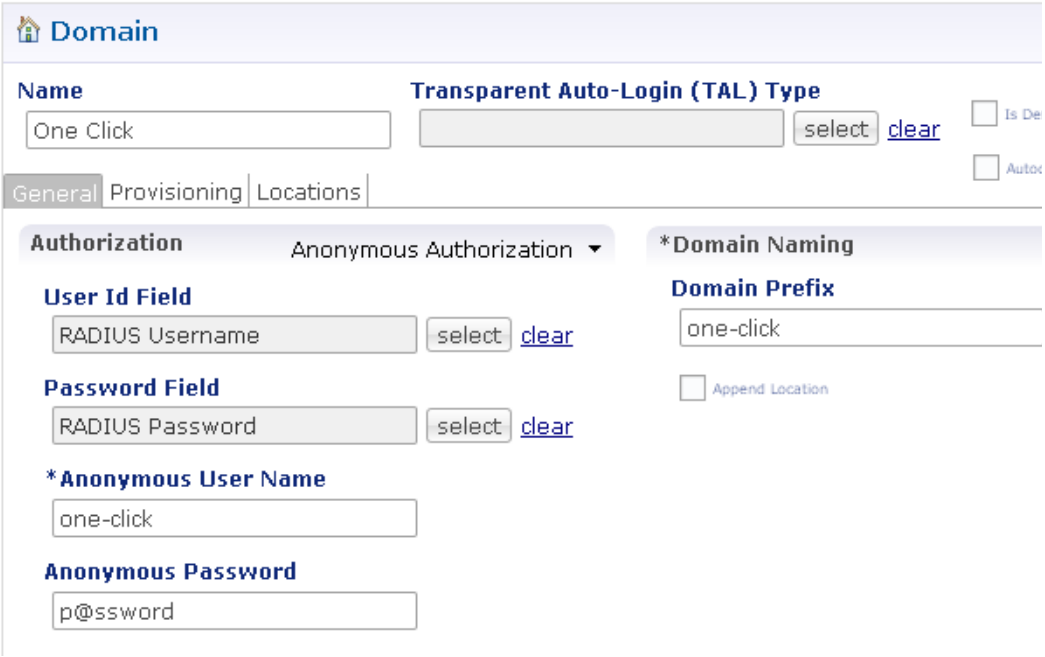

j. Go to Provisioning Subtab.

Provisioning Subtab

Step 3 Configure the Data in the Provisioning Subtab.

a. Click the Provisioning subtab.

b. Click the drop-down menu icon on the right to show the provisioning choices.

c. For Primary Credential click the Select button.

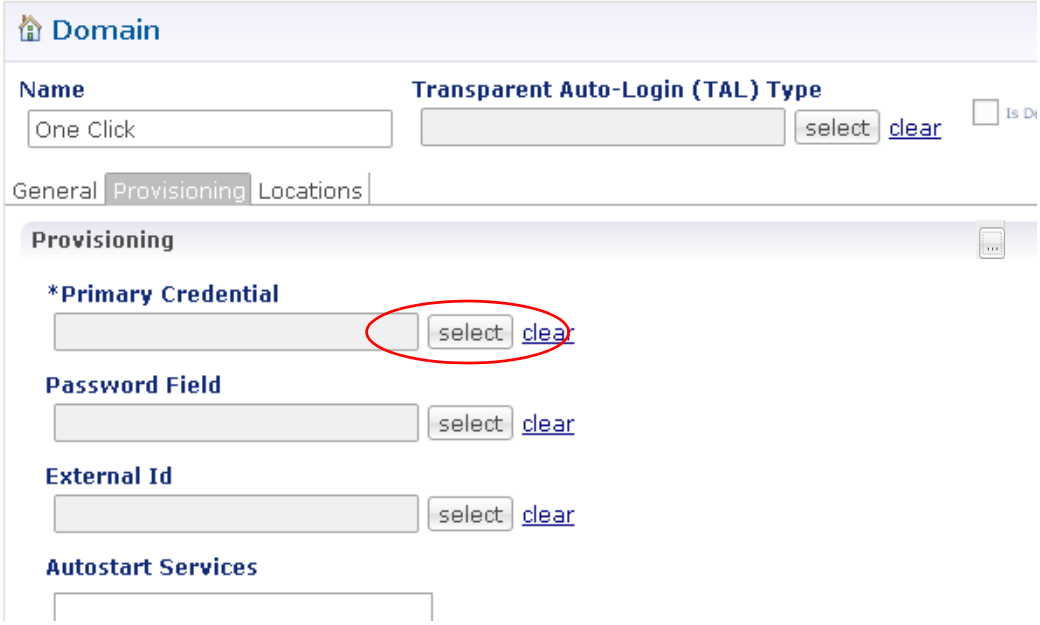

d. Select RADIUS MAC Address(EClass) from the list.

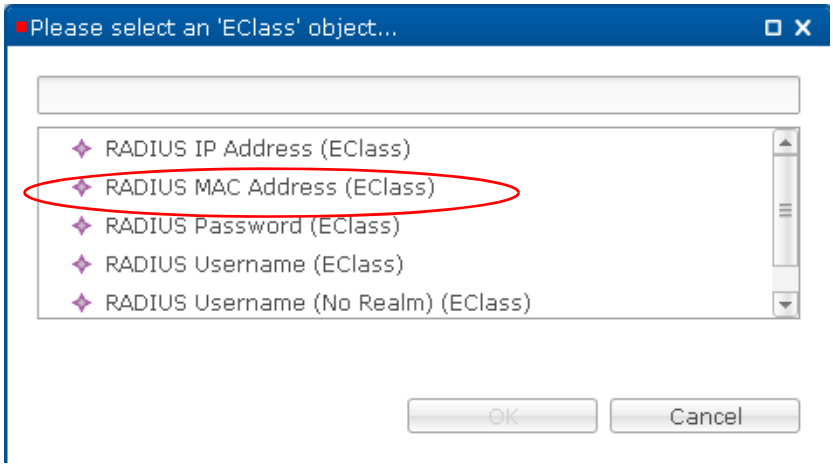

e. Click OK.

f. Click the Add button next to the Autostart Services area.

g. Select the service One Click.

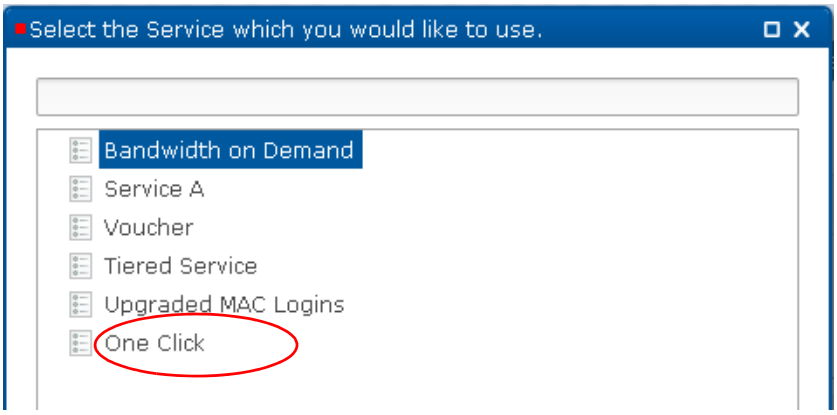

This means that when the subscriber logs in, they automatically uses the One-Click Service that you defined earlier in 3b—Create a Service that Uses the One-Click Service Option.

Your screen should look like this under the Provisioning subtab:

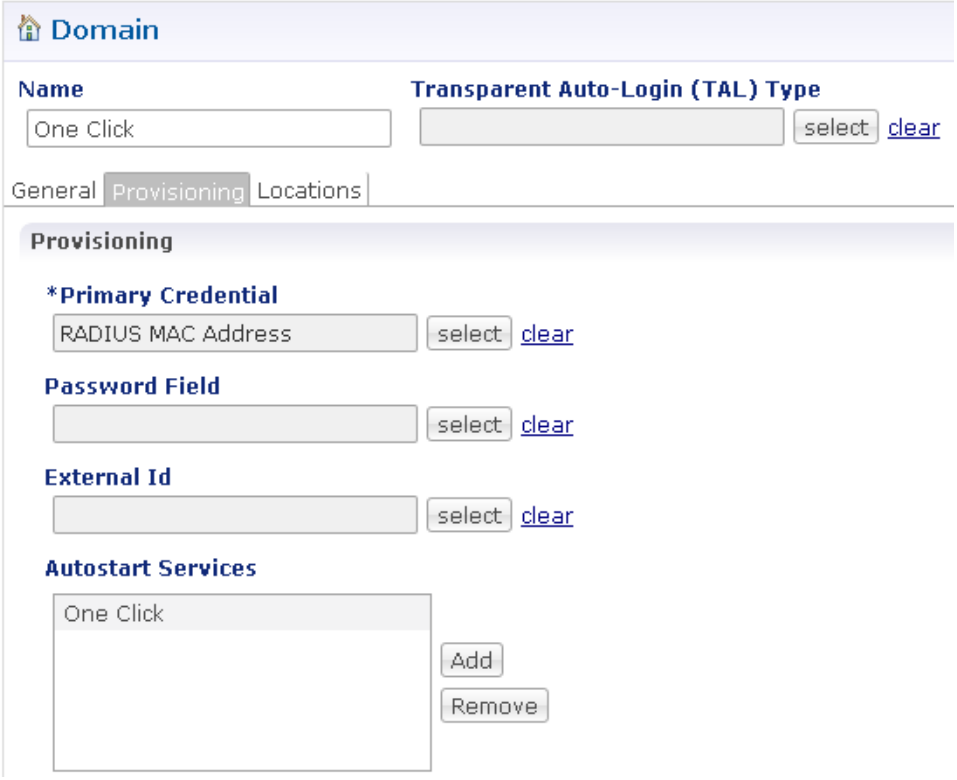

h. Go to Location Subtab.

Location Subtab

Step 4 Configure Data in the Locations Subtab this way:

a. Still on the Domains screen, click the Locations subtab.

b. Click the select button next to Location Matching Type.

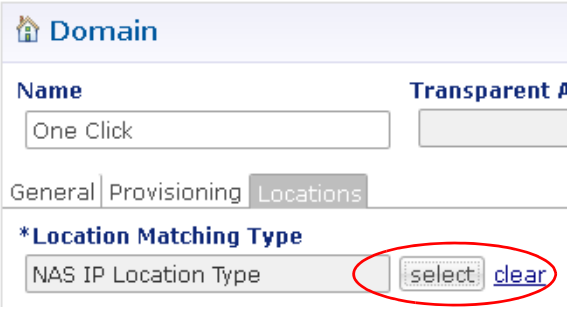

c. Select Framed IP Location Type from the list and click OK.

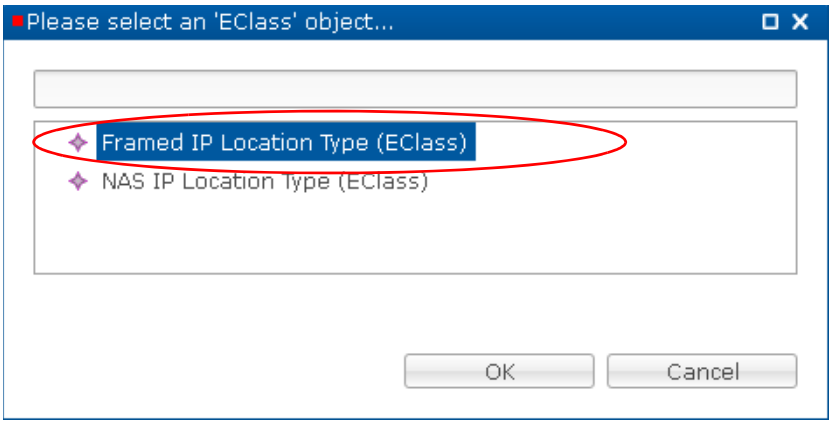

d. Click the Add button at the bottom of the screen to add a row in the Location Matching table.

e. Click in the Name column and enter one-click for the subnet name.

f. Click in the Mapping Values column and then click the drop-down menu icon . This displays the Add Values screen.

g. Enter the IP address subnet for this location, for example, 12.0.0.0/24.

h. Click Add and then OK.

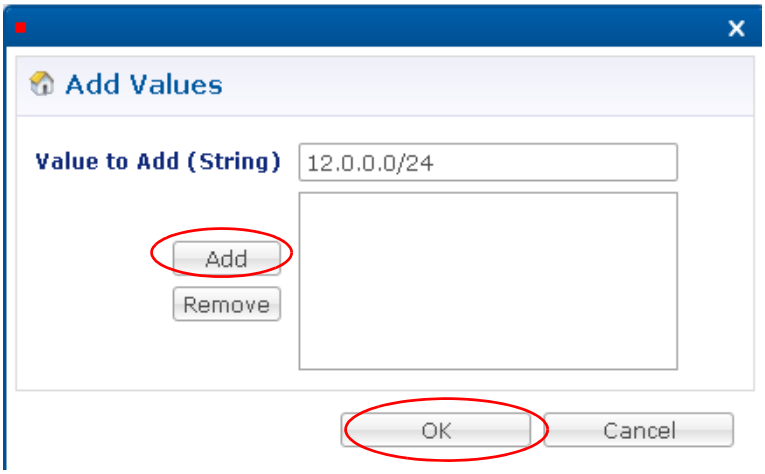

i. Your Domain screen should look like this under the Locations subtab.

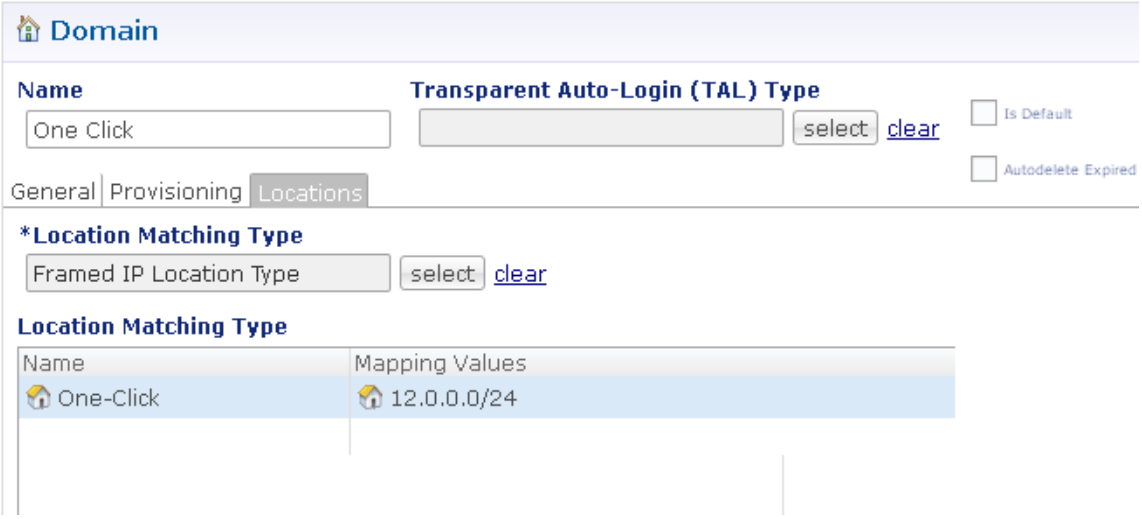

Advanced Rules Subtab

Step 5 Click the Select button next to the field Transparent Auto-Login (TAL) Type.

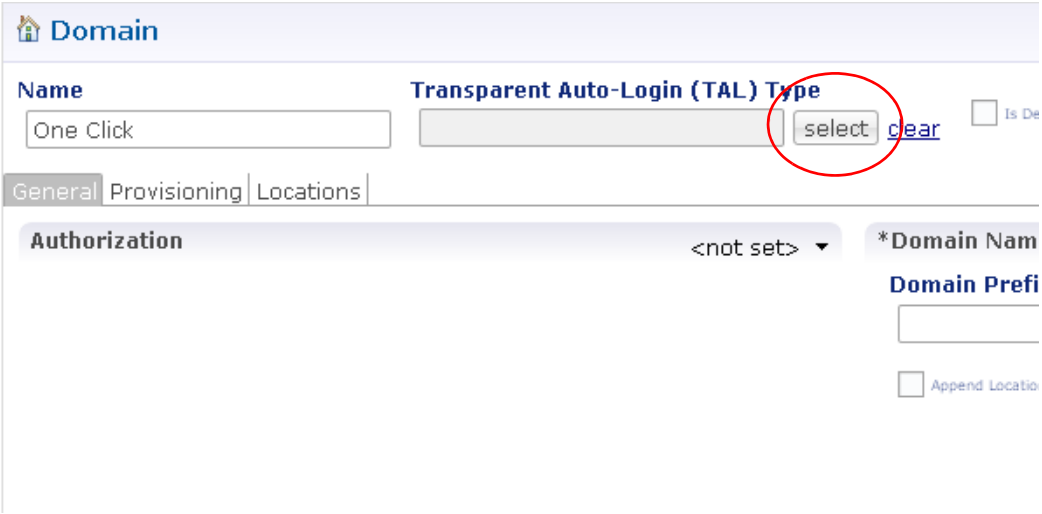

Step 6 From the object list, select RADIUS MAC Address.

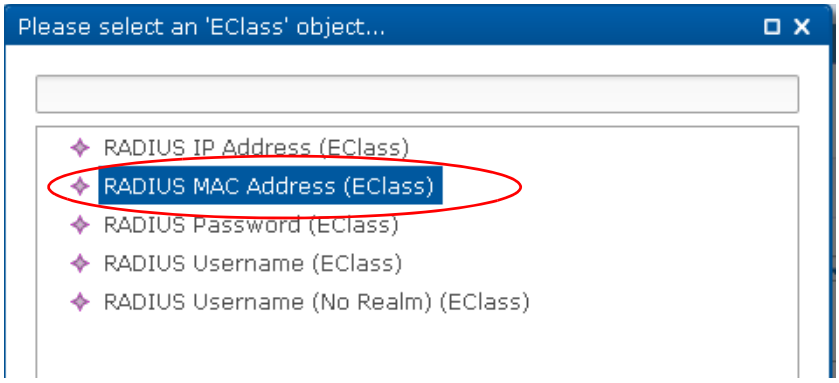

This causes the Domain to store and evaluate the subscriber's RADIUS MAC address.

Step 7 Test your configuration as described in Test MAC-based TAL.

Test MAC-based TAL

Test your service option configuration with these steps:

Step 1 Check that portal redirection occurs when in the subscriber portal.

Step 2 Make sure that MAC credentials are stored in Cisco Unified SuM.

Step 3 Make sure that disconnect, reconnect, and no redirection occur properly in the subscriber portal.

WISPr Use Case

This use case is managed in the subscriber portal and is discussed in that documentation.

Wireless Internet Service Provider roaming, or WISPr, allows a subscriber to roam between wireless Internet service providers, in a fashion similar to that used to allow cellphone users to roam between carriers. A RADIUS server is used to authenticate the subscriber's credentials.

See Subscriber Services Portal User Interface Guide, Document ID 178-727-002-x.x.

Tiered Services Use Case

This use case applies policy rules, such as bandwidth speed, usage limits, number of logons, and more, and bases them on the service plan, on a per tier basis.

Enabling such a use case also segments your subscriber base by offering differentiated services and pricing such as these:

- Premium, high-speed access services for business subscribers and gamers
- Basic access speeds for e-mail and browsing subscribers.

For example, your company may offer all subscribers a base package of Service A, but they also want to offer premium subscribers a better QoS and a higher MAC device limit.

Configuration tasks for a tiered type service are these.

- Define a new RADIUS service template that has better upload and download speeds.
- Create an upgrade service options.
- Create a upgrade service.

Steps

1—Create a Service Template for Upgraded Service

If one is not available, create a service template for the upgraded, or premium, service in the tiered service with these steps.

Step 1 Click Reference Data tab > RADIUS Service Templates node > Service Provider Specific Templates.

Step 2 Select the 512K-Down Service template and review the parameters.

Your premium service must have better upload and download speeds that what is provided.

Recall that the BASE_INTERNET_SERVICE template does not deal with speeds, allowing you the flexibility you are seeking now.

Step 3 Make a note of the QU;100000;D;512000 value. You might copy it to a note pad.

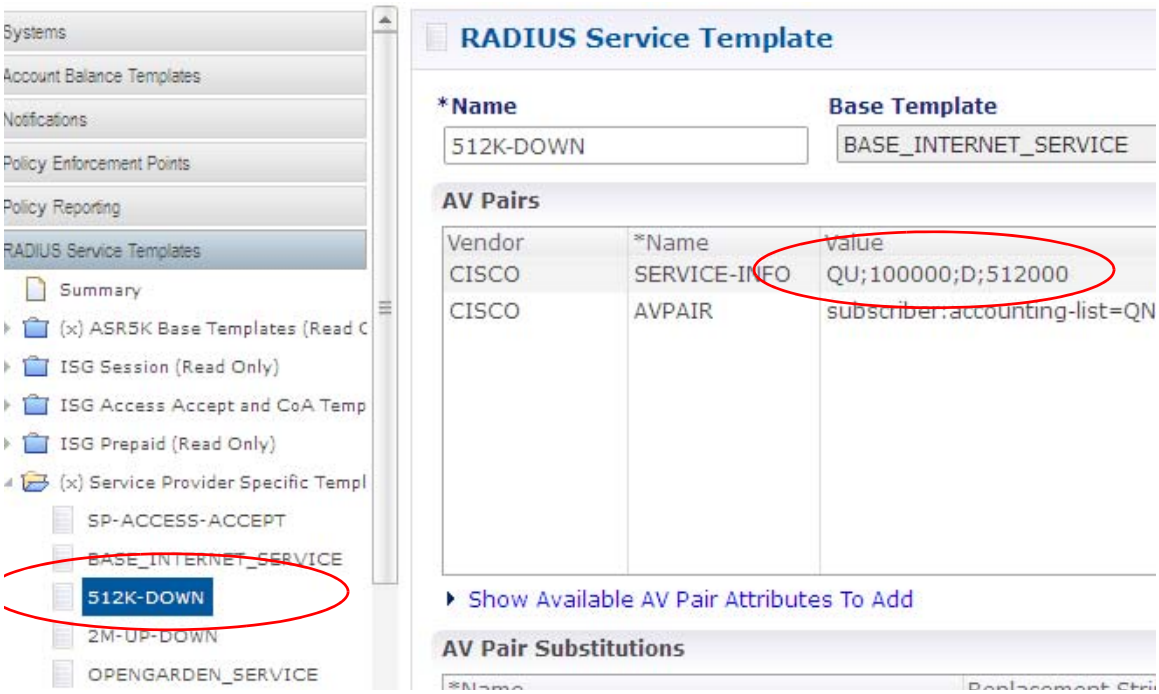

Step 4 In the tree, select Reference Data > RADIUS Service Templates > Service Provider Specific Templates > RADIUS Service Template link.

Fill in the Radius Service Template screen to create an upgraded tier for premium subscribers.

Step 5 Provide the name as 2M-UP-DOWN.

Step 6 Click the select button next to the Base Template field.

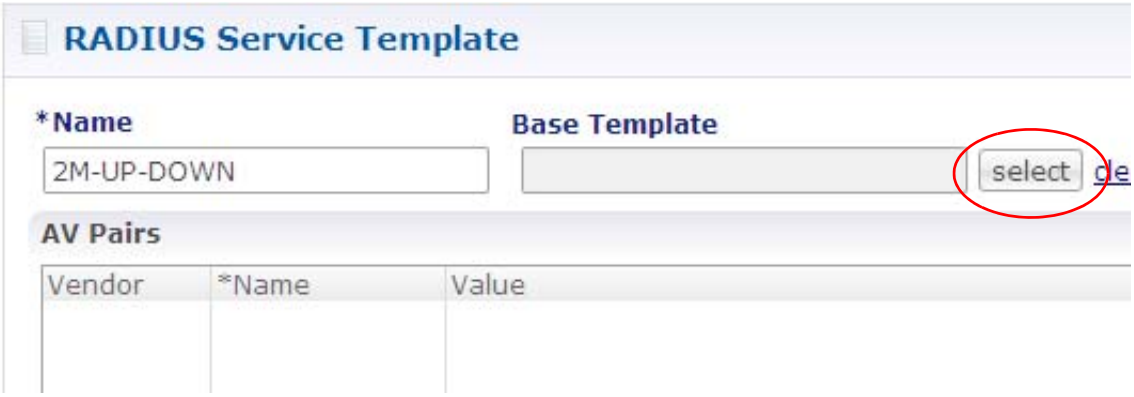

Step 7 Open the Service Provider Specific Templates item.

Step 8 Select to use BASE_INTERNET_SERVICE as the base for this service template.

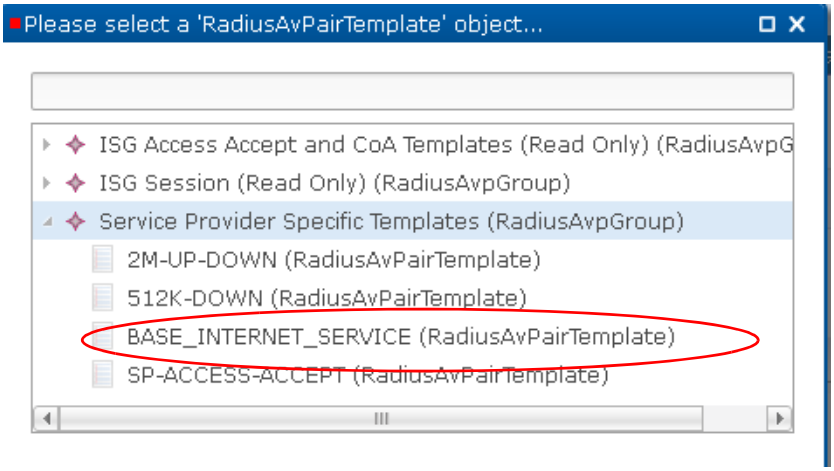

Step 9 Click OK.

Next, add a Cisco AV pair.

Step 10 Click the link Show Available AV Pair Attributes to Add.

Step 11 In the Vendors list, click Cisco.

Step 12 In the Attributes list click AVPair and then click the Add button.

This populates the AV Pairs table above.

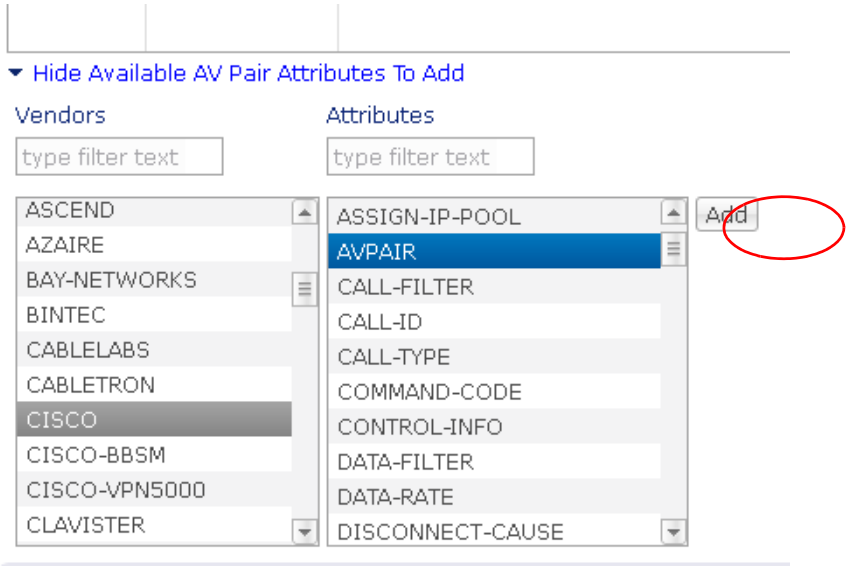

Step 13 Double click in the Value column to enter a QoS value better than the one you noted earlier in step 3.

Step 14 Make this CISCO AVPAIR VALUE be QU;2000000;D;2000000.

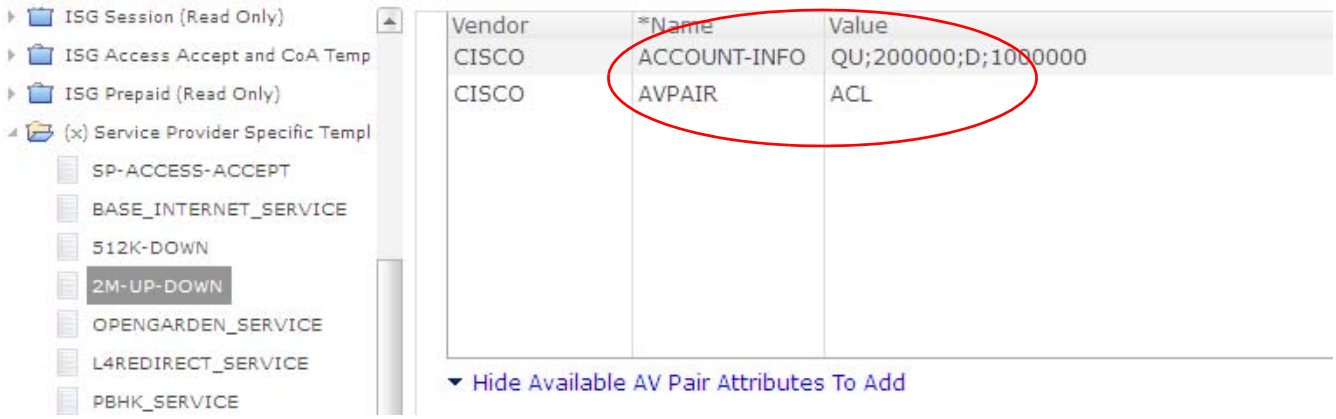

This is your service template for 2 MB up and 2 MB down service that you can use in an upgraded service.

Step 15 Go to 2—Create the Upgrade Service Option.

2—Create the Upgrade Service Option

These steps create the service options you need for an upgraded service, that is, a tiered service.

Step 1 Click Services tab > Services node > Service Options folder > ISG Base Service > Service Option link.

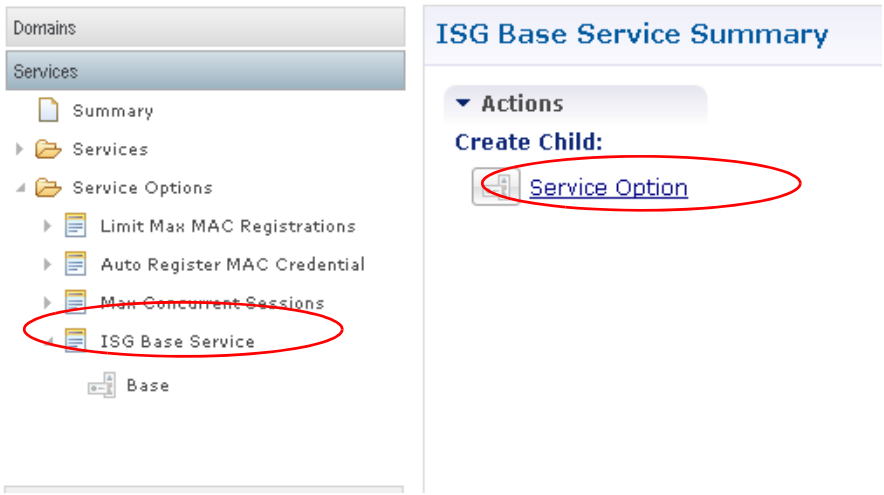

Step 2 Click any of the Service Configurations listed and then click OK to display the Service Option screen.

Step 3 Fill in the Service Option screen as shown in the figures below.

Step 4 Provide the Name field as Upgrade.

Note that this service option is based on the ISG Base Service use case template.

Step 5 In th object list, open the item Service Provider Specific Templates and select the item 2MUP-DOWN previously configured in 1—Create a Service Template for Upgraded Service.

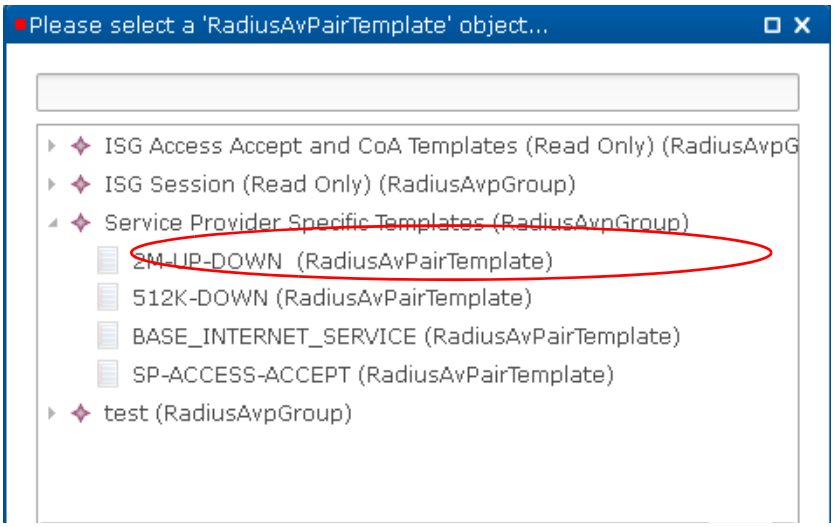

The Base ISG Service service configuration for the Upgrade service option looks like this:

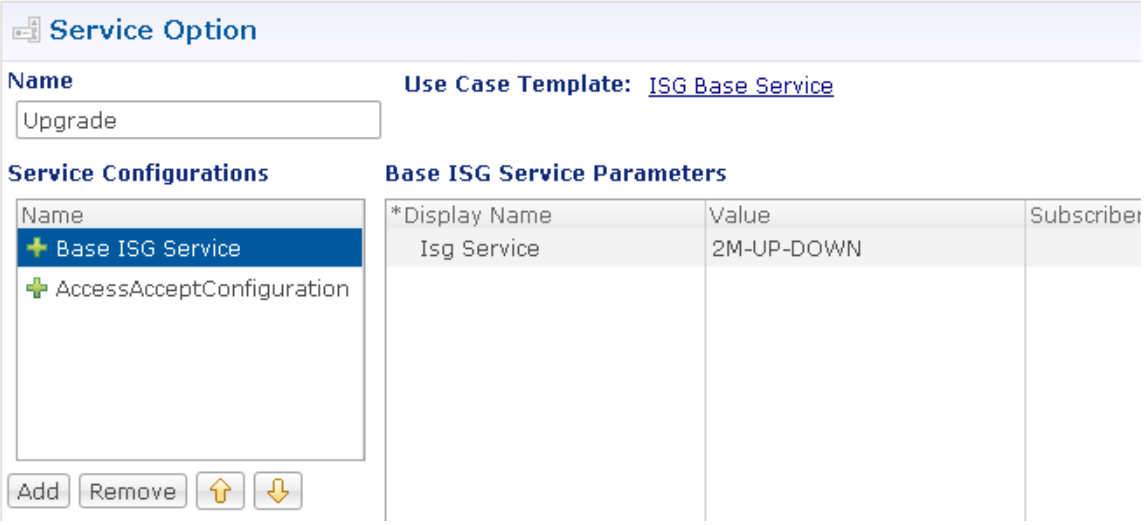

Step 6 In the Service Configurations list, select the AccessAcceptConfiguration to review it, but you make no changes for this service option.

The AccessAcceptConfiguration for the Upgrade service option looks as it did before, like this:

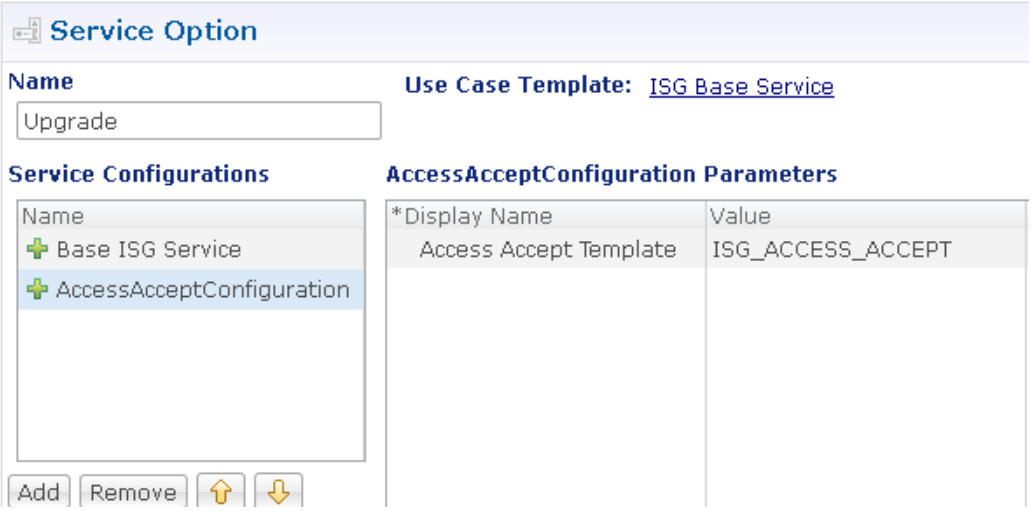

Step 7 Click Services tab > Services node > Service Options folder > Limit Max MAC Registrations.

Note that there is a 10 device maximum service option available in the tree, but for an upgraded service you would like to provide even more device access, that is, raise the maximum limit for a premium subscriber.

Step 8 Click the Service Option link in the main window.

Step 9 In the Service Configurations list select the MAC Limit item and click OK.

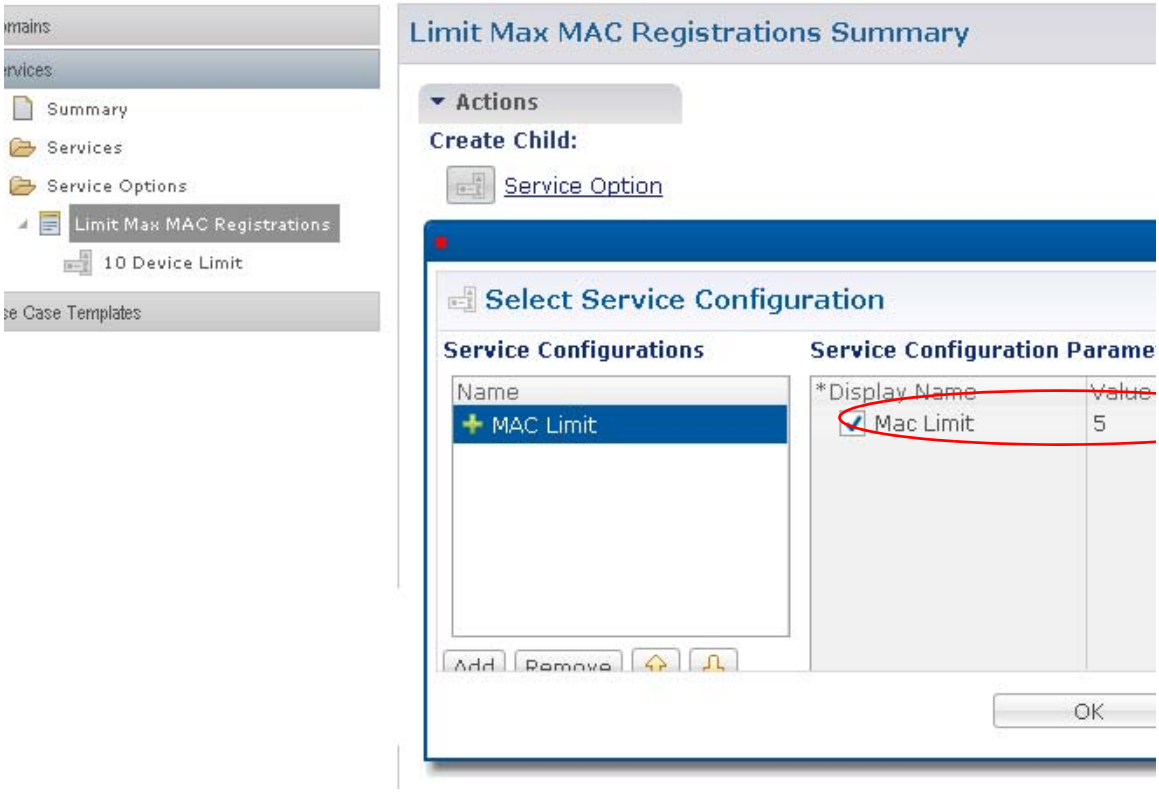

Step 10 On the Service Option screen, provide the Name as 20 Device Limit.

Step 11 Select MAC Limit in the service Configurations list.

Step 12 Click in the Value column and change the value from 5 to 20.

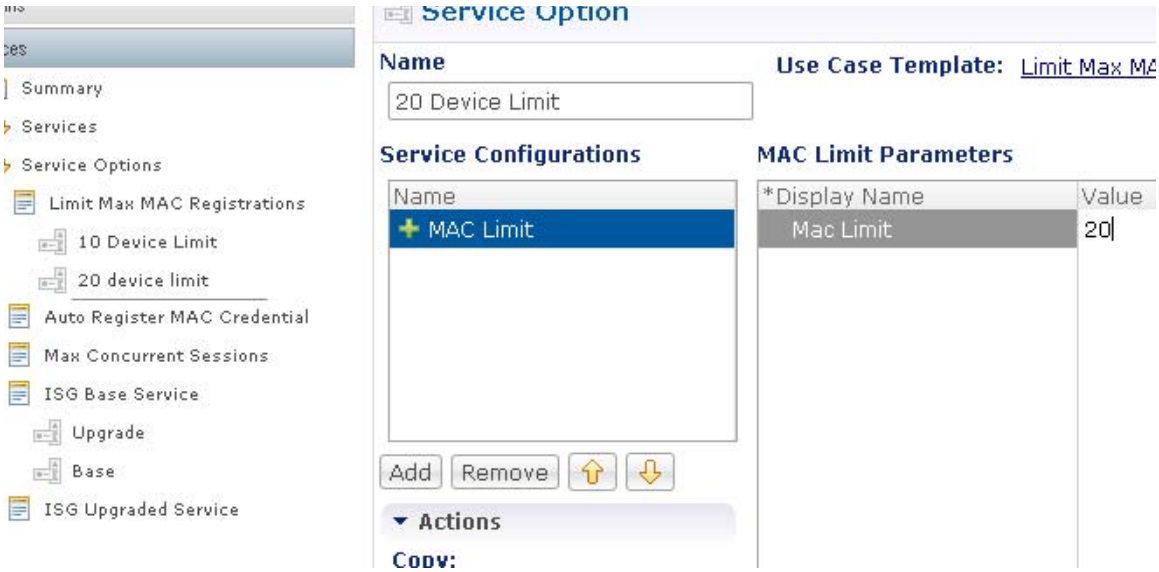

Step 13 Go to 3—Create an Upgraded Service.

3—Create an Upgraded Service

These steps create the upgraded service itself.

Step 1 Click Services tab > Services node > Summary > Service link.

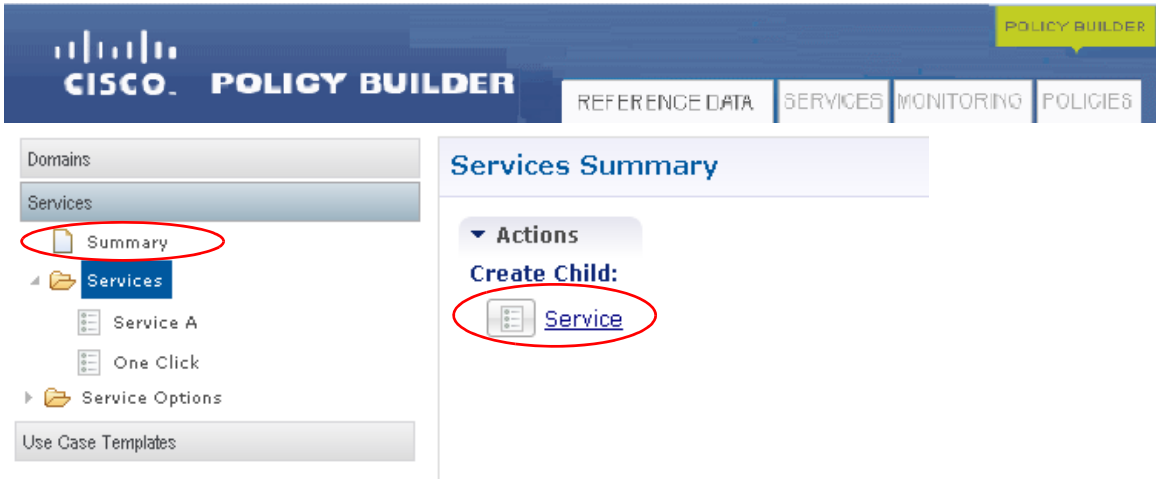

Fill in the Services screen to create the higher tier of service.

Step 2 Provide the Code field as TIERED-SERVICE.

Step 3 Provide the Name as Tiered Service.

Step 4 Select the Enable the check box.

Step 5 Click the Add button to begin to add the service options.

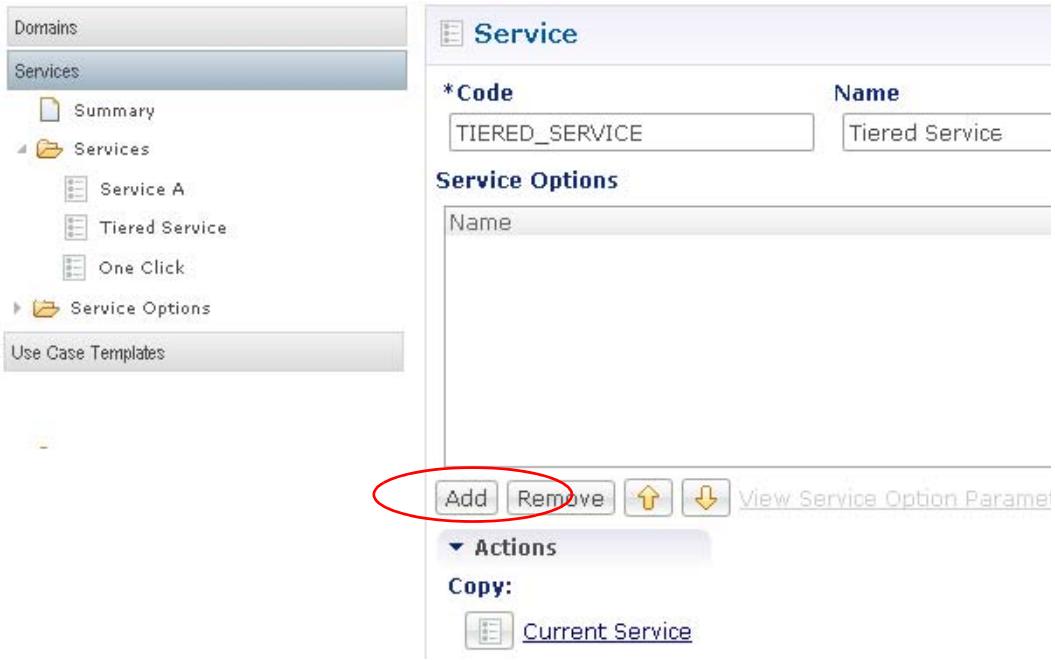

Step 6 The first service option is the 7 day limit on MAC address. Select that and click OK.

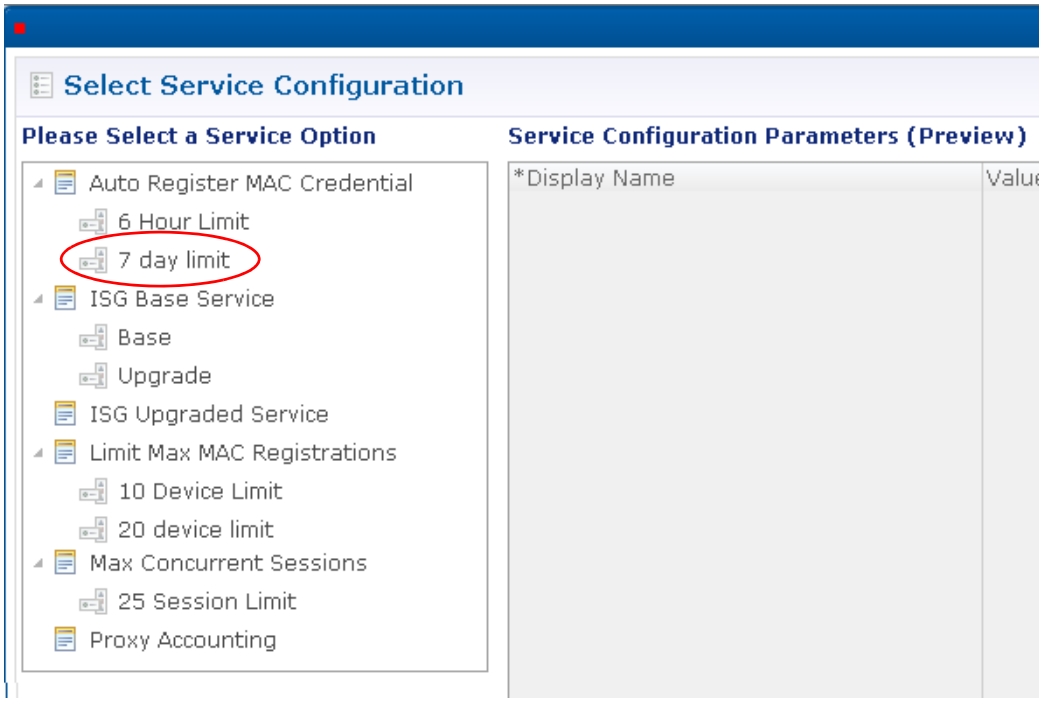

Step 7 Click Add again and select the service option ISG Base Service > Upgrade (affects the QoS, defined earlier).

Step 8 Click Add one more time and select 20 Device Limit.

Step 9 The new service named Tiered Service should look like this:

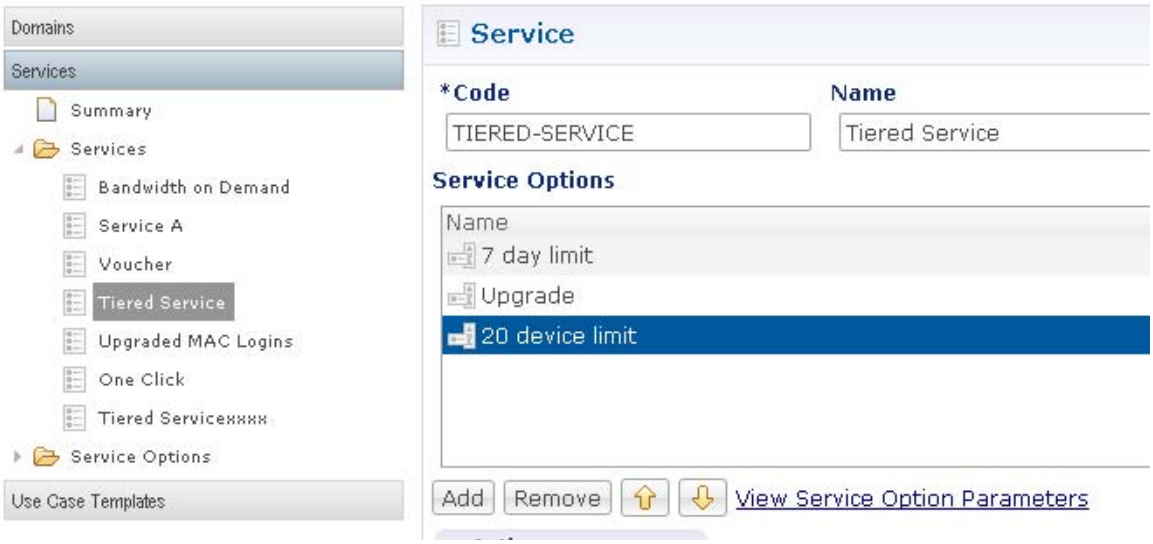

This service is now available for a subscriber to purchase.

Step 10 Check your work with the steps described in the section Test.

Test

Test your configuration with these steps:

Step 1 Log on as a test subscriber on the customer portal.

Step 2 Purchase new service.

Step 3 Make sure the new service installs on the PEP device.

Voucher-based Services

A service that is voucher-based means that the subscriber must have a voucher code and a voucher PIN to sign in. This voucher login is then restricted by things like calendar days, session limits, or device login limits. With a voucher service, there is the ability to develop custom portal web pages, perform redirects, and to provide subscriber assistance by reissuing PINs and time extensions.

Voucher-based services can manage situations such as these:

- One-click voucher
  - The subscriber clicks once on a check box to accept terms and conditions and is then presented with your web page.
  - Untimely guest scenario[1] allows only a two hour session for a 24 hour period, starting with the access time. This usage is tied to one-click authorization and is tracked by MAC.
- Basic voucher
  - Voucher code - the network guest enters a code specifying 1 day of service, 1 week of service, et cetera.
- Convention voucher
  - Perhaps 20 people receive access for 5 days or for 2 hours per session, for example. When the convention is over, no further access is permitted.

In addition, you can regulate vouchers by either of two methods:

- Use the voucher database, accessing that with an API. See Configuring a Voucher-based Service for an example of this more versatile and powerful type of regulation.
- Use a domain to restrict access via a one-click option. This simple and limited method is for setting up a voucher service that is valid for a limited amount of time. All subscribers get the same voucher with this method. See Configuring a Time-based One-click Voucher Service.

Configuring a Voucher-based Service

This example procedure defines a voucher-based service that uses an ISG prepaid service, and relies on the voucher database while calling out other voucher options. Steps also show how to make this configuration differ for non quota-based services as well.

1. A person who habitually takes advantage of others for computer privileges.

Steps

Service Template

Step 1 Create a RADIUS Service Template that provides the name, QoS, and AVPairs as shown below. You uses this template later on page 34.

See the steps Required Service Configuration for specific details on how to create a RADIUS Service Template.

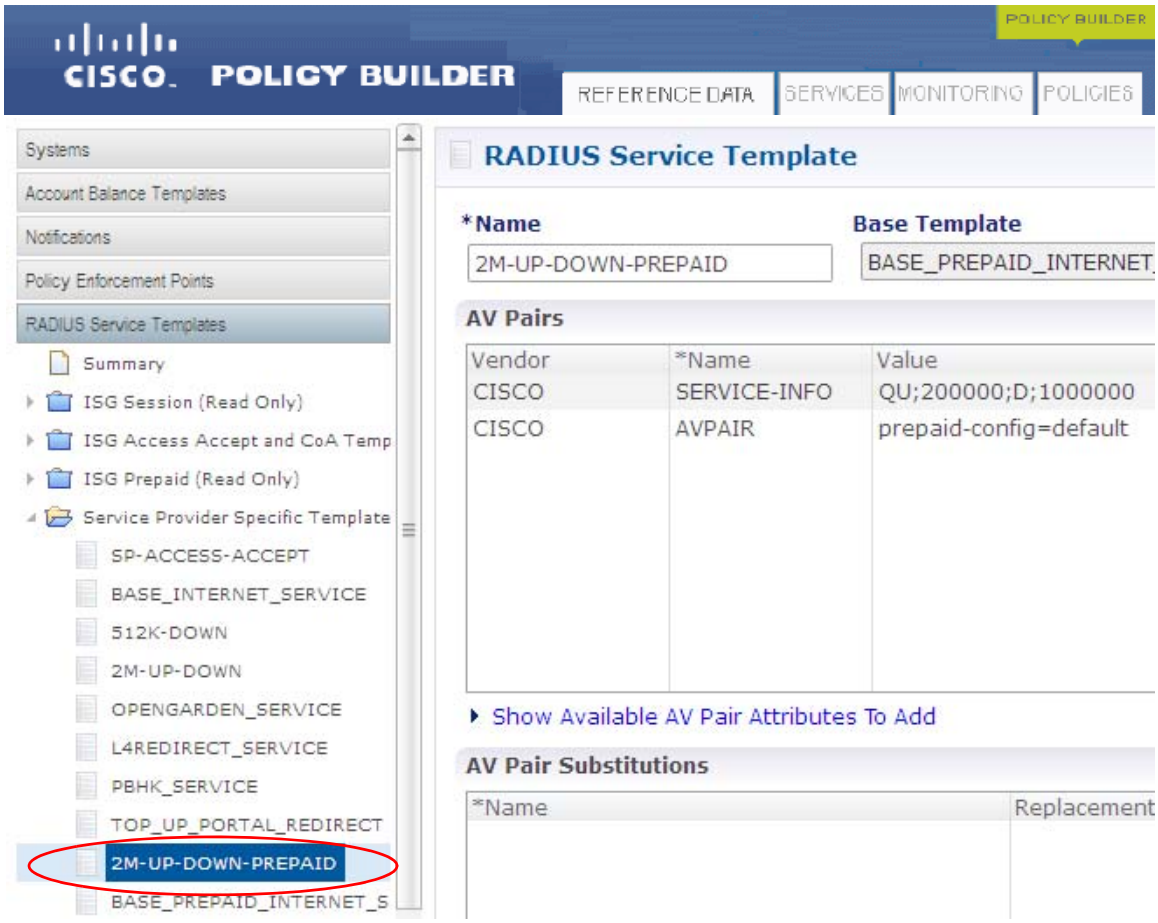

Service Option

Note Next, this procedure creates service option definition for a prepaid data service. This procedure uses that service option to define the quota service element of a voucher-based service.

In some deployments, steps for the Prepaid Data Service service option may have already been performed, but for completeness we include the steps here.

If you do not need to configure a service option for a voucher service, jump to step 11.

Step 2 Click Services tab > Services node > Service Options folder.

Step 3 Click the Prepaid Data Service service option.

Step 4 Click the Service Option link the main window.

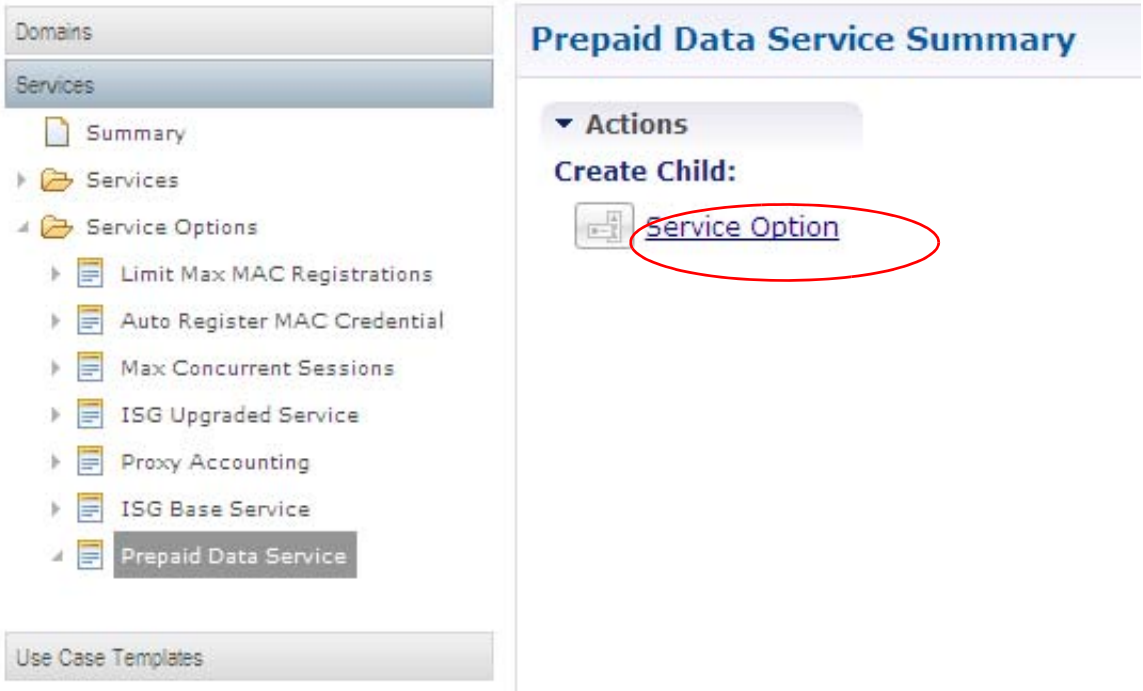

Step 5 Click OK to accept the Service Configuration that appear.

Your Voucher service is based on these parameters. Click OK to move on.

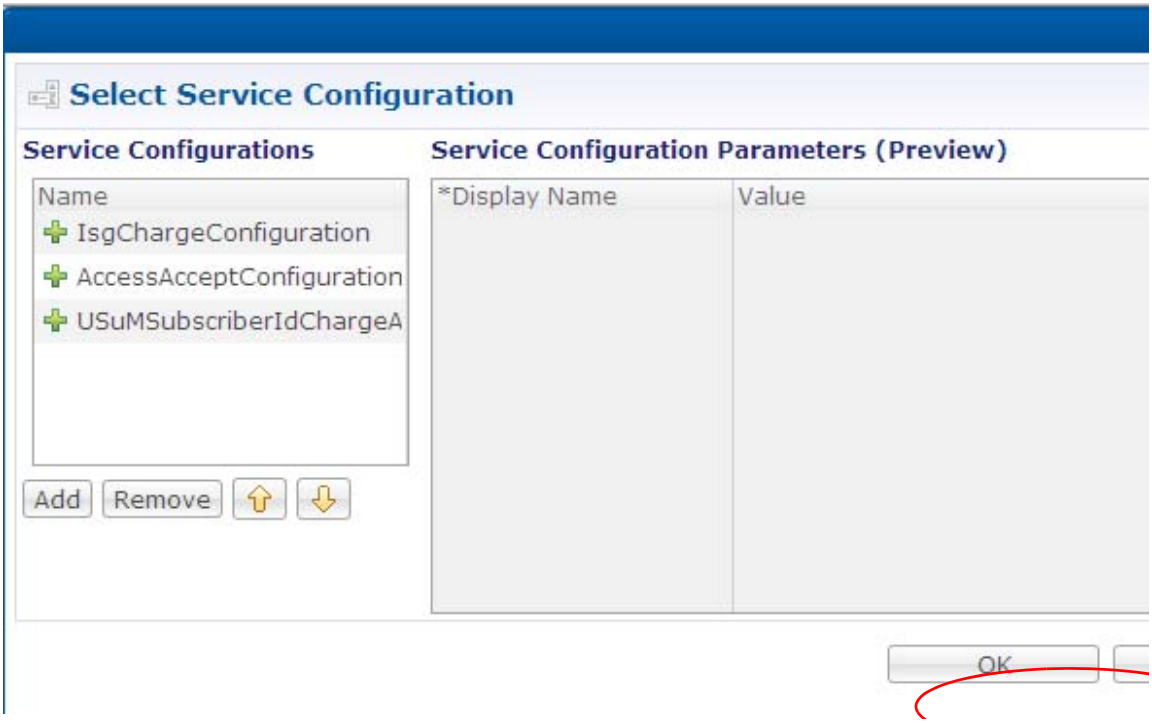

Step 6 Provide the Name as 2M-UP-DOWN-Option.

Step 7 Define the ISG Charge Configuration item in the Service Configurations list as shown in the figure.

Note that the base use case template for this service option is Prepaid Data Service.

Use the green plus symbol to remove items from the Service Configurations list. Use the Remove button to remove items from the IsgChareConfiguration Parameters list.

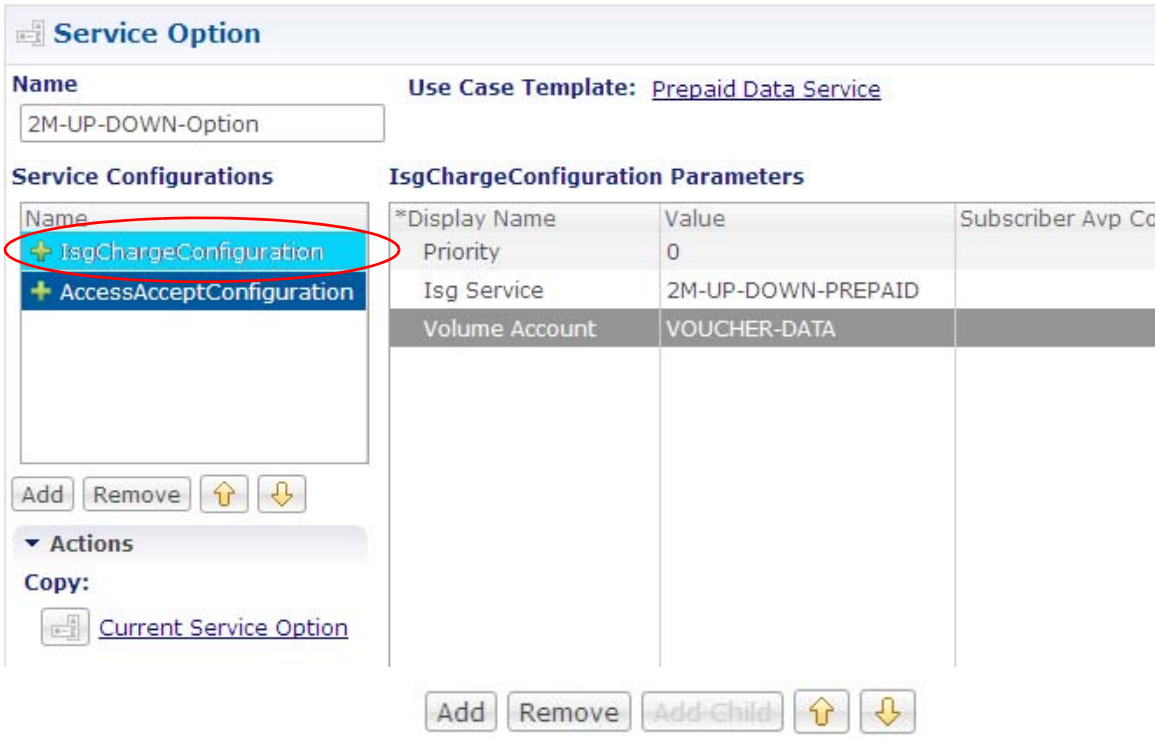

Note Be sure that the ISG Service name must map to a RADIUS Service Template defined on the Reference Data tab, under RADIUS Service Templates > Service Provider Specific Templates, as shown below.

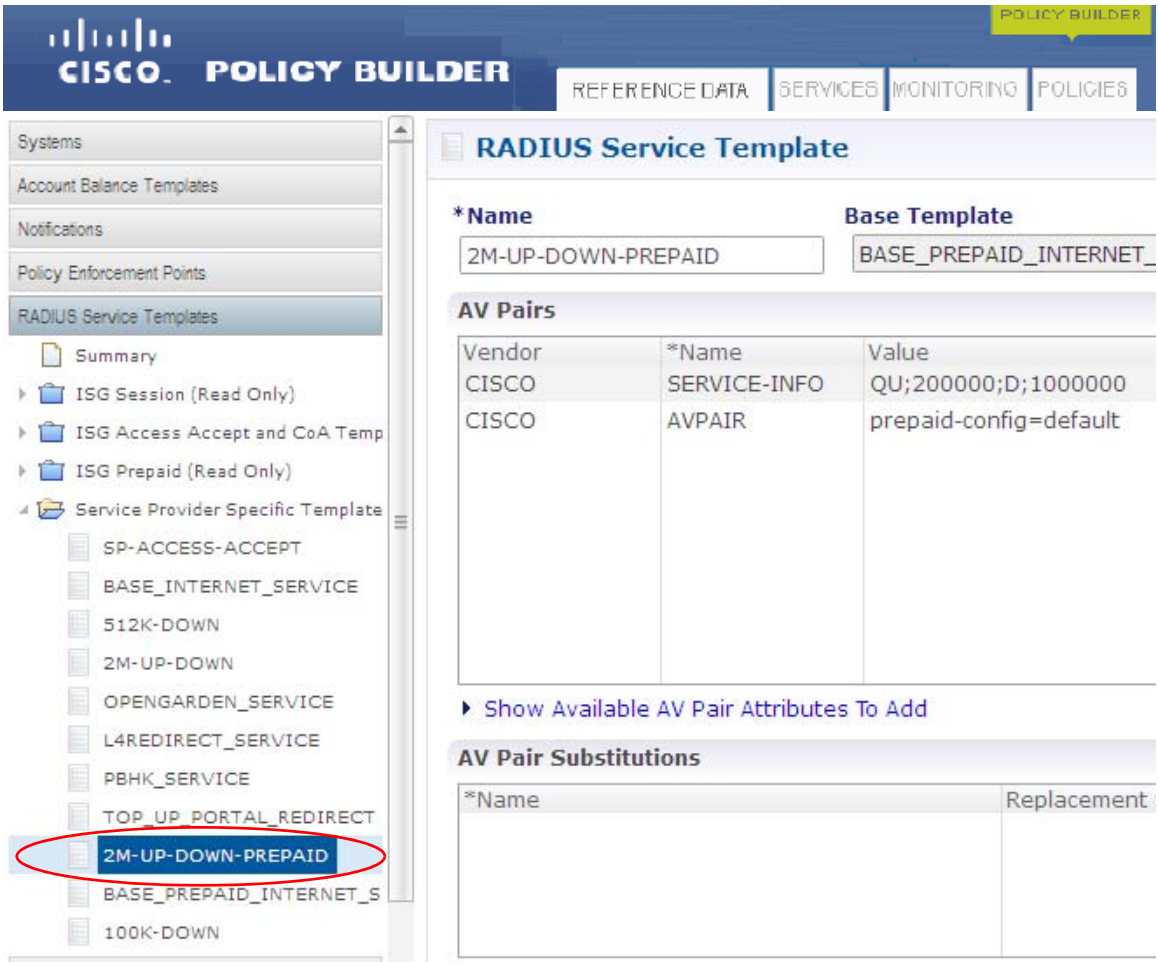

Step 8 Define the AccessAcceptConfiguration item as shown in the next figure.

Step 9 Click the drop-down list icon to display the object list.

Step 10 Your access accept value should be SP-ACCESSS-ACCEPT from the Service Provider Specific Templates object.

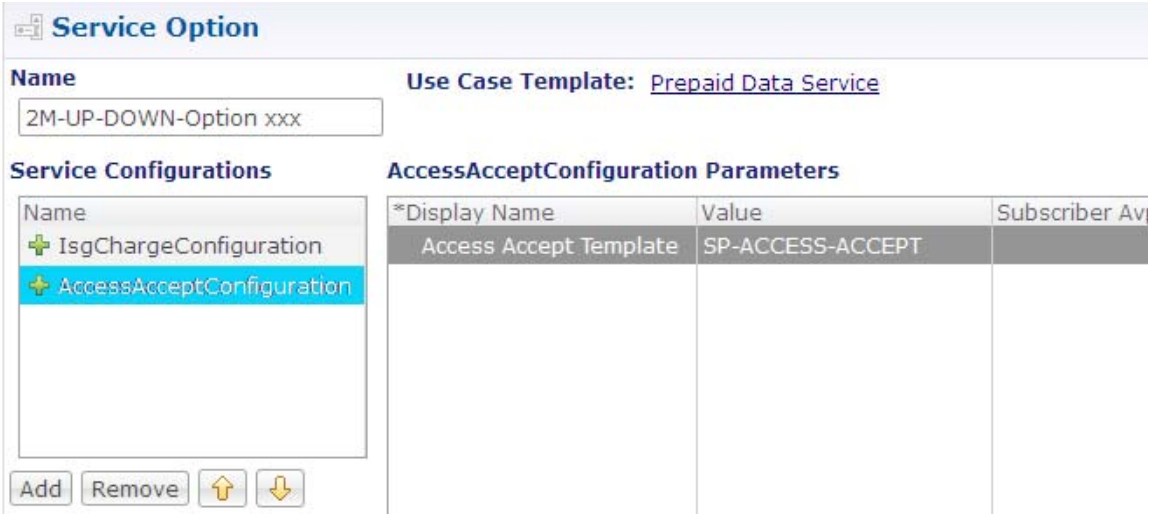

Service

Next, define a voucher-based service that uses this service option.

Step 11 Click Services tab > Services node > Services folder, and open the folder to show its items.

Step 12 Click the Services folder again, and click the Service link in the main window.

Step 13 In the Service screen, provide both the Code and Name as Prepaid Data.

Step 14 Click the Add button, and from the screen that appears, select the 2M-UP-DOWN-Option you created.

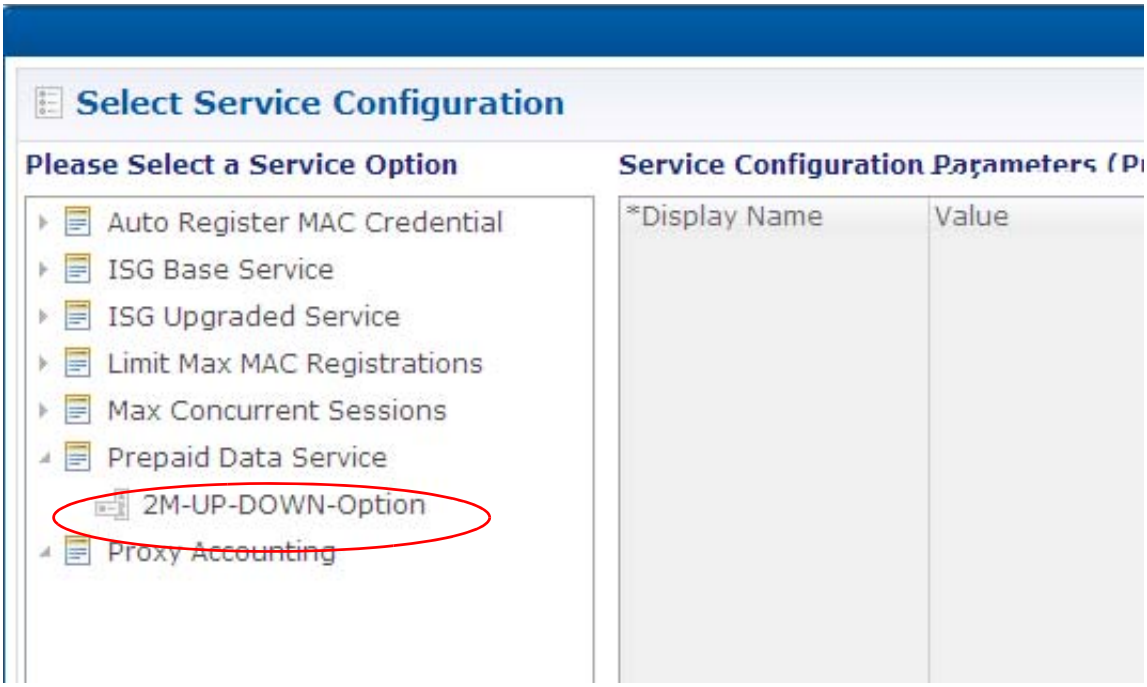

Step 15 Your Service screen looks like this:

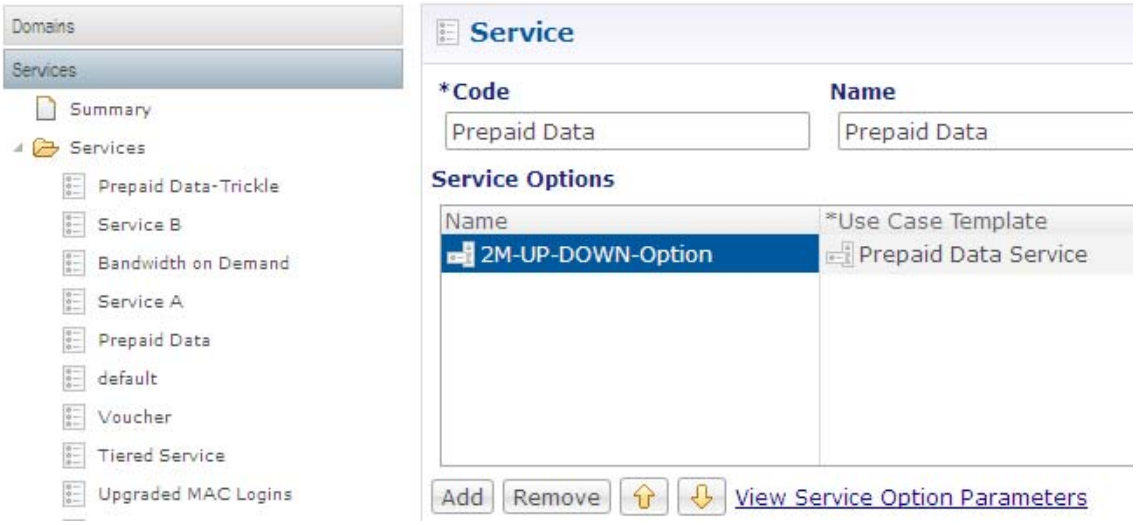

Note If you want to place further restrictions or options on the new voucher service, specify these restrictions by adding service options.

Also, if you do not use prepaid or other quota-based options, instead of specifying this quota-based service option, select other options such as a Time-Limited or Session-Limited service option. The screen below shows a service called VOUCHER that uses three service options to restrict access by voucher, the Base service option and two more restrictive service options.

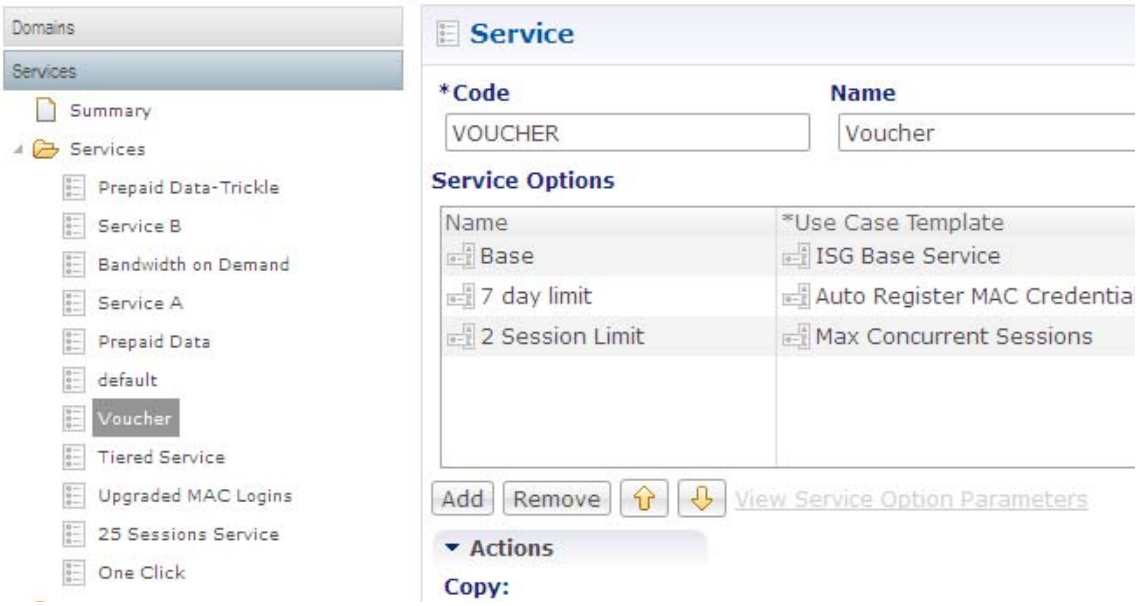

Note The Service Code field must map to the service code defined in the API call to provision the voucher in the voucher database. Refer to the Voucher API documentation for more information about how this operation is performed. However, this provisioning may not require an API call and may be handled by a portal or other web site interface.

Next, define the domain that voucher subscribers uses.

Domain

The service to be applied is not defined in the domain, rather, it is tied to the service defined on the voucher at voucher provision time and then stored in the voucher database. The domain relates the voucher to the subnet where the subscriber logs on.

Step 1 Click Services tab > Domains node.

General Subtab

The General subtab of Domains is discussed in detail at General Subtab.

Step 2 Configure the Domain's General tab to look like this

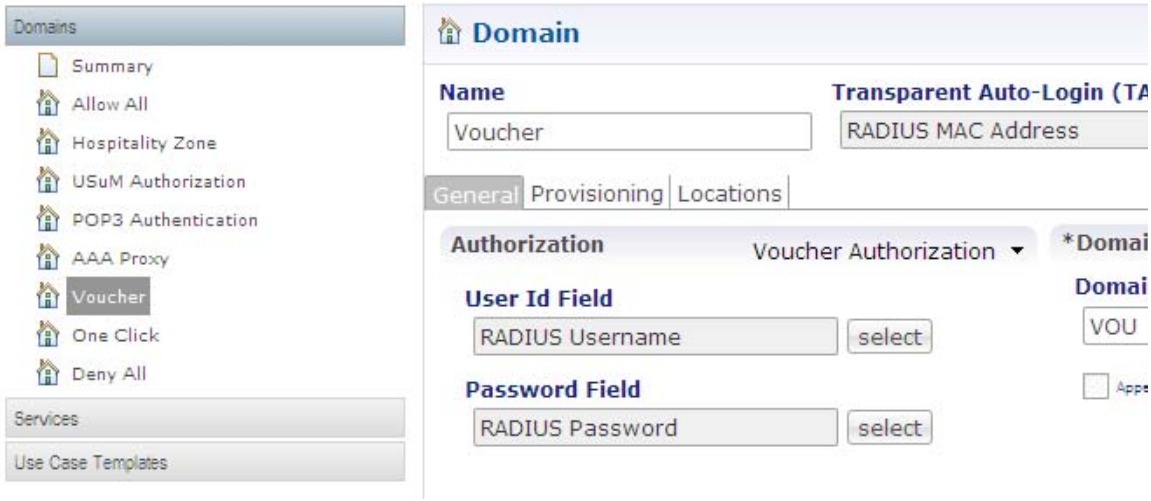

Provisioning Subtab

Step 3 Configure the Provisioning tab of the Domain screen to look like this:

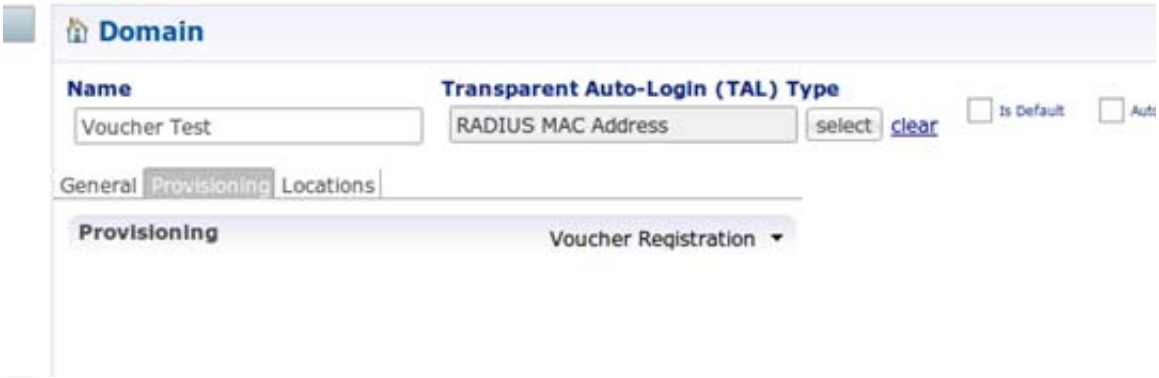

Locations Subtab

Step 4 No specific configuration is required for the Locations subtab. However, you may configure the Locations tab of the Domain screen as needed to limit the domain to particular locations or IP ranges.

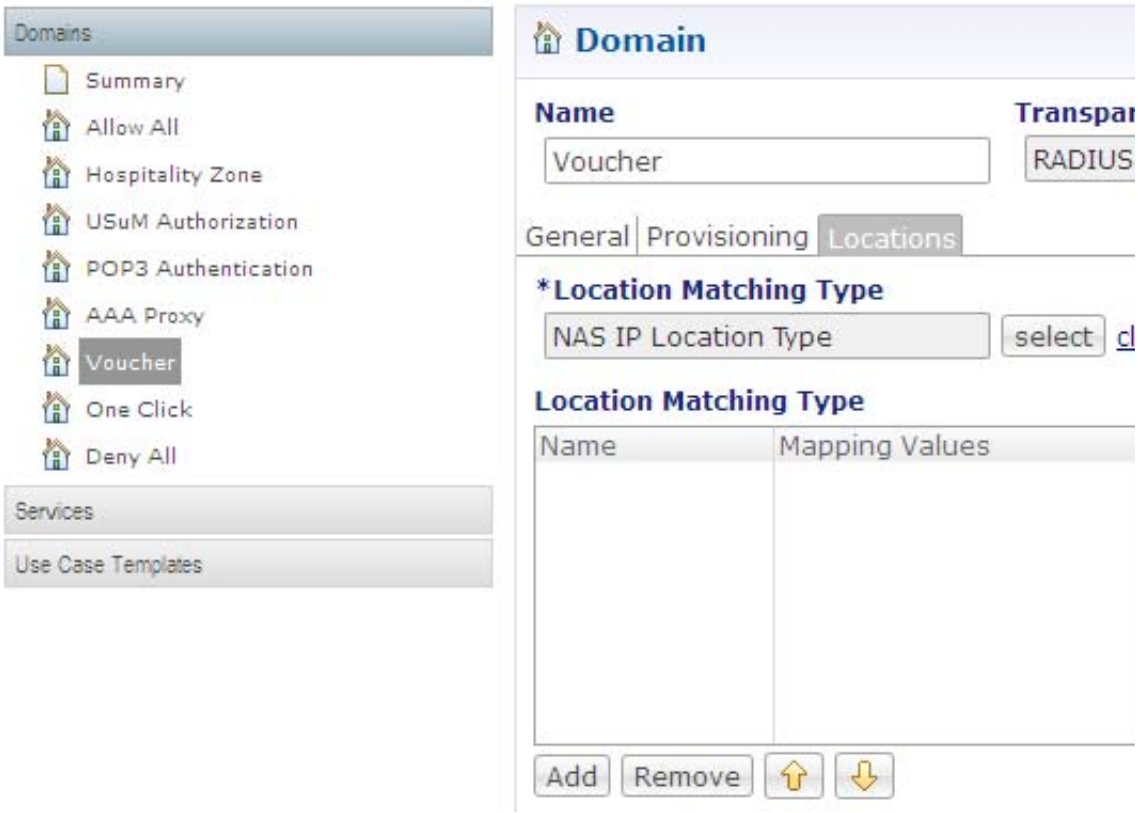

Step 5 Check your work with the steps at Test.

Test

Test your configuration with these steps, working in the subscriber portal:

Step 1 Assign yourself a voucher.

Step 2 Log on to the subscriber portal and use the voucher.

Step 3 Make sure your session requires a RADIUS username and password, and that you are logged out properly.

Step 4 Make sure there is a disconnect after the designated elapsed time.

Configuring a Time-based One-click Voucher Service

This procedure creates a time-limited, one-click voucher service.

This use case grants subscriber access for a specified span of time using a voucher or some other kind of network pass. Time granted is based on elapsed time starting with the instantiation of the pass or voucher. That is, time usage is not interrupted by the subscriber logging off, or disconnecting in any way. This way, you can provide a subscriber access for Monday through Friday, letting them log on and off only during that week.

Steps

1—Check for a BASE Service Option

Make sure you have a time-based service option of BASE for a basic level service option. If you don't have one, see Configuring a Service Option.

2—Create the Time-based Service

Step 1 Click Services tab > Services node > Services Folder > services link.

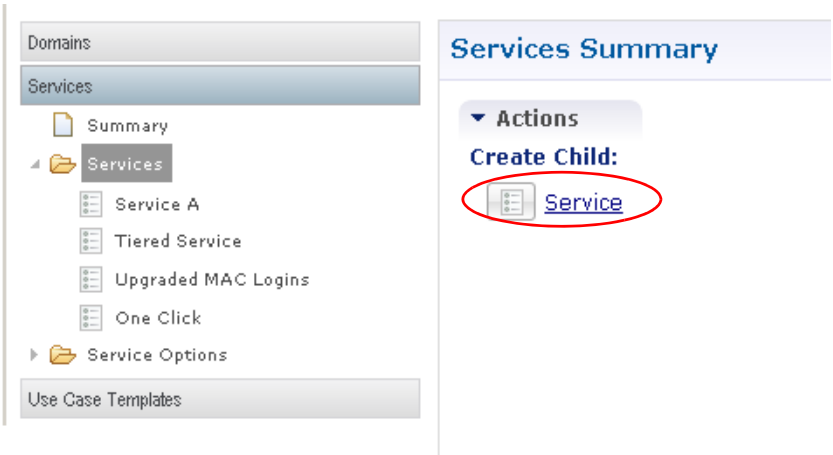

Fill in the Services screen.

Step 2 Provide the Code field as VOUCHER.

Step 3 For the Name field enter Voucher.

Step 4 Select the Enable check box.

Step 5 Click the Add button under the Service Options table to begin populating the service with service options.

Step 6 Select ISG Base Service > Base and click OK.

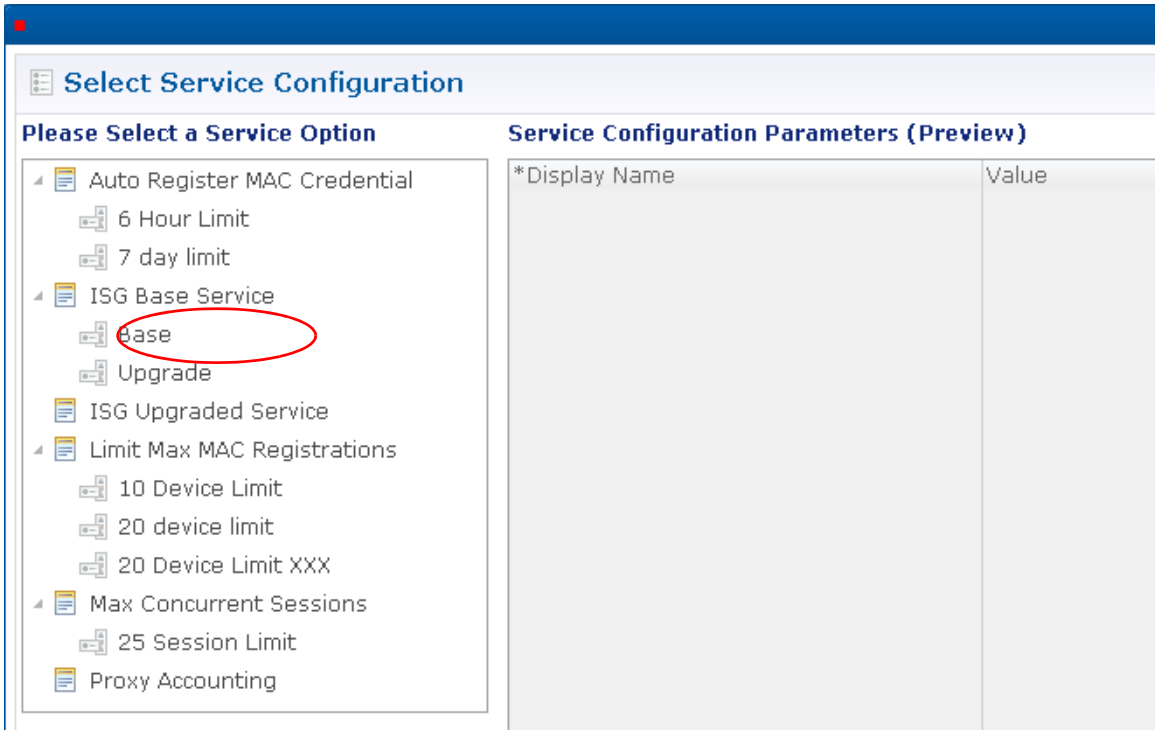

Step 7 Click Add again and select 7 day limit service option.

Now you can give a subscriber with a voucher to the hospitality zone domain a base quality of service and a 7 day limit for access. The subscriber can put their voucher credentials in once and it is valid for the entire time period.

Step 8 Go to 3—Create a Time-based Domain.

3—Create a Time-based Domain

Step 1 Click Services tab > Domains node > Summary > Domain link.

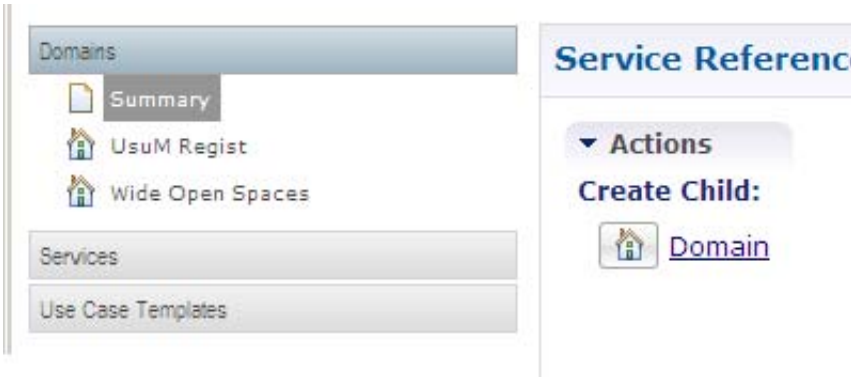

Step 2 On the Domains screen, provide the Name as Hospitality Zone.

General Subtab

Step 3 Click the Authorization drop-down menu arrow and select Voucher Authorization as the method of authorization.

This type of authorization requires a User ID and Password the subscriber finds on his voucher.

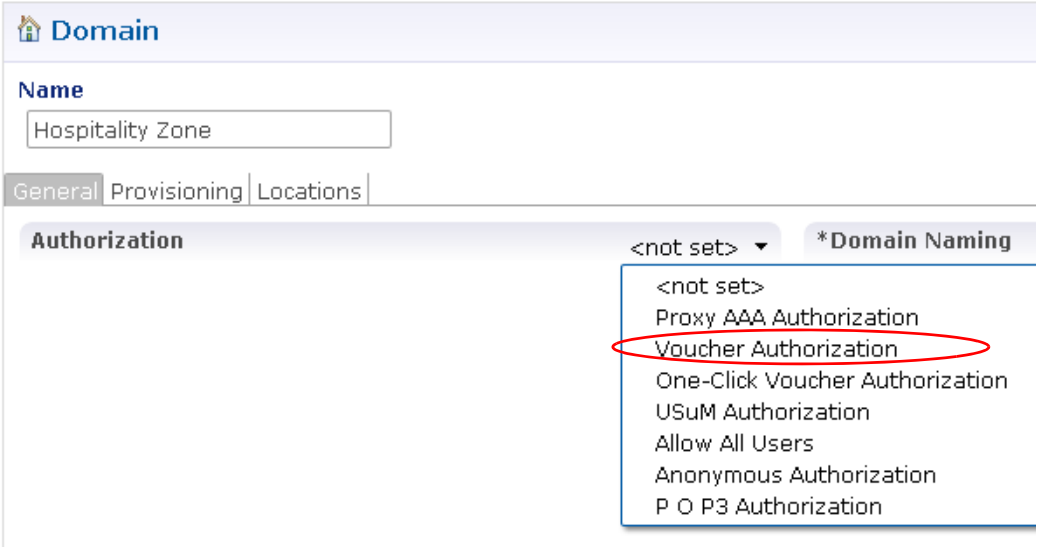

Step 4 Click the select button next to the User Id Field and select RADIUS Username from the object list and click OK.

Step 5 Click the select button next to the Password Field and select RADIUS Password from the object list and click OK.

Step 6 Set the Domain Prefix to hospitality and click the Append Location check box.

Step 7 Fill in the rest of the fields on the General subtab similar to the figure below.

Provisioning Subtab

Step 8 Click the Provisioning subtab.

Step 9 Click the drop-down menu to display the Provisioning items.

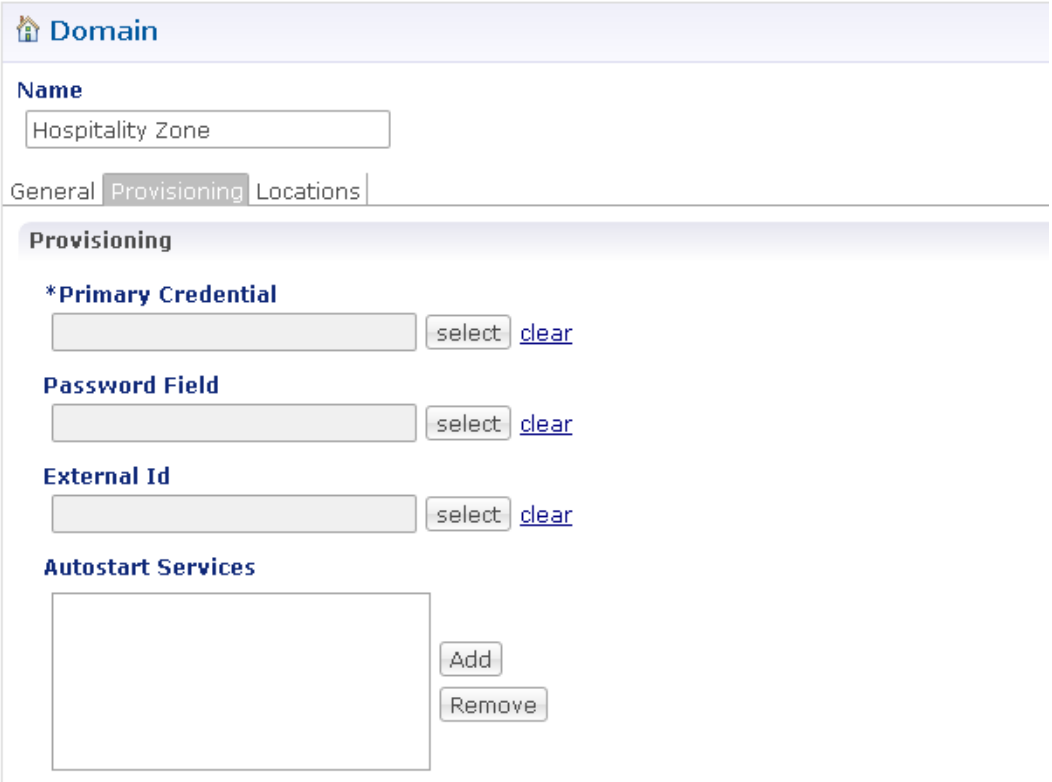

Step 10 Click the select button next to the Primary Credential field and select RADIUS Username (EClass) from the object list.

Step 11 Click the Add button next to the Autostart Services area and select the Voucher service from the list and then click OK.

Your Provisioning subtab should look like the figure below.

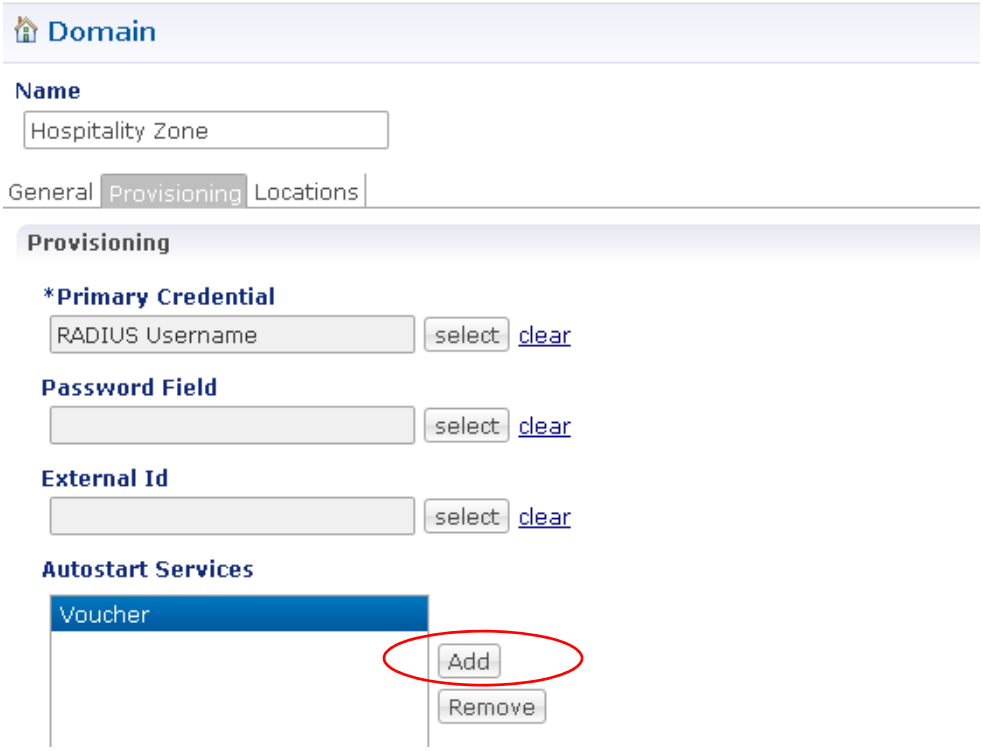

Additional Profile Data Subtab

No configuration is necessary under this subtab.

Locations Subtab

Step 12 Click the Locations subtab.

Step 13 Click Add to add a row to the Location Matching table.

Step 14 Select Framed IP Addresses for the Location Match type.

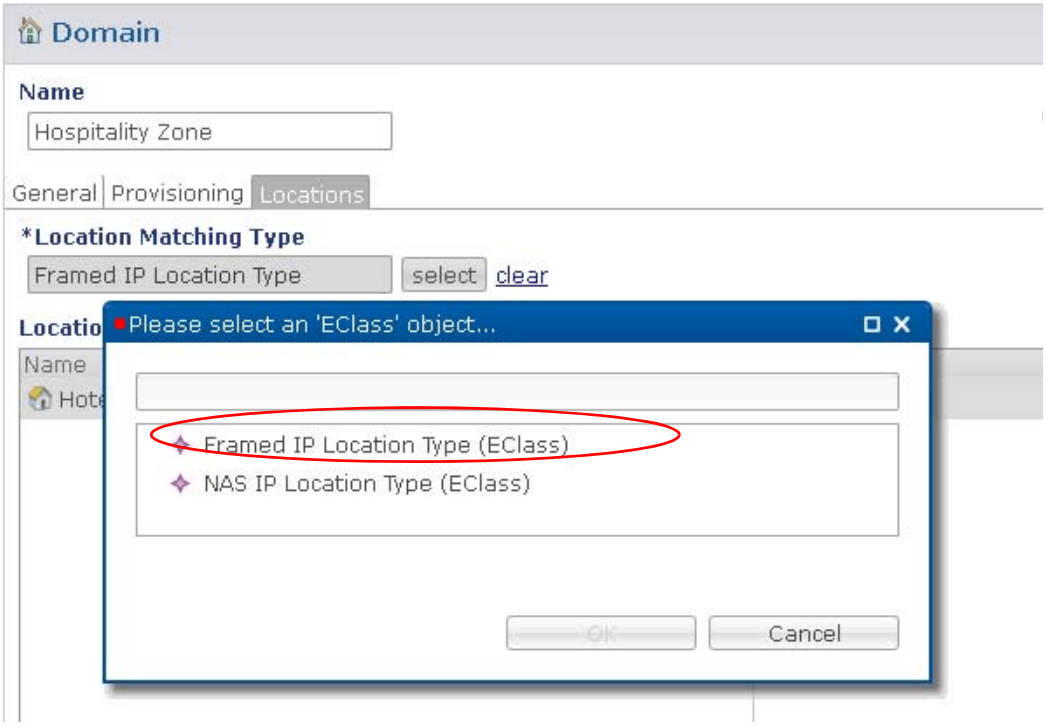

Step 15 Click the Add button to add a row to the Location Matching table.

Step 16 In the Name column, name this location Hotel1.

Step 17 In the Mapping Values column, click the drop-down menu icon to display the Add Values window.

Step 18 Enter an IP address range.

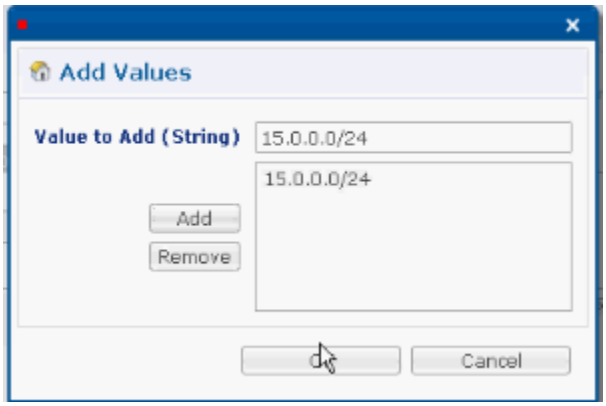

Step 19 Click Add and then click OK.

Your Domain Locations subtab should look like this:

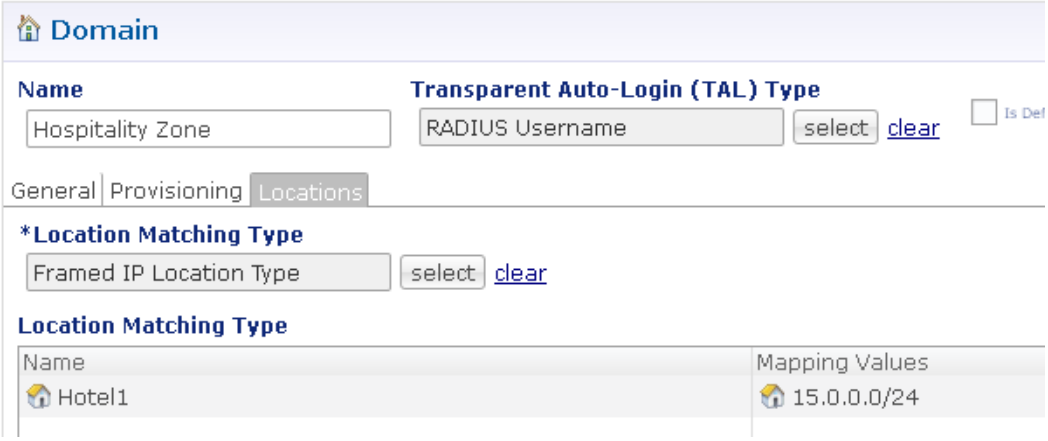

Advanced Rules Subtab

Step 20 Click the select button next to the TAL Type field and select RADIUS Username.

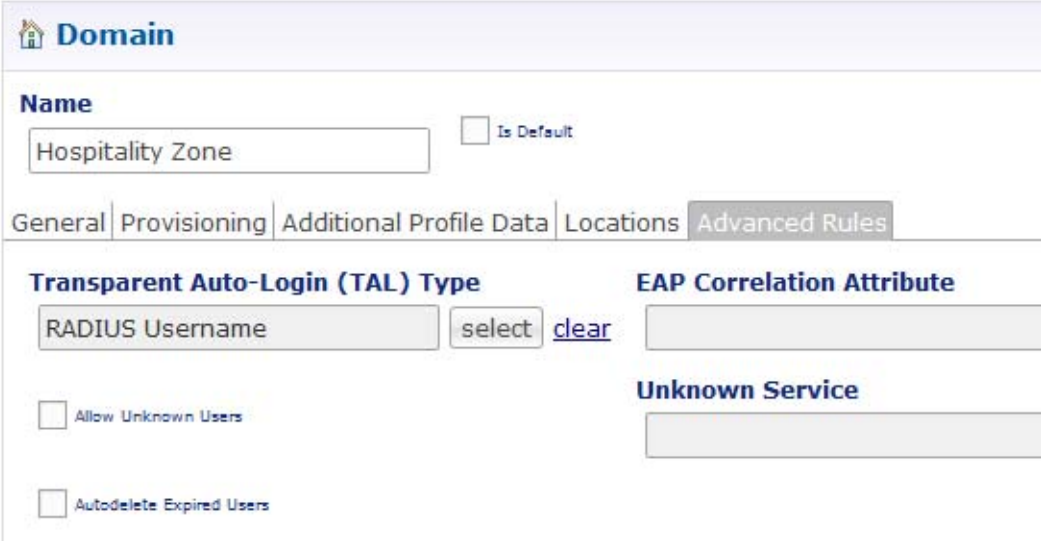

Now, CPS can accept time-based authorizations from users on the domain prefix with specific subnet IP addresses.

Step 21 Check your work with the steps at Test.

Test

Test your domain configuration with these steps, working in the subscriber portal:

Step 1 Purchase new service.

Step 2 Make sure the new service installs properly on the PEP device.

Step 3 Make sure there is a disconnect after the designated elapsed time.

Concurrent Logons Service Option

This procedure creates the Upgraded MAC Logins service which uses the Limit Max MAC Registrations service option.

This service option permits multiple device logons under a single subscriber ID, up to the maximum concurrent logons for the specified service plan.

Usually, a MAC-based authentication limits the number of MACs registered concurrently, but does not prevent users from logging in to the system concurrently.

This addresses the situation when different people use the same user name or other credential to log on. Use this service and service option to make sure a subscriber does not exceed 25 sessions.

Steps

Step 1 Click Services tab > Domains node > select a domain already present.

See Defining a Default Domain if you need to create a domain.

a. On the Domain screen, make sure that one of these authorization schemes is used:

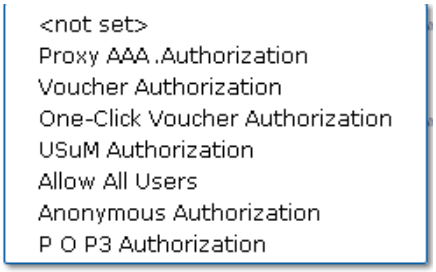

Step 2 Create a service option named 25 Session Limit.

Use the steps at Configuring a Service Option for adequate details.

Step 3 Click Services tab > Services node > Service link.

Step 4 Enter a Code and a Name for this service that shows it limits the number of sessions to 25.

Step 5 Click the Add button and select the service option 25 Session Limit.

Your Service screen should look like this:

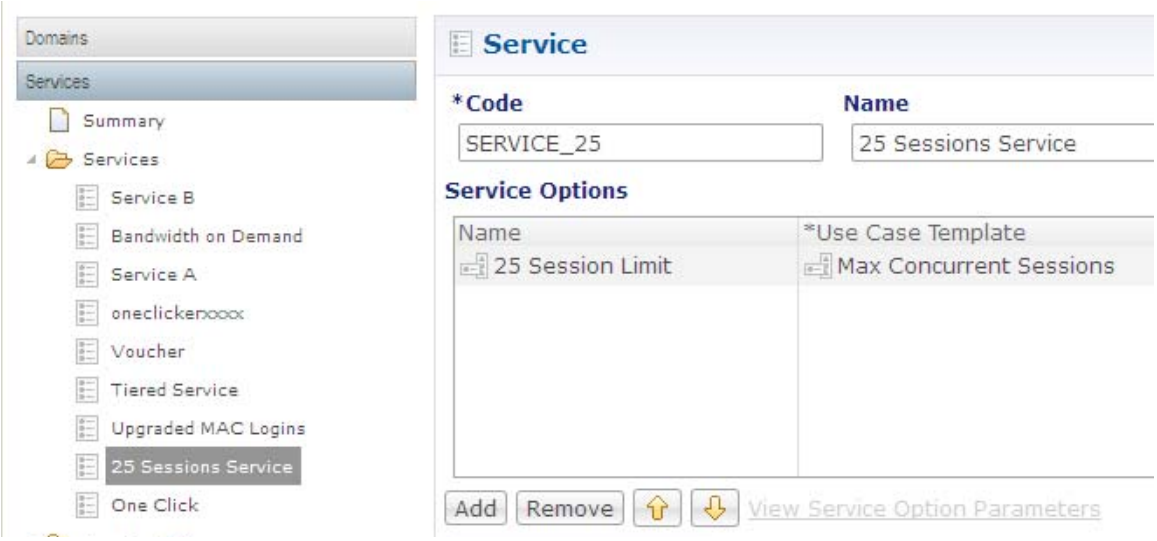

Step 6 Check your configuration with the steps in section Test.

Test

Test your configuration with these steps:

Step 1 Log on with 1 user, log on with 2 user, log on with 25 users.

Step 2 Check for a non-failure when your limit is exceeded.

Bandwidth on Demand Use Case

This service uses the ISG Base Service use case template and adds a premium bandwidth service to it.

In this use case, for a limited and specified period, the subscriber's bandwidth can be increased (or decreased). This limits the higher bandwidth demand to a finite period.

BWoD (bandwidth on demand) may be based on one or more of the following:

- Scheduled defaults for the service plan
- Prescheduled by the subscriber
- On-demand

Bandwidth on demand is similar to the upgraded tiered service. See Tiered Services Use Case.

To create a service that enhances a subscriber's basic service, use the service option Upgrade. For this enhanced service, you first provide the subscriber with a base service, and then add another faster or broader service.

Note The only difference between bandwidth on demand and a tiered service is that tiered service has several options defined, and Bandwidth on Demand has only a premium service *layered* service on top of a basic service. The premium service ceases at a specified time, perhaps the end of the month, when subscriber billing and quotas roll over.

Steps

Step 1 See if you have a service option for Upgrade service. If not, see Tiered Services Use Case.

Step 2 Click Services tab > Services node > Services link.

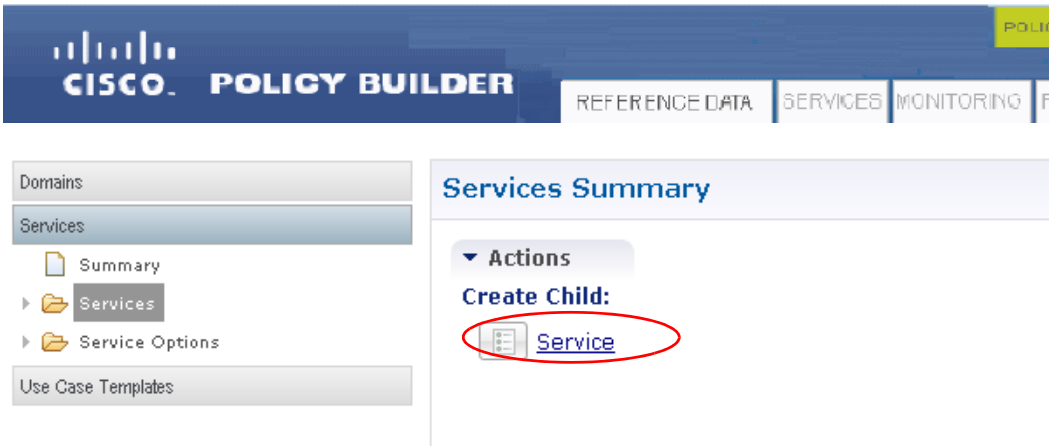

Fill in the Service screen.

Step 3 For the Code field, enter BWOD.

Step 4 For the Name field, enter Bandwidth on Demand.

Step 5 Click the Add button, and from the list of service options, select the ISG Base Service > Upgrade.

Note The steps for creating this service option are found at 2—Create the Upgrade Service Option.

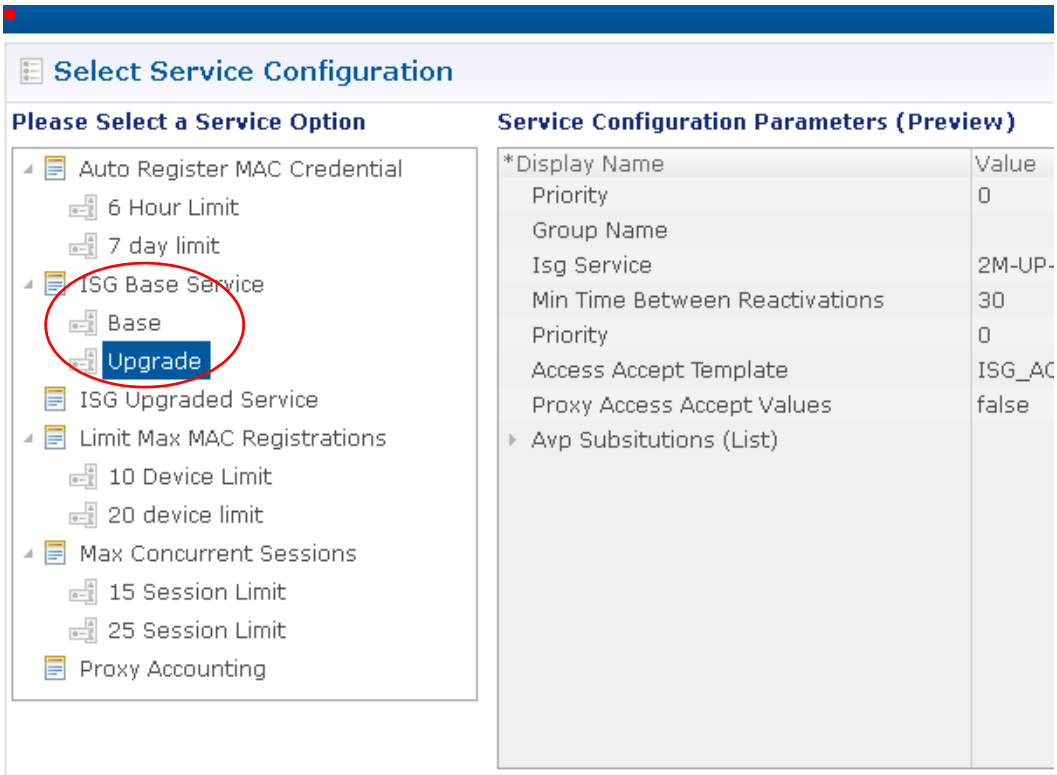

Step 6 Review the main portion of the screen to check that you have the desired parameters.

Note that this example provides 2 MB upload and download speed.

Step 7 Click OK.

Step 8 Make sure your BWoD Service screen looks like this.

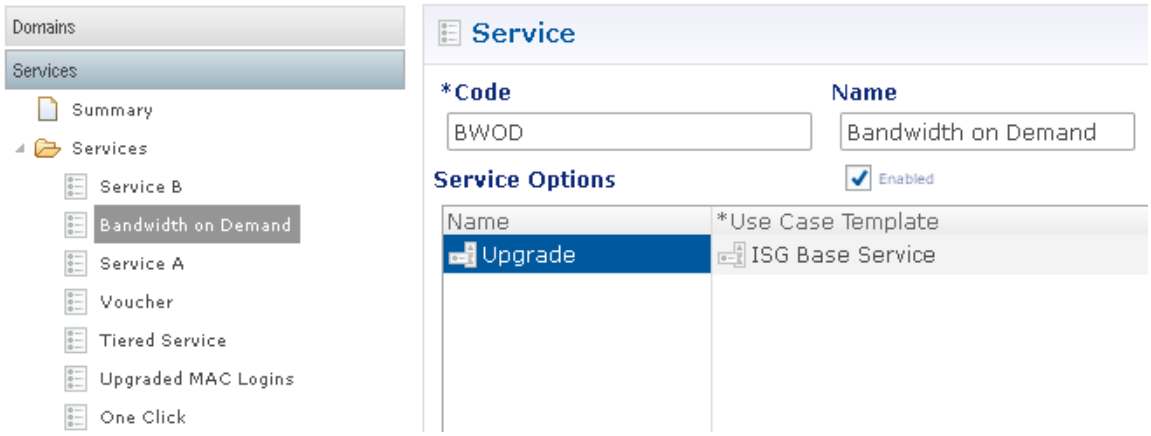

Step 9 Test your service with the steps at Test.

Step 10 Next, go to your subscriber portal to define and manage this service offering there.

Test

Test your configuration with these steps, working in the subscriber portal:

Step 1 log on to the portal as a test subscriber and purchase a new service without any errors.

Note Make sure the new service is installed on the PEP or network device.

Step 2 Consume bandwidth until you exceed 512K.

Step 3 Observe the higher bandwidth take over.

Final Steps

This section shows basic use cases that are almost always wanted. Of course, use cases, service options, and services are developed frequently for many needs.

If you need help going beyond the use cases presented here, please contact your Cisco technical representative.

Test the Configuration

Revised: July 10, 2015

This section shows you how to validate that your configuration tasks are complete.

Validation of the configuration and services involves looking at different components from both a visual level as well as using the command line interface. Because of the complexity of the system, it is necessary to start at the high level view of the system and then pursue the task at a more detailed level.

This chapter covers the following sections:

Validating the Configuration

For configuration validation, the main configuration tools are used to check what is actually configured. Two GUI tools and the command line is available for this task.

- Cisco Policy Builder GUI
- Cisco Subscriber Services Portal GUI
- The command line interface (CLI) tools are all located on Control Center 01 and Control Center 02 and are accessed via the qns user over SSH.

Cisco Policy Builder GUI

Cisco Policy Builder is the main configuration tool used for the policy engine.

Open a browser and enter the correct URL to the Cisco Policy Builder.

```
<IP_address>:7070/pb
```

Display the main page by clicking OK on the **Choose Policy Builder data repository** page.

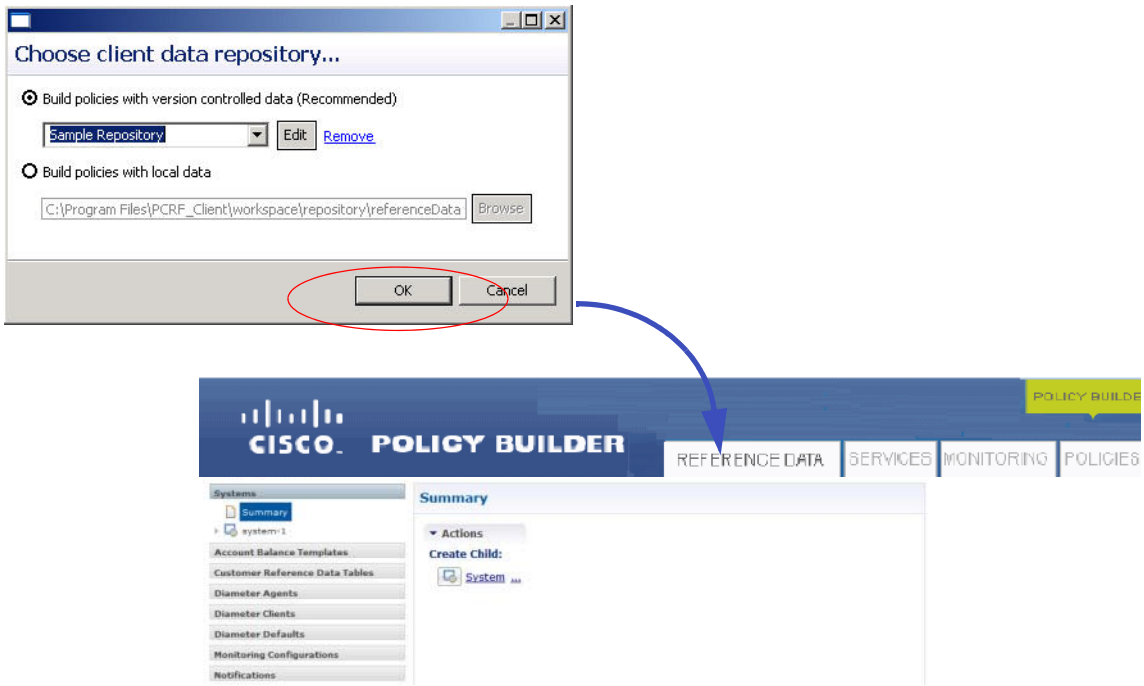

Step 1 First, page through this interface and look for any red Xs in any of the Services, Policies, or Reference Data tabs. The figure shows two examples of error screens, one for the tree and one in a screen:

Any errors denoted with a red X typically cause the polices not to load and so not allow any of the new configuration changes to be deployed.

Step 2 If there are no errors in the configuration, confirm that the changes have been applied by going to the Control Center CLI, and running this command:

```
svn log http://pcrfclient01/repos/run/ --limit 1
```

It may prompt you for a username and password. Enter the same user and password that is used for Publishing from the Cisco Policy Builder.

It then tells you the last time a configuration change has been made to the Cisco Policy Builder.

Here is an example.

```
r283 | broadhop | 2012-05-10 09:03:30 -0600 (Thu, 10 May 2012) | 1 line

Added new IPs for new ISG in region 1. 192.168.181.30 and 10.10.12.11
```

Step 3 If the configurations are published properly, check that they loaded in to the CPS Policy Engine (processing node).

a. In the Control Center CLI, go to /var/log/broadhop/ and run

```
grep "Policies successfully configured" consolidated-qns.log
```

- You should see a message like this upon the last time the policies were published, or if the qns processing node was restarted:

```
2012-05-10 09:04:27,769 [pool-9-thread-1] INFO
   c.b.policy.impl.PolicyConfiguration - 1000:Policies successfully
   configured
```

- If not found, check when the configuration change was made using the svn command above to determine what log it should be in. In some cases, the logs may have already rolled over and this message is not found.

In that case, every 5 minutes these type of messages should be printed out in the consolidated-qns.log as well. These show the state of the system.

```
2012-05-10 09:04:27,785 [pool-9-thread-1] INFO
   c.b.s.w.u.USuMConfigurationManager - successfully initialized the
   usum portal

2012-05-10 09:05:00,066 [pool-10-thread-1] DEBUG
   c.b.d.impl.DiagnosticController - Starting diagnostics

2012-05-10 09:05:00,067 [pool-10-thread-1] DEBUG
   c.b.radius.impl.RadiusDiagnostics - Running radius diagnositics

2012-05-10 09:05:00,158 [pool-10-thread-1] DEBUG
   c.b.d.impl.DiagnosticController - Ending diagnostics

2012-05-10 09:05:00,158 [pool-10-thread-1] INFO
   c.b.d.impl.DiagnosticController - Diagnostic test passed -->
   Policy: 6:Last policy configuration was successful

2012-05-10 09:05:00,159 [pool-10-thread-1] INFO
   c.b.d.impl.DiagnosticController - Diagnostic test passed -->
   Expiration Manager: 8:Current node is expiration manager

2012-05-10 09:05:00,159 [pool-10-thread-1] INFO
   c.b.d.impl.DiagnosticController - Diagnostic test passed -->
   Session Datastore: 7:lab:27017 session management node is enabled

2012-05-10 09:05:00,159 [pool-10-thread-1] INFO
   c.b.d.impl.DiagnosticController - Diagnostic test passed --> Common
   Services: 1:QNS server is alive

2012-05-10 09:05:00,159 [pool-10-thread-1] WARN
   c.b.d.impl.DiagnosticController - Diagnostic is abnormal. A problem
   may exist with the system --> Common Services: 4:Session creation
   allowed. 100.0 % used
```

```
2012-05-10 09:05:00,160 [pool-10-thread-1] INFO
   c.b.d.impl.DiagnosticController - Diagnostic test passed --> Common
   Services: 2:Memcached server is operational

The "Policy: 6:Last policy configuration was successful" section shows
   that the last time changes were made to the configuration, they were
   successful.
```

If this is not found then the system may not be running properly and needs to be diagnosed with help from your Cisco technical representative.

Subscriber Services Portal GUI

To access the administrator-facing screens for the portal at admin user level, open a browser and use this URL:

http://www<xx.xx.xx.xx>/users/login>

replace *xx.xx.xx.xx* with the IP address of lbvip01 from /etc/host.

Enter a username and password to log in as admin.

The SSP GUI is the main tool to configure the functionality of the subscriber services portal. This includes setting up locations, the text content, and the behavior of the portal.

Note Because the SSP GUI immediately posts the information into the policy engine, these configuration changes become effective as soon as you save them from this GUI.

Step 4 To determine your configuration, select your location and make sure it has the proper information on the Behavior tab and the Network tab.

Note Be sure that there are no overlapping network mappings in the Network tabs for each Location.

Services Validation

Validating services consists of using either the Zabbix or SNMP system to determine that all of the services on all of the components are running properly. If there is an issue, then you can access the individual components to determine if the application services are running properly or not.

This section uses command lines as **qns user** with an SSH session.

Zabbix/SNMP

Open Zabbix in a browser window and log in.

```
<IP_Address>/zabbix
```

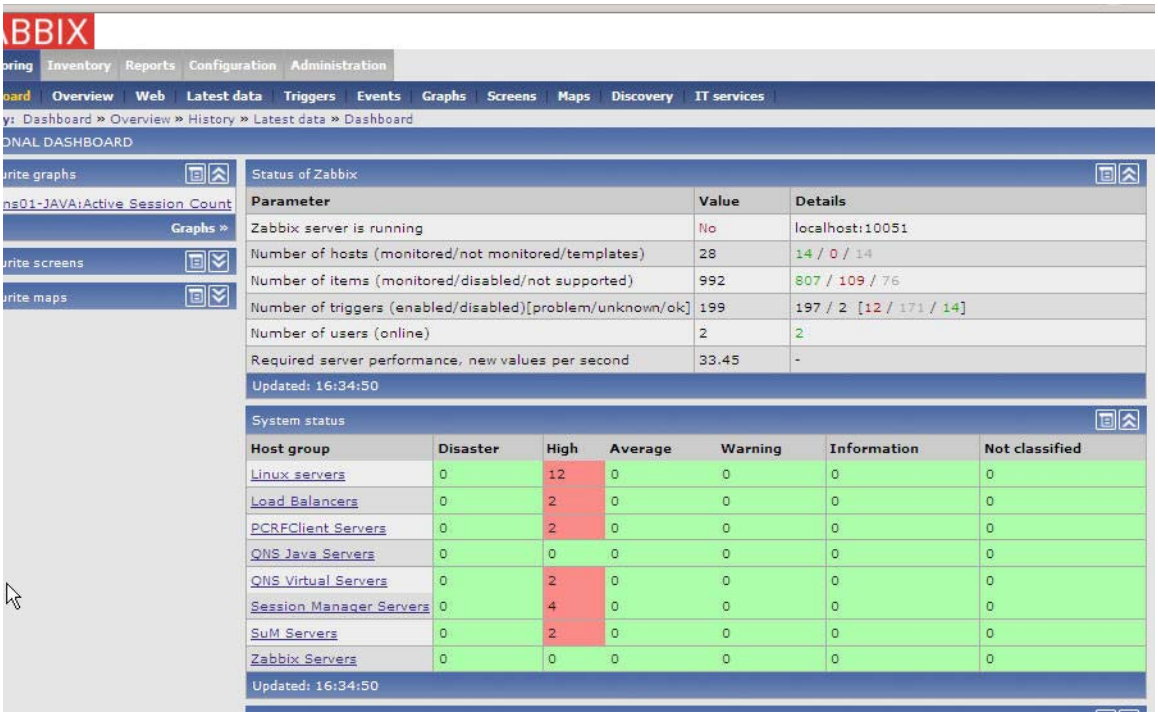

Make sure all the backgrounds are green. Any red areas mean that a service may not be running.

Note Zabbix runs on ControlCenter01 and 02 so if Zabbix is not running, use the checks for the services in that section below.

SNMP should provide something similar as a dashboard, but that is operator specific.

LB01/02

Heartbeat, memcached, and qns (IOMgr) must be running on both load balancers, always.

```
ps -ef | grep java
```

```
ps -ef | grep memcached
ps -ef | grep heartbeat
```

The one that is the active node (has the VIPs), must have Proxy.

```
ps -ef | grep haproxy
```

QNS0x

Java must be running for the Policy Engine to run.

```
ps -ef | grep java
```

SessionMgr0x

The sessionMgr database must be running. Depending upon the solution, there can be between 1 and 6 sessionMgr databases running.

ps -ef | grep mongo

One is for Sessions on port 27717 and is mandatory.

One is for Balance on port 27718.

Two are for SPR (USuM) on port 27719 and 27720.

One is for portal on port 27730.

PortalLB01/02

Heartbeat must be running.

```
ps -ef | grep heartbeat
```

The active node (which has the VIPs), must have sTunnel and HAProxy running.

```
ps -ef | grep stunnel
ps -ef | grep haproxy
```

Portal0x

Apache and Tomcat must be running.

```
ps -ef | grep httpd
ps -ef | grep java
```

ControlCenter01/02

Apache, Cisco Policy Builder, CPS, and MySQL must be running.

```
ps -ef | grep httpd
```

```
ps -ef | grep java | grep pb
ps -ef | grep java | grep pcrf
ps -ef | grep mysql
```

Checking Access

When you are confident that the installation and configuration tasks are complete and processing properly, try running a small amount of test traffic, following it through the system.

Here are three ways to ascertain correct process of access from a subscriber perspective.

Testing Subscriber Access with 00.testAccessRequest.sh

00.testAccessReqest.sh is a test script used to test subscriber access to the ISG and CPS system. You cna find the 00.testAccessReqest.sh in `/opt/broadhop/installer/isg/troubleshooting` directory on the CPS server.

To configure the subscriber used, edit `/opt/broadhop/installer/isg/troubleshooting/config.ini` file.

Step 5 In the `config.ini` file, change the User-Name and Password fields.

Note You may need to change some of the other parameters in order to match your configuration. The other main attributes to change will be the NAS-IP-Address and Framed-IP-Address.

Step 6 Run the script from a command line. No arguments are necessary:

```
00.testAccessRequest.sh
```

Upon success, this output displays:

```
ot@lab:/opt/broadhop/installer/isg/troubleshooting
t@lab troubleshooting]# 00.testAccessRequest.sh
-IP-Address=192.168.181.29,Framed-IP-Address=70.1.2.43,User-Name=banner,Use
sword=banner,NAS-Identifier=QA-Test,Calling-Station-Id=0200.0000.0001,Acct-
on-Id=00000626" | radclient 192.168.181.112:1812 auth aaacisco -r 1 -x
ing Access-Request of id 54 to 192.168.181.112 port 1812
    NAS-IP-Address = 192.168.181.29
    Framed-IP-Address = 70.1.2.43
    User-Name = "banner"
    User-Password = "banner"
    NAS-Identifier = "QA-Test"
    Calling-Station-Id = "0200.0000.0001"
    Acct-Session-Id = "00000626"
recv: Access-Accept packet from host 192.168.181.112 port 1812, id=54, leng

    Cisco-Account-Info = "A2M-UP-DOWN"
    Cisco-AVPair = "accounting-list=QNS_ACCT_LIST"
```

Testing Subscriber Access with soapUI

This procedure tests end subscriber access to your system.

Step 1 Download soapUI from here: http://www.soapui.org/

You only need the freeware version (not soapUI Pro).

Step 2 Launch soapUI.

Step 3 Right click on projects, select New soapUIProject from the drop-down menu.

Step 4 Name your project and enter into Initial WSDL/WADL the appropriate WSDL URL (you may have to replace the IP in display with your own IP) and select OK:

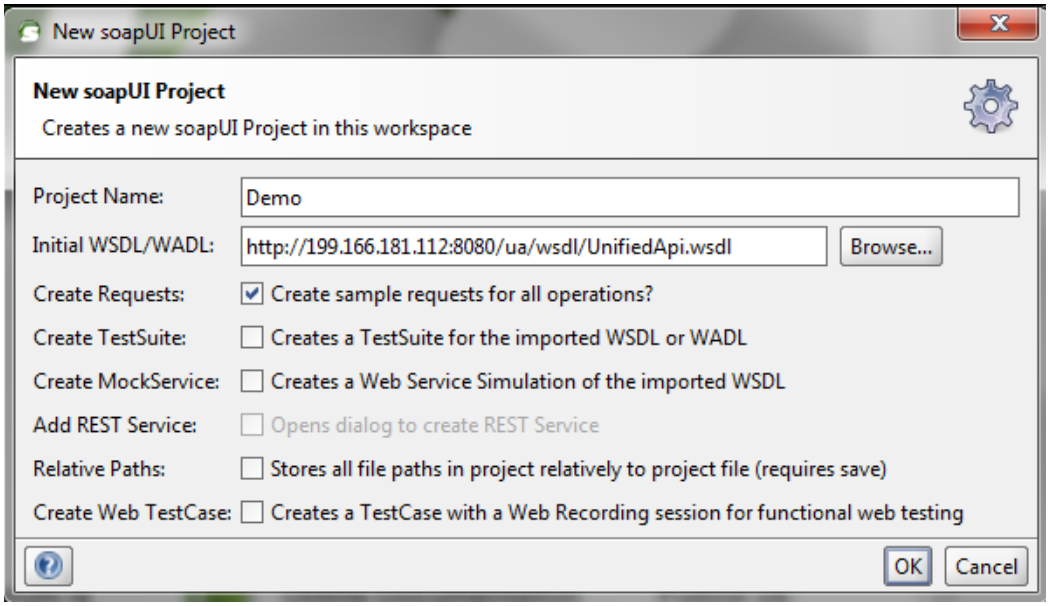

Step 5 In the tree click Demo > SubscriberBinding > GetSubscriber > Request 1, as shown in the figure on the right.

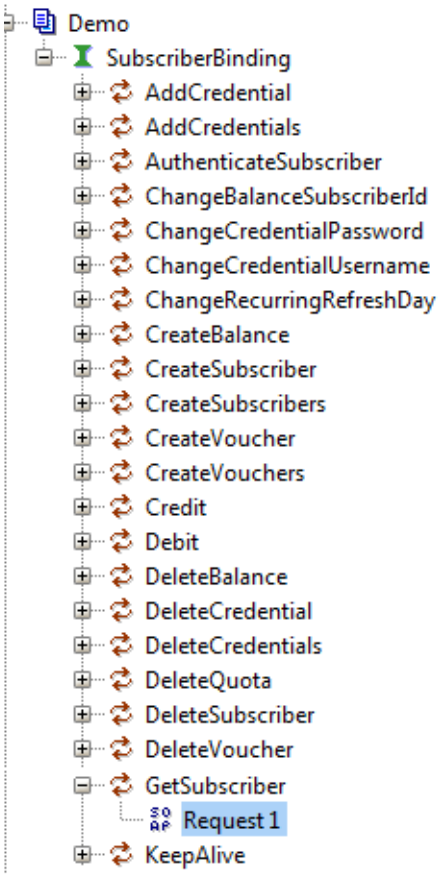

Step 6 Edit the End Point by selecting from the drop down: [edit current...]. Enter the appropriate IP.

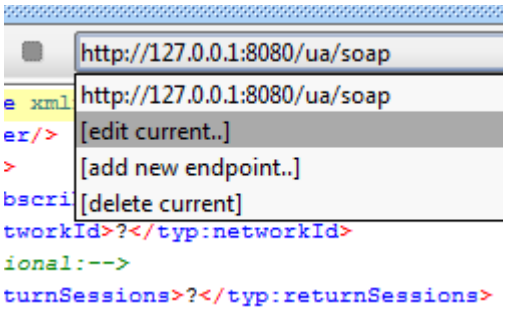

Step 7 In the XML file:

- Replace the ? in <typ:networkId>?</typ:networkId> with the appropriate credential or network Id.
- Replace the ? in <typ:returnSessions>?</typ:returnSessions> with "True".

- Replace the ? in <typ:returnBalance>?</typ:returnBalance> with "True".

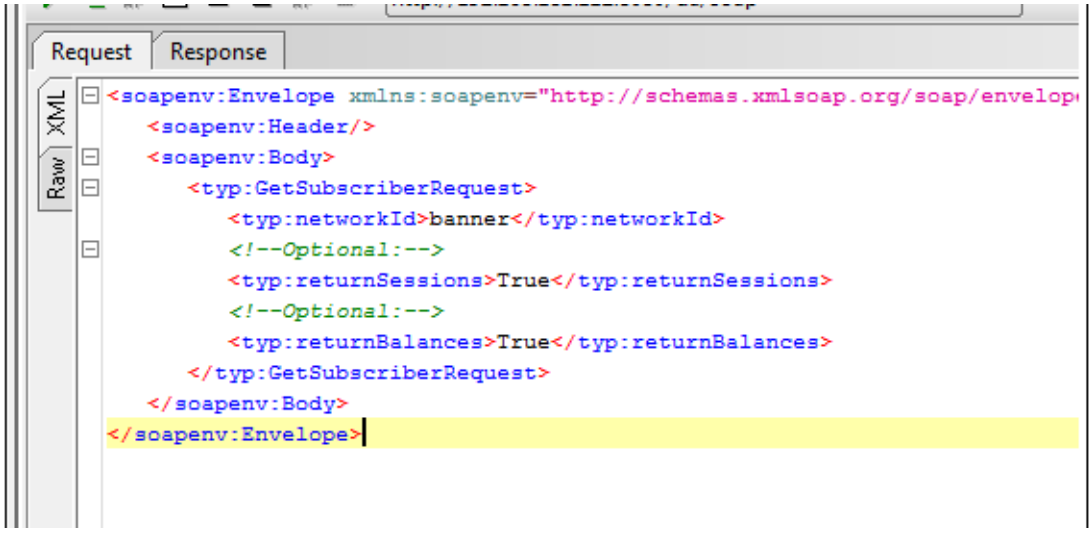

Step 8 Click on the green arrow (underneath “Request 1").

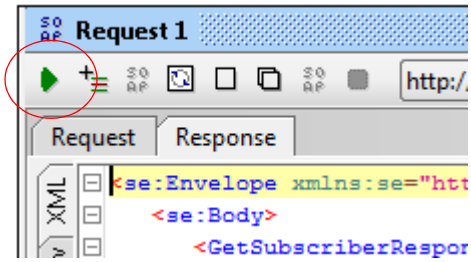

Step 9 Check the resulting XML output. Pay special attention to the relevant subscriber information:


```
se:Envelope xmlns:se="http://schemas.xmlsoap.org/soap/envelope/">
  <se:Body>
     <GetSubscriberResponse xmlns="http://broadhop.com/unifiedapi/soap/types">
        <errorCode>0</errorCode>
        <errorMessage>Request completed successfully</errorMessage>
        <subscriber>
           <id>4fb54d03e4b01e8478d309c2</id>
           <name>
              <fullName>Bruce Banner</fullName>
           </name>
           <credential>
              <networkId>banner</networkId>
              <password>banner</password>
           </credential>
           <credential>
              <networkId>0200.0000.0001</networkId>
              <expirationDate>2012-05-17T13:17:07.020-06:00</expirationDate>
           </credential>
           <service>
              <code>SERVICE_A</code>
              <enabled>true</enabled>
           </service>
           <session>
              <sessionKey>
                 <code>UserIdKey</code>
                 <primary>false</primary>
                 <keyField>
                    <code>userId</code>
                    <value>banner</value>
                 </keyField>
              </sessionKey>
              <sessionObject>
                 <entry>
                    <string>tags</string>
                    <list>
```


Testing for ISG Functionality and Connectivity with test aaa Scripts

The four scripts described here test ISG functionality and connectivity.

Step 1 Connect to the ISG with username and password.

Step 2 Type the 'en' command

Step 3 Enter 'cisco' as the password to the en command.

From here use the four "test aaa" scripts to verify correct ISG functionality and connectivity. No IP addresses or any other arguments are needed.

```
test aaa group QNS_AAA PBHK_SERVICE servicecisco legacy
```

```
test aaa group QNS_AAA L4REDIRECT_SERVICE servicecisco legacy
test aaa group QNS_AAA OPENGARDEN_SERVICE servicecisco legacy
test aaa group QNS_AAA BroadHop BroadHop legacy
```

If functioning correctly, each script returns this message:

Attempting authentication test to server-group QNS_AAA using radius

User was successfully authenticated.

Reference Data Configurations

Revised: July 10, 2015

Under the Reference Data tab, several important configuration tasks are available.

Click the links below to go to sections that discuss the Reference Data tab's configuration tasks.

- Systems:

  System Configuration, page B-1

  The Systems node in the Reference Data tree shows the environment inside the Cisco Policy Builder deployment, that is, the configuration of the Cisco Policy Builder servers, clusters, and instances.
- For custom data structures that import spreadsheet data:

  Customer Reference Data Tables, page C-1
- Using Notifications to send messages to subscribers:

  Notification Configurations, page D-1

  The Notifications node sets up communications to subscribers or administrators with SMS, emails, or Apple push messages to iPhones.
- Defining Your PEPs:

  Policy Enforcement Point Tree, page E-1
- Using Subscriber Data Sources:

  Subscriber Configuration, page F-1
- Managing balances and quotas:

  Account Balance Template Configuration and Tariff Time Configuration, page J-1
- Setting Tariff Times:

  Account Balance Template Configuration and Tariff Time Configuration, page J-1

System Configuration

Revised: July 10, 2015

This appendix covers the following sections:

Overview

The Systems node in the Reference Data tree shows the environment inside the Cisco Policy Builder deployment, that is, the configuration of the Cisco Policy Builder servers.

You always have at least one system defined in Cisco Policy Builder. A system represents a complete Cisco Policy Builder runtime environment. That system is comprised of at least one cluster, and the cluster contains instances.

After installing your VMs, use the Systems node under the Reference Data tab.

Other tasks in the Reference Data tree are discussed at Reference Data Configurations and in the individual blueprint documents required for your deployment.

Hierarchy and Precedence

In the Systems node, settings and configuration attributes may be changed at any of the three levels of hierarchy of system, cluster, and instance. That is, the settings and attributes cascade down from the system level to the cluster level to the individual instances.

Adding the First System Using Default Data

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.

You can modify the example in the next steps.

To create a system without default data, see Making a System Without Default Data.

Defining a System

You can use blank screens and fill them in as you know how, or start with default data to guide you in filling in the screens.

Step 1 Click Reference Data tab > Systems node to display the systems tree.

Step 2 User can create a System with pre-populated data or can create a new system.

Step 3 Click ... to open Create From Example screen.

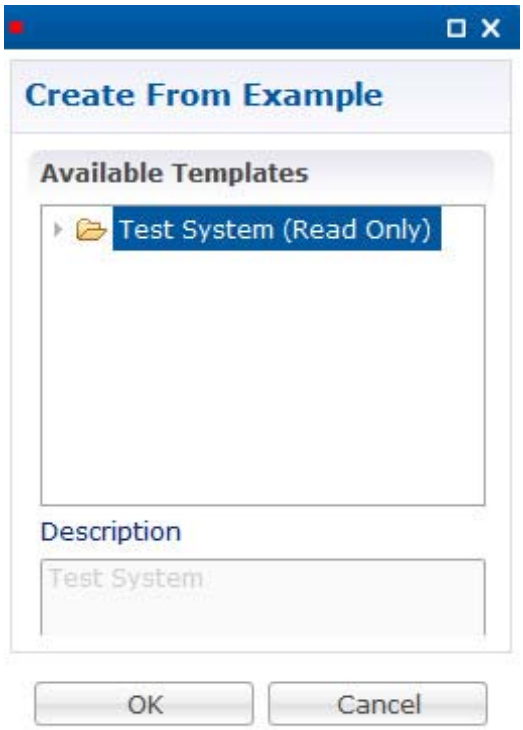

a. Select Test System to use the System with pre-populated data.

b. When you select an test system, configuration plug-ins are provided but you have to select them, depending on your environment and business purposes.

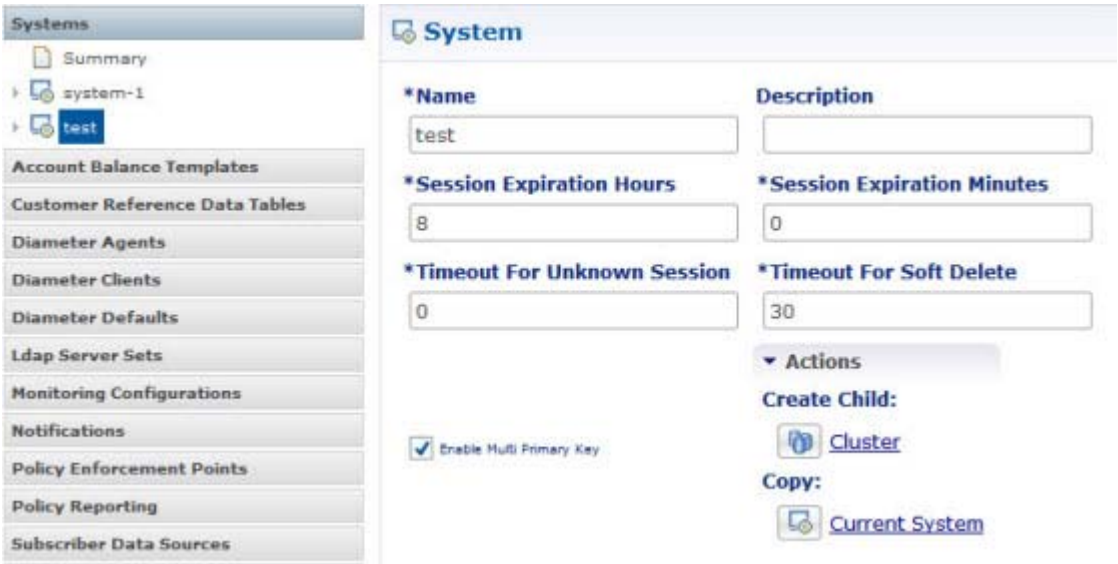

c. In the right pane, begin to work through the screens and change the example data to conform to your needs.

Step 4 Click System to open default system to start with blank screen.

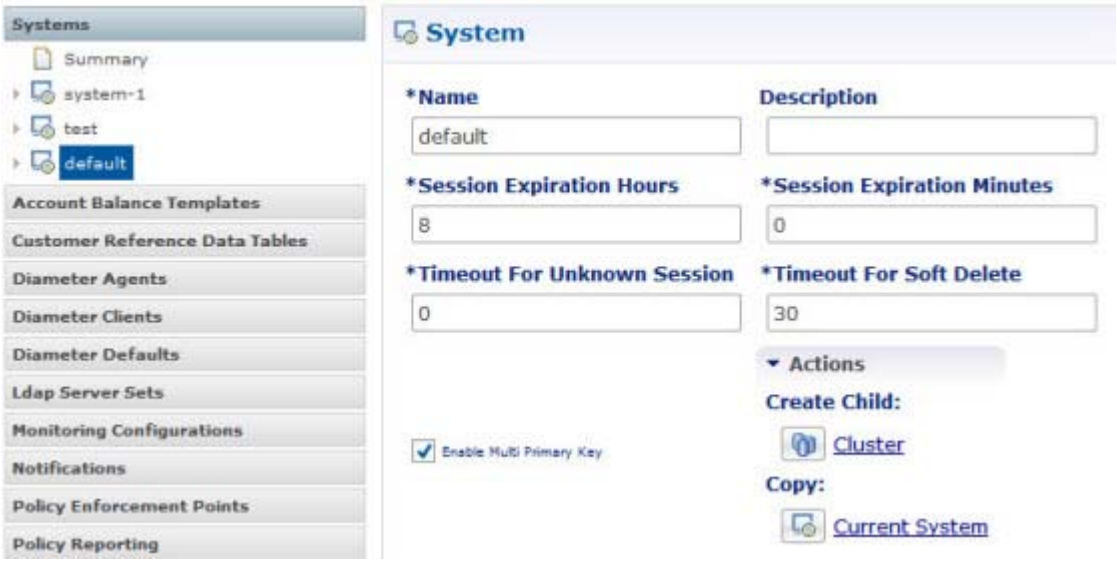

Step 5 Change Name field and add a description.

Note This name must correspond to the system ID used in the configuration files for Cisco Policy Server. See Preface for the reference for the Cisco Policy Server. For example, system-1 is provided at install. If you add a system, or change the name of system-1, you must go to the qns.conf file and make the change there as well.

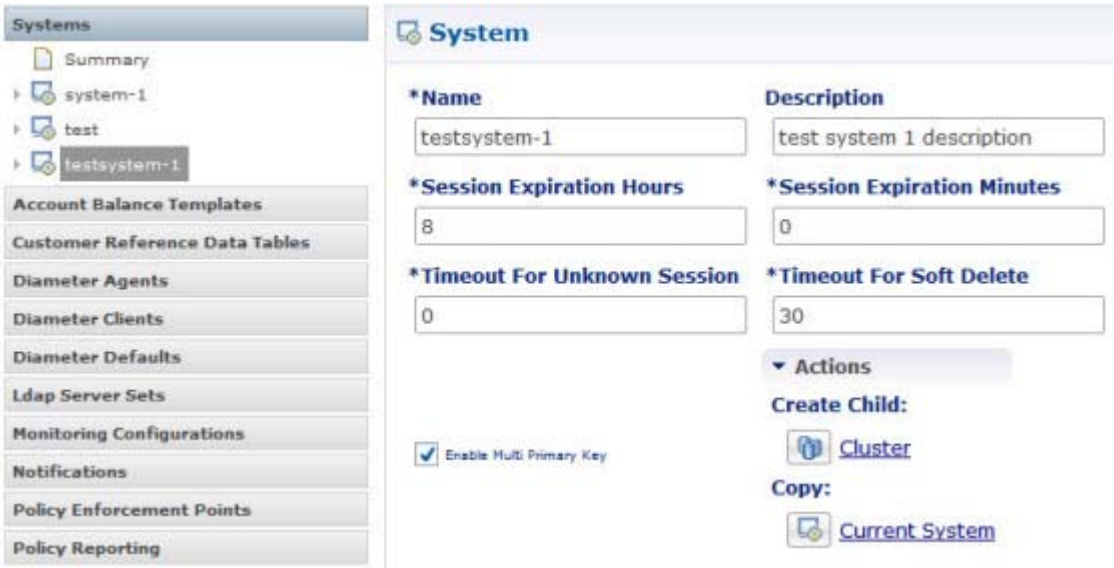

Name	The name of the CPS system.
Description	The description of this entire system.
Session Expiration Hours and Session Expiration Minutes	The amount of time where CPS doesn't received a message for the session after which the session will be removed from session database. Every time a new message is received for a session, the session expiration counter is reset to the full value on a per session basis. This enables CPS to tear down 'orphaned' sessions in cases where a network device is hard reset without sessions being torn down, etc.
Timeout for Unknown Session	An unknown session is one in which the authorization fails but the 'unknown' service is set on the Domain allowing unauthenticated users to have service. The unknown session timeout occurs if we don't receive any messages for the session during this period. Every time a new message is received for a session, the session expiration counter is reset to the full value on a per session basis. This setting is in seconds.
Timeout For Soft Delete	Soft delete sessions are used for scenarios where the session 'stop' message is not the last message in the call-flow. Typically, the reasons for these sorts of call-flows is that a quota used message will come in after the stop message (accounting stop, etc) to provide accurate accounting up to the stop message. In CPS, soft delete session changes the primary key of the session to a UUID and removes all secondary keys except the ones which are unique to the session. This enables a new session to start while a soft delete exists. This timeout setting instructs CPS how long to keep the soft delete session before deleting. The setting is in seconds.
Enable Multi Primary Key	Select this check box to indicate that in the database, this row, for this system, has a multi-part key to make the row unique.

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 6 Review the values for session expiration and set them to perform a timeout after a set period of non-activity.

Step 7 Review the values for Timeout set them to allow adequate time to recognize an unknown subscriber log in or unknown session.

Step 8 From the Systems tree, open up the node that you just added, testsystem-1, and check the plug-in configurations.

Click an item on the right to add those plug-ins to the system selected under Plugin Configurations node.

Defining a 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 these steps.

Step 1 Begin with a system at the Systems node in the Reference Data tree.

Step 2 Click the Cluster link to set up your first cluster.

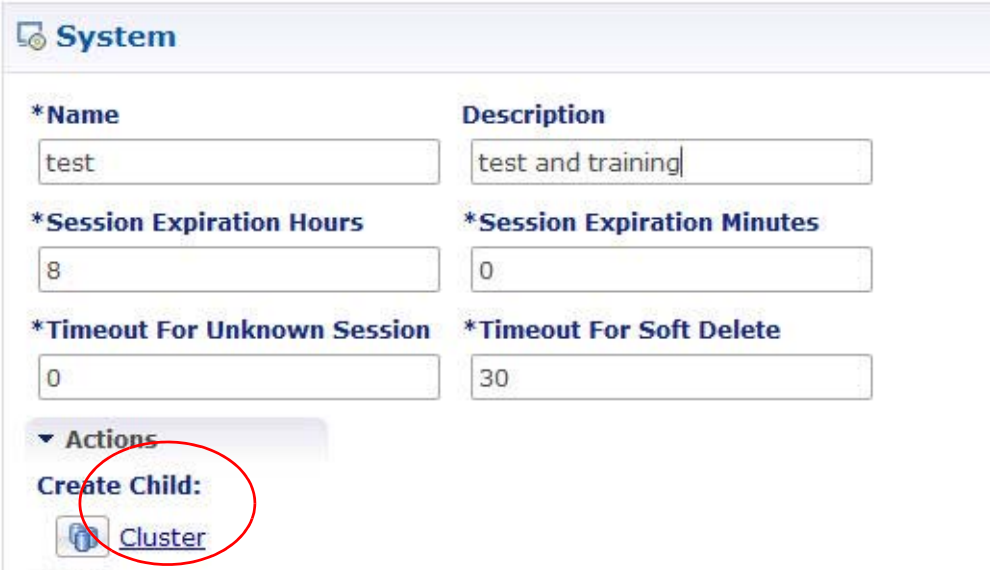

Because some data are relevant at the cluster level, you always have at least one cluster, even if it is a cluster of one instance.

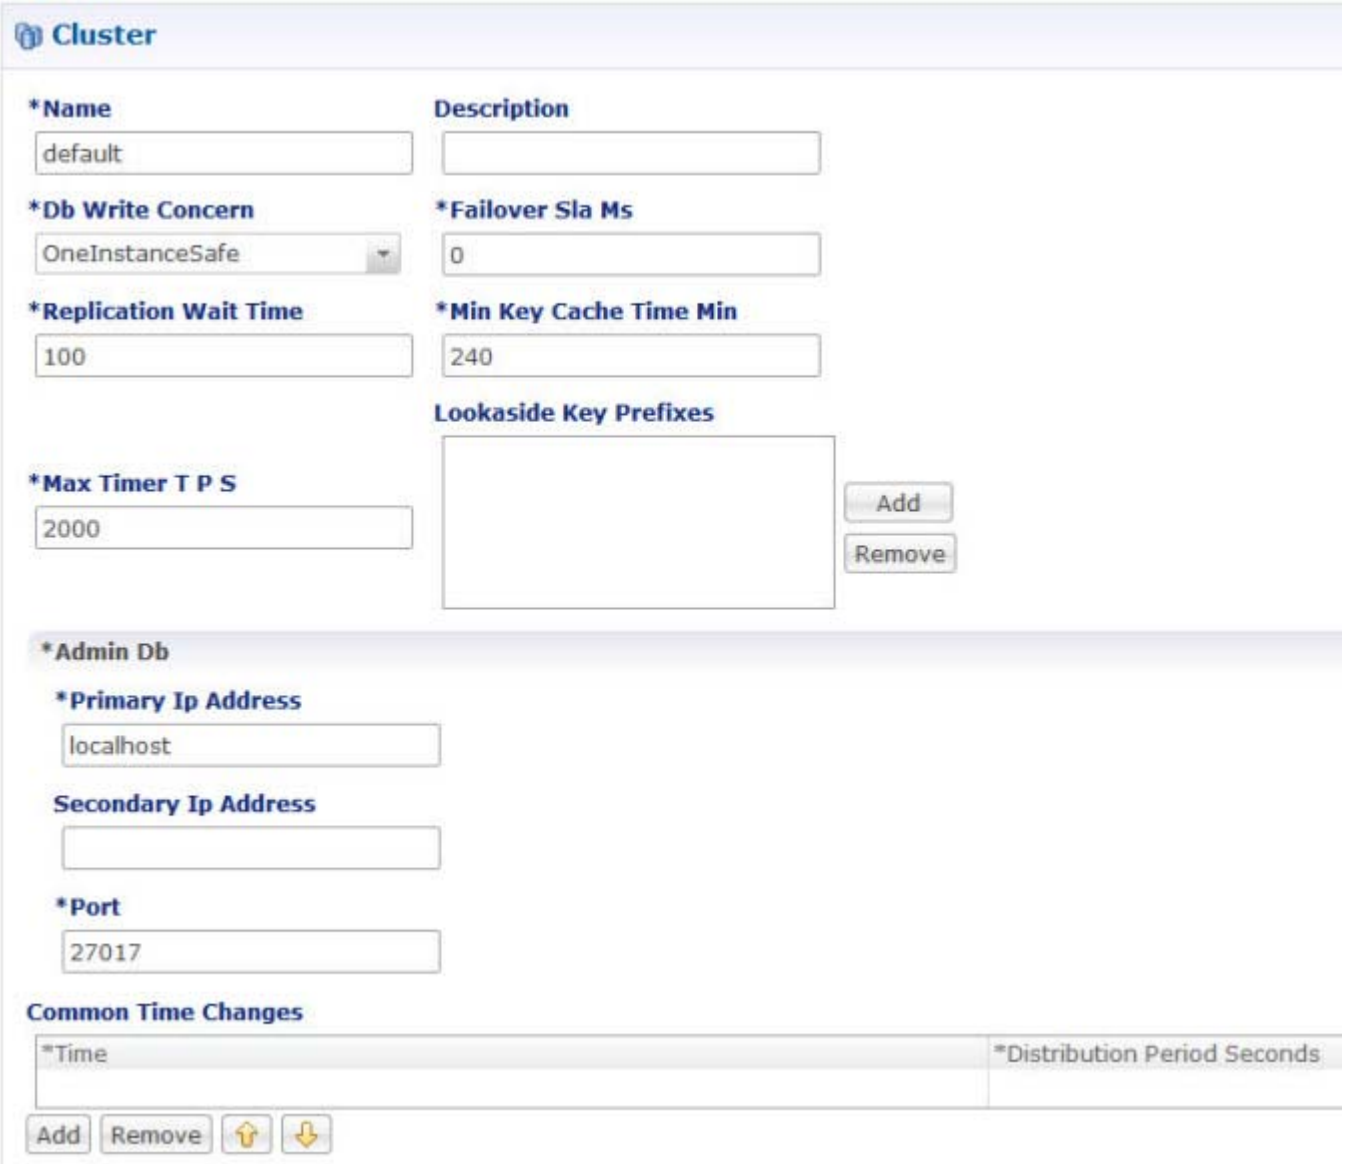

Name	The name of the cluster. This name must correspond to the value stated in the config.ini file on the Cisco Policy Server
Description	A brief description of the cluster.
Primary Ip Address	The IP address of the Session Manager database that holds session information for Cisco Policy Builder and Cisco Policy Server.

Secondary Ip Address	The IP address of the database that provides fail over support for the primary database. This is the mirror of the database specified in the Primary IP Address field. Use this only for replication or replica pairs architecture.
Port	This field is present but deprecated to maintain downward compatibility. Because there is a separate database for session data and logging (Session DB) / reporting (control DB) / diagnostic data, you must provide a port number for the Session database. 27717 is typical.

Click Instance to set up an instance within this cluster.

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

Any of the configurations you specify here are used at the cluster level only 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 on 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.

See Plug-in Configurations for instructions on changing or deleting plug-in configurations.

Defining an Instance

Step 1 Begin with a cluster at the Systems node in the Reference Data tree as discussed in Defining a Cluster.

Step 2 Click the Create Child Instance link on the Cluster screen.

Step 3 View the Instance screen.

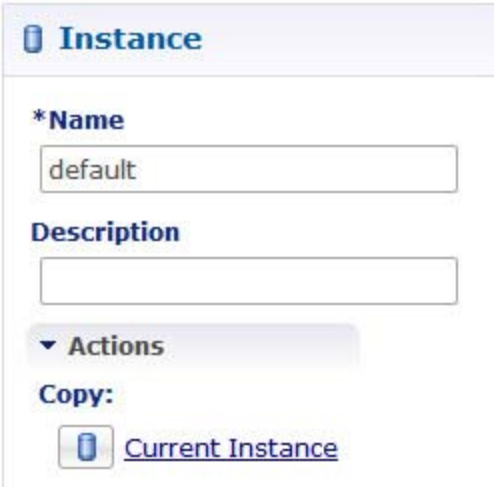

Step 4 Fill in a name and a description.

Step 5 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 on any one of them 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 is used at the instance level, overriding any plug-ins set at the system level or the cluster level.

See Plug-in Configurations for instruction on changing or deleting plug-in configurations.

Step 6 Check the Systems tree for red Xs like this and make sure you have none. Investigate and correct them if you do.

Making a System Without Default Data

After a successful installation, use this procedure to manually configure the CPS system via the Cisco Policy Builder. No default data is provided with this method.

Step 1 Click Reference Data >Systems node to display the Systems tree.

Step 2 Click the System link to fill in the fields, starting with a blank form.

Step 3 Fill in the Name field and provide a description of this system.

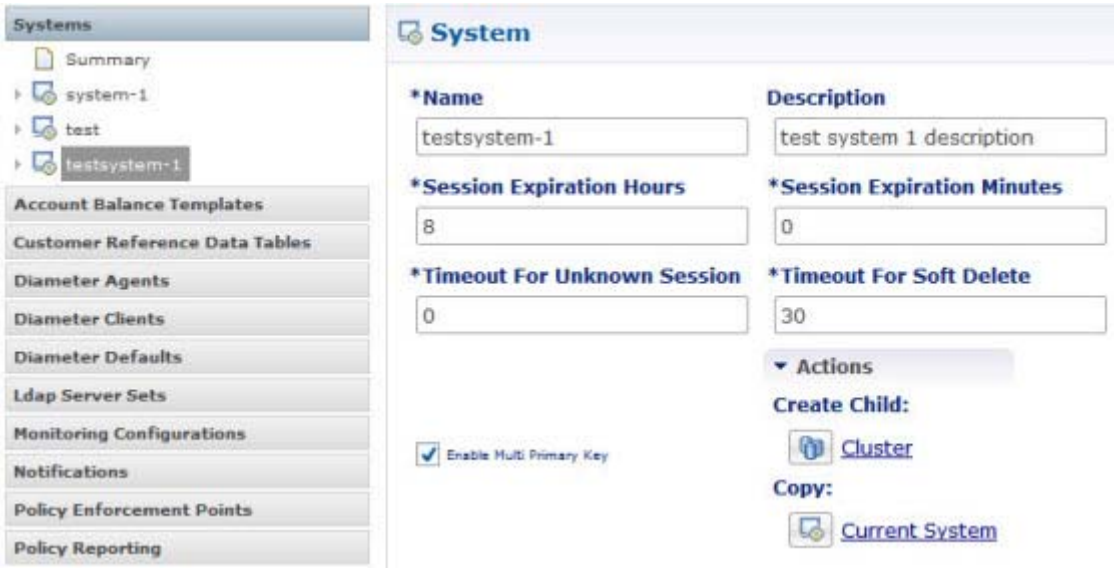

Name	The name of the CPS system.
Description	An apt description of this entire system.
Session Expiration Hours and Session Expiration Minutes	The amount of time where CPS doesn't received a message for the session after which the session will be removed from session database. Every time a new message is received for a session, the session expiration counter is reset to the full value on a per session basis. This enables CPS to tear down 'orphaned' sessions in cases where a network device is hard reset without sessions being torn down, etc.
Timeout for Unknown Session	An unknown session is one in which the authorization fails but the 'unknown' service is set on the Domain allowing unauthenticated users to have service. The unknown session timeout occurs if we don't receive any messages for the session during this period. Every time a new message is received for a session, the session expiration counter is reset to the full value on a per session basis. This setting is in seconds.
Timeout For Soft Delete	Soft delete sessions are used for scenarios where the session 'stop' message is not the last message in the call-flow. Typically, the reasons for these sorts of call-flows is that a quota used message will come in after the stop message (accounting stop, etc) to provide accurate accounting up to the stop message. In CPS, soft delete session changes the primary key of the session to a UUID and removes all secondary keys except the ones which are unique to the session. This enables a new session to start while a soft delete exists. This timeout setting instructs CPS how long to keep the soft delete session before deleting. The setting is in seconds.
Enable Multi Primary Key	Select this check box to indicate that in the database, this row, for this system, has a multi-part key to make the row unique.

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

If your system has many servers, you may want to group instances of Cisco Policy Builder clusters. Multiple clusters provide Cisco Policy Builder services, redundancy, and scalability for your system. Often, a blade server is a logical cluster, with each blade as an instance.

Adding Another Cluster

Use these steps to add another cluster to a system that already exists.

Step 1 Click Reference Data > Systems to open up the Systems tree.

Step 2 Select the system you want to add the cluster to.

Step 3 In the right pane, check the name and description of the system and make sure you have the right one.

Step 4 Click the Cluster link to create your cluster.

Step 5 For the rest of this procedure, go to step 2.

Adding Another Instance

This procedure explains how to add an instance to a system and cluster that is already defined.

Step 1 Click Reference Data > Systems label in the tree.

Step 2 In the Systems tree, select the system and the cluster you want to add an instance to.

Step 3 Proceed to create the instance of Cisco Policy Builder. Click the Create Child Instance link.

For the rest of this procedure, go to step 2.

Plug-in Configurations

Cisco Policy Builder provides core plug-ins for customizing and optimizing your installation. Additionally, other plug-ins become available as the product matures.

- Configurations set at the system level are system-wide except as noted in the bullet items below.
- Configurations set at the cluster level apply to that cluster and the instances in it. A value set here overrides the same value set at the system level.
- Configurations set at the instance level apply to the instance only and override the same value set at the cluster or system level.

Step 1 Select the Create Child action in a Plug-in Configuration node in the Systems tree to define them.

You may change any of the variables from the default, or choose not to use a plug-in, as necessary.

Do not use the Threading Configuration unless instructed to do so by Cisco Technical Support.

Notifications Node

The Notifications node in the Reference Data tree lets you send messages from and about Cisco Policy Builder to subscribers. See Notification Configurations for all the ways you can send messages.

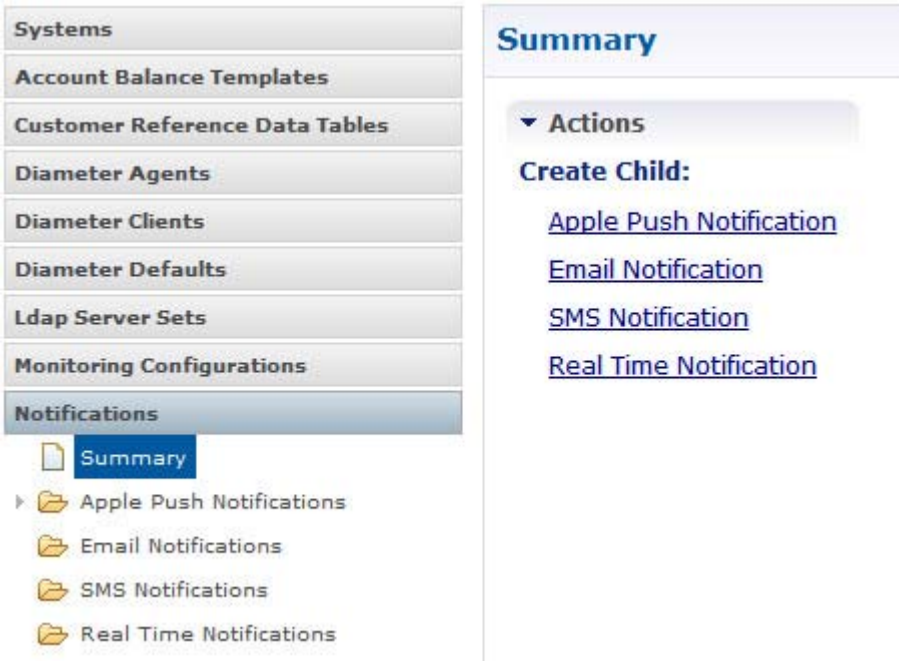

Plug-ins at the Cluster and Instance Level

When you create a system from the example, these configuration stubs appear at the cluster and instance level.

Threading Configuration

A threading configuration utility is provided for advanced users and future development.

If you feel you need to use this screen, please contact your Cisco technical representative.

Thread Pool Name	Name of the Cisco thread pool. Examples include default.
Threads	Threads to set in the thread pool.
Queue Size	Size of the queue before they are rejected.
Scale By Cpu Core	Select this check box to scale the maximum number of threads by the processor cores.

Async Threading Configuration

You are always required to select this configuration, but no changes to it are necessary.

Note Always select the link for Async Threading Configuration to configure your CPS system.

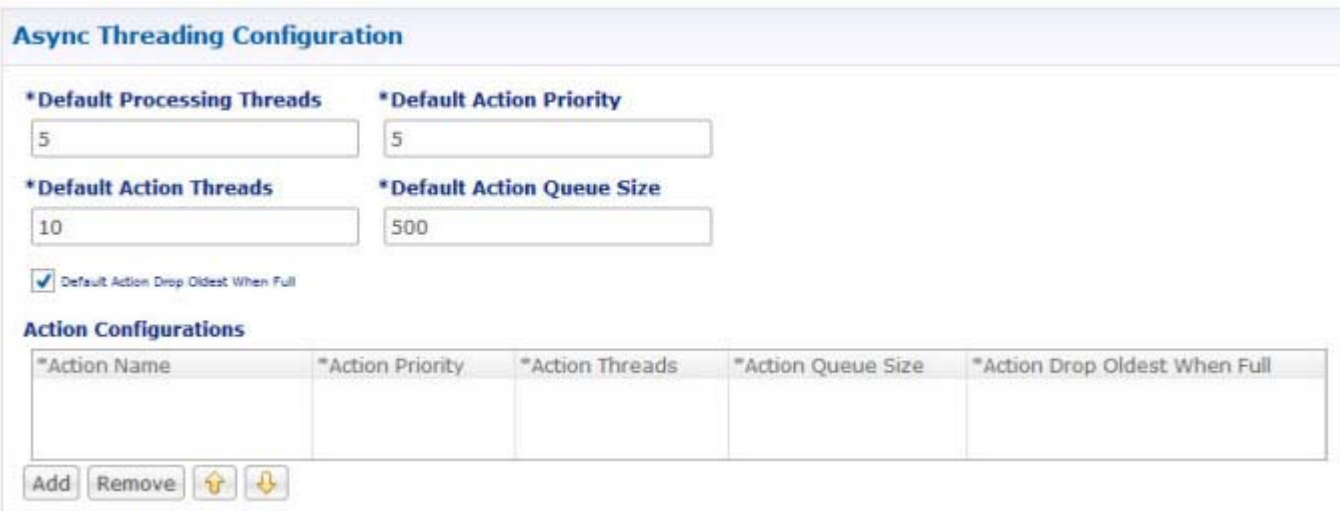

Default Processing Threads	The number of threads that are allocated to process actions based on priority.
Default Action Priority	The priority assigned to an action if it is not specified in the Action Configurations table.
Default Action Threads	The number of threads assigned to process the action if it is not specified in the Action Configurations table.
Default Action Queue Size	The number of actions that can be queued up for an action if it is not specified in the Action Configurations table.

Default Action Drop Oldest When Full	When checked, the oldest queued action is dropped from the queue when a new action is added to a full queue. Otherwise, the new action to add is ignored. This check box applies to all the threads specified in the fields above. To drop a specific thread, leave this unchecked and use the Action Configurations table.
Action Configurations Table	
Action Name	The name of the action. This must match the implementation class name.
Action Priority	The priority of the action. Used by the default processing threads to determine which action to execute first.
Action Threads	The number of threads dedicated to processing this specific action.
Action Queue Size	The number of actions that can be queued up.
Action Drop Oldest When Full	For the specified action only: When checked, the oldest queued action is dropped from the queue when a new action is added to a full queue. Otherwise, the new action to add is ignored.

Mongo SPR Connection Tuning Parameters

The following parameters need to be specified in `/etc/broadhop/pcrf/qns.conf`.

-Dspr.mongo.connection.timeout=1000 (default)

-Dspr.mongo.threads.allowed.to.wait.for.connection=1500 (Default)

-Dspr.mongo.no.auto.connect.retry=true/false (default false)

-Dspr.mongo.no.keep.alive=true/false (default false)

-Dspr.mongo.socket.timeout=60000 (60 seconds)

-Dspr.mongo.thread.maxWaitTime=0 (default)

-Dspr.mongo.max.auto.connect.retryTime=5 (default)

You can tune the above parameters as per your performance need. The default parameters are not required to be defined in qns.conf.

After tuning the paramaters, synchronize the configuration by executing the command `#syncconfig.sh`.

USum Configuration Performance Tuning Parameters

Step 1 Login to pcrfclient01.

#cd /etc/broadhop/pcrf

Step 2 Set below parameters in qns.conf.

-Dspr.mongo.socket.timeout=5000

-Dspr.mongo.thread.maxWaitTime=2000

Step 3 Execute `#syncconfig.sh`.

Step 4 Change thread count to 50.

Step 5 Increase “connections per Host" to 20 and DB read preference to “SecondaryPreferred”.

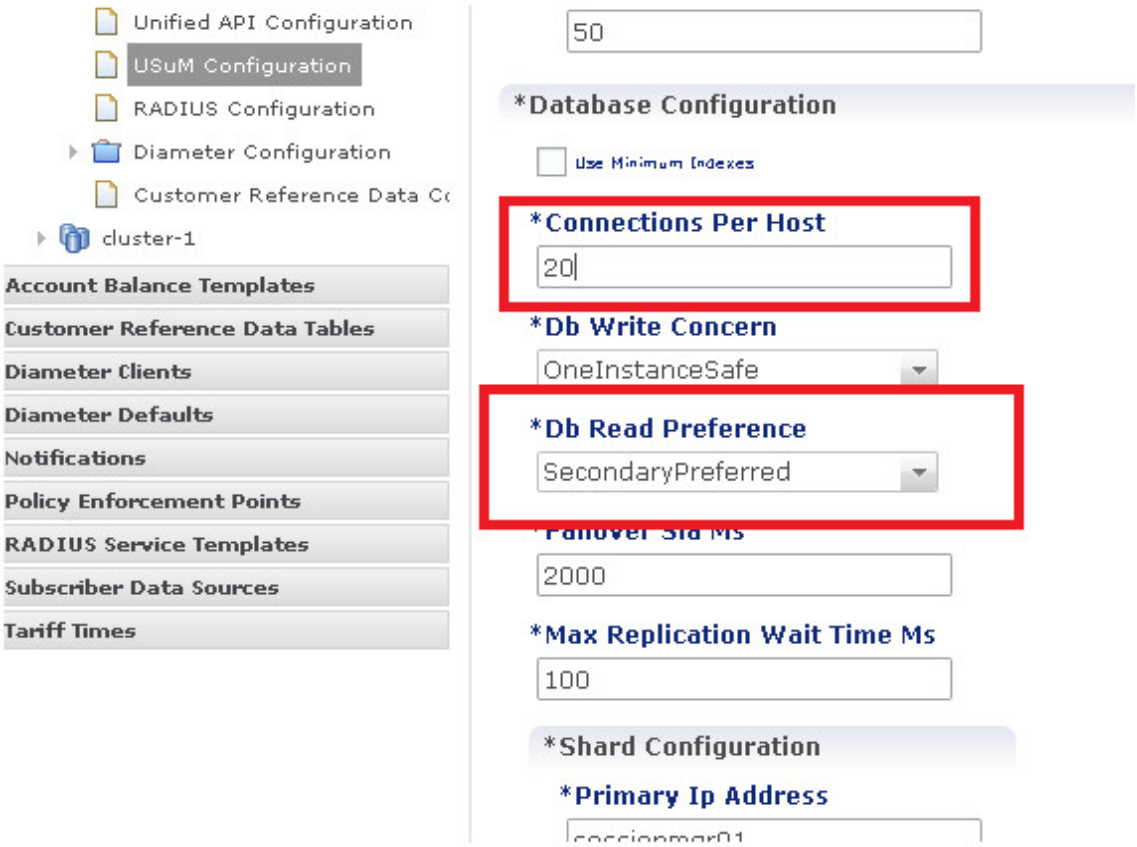
Unified API Configuration
USuM Configuration
RADIUS Configuration
Diameter Configuration
Customer Reference Data Co
cluster-1
Account Balance Templates
Customer Reference Data Tables
Diameter Clients
Diameter Defaults
Notifications
Policy Enforcement Points
RADIUS Service Templates
Subscriber Data Sources
Tariff Times
50
*Database Configuration
Use Minimum Indexes
*Connections Per Host
20
*Db Write Concern
OneInstanceSafe
*Db Read Preference
SecondaryPreferred
2000
*Max Replication Wait Time Ms
100
*Shard Configuration
*Primary Ip Address

Customer Reference Data Tables

Revised: July 10, 2015

In CPS, reference data is considered information that is needed to operate the policy engine, but not used for evaluating policies. For example, under the Reference Data tab in Cisco Policy Builder, are the forms used to define systems, clusters, and instances, and to set times and dates used for tariff switching. The policy engine needs to refer to these data only to process policies correctly, but they do not define the policy itself.

Customer reference data is considered reference data that is specific to a service provider, perhaps the names and characteristics of their networks or cell sites.

Such customer reference data is stored in the data structure of a table, with the columns and field attributes defined by the service provider for their specific use. The Cisco Policy Builder interface does not provide screens for such customized information because it would be so limiting.

Rather, in CPS, customer reference data tables provide a way for service providers to create their own data tables and to populate them.

The resulting customer reference data tables can then be used within Cisco Policy Builder as criteria to use when escalating policy decisions.

Two interfaces are used to construct and populate customer reference data tables:

- Cisco Policy Builder 5.3.5 or greater

  Data table structures are managed in Cisco Policy Builder.
- Cisco Control Center 3.0 or greater, either the full admin or read only privileges

  Data table content is managed in Cisco Control Center.

The information in customer reference data tables handles special considerations such as these:

- Specialty area codes.
- A list of device parameters.
- Location data mapping, to map network sites and cell sites into the subscriber's home network, roaming network, or preferred roaming network.
- IMEI data tagging for smart phone, Apple, or android device, and then use that in policies.

Customer reference data tables allow the service provider to create their own data structures tables and populate them with your own proprietary data.

The data in the tables is then used during policy evaluation.

You can construct your reference data tables yourself, or Cisco builds them for you after installation. If you need any help with customer reference data tables, call your Cisco technical representative.

This appendix covers the following sections:

Customer Reference Data Tables Overview

- The Cisco Policy Builder interface creates and edits the customer reference data table structure, defines its columns, and defines the data type, ranges and size of the fields.
- Cisco Control Center reflects the customer reference data tables constructed in Cisco Policy Builder.
- You provide data to the customer reference data tables in Cisco Control Center.
- Spreadsheets can be imported into a customer reference data table structure. Log in to the Cisco Control Center interface to do so.

When constructing the customer reference data table in the Cisco Policy Builder, look at the spreadsheet you want to import and use the same column names, data types and other attributes.

Steps and Procedures

Use Customer Reference Table with Cisco Policy Builder.

Cisco Policy Builder: Constructing Customer Reference Data Tables

These steps need only be completed once and may already be done.

Use Customer Reference Table with Cisco Control Center.

Cisco Control Center: Populating a Customer Reference Data Table

These steps need to be completed for every customer reference data table you want to create or edit.

Cisco Policy Builder: Constructing Customer Reference Data Tables

There are two tasks needed to create customer reference data tables:

- Setting Up the System Plug-in Configuration
- Creating or Editing a Customer Reference Data Table

Setting Up the System Plug-in Configuration

Before you can create a customer reference data table, configure your system to use the Customer Reference Data Table plug-in configuration.

You only have to do this one time for each system, cluster, or instance. Then you can create as many tables as needed. The steps below configure an example system.

Step 1 Log in to Cisco Policy Builder and click

Reference Data tab > Systems node > the system of your choice > Plugin Configurations.

Step 2 Click Customer Reference Data Configuration in the main pane.

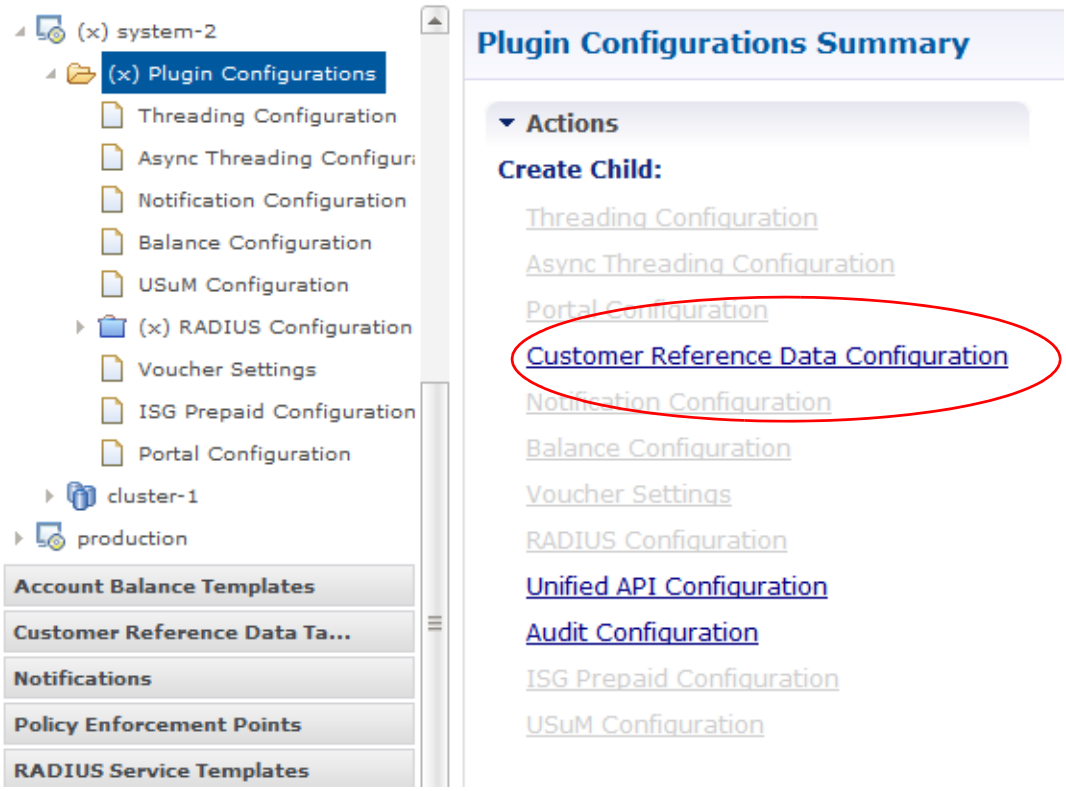

The tree on the left is populated with the configuration.

Fill in the Customer Reference Data Configuration screen that appears.

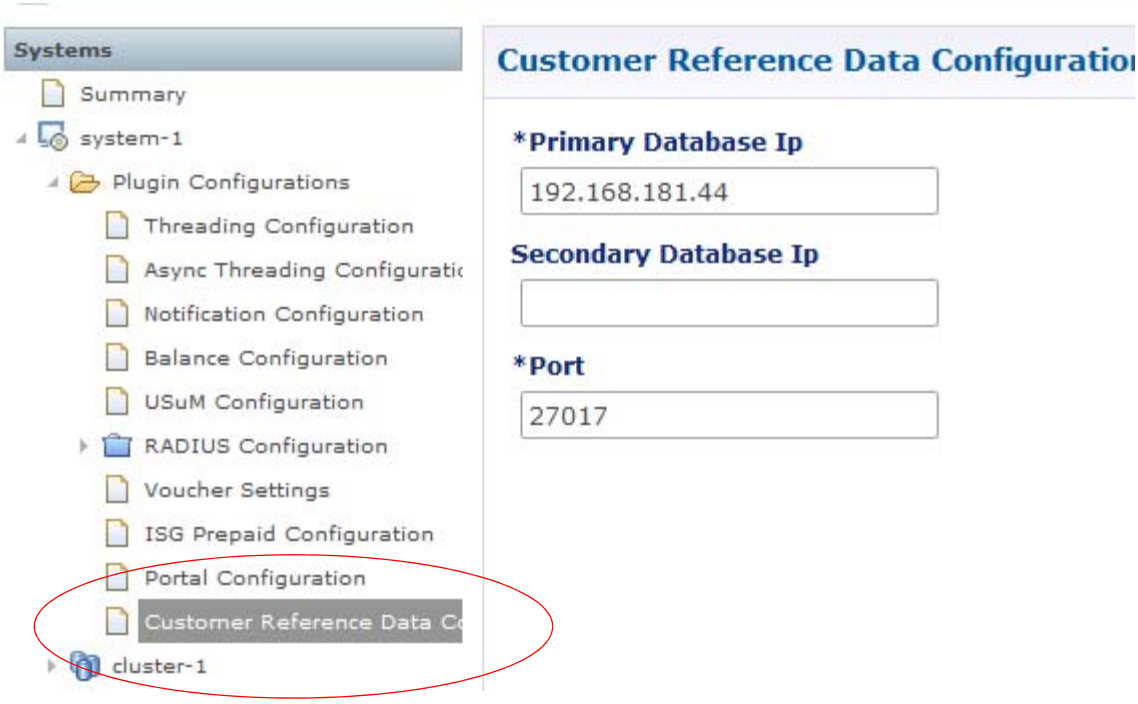

Parameter	Description
Primary Database IP	IP of the sessionMgr database.
Secondary Database IP	Optional, this field is the IP address of a secondary, backup, or failover sessionMgr database.
Port	Port number of the sessionMgr. It should be the same for both the primary and secondary databases.

Step 3 Create data table structures as described in Creating or Editing a Customer Reference Data Table.

Creating or Editing a Customer Reference Data Table

Note Before you begin, be sure that you have configured the plug-in as described at Setting Up the System Plug-in Configuration. Click Reference Data > Systems > Plugin Configuration and make sure you see the Customer Reference Data Config in the tree.

Step 1 In Cisco Policy Builder, click Reference Data tab > Customer Reference Data Tables node.

Step 2 To create a new data table, click Summary > Customer Reference Data Tables link in the main pane.

To edit a table, click a table name in the tree to edit a table.

Note Simply editing the name of a table results in a loss of the table. If you want to change a table name, make a copy first, and then change the name of the copy. There is no data in the newly copied table but you have the new name.

Step 3 Fill in the Data Table screen that appears.

The example shows the columns defined for a table called Countries.

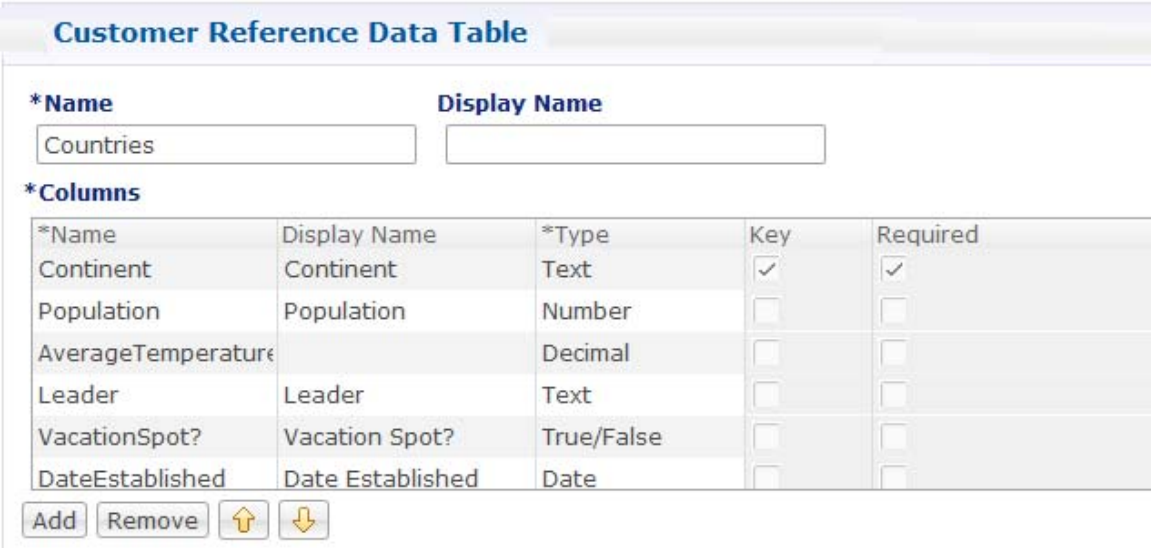

Parameter	Description
Name	The internal name for the column. Best practice is to name the column so that it is recognizable as a key.
Display Name	The table name as you want it to display in Cisco Control Center.

Parameter	Description
Type	The data type of the column, that is: • Text: This is converted to a Java String object. • Number: This is converted to a Java Long object. • Decimal: This is converted to a Java Double object. • True/False: This is converted to a Java Boolean object (1, 1.0, -1, -1.0, true, yes, get converted to true while 0, 0.0, false, no, off get converted to false). • Date: This is converted to a Java Date object using the "MM/dd/yyyy HH:mm:ss" format string. • DateTime: This is converted to a Java Date object using the "MM/dd/yyyy" format string.
Key	If checked, specifies that each row in the table must be unique. If not, when entering data in Cisco Control Center, that admin receives an error indication. Note that several columns can contribute to the whole key. Best practice is to name the column so that it is recognizable as a key.
Is Required	If checked, this box forces the Cisco Control Center admin to enter some value when populating data in the table.
Columns Area	You must change the first row to enable fields in the rest of the screen. Just click in the row to access it.

Step 4 Click on a row in the Columns definition table at the top, and access the fields on the bottom, working your way through all of the column names listed.

In the example, at the top, the row selected is the column called Continent.

In the Valid Values area, only the names of continents in the list are permitted as valid entries. Of course, some continent names have been omitted.

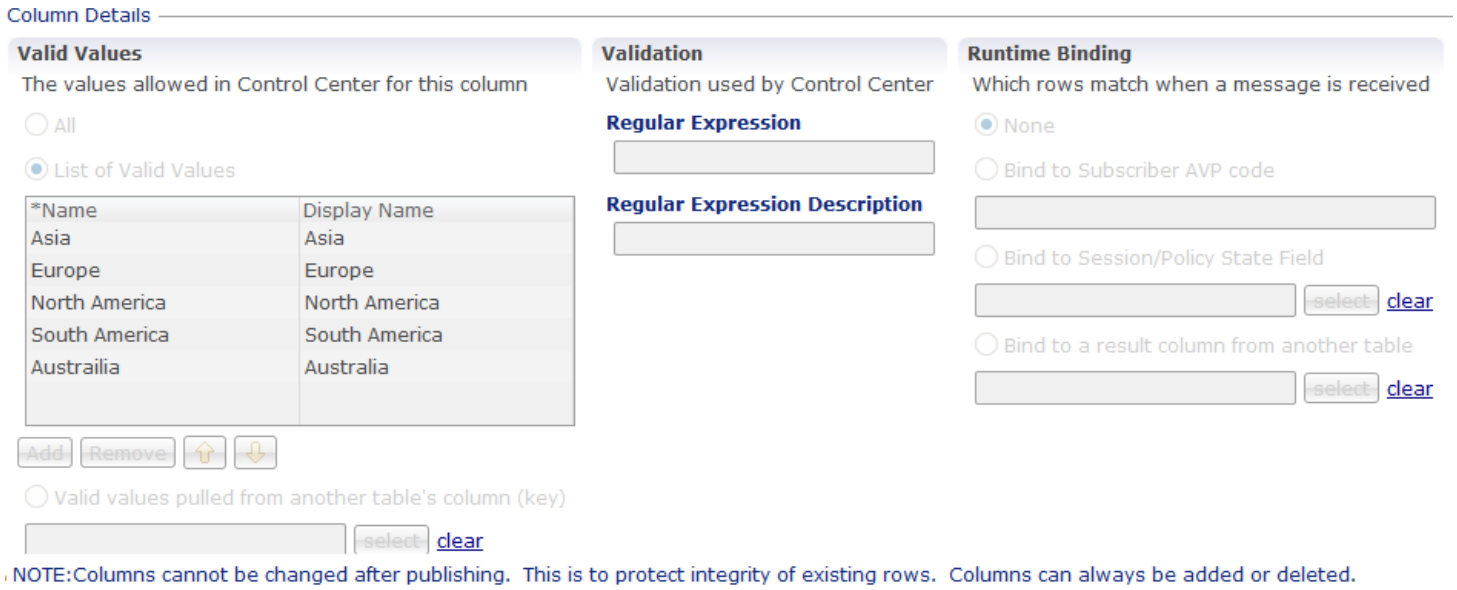

Parameter	Description
Regular Expression	Define a regular expression to define a specific data format. For example if you define a regular expression here, you ensure that a phone number is entered in a phone number format, or a date is entered as a date format when the Cisco Control Center admin is entering data into this column.
Regular Expression Description	A description of what you are trying to achieve above. In Cisco Control Center, this appears as user help or as a tip to help those administrators enter data in the correct format.
Bind AVP Code	Cisco Control Center and Cisco Policy Server may correlate on the basis of AVP codes. If so, specify these codes here.
Bind to Session/Policy State Field	Bind fields are the specific points of intersection to the policy engine, and are optional. Cisco Control Center may not have any specific or defined 'hooks' into the Cisco Policy Server. Click the Select button to display a list of objects that can be used to relate Cisco Control Center with Cisco Policy Server at specific data points. 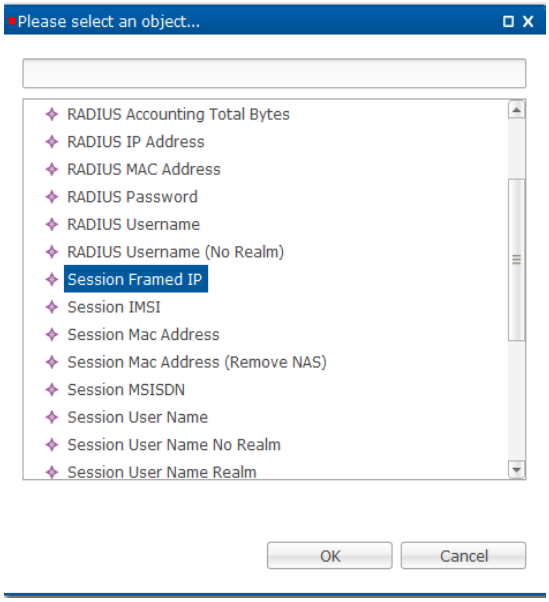

Editing Reference Data Table Structures

See Creating or Editing a Customer Reference Data Table.

Deleting Reference Data Tables

Note that the Cisco Policy Builder does not have direct access to the database. Therefore, when deletions are made via the Cisco Policy Builder, those items are truly not deleted from the database. Extreme caution should be exercised and users should contact their Cisco Systems technical representative for assistance when deleting tables and columns.

The same caution should be exercised when changing table names.

Final Tasks

Now populate or edit the contents of these table columns. provide them with rows and rows of data.

Cisco Control Center: Populating a Customer Reference Data Table

Note The customer reference data table structures are created and edited in the Cisco Policy Builder interface. See Cisco Policy Builder: Constructing Customer Reference Data Tables for those procedures.

Importing Data from a Spreadsheet

If you have a spreadsheet that already has data rows, you can import the spreadsheet into a customer reference data tables. Of course, the data table itself must be created in Cisco Policy Builder. However, importing the data occurs in Cisco Control Center.

Note When the tables are created in Cisco Policy Builder, have your spreadsheet open to make sure you match the data type, for example string, date, or number.

Step 1 Log in to Cisco Control Center.

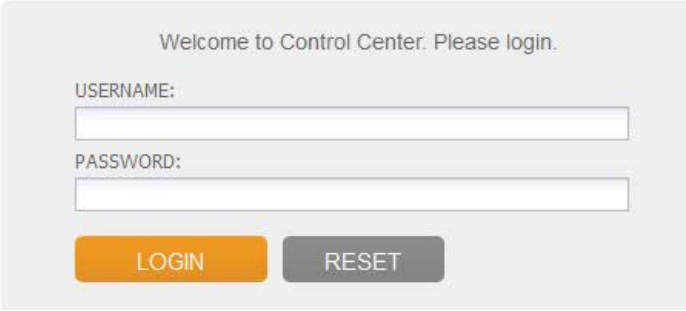

Step 2 Configuration tab > Reference Data folder > table list.

Step 3 Select the table that receives data from the spreadsheet.

The example here uses the AccessTechnology table.

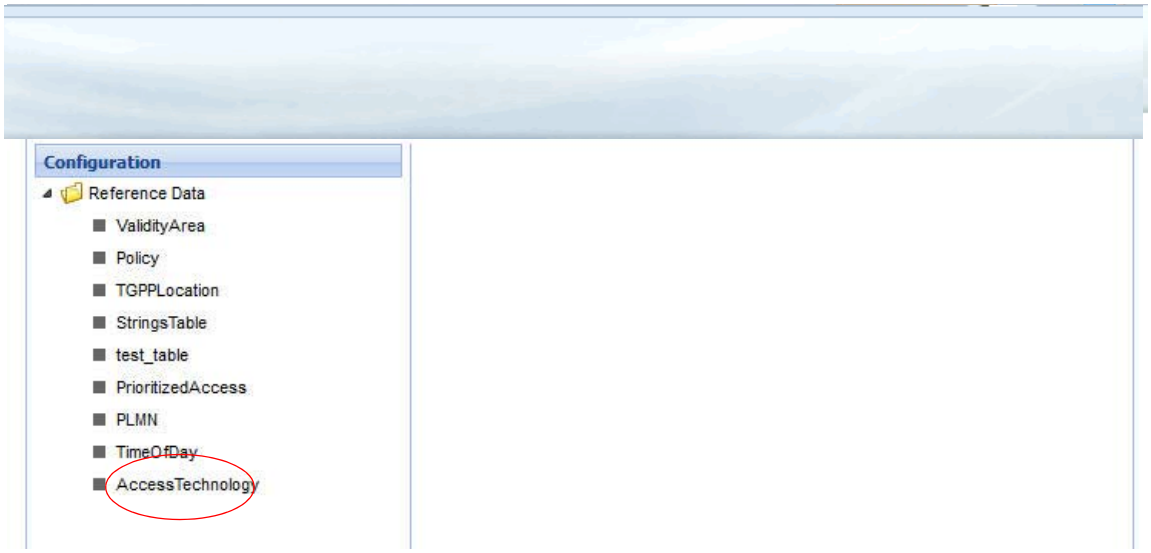

Step 4 Click the import link in the upper right corner.

Step 5 The table you selected under the Reference Data tree appears in the Import in Table field. You do not have to type it in. Make sure this is the table that is to receive data from the spreadsheet.

Step 6 Either name or browse for the .xls or .xlsx file you want to import, the one that populates the table.

Step 7 Click the Next button in the lower corner.

The Map file to Table Columns screen lets you select the columns in the spreadsheet that you want to use in the Access Technology table.

Step 8 Check Exclude first row if your .xls spreadsheet has headers.

Step 9 Check Delete all rows if you want to completely empty the Access Technology table and bring in all new rows.

Step 10 In the Select Column drop-down lists, specify what column of the Access Technology table you want to place your spreadsheet columns in.

Note The Select Column list shows the column names in Access Technology table as they were created in Cisco Policy Builder. This drop down does not show the column names in your spreadsheet.

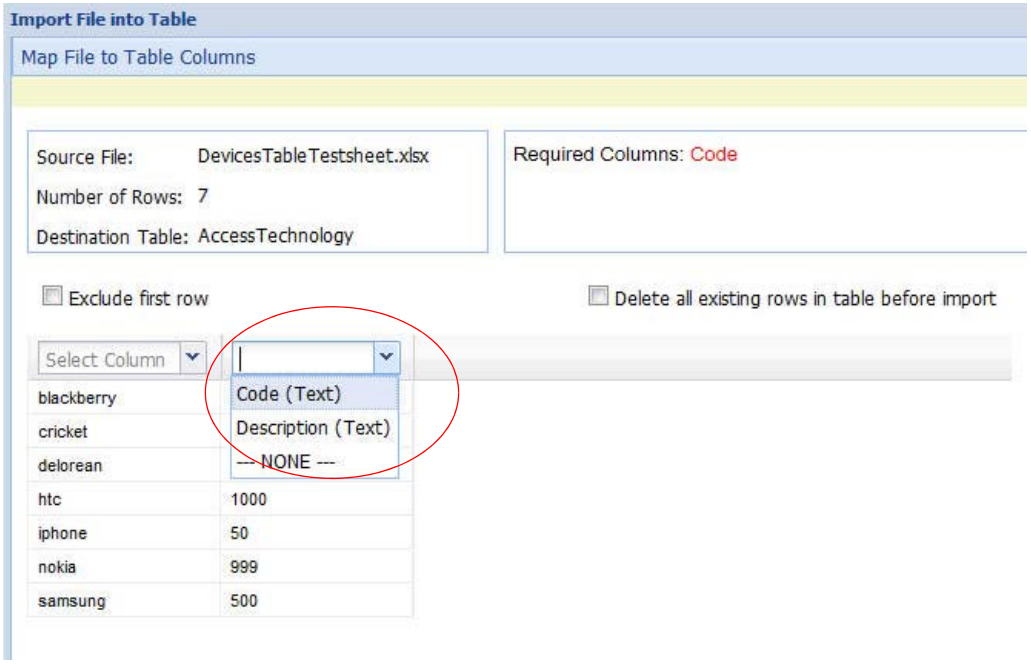

In our example, we want to place the device names in the A String column and the numeric values in the A Valid Values column.

Step 11 In the lower right, click the Import link to proceed, or the Prev link to go back.

Notice the Import Successful screen.

Step 12 Select Open AccessTechnology in the lower left to check your import, then click Done.

The newly populated table appears for you to check.

The AccessTechnology table has the new rows added, in the proper columns.

Entering Data Manually

For data tables that are small, or are not derived from a spreadsheet use the Cisco Control Center interface to enter data manually into the rows of customer reference data tables.

Step 1 Log in to Cisco Control Center.

Step 2 .Click Configuration tab > Configuration node > and open the Reference Data folder.

The items under the Reference Data folder are the table names created in Cisco Policy Builder.

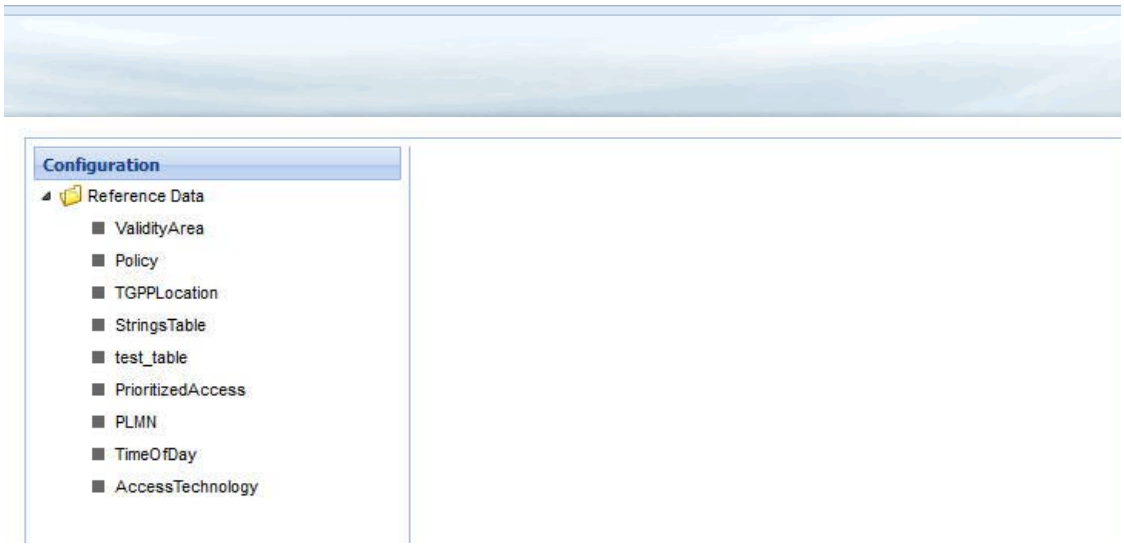

Note If you do not see the table name that you want, go to the Cisco Policy Builder and create it.

Step 3 In the tree, click a table name to open it up.

The table window shows the table structures and columns.

Our example uses AccessTechnology, which has several rows of data already in it.

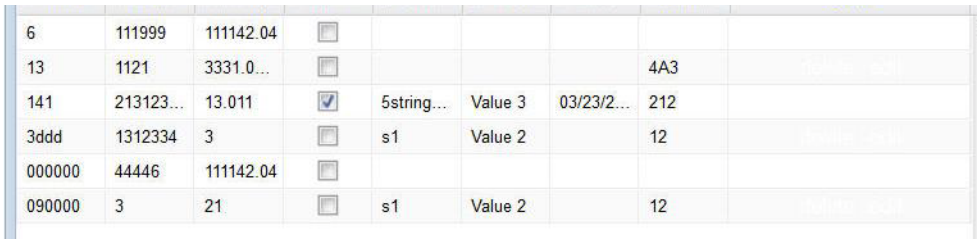

6	111999	111142.04	☐				
13	1121	3331.0...	☐				4A3
141	213123...	13.011	☑	5string...	Value 3	03/23/2...	212
3ddd	1312334	3	☐	s1	Value 2		12
000000	44446	111142.04	☐				
090000	3	21	☐	s1	Value 2		12

Note Add to and edit rows in a table with these tips:

- Click on a row and it becomes editable.
- You can enter stuff in fields, select check boxes, use the drop downs.
- You can build a table of values to use in another table.
- Date fields use the calender widget.
- Save or cancel the row. Save persists the rows to the database.
- Cancel lets the table data revert with no change.
- Carefully delete a row, there is no confirmation.
- Click the add link to add new blank rows.

Adding a Row

Step 1 Open a table.

Step 2 Click the add link in the upper right corner.

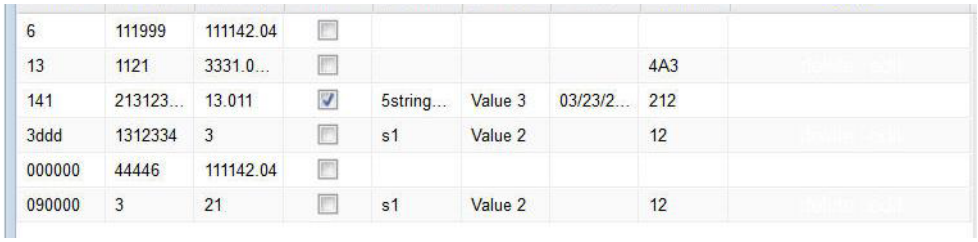

6	111999	111142.04	☐					
13	1121	3331.0...	☐				4A3	
141	213123...	13.011	☑	5string...	Value 3	03/23/2...	212	
3ddd	1312334	3	☐	s1	Value 2		12	
000000	44446	111142.04	☐					
090000	3	21	☐	s1	Value 2		12	

Step 3 Note that the new rows is added to the bottom of the table for you to fill in.

You cannot insert a row at a specific place.

Step 4 Row count appears in lower right corner.

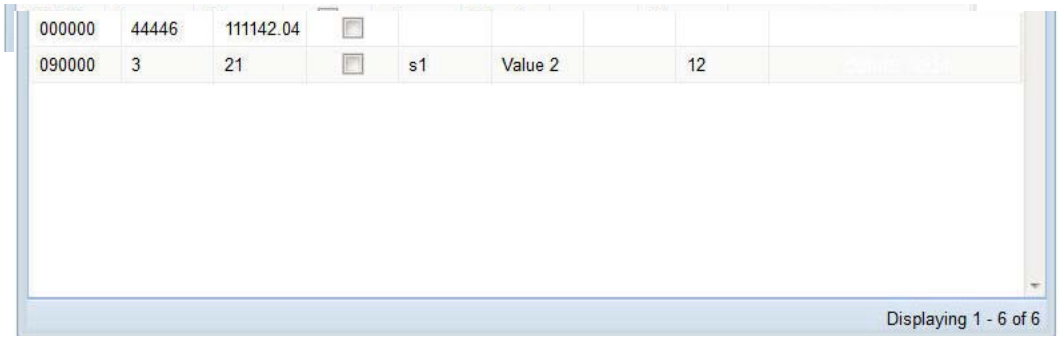

000000	44446	111142.04	☐				
090000	3	21	☐	s1	Value 2		12

Displaying 1 - 6 of 6

Note In this interface, the columns do not indicate which are Key columns or Key fields. Best practice is to name the columns with this attribute when creating them in Cisco Policy Builder.

Note Any required fields display with red margins to let you know that you must enter data.

Step 5 Use any of the save methods mentioned above to save the row.

However, you cannot save a row that has errors.

Editing a Row

There are several ways to save row data and so prevent data loss.

- Click the save link in the column on the far right.
- Press Enter after you have finished in a field.
- Tab over from one field to another until the row is saved.
- Finish in your field and click on another row.

There is no way to promote or demote rows.

There is no way to sort on columns.

There is no filter or search feature at this time.

Fixing Errors in a Row

Errors are denoted as the field margined in red.

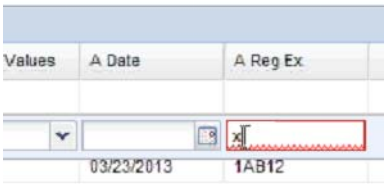

Help text displays to the far left or far right.

Only the cancel link is available until you make a correction.

Deleting a Single Row

Step 1 In the Cisco Control Center interface, choose your table.

Step 2 Hover over the row you want to delete.

Step 3 Click the delete link to the right

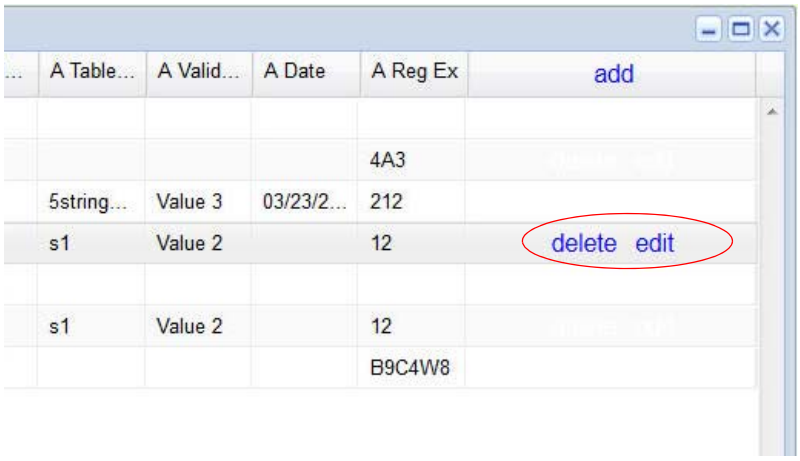

The data row is deleted immediately. No refresh needs to occur.

Typical Tasks for Everyday

These navigation tasks and activities are specific to the customer reference data tables in Cisco Control Center. For other every day type of interface tasks see Tips and Best Practices.

Refreshing the Screen

- The refresh link in the upper right corner.

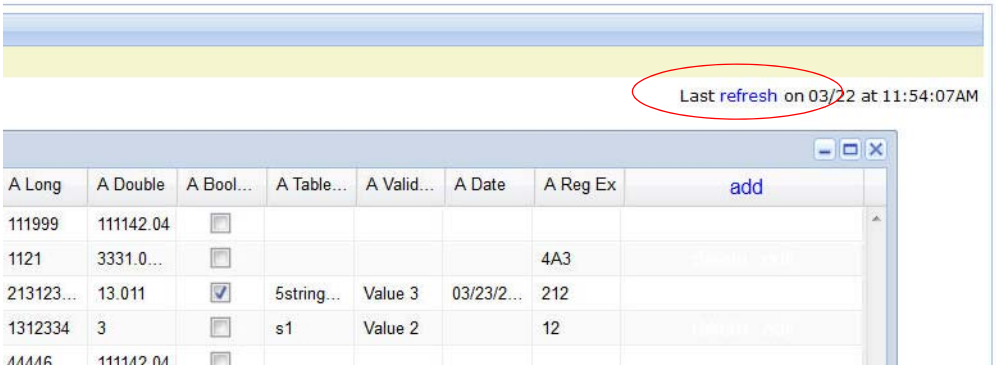

Click the refresh link to force a refresh.

- A successful refresh message appears on the left.

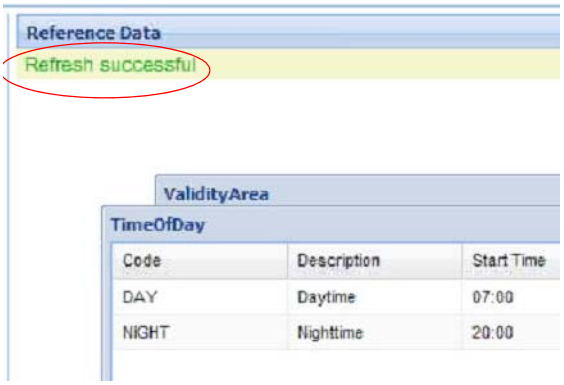

The usual refresh time is every 5 minutes, and is refreshed from the server.

- In the middle of updating, refresh is deferred.
- Upon close and reopen a table, the refreshed table is displayed.
- The refresh interval cannot be changed

Navigating the Table Screens

- Multiple tables can be open at one time.
- Icons in the tree to show the open and active focus.
  - Circle with hyphen—open table
  - Circle with arrows—active table focus
  - Square—table not open
  - Blue square—table open and in the tray at bottom

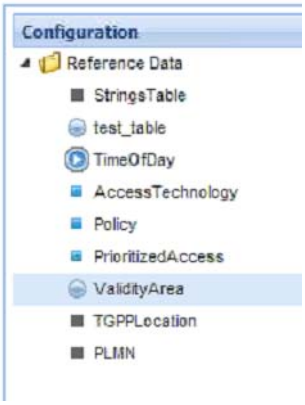

- Icons in the upper right corner of a table minimize, maximize, or close the table.

  When minimized, the table heading displays at the bottom of the screen in the tray.

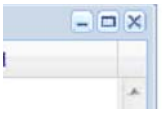

- To restore a table to view,
  - click the restore button of the table in the tray,
    or
  - double click the table tab in the tray.
- Resize a table for viewing by dragging the edges.
- Narrow or widen a column by dragging the margin of the column head.

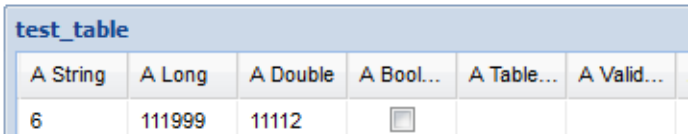

Navigating in a Row

- If you have many columns, they cannot all display in the browser. You must scroll to the far right to get to the delete, edit, save, cancel links.

Notice that in the bottom right corner, there is a indicator of how may rows you are viewing at the moment, and how many rows there are total in the table.

Customer Reference Data APIs Usage

This section covers the following topics:

Introduction

The Customer Reference Data (CRD) APIs exists to allow query, creation, deletion, and update of CRD table data without the need to utilize the Control Center application. The CRD APIs are available via a REST interface. The specific APIs are Query, Create, Update, and Delete.

The CRD component became available in CPS 5.5, but the official CRD APIs are not available until CPS 6.1.2 and higher.

Limitations

These APIs allow maintenance of the actual data rows in the table. They do not allow the creation of new tables or the addition of new columns. Table creation and changes to the table structure must be completed via the Policy Builder application.

Table names must be all in lowercase alphanumeric to utilize these APIs. Neither spaces nor special characters are allowed in the table name.

- Table names containing uppercase characters will return code 400 Bad Request.
- Spaces in the name are also not allowed and will be flagged as an error in Policy Builder.
- Special characters even when escaped or encoded in ASCII cause problems with the APIs and should not be used.

Setup Requirements

This section covers the following topics:

Policy Server

The feature `com.broadhop.custrefdata.service.feature` needs to be installed on the Policy Server.

In a High Availability (HA)/Distributed CPS deployment, this feature should be installed on the PCRF (QNS0x) nodes. No feature is needed on the IO Manager nodes.

Policy Builder

The feature `com.broadhop.client.feature.custrefdata` needs to be installed in Policy Builder.

Step 1 Login to the Policy Builder.

Step 2 In Reference Data > Systems > Plugin Configurations (or a sub cluster or instance), a Customer Reference Data Configuration plugin configuration is defined.

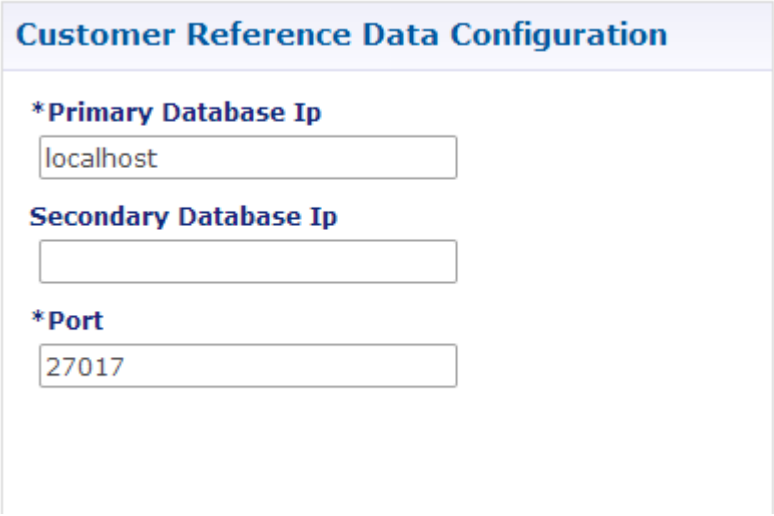

Parameter	Description
Primary Database Ip	The IP address or DNS name of the MongoDB instance that stores the CRD tables and data. Default value is blank, but must be specified for proper function.
Secondary Database Ip	The IP address or DNS name of the secondary (slave) MongoDB instance that stores the CRD tables and data. Default value is blank.
Port	The port instance of MongoDB that contains the CRD tables and data. The default value is 27717; however, the recommended value is 27720 which is the standard location for these tables in an HA/Distributed CPS deployment. The value should be changed to 27017 for an All-In-One/Single Virtual Machine (VM) deployment.

Step 3 In Reference Data > Customer Reference Data Tables, at least one Custom Reference Data Table must be defined.

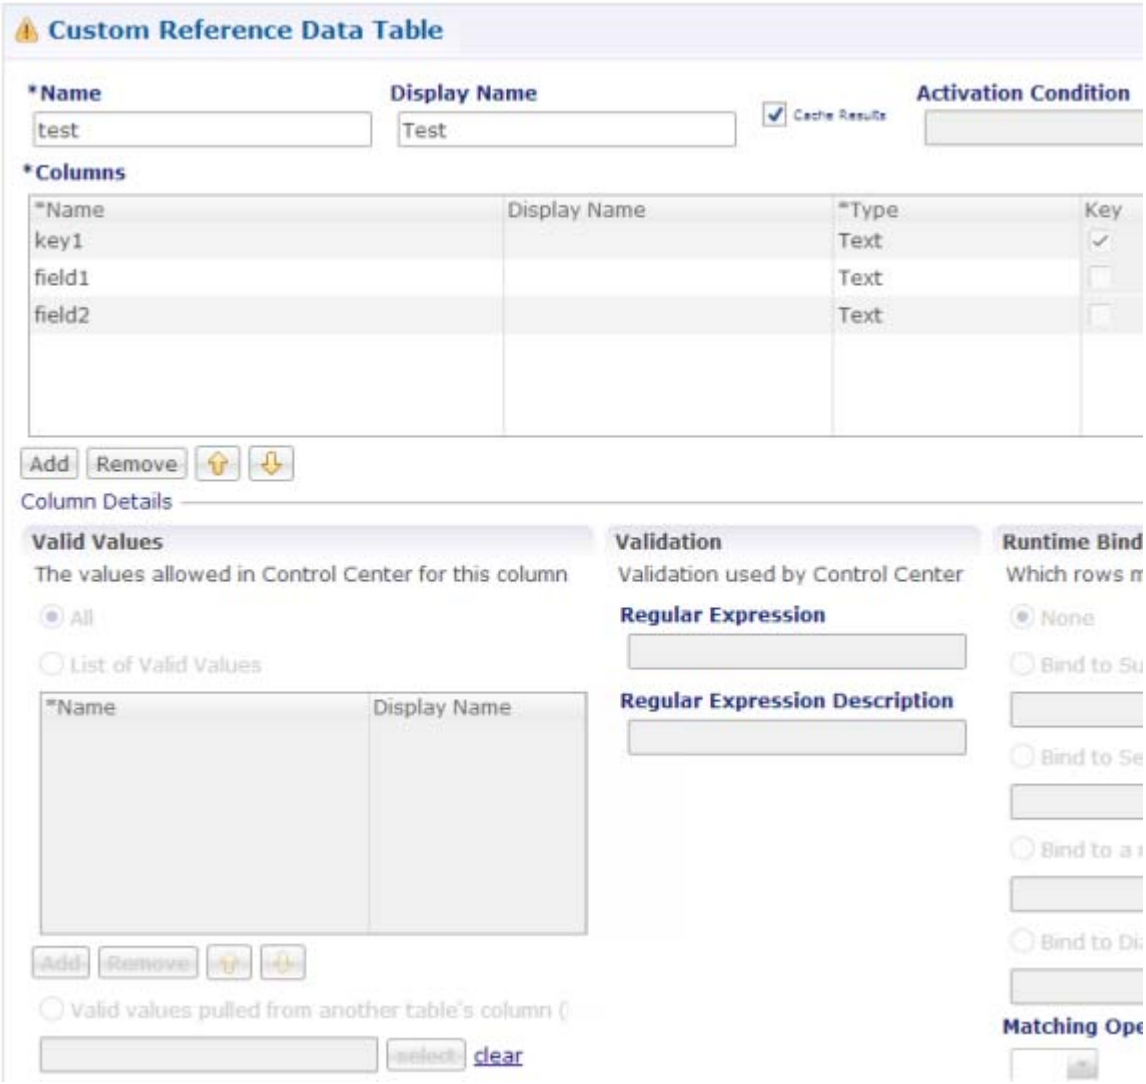

Parameter	Description
Name	This name is used in the endpoint of the APIs so that the API knows which table you wish to take action on. The table name must all be lowercase alphanumeric and must not contain spaces or special characters. Refer to Limitations for more information.
Columns	
Name	This name is used for the Code field in the XML that is submitted as part of the Create and Update APIs.

Parameter	Description
Type	This type defines the data type that the Value field in the XML has that is submitted as part of the Create and Update APIs. Types supported are: • Date • DateTime • Decimal • Number • True/False
Key	Specifies this row is the key field in the table and is the value that should be used for the key code XML tag.
Required	Specifies which columns, or codes in the XML, must be specified to properly populate the row in the CRD table.
Cache Results	Indicates whether or not caching should be applied to this table; that is, whether the table is loaded once from the database and cached until CPS restart or a write occurs to the CRD table versus the database being queried every time the data is needed. Refer to Caching under Architecture for more information on how and when to use this setting.

Architecture

This section covers the following topics:

- MongoDB
- Caching

MongoDB

The MongoDB database containing the CRD tables and the data is located in the MongoDB instance specified in the CRD plugin configuration.

The database is named customer_ref_data.

Two system collections exist in that database and do not actually contain CRD data:

- system.indexes — used by MongoDB. These are indices set on the database.
- crdversion — contains a document indicating the version of all the CRD tables you have defined. The version field increments by 1 every time you make a change or add data to any of your CRD tables.

A collection is created for each CRD table defined in Policy Builder.

- This collection contains a document for each row you define in the CRD table.
- Each document contains a field for each column you define in the CRD table.
- The field contains the value specified for the column for that row in the table.
- Additionally, there is a _id field which contains the internal key used by MongoDB and _version which is used by CPSCPS to provide optimistic locking protection, essentially to avoid two threads overwriting the other's update, on the document.

For example,

```
MongoDB shell version: 2.4.10
connecting to: test
> show dbs
balance_mgmt    0.203125GB
cust_ref_data   0.203125GB
local   0.078125GB
policy_trace    1.203125GB
portal  0.203125GB
radius  0.203125GB
session_cache   0.203125GB
sharding        0.203125GB
spr     0.203125GB
> use cust_ref_data
switched to db cust_ref_data
> show collections
crdversion
system.indexes
test
> db.test.find()
{ "_id" : ObjectId("53e63469a074572ba1b5e1bd"), "_version" : 1, "field2" : "field2examp
{ "_id" : ObjectId("53e634a9a074572ba1b5e1be"), "_version" : 1, "field2" : "field2examp
{ "_id" : ObjectId("53e64be2a074572ba1b5e1bf"), "_version" : 1, "field2" : "testee", "k
>
```

Caching

Setting the Cache Results to true (checked) is the default and recommended settings in most cases as it yields the best performance. Use of the cached copy also removes the dependency on the availability of the CRD database, so if there is an outage or performance issue, policy decisions utilizing the CRD data won't be impacted.

The cached copy of the table is refreshed on CPS restart and whenever the API writes a change to the CRD table, otherwise the cached copy is used and the database is not accessed.

API Endpoints and Examples

http://<IP or DNS name>:8080/custrefdata/<tablename>/_<operation>

This section covers the following topics:

- Query API
- Create API
- Update API
- Delete API
- Tips for Usage

Query API

Purpose

Returns all rows currently defined in the specified table.

HTTP Operation Type

GET

Example URL

http://localhost:8080/custrefdata/test/_query

Example URL with Filtering

http://localhost:8080/custrefdata/test/_query?key1=Platinum

Payload

None, although parameters can be specified on the URL for filtering

Response

Success returns code 200 Ok; XML indicating rows defined is returned. If the table does not exist, code `400 Bad Request` is returned.

Example Response without Filtering

```
<rows>
   <row>
      <field code="field1" value="1004"/>
      <field code="field2" value="testee"/>
      <field code="key1" value="Platinum"/>
   </row>
   <row>
      <field code="field1" value="1004"/>
      <field code="field2" value="testee"/>
      <field code="key1" value="Platinum99"/>
   </row>
   <row>
      <field code="field1" value="field1example1"/>
      <field code="field2" value="field2example1"/>
      <field code="key1" value="key1example1"/>
   </row>
   <row>
      <field code="field1" value="field1example2"/>
      <field code="field2" value="field2example2"/>
      <field code="key1" value="key1example2"/>
   </row>
</rows>
```

Example Response with Filtering

```
<rows>
   <row>
      <field code="field1" value="1004"/>
      <field code="field2" value="testee"/>
      <field code="key1" value="Platinum"/>
   </row>
</rows>
```

The response returns keys with the tag "field code". If you want to use the output of Query as input to one of the other APIs, the tag needs to be changed to "key code". Currently using "field code" for a key returns code `404 Bad Request` and a `java.lang.NullPointerException`.

Create API

Purpose

Create a new row in the specified table.

HTTP Operation Type

POST

Example Endpoint URL

http://localhost:8080/custrefdata/test/_create

Example Payload

```
<row>
   <key code="key1" value="Platinum"/>
   <field code="field1" value="1004"/>
   <field code="field2" value="testee"/>
</row>
```

Response

Success returns code 200 Ok; no data is returned. The key cannot already exist for another row; submission of a duplicate key returns code `400 Bad Request`.

Update API

Purpose

Updates the row indicated by the key code in the table with the values specified for the field codes.

HTTP Operation Type

POST

Example Endpoint URL

http://localhost:8080/custrefdata/test/_update

Example Payload

```
<row>
   <key code="key1" value="Platinum"/>
   <field code="field1" value="1005"/>
   <field code="field2" value="tester"/>
</row>
```

Response

Success returns code 200 Ok; no data is returned. The key cannot be changed. Any attempt to change the key returns code `404 Not Found`.

Delete API

Purpose

Removes the row indicated by the key code from the table.

HTTP Operation Type

POST

Example Endpoint URL

http://localhost:8080/custrefdata/test/_delete

Example Payload

```
<row>
   <key code="key1" value="Platinum"/>"/>
</row>
```

Response

Success returns code 200 Ok; no data is returned. If the row to delete does not exist, code `404 Not Found` is returned.

Tips for Usage

The Query API is a GET operation which is the default operation that occurs when entering a URL into a typical web browser.

The POST operations, Create, Update, and Delete, require the use of a REST client so that the payload and content type can be specified in addition to the URL. REST clients are available for most web browsers as plug-ins or as part of web service tools, such as SoapUI. The content type when using these clients should be specified as application/xml or the equivalent in the chosen tool.

Notification Configurations

Revised: July 10, 2015

Notification in Cisco Policy Builder relates to pushing messages from Cisco Policy Builder to subscribers. Use messages to alert the subscriber to issues as well as opportunities on their network. Not only can you alert subscribers, but you can also send messages to any address you wish, perhaps system monitoring addresses.

Currently, Cisco Policy Builder offers following notification types:

- Apple iOS devices/iPhone® push (iOS devices)
- Email (IMAP only)
- SMS notification (SMPP v 3.4)
- Realtime Notification

Note You can configure one or all notification types. No notification configurations are set up for your deployment until you do so.

This appendix covers the following sections:

Notifications - An Introduction

Setting up notifications to your subscribers has three parts:

- Set up connectivity

  Set up the connections to the external agents and servers that communicate with the subscriber, see Set Up Connections.

- Create messages of your own.

  Create the messages and notices you want to send, see Subscriber Notifications.
- Define a policy action to send the notification

  Make notification an action result for a condition in a policy that evaluates as true, see Create the Policy Action.
- Subscriber Event Processing Framework

  Subscriber Event Processing Framework enables the Cisco Policy Suite to perform the following:
  - Offline Notifications: Notify subscribers of upcoming events, e.g. “your plan expires in 24 hours”, when the subscriber does not have an active network session in the CPS.

The information related to notifications is provided in the following feature files:

- cat /etc/broadhop/pb/features
- com.broadhop.client.feature.notifications
- cat /etc/broadhop/pcrf/features
- com.broadhop.notifications.service.feature
- com.broadhop.notifications.local.feature

Note This configuration is for CPS 5.3.5 and higher.

When the Notifications component is deployed in a distributed environment, the three notifications component features should be installed as follows:

- QNS0x VMs (pcrf_A and pcrf_B)
  - com.broadhop.notifications.local.feature
- Policy Builder (pb)
  - com.broadhop.client.feature.notifications
- IOMGR0x
  - com.broadhop.notifications.service.feature

Set Up Connections

The first part of setting up subscriber notifications is to define the connections to the various servers that communicate with the subscriber.

These steps explain how to set up connectivity for Cisco Policy Builder to send notifications using an external server such as an email server or SMSC.

These steps configure the plug-in configuration for notifications. This configuration defines the connections to the appropriate servers for email, Apple iOS, or SMS.

Step 1 Click Reference Data > Systems > a defined system > Plugin Configurations.

Note You can also go one level deeper to define notifications at the cluster level. Click on a cluster name to create a Notification Confirmation that applies to that cluster only.

Step 2 Open the node in the tree.

Step 3 In the detail pane, click Notification Configuration.

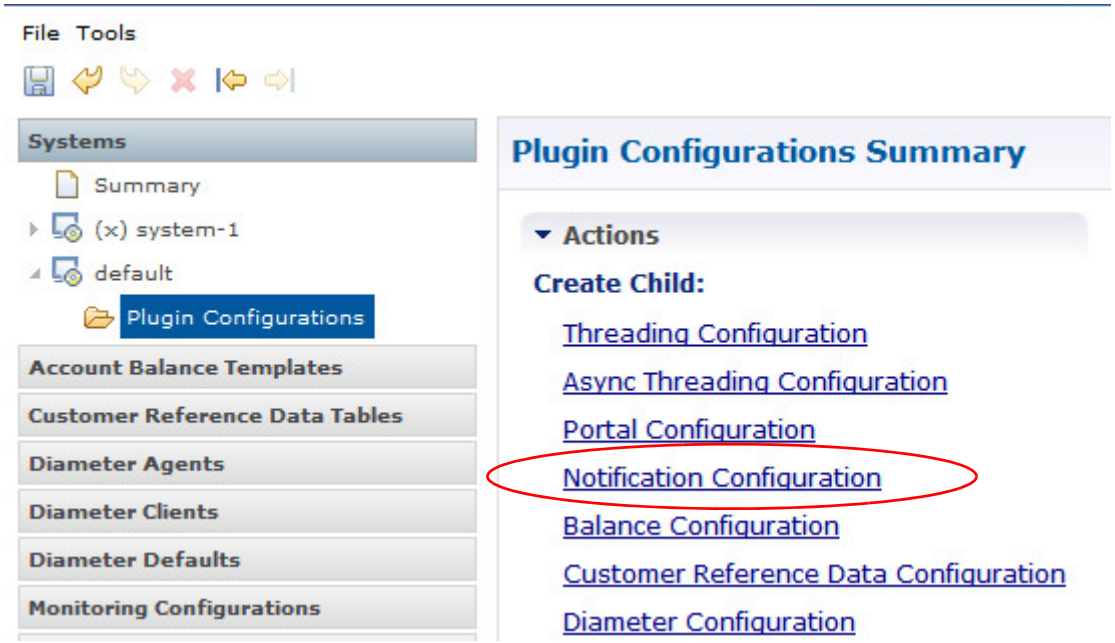

Step 4 Click the new Notifications Configuration node in the tree.

Step 5 Select a check box for the type of notification you want to set up.

You can use one, two, or all of the notification types if you need to.

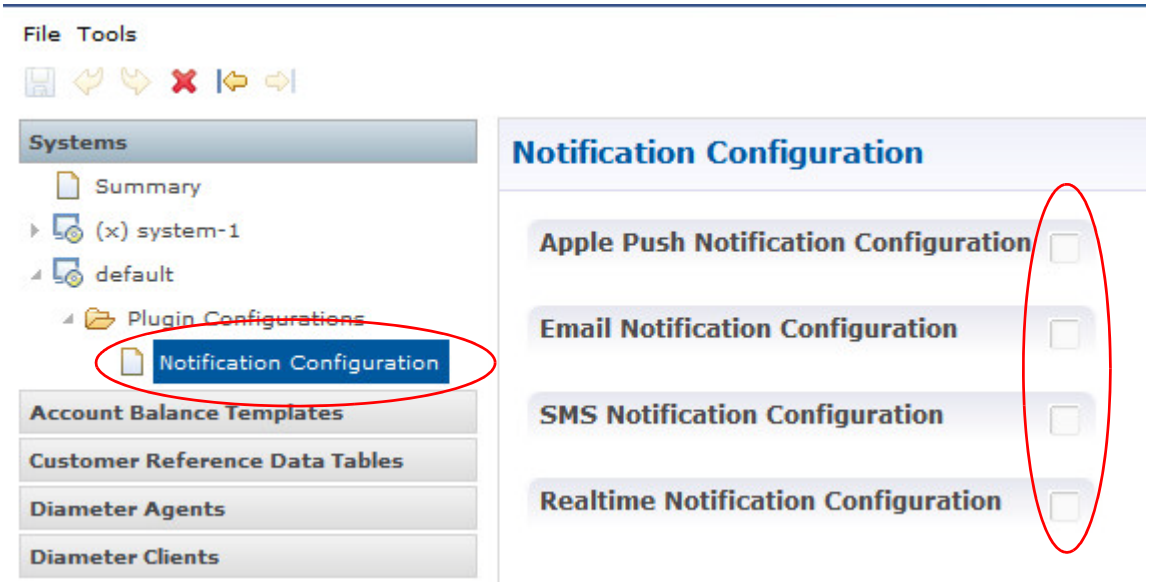

Step 6 To configure the connection for a push to an Apple iOS device or iPhone, go to the section titled Configure Apple iOS Device/iPhone Notification.

Step 7 To configure the connection for an email notification, go to the section Configure Email Notification.

Step 8 To configure the connection for a text message, go to Configure SMS Notification.

Step 9 To configure the connection for realtime notification, go to Configure Realtime Notification.

Configure Apple iOS Device/iPhone Notification

Use this procedure to send a message to a subscriber with an Apple iPhone or other iOS device.

Step 1 From the main screen, select the Reference Data tab > Systems node > a system or a cluster > Plugin Configurations > Notification Configuration.

Step 2 Click the check box next to Apple Push Notification Configuration.

Step 3 View the Notification Configuration screen that drops down.

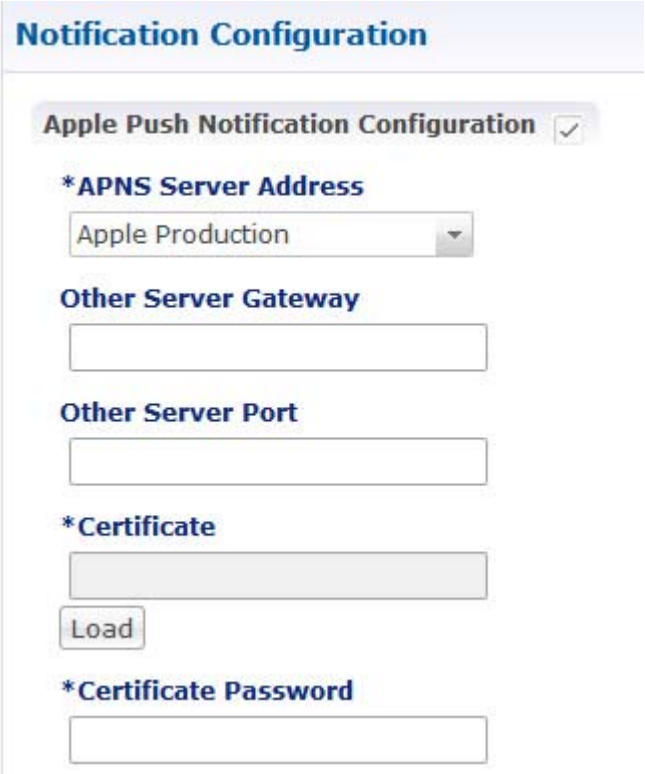

APNS Server Address	• Apple Production—Connects to gateway.push.apple.com on port 2195 • Apple Test —Connects to gateway.sandbox.push.apple.com on port 2195. **Recommended Action** Other—use a server address other than the standard Apple ones, and uses the other gateway and server port fields below.
Other Server Gateway	Name of a gateway if Other selected.
Other Server Port	Port number if Other selected.
Certificate	The certificate of authorization to the gateway. You must provide a certificate file that is loaded into CPS.

Step 4 Go to Messages for the Apple iOS and iPhone® Notification.

Configure Email Notification

Use this procedure to configure an email notification from Cisco Policy Builder to the subscriber.

Step 1 From the main screen, select the Reference Data tab > Systems node > a system or a cluster > Plugin Configurations > Notification Configuration.

Step 2 Click the check box next to the label Email Notification Configuration.

Step 3 View the Notification Configuration screen for email notification that drops down.

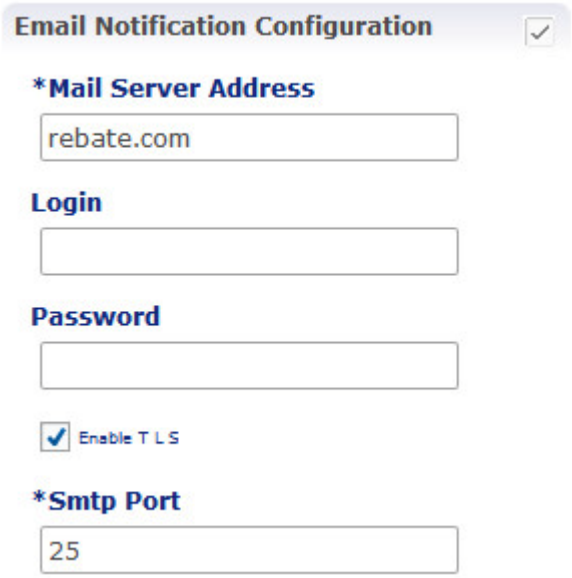

Mail Server Address	URL of the mail server that the Cisco Policy Builder email goes through. Only IMAP email is supported at this time.
Login Password	Enter any login and password information needed.
Enable TLS	Enables transport layer security. This option is used for connecting to services other than basic IMAP services, such as Google's Gmail.
Smtp Port	Specifies the SMTP port. This option is used for connecting to services other than basic IMAP services, such as Google's Gmail.

Step 4 Go to Messages for Email Notification.

Configure SMS Notification

Step 1 From the main screen, select the Reference Data tab > Systems node > a system or a cluster > Plugin Configurations > Notification Configuration.

Step 2 Click the check box next to SMS Notification Configuration.

Step 3 View the Notification Configuration screen for SMS notification that drops down.

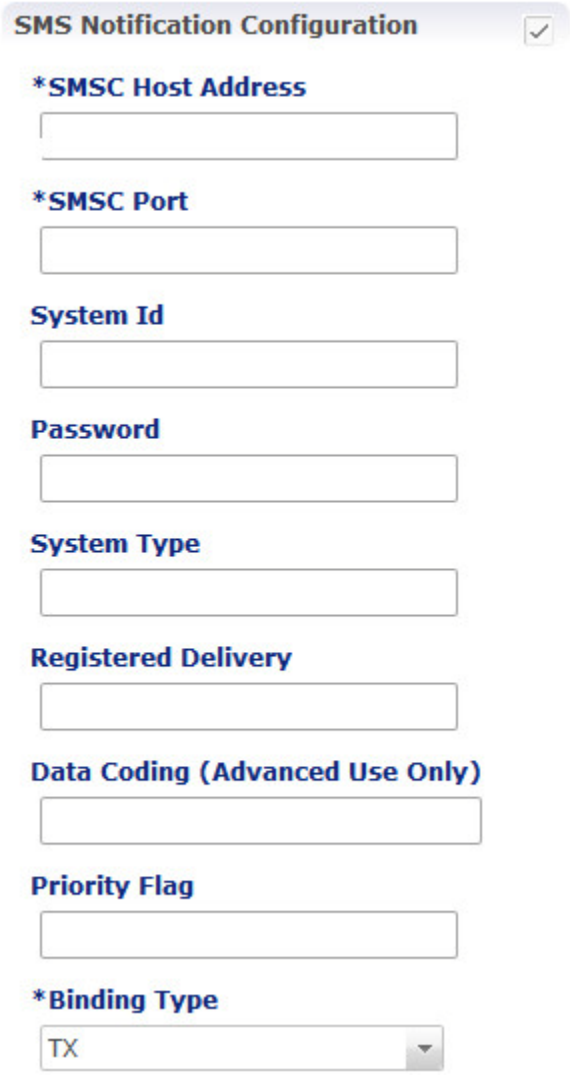

Parameter	Description
SMSC Host Address	The TCP/IP address or host name of the SMPP server, that is, the URL of the SMSC host that pushes the SMS message.
SMSC Port	The TCP/IP port on the SMPP server to which the gateway connects.
System ID	The user name for the gateway to use when connecting to the SMPP server
Password	The password for the gateway to use when connecting to the SMPP server.
System Type	An optional login parameter used only if required by the SMPP server. The SMPP system administrator provides this value, usually a short text string.
Registered Delivery	Optional field for some custom deployments.

Parameter	Description
DataCoding (Advanced Use Only)	Optional field for some custom deployments. Data Coding can be used instead of the Message Encoding and the other DCS fields on the Notification message definition screen, but should be used with care. See the details in section Addendum A: Data Coding. **Note** If it is necessary to use a specific value in this field for data coding, then the message alphabet information should be included in the Data Coding value per the SMS specification as well as any other necessary data coding information. The result is a combination of the capabilities of the SMSC and is not totally controlled by CPS. Please refer to your specific Core version below in the Addendum A: Data Coding section for more information on the functionality and behavior.
Priority Flag	Optional field for some custom deployments.
Binding Type	TX (transmit) or TRX (transceiver)

Go to Messages for SMS Notification.

Configure Realtime Notification

Realtime Notifications allows you to send SOAP/XML messages to a defined server when policy thresholds are breached. The information related to realtime notification is provided in the following feature files:

- **For AIO Setup:**

  In /etc/broadhop/pb/features:

  ```
  com.broadhop.client.feature.notifications
  ```

  In /etc/broadhop/pcrf/features:

  ```
  com.broadhop.notifications.local.feature
  com.broadhop.notifications.realtime.service.feature
  com.broadhop.notifications.service.feature
  ```

- For HA Setup:

  Place the `com.broadhop.notifications.realtime.service.feature` in `/etc/broadhop/iomanager01/02`.

  The other additions for Policy Builder and PCRF features files are same as in AIO.

  – Update CPS to finalize the install. Make sure to execute synchronization command before updating if HA:

  syncconfig.sh

– Execute the following commands in HA setup:

```
/opt/broadhop/installer/support/set_permissions.sh

su qns -c update_qns.sh –f

restartall.sh
```

The following will be installed in CPS 7.0:

```
[root@lab pcrf]# list_installed_features.sh | grep notific

com.broadhop.notifications.local.feature=5.8.0.release

com.broadhop.notifications.realtime.service.feature=5.8.0.release

com.broadhop.notifications.service.feature=5.8.0.release
```

Step 1 From the main screen, select the Reference Data tab > Systems node > a system or a cluster > Plugin Configurations > Realtime Notification Configuration.

Step 2 Click the check box next to Realtime Notification Configuration.

Step 3 View the Notification Configuration screen for realtime notification that drops down.

Parameter	Description
Failed XML Directory	File system path where failed notifications are stored. So when CPS is not able to send notification on both HTTP URL and HTTP Fallback URL then that notification is stored in this path. The path to the failed XML directory needs to be created manually on lb's (lb01 and lb02).
Max Storage allowed for failed XMLs (in MB)	Maximum size up to which CPS can store failed notifications in the XML failed directory.

Step 4 Login to Policy Builder and select Services tab.

Step 5 To configure Realtime Notification message, refer to Message for Realtime Notification.

Step 6 Under Use Case Templates, select Summary and click Use Case Template from right pane to open a new window.

Step 7 Enter the name in Name field and click Use Case Initiators tab.

Step 8 Add the Service Initiators. An example is shown.

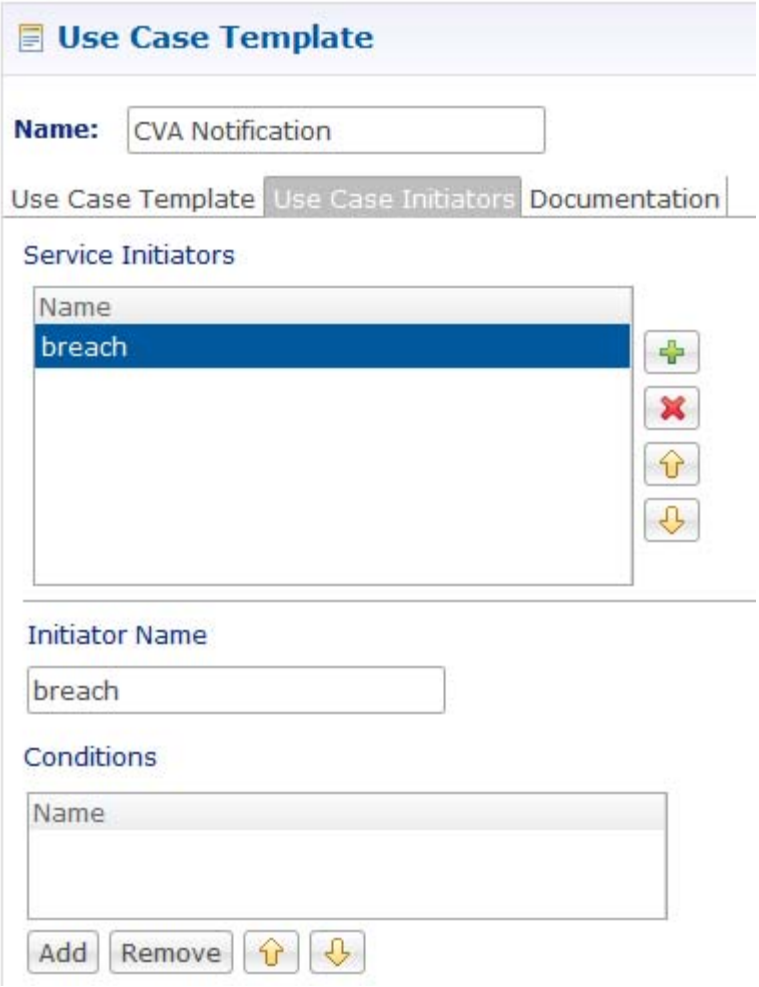

Step 9 Under Conditions pane, click Add to open a dialog box.

Step 10 Select the required condition and click OK to add it under Conditions pane. An example is shown below.

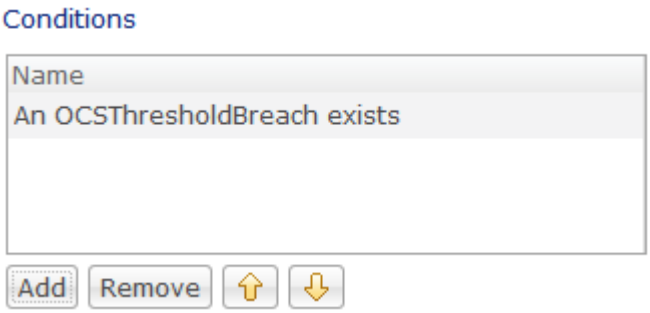

Step 11 Select the conditions and configure the attributes according to your requirement. An example is shown.

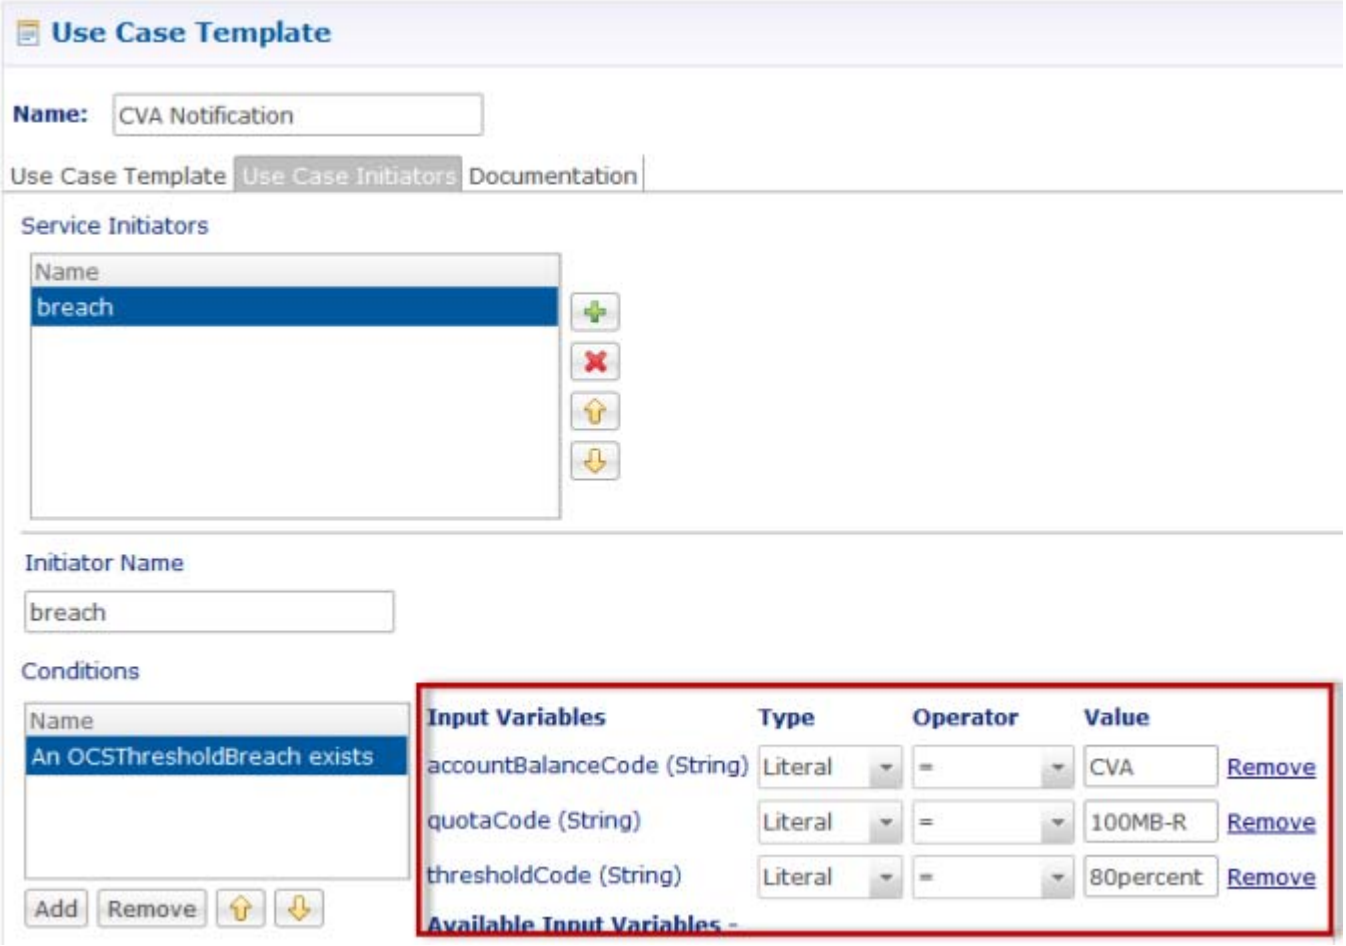

Step 12 Select Use Case Template tab and click Add under Service Configuration to add a service. An example is shown.

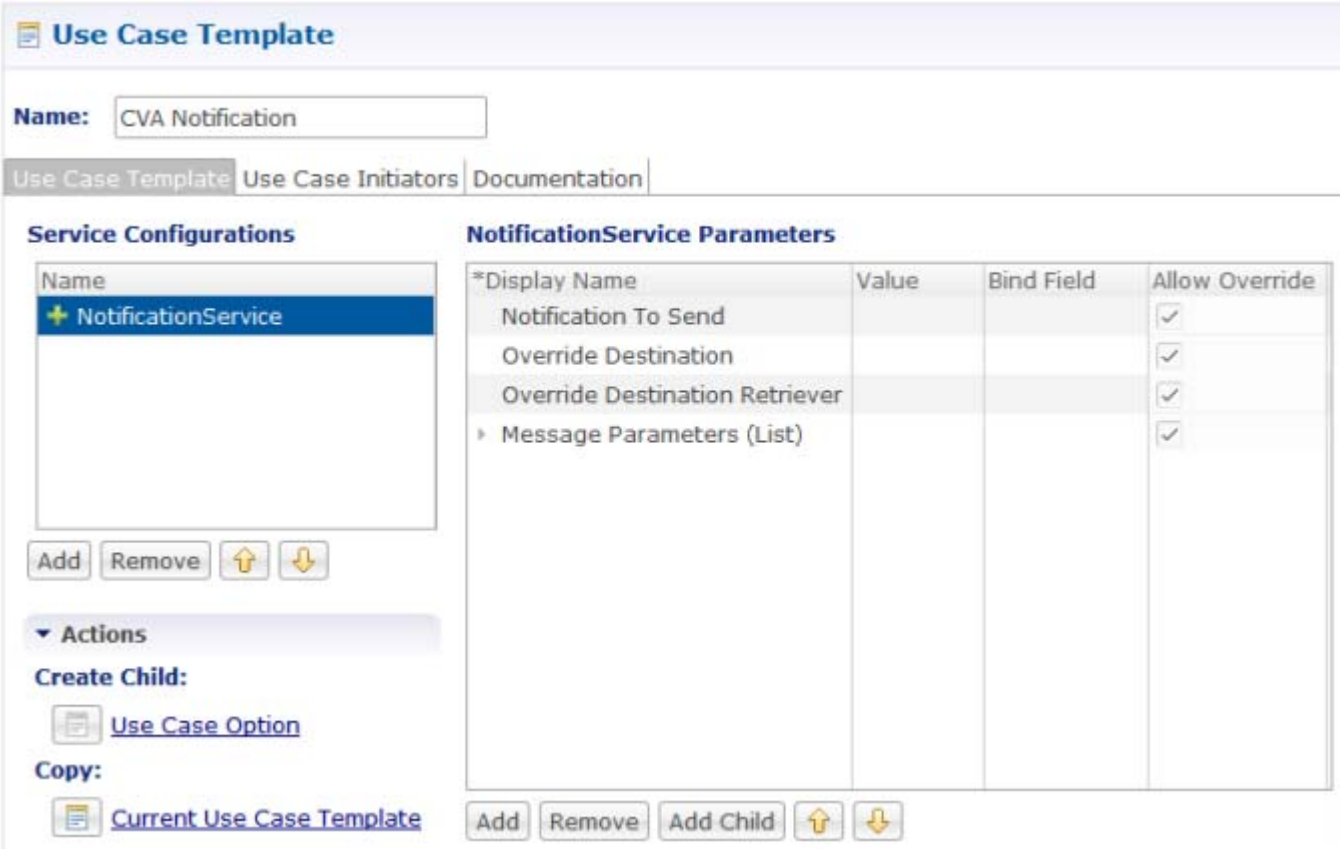

Step 13 Under Services, create a Service Option which is configured with variables that can be added to the notification XML Template. An example is shown.

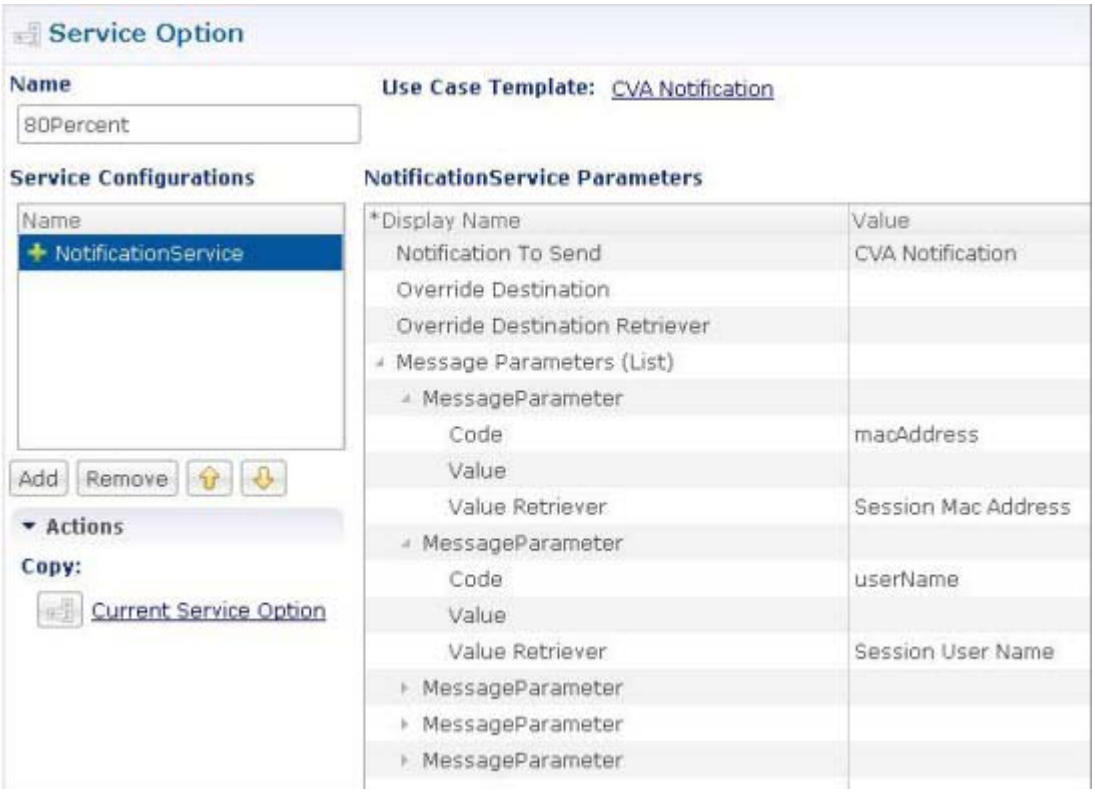

Subscriber Notifications

Note For other information about configuring notifications to subscribers, see Subscriber Configuration.

The second part of setting up subscriber notifications is to create standard, reusable messages. Use this section to create the text in the messages your subscriber receives.

Step 1 Click Reference Data > Notifications.

The Notification node in the Reference Data tree displays the notification types available for configuration.

Step 2 Select the notification types you need, based on the notification type you selected in Set Up Connections.

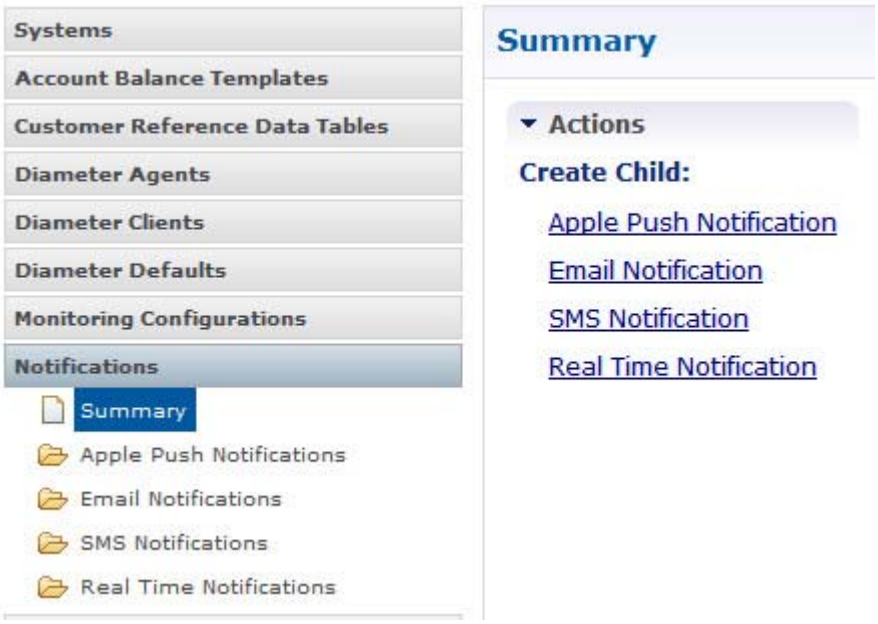

- To write a message for an Apple iPhone or other iOS device, go to Messages for the Apple iOS and iPhone® Notification.
- For an email message, go to Messages for Email Notification.
- For a text message notification, go to Messages for SMS Notification.
- For realtime notification, go to Message for Realtime Notification.

  Some notifications, like SMS and Apple push, have a limit to their length. After you have entered the maximum number of characters, the rest of the message you create here is manipulated in one of two ways.

- Apple Push is truncated.
- SMS is split into multiple messages that are submitted to the SMSC individually. The header information contains multi-part message information if your SMSC supports multi-part messaging. The length of each SMS message is determined by the selected encoding and character set.
- Emails are not limited in size.

  The size limits are important to consider when you include substitution values in your message using the $ notation. Such substitutions affects the size of the message.

  For example if you use a $Name variable in the body of your message, and the subscriber has a very long name as found in their log file for session information, the number of letters in the $Name substitution are counted against the maximum capacity of the message, perhaps causing truncation. More typically, this can occur when there are several to many substitution variables found in a message.

  Go to Define the Policy Action.

Messages for the Apple iOS and iPhone® Notification

Use this procedure to create the messages for a subscriber's Apple iPhone or the Apple iOS operating system.

Step 1 Select Reference Data > Notifications node > Apple Push Notifications.

Step 2 Fill out the fields in the detail pane.

Parameter	Description
Name	The name of the notification message. This name is used in the action phrase in the policy definition. Best practice is to make this short, but meaningful and unique.
Alert	This is the text that appears on the subscriber's iPhone. If the message is too long, it is simply truncated. Test your messages before you place into production. If you want to use a string, and substitute session information, use the syntax $Name, for example, to insert the receiver's name in the email. See Define the Policy Action, step 11. When a policy action phrase uses a notification with a substitution value, it requires a Setup Action. See Actions Subtab on the Policy Screen. Alerts are limited to 160 characters. Alerts longer than that are truncated.

Step 3 Create policy action as described in Create the Policy Action.

Messages for Email Notification

The Cisco Policy Suite subscriber notifications has two methods:

- Online notification that occurs when the subscriber is logged in

- Offline notifications and SPR Cleanup

Subscriber notification lets you define event and action pairs in a similar fashion to the existing policies.These pairs are then evaluated on a defined schedule. Events are similar to a policy's conditions, and actions are like a policy's actions.

As an example, some events and their resulting actions are listed below.

Event:

- Subscriber's quota on balance code Data expires in 1 hour.
- Subscriber has Service A assigned.

Action:

- Send an SMS Notification to the subscriber.
- Send an email to the subscriber.

Event:

- Subscriber has no services assigned in SPR.

Action:

- Delete the subscriber from SPR.

Note Cisco Policy Suite's Notification Center to support MIME multipart/alternative message formatting. Multipart message encoding enables the sending of both an HTML version and a plain text version in a single multipart/alternative message. The user's email client decides the best version to display based on the recipient's individual email client capabilities and settings.

Steps

Use this procedure to create notification messages for email delivery. This message definition is a template. The email address to send to is not defined here. The address to send to is defined in the policy that sends the email, which is defined later in Policies. See Create the Policy Action.

Step 1 Select **Reference Data** tab > Notifications node > Email Notifications.

Step 2 Fill out the fields in the detail pane.

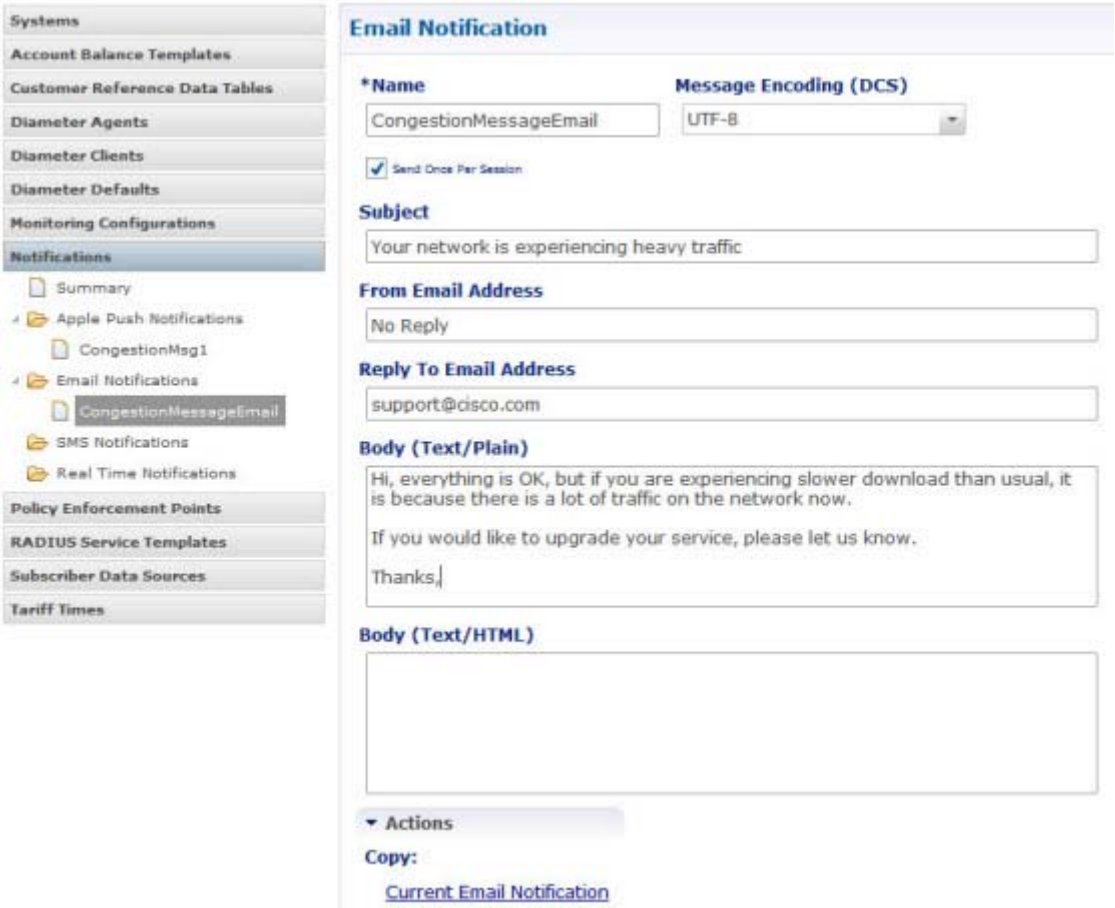

Parameter	Description
Name	Name of the message.
Message Encoding (DCS)	Select the required message coding from drop-down list. Valid values are ISO-8859-1, US-ASCII, UTF-16 (UCS-2) and UTF-8. Default value is UTF-8.
Subject	This is the subject line of the email to the subscriber.
From Email Address	The From field in the email.
Reply To Email Address	Who the subscriber may reply to.
Body (Text/Plain)	The text of the email the subscriber receives in plain format.
Body (Text/HTML)	The text of the email the subscriber receives in HTML format.

Step 3 Go on to Create the Policy Action.

Messages for SMS Notification

Use this procedure to create the text messages for a subscriber's mobile device.

For advanced explanations about data coding and other details, see Addendum A: Data Coding.

Step 1 Select Reference Data tab, Notifications node, SMS Notifications.

Step 2 Fill out the fields in the detail pane.

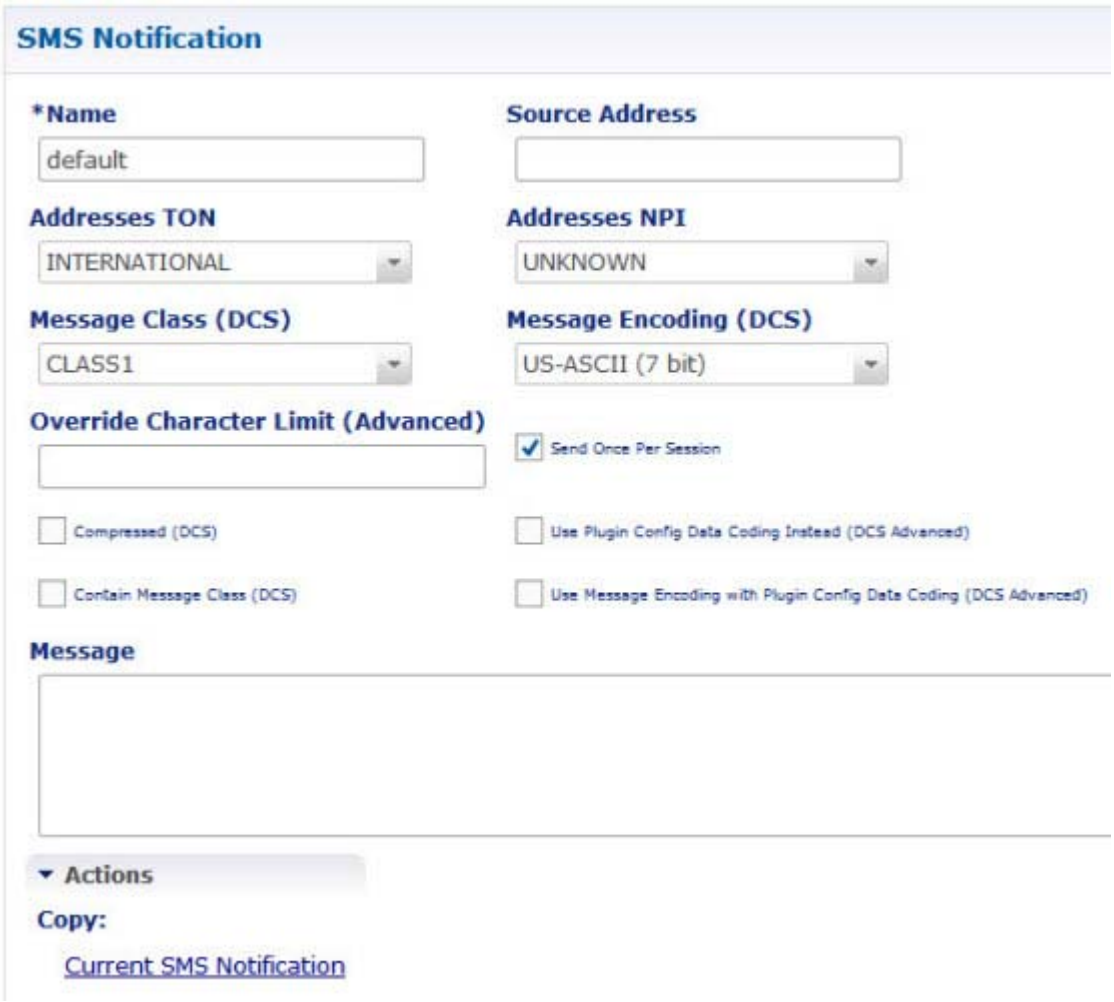

Parameter	Description
Name	Name of the notification message. This name is used later in the policy definition to send the SMS.
Source Address	Source address of the SMS message.

Parameter	Description
Address TON	Type of Number for the source, that is, 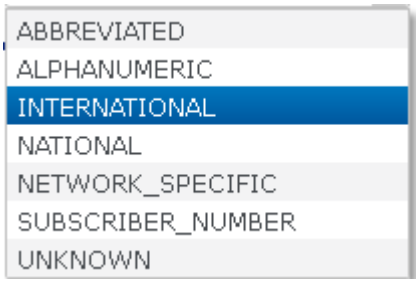
Address NPI	Numbering Plan Indicator, that is: 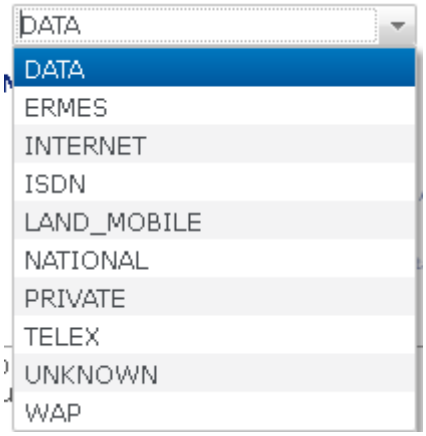
Message Class (DCS)	The message class per the SMPP specification. Valid values are CLASS0, CLASS1, CLASS2. Default value is CLASS1.
Message Encoding (DCS)	Defines the alphabet and byte encoding used for the message. Valid values are US-ASCII (7 bit), ISO-8859-1 (8 bit), and UTF-16 (UCS-2) which is 16 bit. Default value is US-ASCII.
Override Character Limit (Advanced)	Some SMSCs create multi-part messages for long SMS messages instead of having CPS create the multiple messages. This option provides such behavior by overriding the default single message size. This option is for advanced use only and should be used with care. The reason for such care is that if space in the message submitted from CPS does not allow for header information, such as the User Data Header (UDH), then many SMSC are not accepted the messages at all.
Compressed (DCS)	Select this check box to set whether compression is used per the SMPP specification. Default is false.
Use Plugin Config Data Coding Instead (DCS)	Select this check box when you want to use the value specified in Data Coding field in the Notifications Configuration screen instead of the Message Class, Message Encoding, Compressed, and Contain Message Class values on this screen.
Contain Message Class	Select this check box to set whether the contain message class options is used per the SMPP specification. Default is false.

Parameter	Description
Use Message Encoding with Plugin Config Data Coding	Select this check box when the “Use Plugin Config Data Coding Instead” check box above is checked. The check box "Use Plugin Config Data Coding Instead” must be true to use this value. This check box allows the Message Encoding value on this screen to define the byte conversion method that is used in conjunction with the Data Coding value in the Notifications Configuration screen. By default, the byte conversion method is US-ASCII regardless of the Plugin Configuration’s Data Coding value. Other UTF-16 conversions may use Big Endian, Little Endian or Byte Order Mark (BOM). This field is also important for ensuring the proper division of messages, particularly for non-English languages, for multi-part SMS message support.
Message	The text that the subscriber receives. SMS messages have character limits dependent on the selected DCS values. Text in excess of this limit triggers the submission of the multi-part messages to the SMSC.

Step 3 Go on to Create the Policy Action.

Message for Realtime Notification

Use this procedure to create the messages for a realtime notification.

Step 1 Select Reference Data > Notifications node > Real Time Notifications.

Step 2 Fill out the fields in the detail pane. An example is shown.

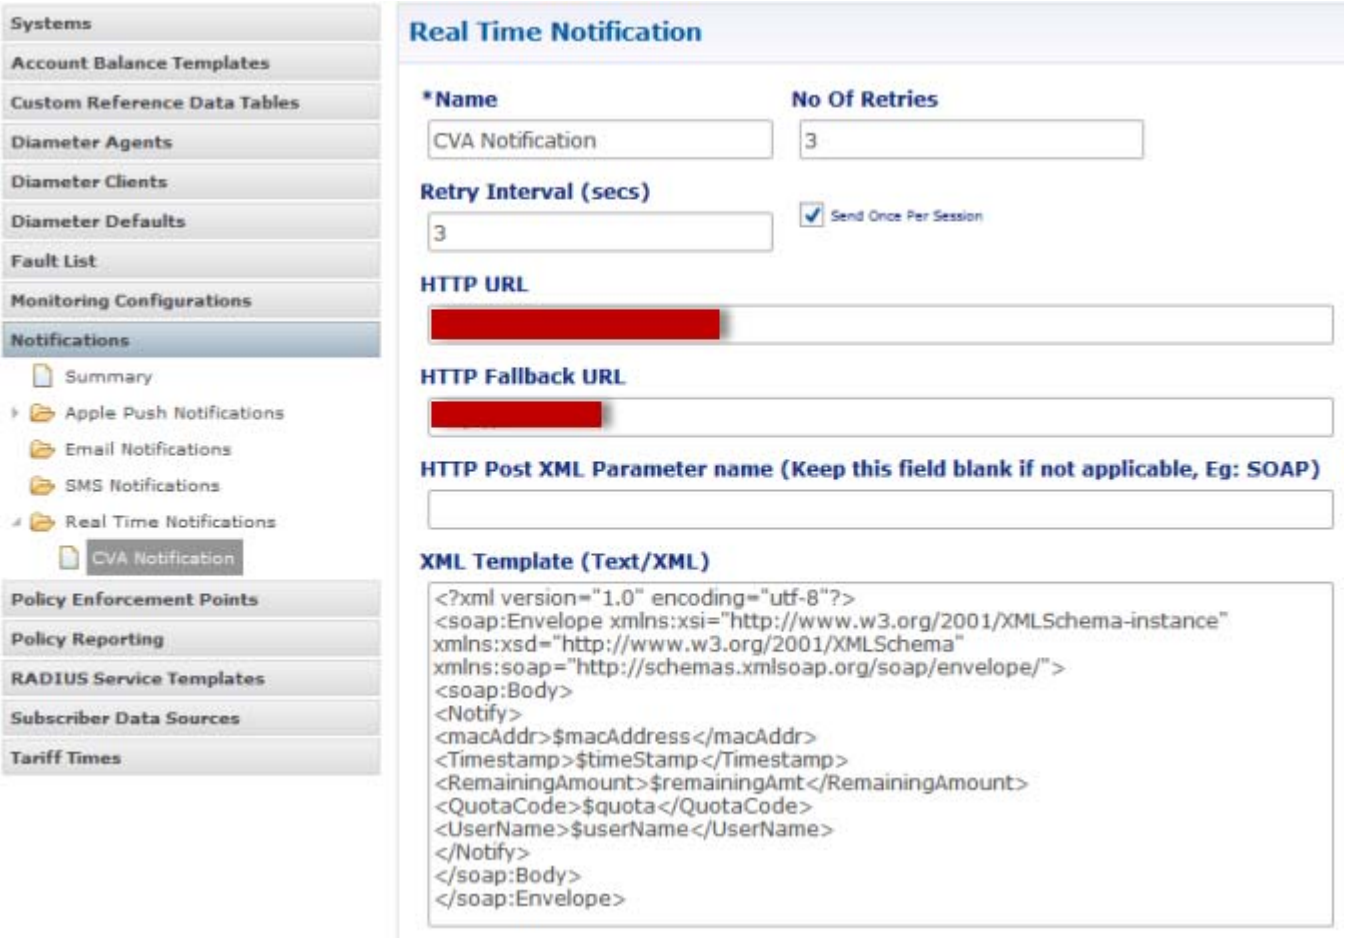

Parameter	Description
Name	Name of the realtime notification message.
No of Retries	When CPS sends realtime notification to the provided HTTP URL and if it is not reachable then this field specifies how many times CPS should send the notification. Same is true for HTTP Fallback URL.
Retry Interval (secs)	Interval during two retries.
Send Once Per Session	If checked, realtime notifications are generated for each session and not for all messages within that session.
HTTP URL	Primary URL where CPS sends realtime notifications.
HTTP Fallback URL	When Primary URL is not reachable then CPS tries to send notification to this URL as per configured No of Retries. When number of retries are exhausted then it tries to send notification to the HTTP Fallback URL.

Parameter	Description
HTTP Post XML Parameter name (Keep this field if not applicable, Eg: SOAP)	For SOAP this field is not applicable and hence should be blank. This field specifies HTTP Post parameter name.
XML Template (Text/XML)	This field has XML template, so as per configured template realtime notifications are generated. CPS provides values to the fields specified in the template from the ongoing session and for any field which is specified in the template but no value is found then that field goes as blank in the generated realtime notification.

Step 3 Go on to Create the Policy Action.

Create the Policy Action

The third part of setting up subscriber notifications is to make notification an action of a condition. This section describes specifically about how to set up an action that is a notification.

Recall that actions can be invoked from both a Policy and a Decision Table.

For general information on adding an action phrase to a condition of a policy see Actions Subtab on the Policy Screen.

For general information about policies, conditions, and actions, see Appendix G, "Policy Configurations".

Define the Policy Action

In this part of setting up notifications, you tie the action in a policy or decision table back to the message you created.

Step 1 Click Policies > an extension point > Policy or Decision Table.

Step 2 Select the Actions subtab.

Step 3 Click Add.

Step 4 In the top field, type something like NOTIF to find all the notification-type actions.

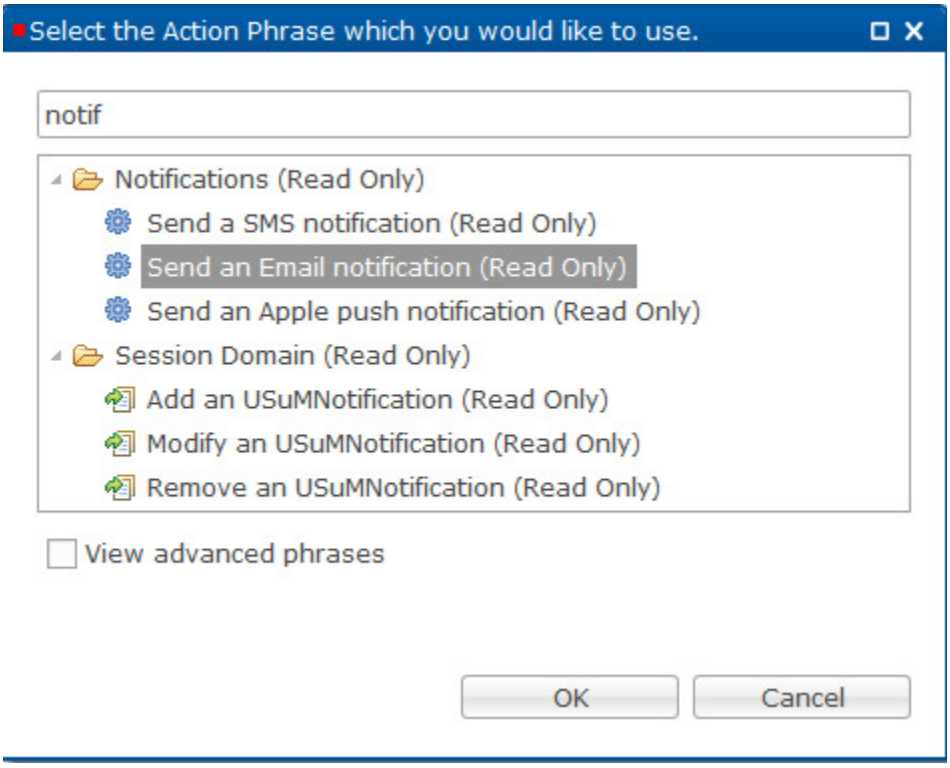

Step 5 Select an Action phrase for the kind of notification you are developing and click OK.

Step 6 On the Actions subtab screen, select the action you just added.

The example uses an email notification action.

The input variables for the action Send an Email notification appear at the bottom for you to enter the required information.

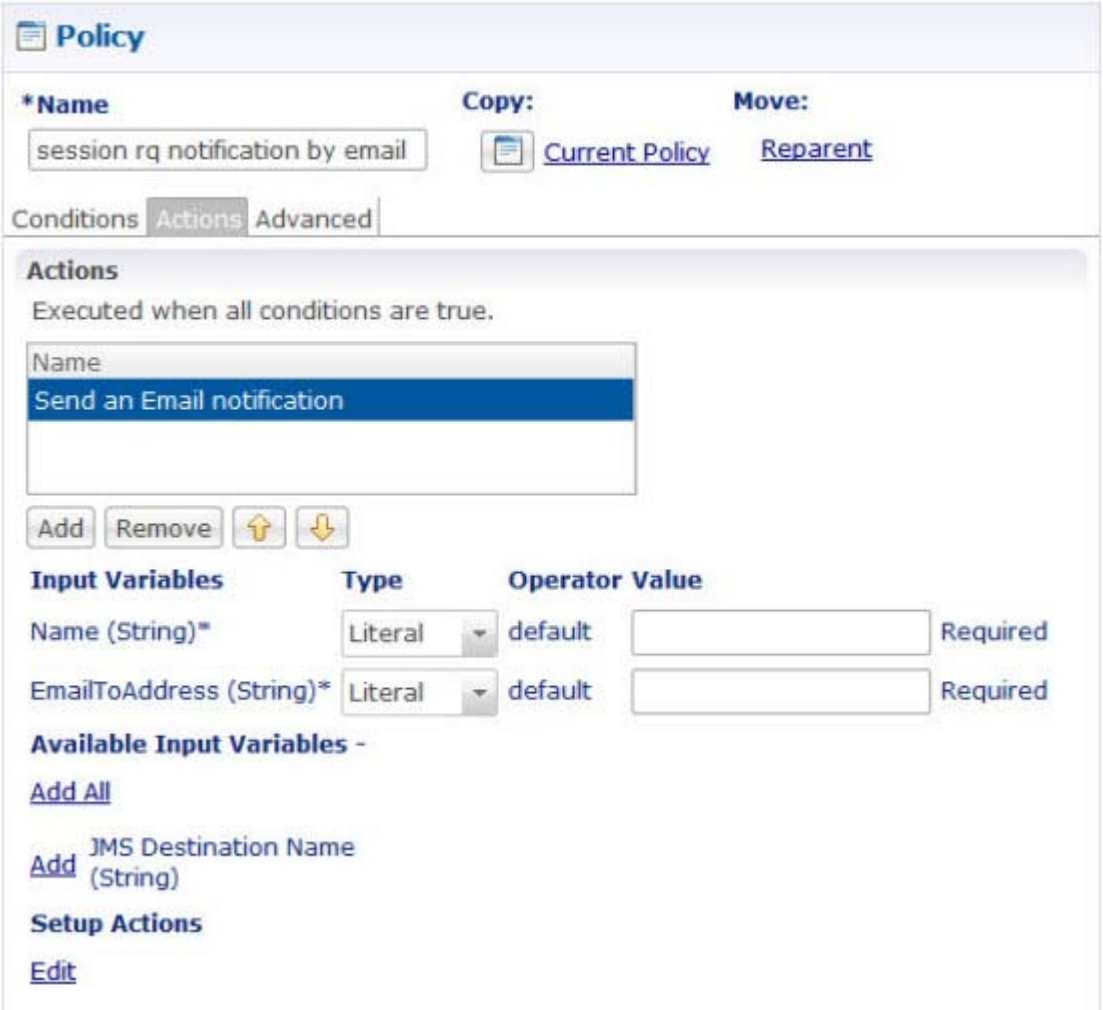

Parameter	Description
EmailToAddress	For an email notification, this is the email address to send to. This comes from a policy output variable or from session information. This can also be a specific, hard coded email address, but this usage is atypical.
Name	This is the name of the notification message as you defined it in the Name field of any of the notification types when you composed the notification message. The name you specify here must exactly match the spelling, spacing, and capitalization of what you used to create the notification message. This is what appears in the Reference Data tree. See the Name field in Messages for the Apple iOS and iPhone® Notification. See the Name field in Messages for Email Notification. See the Name field in Messages for SMS Notification.

Parameter	Description
Device Token	For notification to an Apple iPhone or other iOS device, this is the unique identifying number of the iPhone. Most likely, this comes from a policy output variable or from session information. Obtain output values such as session information by specifying conditions that contain the desired information, for example, “There exists a subscriberProfile”. The Device Token can be hard coded, but that method is not typical and only allows the message to be sent to a single device.
Destination Address	For an SMS notification, this is the NPI address of the subscriber. This comes from a policy output variable or from session information. Obtain output values such as session information by specifying conditions that contain the desired information, for example, “There exists a subscriberProfile” The Destination Address can be hard coded, but that method is not typical.

Step 7 Fill in the EmailToAddress and the Name fields.

If you are not performing any variable substitutions, you are done. Your email message is delivered.

- If you want text substitution in your message, work through steps 8- 11.
- If you want to use output values, skip steps 8 - 11 and go to steps 12 - 13.

Text Substitution

Step 8 Click Edit and display the Setup Actions screen.

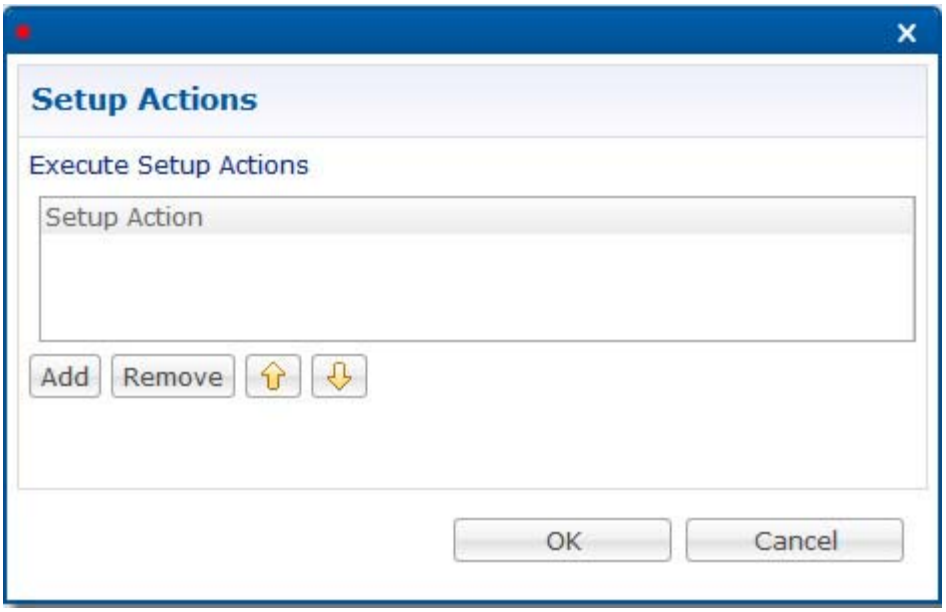

Step 9 Click Add and select an action from the Available Setup Actions list, Set Substitution Value.

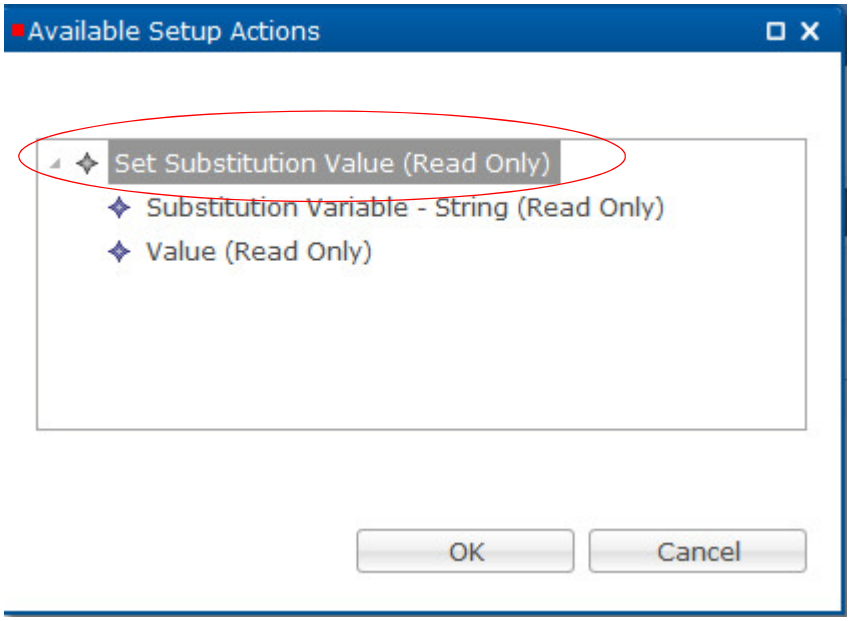

Step 10 Click OK.

The bottom of the Setup Actions screen is populated with the substitution variables.

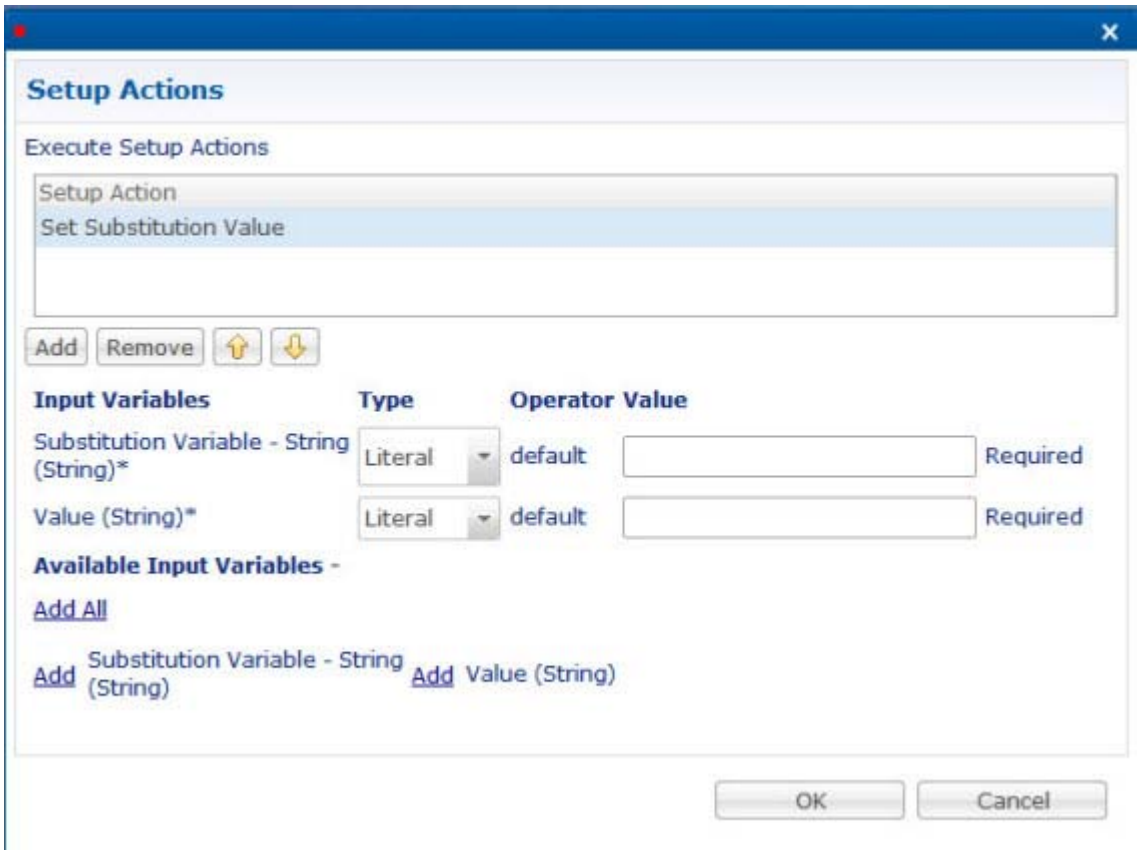

Step 11 In the Substitution Variable - String, use the type Literal and enter a Value of Name for example. Remember that in your message, discussed on page Messages for Email Notification, you use the syntax $Name in your message. The recipient's name are substituted in the message. The recipient's name is most likely obtained from the session information.

Output Values

Step 12 In the variable for Value (String), select the type of Output from the drop-down list, and enter a value of sessionSubscriberName.

Step 13 If you want to add other variable information to your message and make more substitutions, click Add All to insert another variable-value pair to the screen for definition.

Notification Performance Tuning Parameters

CPS supports configurable parameters for email socket timeout and socket connection timeout. Also for SMS, smsResponsetimeout indicating how long we wait for the response from SMPP server is supported.

The following parameters can be added in `/etc/broadhop/iomanager01/qns.conf` and `/etc/broadhop/iomanager02/qns.conf`:

- mail.socket.timeout

  Inactivity period for a socket after which it is closed.

  Default value: 30000 ms
- mail.socket.connection.timeout

  Timeout for connection to be established between CPS and SMTP server.

  Default value: 5000 ms
- sms.response.timeout

  How long CPS should wait for the response from SMPP server.

  Default value: 2000 ms

For higher TPS notification testing, tune the above mentioned parameters as well as Queue Size and processing threads in System > AsyncThreading Configuration PB. An example is shown below.

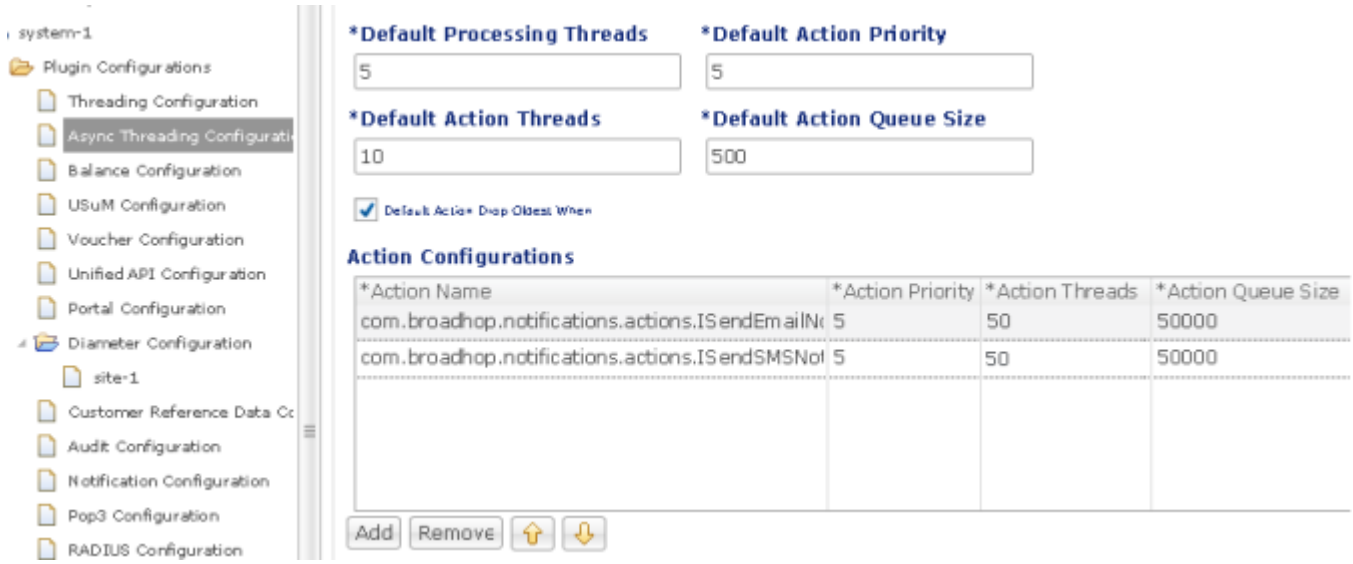

We can increase the Queue Size and processing threads for the following actions as shown in above example:

- com.broadhop.notifications.actions.ISendEmailNotificationRequest
- com.broadhop.notifications.actions.ISendSMSNotificationRequest

Addendum A: Data Coding

General

The Notifications feature is packaged with Core 5.3.3 and lower versions.

Starting in Core 5.3.4, Notifications is its own component and must be installed separately. The first version of this new Notifications component is 5.3.4; however, the versioning are not necessarily remain in sync with the Core version in the future.

Standards and Libraries

CPS supports sending notifications via the following mechanisms:

- Email (SMTP)
- SMS (SMPP v.3.4)
- Apple Push (for Apple iOS devices)

SMS

CPS binds via SMPP to the SMSC as a Transmitter (TX) (or Transceiver (TRX) - Core 5.3 and higher only). Note that when bound as a Transceiver, CPS still does not have the ability to receive SMS messages.

Notification Plugin Configuration

Located in Policy Builder under Reference Data > System > Plugin Configuration > Notifications Configuration in Policy Builder.

Notification Plugin Configuration Fields

Parameter	Description
<Required> “SMSC Host Address”	Specifies the user name (sometimes called User Name) for the gateway to use when connecting to the SMPP server.
<Required> “SMSC Port”	Specifies the TCP/IP port on the SMPP server to which the gateway should connect.
“System Id”	Specifies the user name (sometimes called User Name) for the gateway to use when connecting to the SMPP server.
“Password”	Specifies the password for the gateway to use when connecting to the SMPP server.
“System Type”	An optional login parameter that should be set only if required by the SMPP server. The SMPP system administrator provides this value, which when required, is usually a short text string.
"Registered Delivery" "Data Coding," and "Priority Flag"	An optional fields that are applicable for some customer deployments. A note about these three fields as well is that while they are defined in the Plugin Configuration, they can actually be set on a per message/notification basis, but CPS does not currently support this functionality as it was not needed by customers. Other fields, such as Message Encoding, exist on the SMS message definition screen which allow much of the same functionality as Data Coding.

Parameter	Description
"Data Coding"	Use instead of the Message Encoding and the other DCS fields on the SMS message definition screen and should be used with care, see the details in Data Coding section. If it is necessary to use a specific value in this field for Data Coding, then the message alphabet information/byte encoding type should be included in the Data Coding value per the SMS specification as well as any other necessary data coding information, or the "Use Message Encoding with Plugin Config Data Coding" check box should be used and the Message Encoding should be selected on the SMS message definition screen. What exactly happens is also a function of the capabilities of the SMSC and is not totally controlled by CPS. Please keep this fact in mind! Please refer to your specific Core version below in the Data Coding section for more information on the functionality and behavior.
"Binding Type"	Specifies if CPS should bind to the SMSC as Transmitter (TX) or Transceiver (TRX). TX is the legacy default. CPS can bind as TRX to support other SMSC configurations; however, please keep in mind CPS cannot receive or process SMS messages, therefore selecting TRX just determines the bind type, not the functionality offers. CPS can only send SMS messages to the SMSC regardless of the setting selected.

Fields Hidden/Removed

Address Range and Timeout are not currently implemented and/or have been defaulted and so have been removed from Policy Builder in the most recent versions to avoid confusion. They can be added back if necessary for a customer implementation. The fields have only been hidden in Policy Builder in the event a customer has populated the data in the past; however, no customer is known to have populated or used these fields as of their hiding.

The fields were also hidden since, in some cases, they weren't even applicable to what CPS does with SMS. For example, “Address Range” is a parameter used primarily when receiving messages. CPS does not receive SMS messages.

SMS Notification Definition Fields

Parameter	Description
Name	This is the name you use to refer to this message in your Policy or other setup. It should be unique across the SMS Notifications names.
Source Address	This field does not affect processing or encoding in CPS, but it is passed to the SMSC. This field is typically what the SMSC puts in the From field on the message.

Parameter	Description
Addresses TON and Addresses NPI	These fields are the Type Of Number and Numbering Plan Identification that define the format of the phone numbers/addresses. Binding to the SMSC is initially made with both of these set to Unknown. The actual Send SMS submissions to the SMSC use the values defined in these fields. The values in these fields are applicable for both the source and destination addresses.
Message	This field contains the actual message to be sent in the SMS.

For fields on this screen related to message encoding, see the later section SMS Notification Definition Fields Related to Message Encoding.

Data Coding (DCS)

The Message Encoding field determines what byte encoding is used to translate the message text to bytes to send to the SMSC.

Use of the Data Coding field in the Plugin Configuration can also be used in conjunction with the Message Encoding field to determine the encoding and byte encoding. Please see the Plugin Configuration information above if you are specifying a Data Coding value.

Current Versions (Core 5.1.10 and higher, 5.2.8 and higher, and 5.3.4 and higher)

Note Core versions 5.1.9, 5.2.7, and 5.3.0 - 5.3.3 only contain some of these new fields. If a particular field mentioned is needed for your deployment or you are experiencing a bug with your data coding, please upgrade to the latest version of Core.

Your choice for the DCS fields is driven by the following:

- Language your Message is in.
- Capabilities of your SMSC.
- 7 bit is what is used for typical U.S. text messaging.
- UTF-16 is required for some foreign languages, i.e. Arabic, Cyrillic.

Data Coding Choices:

- US-ASCII (7 bit)
  - 160 character limit for a single message.
  - 152 character limit for each part of a message than contains multiple parts (long SMS).
    - The reduced character limit is required to accommodate the User Data Header (UDH) needed for multi-part messaging.
  - Uses JSMPP alphabet of ALPHA_DEFAULT.
  - Uses Java text to byte conversion character set of US-ASCII.
- ISO-8859-1 (8 bit)
  - 140 character limit for a single message.
  - 133 character limit for each part of a message than contains multiple parts (long SMS).

    - The reduced character limit is required to accommodate the User Data Header (UDH) needed for t is required to accommodate the UDH.
  - Uses JSMPP alphabet of ALPHA_8_BIT.
  - Uses Java text to byte conversion character set of ISO-8859-1.
- UTF-16 (UCS-2)
  - 70 character limit for a single message.
  - 66 character limit for each part of a message than contains multiple parts (long SMS).
    - The reduced character limit is required to accommodate the UDH.
  - Required for non-English languages, such as Arabic.
  - Uses JSMPP alphabet of ALPHA_UCS2.
  - Uses Java text to byte conversion character set of UTF-16.

Note UCS-2 is the older specification, but is referred to frequently in SMS documentation, so it is included here for legacy understanding. Per unicode.org, "UCS-2 does not define a distinct data format, because UTF-16 and UCS-2 are identical for purposes of data exchange. Both are 16-bit, and have exactly the same code unit representation."

Default Data Coding is:

- Compressed: False
- Contain Message Class: False
- Message Class: Class 1
- Message Encoding/Alphabet: US-ASCII (7 bit)
- Default Data Coding is used in two cases:
  - If Use Plugin Config Data Coding Instead on the SMS Notification message definition screen is checked and the notifications plug-in configuration Data Coding value is null, a warning message to the effect that the Default Data Coding is being used also appears in the CPS log.
  - The default selections for Message Class, Message Encoding, Compressed, and Contain Message Class on the SMS Notification message definition screen also yield the Default Data Coding.

Locations to set the Data Coding (DCS) value to be used for SMS messages:

- Individual Compressed, Contain Message Class, Message Class, and Message Encoding (Alphabet) fields on the SMS Notification message definition screen.
  - Recommended for most implementations.
  - Each message can use different encoding if desired.
  - Data coding values are set for each message.
  - Use Plugin Config Data Coding Instead is NOT selected on the SMS Notification message definition screen (default, recommended).
- Data Coding value in an integer representation on the Notification Plugin Configuration screen.
  - Not recommended for most implementations - Advanced Use Only.
  - Retained for legacy support and handling of advanced situations which may not be handled by the drop down lists and check boxes alone.

- Also retained for explicit specification of a data coding value per the SMPP Specification 3.4, section 5.2.19, but as an integer value, which may be needed in special or edge cases.
- The "Use Plugin Config Data Coding Instead" check box needs to be selected on the SMS Notification message definition screen to use this option.
- Use of the "Use Message Encoding with Plugin Config Data Coding" check box.
  - Selecting allows the Java text to byte conversion character set to be set according to the value selected in the Message Encoding drop-down list. If this option is used, this value must correspond to the Data Coding value defined according to the SMPP Specification or undesirable behavior may result. In this case, the Message Encoding drop-down list ONLY drive the Java text to byte conversion and multi-part message character break limit. The actual Data Coding or DCS value still drives the plug-in config Data Coding value.
  - Deselected (default condition), when "Use Plugin Config Data Coding Instead" is also selected, uses the following text to byte conversion and multi-part message length:

    - Java text to byte conversion character set of US-ASCII

    - 160 char limit/msg before multi-part message occurs

Note When using the notifications plug-in configuration data coding value, message length is assumed to be 160 characters. Consequently, if alternate alphabets or multi-part messages are desired, it is recommended that the fields (compressed, contain message class, message class, message encoding) on each individual notification message definition is utilized and not the plug-in config data coding value. Additionally and recommended, the Use Message Encoding with Plugin Config Data Coding check box can be used.

SMS Notification Definition Fields Related to Message Encoding

Message Class – The message class per the SMPP specification. Valid values are CLASS0, CLASS1, CLASS2. Default value is CLASS1.

Message Encoding – Defines the alphabet and byte encoding used for the message. Valid values are US-ASCII (7 bit), ISO-8859-1 (8 bit), and UTF-16 (UCS-2) which is 16 bit. The default values is US-ASCII.

Use Plugin Config Data Coding Instead – Use the value specified in Data Coding in the Notifications Plugin Configuration instead of the Message Class, Message Encoding, Compressed, and Contain Message Class values on the message definition screen

Compressed – Whether compression is used per the SMPP specification. Default is false.

Contain Message Class – Whether the contain message class options is used per the SMPP specification. Default is false.

Use Message Encoding with Plugin Config Data Coding – “Use Plugin Data Coding Instead” must be true to use this value. Allows the Message Encoding value on the message definition screen to be used to define the byte conversion method that is used in conjunction with the Data Coding value. By default, the byte conversion method is US-ASCII regardless of the Plugin Configuration’s Data Coding value. This field is also important for ensuring the proper division of messages, particularly for non-English languages, for multi-part SMS message support.

Multi-part Long Messages

Please keep in mind that support for multi-part long SMS messages is not only based on the capabilities of CPS. The abilities and configuration of the SMSC are also a factor.

Current Versions (Core 5.1.10 and higher, 5.2.8 and higher, and 5.3.4 and higher)

Note Core versions 5.1.9, 5.2.7, and 5.3.0 - 5.3.3 do not contain the Override Character Limit field. Additionally those versions may have issues with long SMS support in some cases. If problems occur, please upgrade to the latest revision of your Core version (currently 5.1.10, 5.2.8, and 5.3.4). Please refer to the release notes of the specific versions for details regarding issues and bug fixes.

Multi-part or long SMS messages are supported in these versions. See the Message Encoding section for information about character limits and character sets and how to use them.

When in Debug mode, a log message of "sendSMSNotification: Message submitted to SMSC, message_id is <messageId>" appears in the CPS log for single part SMS messages.

When in Debug mode, the log message is "sendSMSNotification: Part: <msg#> of a <#OfParts> part multi-part message submitted to SMSC, message_id is <messageId>" for multi-part/long messages.

A new field was added on the SMS message definition screen that deals with long SMS support; however, it is for advanced use only and is not utilized in most deployments:

Override Character Limit (Advanced)

This is an override field for advanced use only and should be blank for most deployments. Zero and negative values are treated as blank as well. This field is the number of characters used to determine how many parts the message should be broken into. For example, if your defined message is 180 characters and you have specified UTF-16 Message Encoding, which has a character limit of 66 characters(140 bytes, up to 7 bytes is used for User Data Header (UDH)), then by default, when this field is blank, the message is submitted to the SMSC in 3 Submit SMS requests. If you put a value of 50 in the Override Character Limit, then the message is submitted to the SMSC in 4 Submit SMS requests. CAUTION: Be aware that if your value in Override Character Limit is greater than the character limit defined for a single message for the Message Encoding you selected, i.e. you put 100 in Override Character Limit and are using UTF-16 which has a single message limit of 66 characters (140 bytes including UDH) in the SMS specification, then it is the responsibility of the SMSC, or other downstream system, to handle any multi-part messaging. Many SMSC's does not accept messages that exceed 140 bytes for the message. Remember to plan for the 6 to 7 bytes of UDH.

Email

CPS Email Notification support allows an email to be sent to any email address based on a policy condition(s) being met via a IMAP email server.

iOS Push

CPS iOS Push Notification support allows a message to be pushed to an iOS device based on a policy condition(s) being met.

Notifications Manager Logging

Current Versions (Core 5.1.10 and higher, 5.2.8 and higher, and 5.3.4 and higher)

Note Core versions 5.1.9, 5.2.7, and 5.3.0 - 5.3.3 contain the majority of these logging enhancements; however, a few of the below examples may not be implemented in these versions.

As of implementation of CPS-901, the logging on the SMSC connection/binding was enhanced and differs from the past as follows:

- When initially binding, the binding is lost, or the reconnection of the binding fails, ERROR level log messages are returned.
- When initially binding or the reconnect binding is successful, INFO level log messages are returned.
- Plugin configuration missing is now ERROR level as well as throwing the exception as always.
- Failed connection to the Apple Push server is now ERROR level as well as throwing the exception as always.
- All exceptions (Email and Apple Push, not just SMS) due to Velocity and JSMPP library exceptions are now ERROR level as well as throwing the exception as always.
- Additional messages, such as the state change of the SMSC session and the Apple Push Device Token used, are returned as DEBUG level messages.

Log messages from NotificationsManager are prefaced as applicable by one of the following:

- NotificationsManager
- sendEmailNotification
- sendSMSNotification
- sendApplePushNotification

Log messages additionally include (VelocityEngine) after the above prefixes if the message relates to the Velocity Engine (string substitutions). So for example: sendSMSNotification(VelocityEngine).

Addendum B: Connections and Auto Reconnections

SMSC Connection and Binding

The Notifications Service connects to the SMSC per the connection parameters specified in the Notifications Plugin Configuration in Policy Builder under Systems. This connection is made when the QNS process is started or at any time the com.broadhop.notifications.service bundle is updated ((re)started).

In later versions of CPS, reconnection to the SMSC is performed automatically.

Auto Reconnection to SMSC

This logic was implemented in Core 5.1.9 and higher, 5.2.7 and higher, and 5.3.2 and higher. That said, a more refined version exists in 5.1.10, 5.2.8, and 5.1.10, so if you experience issues, please upgrade your CPS Core version.

A SMSC connection state listener was implemented. If this state listener detects a connection state change to closed or not bound, then CPS attempts to reconnect to the SMSC immediately and then every 5 minutes until the connection is again bound. Please see the logging section for information on how this loss of connection and reconnection appears in the logs.

Note The reconnection is attempted only if the original connection to the SMSC on CPS server start (or c.b.notifications.service bundle restart) was successful. This behavior is because, as mentioned, the reconnection logic is only triggered when a SMPP session state change results in a closed or binding lost state. Therefore, if the connection starts out as not bound or closed, then there is no state change to detect, and thus the reconnection logic is not triggered. This behavior is desirable so that when SMSC connection reference data is incorrect, CPS does not continually try to connect to something that isn't there. The point of the reconnection logic is to reconnect when the connection is lost when CPS already knows the connection is valid.

More About the Listener and Behavior

The listener listens for any changes to the state of the connection. If the connection state does not change, the listener does nothing. If the state becomes null or closed, then CPS attempts to reconnect immediately and then every 5 minutes until reconnected. Please note the CPS reconnect code checks the connection status one more time 5 minutes after it reconnects successfully before exiting the retry loop. CPS then stops checking the connection every 5 minutes once it determines the connection is not null or closed. CPS then goes back into listening mode until another connection state change occurs. And even if a state change occurs, CPS only reacts to a null or closed connection.

Again the CPS reconnect code does check the connection status again 5 minutes after it reconnects successfully before exiting the retry loop. However, as long as the connection has not again become null or closed during that 5 minute period, it exits the loop before trying to reconnect. This point could be confusing in the logs which is why this description provides further clarification.

Enquire Link Request/Response

Reply to enquire link messages is handled by a third-party library that is utilized by CPS for SMPP communication.

From the library's documentation, the library "… hides the complexity of the low level protocol communication such as automatically ..." managing "... enquire link request-response."

The default library value for the enquireLinkTimer is 5000 and the value is in milliseconds. The library does expose an advanced mechanism for overriding this value; however, CPS does not expose the value for override at this time.

The library handles both sending enquire_link requests and responding to enquire_link requests from the SMSC, although SMSCs do not typically send enquire_link requests to clients based on our research.

Addendum C: Logging

This section does not cover all logging, but it does provide some typical scenarios with respect to notifications.

There are many options for debugging and logging information for SMS Notifications at the Error, Warn, Info, Debug, and Trace levels, depending on the level of troubleshooting desired.

Turn logging up too high and performance is affected.

Searching the logs for NotificationsManager directs you to the majority of Notifications-related logging.

Notifications Manager Logging CPS 5.3.1

The majority of messages are returned at the DEBUG level.

The messages seen at CPS Start are:

- “Notification configuration missing” if the plug-in configuration isn't there.
- “Attempting to connect to SMSC, smscHostAddress: <smscHostAddress> smscPort: <smscPort>”.
- “NotificationsManager - Failed connect and bind to SMSC host:...” if the bind fails due to an exception.

Exceeding the length of an Apple Push Notification Payload (256 bytes) is a WARN level log message.

Exceptions are thrown as expected.

Notifications Manager Logging CPS 5.3.2

When initially binding, the binding is lost, or the reconnection of the binding fails, ERROR level log messages are returned.

When initially binding or the reconnect binding is successful, INFO level log messages are returned.

Plugin configuration missing is now ERROR level as well as throwing the exception as always.

Failed connection to the Apple Push server reports at the ERROR level as well as throws the exception.

All exceptions (Email and Apple Push, not just SMS) due to Velocity and library exceptions report at the ERROR level as well as throw the exception.

Additional messages, such as the state change of the SMSC session and the Apple Push Device Token used, report as DEBUG level messages.

Log messages from NotificationsManager are prefaced, as applicable, by one of the following:

- NotificationsManager
- sendEmailNotification
- sendSMSNotification
- sendApplePushNotification

Log messages additionally include (VelocityEngine) after the above prefixes if the message relates to the Velocity Engine (string substitutions). For example: sendSMSNotification(VelocityEngine).

Policy Enforcement Point Configurations

Revised: July 10, 2015

A Policy Enforcement Point, or PEP, is a component of policy-based management that might be a network access system (NAS). PEPs are not limited to NAS devices however.

Consider, when a user tries to access a file on a network or server that uses policy-based access management, the PEP describes the user's attributes to other entities on the system. The PEP gives the Policy Decision Point (PDP) the job of deciding whether or not to authorize the user based on the description of the user's attributes.Applicable policies are stored on the system and are analyzed by the PDP. The PDP makes it's decision and returns the decision. Then, the PEP lets the user know whether or not they have been authorized to access the requested resource.

This appendix covers the following sections:

Policy Enforcement Point Tree

Upon installation of Cisco Policy Suite, the Reference Data policy enforcement points (PEP) selection resembles this:

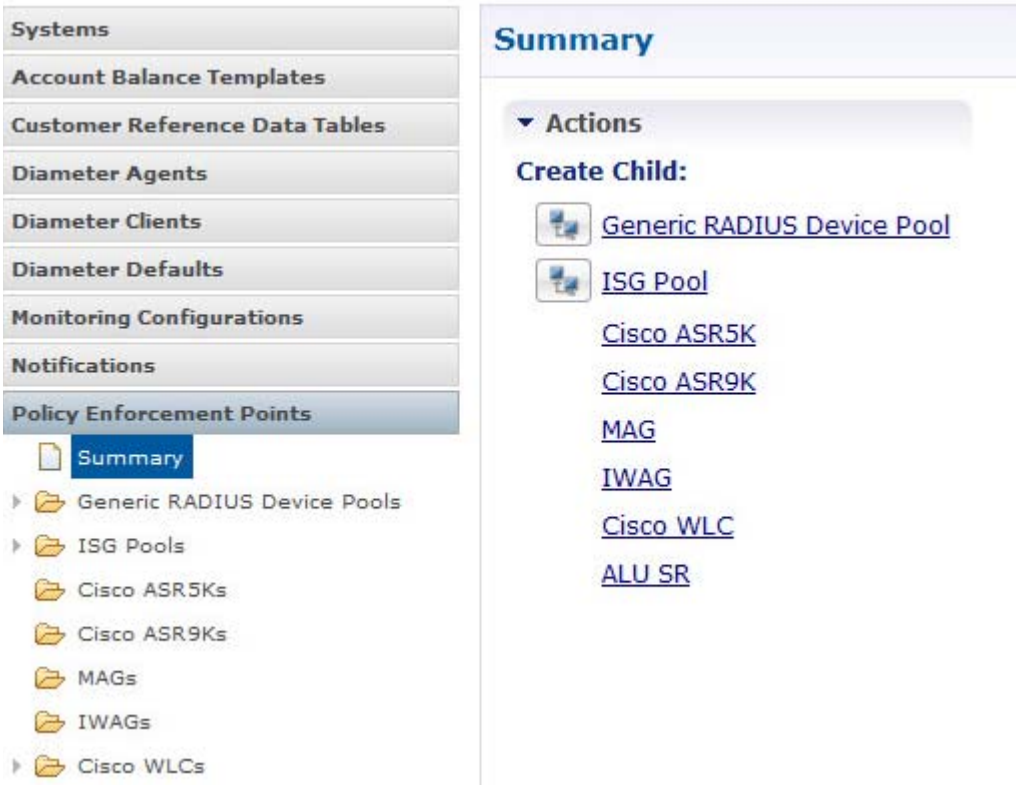

At install time, you need to determine what policy enforcement points your installation use and what features you need to install.

PEPS might be:

- Cisco ISG pool
- Cisco ASR 5K
- Cisco ASR9K
- MAG
- IWAG
- Cisco WLC
- SCE Device Pool
- RADIUS AAA server or device pool
- Procera
- Allot
- PDSN
- PCEF

Consult your Cisco technical representative for configuring a custom site.

Adding a Policy Enforcement Point

This example shows you how to add a Cisco Systems ISG network device as a policy enforcement point. Your PEP may be different, but you can easily follow this example.

Step 1 Click Reference Data tab > Policy Enforcement Points node.

Step 2 Choose the link from the main window that matches your type of PEP. For this example, select ISG Pool.

You might open up the ISG Pool folder to see if it has any PEPs already.

Defining a Policy Enforcement Point

Step 1 From the main screen click Reference Data > Policy Enforcement Point.

Step 2 In the right pane, under Create Child, select Generic RADIUS Device Pools.

Step 3 The first time you set up a device, the device has a name of 'Plugin default'. You should rename it.

Step 4 Fill in the RADIUS Device Pool screen.

The fields in the top area of the screen apply to all the devices. To use other addresses or secrets, specify such an exception in the bottom of the screen, in the Devices area.

Step 5 If you have a RADIUS device that uses different values from the ones displayed in the top area, create another device pool to accommodate that information.

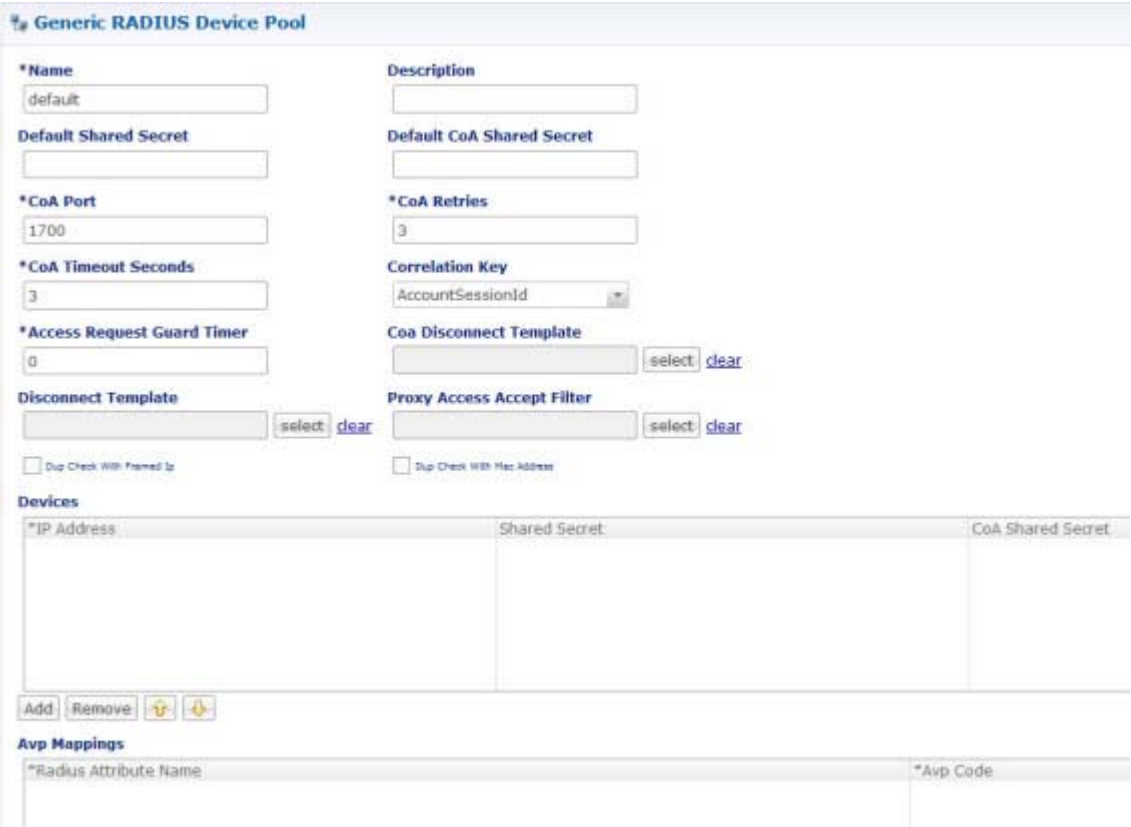

General Information	The fields in this area of the screen apply to all of the RADIUS devices defined except for those in the Device table at the bottom. If you have a RADIUS device that uses different values from the ones displayed in this area, create another RADIUS device pool to accommodate that information.
Name	Name of the RADIUS device pool. This name does not have to be unique, but best practice is to make it unique.
Description	Helpful information about the device pool.
Default Shared Secret	The shared password or phraseword between Cisco Policy Builder and the ISG device.
Default CoA Shared Secret	This shared secret is used between Cisco Policy Builder and the RADIUS devices unless a different one is specified in the Devices table below.
CoA Port	The hardware port on the RADIUS that listens for authentication tries. The default CoA port is 1813.
CoA Retries	The number of times that Cisco Policy Builder tries to authenticate with the RADIUS device in the list below.
CoA Timeout Seconds	The number of seconds that CPS tries to authenticate with an ISG device.

Correlation Key	This is the key that correlates between the subscriber authentication request and the rest of the requests. Your choices are these: • AccountSessionId • callingStationId • Tgpp2CorrelationId • UserId
Access Request Guard Timer	Enables the number of seconds between an Access-Accept being sent and the accounting start being received. If the Accounting start is not received before the timer expires, then the session is dropped.
CoA Disconnect Template	What you select here determines the RADIUS template used when a CoA message is sent to terminate a subscriber session on the RADIUS device.
Disconnect Template	Your selection here determines the disconnect template that is used when using the Packet of Disconnect message to terminate a subscriber session on the RADIUS device. Your RADIUS device should support either CoA or PoD.
Dup Check With Framed Ip	Select this check box to look for a CPS session with the same IP address on the Access Request or Accounting Start. If there is a session up with the same framed IP, that session is removed so that the new session can be created.
Dup Check With Mac Address	Select this check box to look for a CPS session with the same MAC address on the Access Request or Accounting Start. If there is a session up with the same MAC, that session is removed so that the new session can be created.
Devices	This list identifies the individual RADIUS devices in this RADIUS pool.
IP Address	The IP address of a RADIUS device you are using.
Shared Secret	The shared password or phraseword between Cisco Policy Builder and the RADIUS device. If no secret is specified here, the value in the Default Shared Secret field is used.
CoA Shared Secret	The shared password of phraseword between Cisco Policy Builder and the RADIUS device for purposes of authentication. If no secret is specified here, the value in the Default CoA Shared Secret field is used.
Loopback Addresses	Loopback addresses are set here. You cannot use the management address of the ISG. If loop back address are not set properly here, the system does not function.
AVP Mapping Table	This table area is used for generic mappings between subscriber session AVPs and an AccessAccept for the subscriber's authentication. Information you can map is the RADIUS attribute, AVP code, and the replacement value that you wish.

Edit a Policy Enforcement Point

To edit a PEP you already have in place, follow these steps:

Step 1 From the main screen, click Reference Data > Policy Enforcement Points.

Step 2 Select the device pool that holds your device.

Step 3 Make your changes to the Device Pool screen.

Step 4 Save your work to the local directory by clicking on the diskette icon or CTRL S.

Step 5 If you are ready to commit these changes to the version control software select File > Save to Repository.

Removing a Policy Enforcement Point

At times in building out your Cisco Policy Suite deployment, or perhaps due to network reconstruction, you may want to remove a device or a device pool.

To remove the entire node, highlight the node in the tree, and then click the red X at the top.

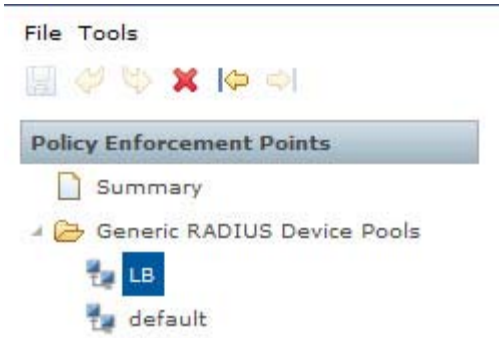

Use these steps to delete an individual instance from a pool.

Step 1 From the main screen, click Reference Data > Policy Enforcement Points.

Step 2 Scroll through the tree on the left until you find the pool or device you want to delete.

Step 3 To delete a device that is part of a pool, find the device pool and find the device in the device table at the bottom.

Step 4 Select the device as shown in the figure.

Step 5 Click the Remove button.

IWAG Configuration

Configure Policy Enforcement Point

Step 1 From the main screen, click Reference Data > Policy Enforcement Point > Summary.

Step 2 In the right pane, under Create Child, select IWAG or user can click IWAGs from the left pane.

Step 3 The first time you set up a device, the device has a name of 'Plugin default'. You should rename it.

Step 4 Provide the configuration details in the IWAG screen. The example configuration is shown below which is similar to other end point configuration.

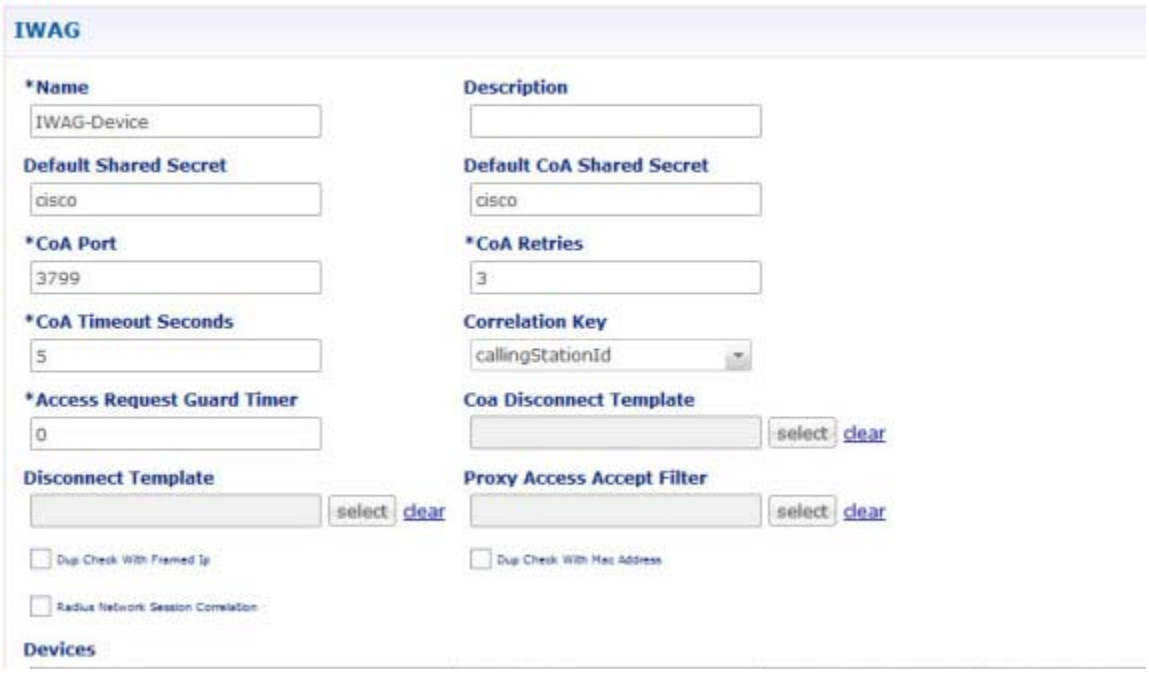

Configure Access Accept Template

Step 1 From the main screen, click Reference Data > RADIUS Service Templates.

Step 2 In the right pane, create a child in IWAG Access Accept template and configure as below. This configuration is also same as other Access Accept template.

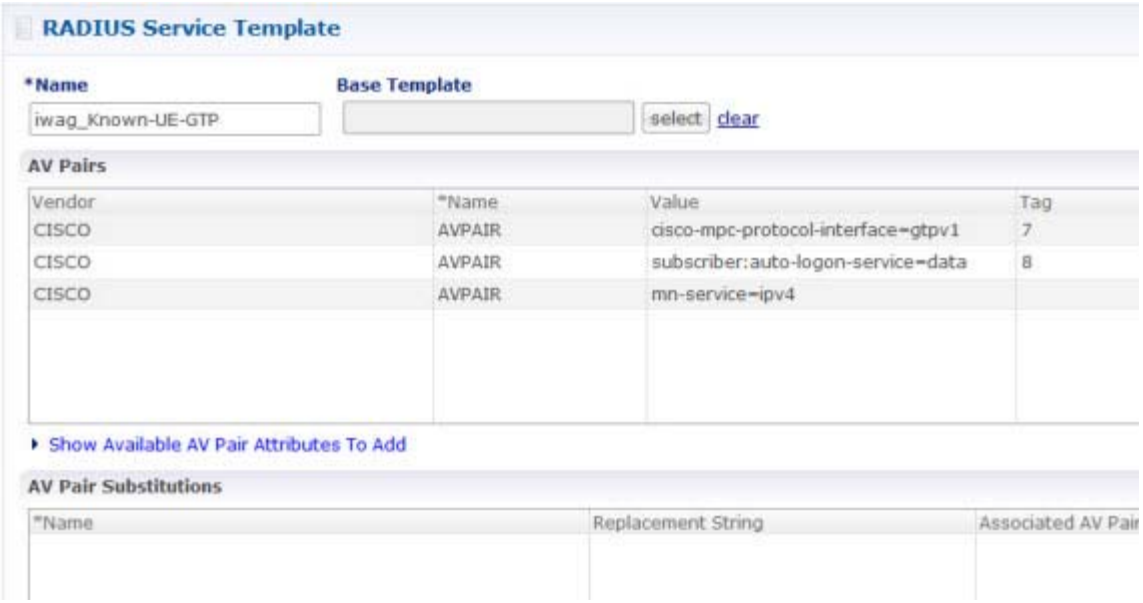

Configure Use Case Template

Step 1 From the main screen, click Services > Use Case Templates.

Step 2 Create a Use Case Template for IWAG Access Accept Configuration. The sample configuration is shown below.

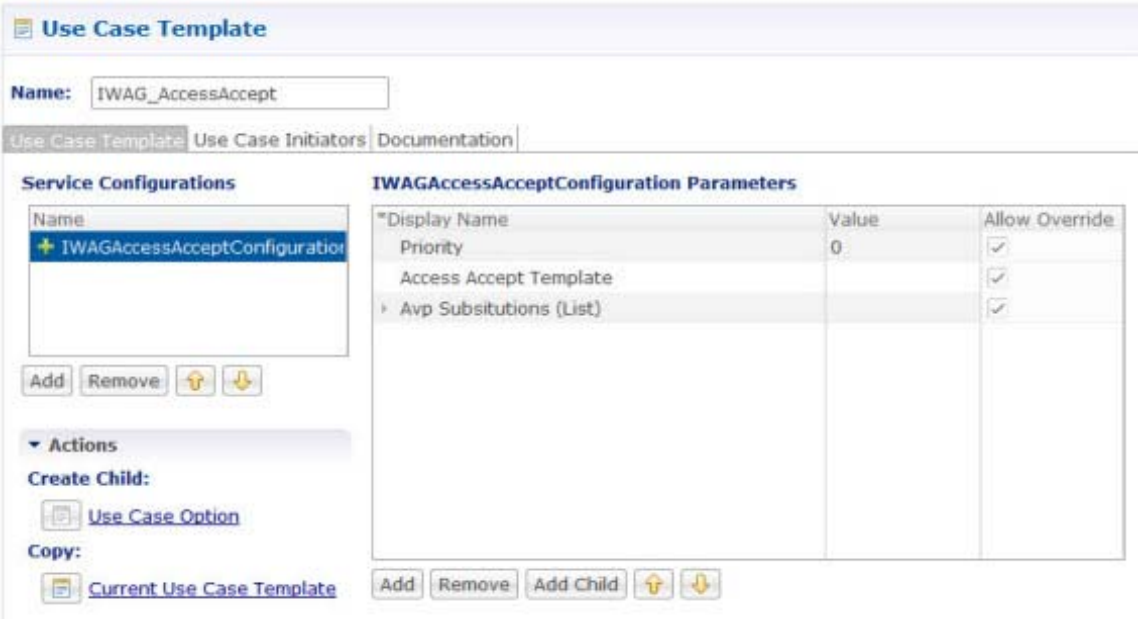

Configure Service Option

Step 1 From the main screen, click Services > Services.

Step 2 Create a Service Option using the Use Case Template created in Configure Use Case Template. The sample configuration is shown below.

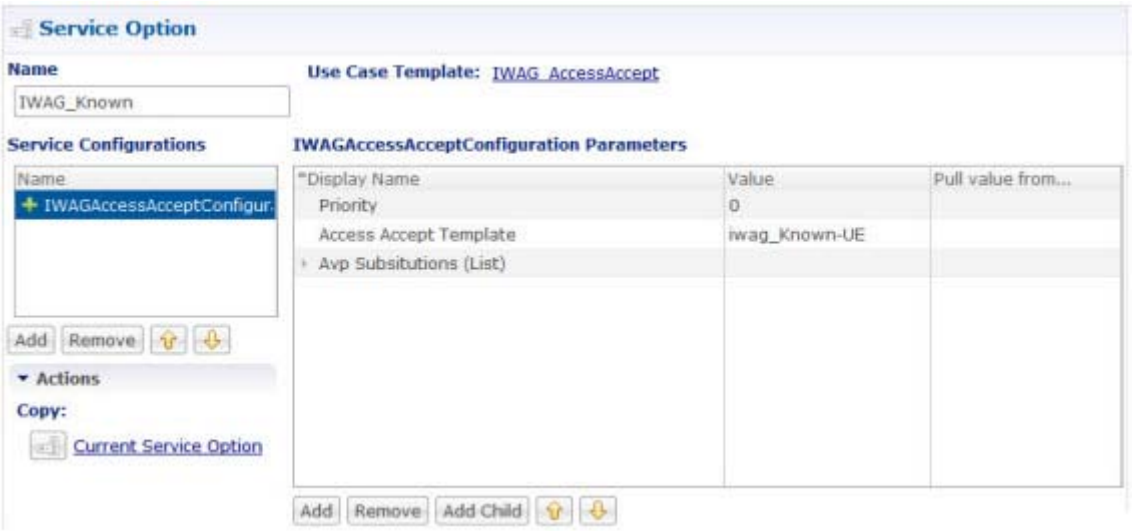

Step 3 Create a Service which uses the service option created in Step 2. The sample configuration is shown below.

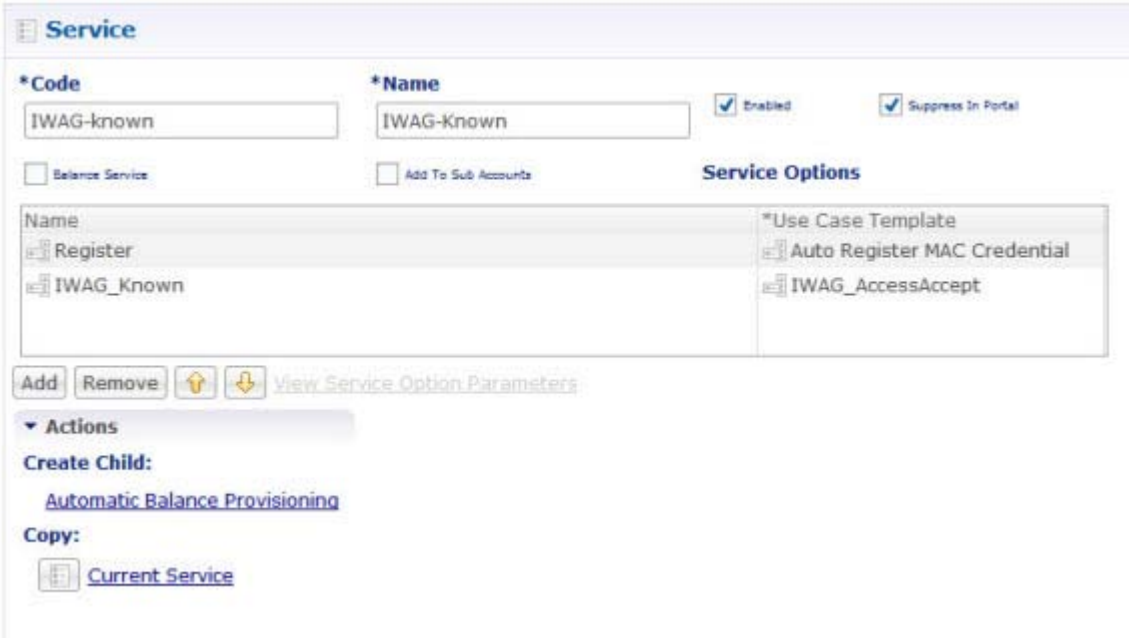

Step 4 Publish the configuration and associate this service with the subscriber in Control Center.

Subscriber Configuration

Revised: July 10, 2015

Cisco Policy Suite adapts to a variety of sources for subscriber data.

Possible subscriber profile repositories (SPR) that may be available to you are:

- Cisco's Cisco Control Center interface component of CPS
- Cisco's Unified Subscriber Manager (Cisco Unified SuM) component of CPS
- Cisco's AAA server component of CPS
- LDAP
- AAA

Note This flexibility lets you include either an external subscriber management system in your Cisco Policy Builder architecture or the internal, integrated Cisco Unified SuM.

Subscriber management schemes vary and are particular to an individual network.

Because of this, the procedures for obtaining subscriber data are discussed in the specific documents that matches your network architecture. See your specific document.

Subscriber Notifications

For more information about configuring notifications to subscribers, see Notification Configurations.

The Cisco Policy Suite subscriber notifications has two methods:

- Online notification that occurs when the subscriber is logged in
- Offline notifications and SPR Cleanup

Subscriber notification lets you define event and action pairs in a similar fashion to the existing policies. These pairs are then evaluated on a defined schedule. Events are similar to a policy's conditions, and actions are like a policy's actions.

As an example, some events and their resulting actions are listed below.

Event:

- Subscriber's quota on balance code Data expires in 1 hour
- Subscriber has Service A assigned

Action:

- Send an SMS Notification to the subscriber
- Send an email to the subscriber

Event:

- Subscriber has no services assigned in SPR

Action:

- Delete the subscriber from SPR

Subscriber notifications via email can use a variety of alphabet systems, including Cyrillic and Arabic.

Policy Configurations

Revised: July 10, 2015

Policies are typically reflected in services and service options, under the Services tab. This chapter is included for completeness.

A special policy can still be used to accommodate a specific, discrete use case, but this is rarely necessary.

This appendix covers the following sections:

Blueprints

Many policies and reference data are packaged in the Cisco Policy Builder as a kind of blueprint for you to follow. You can load blueprints easily from the Cisco Policy Builder screens. After you have a policy blueprint loaded, it is a simple matter to adjust or add to the policies they contain. You may not need to make any changes at all.

Blueprints are created by Cisco for your use. Blueprints are packages of policies and reference data relating to a given scenario or protocol. Blueprints can build on each other to support highly custom scenarios. The most basic blueprint is called the Initial Blueprint, which makes basic assumptions and is utilized in all but the most custom implementations.

The Policy Tree

The original nodes in the policy tree are part of the Initial Blueprint. You can add nodes to the policy tree to meet your network needs, causing it to look very different from any example we could provide here.

The Initial Blueprint is a root-level blueprint. This blueprint does not depend on data from another blueprint to exist. The Initial Blueprint uses these screens:

- Configured Blueprint Screen that relates to common, Cisco-provided blueprints.
- Configured Extension Point Screen that relates to a given piece of functionality.
- Configured Trigger Extension Point Screen that relates to a given piece of functionality.
- Policy Screen where you set up conditions and actions that make a policy.
- Decision Table Screen where you set up a table of inputs that feeds a condition of a policy.
- Policy Group Screen where you group several policies to improve processing.

Note The Root Configured Blueprint in the Policies Summary screen becomes the Initial Blueprint in the Policies tree when you select and use it.

In the Policies tree, the policies and extension points listed are in the order in which they are processed. That is:

- Initial Blueprint
- Network Session Folder—holds policies and objects related to session data. This folder is typically empty at initialization. This folder holds the Calculated Policies Folder, which holds calculated policies that track inter-session data. This folder is typically empty at initialization.
- Autowire
- ISG Portal Login

- USuM
- POP3 Authentication
- Policy Reports
- Configured Extension Points

The configured extension points that comprise the Initial Blueprint provide typical points where you may want to extend with your own policies or add new policies. Extension points are typically empty at initialization.

The figure below shows these Configured Extension Points:

- Set Inactive Session Retention Rules
- Store Inactive Attributes
- Setup Network Access Policies

The Root Configured Blueprint and the Initial Blueprint

The Initial Blueprint is the starting blueprint for a policy framework. In general, this is *always* the starting blueprint used when configuring the policies.

High Level Design

The Initial Blueprint executes the following policy flow.

- Pre-Session Policies. These are policies not associated with a subscriber session. They are defined in the "Pre-session policies".
- Load Session. Upon receiving a policy message, the load session policies attempt to load the session using keys that are retrieved from the input message.
- Stop Session. Upon loading a session, the session can be stopped if "Stop session" criteria is fulfilled (for example, a RADIUS stop message can be a stop session criteria).

- Start Session. If a session does not exist, then a new session can be started if the "Start session" criteria is fulfilled (for example, a RADIUS start message can be a start session criteria).
- Active Session Policies. If a session is active then the active session policies are initiated. The active session policies are executed in the following order:
  - Map session data from input. This maps data from the input record to the network session (for example, mapping the user ID from a RADIUS record).
  - Load inactive session. This retrieves any data (AVPs for example) that were stored when the previous subscriber's session expired or was stopped. For example, a Turbo AVP could have been added by a web portal and this can be stored when the subscriber session disconnects.
  - Remove expired AVPs. Any expired AVPs are removed at this point.
  - Load subscriber data. This retrieves any AVPs from the subscriber's data source. For example, retrieving the AVPs from an AAA server or the Cisco SuM connector.
  - At this point there are no AVPs. The subscriber's session is removed.

Note If the session is going to continue, then the post subscriber load policies can be executed and pre-configured network access policies are executed. These are used to perform changes to the subscriber's profile before the network access policies are set.

  - Setup network access policies. These are used to setup the network access for a given subscriber. **Note**: this does not send any outbound messages.
- Send outbound messages. The Send Outbound Messages extension point is where the reply and /or any additional messages are sent. For example, RADIUS accounting responses can be sent at this point.
- After all policies are executed, the session keys and the next evaluation time for the session is calculated.
- Use the Policies tree to add the Initial Blueprint. Perform this procedure the first time you configure your CPS application.
- Click the Policies tab and click the Root Configured Blueprint in the right pane to display the Initial Blueprint.

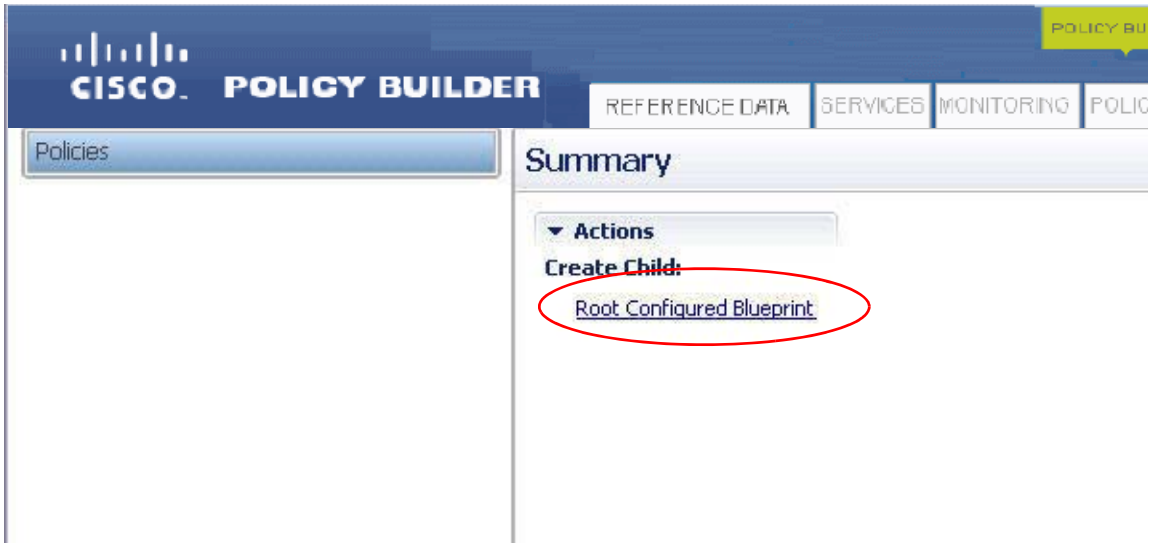

Choose the Initial Blueprint to start.

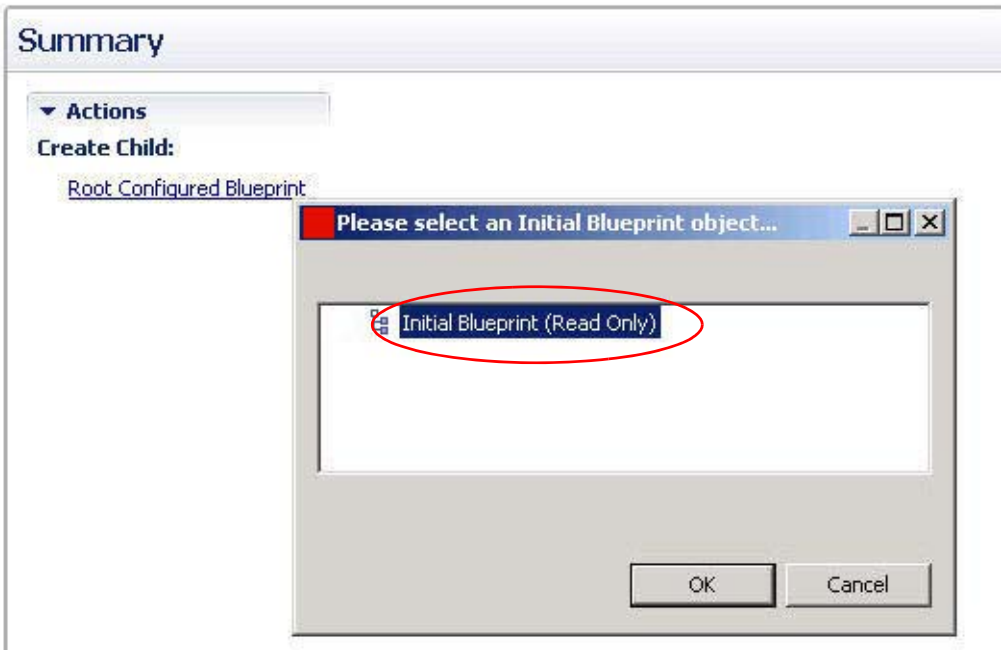

Caution if you choose a new root blueprint, all data on your Policies tab is lost as it exists under the root blueprint which are replaced. Unless you want to start over from scratch, you should never do this. The only way to recover from this loss of data is to use a backup or subversion commands to restore a previous Policy Builder configuration.

The Initial Blueprint policy has loaded into the Policies tree.The expanded view is shown.

Take a moment now to see all the policies provided for you by performing this single step.

Step 1 Click Initial Blueprint > View All Policies.

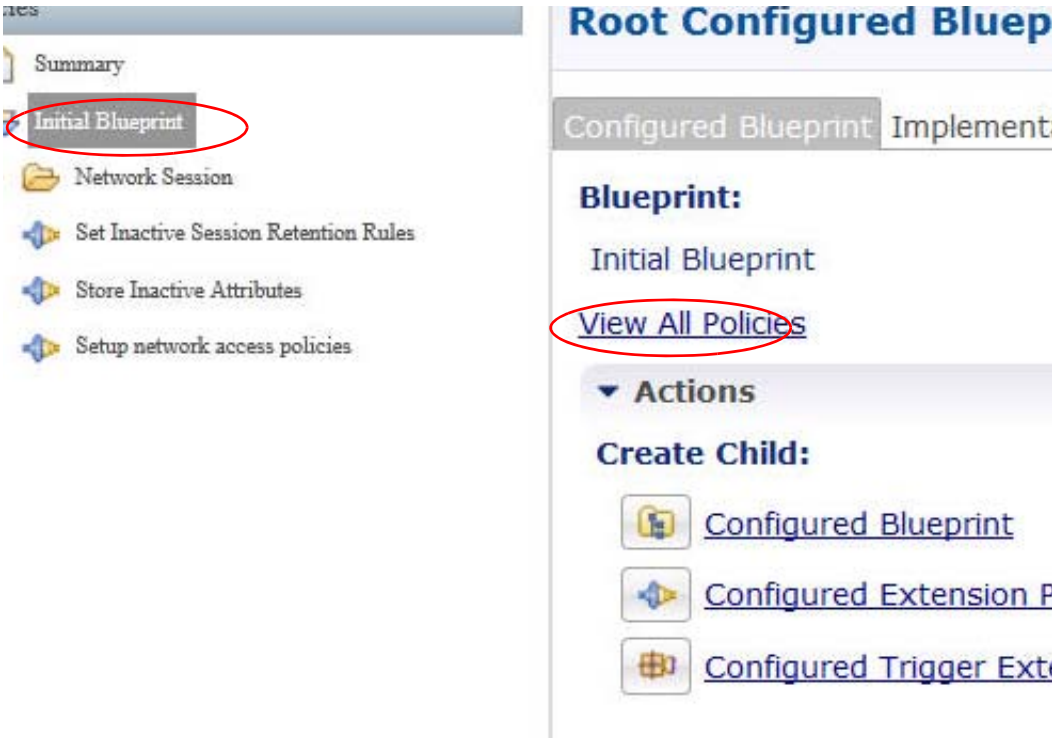

Configured Blueprintreparent	Configured blueprints provide policies that are already collected and configured. This way, you do not have to build a policy set yourself. Rather, we provide them for you.
Configured Extension Points	Configured extensions provide a place for you to add policies to your system, augmenting the standard policies the Cisco Policy Builder provides in a blueprint.
Configured Trigger Extension Points	Trigger points let you specifically start and stop a session, based on conditions. The Configured Trigger Extension Points in the example above are these: • Pre-session policies • Set Inactive Session Retention Rules • Store Inactive Attributes • Setup Network Access Policies • These are points where you may want to develop your own specific polices, conditions, and actions.

Scroll through this list to see what is already provided in your CPS blueprint.

Policies have configuration data available that affect the session domain, extensions, and triggers. The Initial Blueprint starts out with the nodes described in this list. You can populate them with policies that meet the needs of your business rules.

This is the complete list of policies provided by the blueprints you have configured under the Policies tab. This figure shows the processing order for all policies.

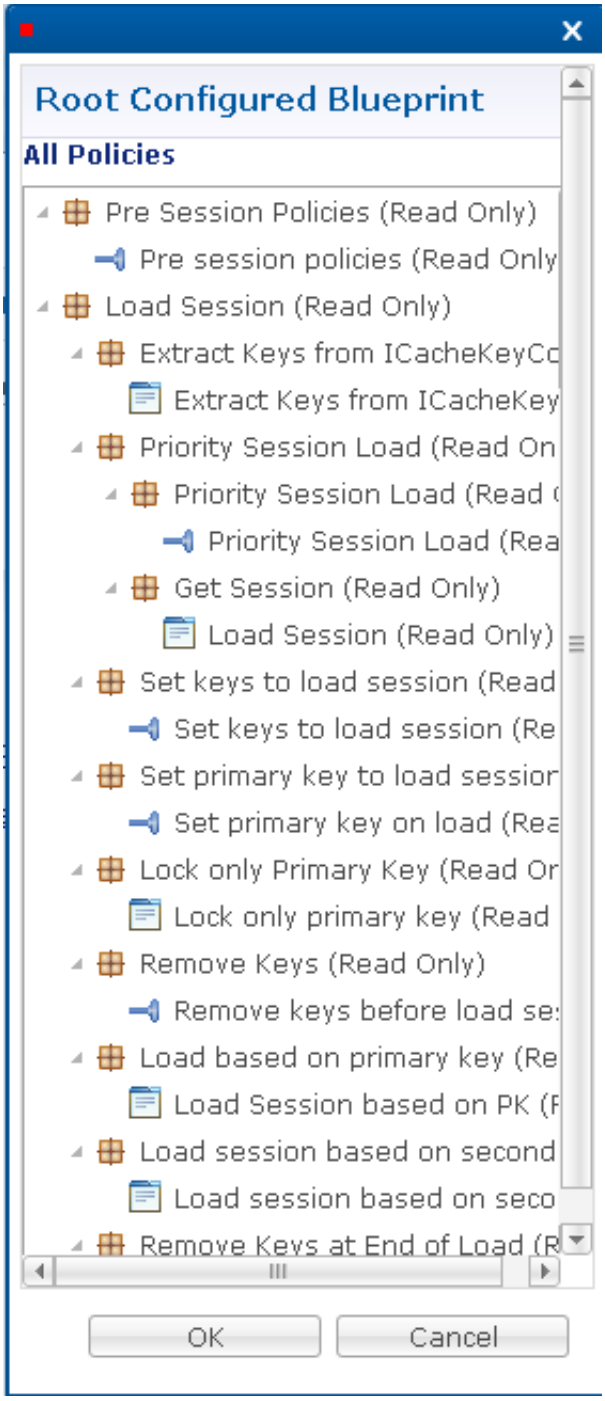

Customizing the Initial Blueprint

The Initial Blueprint comes with several configured extension points available. If you want to add additional configured extension points, use these steps.

Step 1 Display the Initial Blueprint in the Policies tree and open it to display its nodes.

Configured Blueprint Subtab

The Configured Blueprint subtab, available when you select the Initial Blueprint node, lets you create children to the Initial Blueprint with several types of extension types. The children you can create are of three types:

- Configured Blueprint type
- Configured Extension Point type

- Configured Trigger Extension Point type

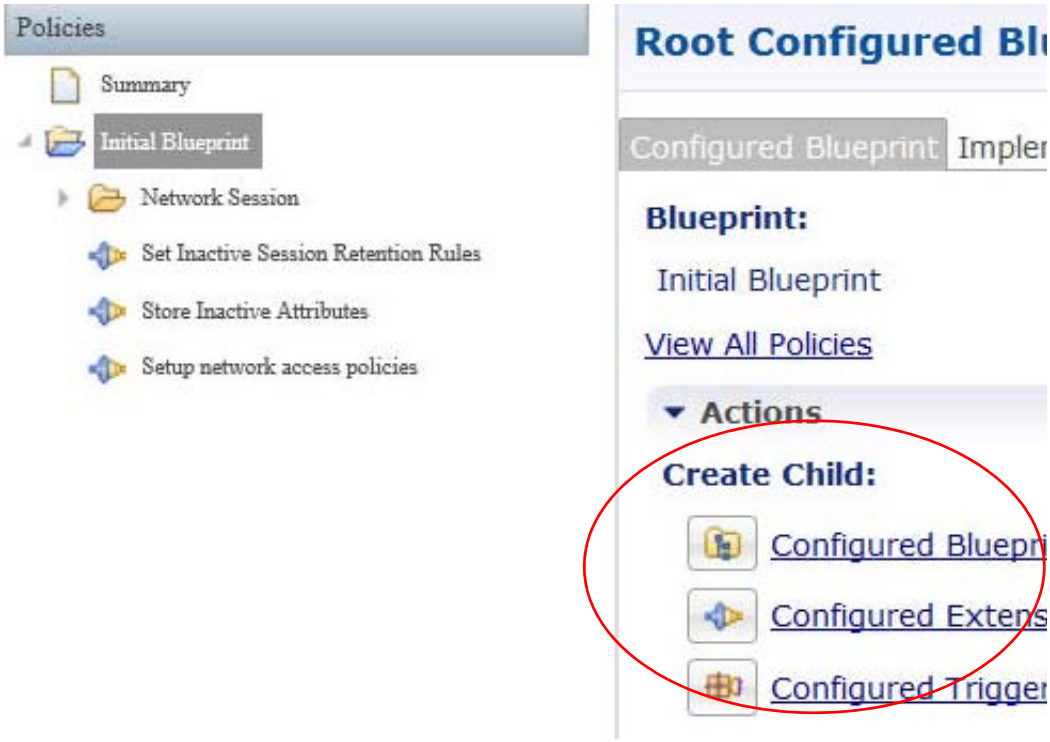

Step 1 Click the View All Policies link to see all the policies that comprise the Initial Blueprint. Then click either OK or Cancel to close the window.

Step 2 From the Initial Blueprint node, make customized additions to your Initial Blueprint with these types of extensions:

- Configured Blueprint Screen
- Configured Extension Point Screen
- Configured Trigger Extension Point Screen

Implementation Notes Subtab

With implementation notes, you can document what you are doing and planning when developing your policies structure.

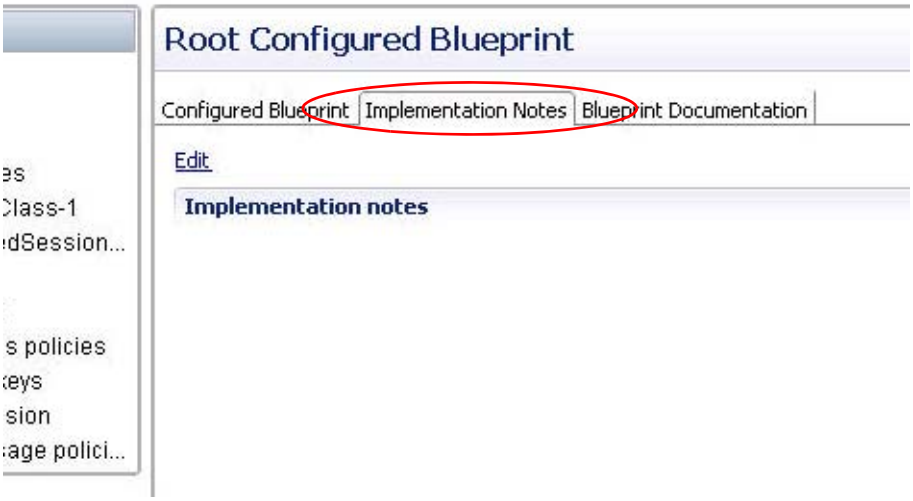

Step 1 Click Edit to open a note pad session and enter your comments into the Wiki Documentation Editor, the pane on the left. Your notes appear on the right.

Step 2 Use the square ☐ in the corner of either pane to make that pane full-screen.

Click the square again to return to the original screen.

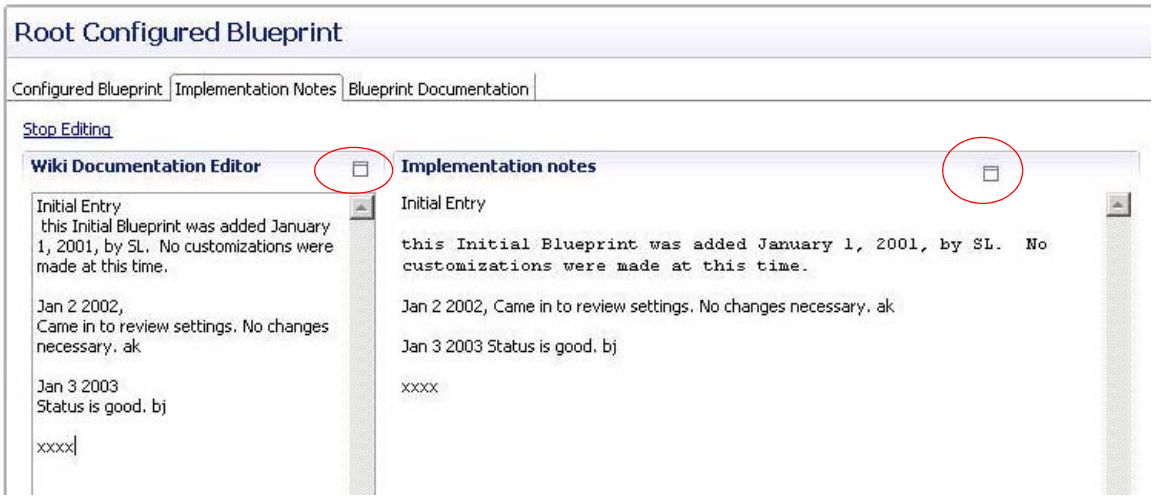

When entering text in the Wiki Documentation Editor, the left pane, the following mark up tags are available to help format your text:

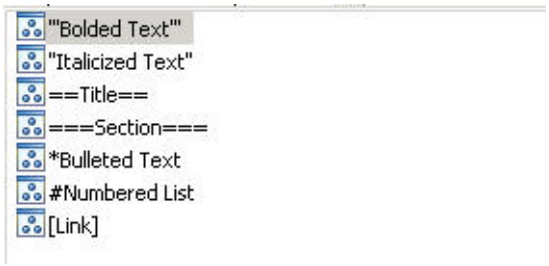

Step 3 Display this help list by pressing CTRL-space. To embed the special markup tags, you can select from the help list or type them in manually. You can also use a limited number of bracketed HTML tags.

This figure shows the result of using some of the formatting tags.

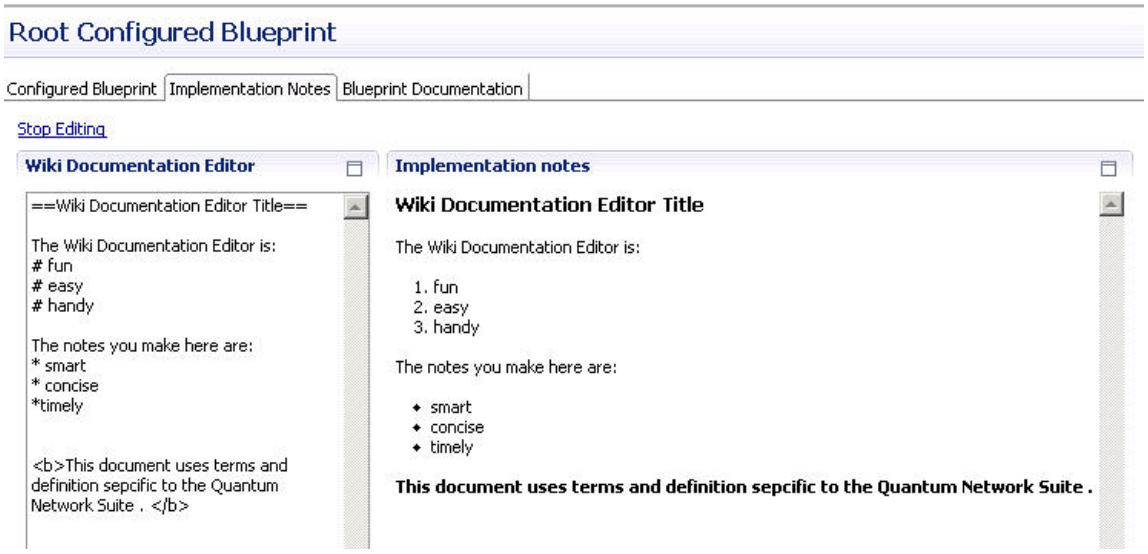

Blueprint Documentation Subtab

The information in this subtab describes what is in the blueprint and how you can use it.

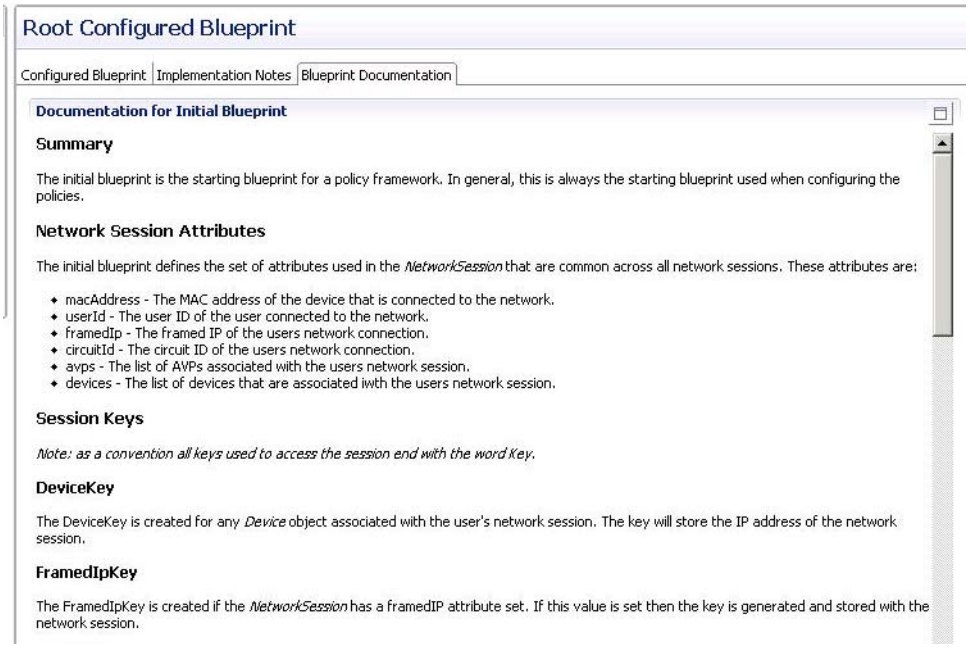

Configured Blueprint Screen

Add another configured blueprint to the Initial Blueprint with this screen. With an additional blueprint, the Initial Blueprint is augmented with all the policies, extension points, and triggered extension points already present in the one you add. All its conditions and actions execute along with those originally in the Initial Blueprint.

Step 1 In the Policy tree, select the Initial Blueprint.

Step 2 Click the Configured Blueprint subtab.

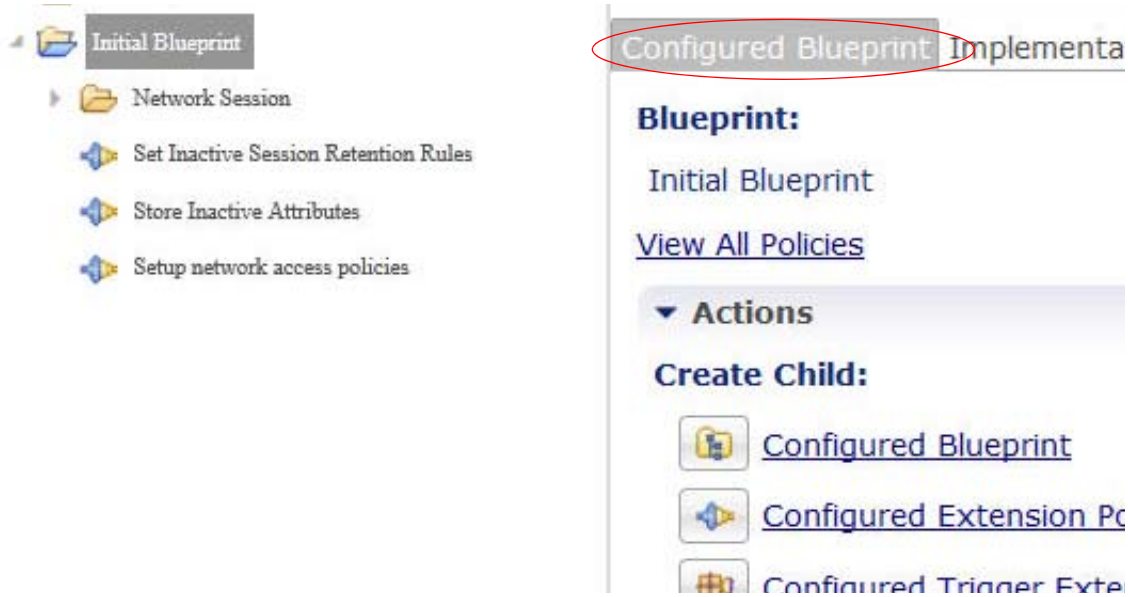

Step 3 Click the Configured Blueprint link to display the available Extension Blueprint objects.

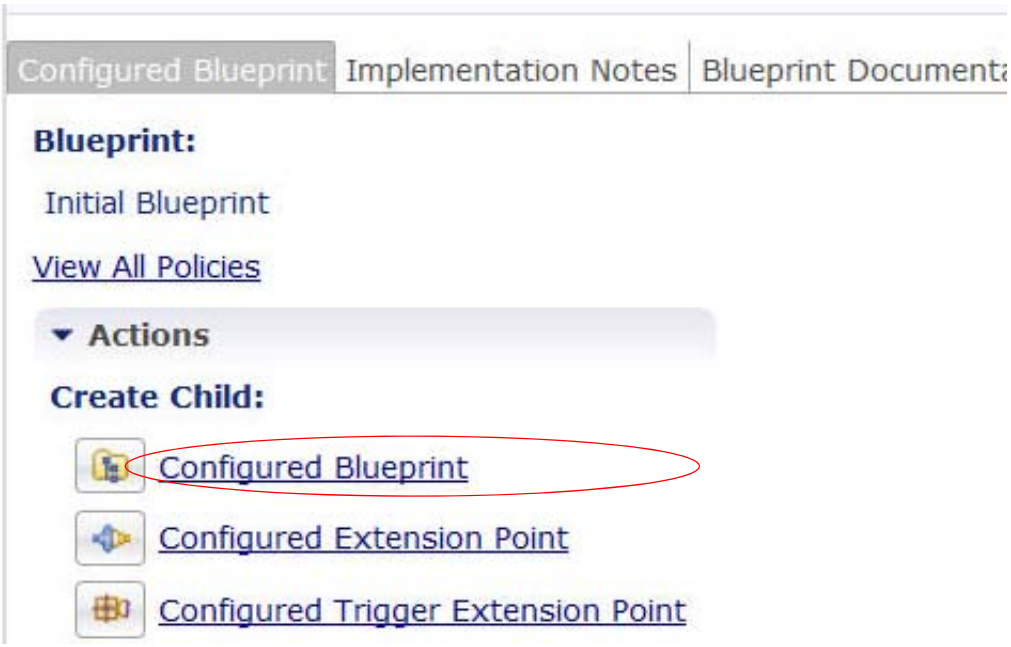

Step 4 Select an object from this list, RADIUS AAA for example.

Note Note that in the Policies tree, that at first, your new blueprint extension is placed at the bottom. It automatically corrects itself and assumes the proper processing order when the page refreshes.

Configured Extension Point Screen

This screen adds a configured extension point to a blueprint to more specifically define its behavior. Select the Configured Extension Points link to display the extension points provided.

This screen shows all the extension point available.

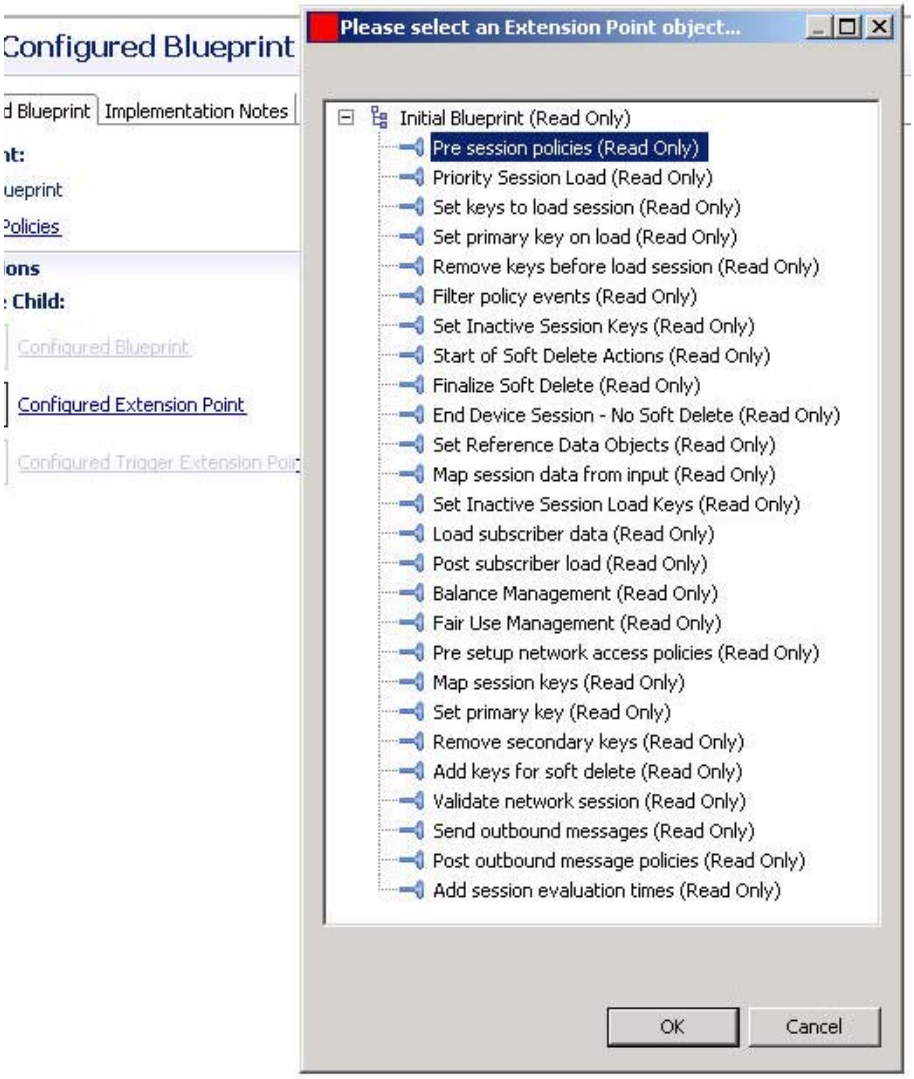

The Policies tree displays your addition immediately.

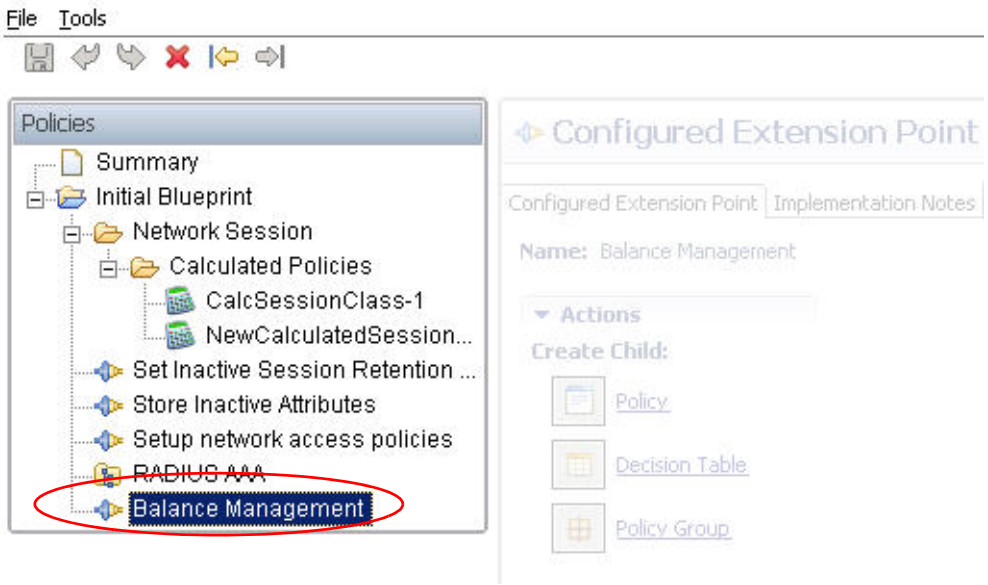

After the extension point is added to the Policies tree, you can use the instruct ions at Configured Extension Point Screens to customize it further.

Configured Trigger Extension Point Screen

You may add a triggered extension to the new blueprint to more specifically define its behavior. Select Configured Triggered Extension Points to display the trigger type extension points already created.

This screen shows all the trigger extension points available in the Initial Blueprint. Select a trigger extension point present to place it in your own Initial Blueprint, perhaps Start Session Criteria.

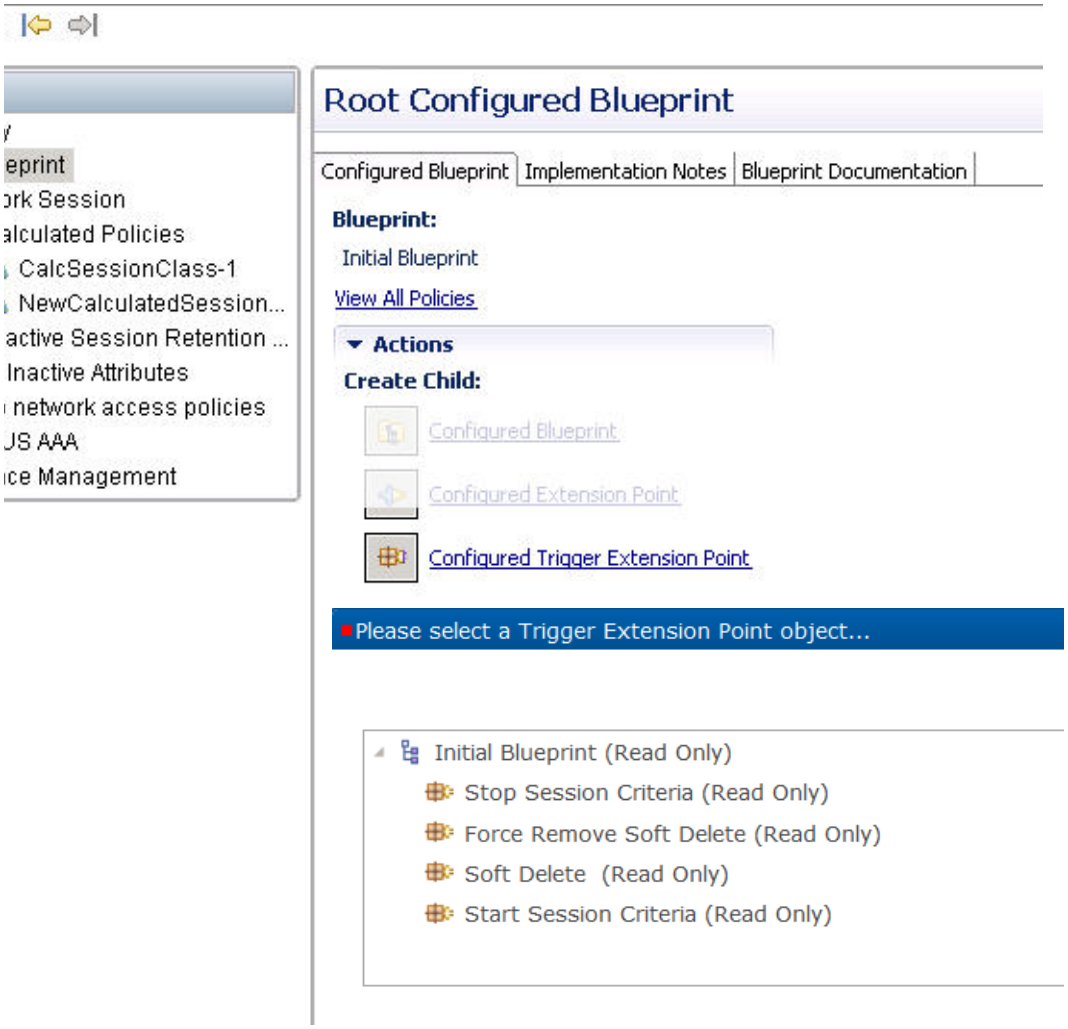

The Policies tree shows your addition immediately.

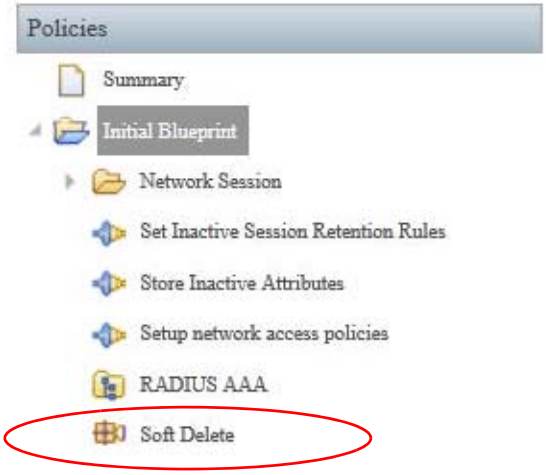

To further define or modify the newly added trigger extension point, see the instructions at Configured Extension Point Screens.

Network Session Screen

The Network Session node is always part of the Initial Blueprint. This node describes the data you want to capture for each subscriber's session.

The Initial Blueprint defines the set of attributes used in the NetworkSession that are common across all network sessions. These attributes can be:

- macAddress—the MAC address of the device connected to the network
- userId—the user ID of the subscriber connected to the network
- framedIp—the framed IP of the subscriber's network connection
- circuitId—the circuit ID of the subscriber's network connection
- avps—the list of AVPs (attribute value pairs) associated with the subscriber's network session
- devices—the list of network devices associated with the subscriber's network session

Note Note that attributes vary from deployment to deployment.

Configured Extension Point Screens

Step 1 Click any Configured Extension Point link, expanding the Initial Blueprint if necessary, and select an extension point object from the list.

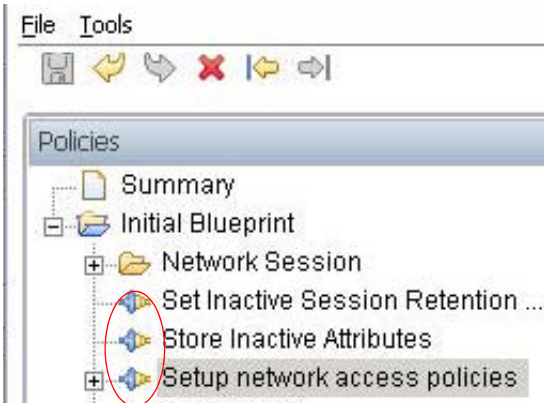

Some extension points are provided in the Initial Blueprint and are typically used and shown automatically. Less frequently used extension points can be added later as you fine tune your network operations.

For example, in the list of tasks performed in setting up a session, you can use a configured extension point to execute other policies at your session start. The next figure shows the built-in extension points in the Initial Blueprint available for your use.

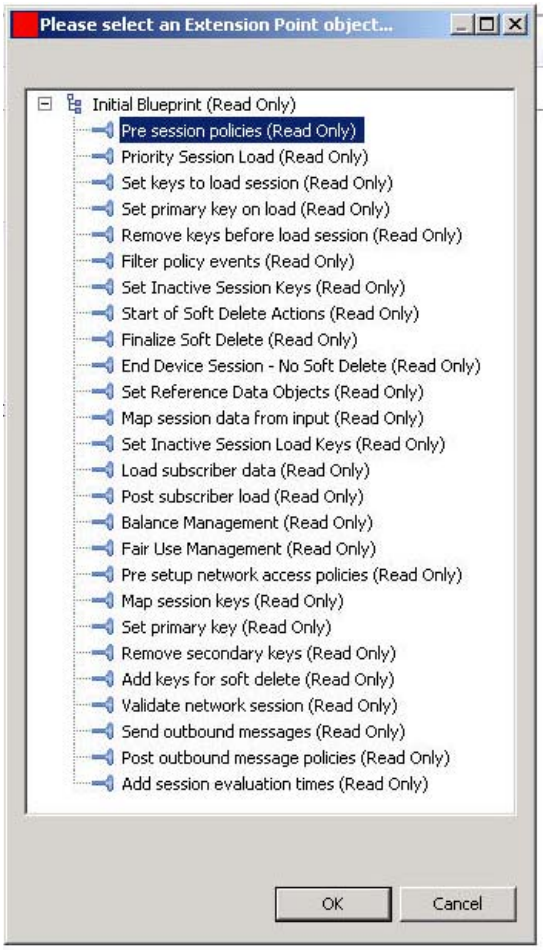

Configured Extension points use these screens:

- Policy Screen
- Decision Table Screen
- Policy Group Screen

Policy Screen

In the Policy screen, if all the conditions listed in the Conditions tab evaluate as TRUE, the actions in the Actions tab are executed. Select the conditions from a list of prepared condition phrases.

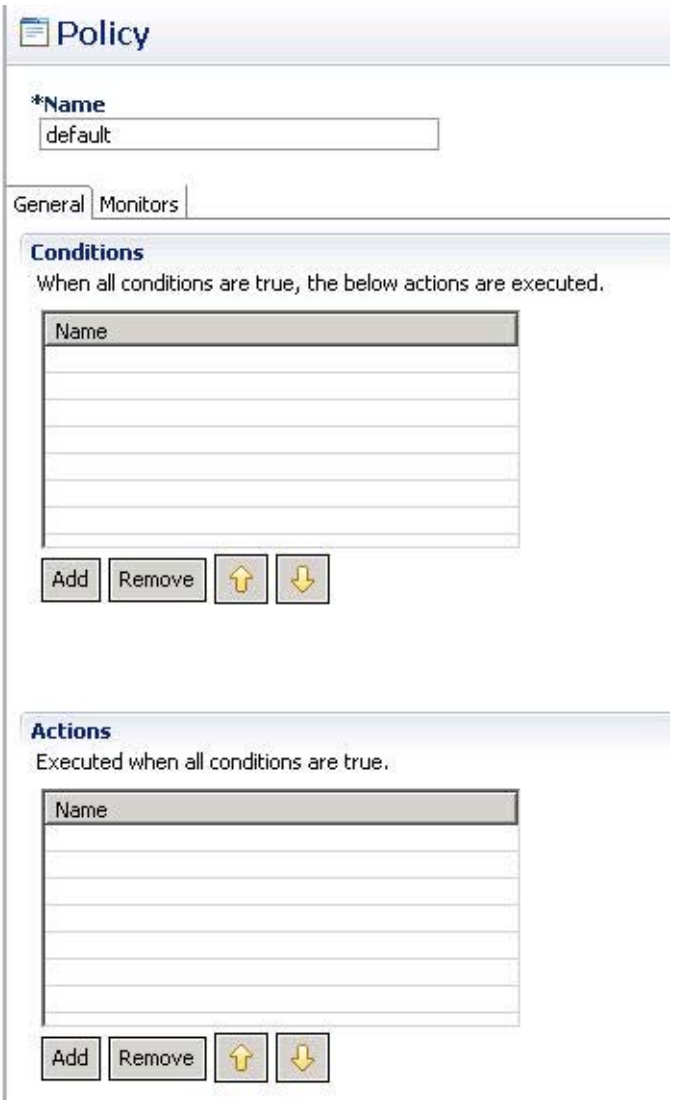

For example, Conditions serve two purposes:

- A condition's inputs determine if it is evaluated as true or false.

If a condition is true, it may produce output, which can be used by subsequent conditions or actions.

Reparent Link

The Reparent link lets you easily move the displayed policy to other locations in the Policy tree.

Click this link to move a policy group to another extension point in the policy tree. The Policy group is moved, not copied to the new location.

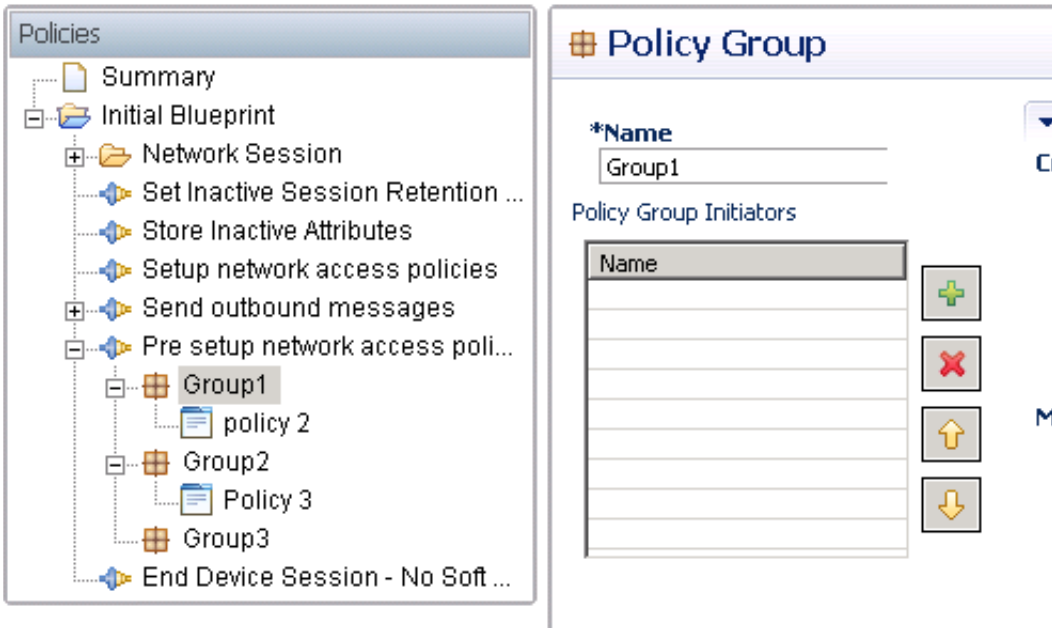

To perform a reparent:

Step 1 In the Policies tree, select the policy group you want to move. All it's policies travel with it.

Step 2 Click Reparent.

Step 3 Open the Root Configured Blueprint in the dialog window.

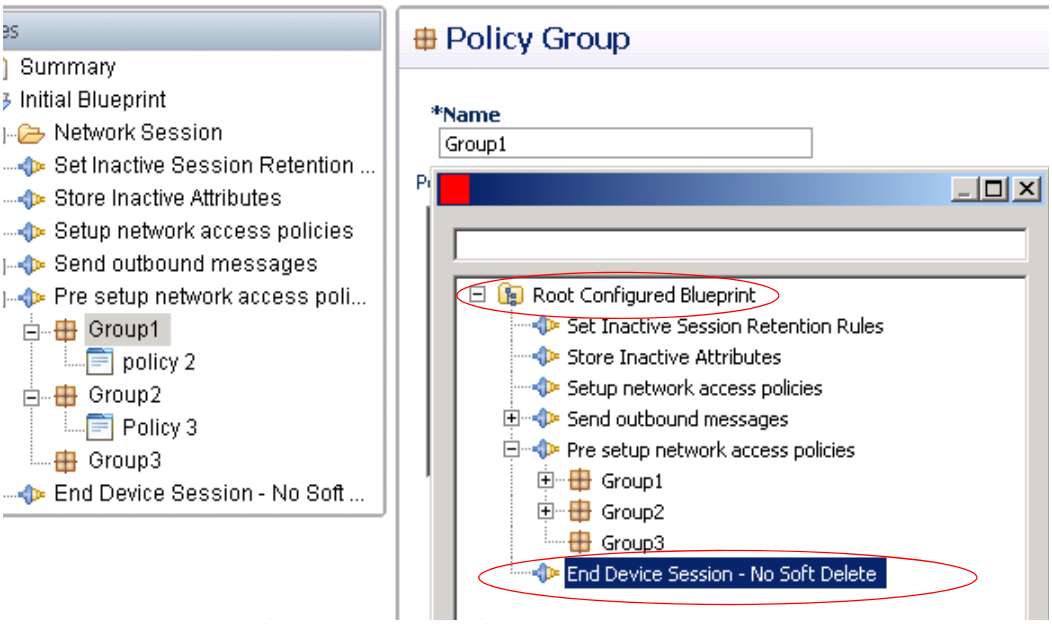

Step 4 Select the extension point you want to receive the object.

Note If you try to move a policy group A to another policy group B, policy group A is deleted.

Condition Subtab on the Policy Screen

The Conditions tab in the Policy screen lets you specify what conditions must evaluate to true and then invoke the actions you have specified in the Actions tab.

In the Name field, change the name from 'default' and enter your own name for this policy.

Construct the list of conditions that must evaluate as TRUE to invoke the policy's actions.

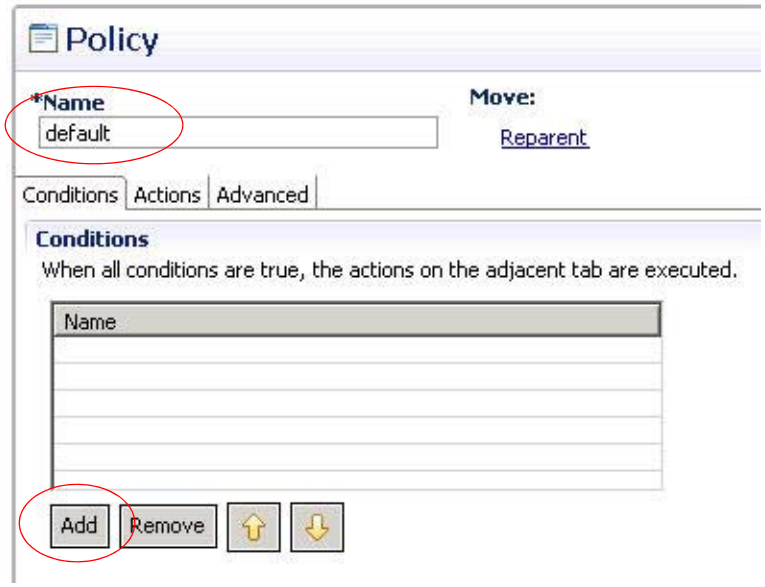

Click Add and select the condition phrase that must be met for this policy to be invoked and click OK.

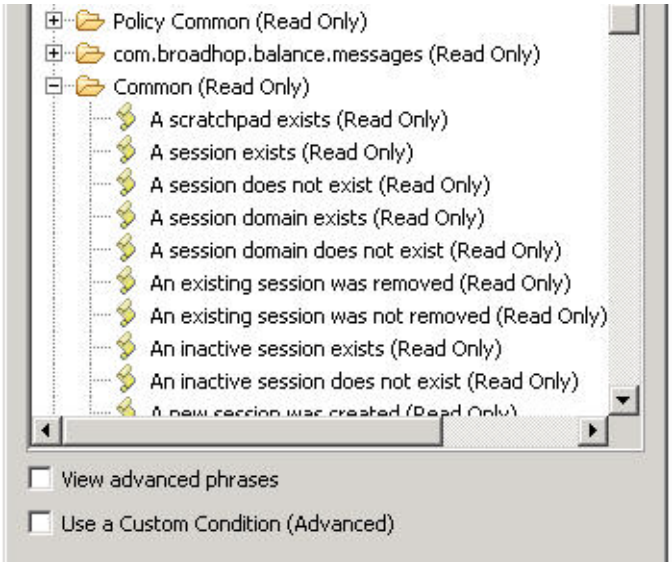

View advanced phrases Use a Custom Condition (Advanced)	These two check boxes let advanced administrators display and develop custom condition phrases. Please call your Cisco technical representative before using.

The Conditions list begins to fill in as you add conditions.

Add as many conditions as you need, then use the up and down arrows to order the list of conditions. The order of conditions in the list determines the order of evaluation. This order may affect or depend on data from a previously listed condition being available to a subsequently listed condition. The benefits of the ordering are discussed in more detail later in this document.

For each condition in the list, select it and display the inputs and outputs for that condition.

Some of the conditions in the Conditions list may have only inputs, only outputs, both, or neither. Every condition available has its own combination of inputs and outputs.

Click the Add All link to display all the variables associated with the condition phrase, or click a specific link in the Input Variables area to display a specific one.

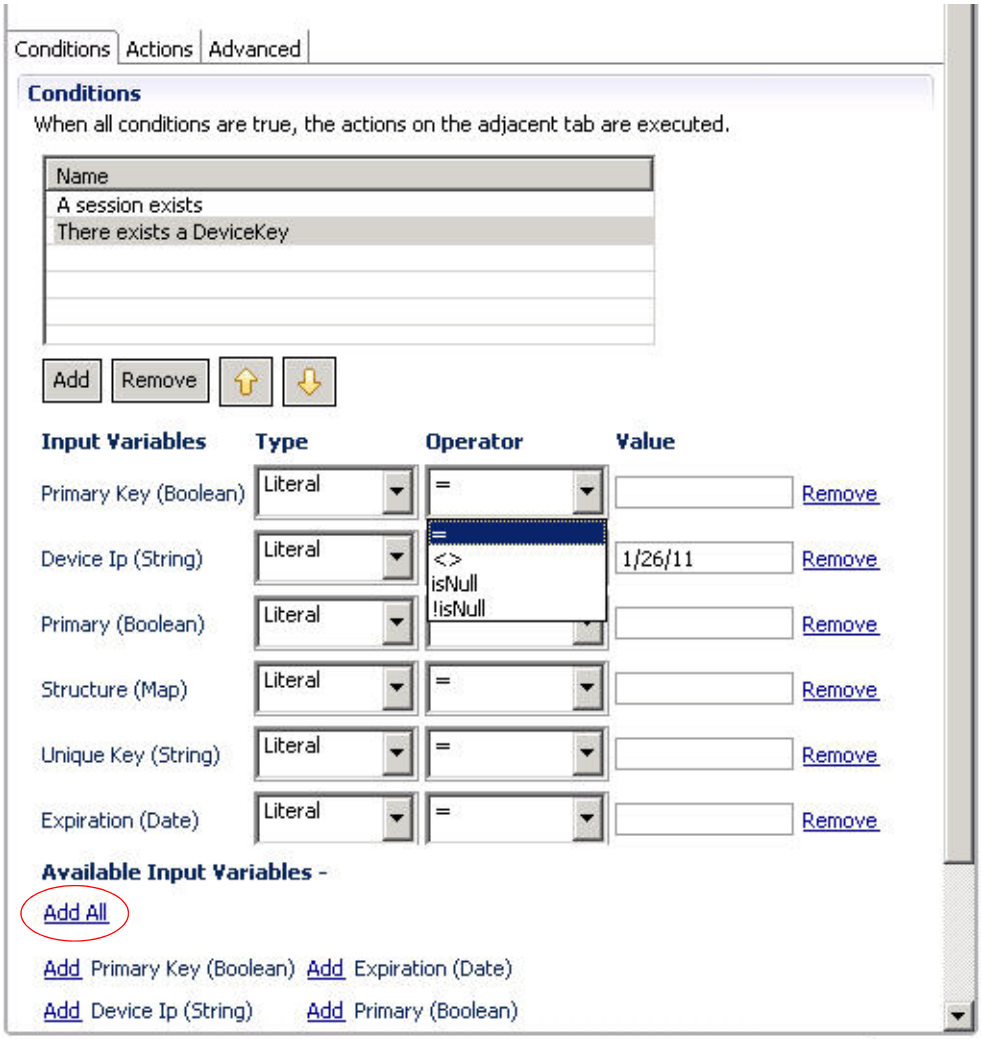

By setting definitions for input variables, you can limit the possibility of the condition itself evaluating for TRUE by having stricter rules of what TRUE means for the condition.

The input variables and outputs for the specific condition phrase appear.

- Each variable has a type of either Literal, Output, or Table.
  - A Literal type of variable contains captured or saved data.
  - An Output variable uses the output of any of the conditions previous to it in the Condition list. If your variable has the output of TypeOfUser, the value for TypeOfUser can be an input variable for the condition below it.
  - A Table type indicates that to perform the evaluation, the Cisco Policy Builder checks the corresponding decision table under the Table tab in the Decision Table screen. You have to set this up previously as described in Decision Table Screen.
- The Operators drop-down list provides numerous relational operators, such as EQUAL (=), NOT EQUAL (<>), and so on.

- The Value field lets you set a value the condition phrase must evaluate against.
- The Remove link lets you remove unnecessary inputs.

Actions Subtab on the Policy Screen

In the Actions subtab, you set the actions that are executed by the policy if *all* of the condition phrases in the Conditions tab evaluate to TRUE.

Click the Actions tab and begin to set actions that result when your conditions have been met.

In the Name field, change the name of the policy if you need to. Note that it is reflected in the Policies tree as you do so.

Construct the list of actions that occurs if all the condition phrases under the Conditions tab evaluate as TRUE.

Step 1 Click Add.

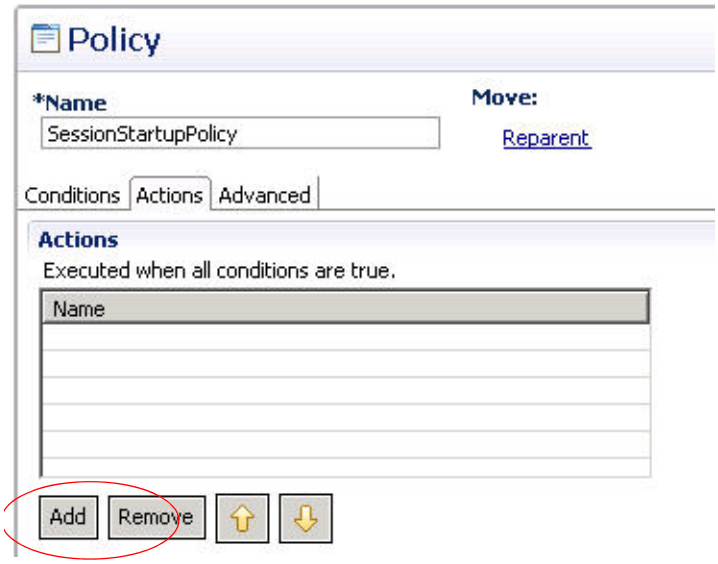

Step 2 Select an action phrase from the list.

Note Use the search-ahead field at the top to help you.

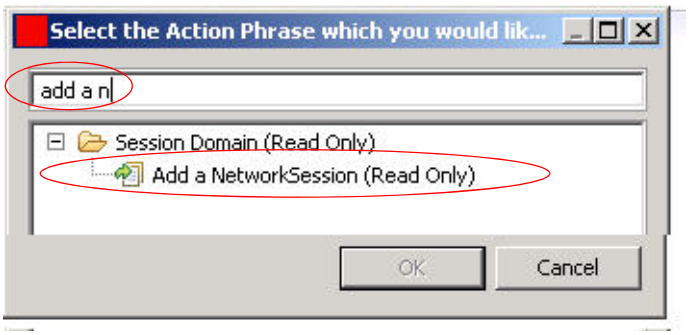

The Actions list begins to fill in as you add actions.

Add as many actions as you need, then use the up and down arrows to order the list of actions. Actions are performed in this order.

For each action in the list, select it and display the inputs and outputs for that action.

Some of the conditions in the Actions list may have only inputs, only outputs, both, or neither. Every action available has its own combination of inputs and outputs.

In the example, a condition that has this action when true, evaluates the session on the date shown in the Value field.

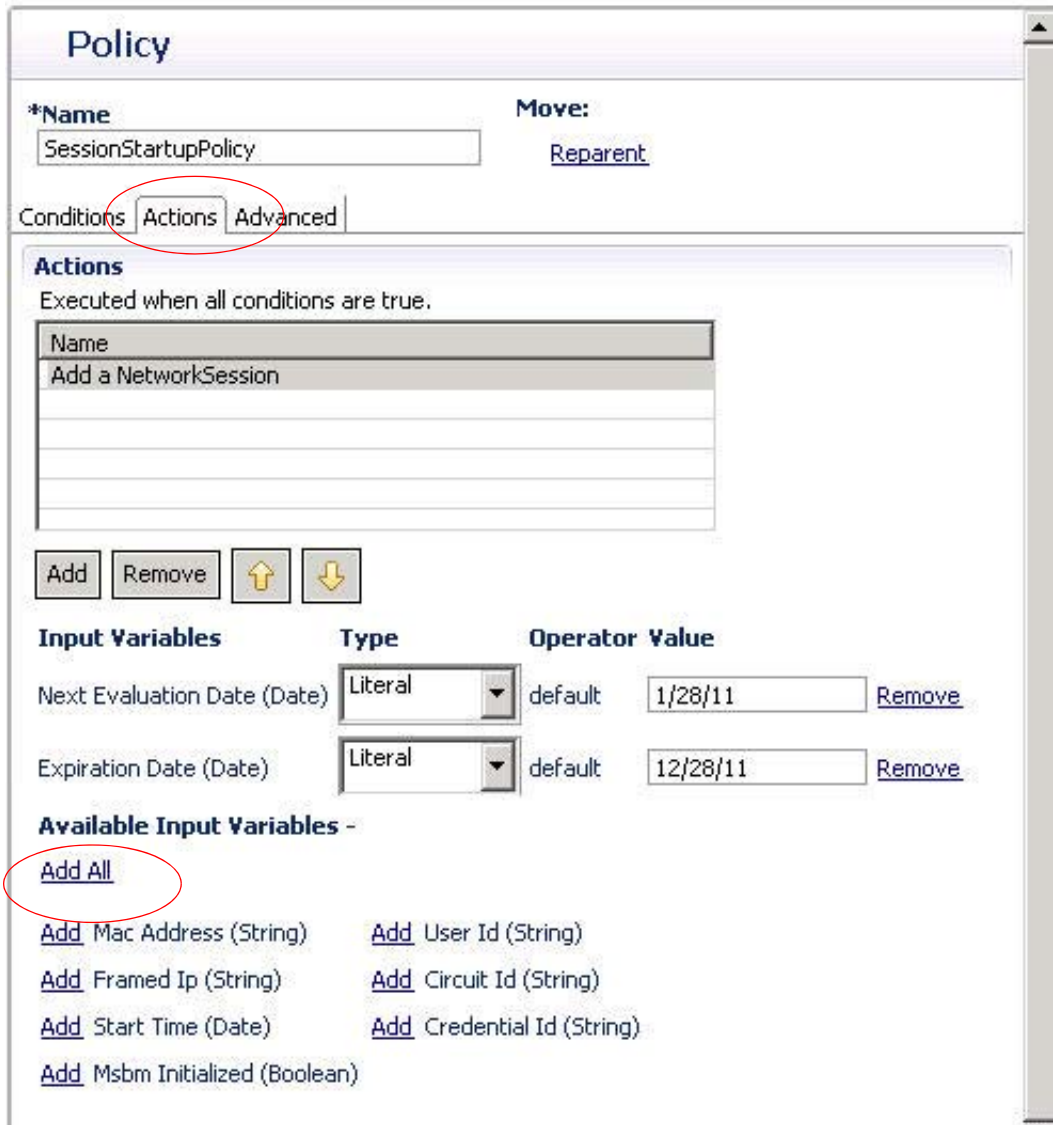

Click the Add All link to display all the variables associated with the action phrase, or click a specific link in the Input Variables area to display a specific one.

The input variables and outputs for the specific action phrase appear.

- Set Up Actions are used for notifications such as email or text, or iPhone push. You must have set these up previously for them to show here.

- Each variable has a type of Literal, Output, or Table.
  - A Literal type of variable contains captured or saved data.
  - An Output variable uses the output of any of the conditions previous to it in the Condition list. If your variable has the output of TypeOfUser, the value for TypeOfUser can be an input variable for the condition below it.
  - A Table type indicates that to perform the evaluation, the Cisco Policy Builder checks the corresponding decision table under the Table tab. You must have set up this table previously.
- The Operators drop-down list provides numerous relational operators.
- The Value field lets you set a value the condition phrase must evaluate against.
- The Remove link lets you remove unnecessary inputs, however, it's best to just not use them but leave them present.

Advanced Subtab on the Policy Screen

Settings in the Advanced subtab affect how the policies are executed in terms of their priority within a policy group.

In the drop-down list, select a priority level for this policy within its policy group.

Select the appropriate triggers with the check boxes.

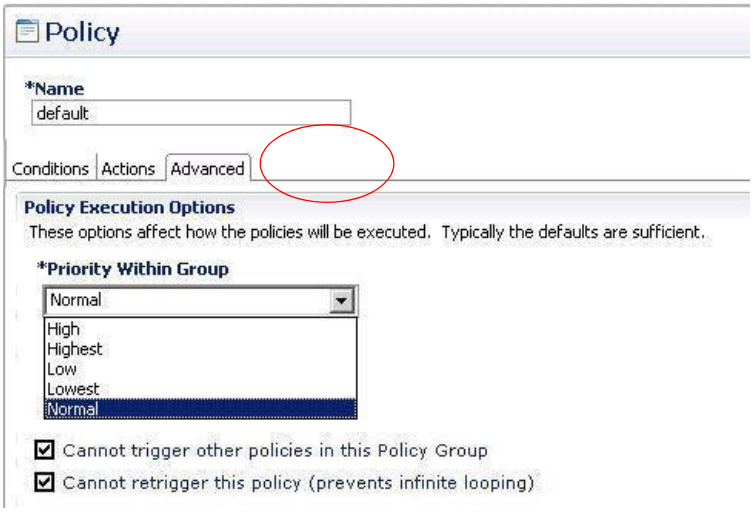

Name	Name of the policy.
Priority	Within GroupIn the drop-down list, select a priority level for this policy within its policy group. By default, all policies in a policy group execute at the same time as Normal priority rather than in any specific order. With the Priority drop-down list, you can specify that this policy executes before other policies in the group by giving it the Highest priority level. Conversely, you can specify that a policy executes only after all of the other policies in the group execute by giving it the Lowest priority. • Highest • High • Normal • Low • Lowest
Cannot trigger other policies	Select this check box to prevent this policy from triggering other policies in this policy group regardless of inputs and output variables. The default is TRUE.
Cannot retrigger this policy	Select this check box to keep this policy from triggering itself. The default is TRUE.

Decision Table Screen

A decision table is a way to execute a policy using a range of input variables.

You may have several policies that are similar but not exactly the same. With a decision table, you can look at the table, and administer policies without inventing new, custom policies.

In this version of CPS, decision table strategies are often handled by defining a new service using the Services tab.

For example, in a decision table you might have three columns: Plan, QoS, and Rate.

The rows of the table contain inputs used when evaluating conditions. If a decision table is used, several use cases can be presented for evaluation by a condition.

With a decision table set like this, you can have a condition that says "If the Plan is Gold, then set QoS to what is defined in the decision table, and set Rate to what is defined in the decision table.

The decision table is an easy and concise way to control the important inputs of the policies. With a decision table, editing or changing policies themselves is needed only infrequently.

Table Subtab on the Decision Table Screen

The Table subtab on the Decision table holds the associated conditions and actions defined in the Policy screen for each entry in the table.

Select an extension point in the Policy tree, and then click the Decision Table link on the right.

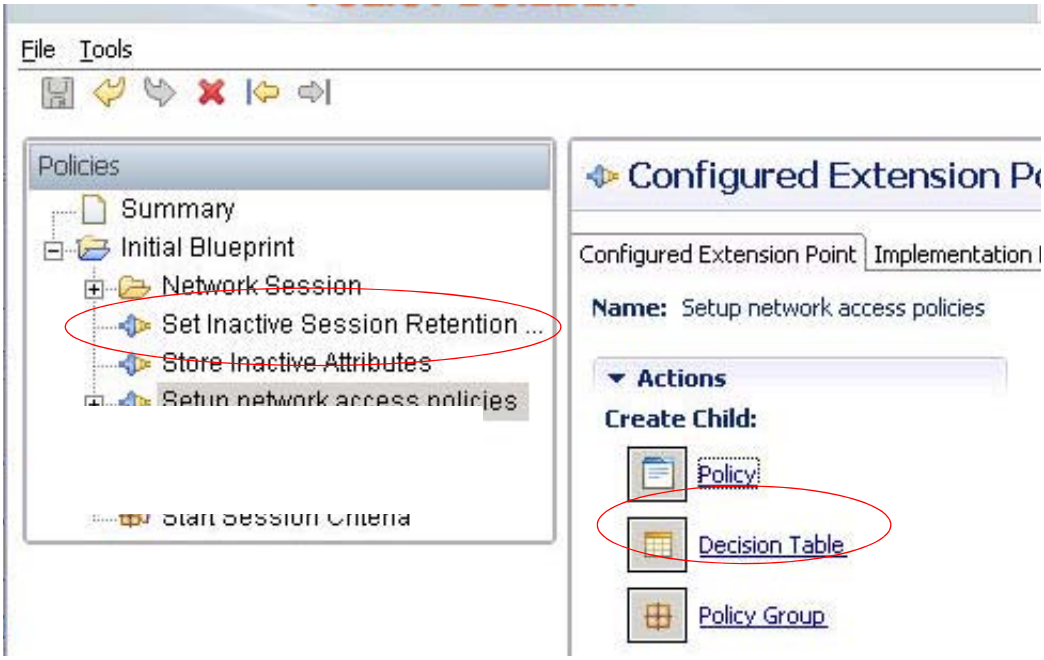

Click the Table subtab.

When you select the Edit Columns link, you can begin to define the name, data type, and data values of the rows in the decision table.

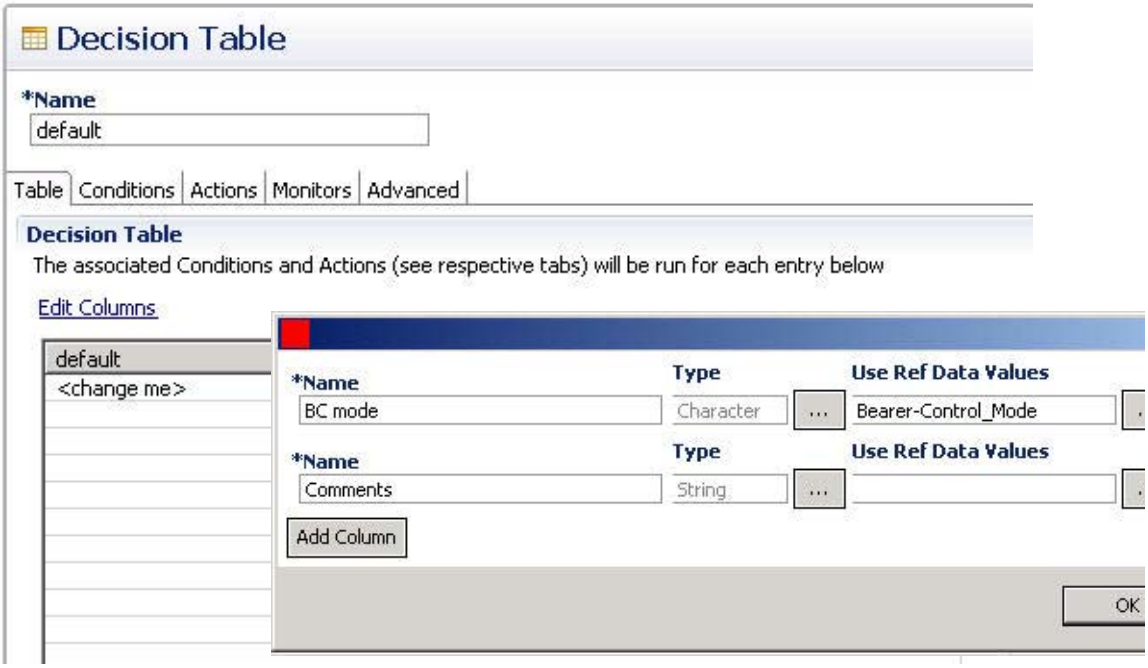

Name	Assign a name to the column.
Type	Assign the data type to the column. You have these choices. 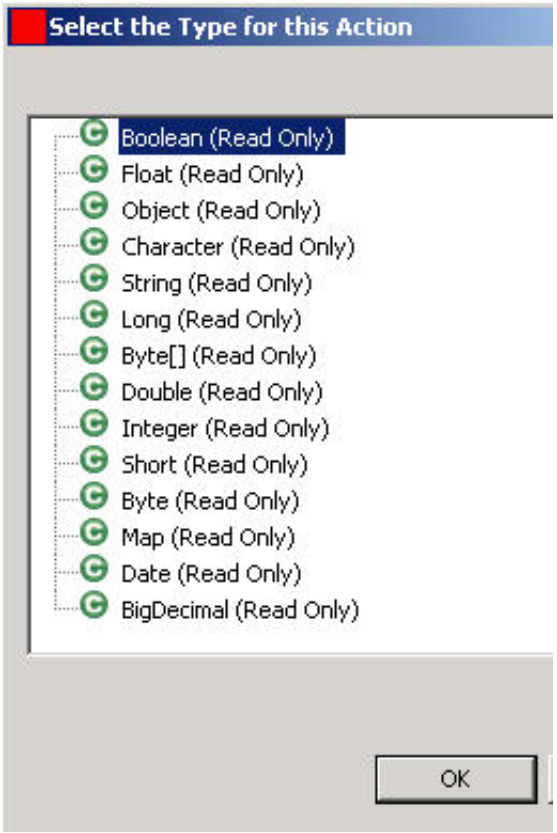

User Ref Data Values	This field sets up commonly used data elements that you do not want the subscriber to invent. Rather, you provide them and the subscribers select from a list. 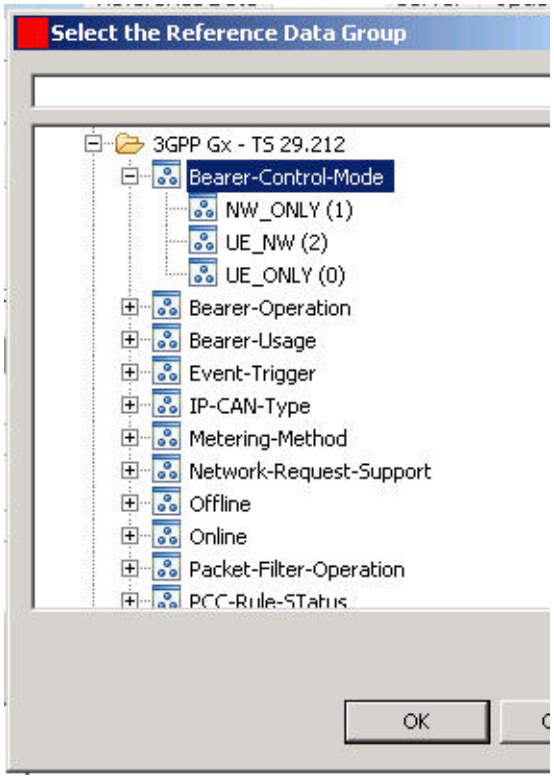
Up Dn	Use the Up and Down button to change the order of the columns as they display in the Decision Table.
Remove	Click Remove to remove a column specification entirely.
Add Column	Click this to add another column heading to the decision table.
OK	When you are finished defining your columns in the decision table, click OK.

Click Add Row on the decision table screen.

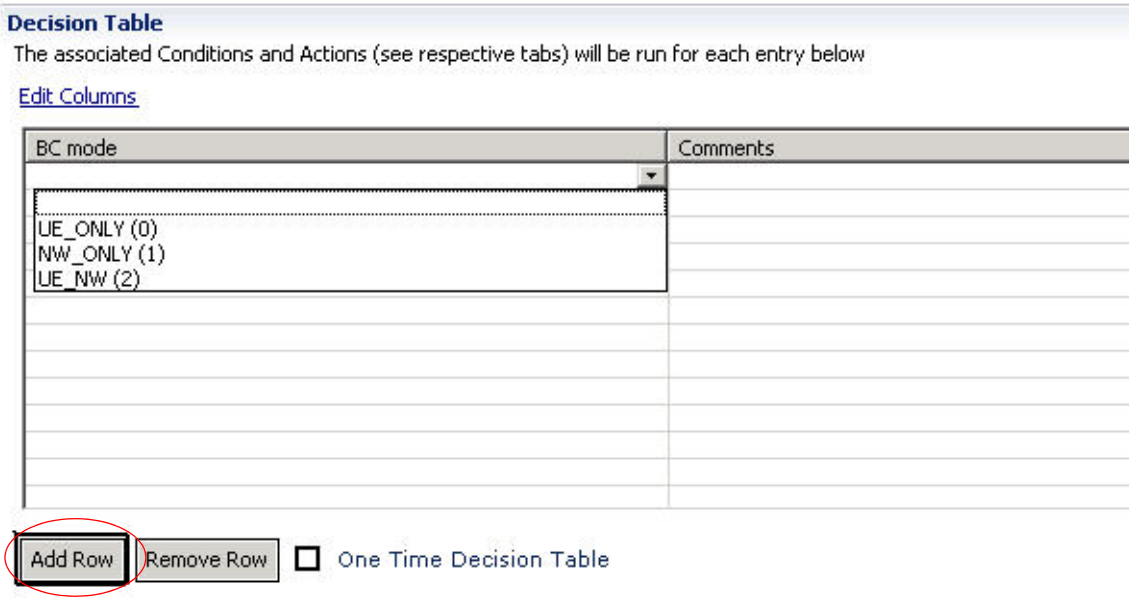

With the drop-down list for the row, begin to populate the decision table.

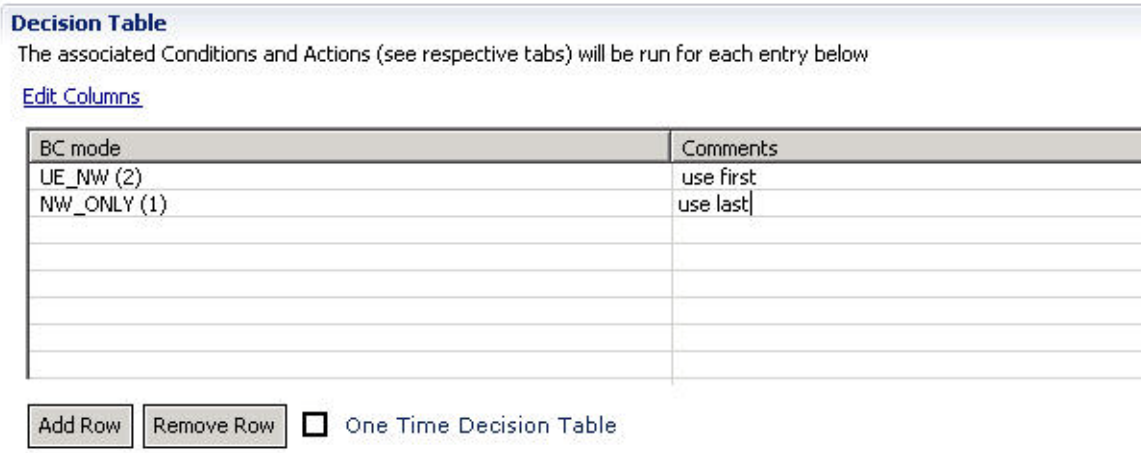

Other Subtabs on the Decision Table Screen

Use the other tabs in the decision table screens as described in these sections:

- Condition Subtab on the Policy Screen
- Actions Subtab on the Policy Screen
- Advanced Subtab on the Policy Screen

Policy Group Screen

Grouping policies together and treating them as one object lets you efficiently invoke or evaluate policies.

Use The Policy Group screen to set several policies to all run without any unnecessary processing or degradation of performance of evaluating the conditions of each policy.

Click any Configured Extension Point link, expanding the Initial Blueprint if necessary, and select an extension point object from the list.

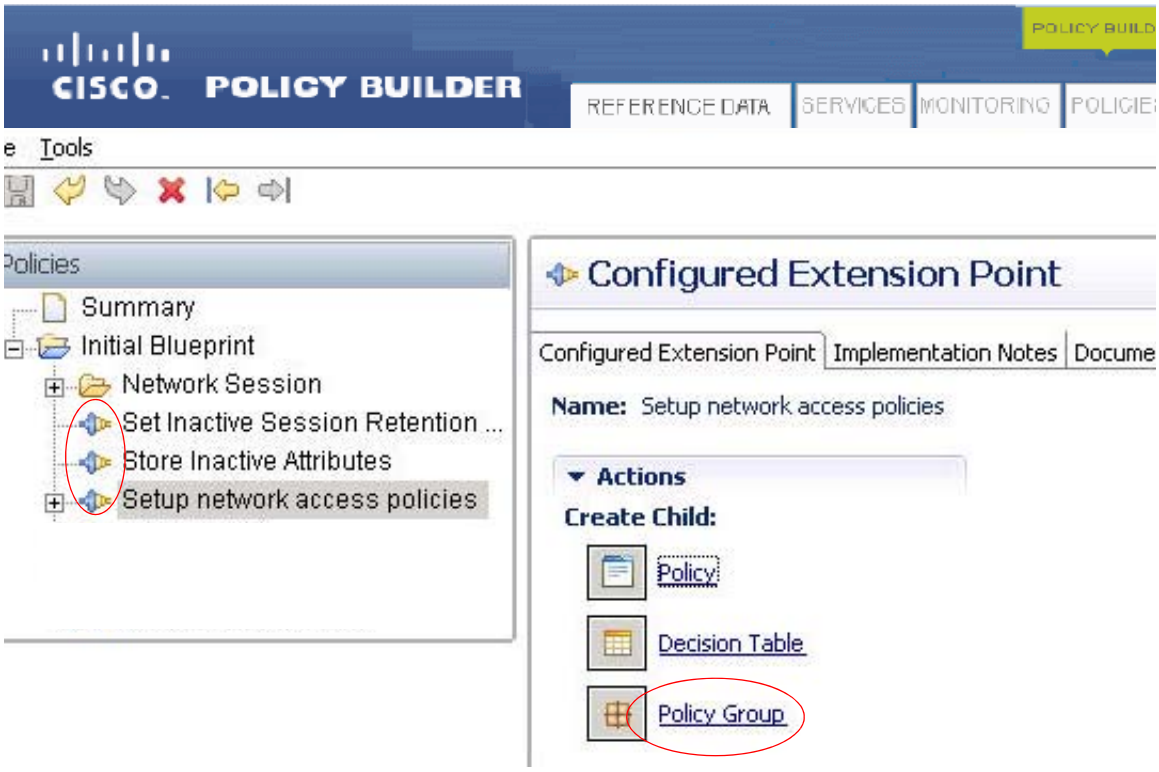

The Policy Group screen appears.

Note Recall that all policies in a group, by default, execute at the same time. The Policy Group Initiators is a list of conditions that says that if all of its conditions are met, then run the policy group's conditions as usual.

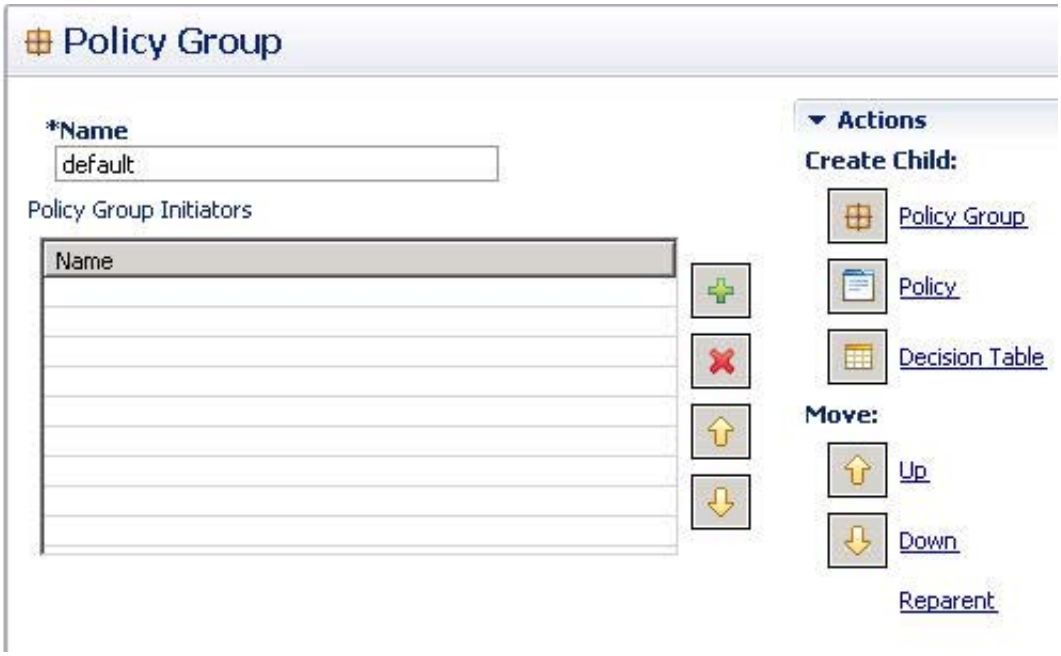

Name	Name of the policy group that appears in the policy tree.
Policy Group Initiator Name List	The logical name of the policy group that are triggered.
Green Plus	Click to add a row to the list of policy group initiators.
Red X	Click to remove a row.
Yellow Up Arrow Yellow Down Arrow	Click the up and down arrows here to change the order of the rows in the criteria table.
Create Child	Use the links in the Create child area to define an initiator of a policy group, a policy, or a decision table within the decision table.
Initiator Name	What you type here populates the Policy Group Initiators list, which provides a name for the various initiator conditions.
Conditions List: Name	Name of the condition you select from a prepared list.

Add Button Remove Button	Click Add or Remove here to add or remove a condition in the list. The condition you set for your policy initiator comes from the prepared Condition Phrase list. 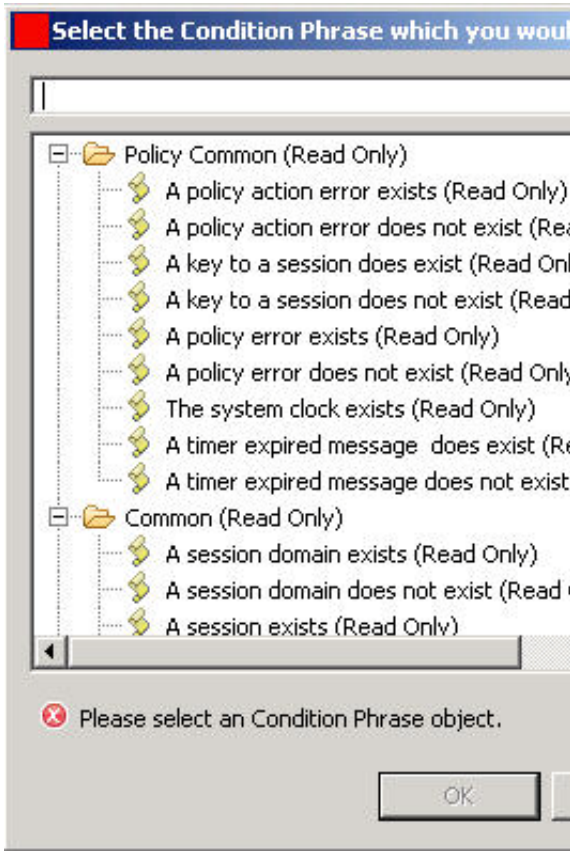
Yellow Up Arrow Yellow Down Arrow	These arrows move the condition up and down in the condition list.
Reparent	See Reparent Link.

Combine Policies into a Group

First create the Policy Group itself.

Click the Policies tab > an extension point > Policy Group Link.

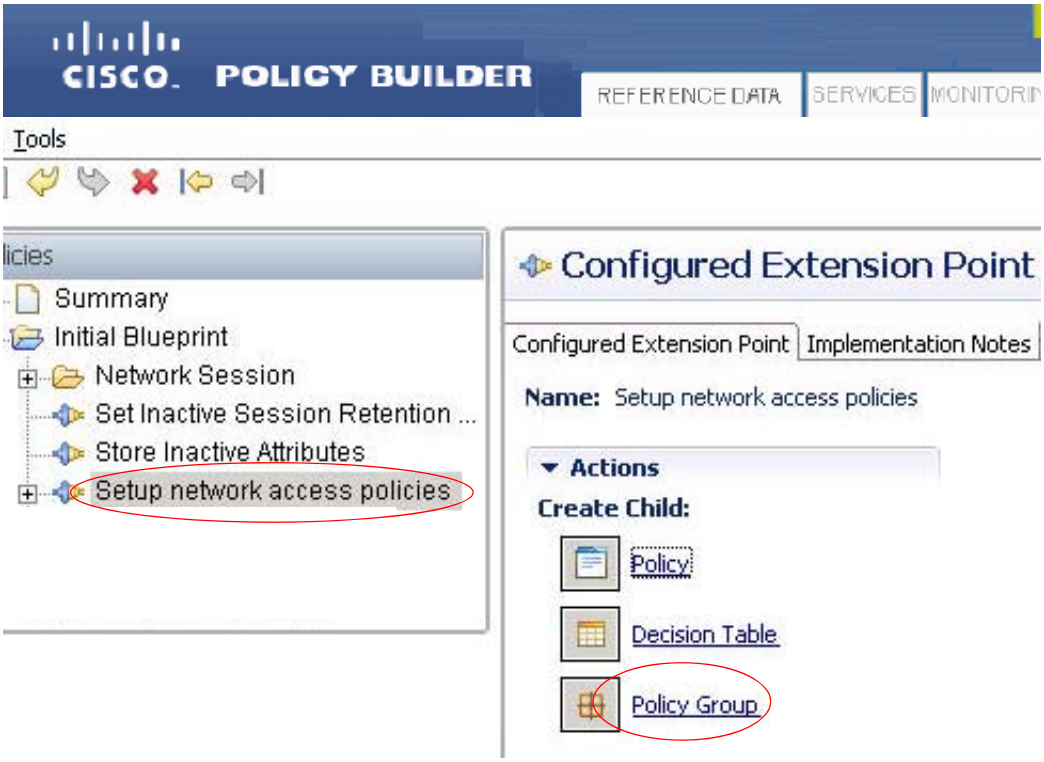

This displays the Policy Group screen where you make the Policy Group itself.

Enter a Name for your Policy Group.

Click the green plus button to add single policies to this Policy Group.

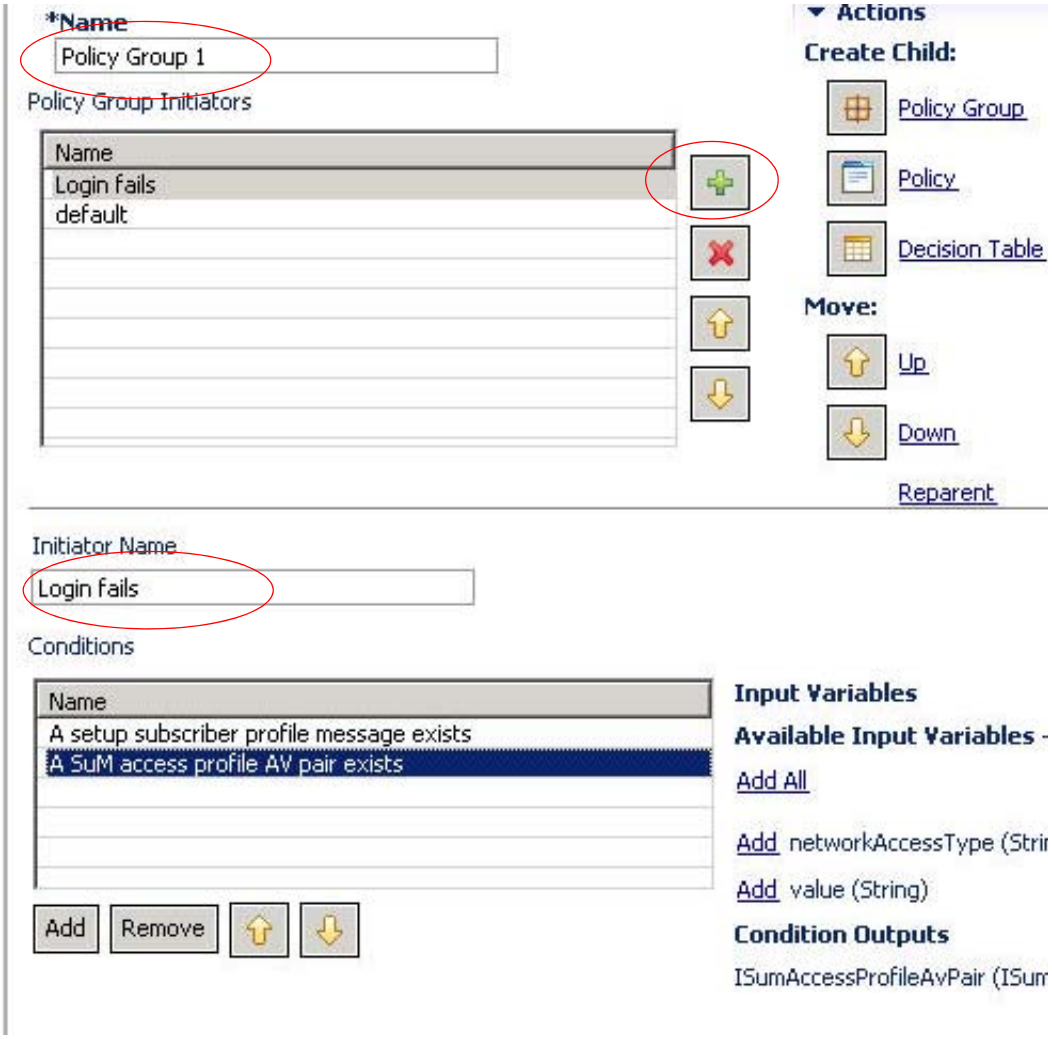

Name	Name of the policy group that appears in the policy tree.
Policy Group Initiator Name List	The name of the policy group that are triggered. The Policy Group Initiator list has no hierarchy or order of precedence. That is, if the policy group is triggered, all initiators are invoked, and not in any particular order.
Green Plus Sign	Click to add a row to the list of policy group initiators.
Red X	Click to remove a row.
Yellow Up Arrow Yellow Down Arrow	Click the up and down arrows here to change the order of the rows in the criteria table. This is for logical purposes, but does not affect the order of execution of policies in the list.

Create Child	Use the links in the Create Child area to define an initiator of a policy group, a policy, or a decision table.
Move Yellow Arrows	These arrows move the condition up and down in the condition list.
Reparent	See Reparent Link.
Initiator Name	What you type here populates the Policy Group Initiators list, which provides a name for the various initiator conditions.
Conditions	Name of the conditions you select from a prepared list. This list presents all condition in the Cisco Policy Builder for you to choose from. 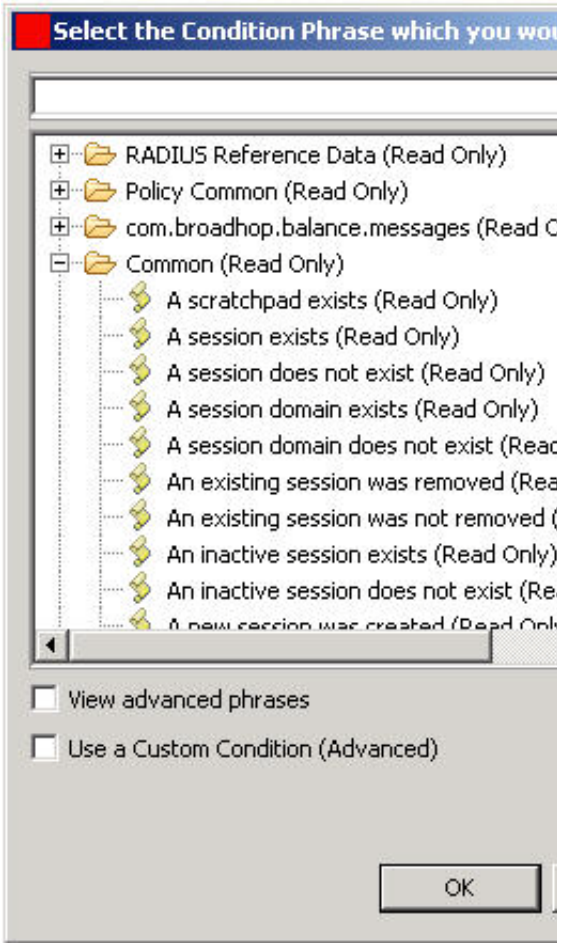
View advanced phrases Use a Custom Condition (Advanced)	These two check boxes let advanced administrators develop custom condition phrases. Please call your Cisco technical representative before using.

Add and Remove Buttons	These Add and Remove buttons affect the list of conditions.
Yellow Arrows	These arrows move the conditions up or down in the list. The list position of a condition here affects processing behavior.

Select a condition to evaluate the input variables and outputs of the condition. See Condition Subtab on the Policy Screen to manipulate the inputs and outputs.

Check the order of your policy groups in the policy tree. Move a policy group up or down in the policy tree with the arrows in the Move area. The order of a policy group in the tree may affect the intended behavior of the policy.

Configured Trigger Extension Points Screens

A triggered extension point is a way to add conditions for triggering the policies inside a policy group.

Session Stop or Start Session Criteria

Almost every policy configuration you create has criteria for stopping and starting the subscriber session. To review or create these criteria, use this procedure.

Click Policies > Initial Blueprint > Stop Session Criteria, which is a configured trigger extension point.

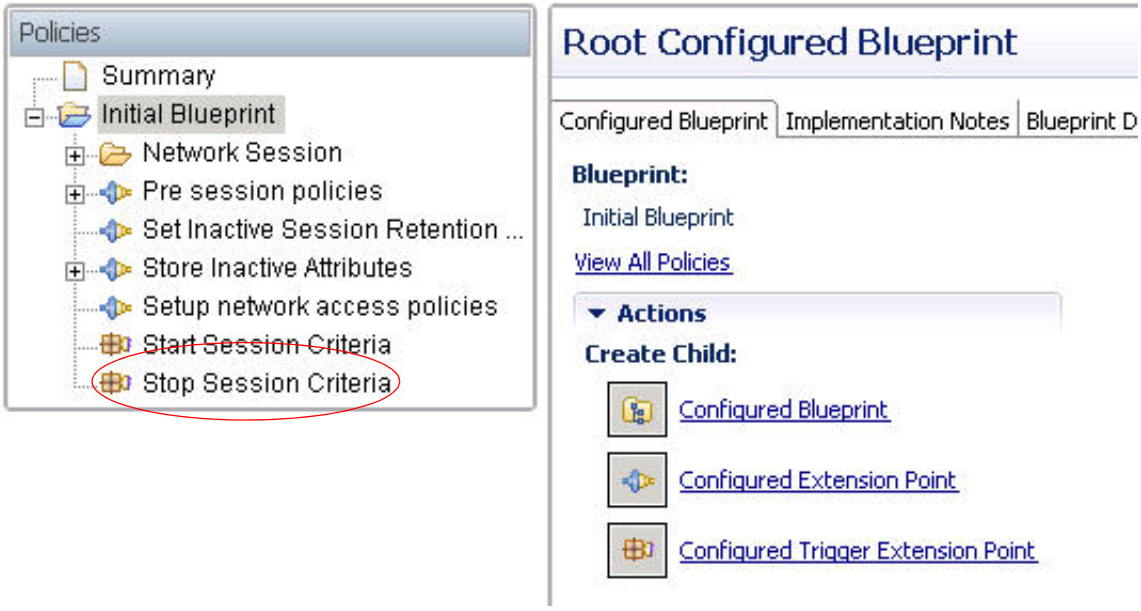

Note Click the Configured Trigger Extension Point link if you need to add Stop Session Criteria to the Policies tree.

View the Configured Trigger Extension Point screen.

Note Recall that the Policy Group Initiators list is OR based. If *any* of the initiators are TRUE, then the policy group is triggered.

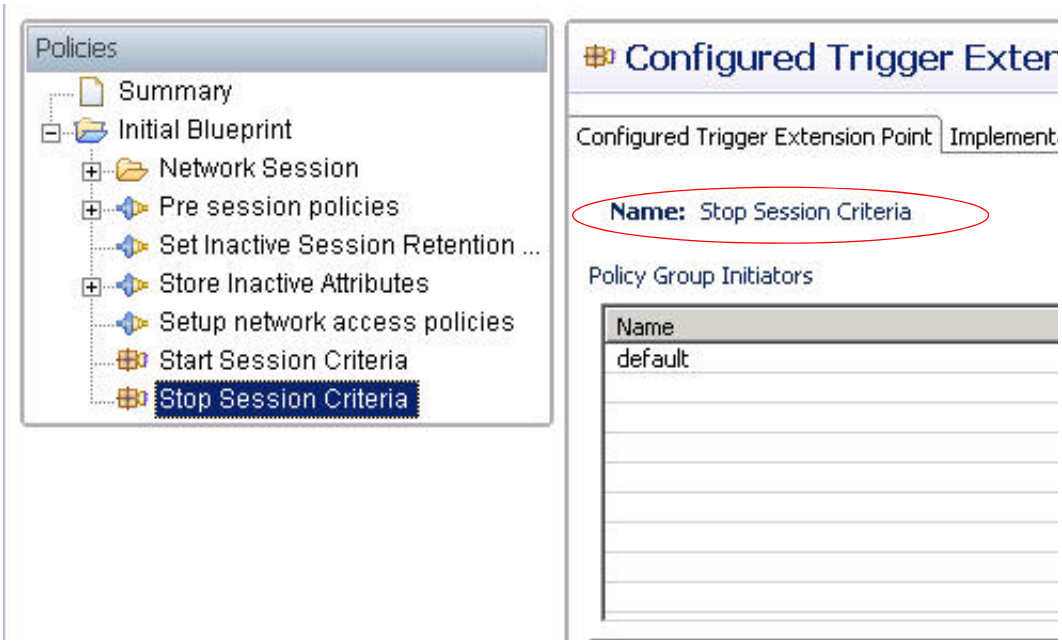

Name	The name of the policy group that are triggered.
Green Plus	Click to add a row to the table of stop criteria.
Red X	Click to remove a row.
Yellow Up Arrow Yellow Down Arrow	Click the up and down arrows here to change the order of the rows in the criteria table. The order affects processing behavior.

Click the green Plus icon to add a policy group initiator to the list. The bottom area of the screen displays additional fields when you do.

Enter the name for this extension point, NoTime, which indicates that the subscriber is trying to start a session, but needs to purchase more time.

The initiator name you type in the field automatically populates the Policy Group Initiators list at the top.

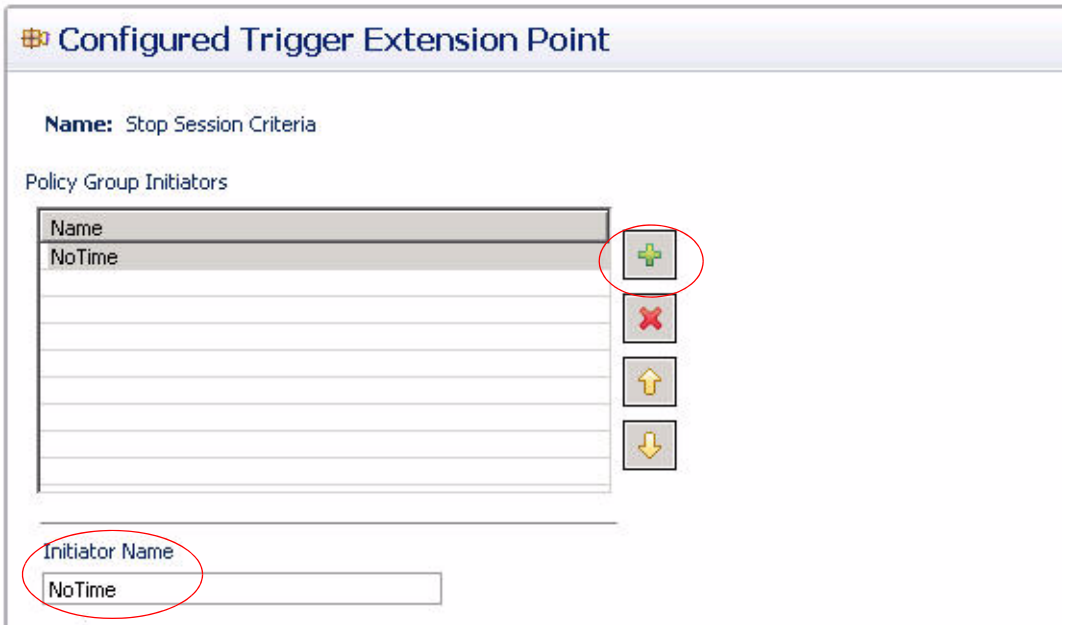

:

Use the fields and icons as described here.

Policy Group Initiators table	This area is actually a list of conditions, any of which can be true, for the configured extension point to have its conditions executed. In the Conditions area of this screen, you can add and remove conditions that must be met to ease or tighten the evaluation of the initiator.
Initiator Name	This is the name of the initiator.
Conditions table	Add, remove, and configure conditions in this table so that Cisco Policy Builder can determine when to start or stop a subscriber session.
Add Links	For each row you create in the condition table, click the links to the right of the table. Input and output fields specifically related to that row are presented for editing. Each condition has it's own input and output variables.
Add Button Remove Button	Click Add or Remove here to add or remove a condition in the table.
Yellow Up Arrow Yellow Down Arrow	These arrows move the condition up and down in the Condition table. The order affects processing behavior.

Use the Add button and populate the Conditions table with these conditions:

- A policy action error exist

- A key to a session does exist
- A policy action error does not exist

Select a condition in the Conditions table at the bottom, perhaps the second condition, A key to a session does exist. This conditions variables are listed to the right.

Click the Add All link to display all of the variables necessary for this condition.

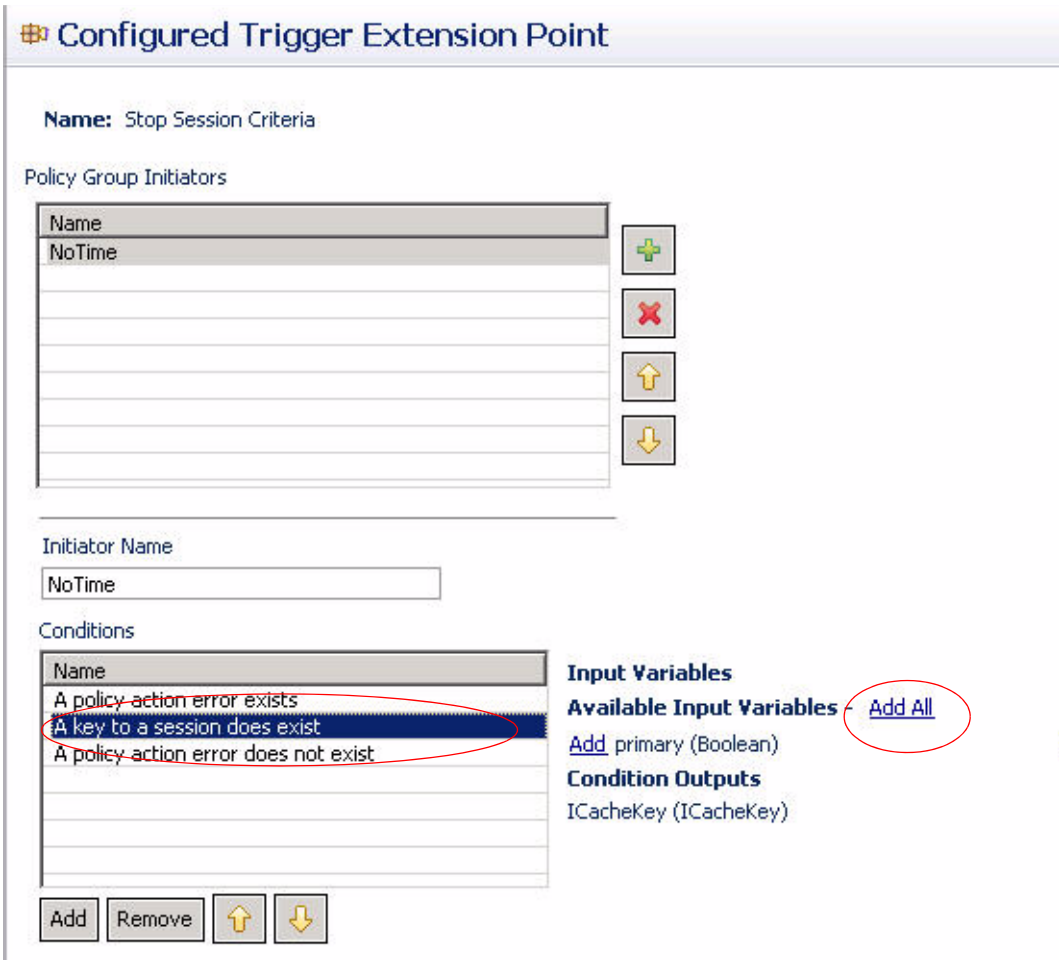

Input Variables	These labels show what you are defining on the right.
Type	Means that the variable uses data as it has stored. Output means that the output of the condition immediately previous in the list is used as the variable.
Operator	Use this drop-down list to choose a relational operator.
Value	The value you are measuring against.
Remove link	Click this link to remove the variable from the condition itself. Not recommended.

The variables for this condition are displayed. No further changes are necessary and the procedure is complete.

Note that the condition outputs the ICacheKey for use elsewhere in your Polices constructs.

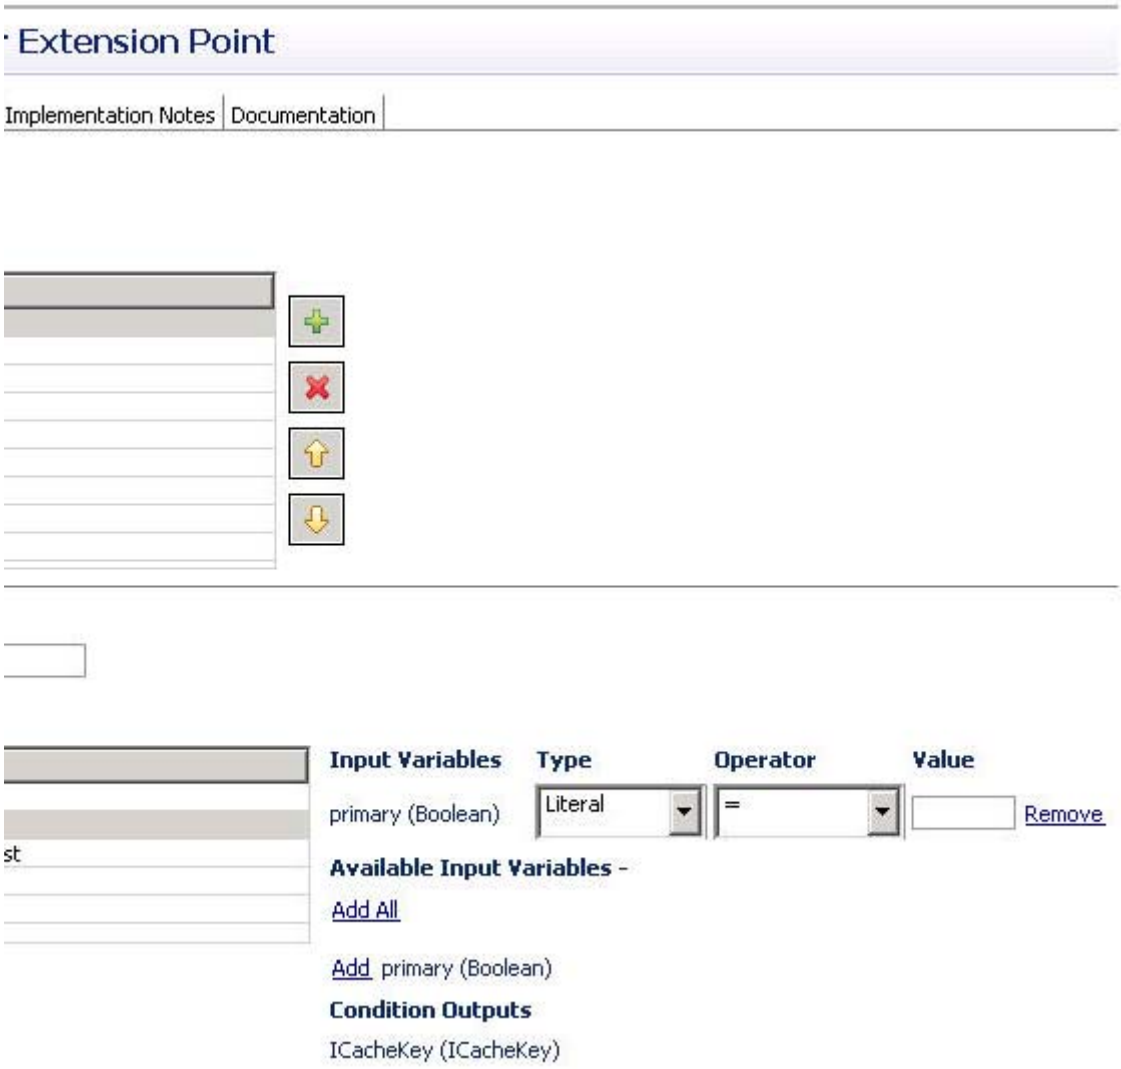

Client Repository Configurations

Revised: July 10, 2015

This appendix covers the following sections:

Configuring a Client Repository

This section discusses how to use a Client Repository for configuration tasks.

The data for the objects you develop in the Cisco Policy Builder are stored in the client repository in these ways:

In the local directory, you can temporarily save them and work on them later without check out.

- When committed to the version control software. This occurs when you save to the client repository with the Save icon or with CTRL S on the keyboard.

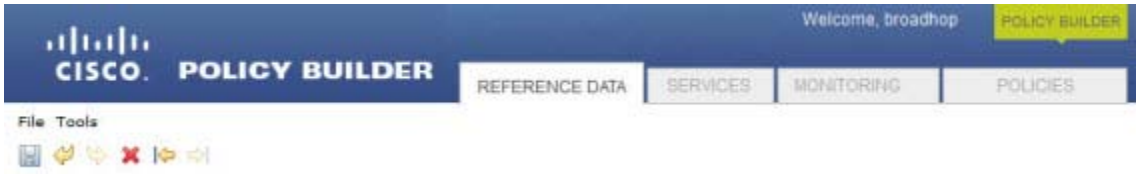

- When you publish, the data in the client repository is checked in to the server repository. Typically, you publish to a Test repository, and when satisfied, publish to a Production repository.

Creating the First Client Repository

Before setting up your client repository sites, you should have version control software installed and available. In this example, a Subversion server is located at this URL:

```
http://svn01/repos/trunk/main
```

Cisco Policy Builder comes with the Subversion version control software. You can use your own copy of Subversion server, but not another version control software product such as IBM® ClearCase® or others.

Adding a Client Repository

Step 1 Start the Cisco Policy Builder interface.

Step 2 At the **Choose Policy Builder data repository...** screen as shown below, use the drop-down list and select the item Add a New Repository.

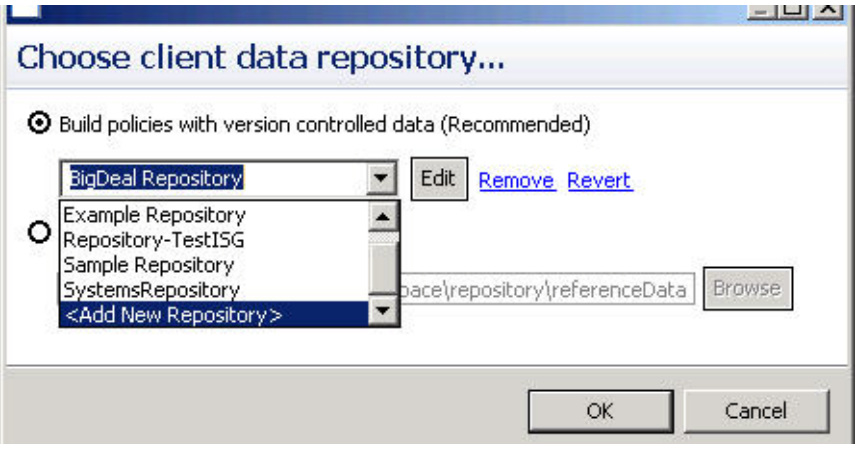

Step 3 The Repository window appears. Fill out the Repository window.

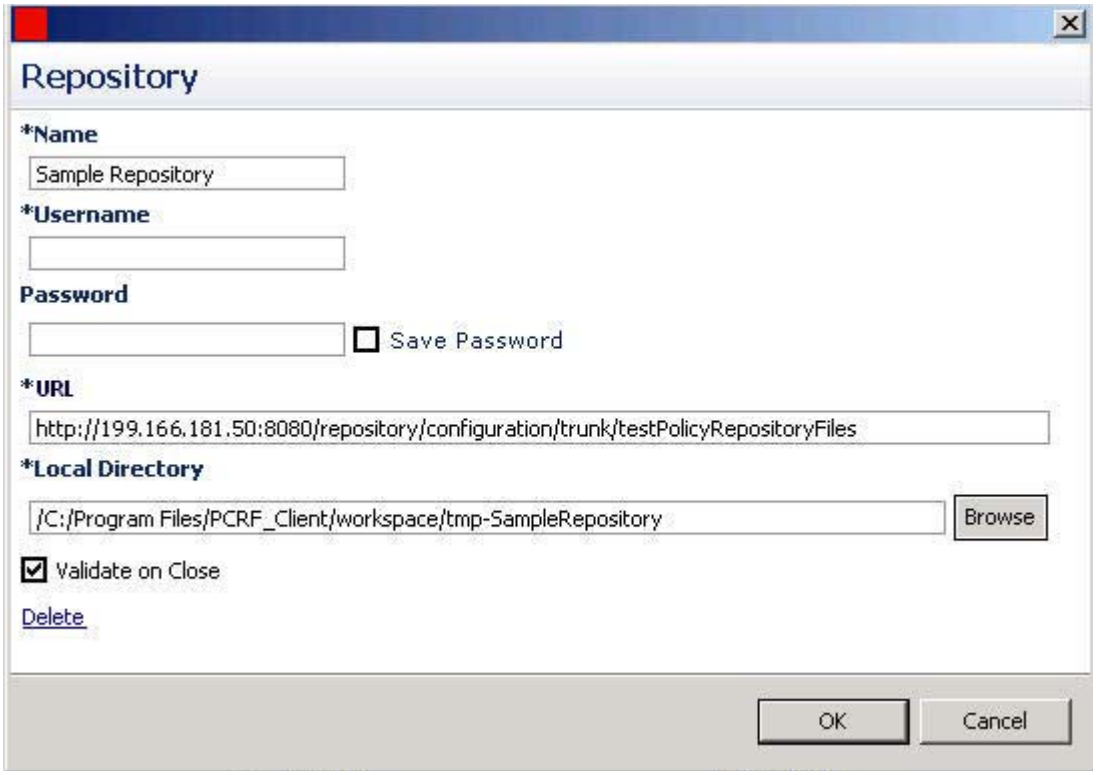

Name	This required field uniquely identifies your repository's site with a friendly name
Username and Password	Enter a username that has read and write access to the repository. A password, used in conjunction with the Username, permits or denies access to make changes to the repository.
Save Password	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.
URL	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. Go to your version control software to create the branch before assigning it in this screen. Enter the URL of the branch of the version control software server that are used to check in this version of the data.
Local Directory	This value need not be changed. This is the location on the hard drive where the Policy Builder configuration objects are stored in version control. When you click either Publish or Save to Repository, the data are saved from this directory to the version control application specified by the URL above.
Browse	Use the Browse button to either locate or create a place on the local drive to temporarily store your GUI object data.
Validate on Close	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.
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 4 Fill in the fields on this screen.

Step 5 Click OK to save your work to the local directory.

Note When you change screens, Cisco Policy Builder automatically saves your work.

Develop the habit of saving your work to the local directory by clicking on the diskette icon or CTRL S on the keyboard.

If you are on the **Choose Policy Builder data repository....** screen, these save steps are not necessary because the data are stored in a manner specific to each client install.

Step 6 If you are ready to commit these changes to the version control software,

View the Policy Builder main screen > select the menu File > select the item Save to Client Repository.

Changing the Details About the Client Repository

Use this procedure to change any of the following details of your Client Repository:

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

Step 1 Start the Cisco Policy Builder application.

Step 2 On **Choose Policy Builder data repository....** screen, click the Edit button.

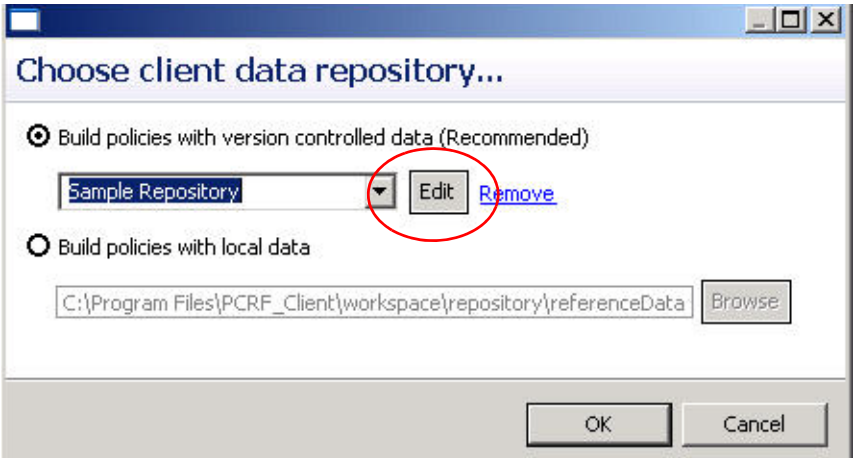

Step 3 View the Repository screen and make your changes.

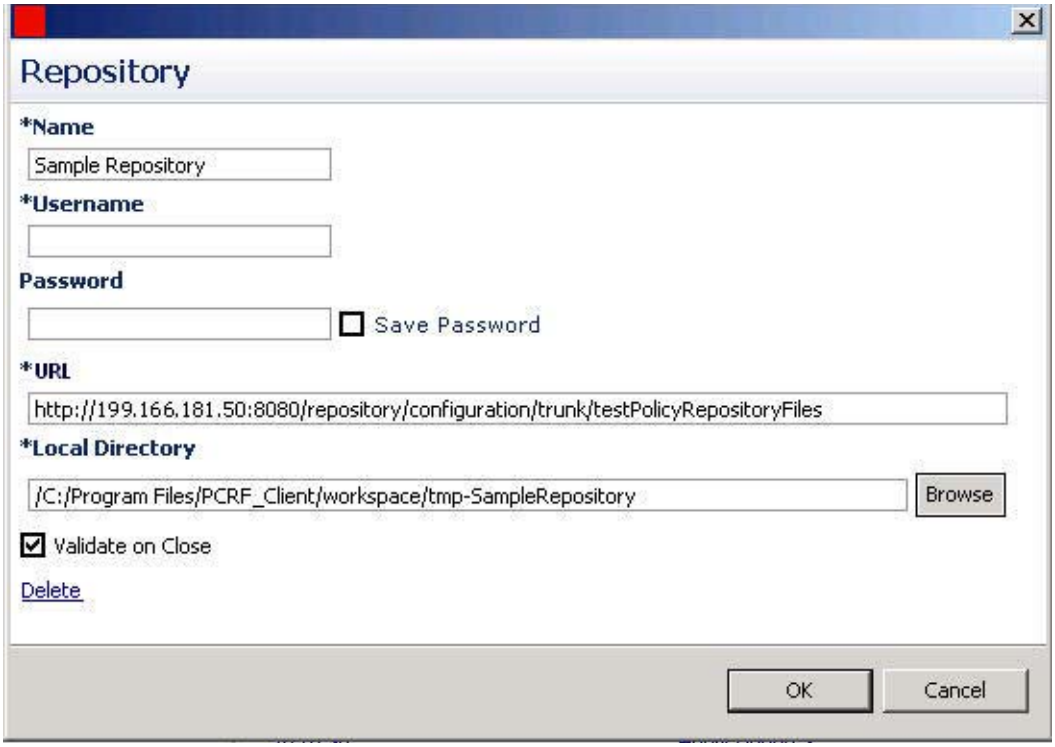

Name	This required field uniquely identifies your repository's site with a friendly name
Username and Password	Enter a username that has read and write access to the repository. A password, used in conjunction with the Username, permits or denies access to make changes to the repository.
Save Password	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.
URL	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. Go to your version control software to create the branch before assigning it in this screen. Enter the URL of the branch of the version control software server that are used to check in this version of the data.

Local Directory	This value need not be changed. This is the location on the hard drive where the Policy Builder configuration objects are stored in version control. When you click either Publish or Save to Repository, the data are saved from this directory to the version control application specified by the URL above.
Browse	Use the Browse button to either locate or create a place on the local drive to temporarily store your GUI object data.
Validate on Close	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.
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 4 Click OK to save your work.

Removing a Client Repository

This procedure removes a repository from Cisco Policy Builder. This procedure does not delete the data residing at that http site. You can still go to the URL, but not from Cisco Policy Builder unless you set up another client repository.

Step 1 From the **Choose Policy Builder data repository....** screen, use the drop-down list and select the repository you want to remove.

Step 2 Click the Remove link.

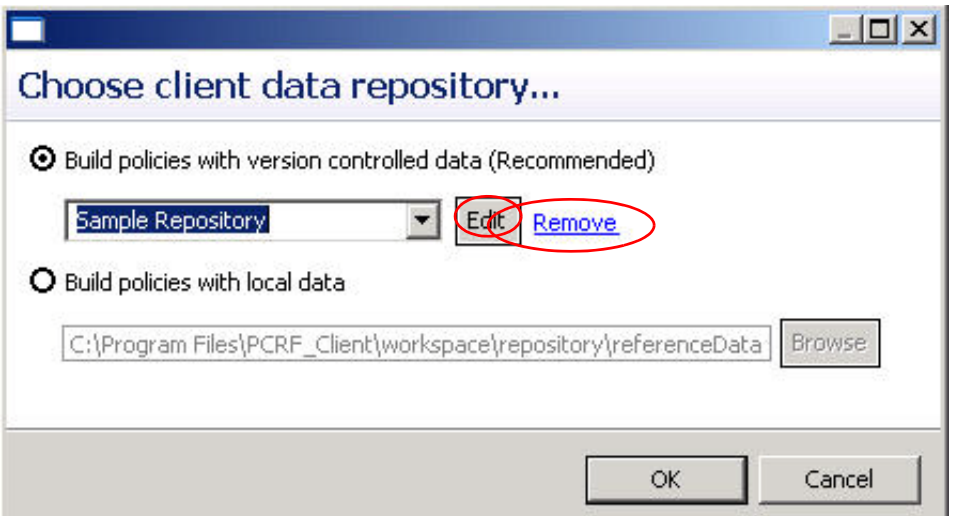

Step 3 Answer the confirmation question to confirm or deny the deletion.

Step 4 Check that you are deleting the repository you want.

Note Currently there is no way to undelete a repository or to undo this action.

Publishing the Client Repository

To put changes into effect and have the Cisco Policy Builder server pick up the configuration changes made in you 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 Cisco Policy Builder unless you are completely satisfied with the configuration data in your client repository.

- Use the Cisco 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 file config.properties.
- 'Unpublish' with an SVN delete and restore.

When you're ready to put your Cisco Policy Builder changes into production, you'll publish them to Subversion. This preserves version history.

Step 1 To publish in Cisco Policy Builder, click File > Publish to Runtime Environment.

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

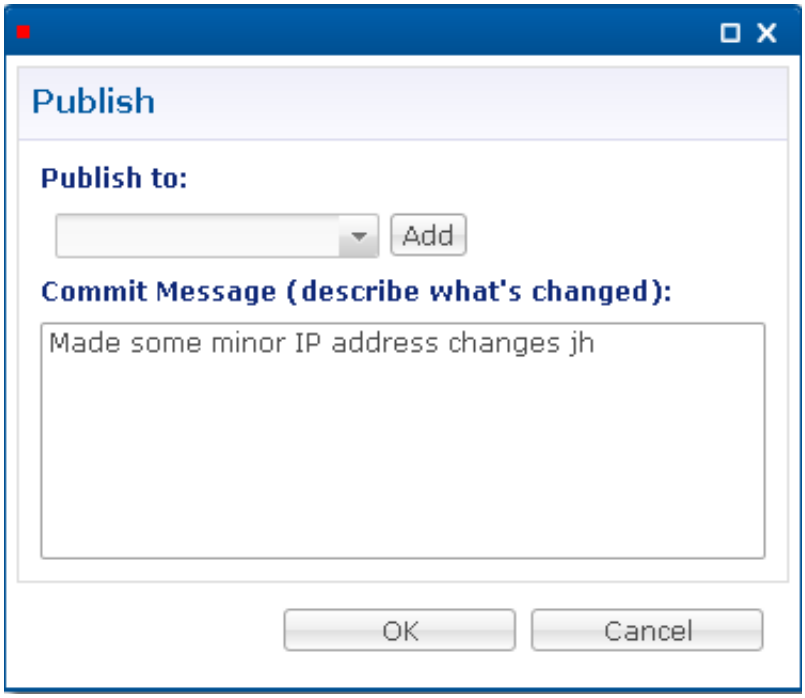

Step 3 If you have not set up the repository, this screen shows the defaults for setting one up. Accept or change, and click OK.

Note On this screen, you might want to change the Name field from the default.
and
You must select the Save Password check box before typing in the password.

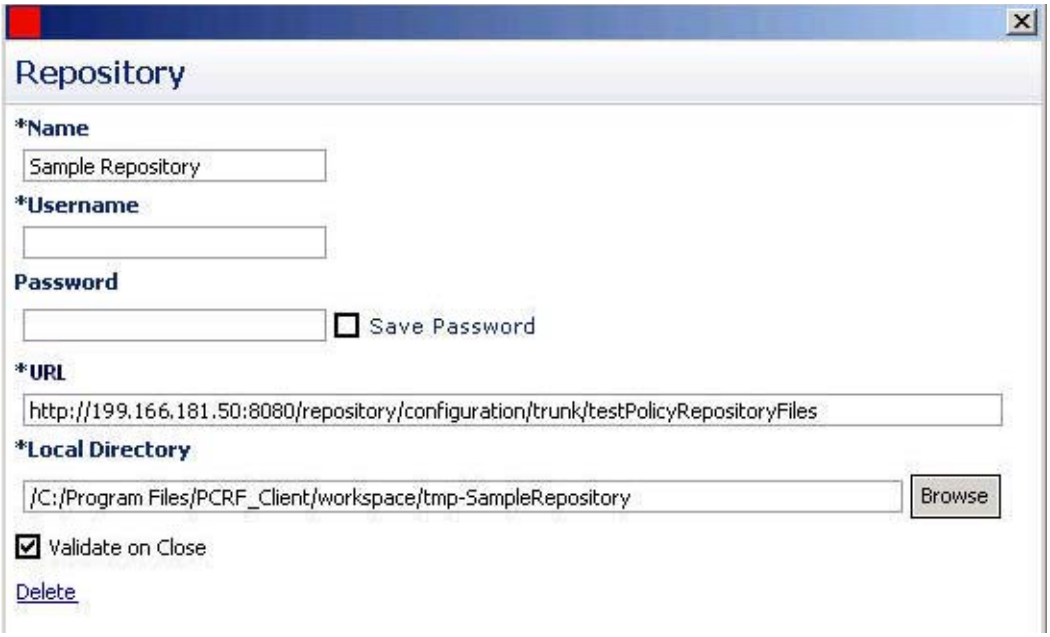

Note Detailed instructions for this screen are found at Creating the First Client Repository.

Step 4 Verify the changes to Production repository:

Remember:

- All changes are published to Subversion, so they are version-controlled and can be rolled back.
- To verify a publish as part of a trouble shooting 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 (broadhop/broadhop by default).
- If a traditional web browser cannot access the system, you can use a command line browser from the CPS VM's URL.

- The most important file to look at is config.properties. This has the date published as well as the location published from.

Rolling Back Changes to Production

All changes are made in Subversion. Use the SVN version control commands to roll back changes if you notice problems.

Step 1 Log in to any of the CPS virtual machines via SSH and take note of the URL you published to (usually `http://pcrfclient01/repos/run`).

Step 2 Look at the previous changes with this command:

```
more | svn log http://pcrfclient01/repos/run
```

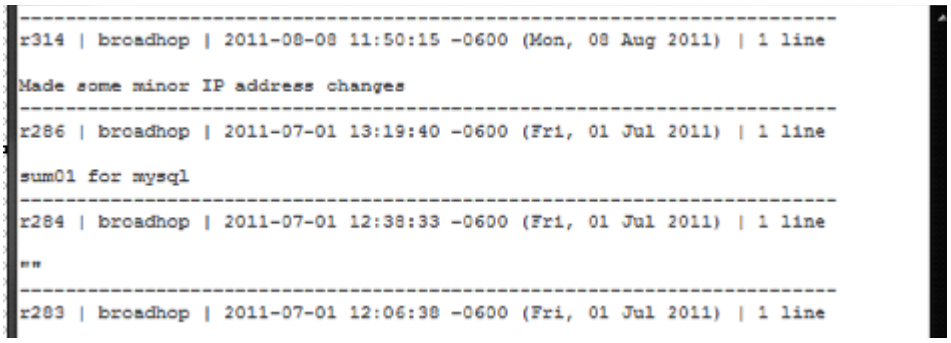

Step 3 Using the commit comments and/or time/date, find the version you want to roll back to. This example undo r314 (Made some minor IP address changes) and roll back to r286.

Step 4 To roll back, do a delete and restore.

```
[root@pcrfclient01 ~]# svn delete http://pcrfclient01/repos/run -m 'deleting for
 rollback'

Committed revision 315.
[root@pcrfclient01 ~]# svn cp http://pcrfclient01/repos/run@286 -m .rolling back
 to 286.
svn: Local, non-commit operations do not take a log message or revision properti
es
[root@pcrfclient01 ~]# svn cp http://pcrfclient01/repos/run@286 -m 'rolling back
 to 286'
svn: Try 'svn help' for more info
svn: Not enough arguments provided
[root@pcrfclient01 ~]# svn cp http://pcrfclient01/repos/run@286 http://pcrfclien
t01/repos/run -m 'rolling back to 286'

Committed revision 316.
```

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

Restore: svn cp http://pcrfclient01/repos/run@286 http://pcrfclient01/repos/run -m 'rolling back to 286'

```
[root@pcrfclient01 ~]# svn delete http://pcrfclient01/repos/run -m 'deleting for
 rollback'

Committed revision 315.
[root@pcrfclient01 ~]# svn cp http://pcrfclient01/repos/run@286 -m .rolling back
 to 286.
svn: Local, non-commit operations do not take a log message or revision properti
es
[root@pcrfclient01 ~]# svn cp http://pcrfclient01/repos/run@286 -m 'rolling back
 to 286'
svn: Try 'svn help' for more info
svn: Not enough arguments provided
[root@pcrfclient01 ~]# svn cp http://pcrfclient01/repos/run@286 http://pcrfclien
t01/repos/run -m 'rolling back to 286'
```

Step 5 Using the SVN log (history commands), we can see the rollback. Notice that the information about minor IP changes is gone.

```
------------------------------------------------------------------------
r316 | broadhcp | 2011-08-08 11:56:18 -0600 (Mon, 08 Aug 2011) | 1 line

rolling back to 286
------------------------------------------------------------------------
r286 | broadhcp | 2011-07-01 13:19:40 -0600 (Fri, 01 Jul 2011) | 1 line

sum01 for mysql
------------------------------------------------------------------------
r284 | broadhcp | 2011-07-01 12:38:33 -0600 (Fri, 01 Jul 2011) | 1 line

""
------------------------------------------------------------------------
r283 | broadhcp | 2011-07-01 12:06:38 -0600 (Fri, 01 Jul 2011) | 1 line

""
------------------------------------------------------------------------
r281 | broadhcp | 2011-07-01 11:02:13 -0600 (Fri, 01 Jul 2011) | 1 line

""
------------------------------------------------------------------------
r280 | broadhcp | 2011-07-01 10:02:09 -0600 (Fri, 01 Jul 2011) | 1 line

:
```

Saving Client Data to a Repository

To save your work into the client repository version control software, click

File > Save to Client Repository.

The best practice is to save to client repository at least once a day to ensure timely and sufficient data capture.

As you develop your Cisco Policy Builder configuration, save your work **locally** as you go with these methods:

- Switch to a new window.
- Use CTRL S.
- Click the disk icon at the top of the tree on the left.

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:

- Use the File menu.

- Use the Choose Policy Builder data repository.... screen.

  When presented with the Choose Policy Builder data repository.... screen, use the drop-down list to select another repository.

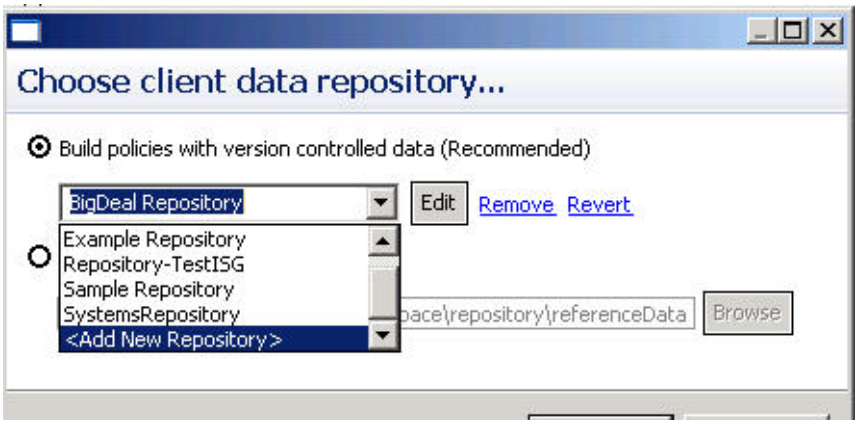

Reverting Repository

If you want to abandon changes to the client repository in the local directory, you can revert to a previous version. The version you revert to is available in the version control software. That version was saved the last time you used the option Save to Repository.

Step 1 Display the Choose Policy Builder data repository.... screen

Step 2 Click the Revert link to begin using the previous version of your work, the one that you saved with the Save to Repository menu item.

Note The Revert link appears only if there are uncommitted local changes.

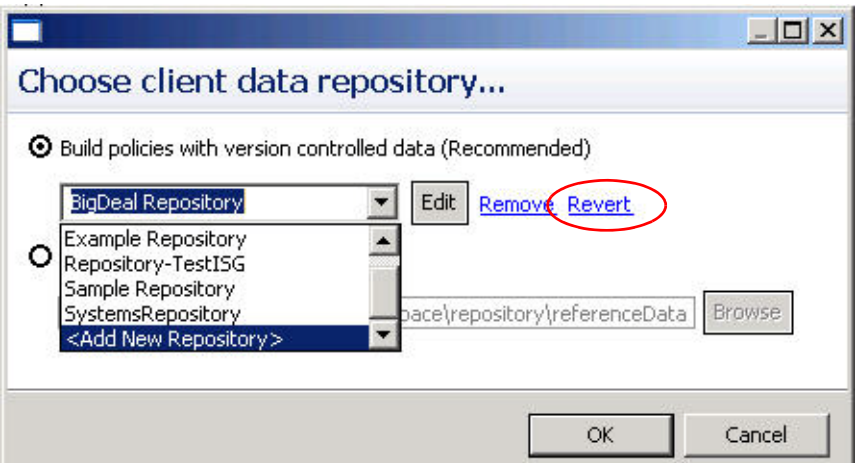

Call Flows

Revised: July 10, 2015

The following call flow diagrams are given to help you troubleshoot and understand CPS deployment.

One-click Call Flow

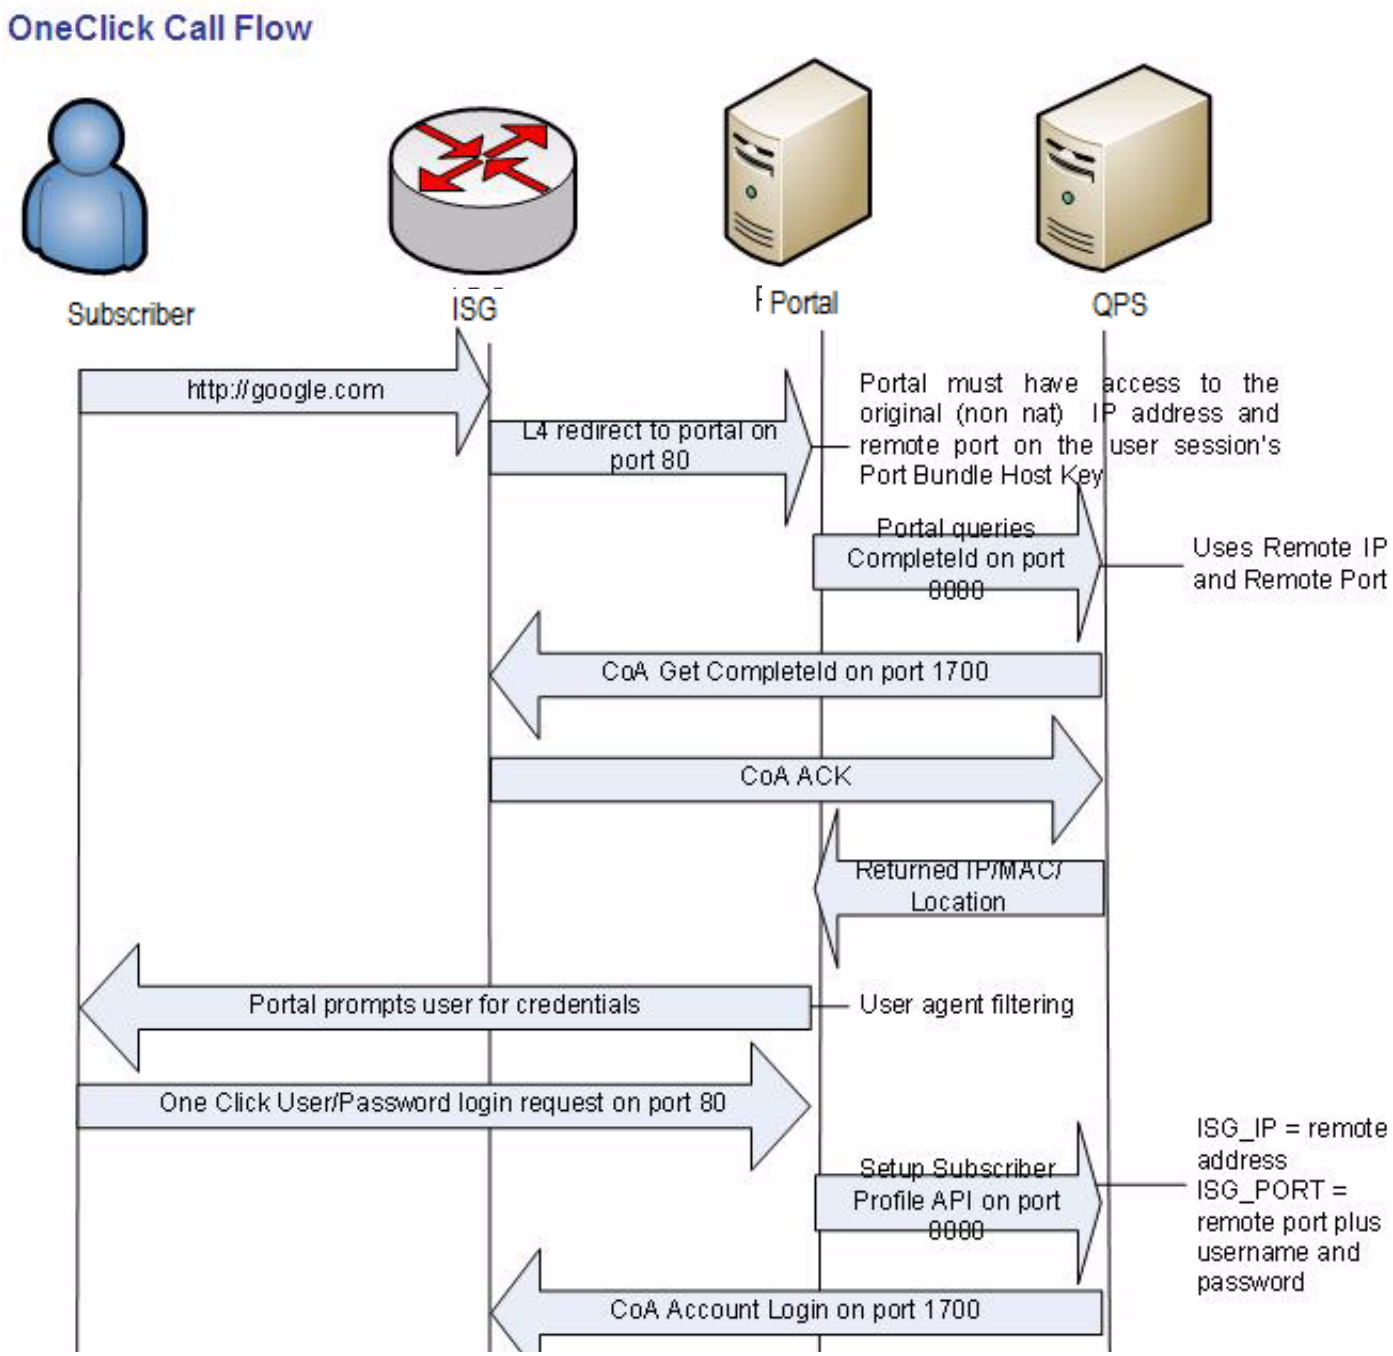

(continued)

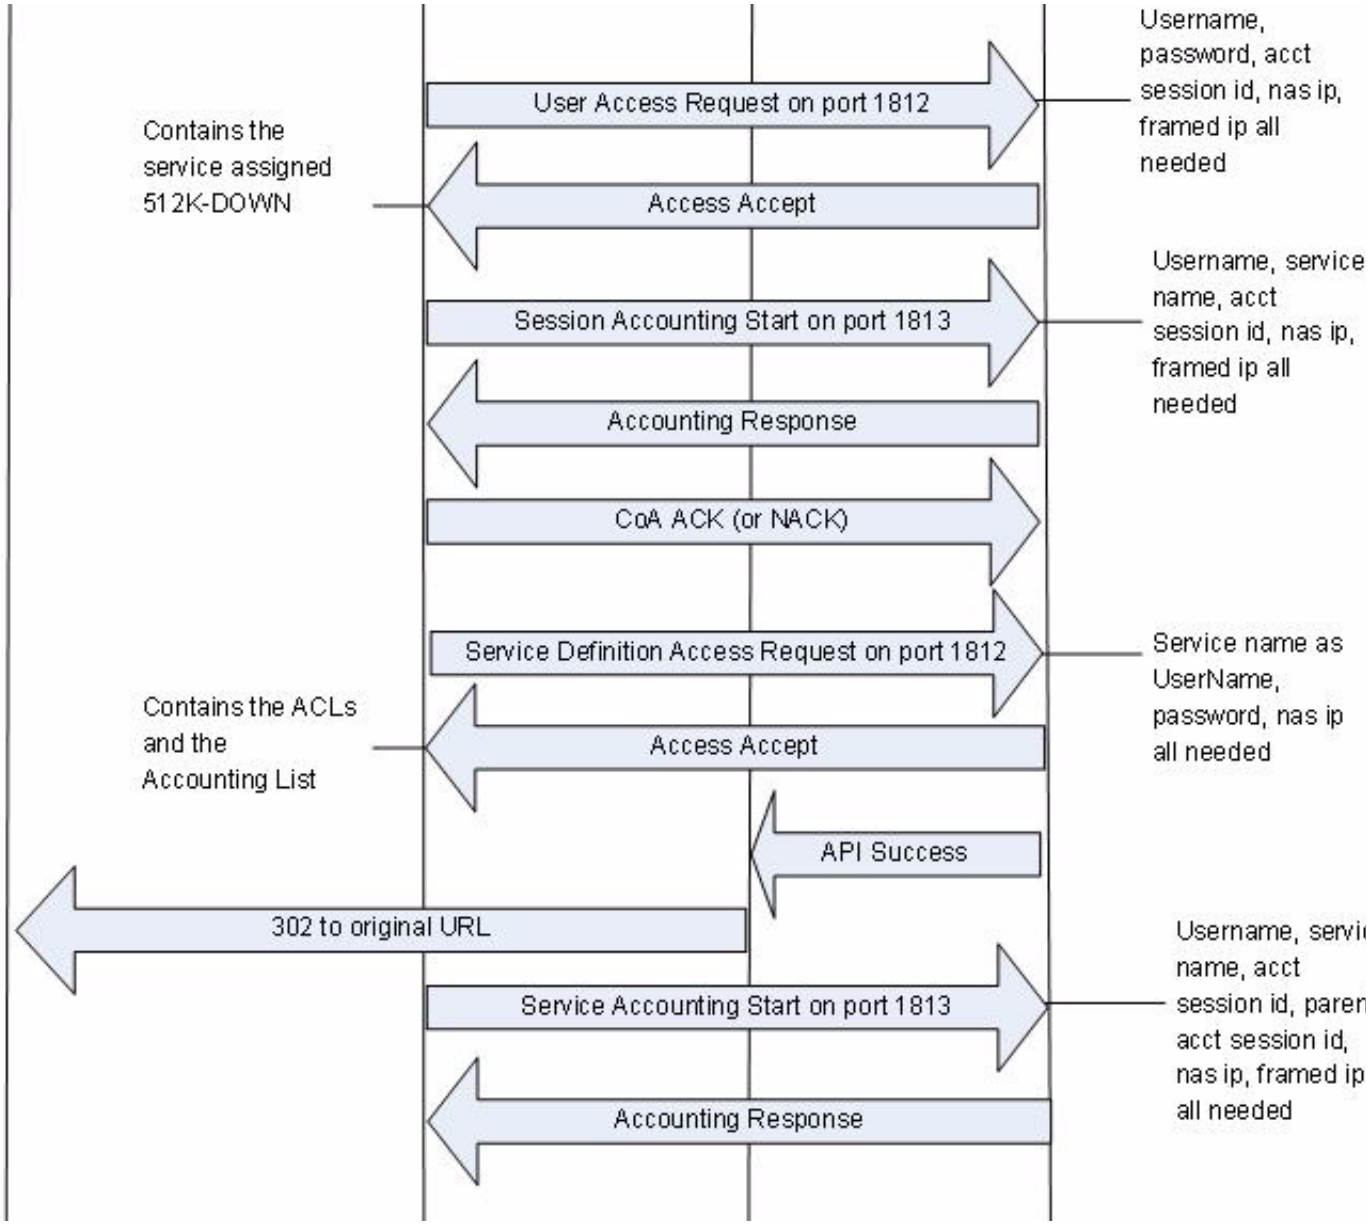
User Access Request on port 1812
Username, password, acct session id, nas ip, framed ip all needed
Contains the service assigned 512K-DOWN
Access Accept
Session Accounting Start on port 1813
Username, service name, acct session id, nas ip, framed ip all needed
Accounting Response
CoA ACK (or NACK)
Service Definition Access Request on port 1812
Service name as UserName, password, nas ip all needed
Contains the ACLs and the Accounting List
Access Accept
API Success
302 to original URL
Service Accounting Start on port 1813
Accounting Response

User/Password Login Call Flow

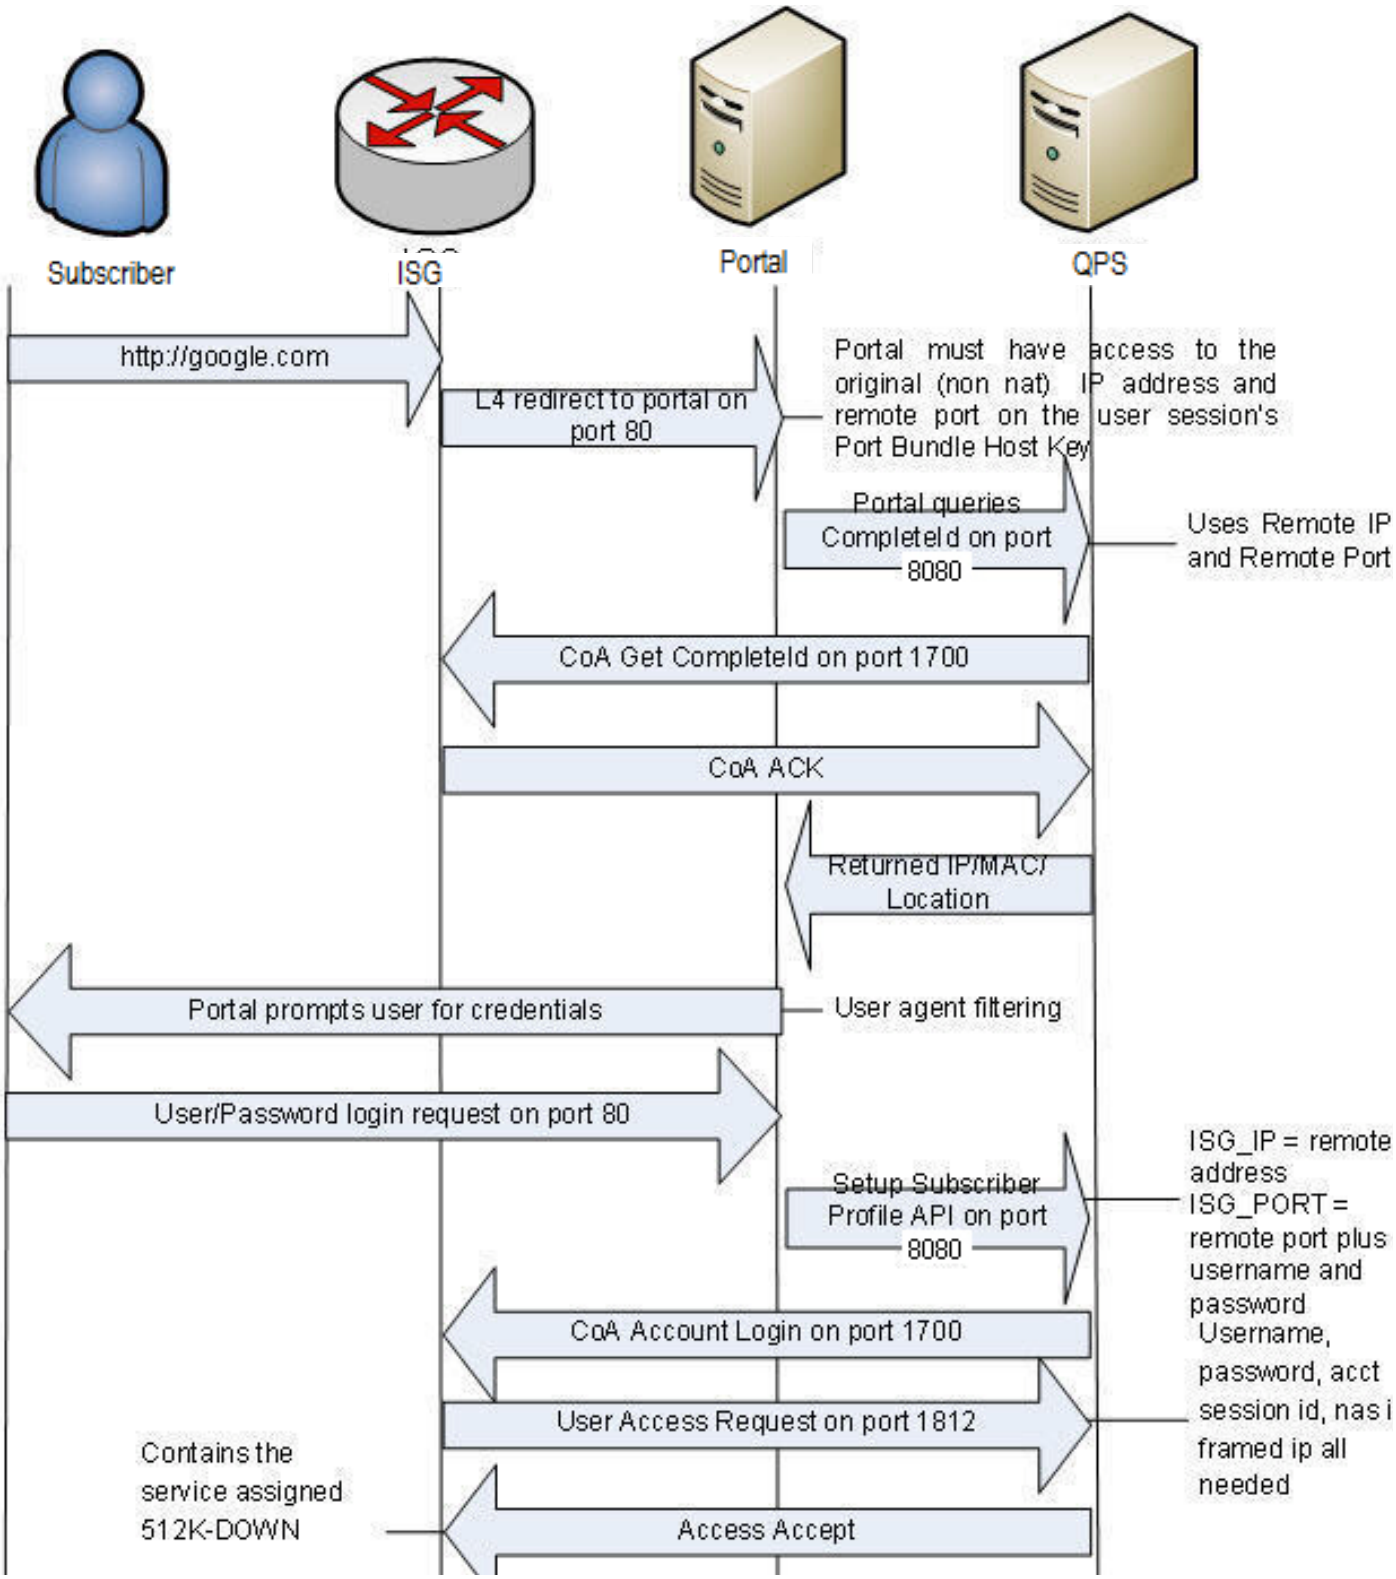

(continued)

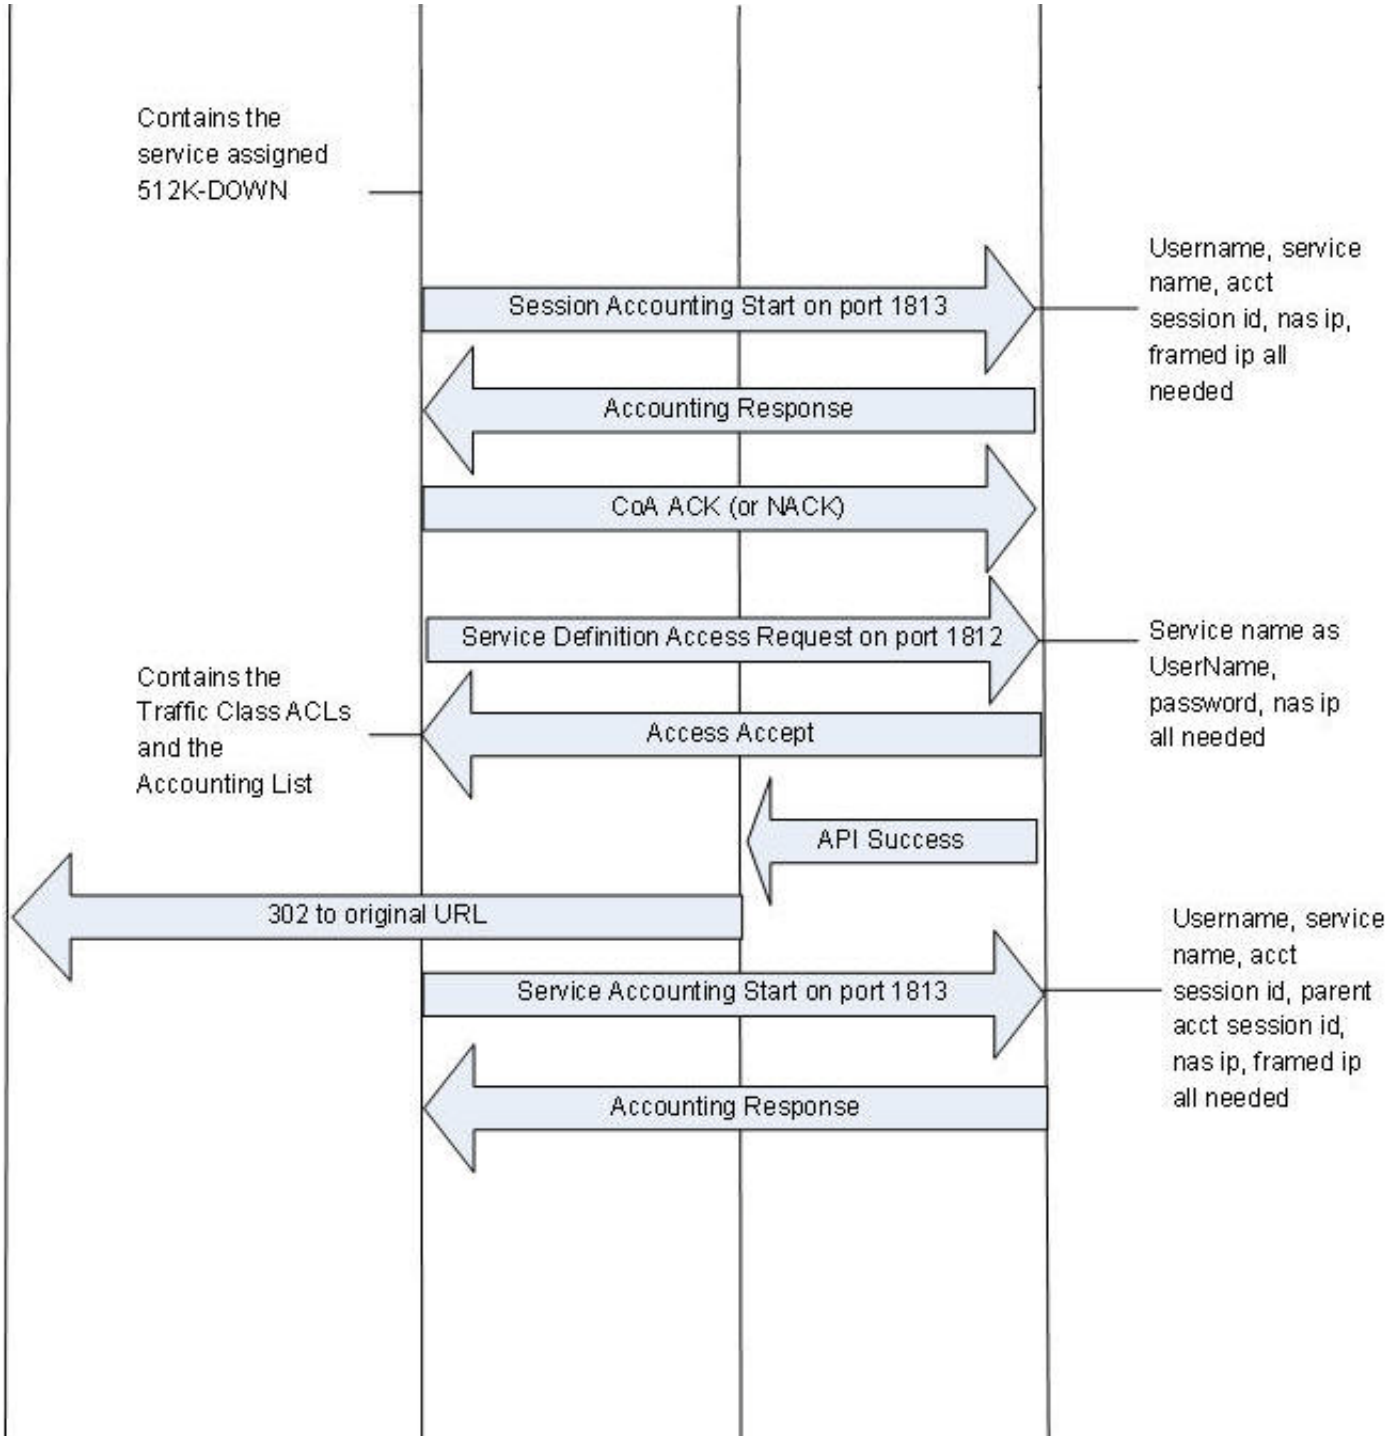
Contains the service assigned 512K-DOWN
Session Accounting Start on port 1813
Username, service name, acct session id, nas ip, framed ip all needed
Accounting Response
CoA ACK (or NACK)
Service Definition Access Request on port 1812
Service name as UserName, password, nas ip all needed
Contains the Traffic Class ACLs and the Accounting List
Access Accept
API Success
302 to original URL
Service Accounting Start on port 1813
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Accounting Response

Data-limited Voucher Call Flow

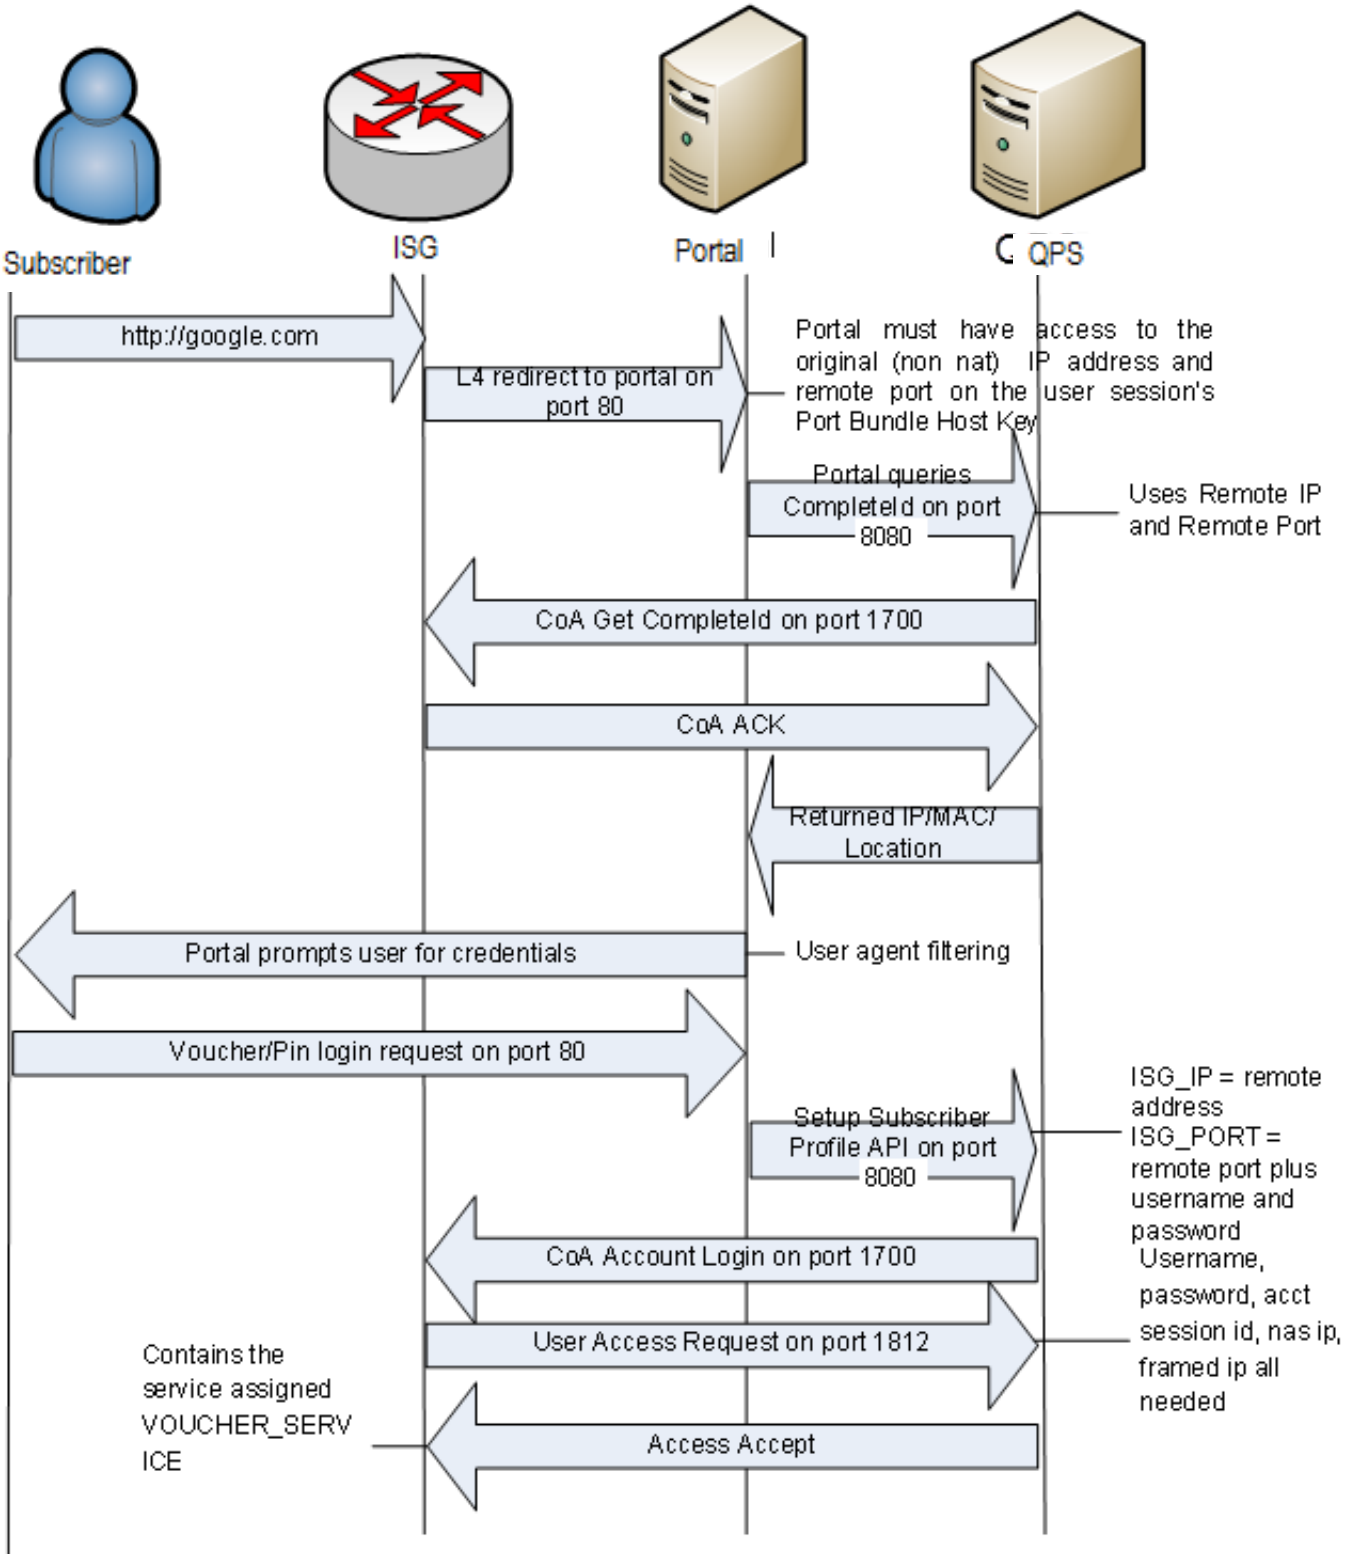

(continued)

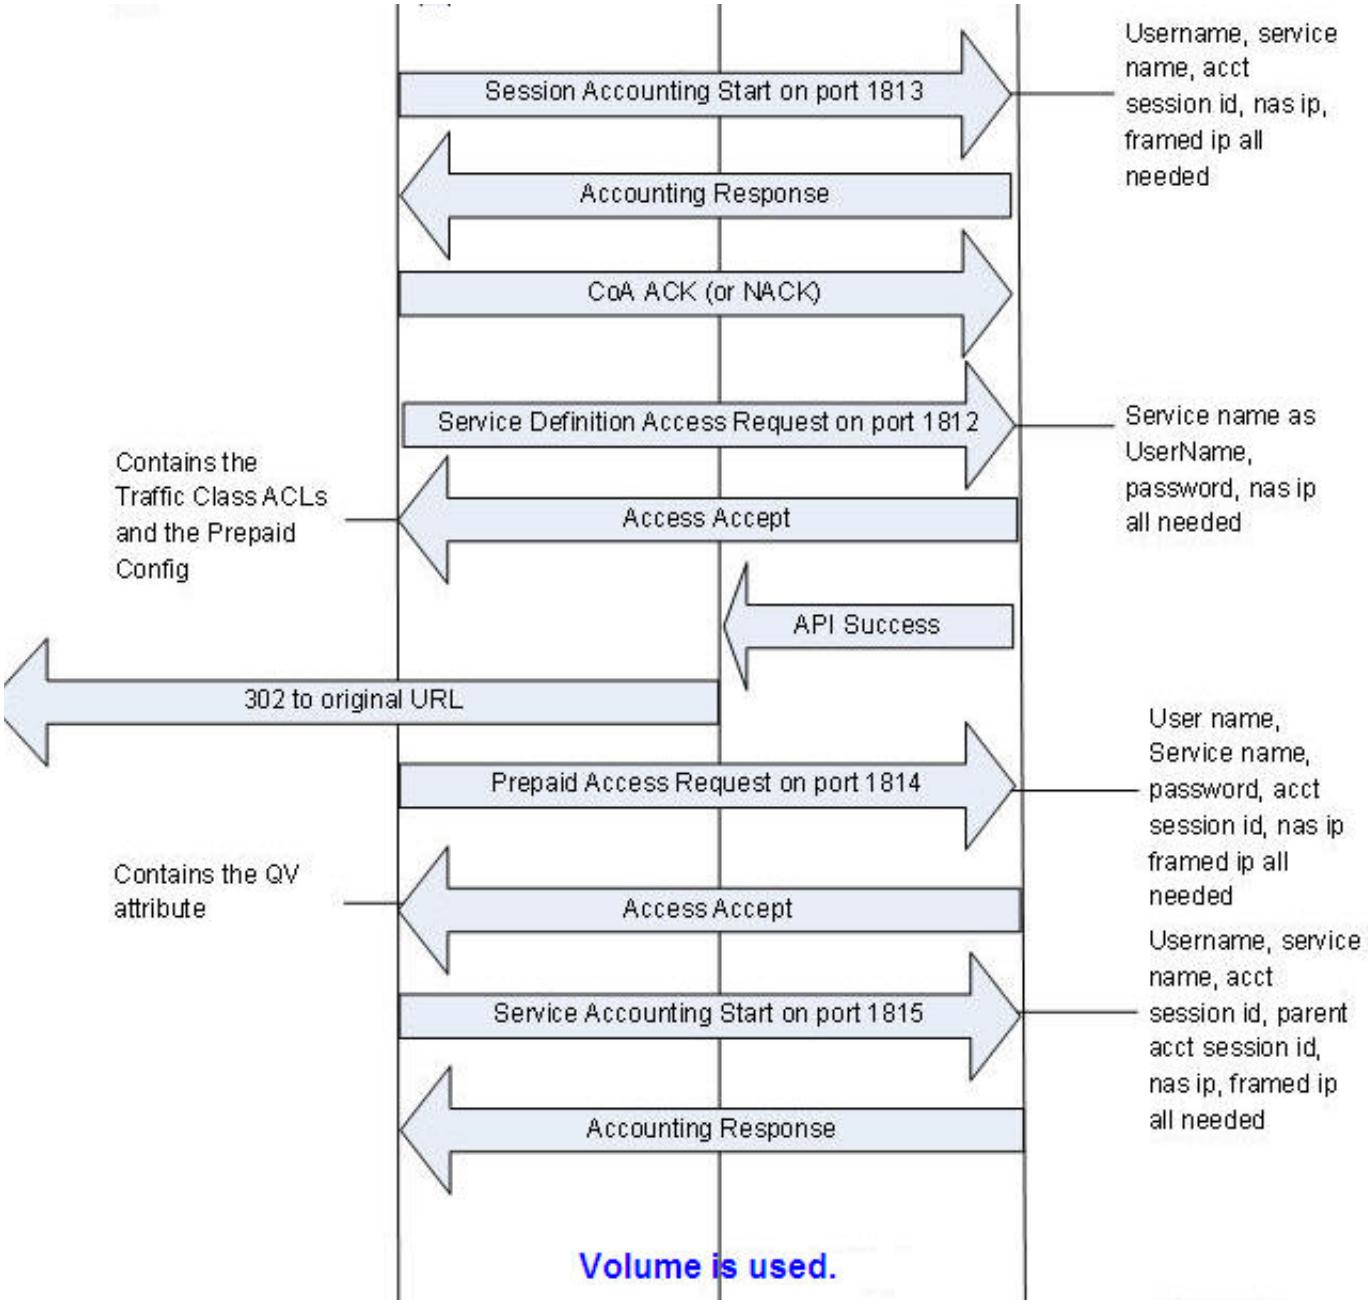
Session Accounting Start on port 1813
Username, service name, acct session id, nas ip, framed ip all needed
Accounting Response
CoA ACK (or NACK)
Service Definition Access Request on port 1812
Service name as UserName, password, nas ip all needed
Contains the Traffic Class ACLs and the Prepaid Config
Access Accept
API Success
302 to original URL
Prepaid Access Request on port 1814
User name, Service name, password, acct session id, nas ip framed ip all needed
Contains the QV attribute
Access Accept
Service Accounting Start on port 1815
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Accounting Response
Volume is used.

(continued)

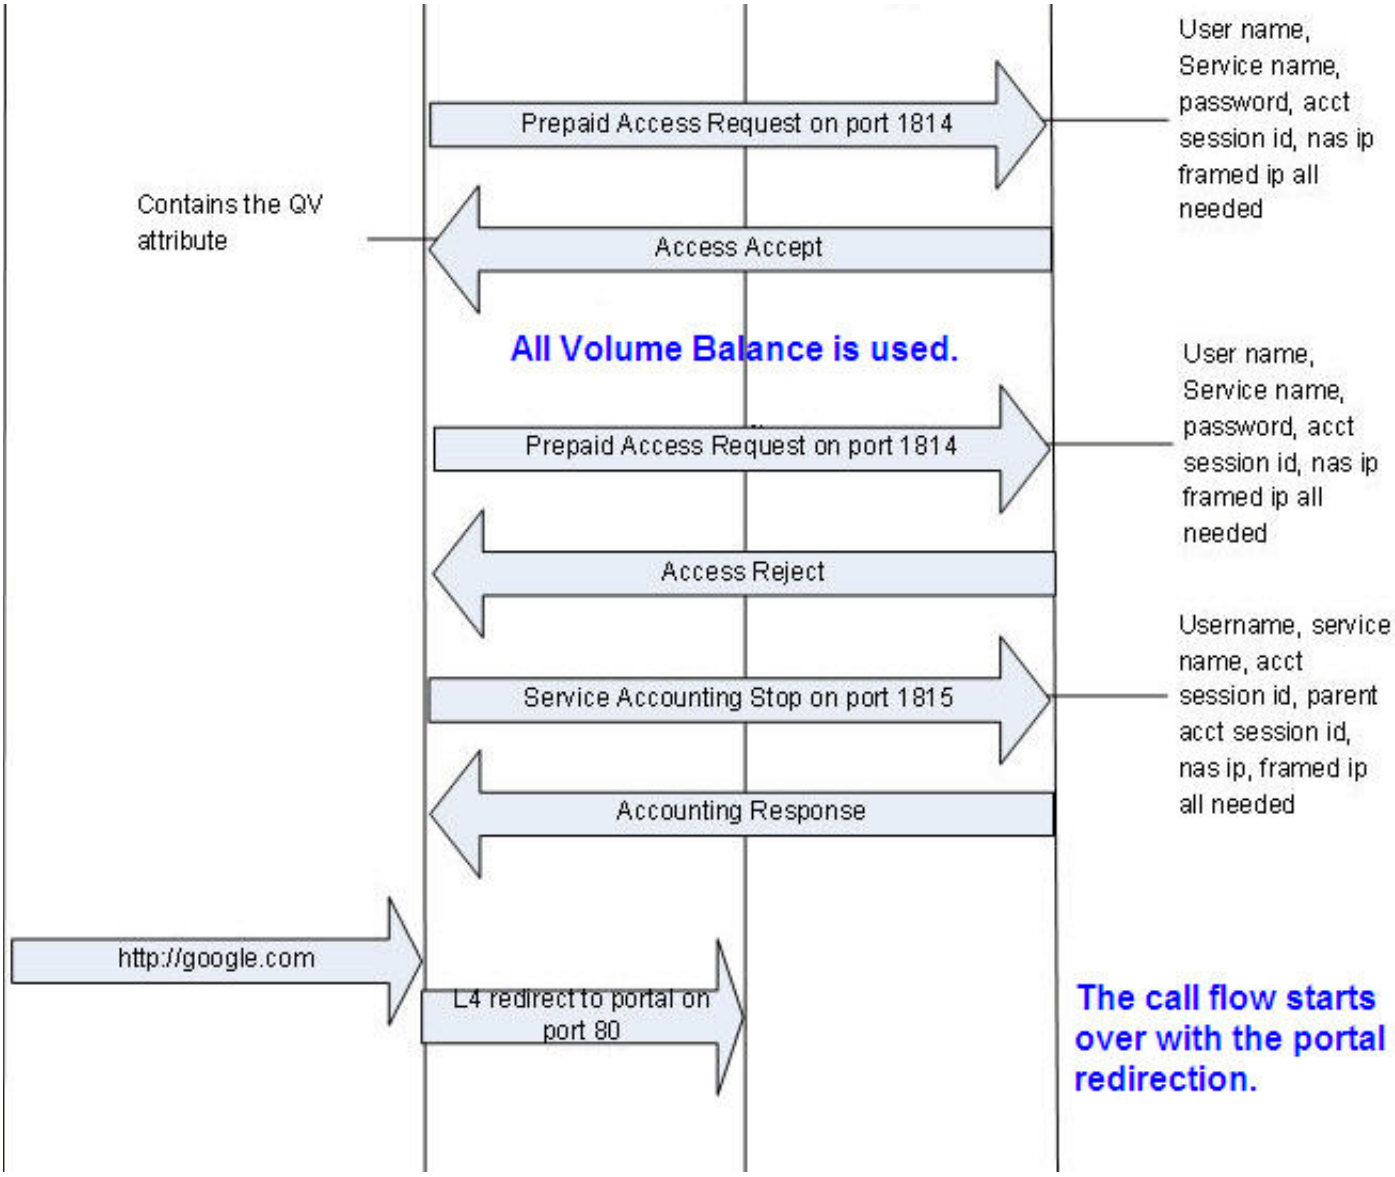
Prepaid Access Request on port 1814
User name, Service name, password, acct session id, nas ip framed ip all needed
Contains the QV attribute
Access Accept
All Volume Balance is used.
Prepaid Access Request on port 1814
User name, Service name, password, acct session id, nas ip framed ip all needed
Access Reject
Service Accounting Stop on port 1815
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Accounting Response
http://google.com
L4 redirect to portal on port 80
The call flow starts over with the portal redirection.

Time-limited Voucher Call Flow

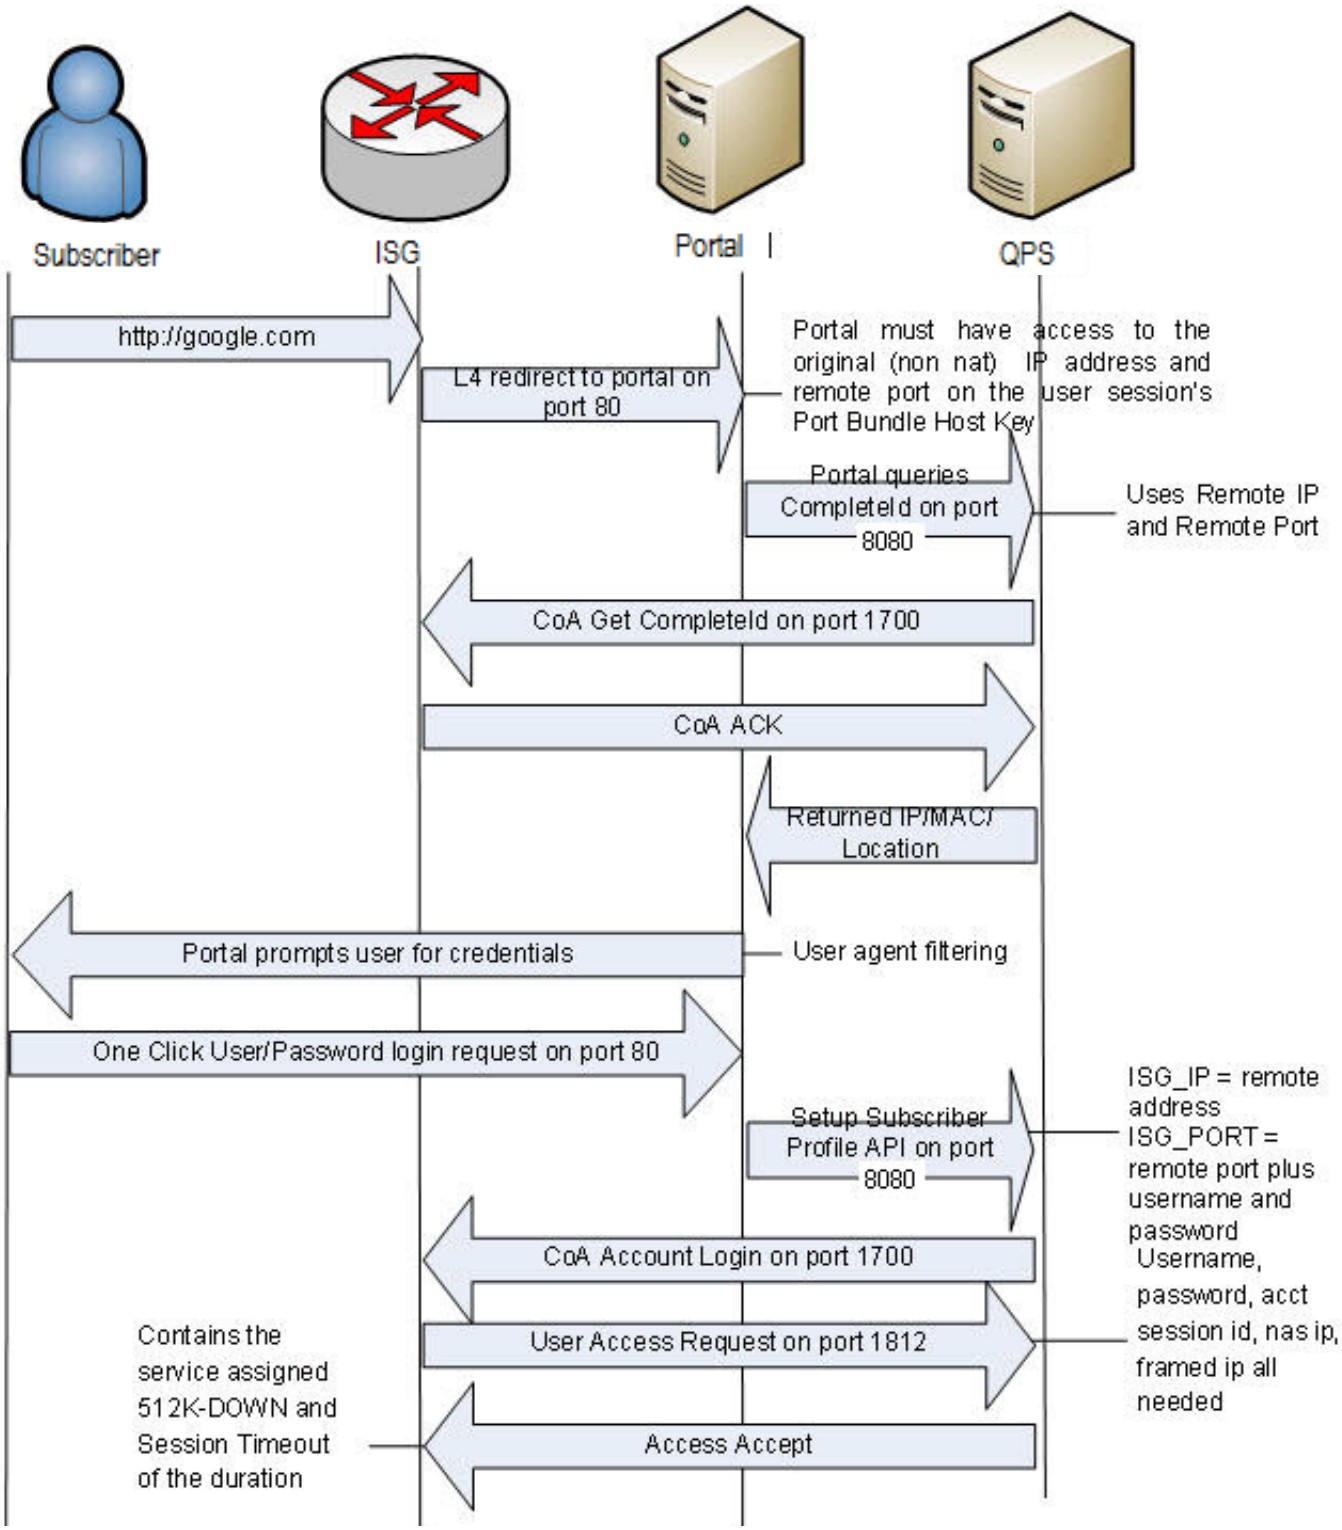

(continued)

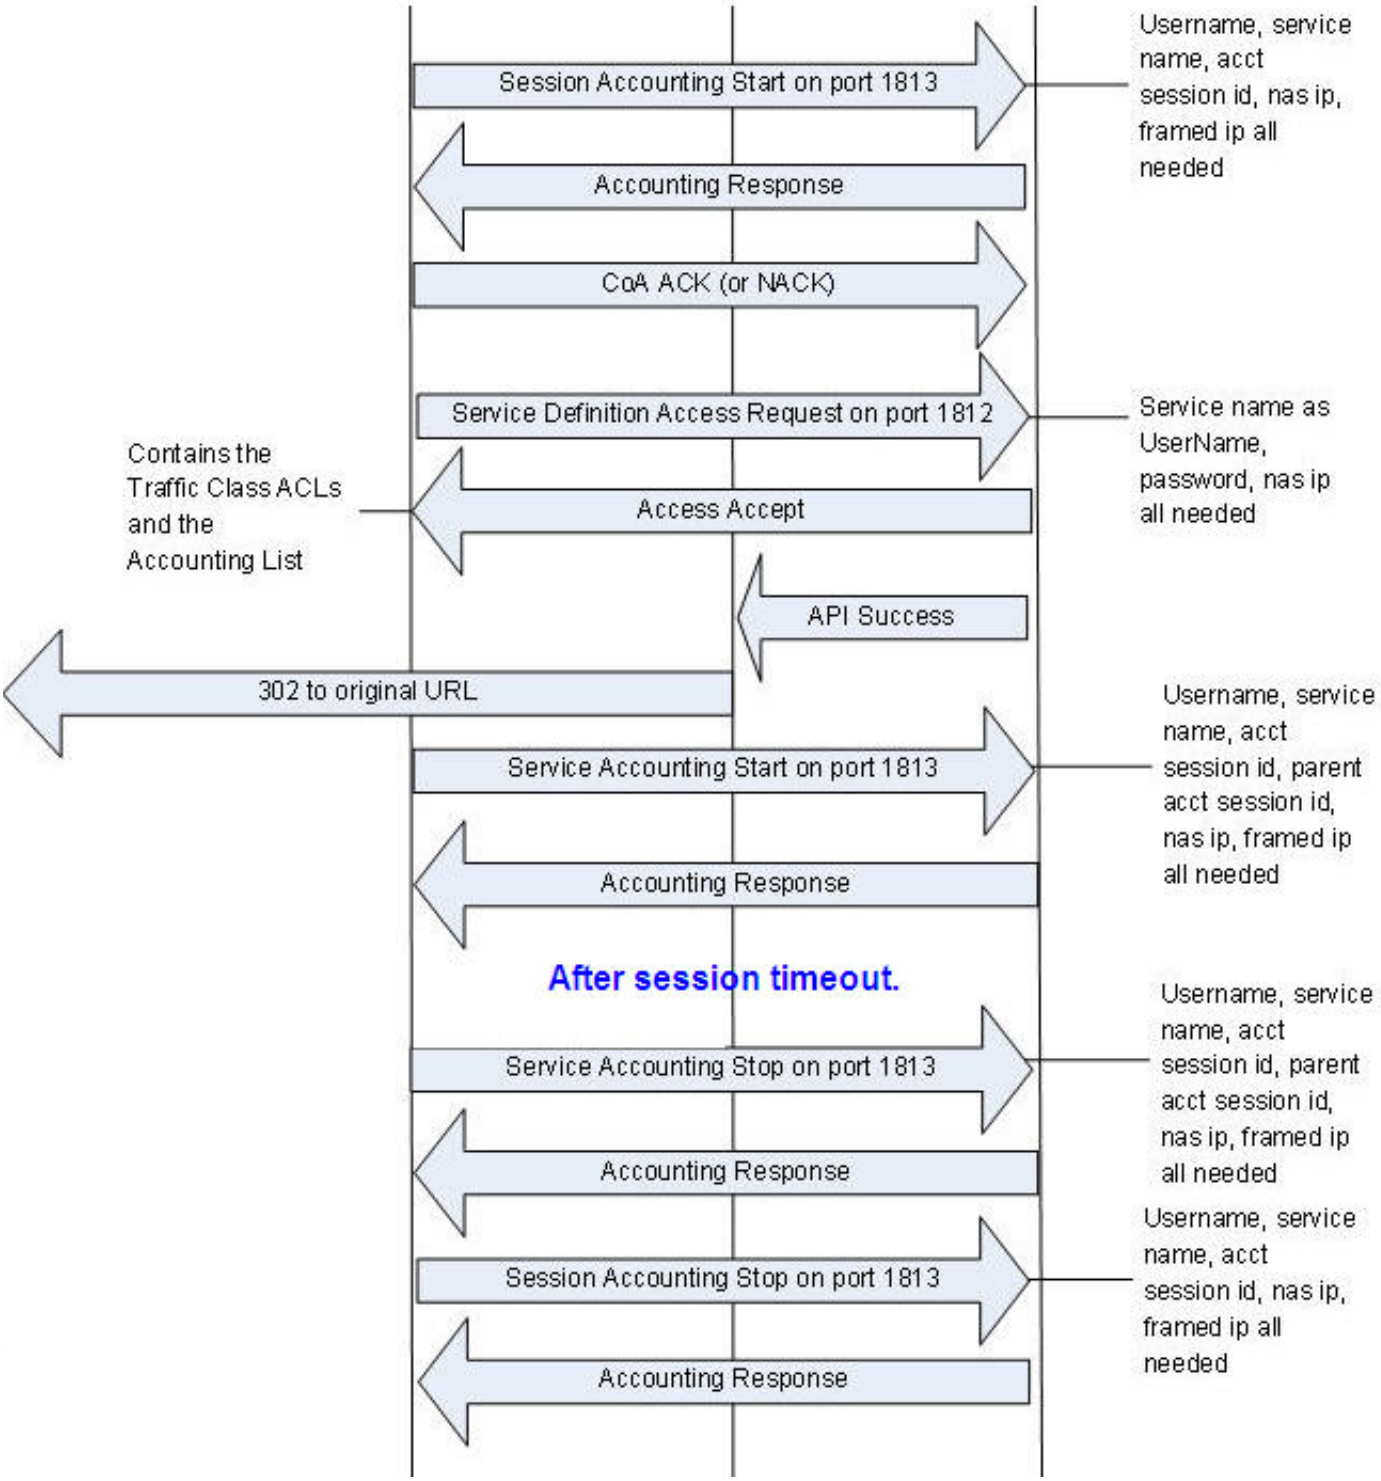
Session Accounting Start on port 1813
Username, service name, acct session id, nas ip, framed ip all needed
Accounting Response
CoA ACK (or NACK)
Service Definition Access Request on port 1812
Service name as UserName, password, nas ip all needed
Contains the Traffic Class ACLs and the Accounting List
Access Accept
API Success
302 to original URL
Service Accounting Start on port 1813
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Accounting Response
After session timeout.
Service Accounting Stop on port 1813
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Accounting Response
Session Accounting Stop on port 1813
Username, service name, acct session id, nas ip, framed ip all needed
Accounting Response

WISPr Call Flow

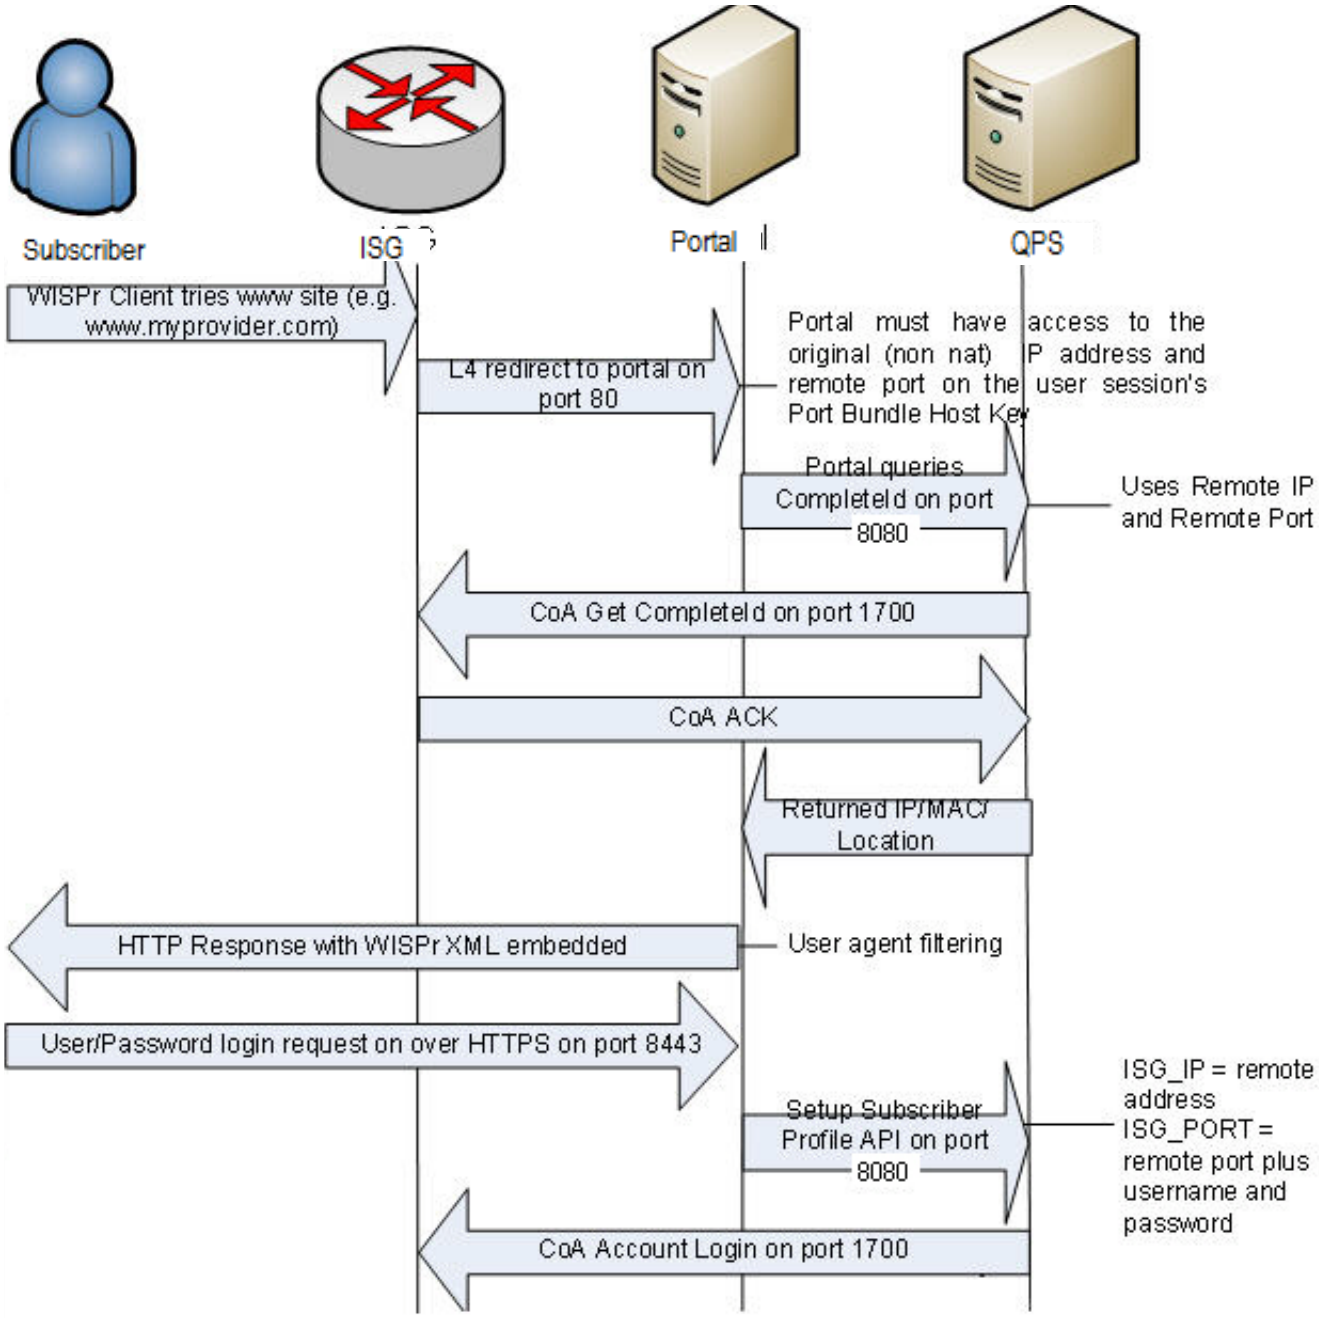

(continued)

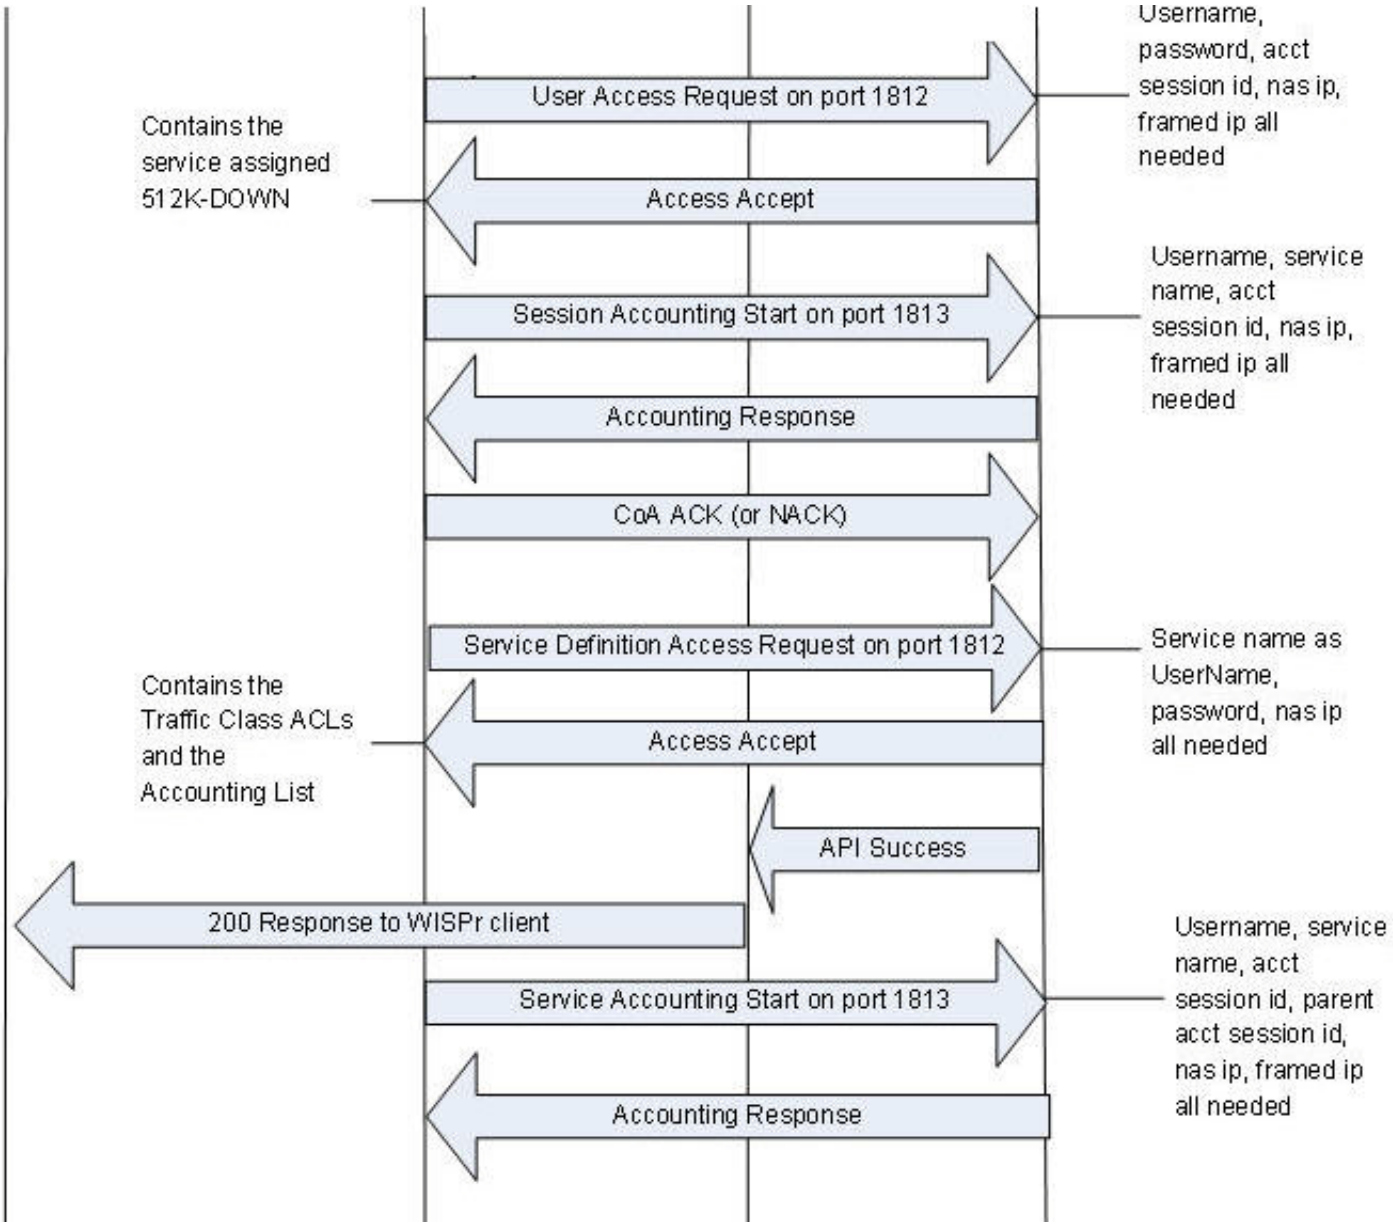
Username, password, acct session id, nas ip, framed ip all needed
User Access Request on port 1812
Contains the service assigned 512K-DOWN
Access Accept
Username, service name, acct session id, nas ip, framed ip all needed
Session Accounting Start on port 1813
Accounting Response
CoA ACK (or NACK)
Service Definition Access Request on port 1812
Service name as UserName, password, nas ip all needed
Contains the Traffic Class ACLs and the Accounting List
Access Accept
API Success
200 Response to WISPr client
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Service Accounting Start on port 1813
Accounting Response

EAP-TTLS Call Flow

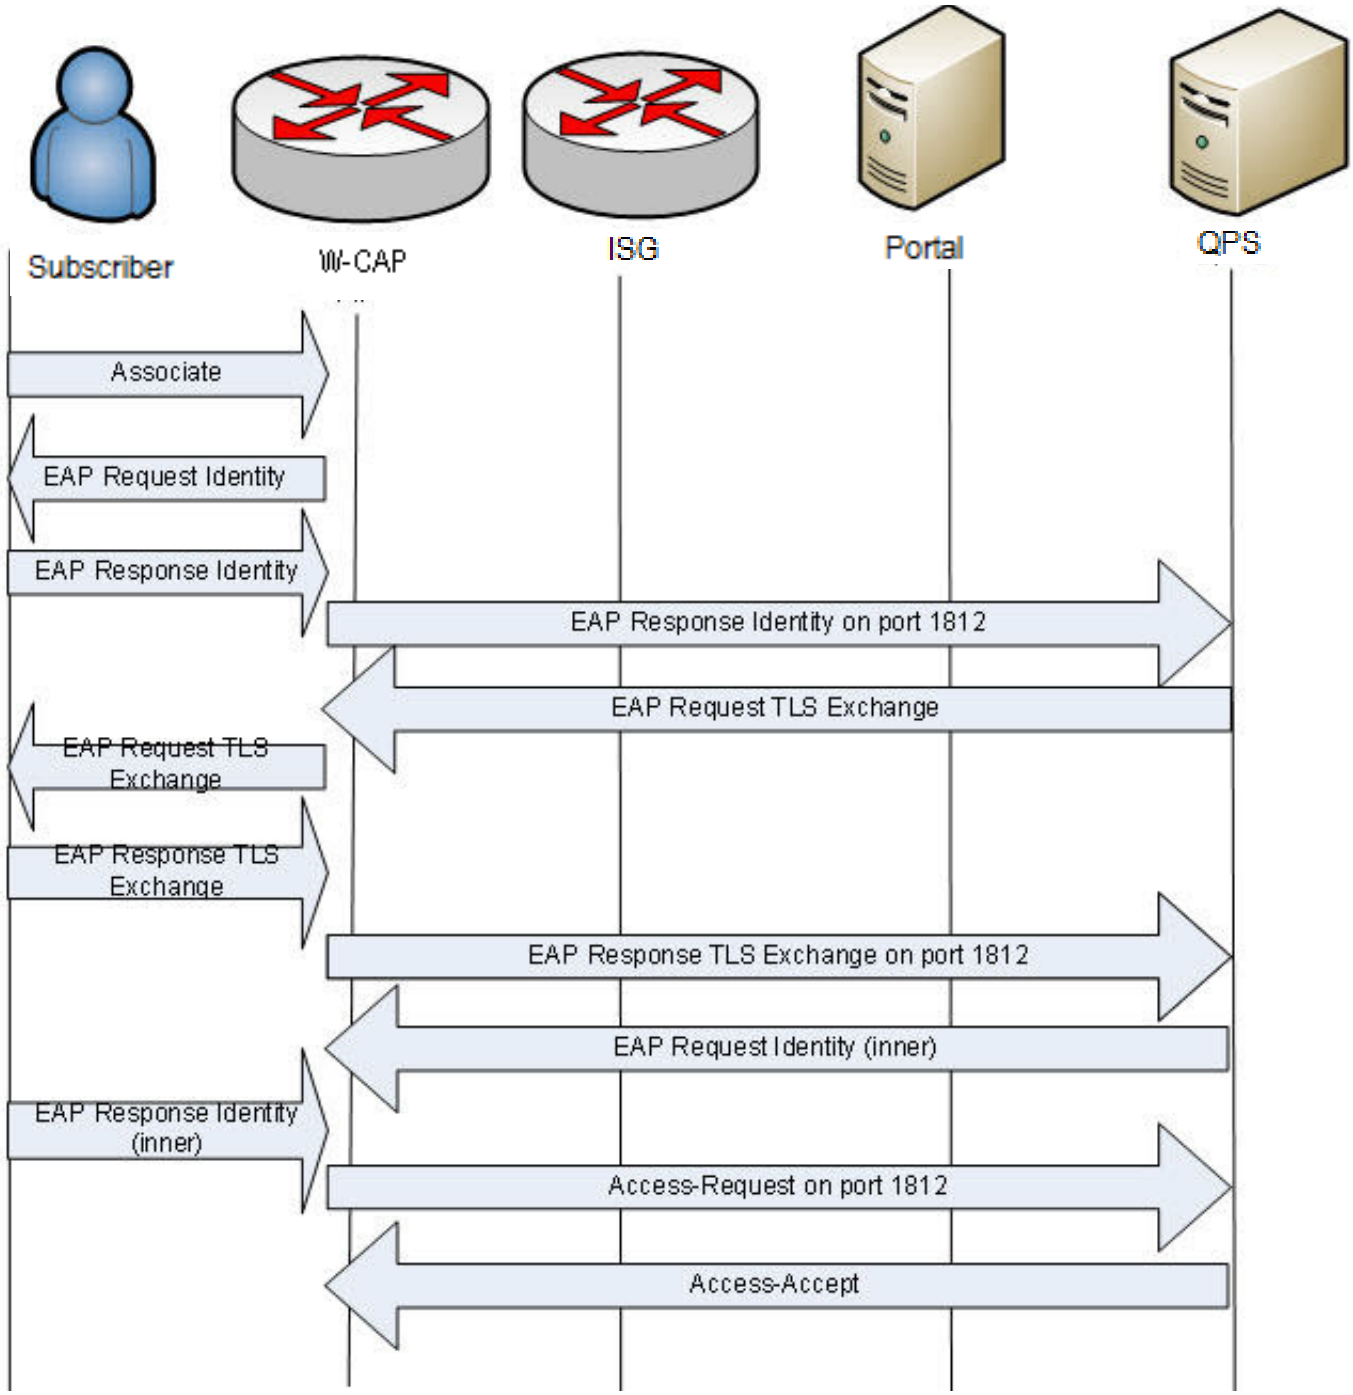

(continued)

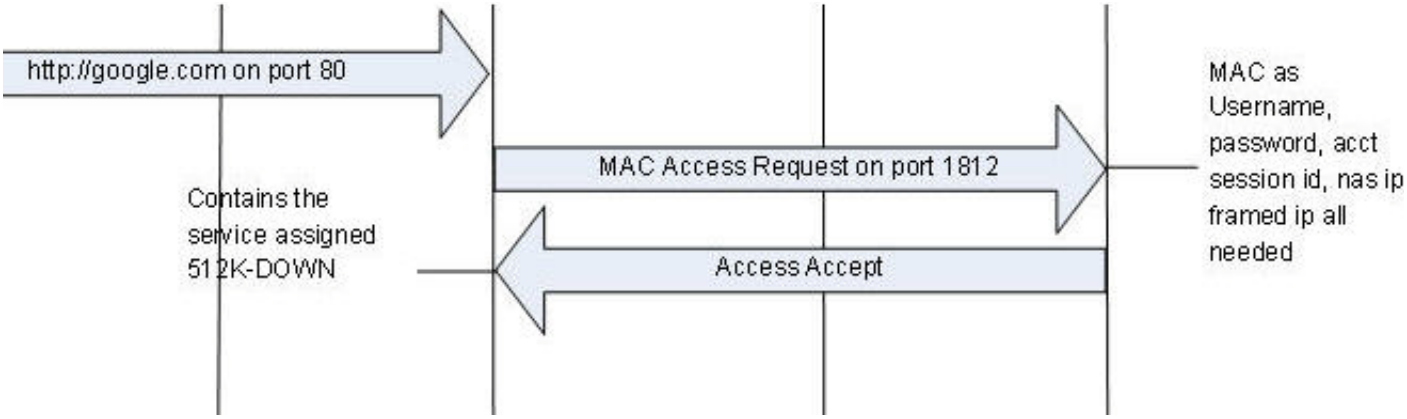
http://google.com on port 80
MAC Access Request on port 1812
MAC as Username, password, acct session id, nas ip framed ip all needed
Contains the service assigned 512K-DOWN
Access Accept

Service Selection Call Flow

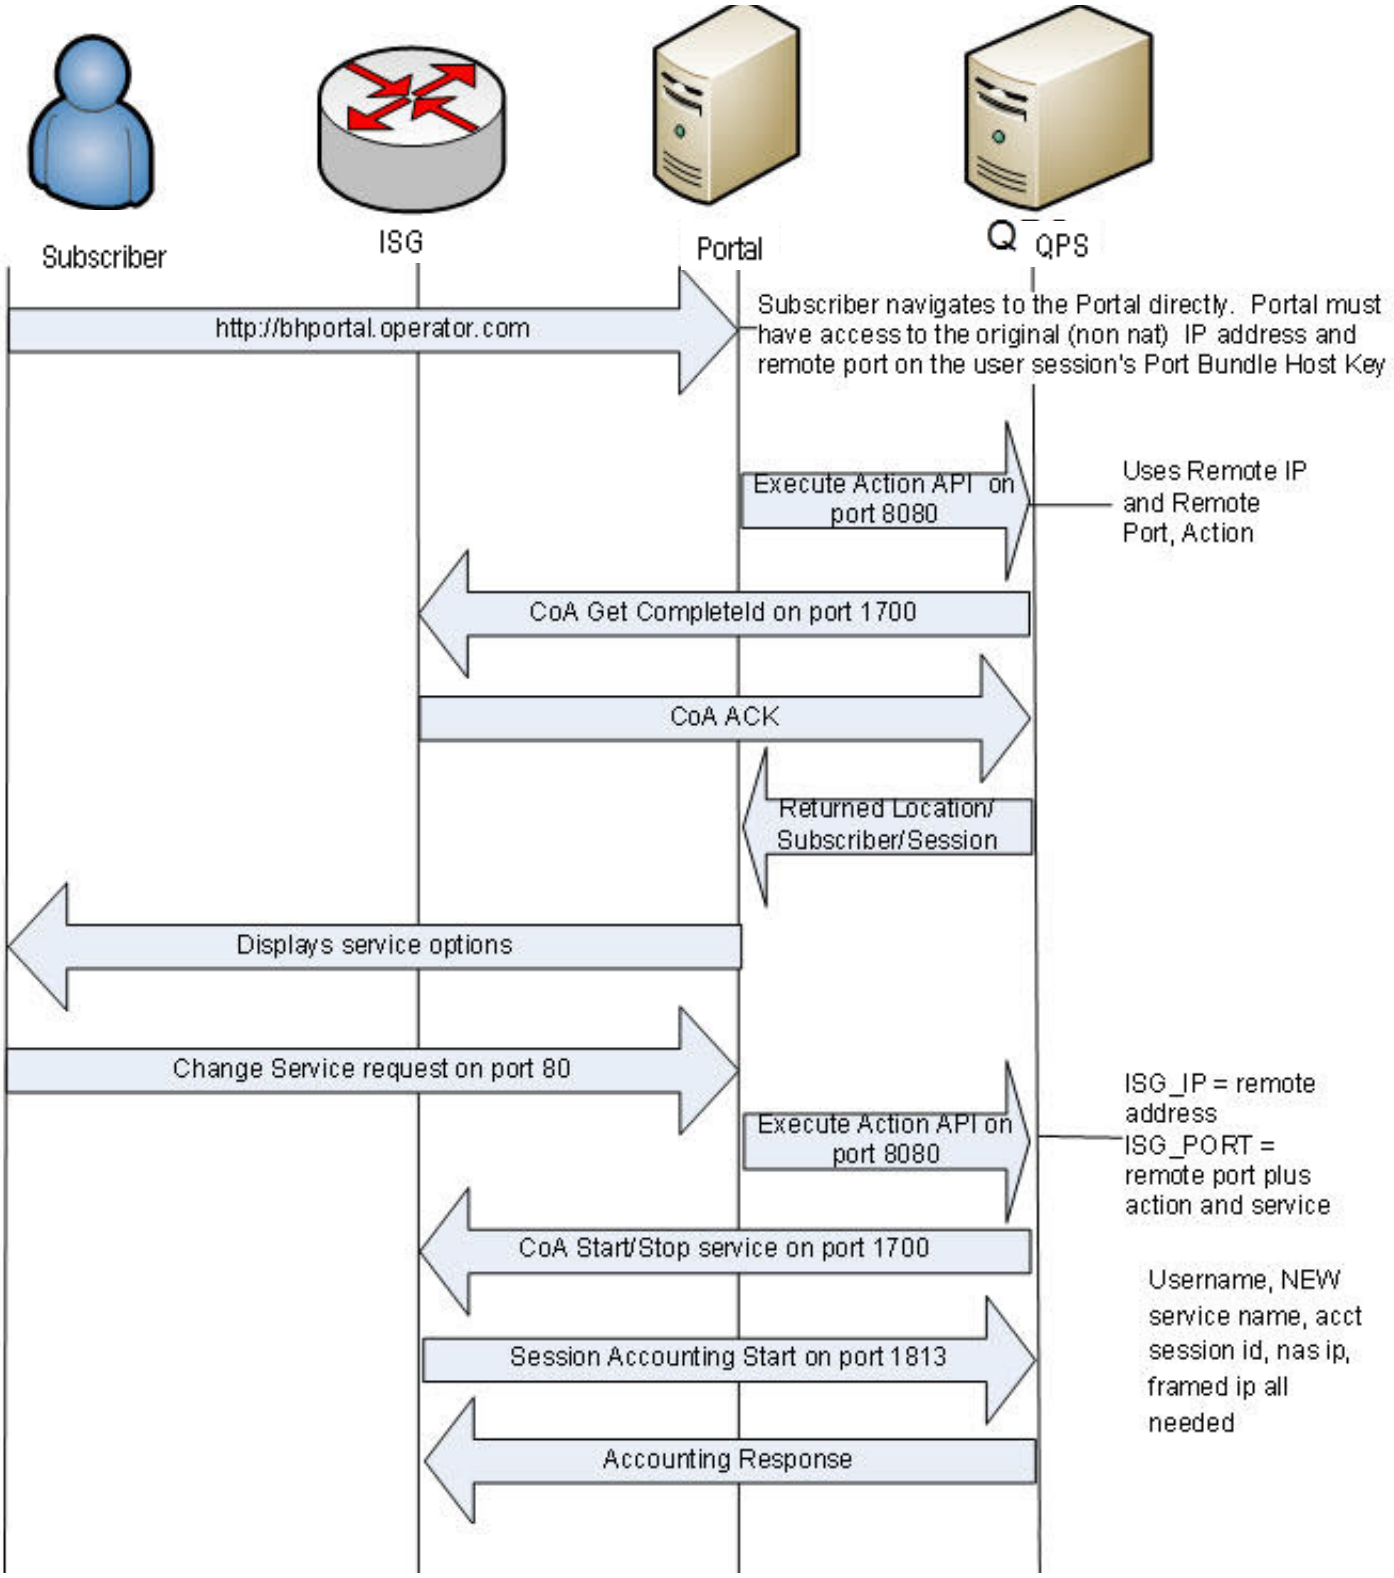

(continued)

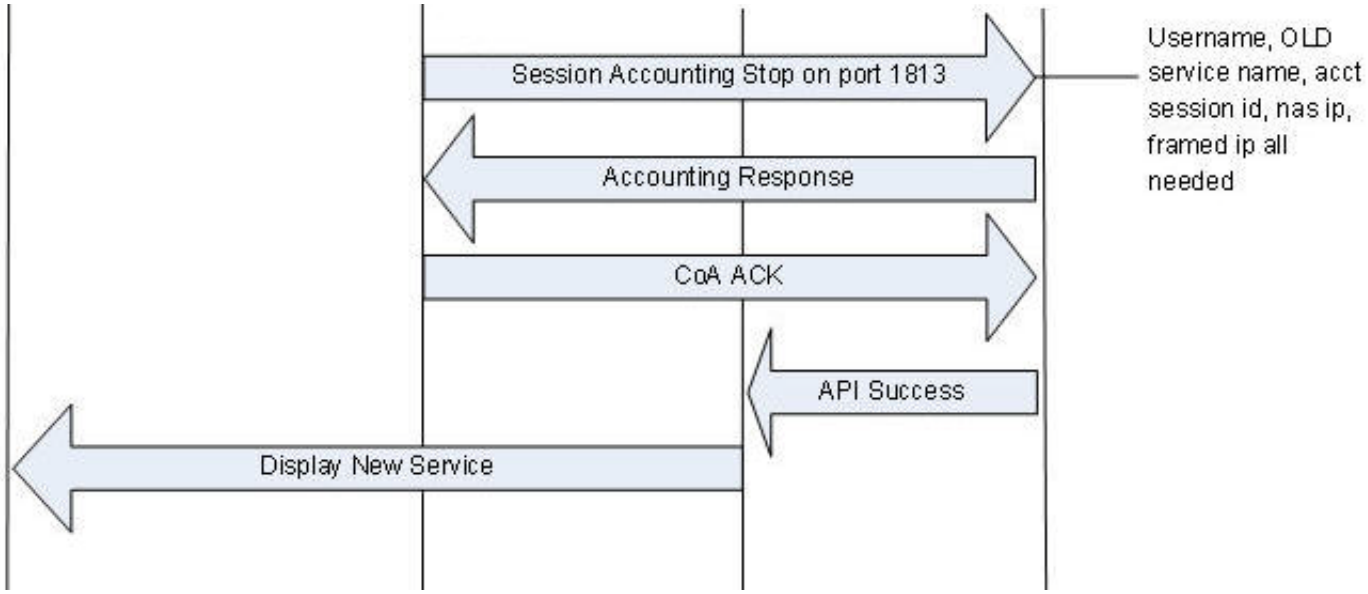
Session Accounting Stop on port 1813
Accounting Response
CoA ACK
API Success
Display New Service
Username, OLD service name, acct session id, nas ip, framed ip all needed

MAC TAL Call Flow

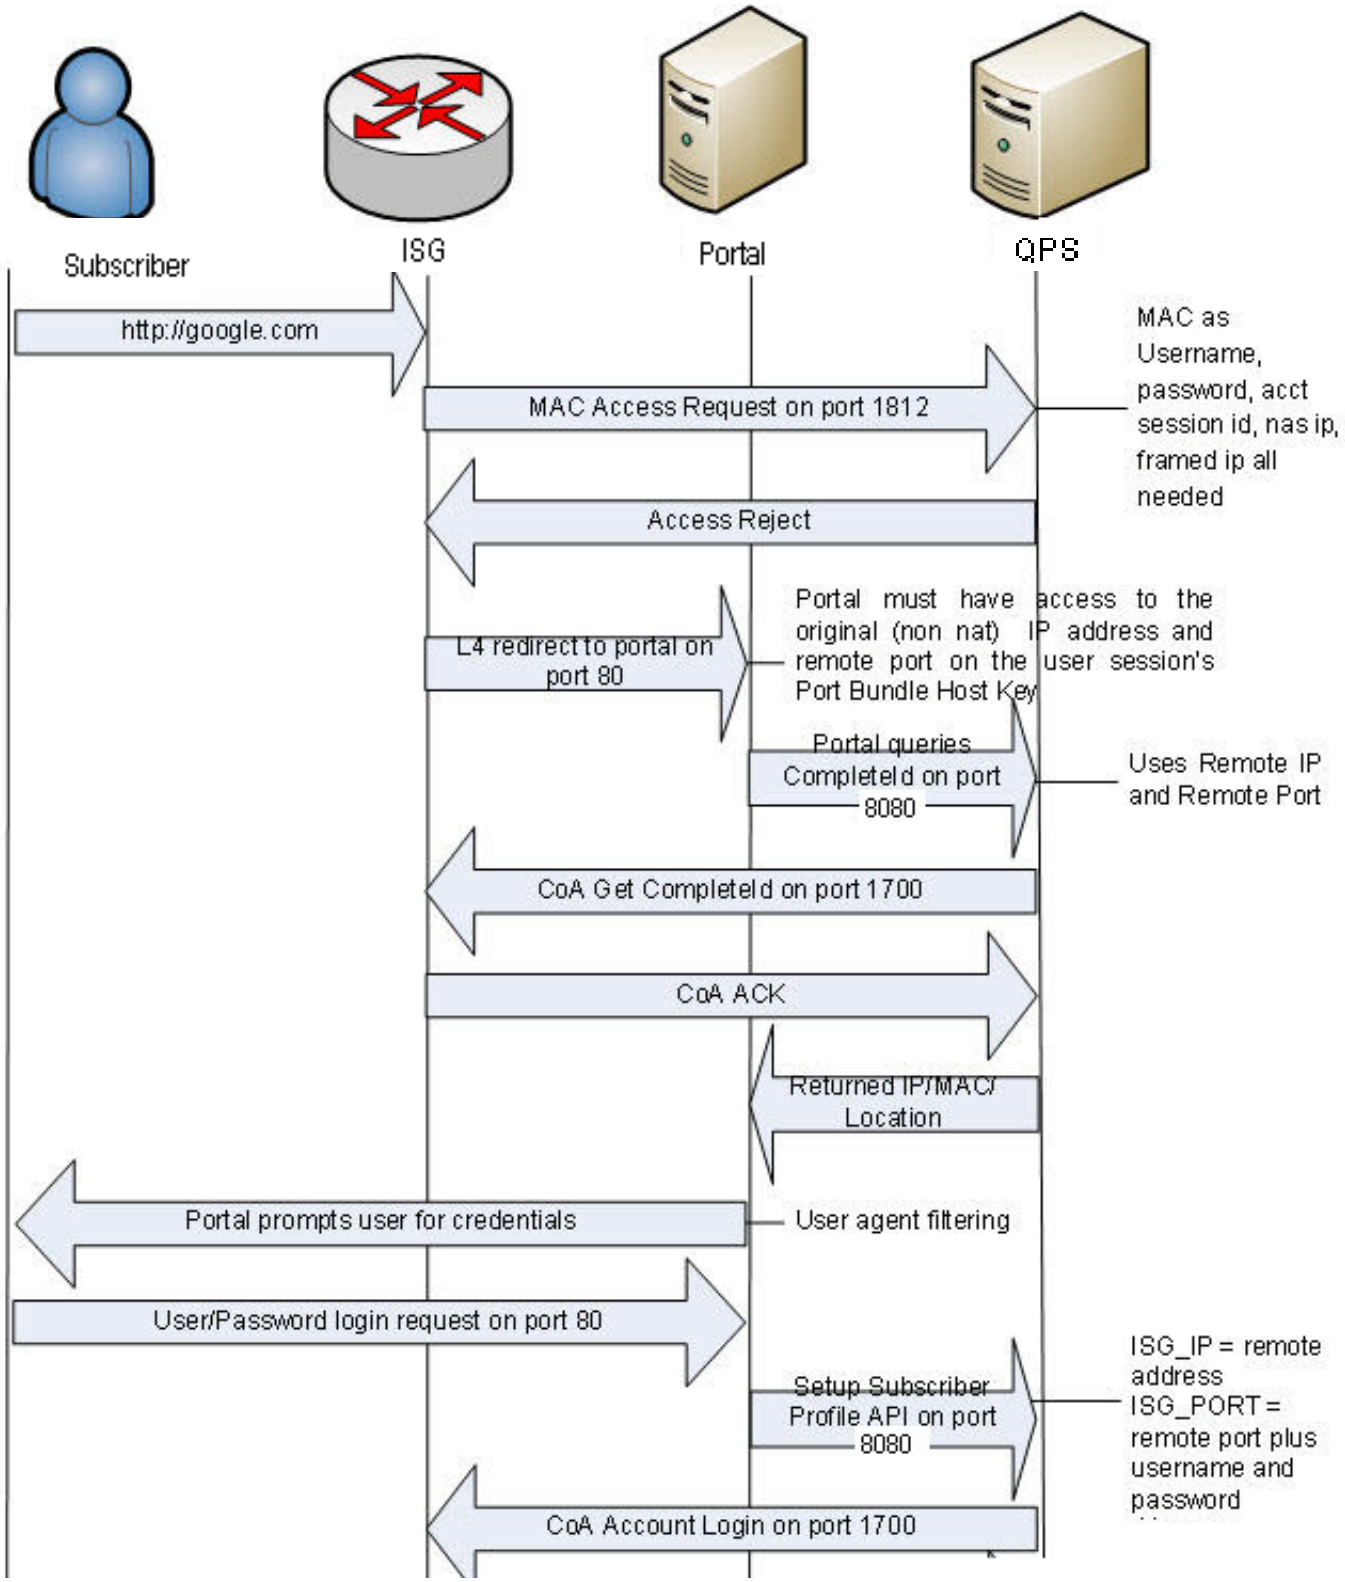

(continued)

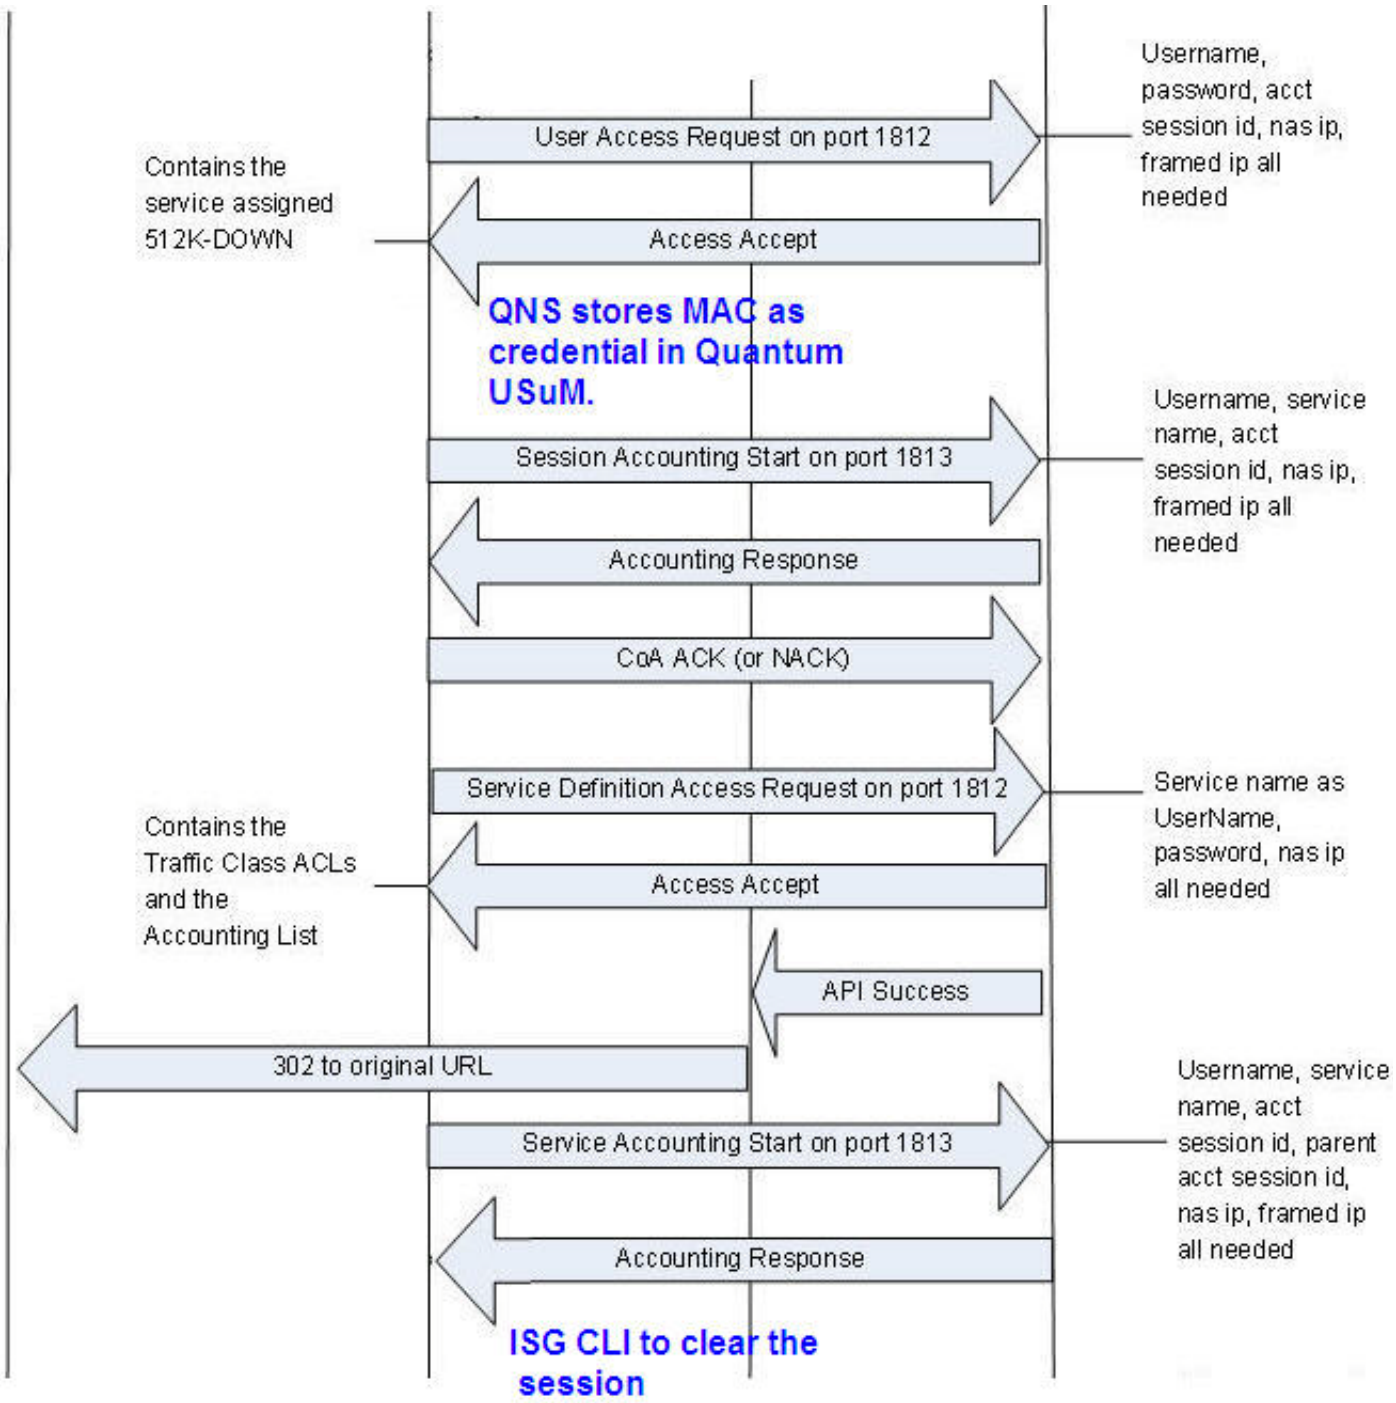
Username, password, acct session id, nas ip, framed ip all needed
User Access Request on port 1812
Contains the service assigned 512K-DOWN
Access Accept
QNS stores MAC as credential in Quantum USuM.
Username, service name, acct session id, nas ip, framed ip all needed
Session Accounting Start on port 1813
Accounting Response
CoA ACK (or NACK)
Service Definition Access Request on port 1812
Service name as UserName, password, nas ip all needed
Contains the Traffic Class ACLs and the Accounting List
Access Accept
API Success
302 to original URL
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Service Accounting Start on port 1813
Accounting Response
ISG CLI to clear the session

(continued)

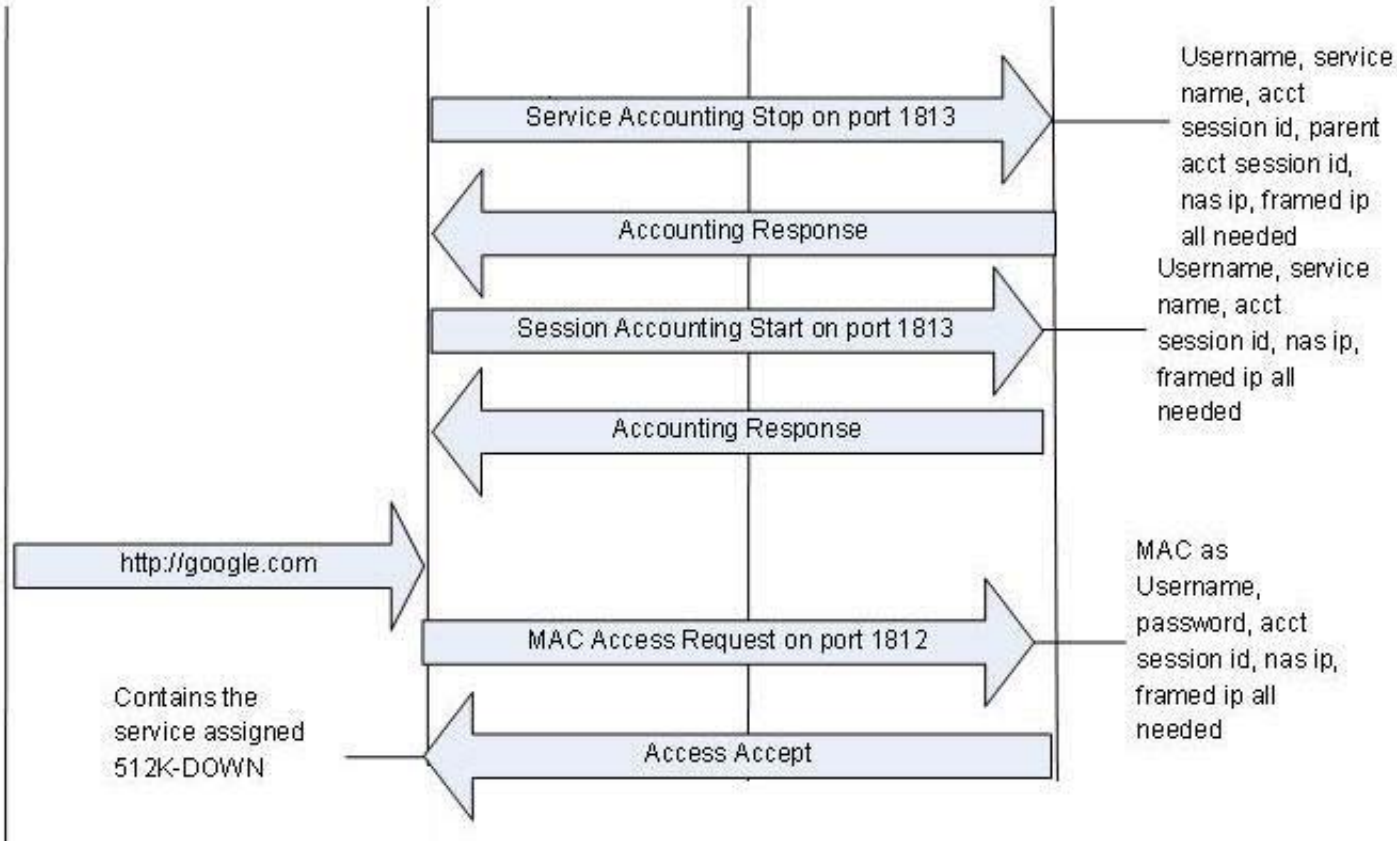
Service Accounting Stop on port 1813
Username, service name, acct session id, parent acct session id, nas ip, framed ip all needed
Accounting Response
Session Accounting Start on port 1813
Username, service name, acct session id, nas ip, framed ip all needed
Accounting Response
http://google.com
MAC Access Request on port 1812
MAC as Username, password, acct session id, nas ip, framed ip all needed
Contains the service assigned 512K-DOWN
Access Accept

Tiered Services Call Flow

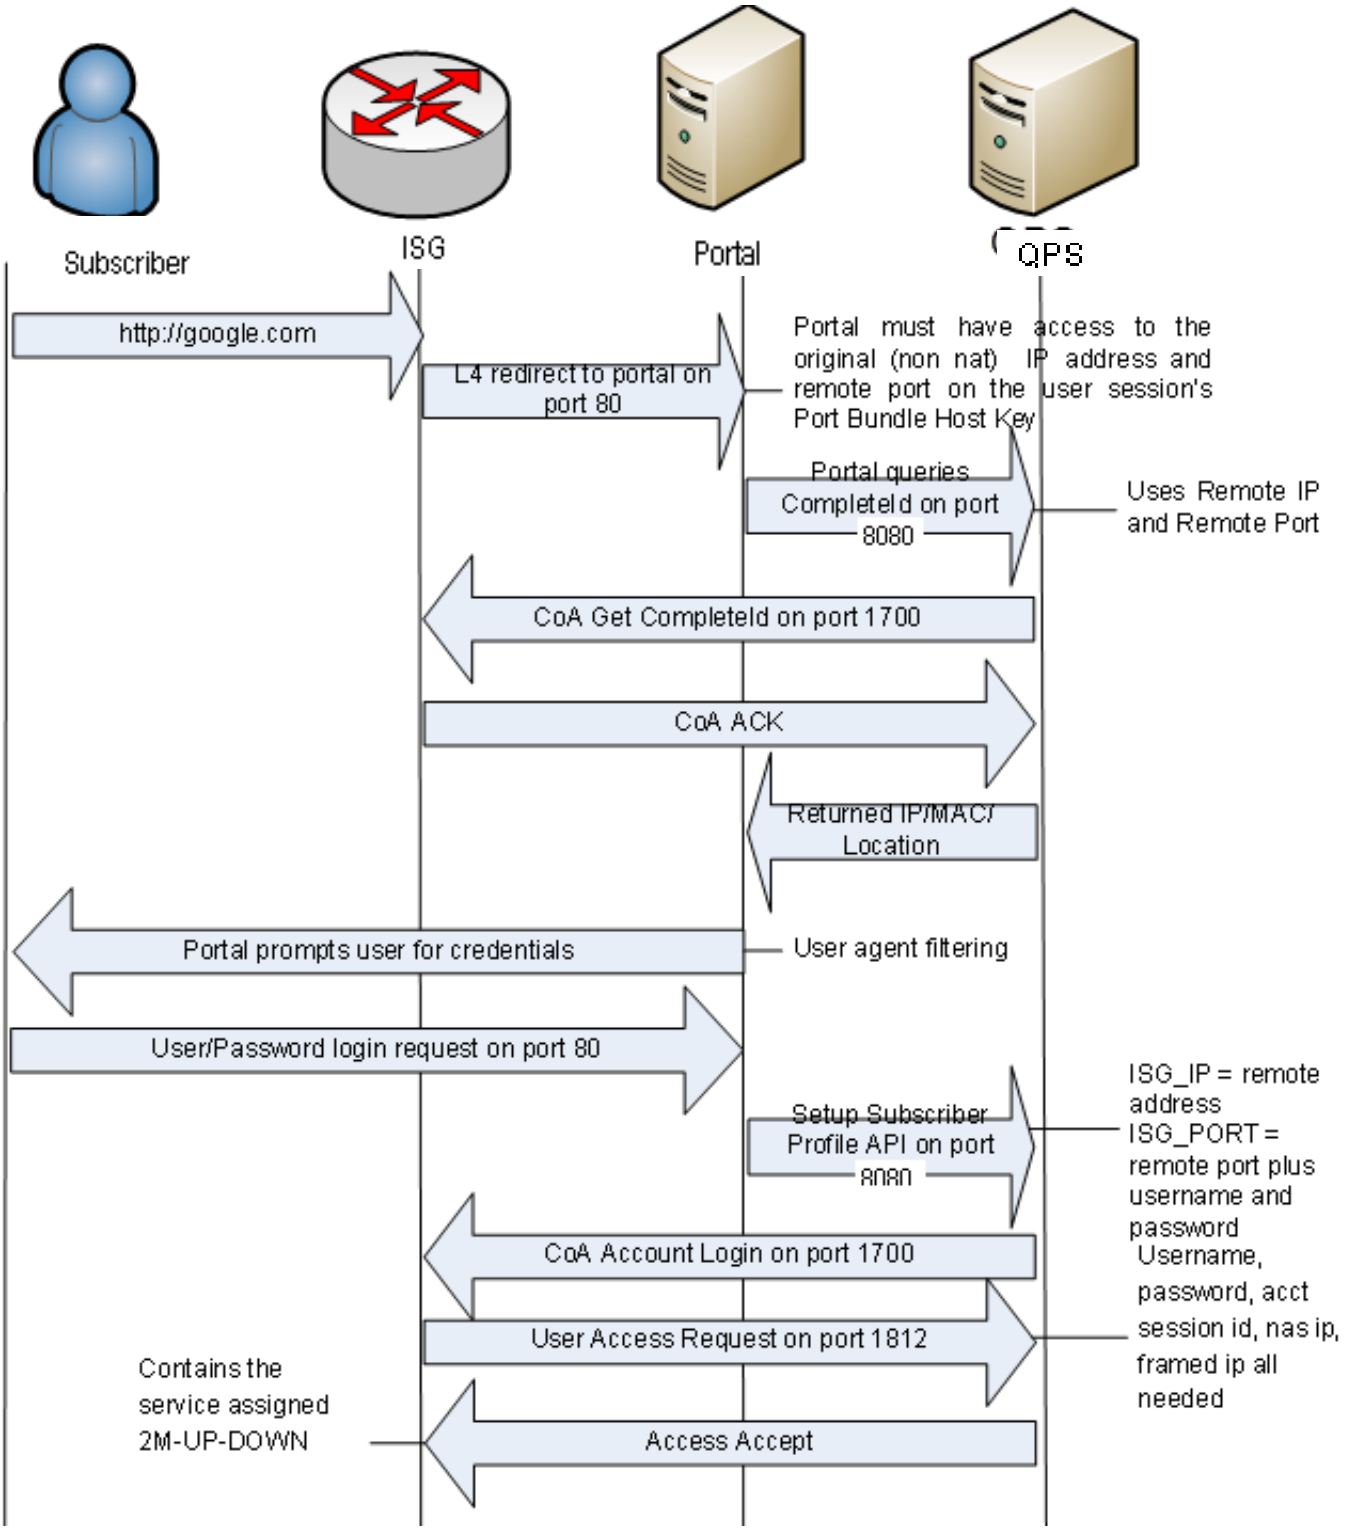

(continued)

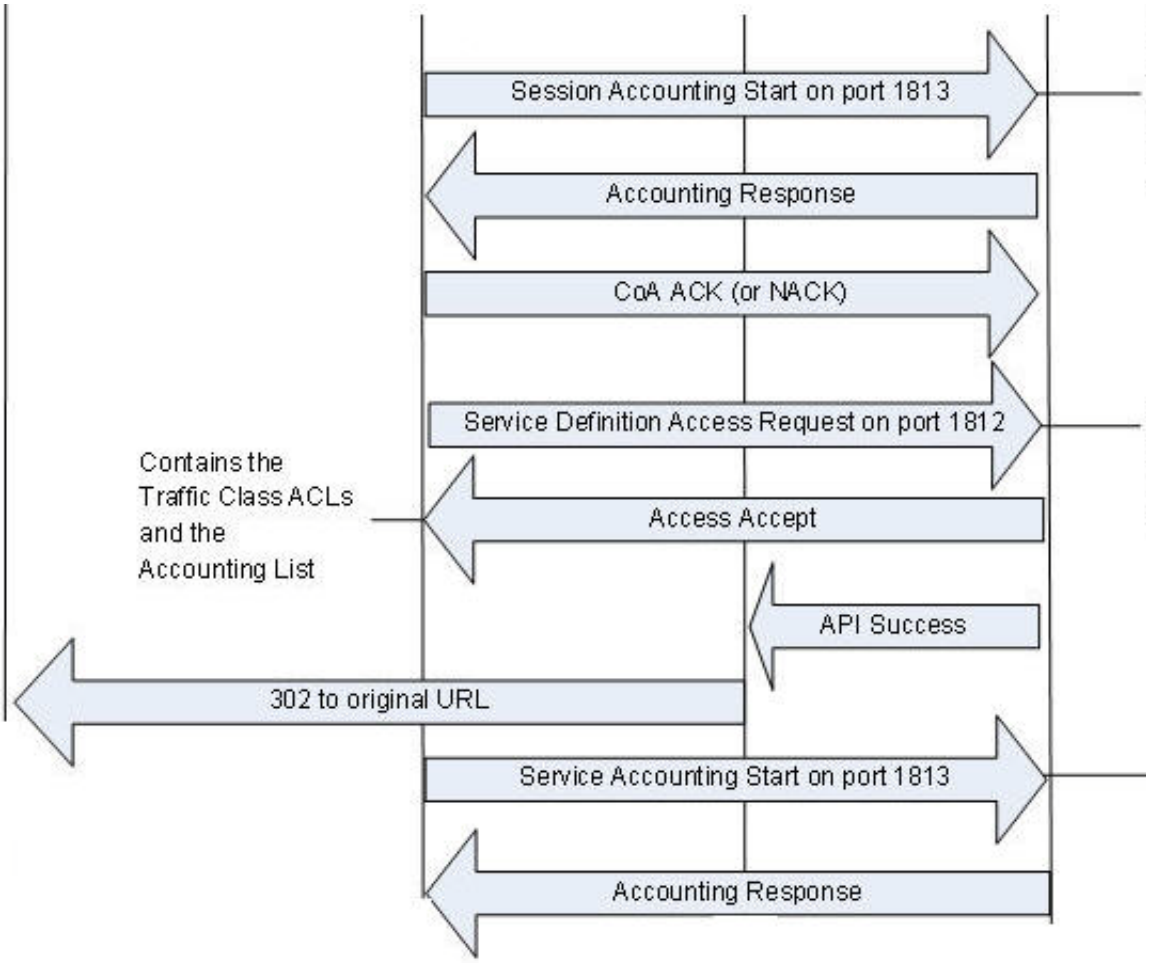

SP WiFi-4.0 Call Flows

Following are the call flows within the CPS system based on the requests received from the network device and the presence of the subscriber information in SPR profile.

Authentication Sequence in WLC Network Device Manager

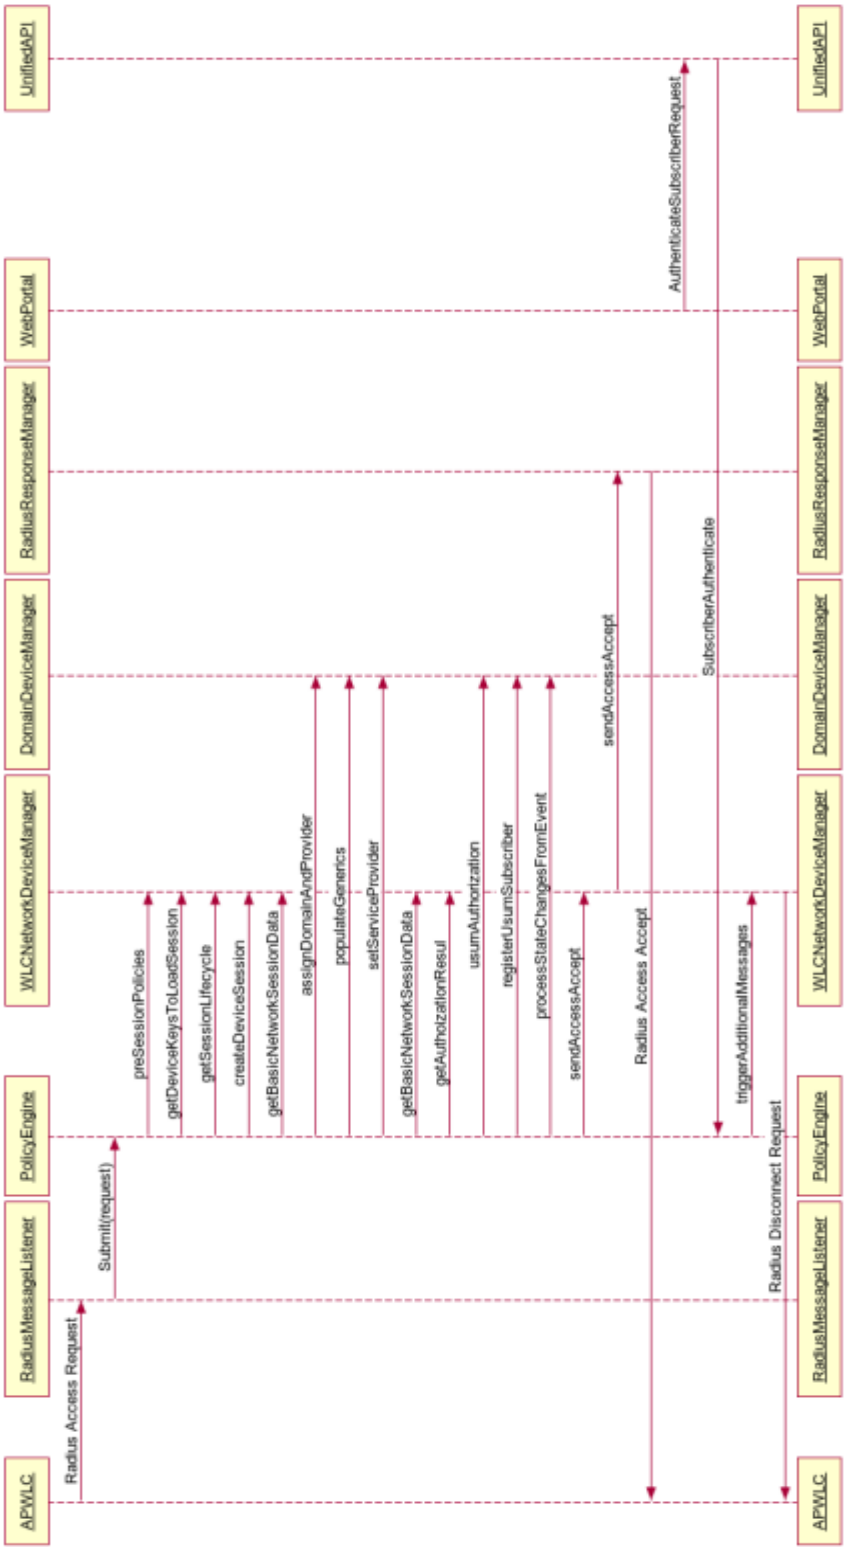

Authentication Sequence for EAP Requests

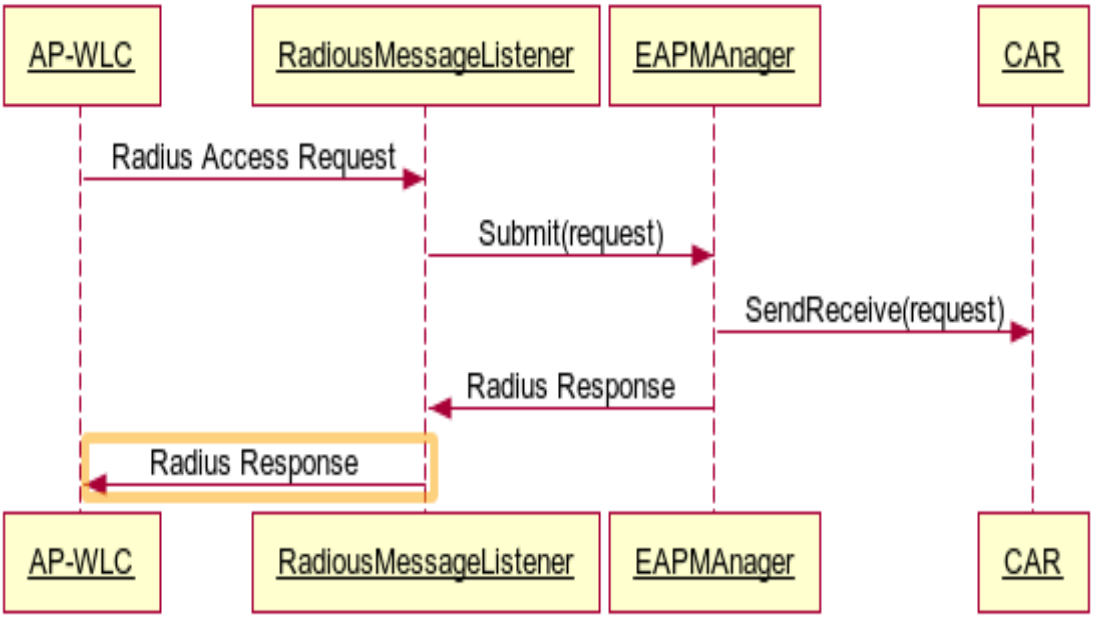

Authentication Sequence in iWAG Network Device Manager

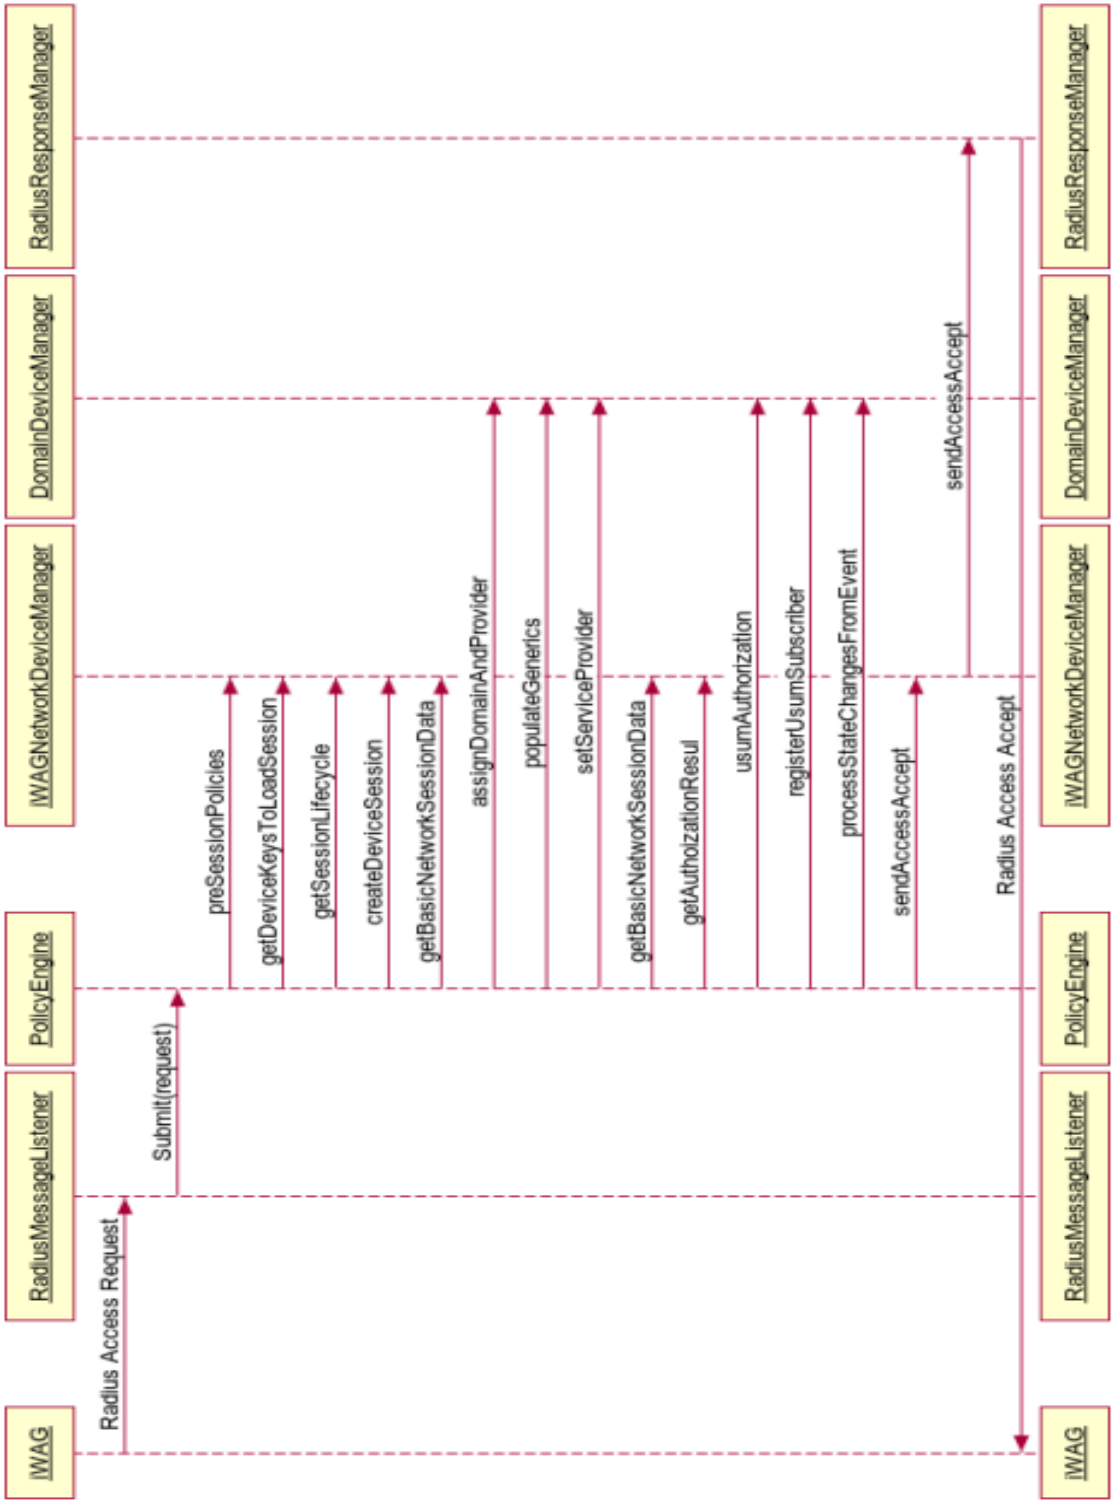

MAG Coupled Web Based Authentication

MAG coupled Web based authentication entails that the unauthenticated client is assigned a temporary IP address. This temporary IP address is used as the source address by the UE while accessing the web portal. After successful authentication, the client is forced to re-associate and that causes the client to go through the IP address assignment phase again. At this stage since the client is already authenticated, the LMA assigns the IP address for the client.

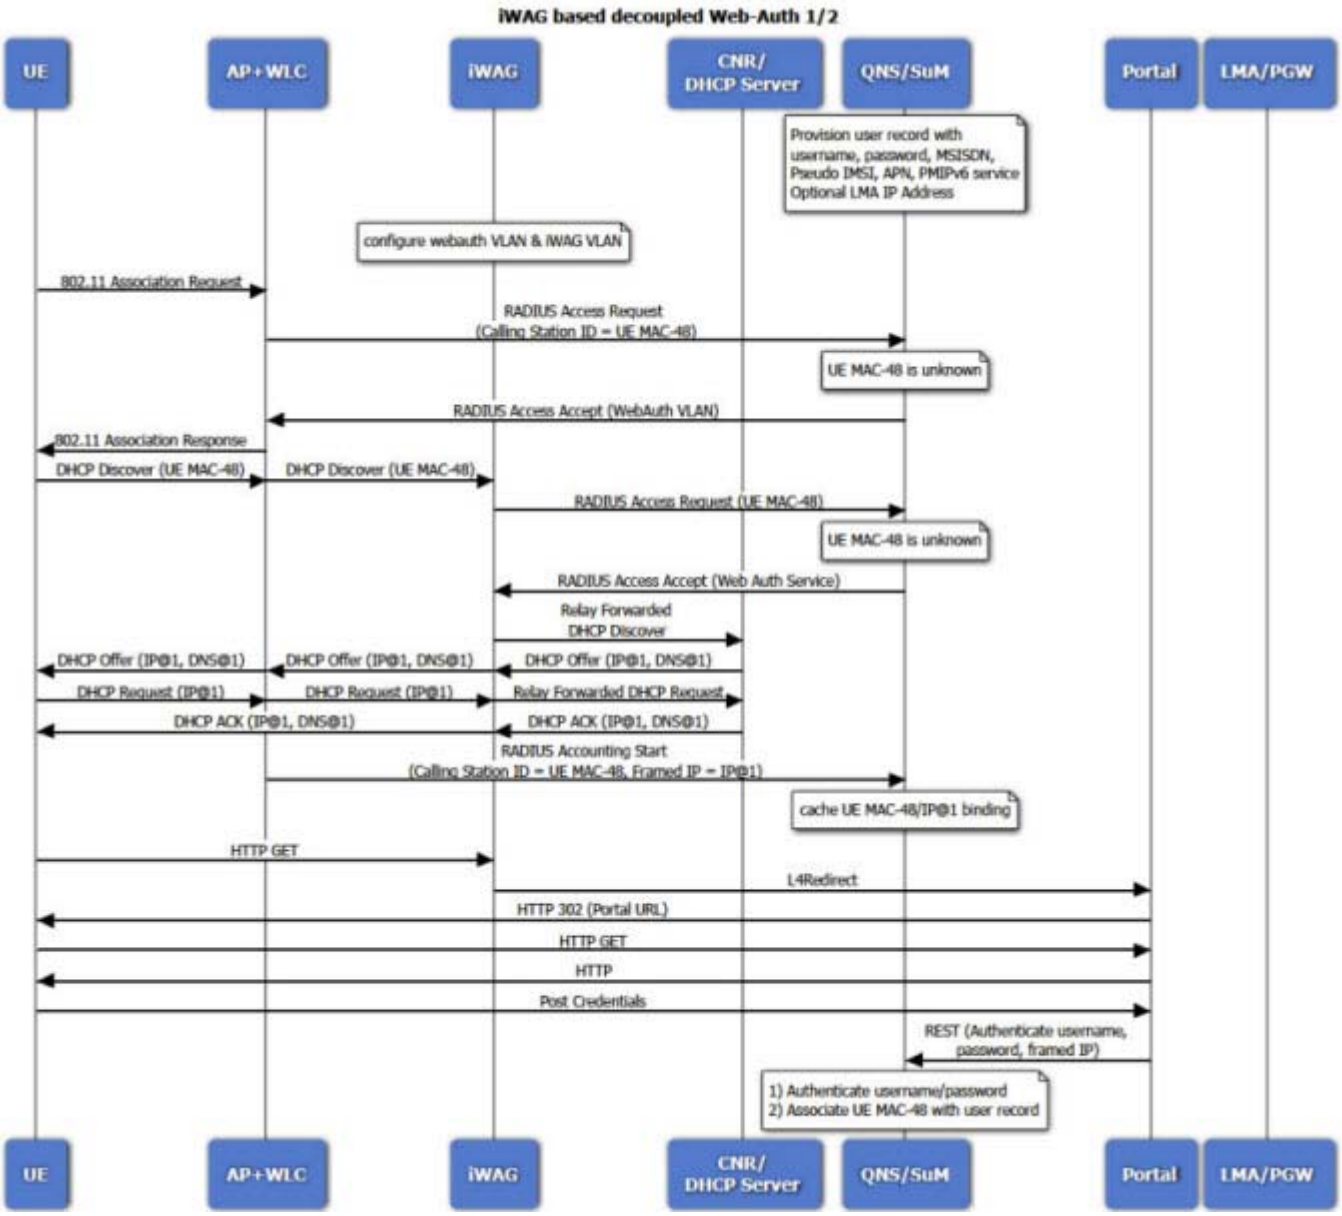

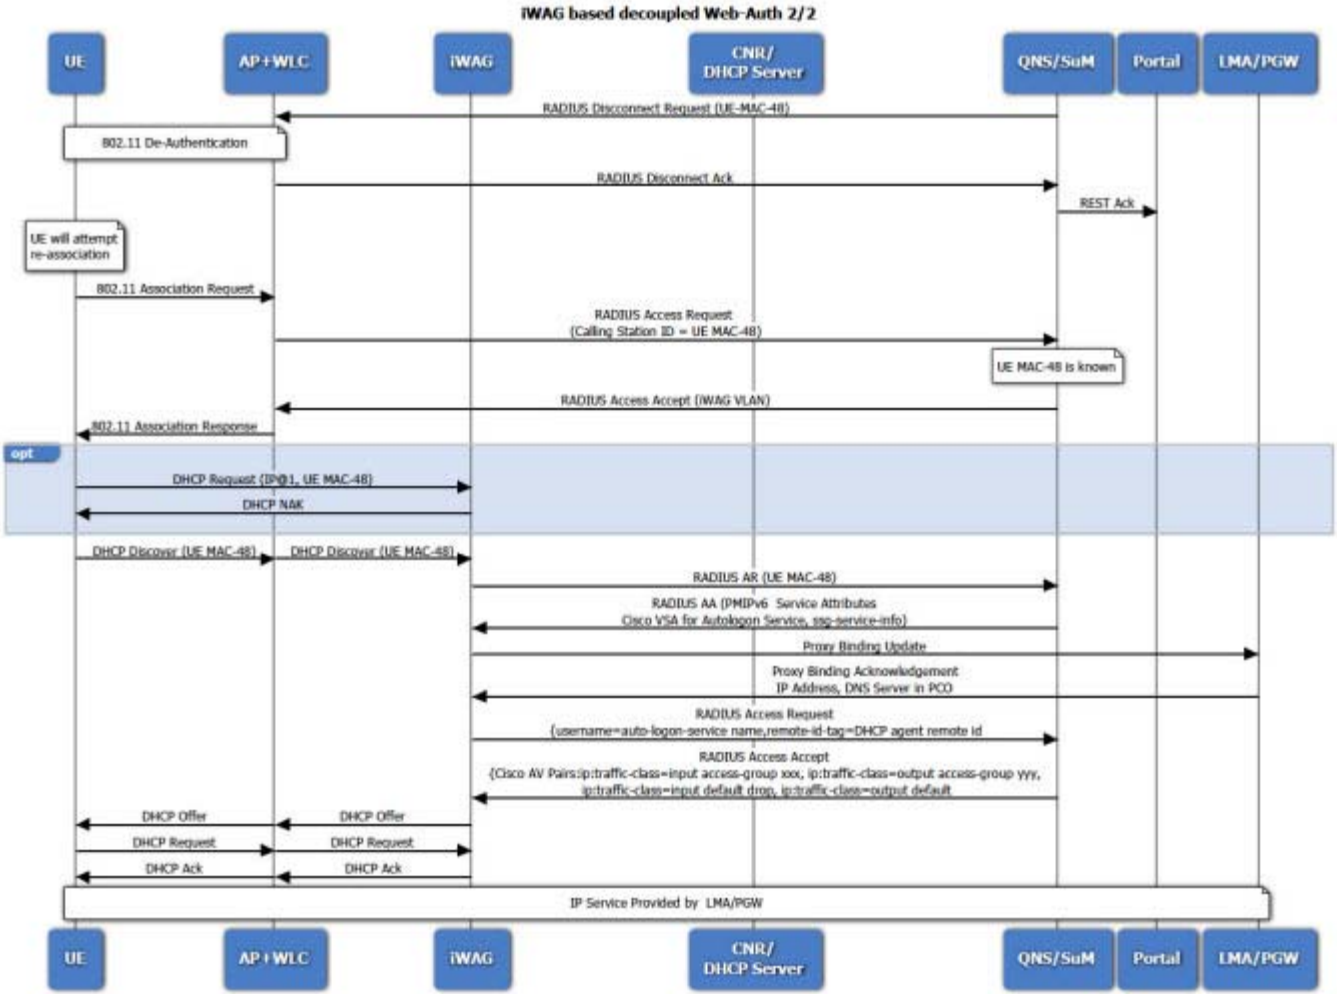
iWAG based decoupled Web-Auth 2/2
UE
AP+WLC
iWAG
CNR/
DHCP Server
QNS/SuM
Portal
LMA/PGW
RADIUS Disconnect Request (UE-MAC-48)
802.11 De-Authentication
RADIUS Disconnect Ack
REST Ack
UE will attempt
re-association
802.11 Association Request
RADIUS Access Request
(Calling Station ID = UE MAC-48)
UE MAC-48 is known
RADIUS Access Accept (IWAG VLAN)
802.11 Association Response
opt
DHCP Request (IP@1, UE MAC-48)
DHCP NAK
DHCP Discover (UE MAC-48)
DHCP Discover (UE MAC-48)
RADIUS AR (UE MAC-48)
RADIUS AA (PMIPv6 Service Attributes
Cisco VSA for Autologon Service, ssg-service-info)
Proxy Binding Update
Proxy Binding Acknowledgement
IP Address, DNS Server in PCO
RADIUS Access Request
(username=auto-logon-service name,remote-id-tag=DHCP agent remote id
RADIUS Access Accept
(Cisco AV Pairs:ip:traffic-class=input access-group xxx, ip:traffic-class=output access-group yyy,
ip:traffic-class=input default drop, ip:traffic-class=output default
DHCP Offer
DHCP Offer
DHCP Request
DHCP Request
DHCP Ack
DHCP Ack
IP Service Provided by LMA/PGW
UE
AP+WLC
iWAG
CNR/
DHCP Server
QNS/SuM
Portal
LMA/PGW

Account Balance Template Configuration and Tariff Time Configuration

Revised: July 10, 2015

Account balances are best defined under the Services tab, with service options and services. This chapter is provided for completeness to explain the Reference Data tab.

Along with concepts about balance, quota, and tariffs, these configuration procedures are provided:

- Changing the Party Billed
- Defining the Systems
- Defining the Account Balance Template
- Setting Up Rates and Ratings
- Tariff Switching Times

Cisco Policy Suite's Multi-service Balance Management (Cisco MsBM) feature stores policy-related balances, such as usage, against application usage, fair use caps, bill shock thresholds, and roaming caps.

One way to view balances or apply top ups to a specific subscriber is with the Cisco Control Center interface.

This appendix covers the following sections:

About Cisco MsBM

Cisco Policy Suite defines the times, rates and balances for policies.

In CPS some policies track subscriber usage. For example, a service plan may have a daily application use cap of 10 KB and the subscriber has requested notification if their usage exceeds a 9KB. In support of such policies, the Cisco MsBM feature stores policy-r elated balances against the cap.

- Multiple recurring balances (for example, reset daily, monthly, or per billing cycle) as well as one-time balances such as an introductory offer might be assigned per subscriber and can be topped-up, as appropriate.
- Group or sub-account balances may also apply to a subscriber's usage, for example, as in family or corporate plans.
- Subscriber-defined thresholds may also be defined, for example, for bill shock, roaming caps, and family or corporate plans.

When to apply a given balance is based on policy rules configured by you, the service provider. For example, the roaming cap balance applies only when the subscriber is roaming. The fair use balance may apply to certain traffic only during peak hours.

Various policy actions may be defined upon thresholds or balance exhaustion in the policy rules, including subscriber notifications, browser redirects, and bandwidth downgrades.

CPS defines the times, rates and balances for such policies.

Balance Management Data Model

Balances may be of these types:

- Recurring—recurring balances automatically replenish on a configured cycle, for example daily, weekly, monthly, or per billing cycle. Any remaining balance is not carried over when the balance is reset.
  - Roll Over—a type of recurring balance. Roll over balances, which are additive, are similar to recurring balances, but any remaining balance is kept when the balance replenishes.
- One-time—one-time balances do not automatically replenish. One-time balances may be configured with an expiration date.
  - Top Up—a type of One-time balance. Top ups add to a subscriber's balance and may be applied to any of the above balance types. Top ups are differentiated from the balance to which they are applied, enabling the service provider to define the order in which top ups, and the underlying balance, is consumed. Top ups may be assigned an expiration date.

Balances can be configured to be automatically allocated in the Cisco MsBM for a subscriber during a balance reserve if a balance does not already exist. The fair use balances for new subscribers is an example. Auto-provisioned balances are created based on rules configured by the service provider, which may refer to parameters such as subscriber class, rate plan, and so on, when provisioning the balance.

Balance Operations

The Cisco MsBM supports the following balance operations:

- Provision—initializes new balance for a subscriber.
- Reserve—sets aside a portion of the subscriber's balance for imminent use.
- Charge Reservation—a charge always releases the reservation but provides the ability to reserve a new amount. The two types of Charge Reservation are Debit Reservation and Release Reservation.

  - Debit Reservation—applies a debit to an outstanding reservation.
  - Release Reservation—releases a remaining, unused reservation and optionally debits an amount of the reservation before release.
- Direct Debit—immediately applies a debit to the subscriber's balance.
- Direct Credit—immediately applies a credit to an existing balance.
- Query Balance—retrieves a subscriber's current balance. Query balance can be divided into recurring top up(s) and reserved balances. A Query Balance operation may be configured to request the network to report current usage against reserved balances.

In a typical call flow, the Reserve operation sets aside a portion of a subscriber's total available balance in the database for use. As the subscriber uses their service, the Policy Control Enforcement Function (PCEF) reports usage, and the debit operation subsequently subtracts that usage from the subscriber's balance.

Balance APIs

Balance operation Web Services APIs enable external systems to integrate with the Cisco MsBM. A custom portal to query and top up a subscriber's balance is an example of this. Reservations, debits, and credits may also be placed against a balance by external systems, such as a content delivery platform.

Tariff Switching

The Cisco MsBM can set time-of-day tariff switches for quota expirations and so minimize network signaling, especially for recurring balances, which typically reset at the same time and date across subscriber groups. In addition, the Cisco Policy Suite supports configuring tariff boundaries to support rating changes. Tariff time switches are supported and can be configured by a combination of these elements:

- Time of day, for example, where 19:00-06:00 is configured as night, 06:00-19:00 is day
- Day of the week for which it is valid, for example, Monday – Friday
- Holidays, for example, Nov. 3, 2014, Nov. 23, 2014, Dec. 23, 2014

Rating

Rating is a way to modify the amount of balance reserved, debited, and/or credited for a given balance operation using a multiplier. This way, CPS is never configured with a specific dollar, yen, or Euro amount.

Although Cisco MsBM may be integrated with an external rating engine, basic real-time rating rules may be configured within the Cisco MsBM itself. Some typical examples of rating are these:

- During the day time, a subscriber's data usage is applied 100% to their current balance. However, at night, only 25% of a subscriber's data usage counts towards their monthly quota. A multiplier of 0.25 is used.
- There is a relationship between data usage and currency (for example 10 MB = $1). Usage can then be related to currency and then stored or charged.
- The first 5 MB of data usage during the day/week/month/billing cycle is free, for example, included in the plan. After that, subsequent traffic is charged (for example 10 MB = $1).

- Due to a subscriber receiving a special promotion, they are charged only 10% of their bandwidth usage at night, rather than 25%.

Policy Conditions and Actions

You can configure Cisco MsBM to take various policy actions on various balance thresholds or on exhaustion. Thresholds may be set as either a predefined value or percentage of a balance, and may also be subscriber-defined. For example, a subscriber's usage reaches…

- 80% of their fair use limit and the subscriber receives an SMS message of their usage level. At 100% usage, the subscriber's bandwidth is downgraded until the fair use limit is reset at the end of the billing cycle.
- 100% of their daily P2P allowance at which the bandwidth allotted for P2P is downgraded.
- 50 ? while roaming, is notified via SMS, and the subscriber's browser is redirected.

Basic Configuration Overview

Configuring the Cisco MsBM consists of these basic steps:

- Load the feature
- Define the repository configuration
- Define extension points and policies
- Publish and load the data
- Test with traffic

Cisco MsBM uses these components and concepts:

- Balance management configuration
- Balance management primitives

  and optionally,
- SCE Quota or ISG Prepaid Balance Management
- Tariff Switching Times
- Thresholds
- Rates and Ratings
- Various Charging Ids
- Subscriber Management (SuM)
- Policies, their conditions and actions
- Notifications

Where to Begin

To implement the Cisco MsBM, have these tasks complete and working:

- Have CPS installed and taking traffic.

- Have HA and LB set up if you use that.
- Open the Policy Builder interface.

Assumptions for This Example

Your business rules are different of course. but the business rules in this example are common:

- Daytime usage is full cost, that is, 100 percent of the specified charging rate.
- Nighttime use is 50 percent of the day rate,
- Weekend use is 25 percent of the day rate.
- One Holiday on March 17 with a daytime tariff switch at 7 a.m. and 7 p.m. After 7 p.m., use the Night rate, of 50 percent.

Defining the Systems

Step 1 Click the Reference Data tab > Systems node > a system.

Step 2 Open the Plugin Configurations folder in the tree.

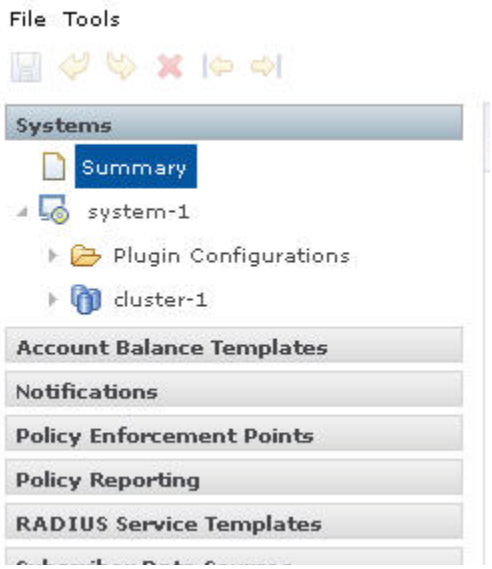

Step 3 In the main window, click Balance Configuration.

The Balance Configuration screen appears.

Step 4 Fill in these fields as explained in the table below the figure.

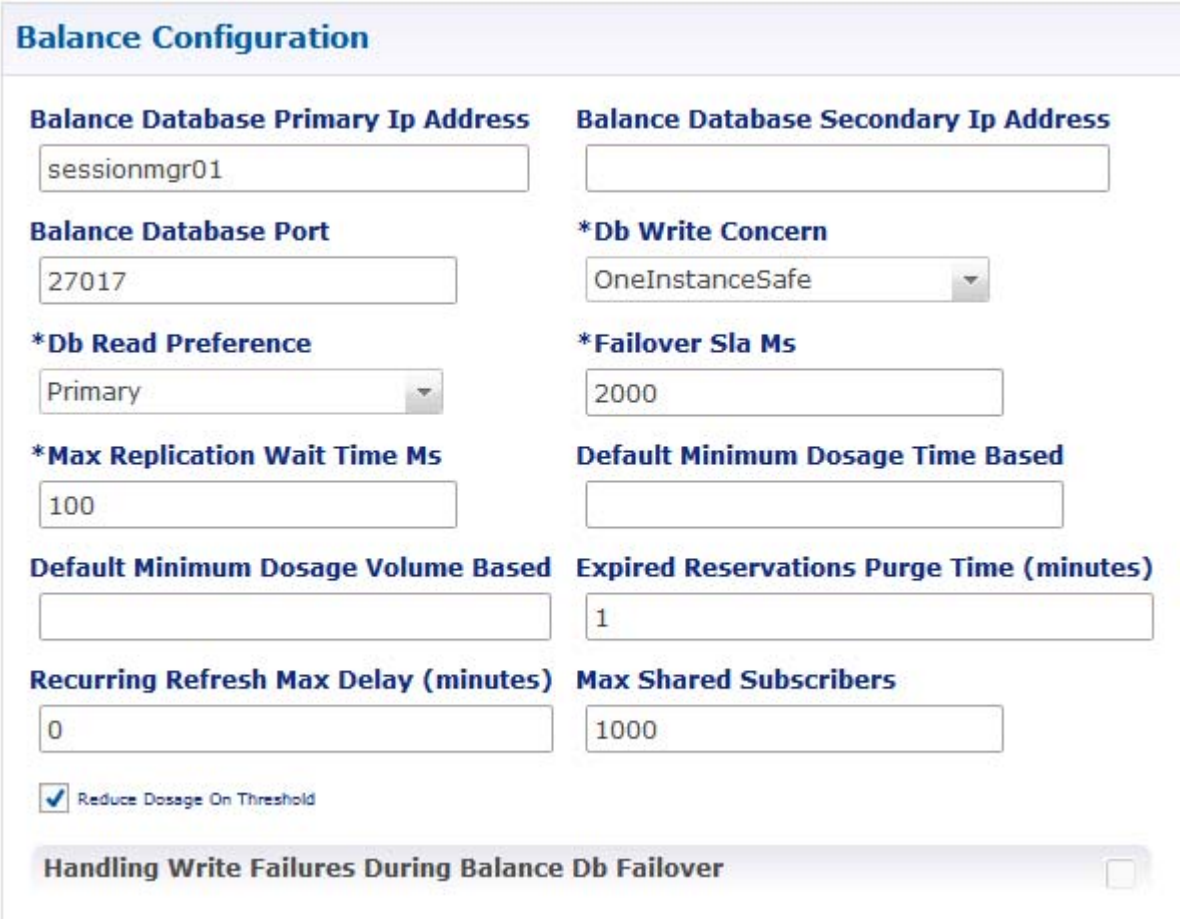

Balance Database Primary Ip Address	This is required. Enter the DNS name of the ou must install and configure a sessionMgr03 for balance data, for example.
Balance Database Secondary Ip Address	This is used only if you have the database that holds your balance data, usually sessionMgr01. If you want to use your sessionMgr01 exclusively for session data, have a HA Session Manager database available. Enter the backup, HA sessionMgr02 database DNS name here.
Balance Database Port	This is required. This is the port the Balance database uses, that is, the port of sessionMgr.
Db Write Concern	Controls the write behavior of sessionMgr and for what errors exceptions are raised. Default option is WriteConcern.Safe.
Db Read Preference	Read preference describes how sessionMgr clients route read operations to members of a replica set. http://docs.mongodb.org/manual/core/read-preference/
Default Minimum Dosage Time Based	This field is optional but recommended. This is the minimum amount of time that is granted for a reservation, assuming quota is not exhausted. If you want to manage subscriber balances on the basis of time used, check with the network device administrator and have this value be slightly larger than the minimum amount of time the network device such as an SCE or ISG accepts for a reservation.
Default Minimum Dosage Volume Based	Using this field is optional but recommended. This is the minimum amount of volume that is granted for a reservation, assuming quota is not exhausted. If you try to make a reservation for 1 KB, and your minimum is 10 KB, the router rejects it because it is too small an amount to bother with.
Expired Reservations Purge Time (minutes)	The amount of time a record of expired reservations is retained and Cisco MsBM attempts to charge them. Note that expired reservations are charged only if sufficient quota is still available; that is, expired reservations do not retain the lock on quota that current reservations do. The default value is 0.
Recurring Refresh Max Delay (minutes)	The amount of time refreshing of recurring quotas are staggered across randomly, for sessions that are not actively using quota but are still established. This parameter is used in cases where subscribers always have a session, but might not be using their quota actively. This allows staggering of recurring refreshes where the customer has set all their subscribers to refresh at the same time, say midnight. It avoids spiking the CPU. The default value is 0.
Reduce Dosage on Threshold check box	When checked, reservation dosages are reduced as an Cisco MsBM threshold is approached. This way, a dosage does not pass a threshold by a large amount before notification of the breach is sent out. When unchecked, normal dosages is granted. Recall that when enabled, messaging becomes much more chatty, but threshold breach accuracy is enhanced.

About Minimum Dosages

To help improve system performance and also manage user accounts in a consistent manner, most network devices support the concept of minimum dosage.

When CPS tries to create a reservation for a subscriber, the network device checks the reservation request against its minimum dosage value. The network device evaluates like this: “If I can’t give you at least this many bytes, I will give you zero.”

Also, when approaching a threshold, if the subscriber doesn’t have enough quota remaining to satisfy the router, based on the minimum dosage, the reservation is not made. Cisco MsBM balance algorithms does not provide an amount of quota that the router surely rejects.

No reservation is made that exceeds a threshold until the subscriber meets the threshold, unless a minimum dosage is specified.

Without Minimum Dosage: Hard Threshold

For example, a subscriber has 100 MB of quota with a threshold of 50 MB. The subscriber has used 45 MB and logs in. The CPS makes a reservation request for 10 MB. Because no minimum dosage is defined, the CPS only creates a 5 MB reservation because of the thresholds set to 5 MB. When the subscriber uses up the 5 MB, the CPS sends a breach notification message. The next reservation is for the amount requested.

With Minimum Dosage: Soft Threshold

Now, consider the same example but with a minimum dosage. The subscriber logs in and the CPS makes a reservation request for 10 MB. This time, because of the minimum dosage value, the CPS creates the 10 MB reservation even though there is a notification set at 50 MB. In this case, the subscriber uses the whole reservation, bringing the total to 55 MB, which then triggers the notification message.

Overcharging

Overcharging a reservation allows the subscriber to consume slightly more quota than they truly have available.

Overcharging occurs because:

> Usually overcharging a quota is permitted because the service provider knows or expects that the subscriber is refreshed or topped up soon. With overcharging, the subscriber is not cut off from service exactly at the end of their quota limit. Rather, they are given a grace amount, avoiding unpleasant support calls from the subscriber.

Overcharging and so over-consumption may occur this way:

> A PCEF uses five minute intervals to check usage when Cisco MsBM sends usage volume and enables usage monitoring. When the threshold-met trigger is sent in CCR Update, a higher used volume is reported than the provisioned threshold. For example, if CPS provisions 20 MB as the threshold, when the threshold is met, the usage reported is around 70 MB because of this.
>
> In this use case, the subscriber is downloading a large file. The download occurs within the five-minute interval. The quota is granted, the download occurs, the evaluation happens, and the resulting overconsumption is noted.

Cisco MsBM handles over-consumption this way:

- If a reservation is over-consumed (that is, the reported usage is greater than the reserved amount) and there is sufficient balance to cover the debit in full:
  - Cisco MsBM debits the reported usage in full.
  - If the debit results in a threshold being breached, Cisco MsBM:

    i. Debits the reported usage in full.

    ii. Notifies the Cisco Policy Server of the breach and the amount by which the threshold was exceeded.
- If the debit overdraws the balance (that is, there is insufficient balance to debit the full reported usage) Cisco MsBM:
  - Debits the balance to exhaustion.
  - Notifies the Cisco Policy Server of the exhaustion event and the amount not debited (that is, the overage).

Defining the Account Balance Template

This section defines the basic account balance information to the CPS system. This procedure sets up an account balance template for data, but you could easily adapt it for voice or IP cable TV. Your account representative is always available to help you do this adaptation if you need it.

Note Cisco MsBM works with several network devices and protocols. This section presents information using screens showing a Cisco SCE, but balance management can use any network device. Make sure that your PEP is installed and properly configured with the packages you want.

Step 1 From the Cisco Policy Builder main page, click Reference Data tab > Account Balance Templates node, and click the Account Balance Template link

Step 2 Check the display of the Account Balance Template form on your screen.

In this screen, use of the Thresholds table is optional, but it is conveniently provided for you.

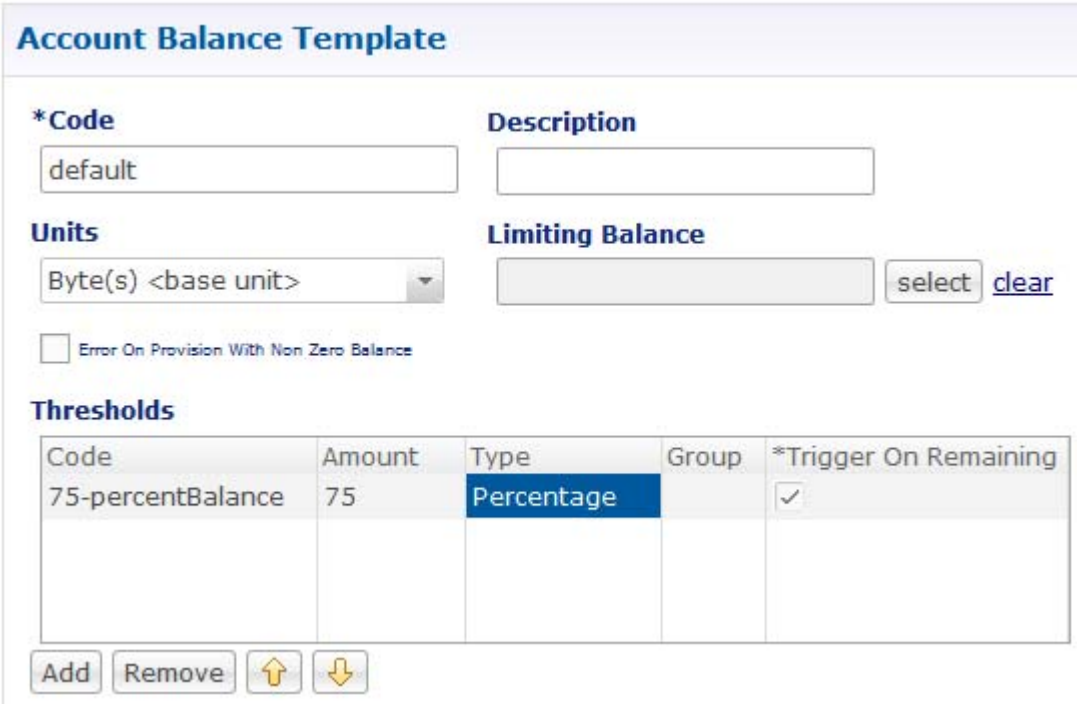

Step 3 Fill in the template using the information in the table below.

Code	The name you want to give to this template. Our example is tracking data usage balances and so has the Code of Data. Code is what you refer to when defining polices.
Description	Provide a description for use by other administrators in this optional field.
Units	You can monitor the balance of subscribers in terms of how much data they consume, how much currency they spend, or how much time they use. For this template, select a unit with which to measure the balance. Currently, following Units are supported: • Byte(s) • Currency • Gigabytes(s) • Hour(s) • Kilobyte(s) • Megabyte(s) • Minute(s) • Other • Seconds(s)

Thresholds Table Area	The Thresholds table lets you specifically set the limit for balance management quota. When a threshold is reached, a message is sent to the policy engine.
Code	Enter a name for the threshold you are going to check against when your policy executes.
Amount	Set an amount. This specifies the cut off point for the Type in the next column.
Type	For what you are measuring, this is the measure of total balance. You can measure the actual amount of quota or a percentage of the quota. To measure time as seconds or any other quantity that is the same as the type of units used by the balance, for instance Euros, use the Other item in this drop-down list. Currently, following Types are supported: • Percentage • Bytes • Kilobytes • Megabytes • Gigabytes • Other
Group	Group thresholds together and take the first one in the table that evaluates as true. That is, force the logic of the table to refrain from notifying the policy engine of any subsequent threshold breaches. By grouping thresholds together, an account that has breached the 80 percent threshold does not trigger a 50 percent, 60 percent or 70 percent notifications.

Shared Quota: Per-User Limits

Family and other shared quota plans often have a shared quota that applies across all the multiple members/devices covered by the plan. The account owner may wish to limit the usage of the shared quota by one or more of the account members. Limits are not an allocation of quota to the user.

MsBM is an existing component in CPS which in conjunction with the SPR component allows the configuration of shared balance usage across multiple subscribers. The shared quota per user limit adds the ability to impose limits on how much of the shared quota an individual subscriber can use from the shared balance amount. Additionally, by leveraging existing threshold functionality in Balance, thresholds are defined against these limits, which when reached trigger notifications such as emails, SMS.

To add a subscriber that is participating in a shared balance, user needs to create two balance accounts:

- One balance account must be their individual account.
  - This account must contain any balances/quotas that are only available to the individual.
  - This account must also contain one balance that is to be used for tracking the per user limit.
- The other Balance account is not owned by the subscriber and contains the shared balance/quotas.

Step 1 Click Reference Data > Account Balance Templates > Summary.

Step 2 Click Account Balance Template to open a window and create an account where the user can allocate the quota limits. An example is shown.

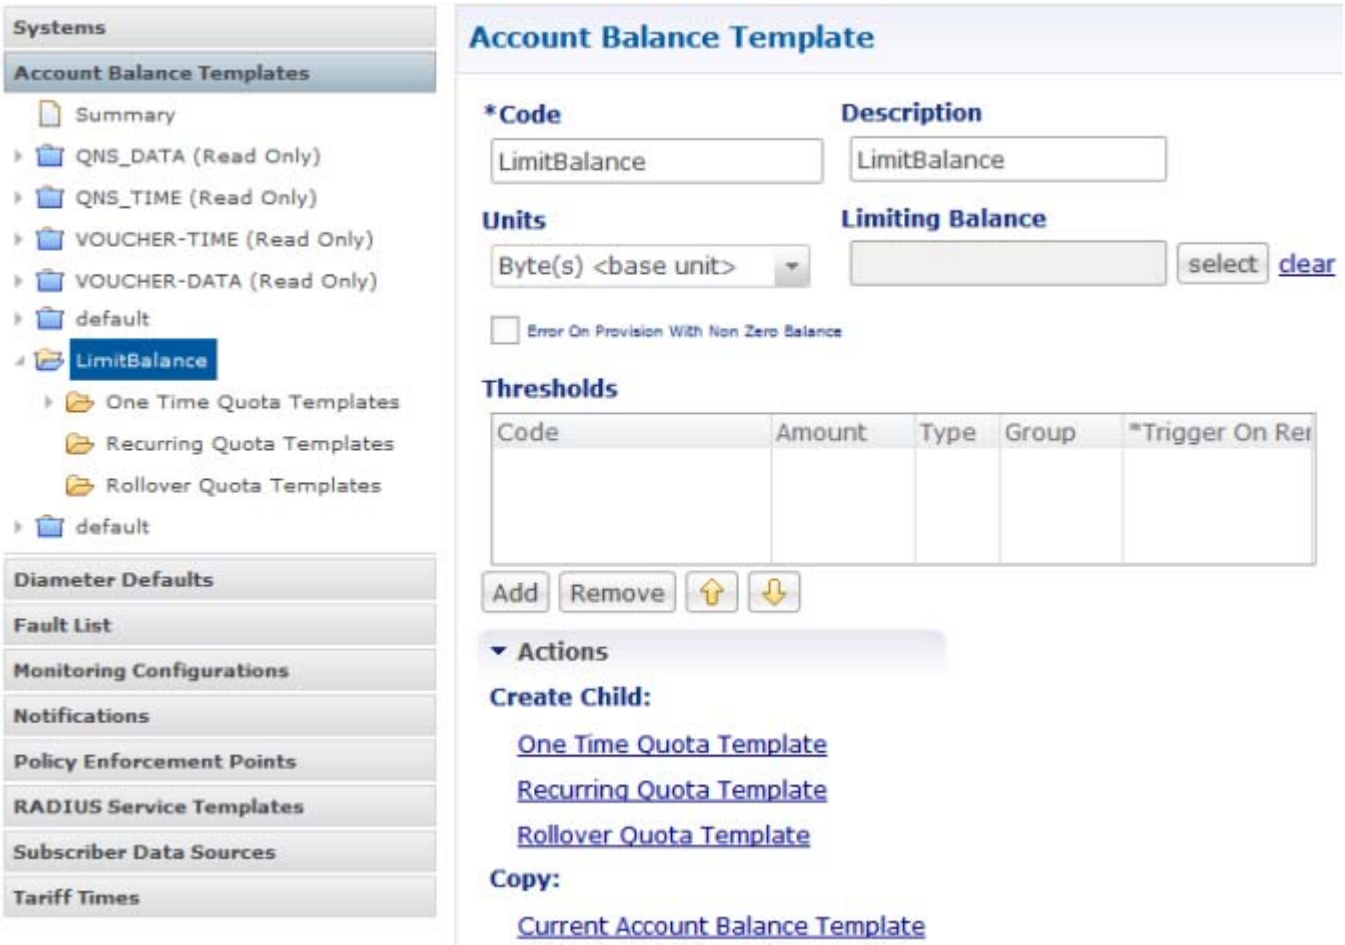

Step 3 Select the Units of quota from Units drop-down list.

Step 4 Under One Time Quota Templates, a quota template needs to be defined for each limit that can be applied to any subscriber.

For example, the user can create a 2MBLimitBalance. An example is shown.

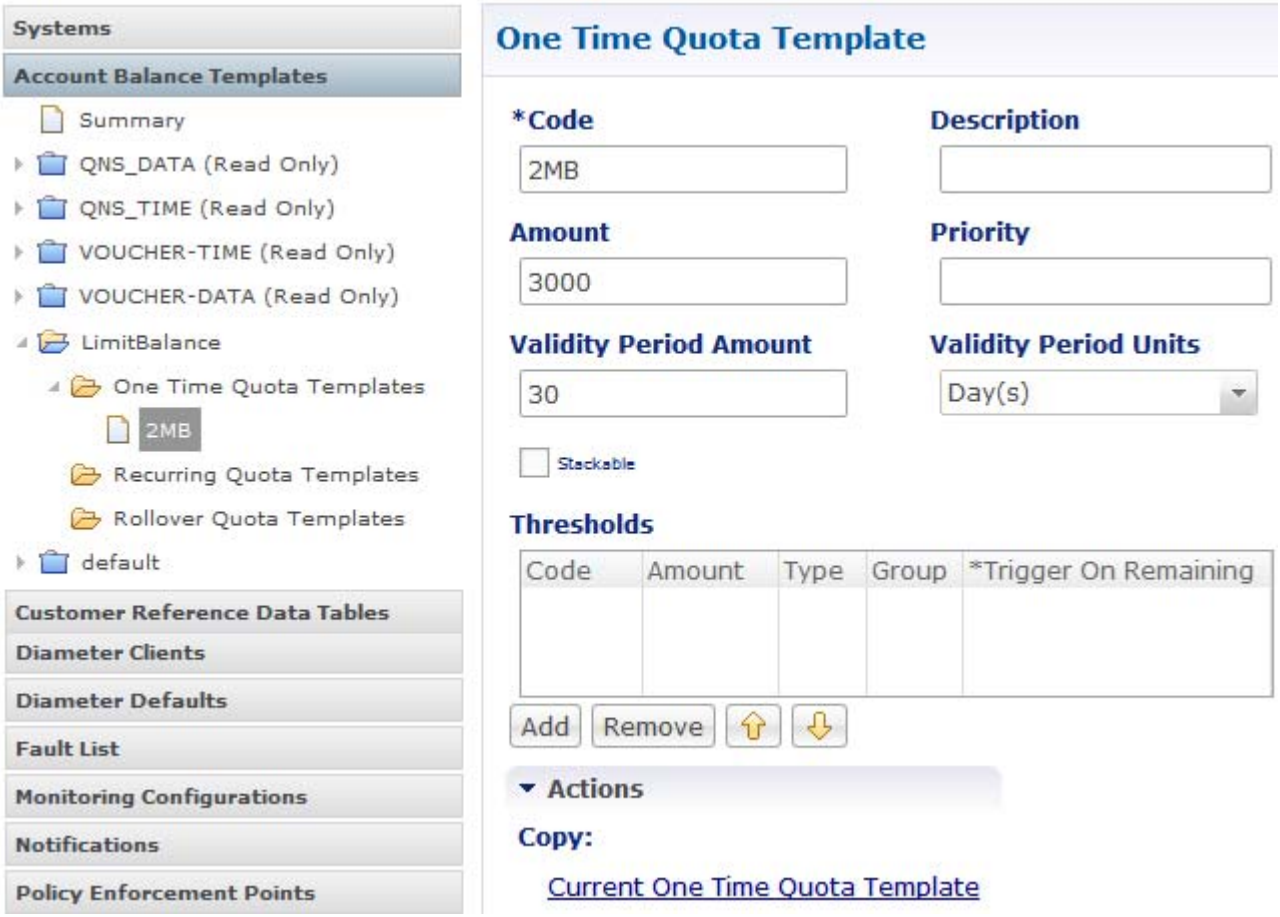

Step 5 Create an account to which the user can assign the quota limits created in Step 4.

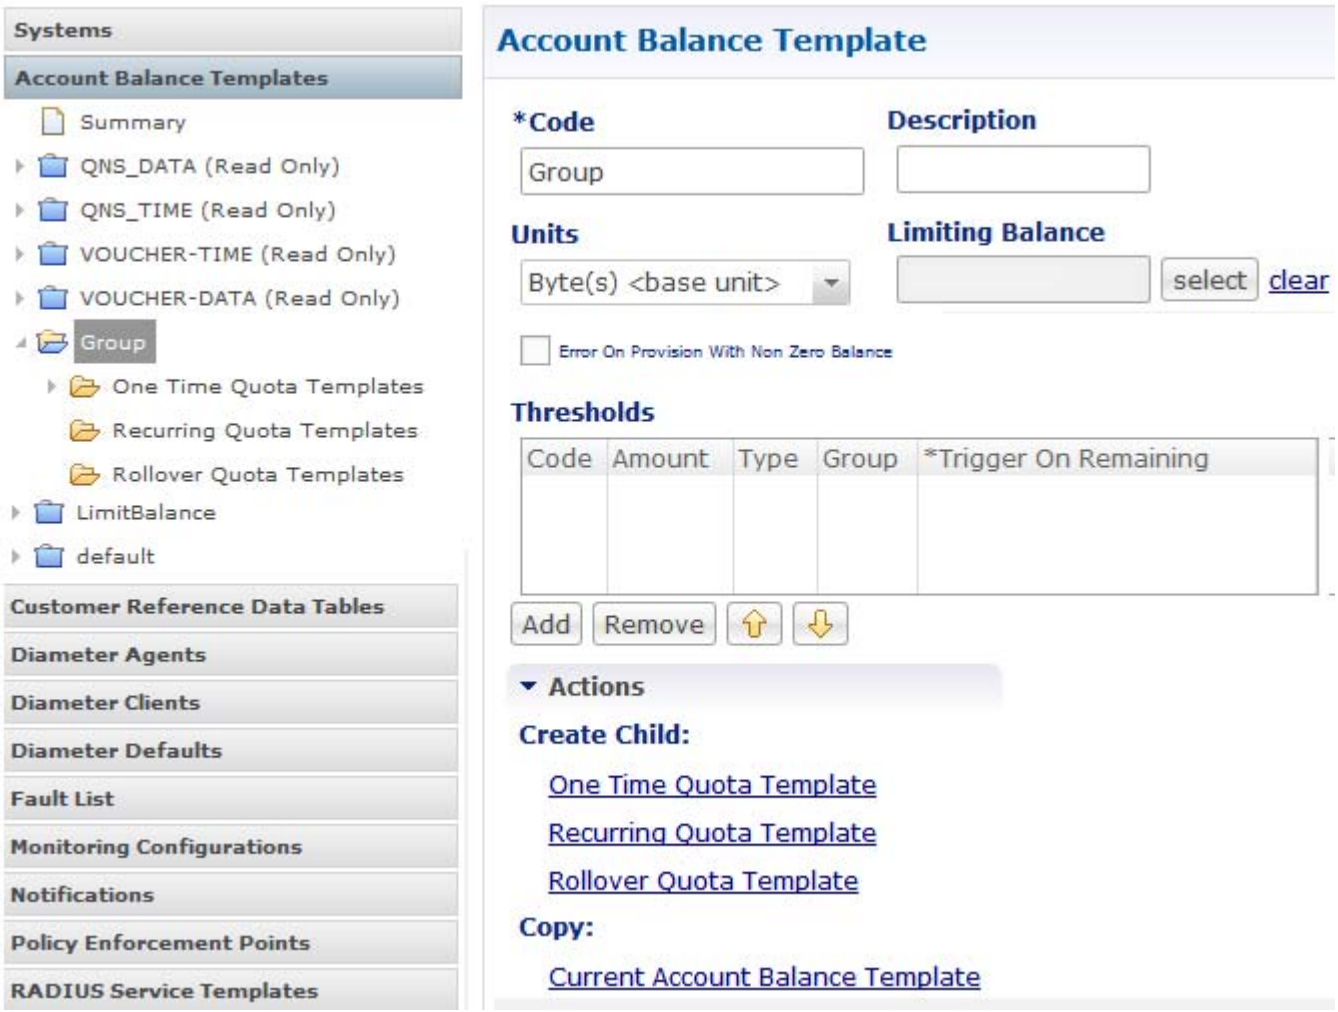

Step 6 Click select near the Limiting Balance field to open ‘AccountBalanceTemplate’ window.

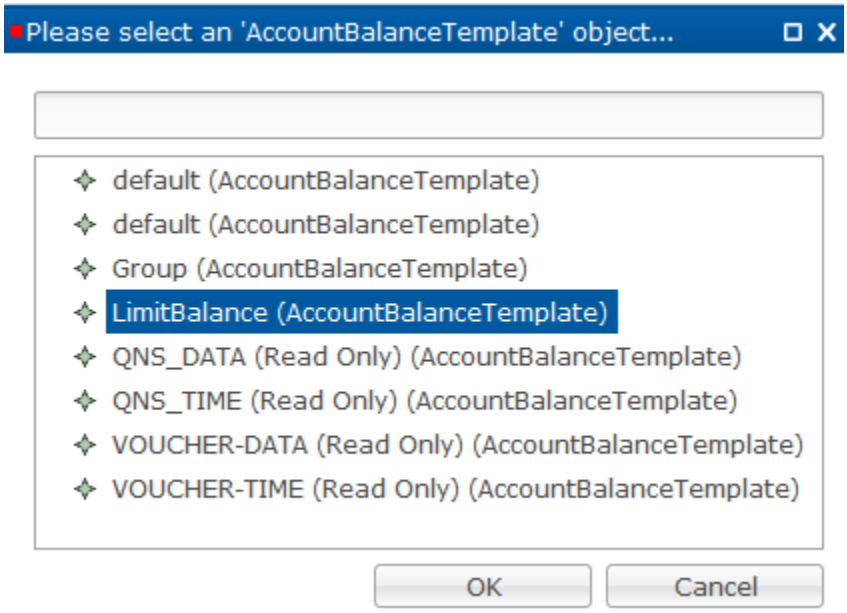

Step 7 Select the quota account created in Step 4 and click OK.

Step 8 By adding the one time quota template to group, subscribers in a group can use this quota as per their limit. An example is shown.

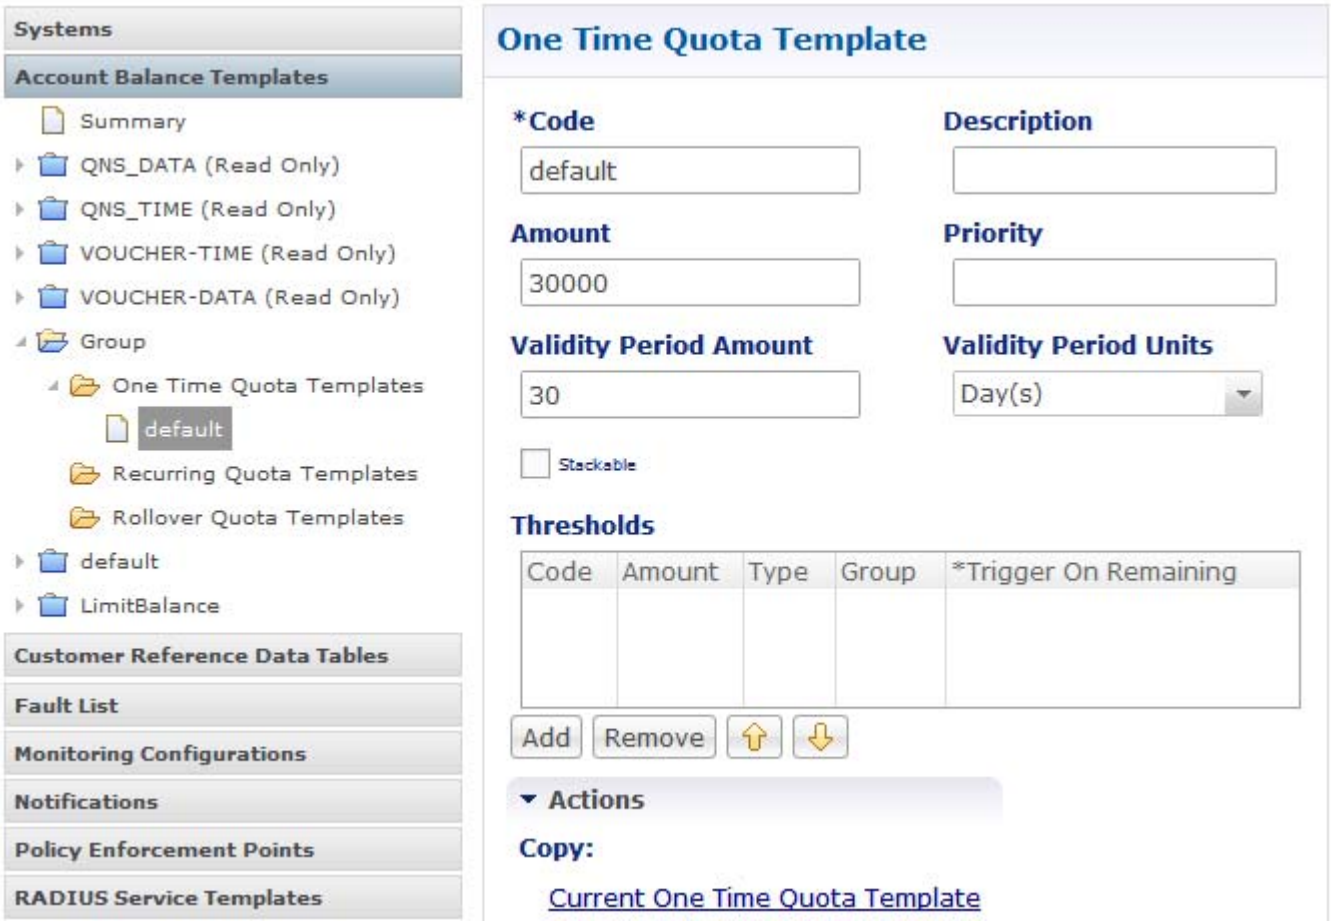

Note To apply thresholds, such as for notifications, as to when a certain amount of a limit is remaining or has been used, you can define the thresholds on the limiting balance and/or quota template as with any other balance/quota in Thresholds pane.

Thresholds

Thresholds are optional. Use thresholds to notify the policy engine of an approaching end of quota. The policy engine can then make a decision, notify the subscriber, or perform some other action.

The decisions for establishing thresholds are based on usage caps.

As administrator, you set limits as to how much quota, in this case data MBs, is consumed. At certain break points, or thresholds, you send a notice to the policy engine.

Set thresholds at the Account Balance Template level or at the one-time level or recurring level. At the one-time level, the policy engine is notified when a subscriber account uses 50 percent of a prepaid card, for example.

Thresholds can be used in numerous ways, that is, the first 50% of data might be charged at half off or perhaps charged double.

Thinking from the opposite perspective, if you set a low threshold, the policy engine can be notified quickly if a subscriber has logged in.

If you set a threshold for 1 MB, you can either provide more service or take away service after reaching the threshold.

Step 1 In the Account Balance Template screen, click the Add button to enter a new row of thresholds to the Thresholds table.

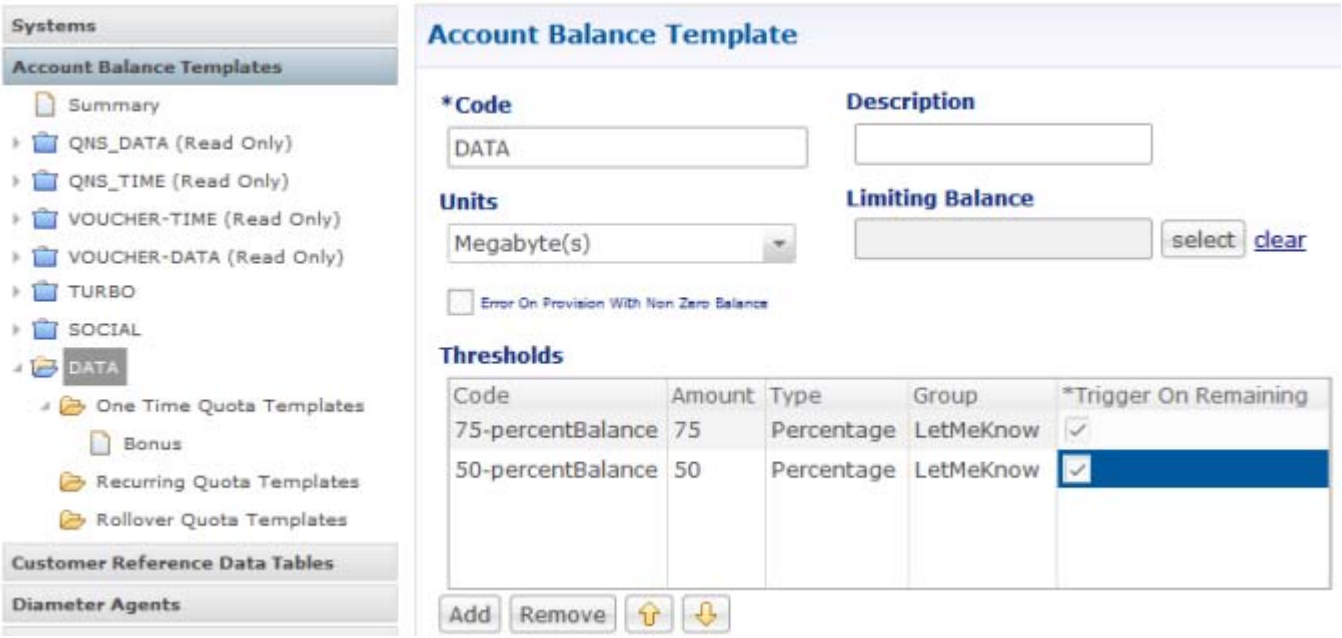

Interpret the rows as follows:

- If the subscriber uses 75 percent of the entire data quota, the policy engine is notified.
- If the subscriber uses 50 percent of the entire data quota, the policy engine is notified.
- Consider that if you breach the 75 percent threshold, you have most certainly also breached the 50 percent threshold and two messages would be sent to the policy engine.

To avoid that, we combine these two rows in a group.

Step 2 So, for these two rows, assign them both to the group LetMeKnow.

When thresholds are grouped, only the first threshold of the LetMeKnow group in the table generates a message to the policy engine. That is, whichever threshold occurs first in the policy group in this table is the one that sends the notification. Do not be fooled that the highest percent number generates the notification. It is the one that is *physically first*.

For example, the policy engine gets a notification when the 50 percent threshold is breached. If the subscriber is at 80% usage, the policy engine gets a notification only for breaching the 75% threshold.

You can order the thresholds with the yellow arrows and so control the order the rows are processed. Perhaps, you want to send a notification to the policy engine immediately, when a small threshold is reached. Move the smallest threshold to the top. This way you can check thresholds from the bottom up instead of top down.

The Thresholds table is now defined for two thresholds.

One-Time Quota Template

One-time quotas are amounts of time or data use that occur only one time and are not renewed. For example,

- If an ISP provider gives free introductory minutes
- If data bandwidth is provided on a tiered system, with the first 100 KB free, and then payment begins
- For a prepaid card use case

Note The three types of templates, One Time, Recurring, and Rollover are created for you. You do not configure them.

Step 1 Account Balance Template node, > Data node > OneTime Quota Templates > OneTime Quota Template link in the right pane.

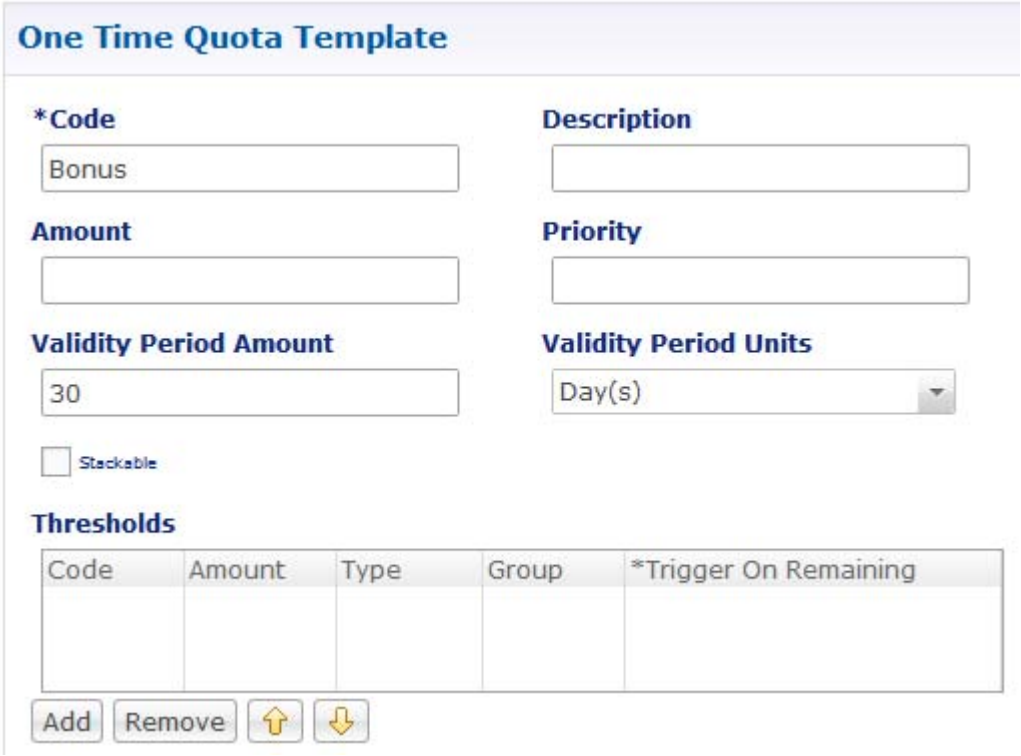

Code	The name you want to give to this template.
Description	Description of the one-time quota evaluation.
Amount	Enter the total amount you want to provide for subscriber use.

Priority	A subscriber may have several balances from which to draw. This number indicates which balance to use up first. If you do not specify a priority, Cisco MsBM checks the balances and uses the one that expires soon, optimizing quota use for the subscriber.
Validity Period Amount and Validity Period Units	These two fields work together to set the time frame for the one-time quota. Validity Period Units lets you select a unit of time from the drop-down menu. This field is required if you provide a Validity Period Amount. Validity Period Amount lets you specify the number of time units, 5 days, one billing period, 15 minutes.
Thresholds Table Area	The Thresholds table lets you specifically set the limit for balance management quota. When a threshold is reached, a message is sent to the policy engine.
Groups	This column lets you collect threshold rows and evaluate them as a group rather than individually, as rows.

Recurring Quota Template

Recurring Quota Templates are those services the account subscriber signs up for and gets refreshed at some recurring period. This procedure describes the first type only, a fully refreshed quota.

Step 1 Click Account Balance Template node > Data > Recurring Quota Templates > Recurring link in the right pane.

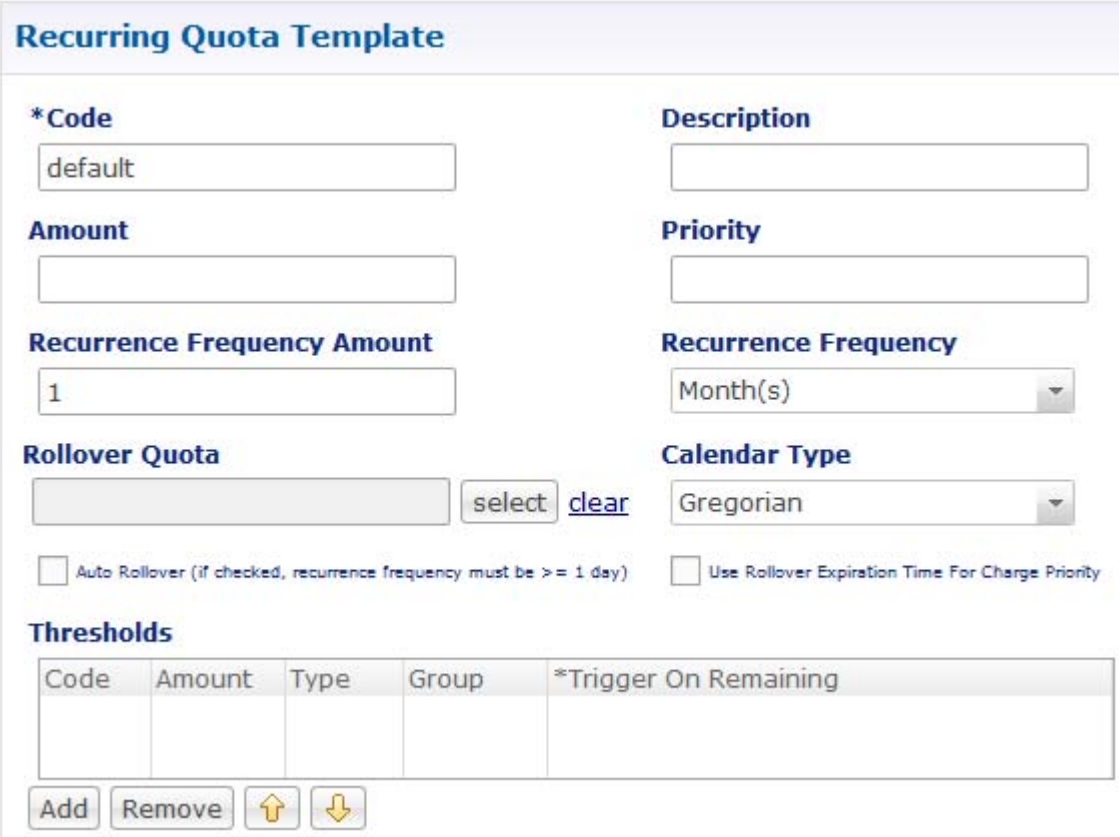

Step 2 Change the Code field to Monthly Recurring.

- Amount—what you want to provision the subscriber account with, give it 500 (MB) according to your business rules.

  This field is required to ensure that the refresh or the recurring frequency operates properly.

Note Because our parent Data template is specified in Bytes as the base unit, we need to set some type of byte measurement here.

- Recurrence Frequency Amount—how often to refresh. Type a 1 here, and also select Month in the Recurrence Frequency field pull down.

  Interpret this Threshold table this way: once a month, the service provider gives this subscriber account 500 MB of data.

Note Recall that in our business rules, there is no rollover of quota.

One month from today, it creates a new credit of 500 MB, which is available for one month.

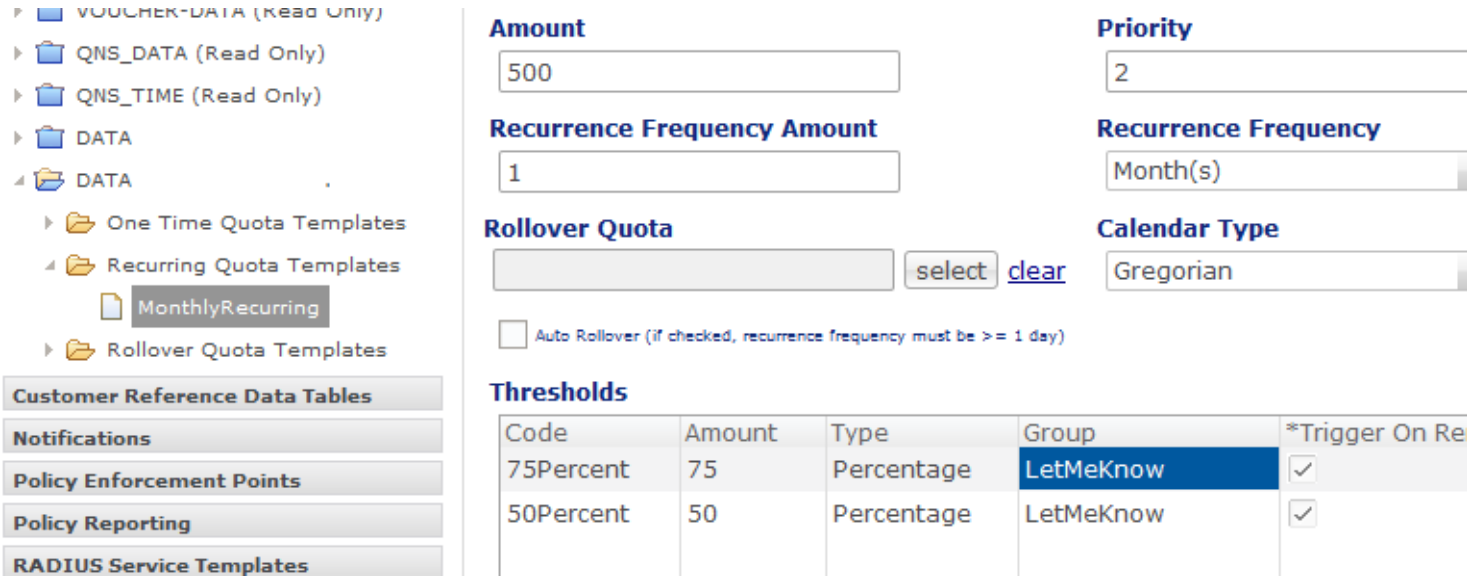

Code	The name you want to give to this quota evaluation.
Description	Provide a description of the recurring quota evaluation.
Amount	Enter the amount you want to provide.
Priority	The order in which you want to debit and credit buckets of quota. This is relative to the other defined quotas, both onetime and recurring, defined under the parent account balance.
Recurrence Frequency Amount and Recurrence Frequency	The two Recurrence Frequency fields work together to set the recurring period for the recurring quota. The Recurrence Frequency lets you select a span of time from the drop-down menu. During this time span, quota is used up and debited. At the point when the time span ends, the quota is refreshed to the amount in the Amount field. The Recurrence Frequency amount lets you specify the number of time span units. So 1 Month selected in these two fields means that quota is debited over a month's time. At the end of the month, the quota is filled up again.
Thresholds Table	These thresholds are calculated when actions that modify a subscriber's quota amount occur: • Debit • Credit • Charge • Notifications of breach • Continued breach • Debreach The results of calculation are sent to the policy engine, allowing actions to be taken on these events.

Tariff Switching Times

Tariff times, or more precisely, tariff switch times, let you specify when you want to change usage rates for a subscriber. The switch time can be a time of day, certain days of the week, or specific days during the year that require a special rate, either higher or lower.

Tariff times let you set a period of less expensive hours, savings for weekend use, or perhaps increase rates for business days.

The time of the tariff switch is at the end of the current tariff period as measured by the system clock.

To determine the current Tariff Switch Time, Cisco MsBM takes the current time (using the Time zone specified on the Tariff Time screen) and checks each switch time in order, from top to bottom, to see if the current time matches a Tariff Switch Time. The first tariff switch time that matches, including the associated valid dates *or* a holiday date, is used. The time of the tariff switch is the end of the current tariff period. The next tariff switch time is calculated, by taking the end of the current tariff switch time, adding one second, and searching each tariff switch time (top to bottom) to find the first one that matches.

The Tariff Time Identifier, (not Name) is expected to be used when determining rates. The Name is a human readable name and is not typically meant to be used directly by a system.

Tips About Tariff Times

CPS sets time-of-day tariff switches for quota expirations to minimize network signaling, especially for recurring balances that typically reset at the same time across subscriber groups. In addition, CPS supports the configuration of tariff boundaries to support rating changes. Configure Tariff Time switches with a combination of these timing elements:

- Time of day, for example
- 19:00-06:00 denotes night time hours, 7 p.m. to 6 a.m.
- 06:00-19:00 denotes day time hours, 6 a.m. to 7 p.m.
- Day of the week for which it is valid, Monday – Friday
- Holidays, Nov. 3, 2010, Nov. 23, 2010, Dec. 23, 2010

Other points to know about tariff times, which are specific to CPS, are these:

- Tariff Times are not allowed to cross over midnight. In practice, this may mean you have to create to two tariff switch times to cover a single logical period. For example 10 p.m. to midnight and midnight to 5 a.m defines your night time tariff time.
- A Start Time of midnight assumes it is midnight today.
- An End Time of midnight assumes it is midnight tomorrow (Start Time and End Time of 00:00 and 00:00 covers the whole day)
- If a Tariff Switch occurs during a daylight savings time forward switch (i.e. between 2:00 a.m. and 3:00 a.m. during 'Spring Forward' in March), an error occurs in processing during that time (2:30 a.m. doesn't exist that day).
- Best practice is to set switch times so that they do NOT occur during such times. Best practice is to add a (first) row in the table with holiday dates for these days if you must, creating different switch times for those days.

- If no tariff switch time is found for a given time, the next tariff time is not populated properly. It only starts working again after a message comes in and has a 'current' tariff time again. It is STRONGLY recommended that you define tariff switch times that handle all times and not have 'gaps'. You can always create a 'default' tariff switch entry as the last row if needed.

This is an example of a tariff time screen in CPS:

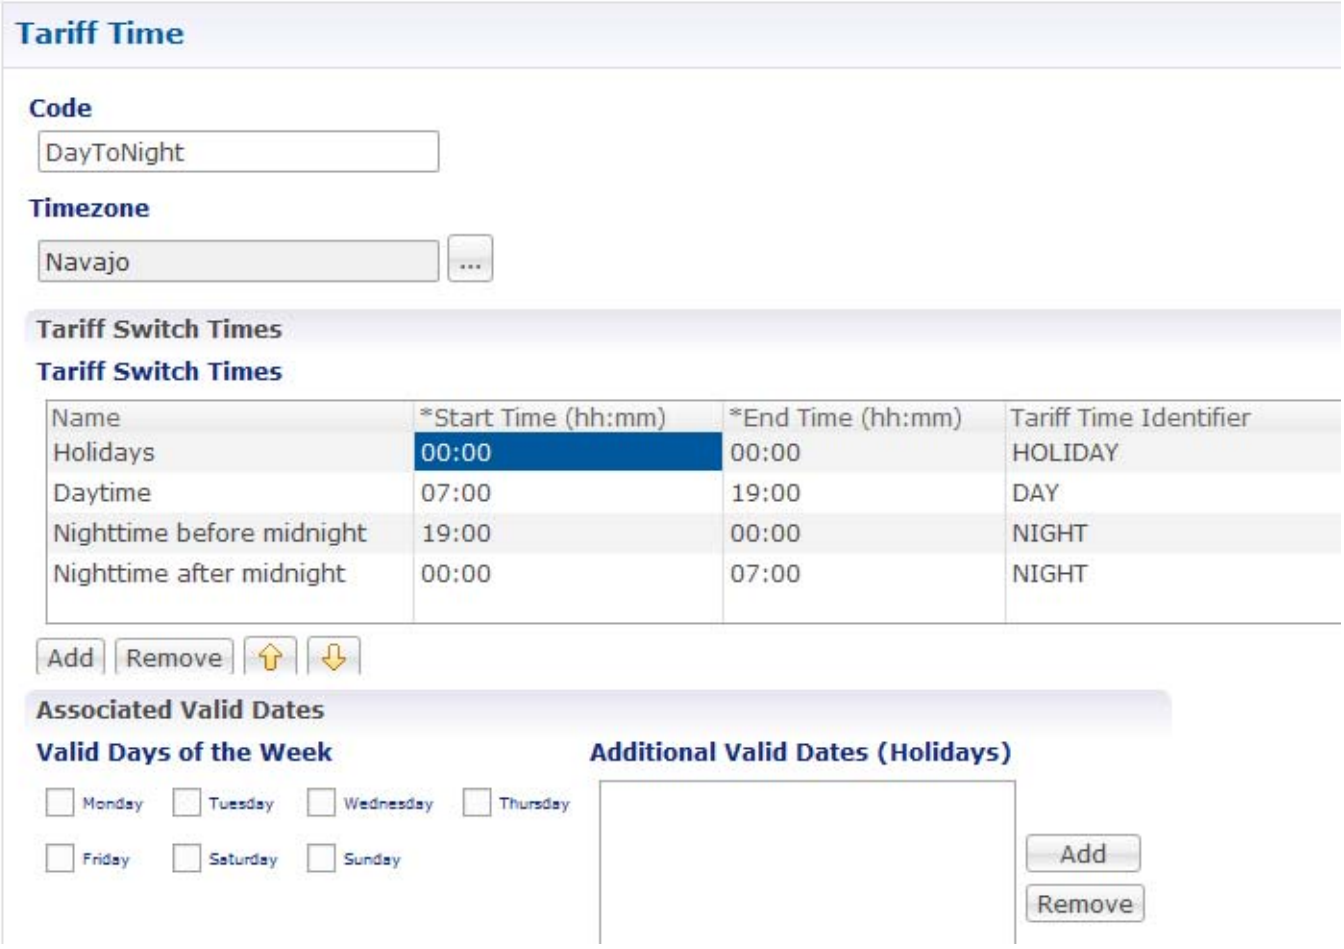

No Days of the week are checked at the bottom because this switch time is only concerned with day and night.

Setting Tariff Switch Times

Tariff times are when Cisco Policy Suite tells the router that at a certain time, something changes. The most common example is that at 7 p.m., we want to switch and use a different charge rate.

For example, on certain carriers, with a rate of 1, the subscriber may have daytime minutes from 7 am to 7 p.m., and with a rate of zero, after 7 p.m. the subscriber has unlimited minutes, a special nighttime rate.

There are a set of policies that belong to Cisco MsBM, such as Create Reservation, Charge Reservation, and so on. The Balance Request is generic, and knows which one of the base functions to call based on what fields are populated.

And, for mobile bandwidth the subscriber may get 100 MB per month, but from 7 p.m. to 7 a.m., the rate switches to one half of the daytime Tariff Time Identifier, encouraging people to use bandwidth at night when traffic is less.

Day of Week

You can associate days of the week to various tariff times, according to your business rules. For example, you may want to offer a discounted rate on Sunday night.

Step 1 Select the first row in the table, Day, and in the Associated Valid Dates area, select the days of the week you want the Tariff Time Identifier to be used.

In our example, we want the DAY Tariff Time Identifier to be used only for usage during the work week, Monday through Friday.

Step 2 Focus on the Valid Days of the Week area and deselect the check boxes for Saturday and Sunday. Leave the weekdays checked.

For our example, nighttime Tariff Time Identifiers are used during the work week as well. Recall that there is a different weekend rate, even for nighttime usage.

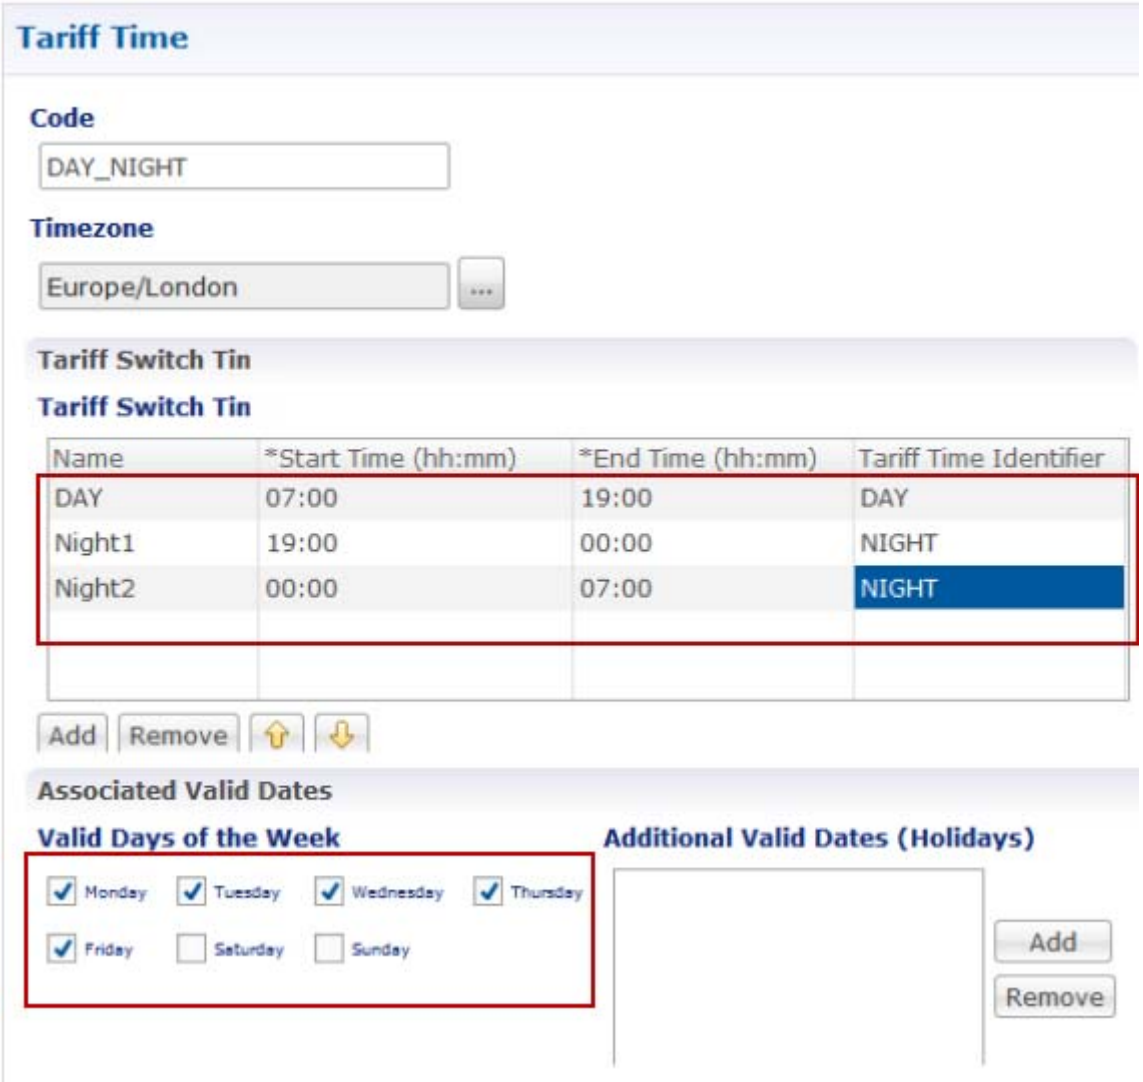

Step 3 Select the Night1 row and select the check boxes for Monday through Friday.

Step 4 Select the Night2 row and select the check boxes at the bottom for Monday through Friday.

With these three rows, you now have two tariff time set:

- During daytime hours, from 7 a.m. to 7 p.m., Monday through Friday
- During nighttime hours, from 7 p.m. to 7 a.m., Monday through Friday

Weekends

Our business rules state to apply a different rate to the subscriber on weekends, so switch the tariff for the subscriber for weekend use now.

Step 1 Click the Add button to add a row to the table.

Step 2 For the Name column, type Weekends, Start Time is 00:00, End time is 00:00, and the tariff Time Identifier is WEEKEND.

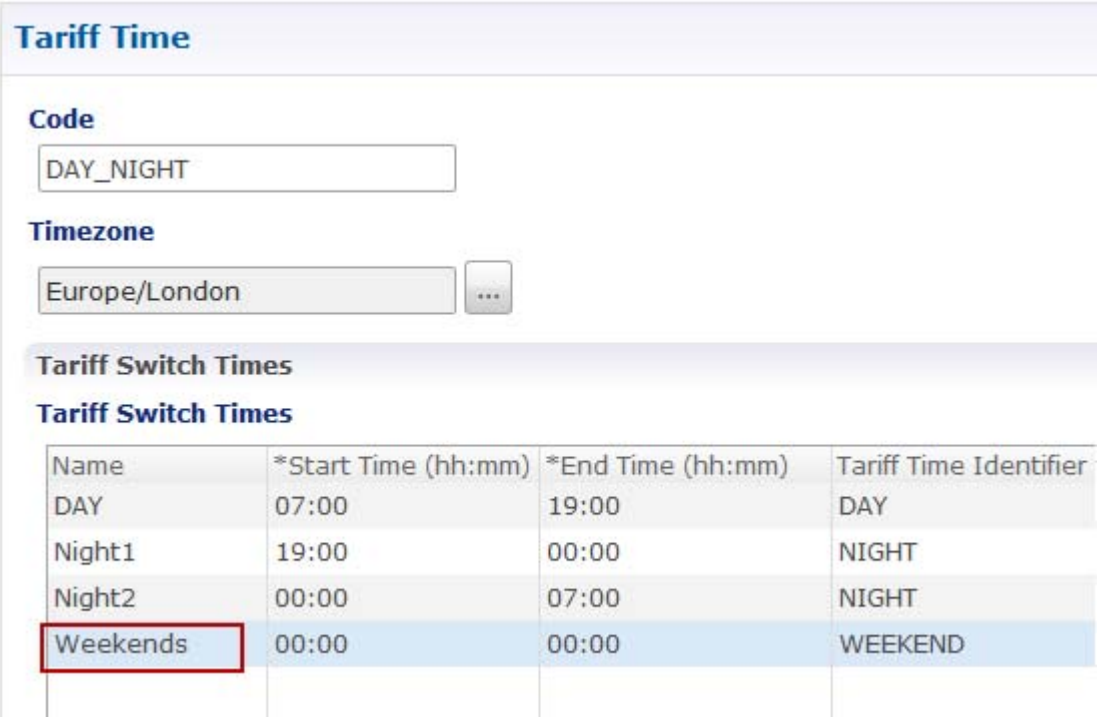

Step 3 For the Weekends row, in the Valid Days of the Week area, select only the Saturday and Sunday check boxes.

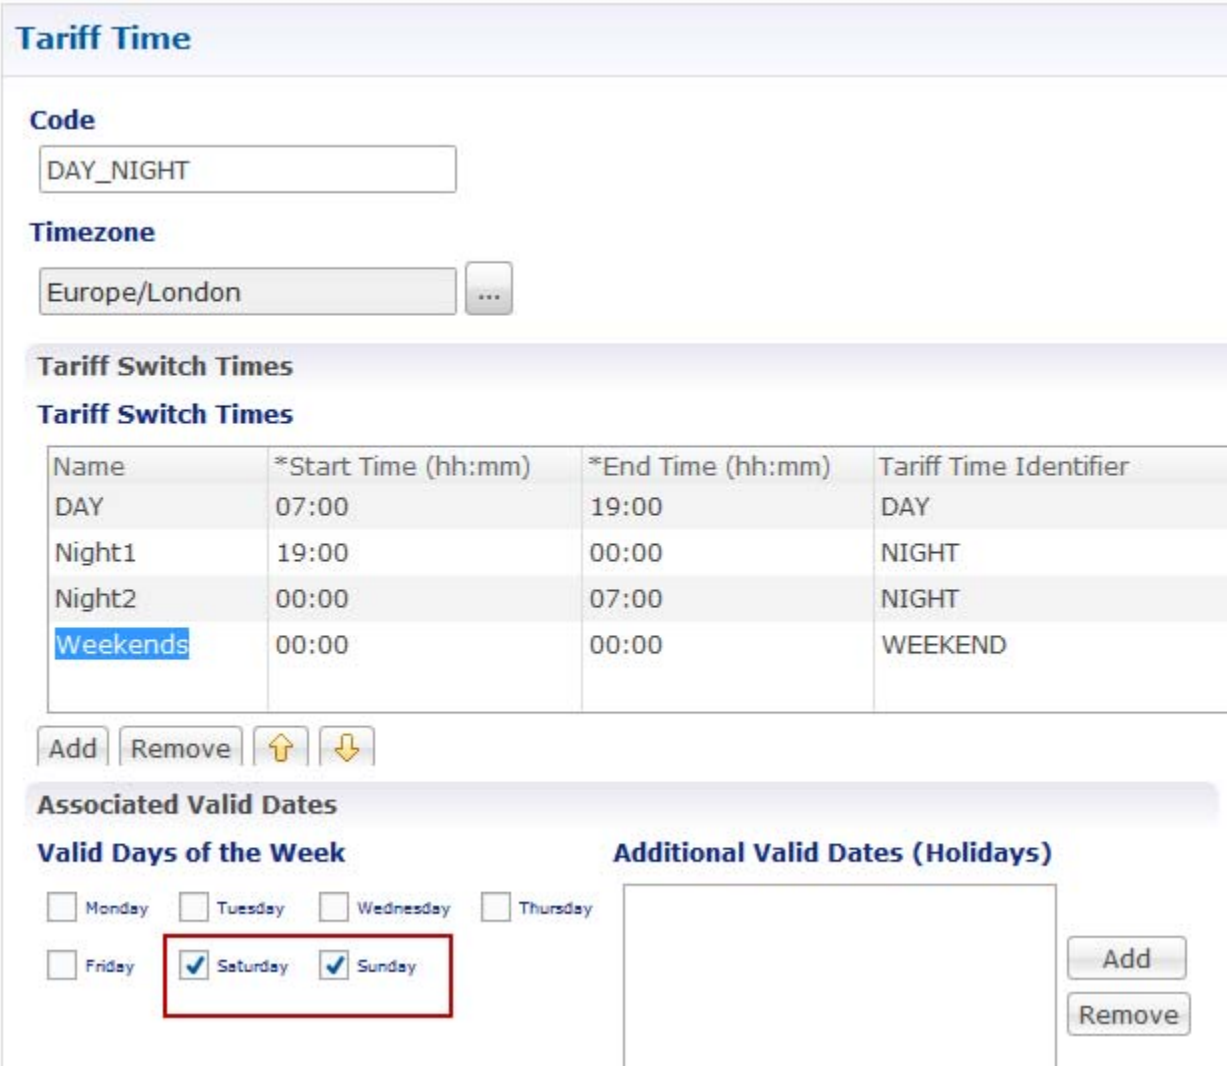

Note Remember that the table rows are evaluated from first row to last row. If it is Tuesday 8:30 p.m., the evaluation stops evaluating on the Night1 row and never encounter the Weekends row. If it is Saturday 2:00 p.m., the table evaluation proceeds until it reaches the last row and executes that.

Holidays

In addition to days of the week, Cisco MsBM can specify holidays during the year to have a tariff switch.

Now we can define a holiday tariff of 0 percent of the Day usage rate. That is, holiday usage is free. Our holiday is March 17.

Because of holiday pricing, the subscriber does not use daytime minutes on March 17, even if it is a Wednesday. In fact, your business rules require two holiday tariff switches, one for Holiday Day, and one for Holiday Night.

Step 1 Click the Add button to add a row to the table.

Step 2 In the Name column, type Holiday Day. For Start time put 7 a.m., for End time put 7 p.m.

Step 3 For the Tariff Time Identifier put HOLIDAY-DAY all caps.

Step 4 In the Valid Days area, select no check boxes. There is no way to predict which day of the week the holiday occurs.

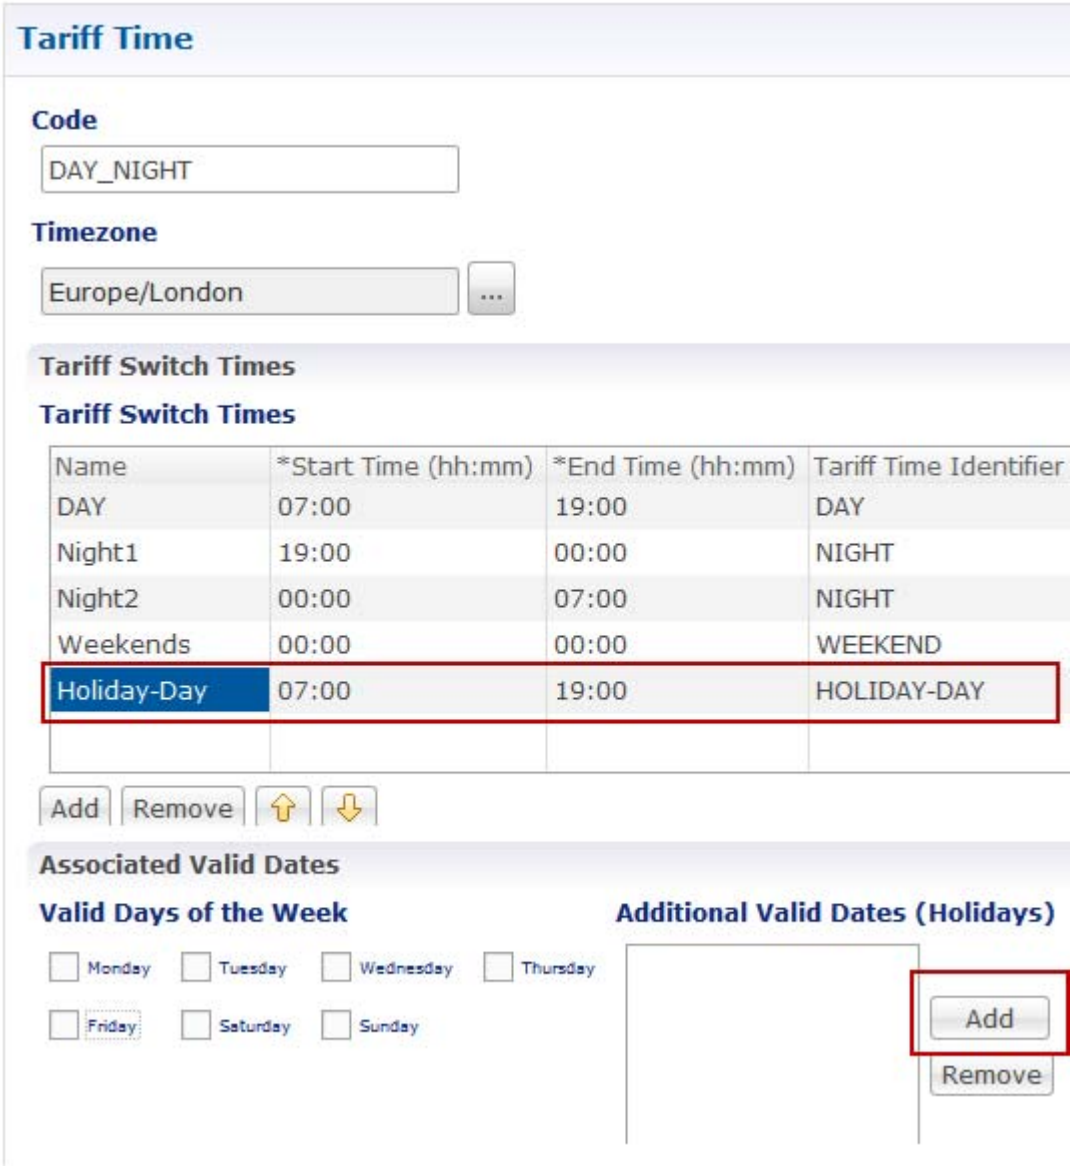

Step 5 In the Additional Valid Dates (Holidays) area, click the Add button to the left and display the Add Values window.

Step 6 Click the down arrow in the Values to Add field and display the calendar.

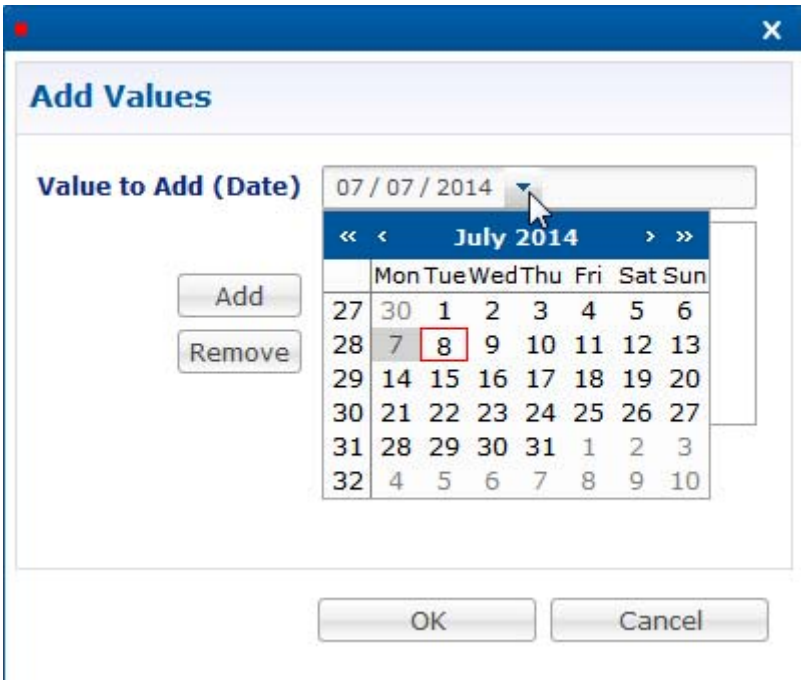

Step 7 Use the arrows in the calendar to display the calendar date you want to specify as your holiday. For our example, we select July 11, 2014.

Our Add Values screen looks like this so far.

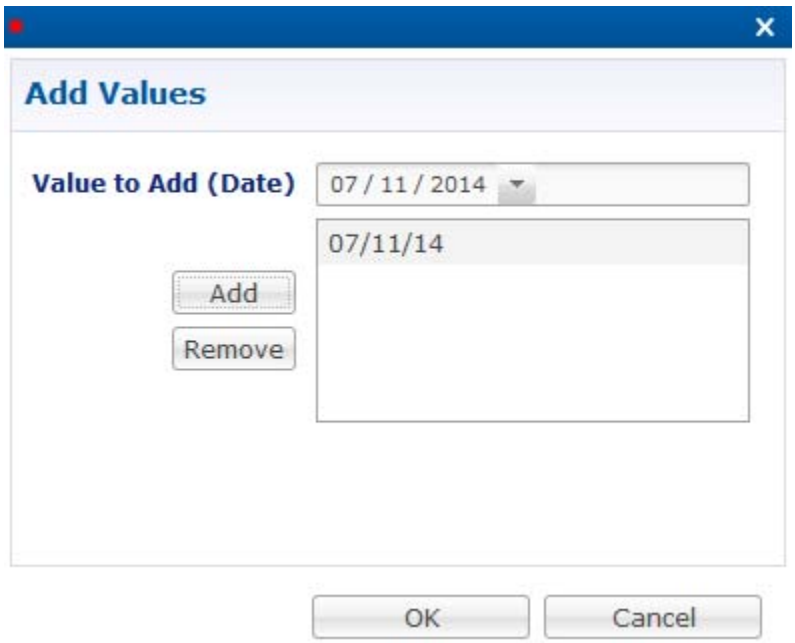

Step 8 Continue and add another holiday.

Step 9 Now add the row for Holiday-Night. Name is Holiday Night. Start Time is 00:00 and End Time is 00:00. Tariff Time Identifier is HOLIDAY-NIGHT, all caps. Our table now looks like this:

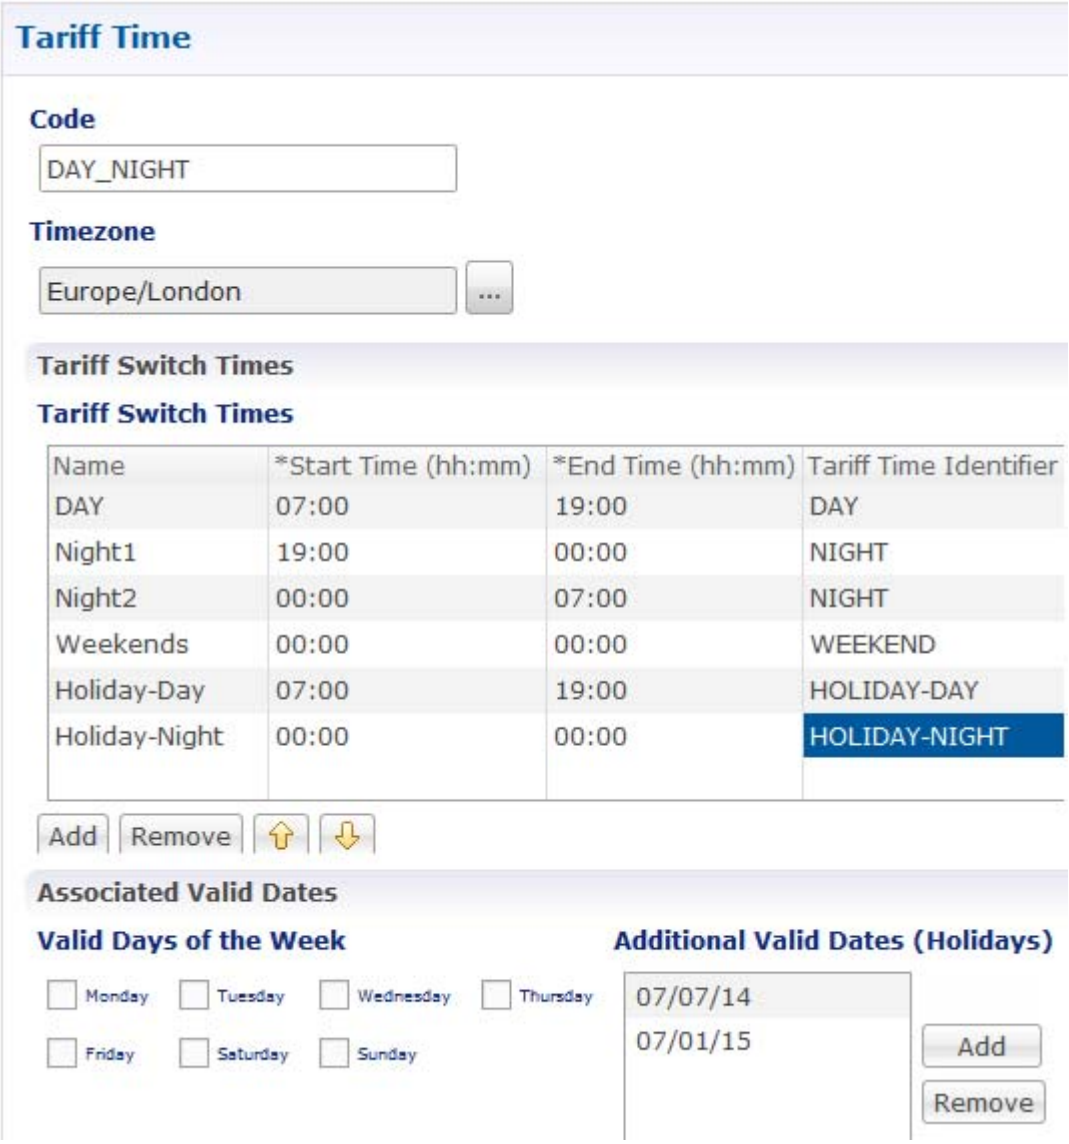

For this holiday, the table is evaluated this way:

- The table evaluates, beginning with the first row.
- The current time is 9.31 a.m. on a Friday, August 13th, 2010.
- Is the current time after the Start Time? Yes.
- Is the current time before the End Time? Yes again.
- Is today one of the Valid Days of the Week or one of the designated Additional Valid Dates? YES, it certainly is.

Note This case helps describe the order of evaluation in the table. First, we can casually set the Start and End Time values as 00:00 and 00:00. Second, the first row is evaluated with the 7 a.m. to 7 p.m. parameters, no day of the week specified. This row evaluates as FALSE. Our case matched the time frame, but not the day of the week specification. The second row is evaluated, and the result is again FALSE. Processing continues down the table until a row evaluates as TRUE, in our case, the third row, DAY. No subsequent rows are examined.

Setting Up Rates and Ratings

Setting Rates and Ratings is optional but is often a best practice, especially when dealing with tariff switch times.

Definition of Rates and Rating

Rating are associated with tariff times. That is, when a tariff switch time occurs, rates determine the change in billing rate for quota used at that time.

Here is an example.

During the day the subscriber gets charged 100 percent of the rate. During the night the subscriber gets charged 50 percent of the rate. On weekends, the subscriber is charged 25 percent of the rate.

Accounting and Quota

Recall that Cisco MsBM does not deal in currency. For the above scenario of business rules and tariff times, there is no monetary billing. Rather, during the day CPS charges quota for all the quota the subscriber uses, one hundred percent of it. During the night, CPS charges the reservation for only half of the quota the subscriber used. And, yes, on the weekend, CPS charges the reservation for only one quarter of the quota the subscriber used. Accounting an receivables are left to others.

If you have tariff times configured, and messages start coming in to the Balance Blueprint, Cisco MsBM creates balance rate objects for you to put balance rates on. This is how Cisco MsBM associates a tariff time, with a rate. That affects the balance requests themselves. When rating is used, Cisco MsBM sends out a message to the policy enforcement point (PEP, or network access device) saying there is going to be a tariff switch at 7 p.m. and you need to keep track of the bytes used before 7 p.m. and after 7 p.m. That way, if the subscriber is in a session at 6.45 p.m. and requests another hour, the PEP calls back Cisco MsBM at 7:45, not at 7:00, the tariff switch time, along possibly millions of other people.

If Cisco MsBM were trying to put the rates on the requests themselves, because they have a rate on them, Cisco MsBM would require duplicate decision tables, one for before 7p.m. and one for after. This would result in a confusing duplication of tables.

To avoid creating duplicate tables, two Balance Rate objects are created by the Balance Blueprint. Both objects use the decision table to populate their Balance Rates. One Balance Rate object is the current time. The other Balance Rate object hold the upcoming tariff time. Having both objects makes processing proceed smoothly.

Checks on the Tariff Time

The Cisco Policy Suite checks to see what the tariff time is, and it corresponds it to a rate, and the tariff time. The first row is DAY. Then Cisco MsBM asks: what is the next tariff time? Answer: 7 p.m. At 7:01 it is Night. The Balance Required Id is still 1, so all that is set there is the rate. By using a table we state the Tariff time ID, of type String, and its rate, big decimal. Here, we set up how we want the usage to be rated. The rate for Day is 1, and for Night the rate is `.5`, and weekend is `.25`. Holiday day is `0`, and holiday night is `0`.

If there is no entry, CPS does not populate the field and it ends up defaulting to `1`.

The basic steps for rates are these:

- Tie the rate to the tariff time.
- Tie the tariff time to the conditions.
- Add this to the balance.
- Modify the rate.
- Select the rate.
- Set the requested rate.

This way, for each balance request there are two rates, one for the current time, and one for the next time, the one after a tariff switch time.

Knowing these rates, we can evaluate both times based on the time and length of the subscriber's session. This way, all the subscribers in a session at the tariff switch time do not need to be refreshed at the same time, 7:00 p.m.

Rates and the PEP

The rate itself is set in the policy enforcement point or PEP. Cisco MsBM provides megabyte intelligence to the PEP. Cisco MsBM notifies the PEP that 100 MB are available to a subscriber for download, and Cisco MsBM is debiting 100 MB. When that 100 MB of quota depletes, the PEP calls back and notify the Cisco MsBM that the subscriber has used up their 100 MB.

Recall that the PEP deals in traffic. The rate modifies how much of the traffic Cisco MsBM debits. In the case of Night, if the subscriber downloads a 100 MB file, the PEP calls back, reporting that the subscriber downloaded a 100 MB file. Cisco MsBM checks and evaluates the time period as Night, and so only debits 50 MB.

If quota, from a business perspective, is understood as currency, you have a way to devise a rate that relates bytes to currency.

Setting Up Ratings

If you need to change the rates or ratings that were implemented at install time, please call your Cisco technical representative for help.

Changing the Party Billed

This step is optional, but this section shows you how to set up family plans or corporate plans. In these types of plans, one person may incur usage, but that usage is charged to the account, not the session user. For example, in a family plan, child Mike may use the Internet and consume data, but his usage is charged against the Januski family plan.

For example, when using an SCE there is an extension point called Quota Charging ID. For example, a Subscriber Id is a component of an overall family plan. In this use case, Mike is a subscriber under the family plan Januski. We want a policy that directs Cisco MsBM to send all Mike's requests and usage to debit against the account Januski's subscriber ID instead of his own. He may not even have or need to have a CPS account.

This means all of the balance requests coming in from Mike must be debited against the Januski account. Cisco MsBM requests that the Subscriber Id gets charged. The service info attribute is stored in SuM, and checks that it is set to the correct value.

Cisco MsBM takes the Subscriber Id from SuM, and puts it on the balance request Subscriber Id, so by the time it gets to Cisco MsBM, Cisco MsBM has the correct Subscriber Id.

Setting Up the Change ChargingID

By default, a policy uses the subscriber ID as identified on the session. This is what must be changed.

Whether the subscriber ID came from SuM or any other SPR, CPS is going to look on the session and look for an AVP that says "Do we have Charging-Id information on our session?"

If the subscriber ID has been set up from the way that you loaded the session, CPS can look that up and say "Do we have Charging-Id data on the session?"

If so, CPS can then modify the balance request action phrase to set a different subscriber's ID to be the one that gets charged.

For example, Jane logs in with a subscriber's ID of jdoe. On the session however, CPS has marked Jane's session as being charged to Cisco. This way, CPS can charge quota to Cisco, not jdoe.

This changes the Charging ID before it is sent into the balance manager to reserve or charge data.

Step 1 In the Policies tree, click SCE Quota Blueprint > Configured Extension Point > Change Charging ID.

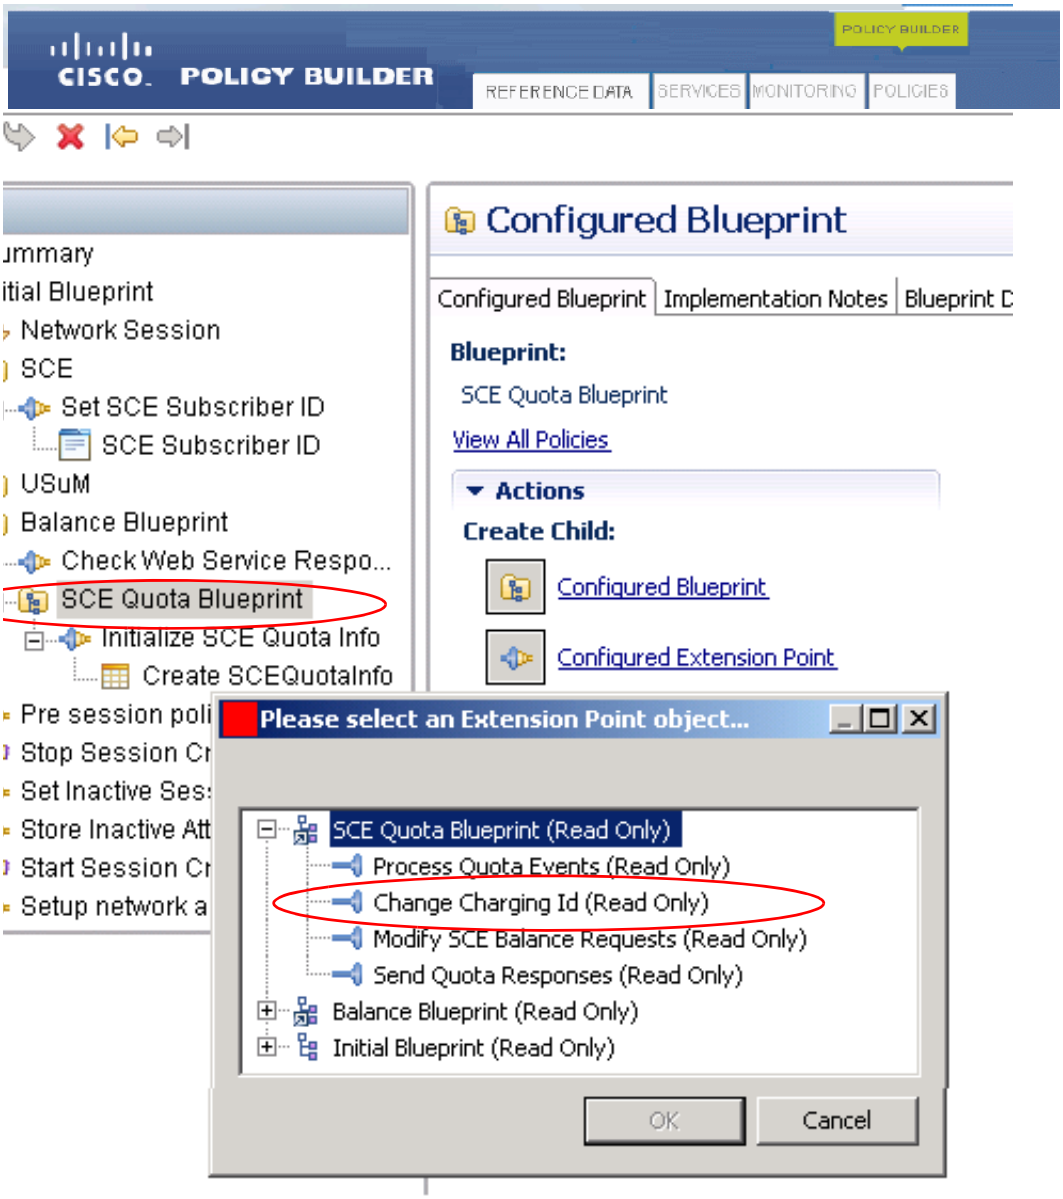

Step 2 The Policies tree now has the ChangeChargingId node in it. Select that.

Step 3 This policy requires no decision table, just click the Policy link.

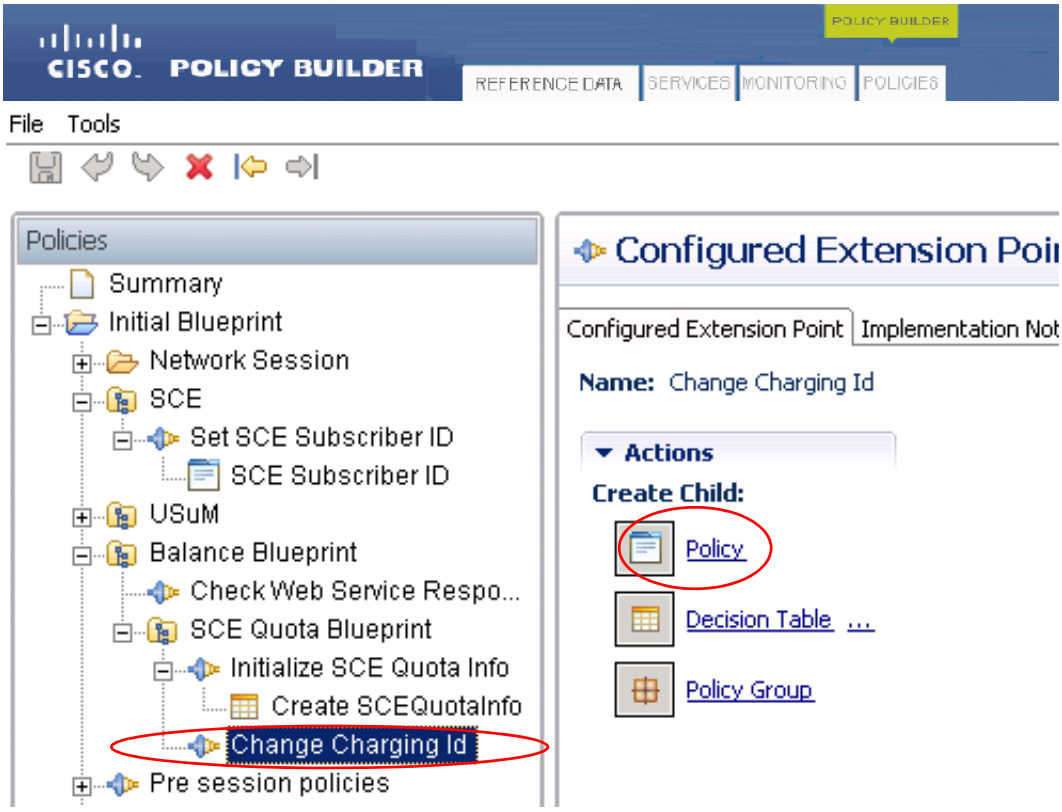

The Policy screen appears.

Step 4 Name this policy Change Charging ID.

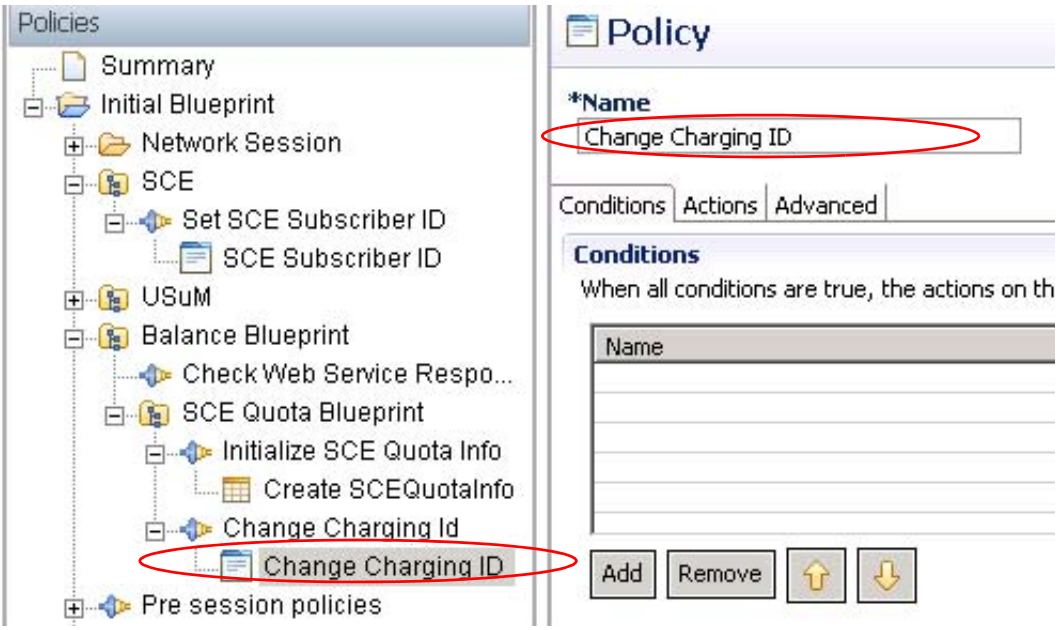

This policy name is reflected in the Policies tree.

Step 5 Under the Conditions tab, click Add and insert the condition phrase. There exists a Balance Request.

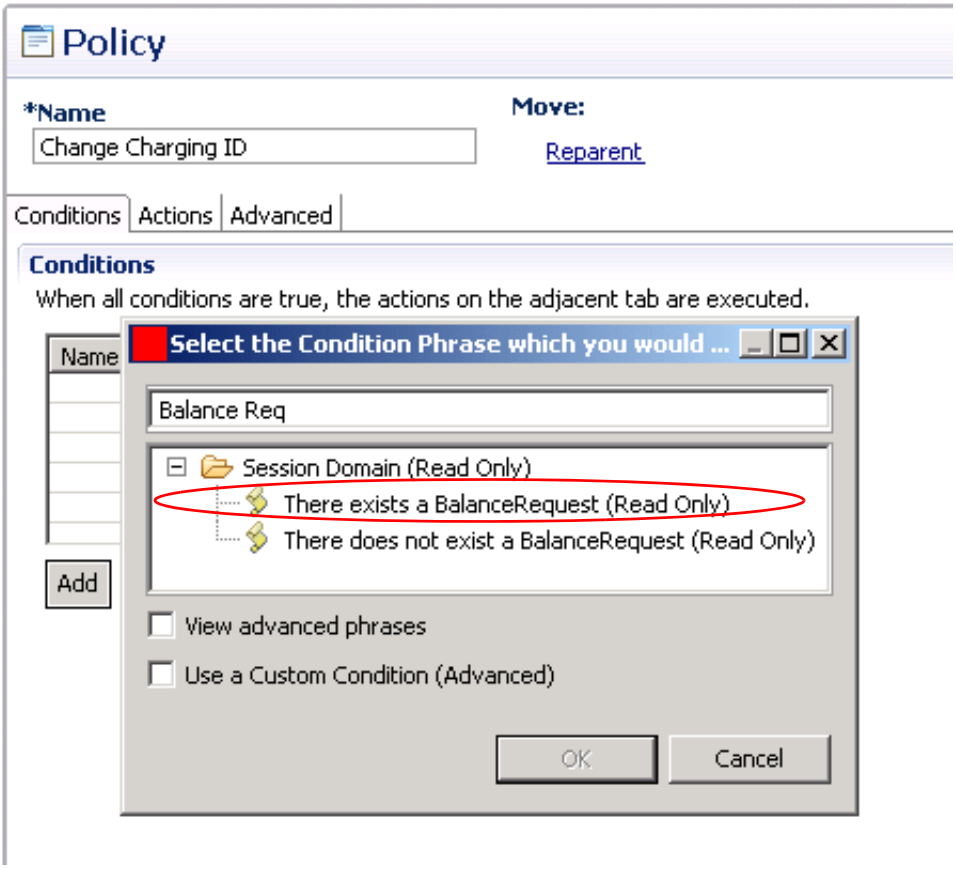

Step 6 Select this phrase in the Condition table to display its variables.

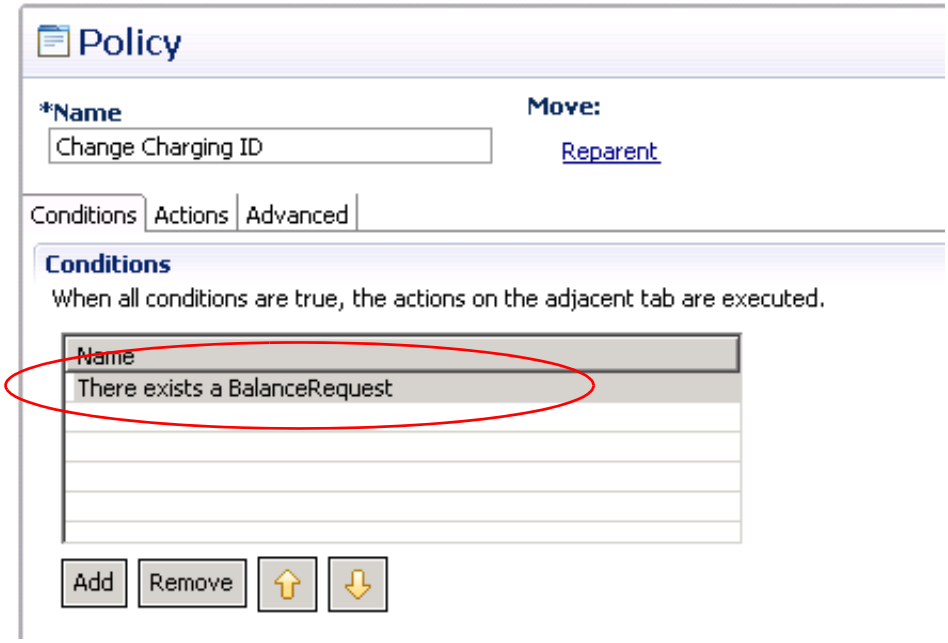

Step 7 Add the variable Request Subscriber Id, leave it of Type Literal. We do not care about the value of it, so do not enter an operator or value.

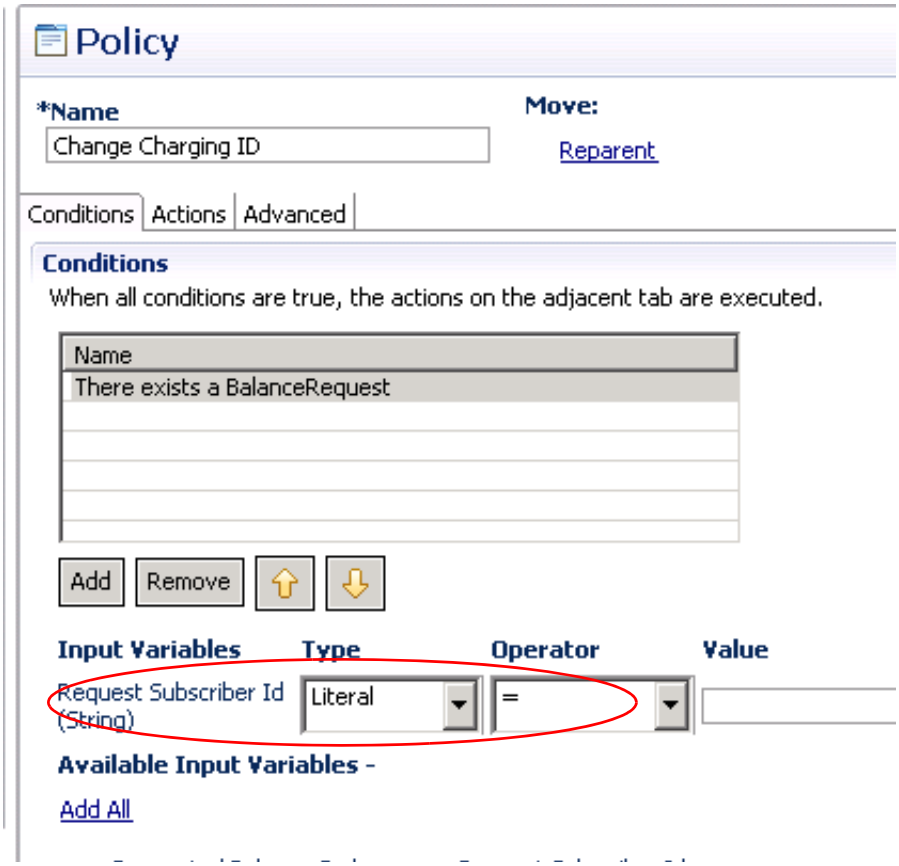

Step 8 Still in the Conditions tab, add the additional condition There exists an Avp.

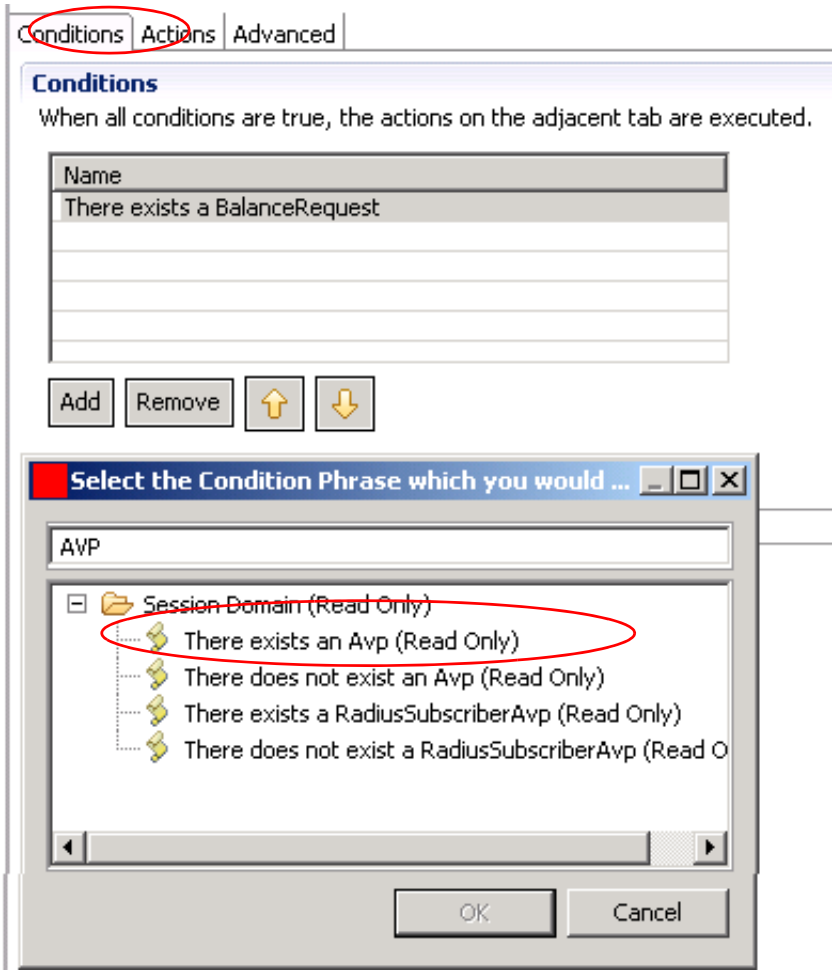

There are now two conditions that this policy must meet.

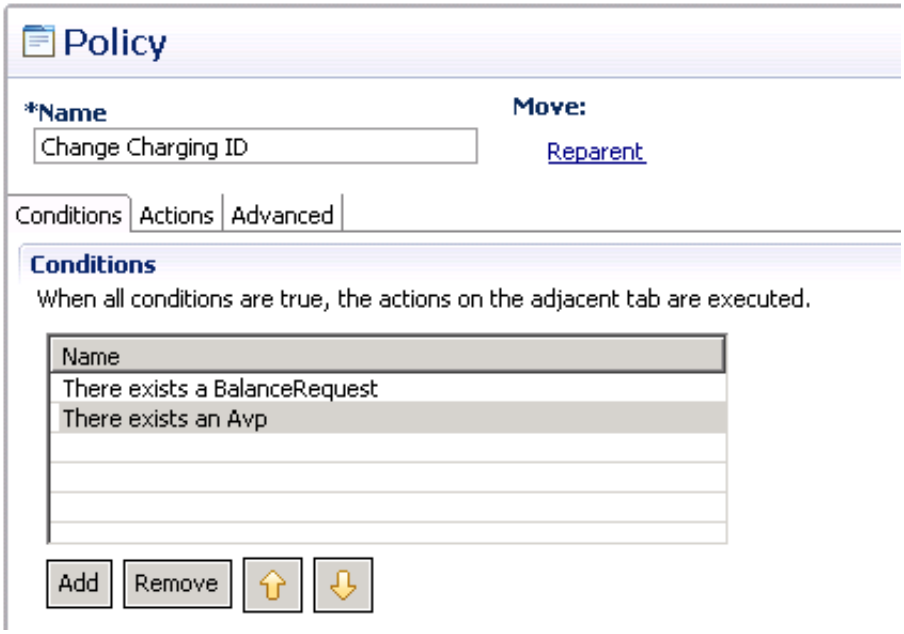

Step 9 Select the condition There exists an Avp, and add the Code variable so that we can manipulate its required variable, Code.

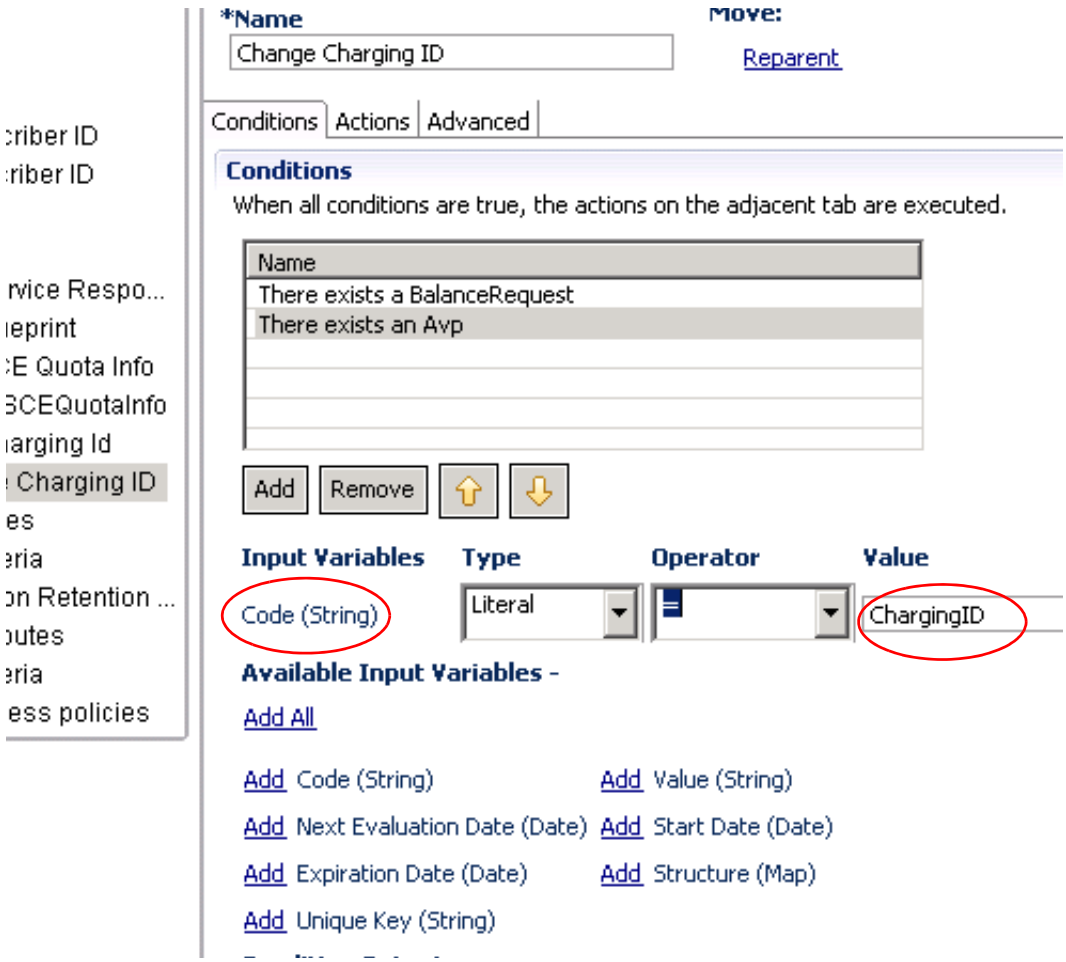

Step 10 Leave the Code as a Literal.

Step 11 Change the Value to be the value of the AVP, Charging-Id.

Step 12 Click the Actions tab for the policy Change ChargingID.

Step 13 Click the Add button to add an Action phrase Modify a Balance Request…

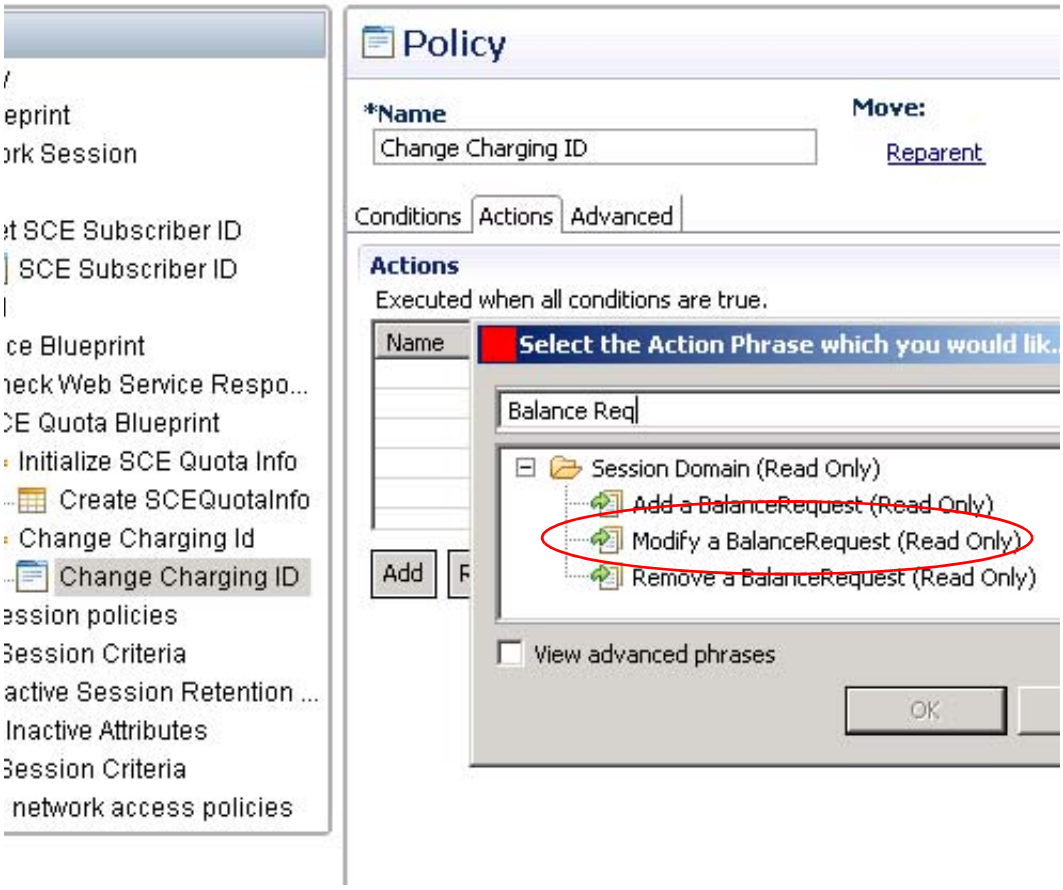

Step 14 Select the phrase in the Actions table and display its variables.

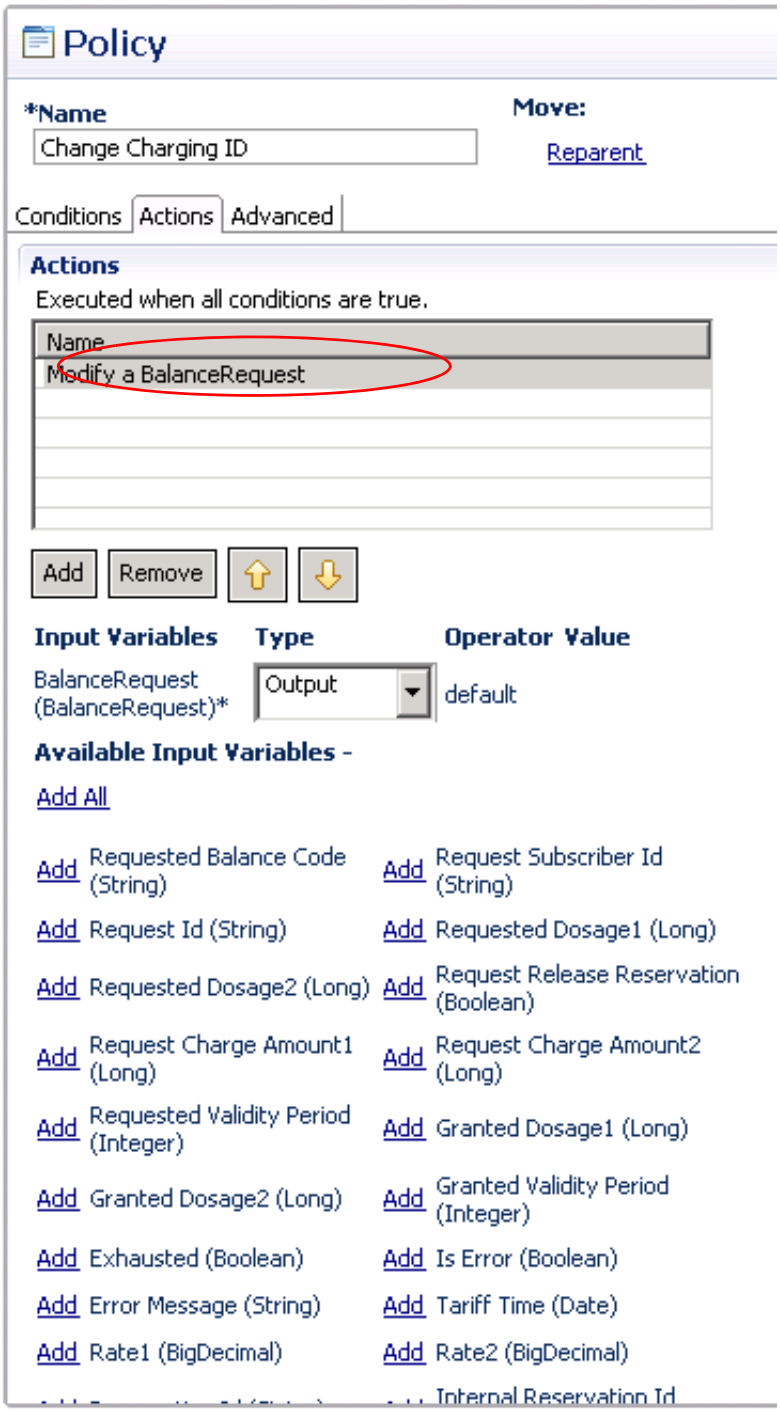

Step 15 For the required BalanceRequest variable, click Output in the drop down list for Type.

Step 16 From the Output variables list, select the one present, Output Variable:balanceRequest.

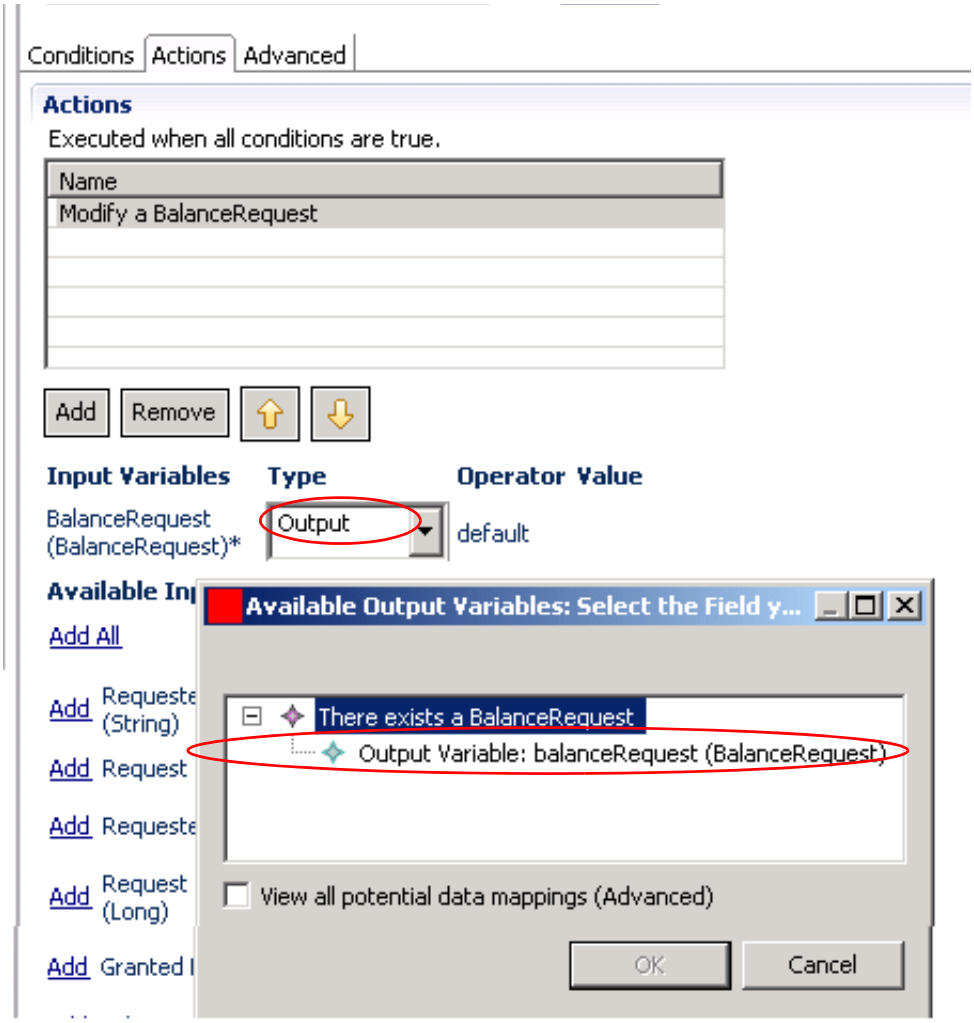

Step 17 Click to Add the variable Request Subscriber Id.

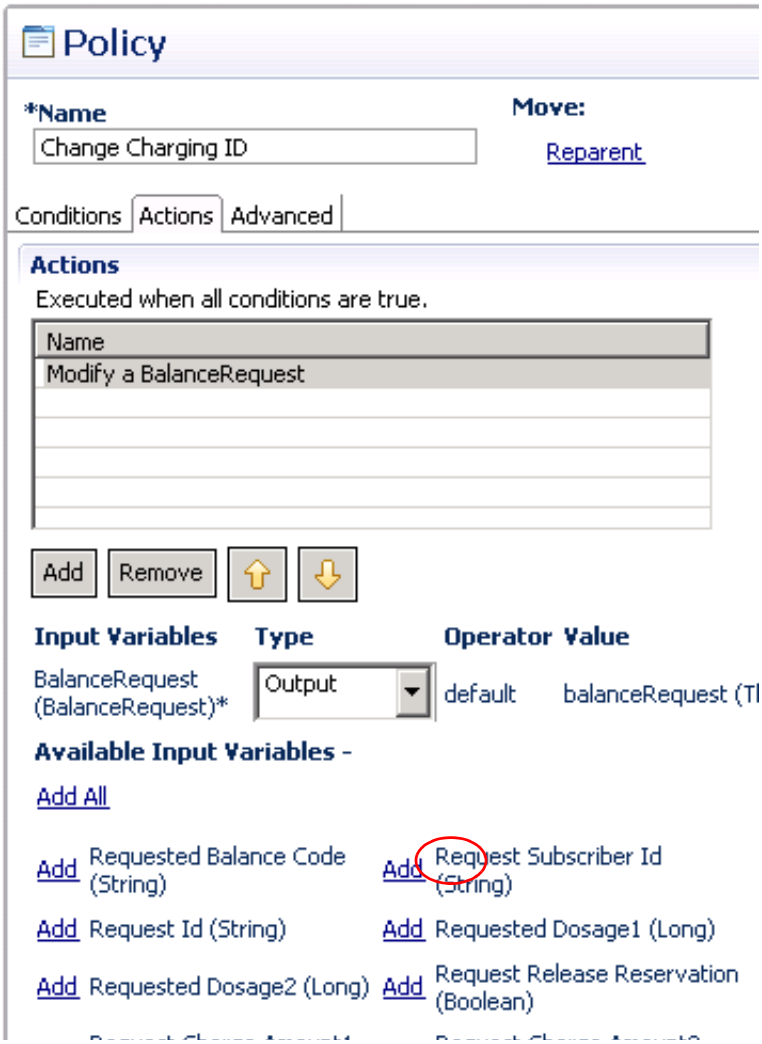

Step 18 Select Output in the Type drop-down and select Value under the AVP portion of the tree in he Output list.

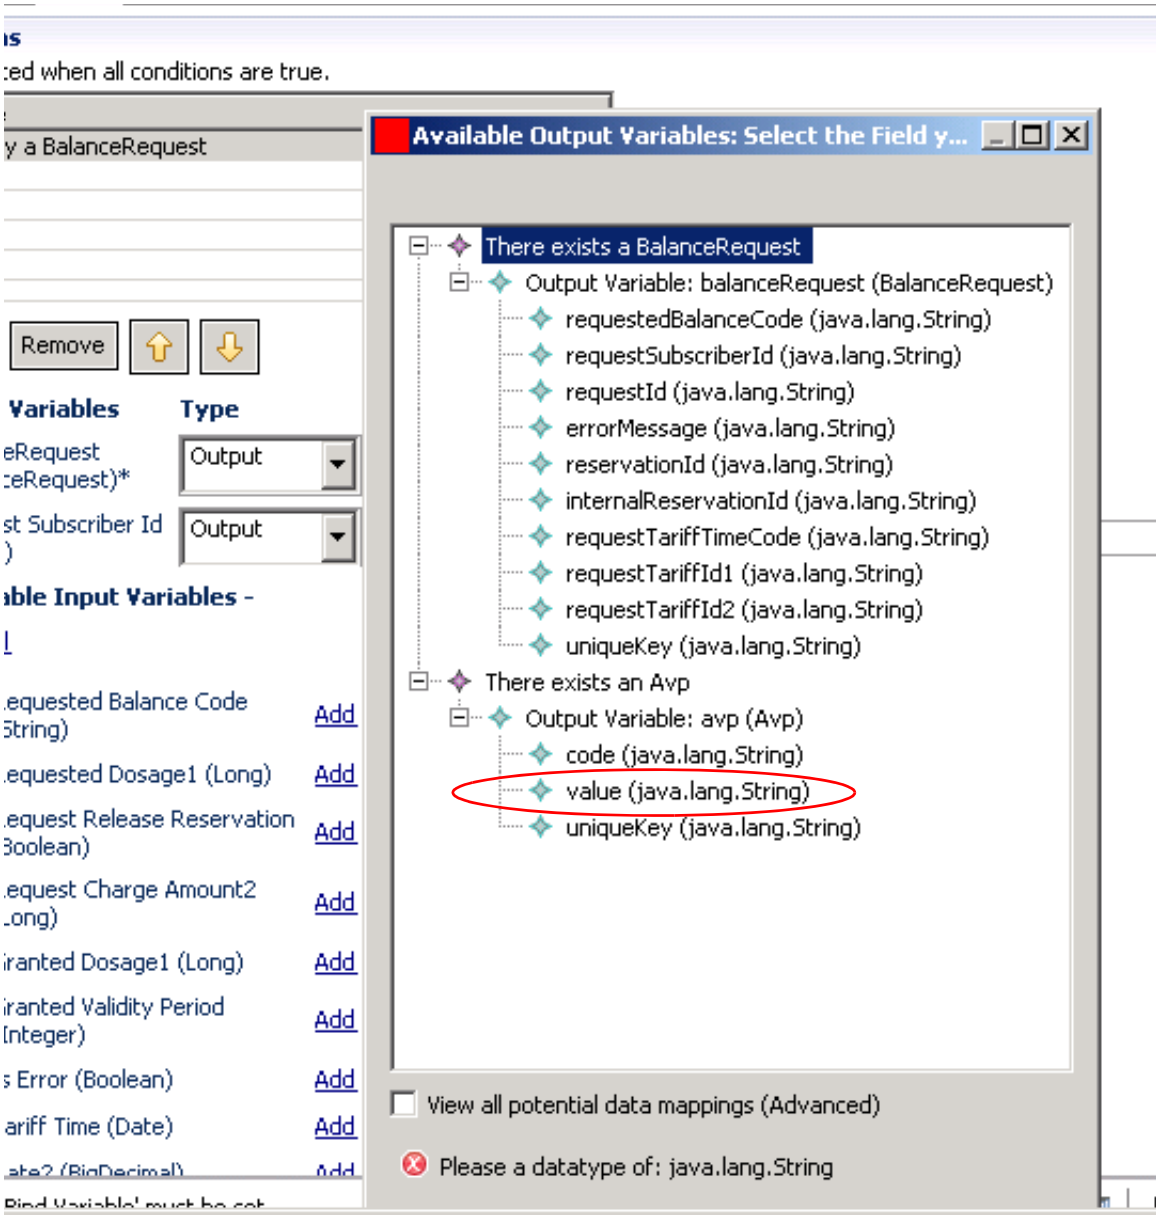

The final action table for the Change Charging ID policy looks like this:

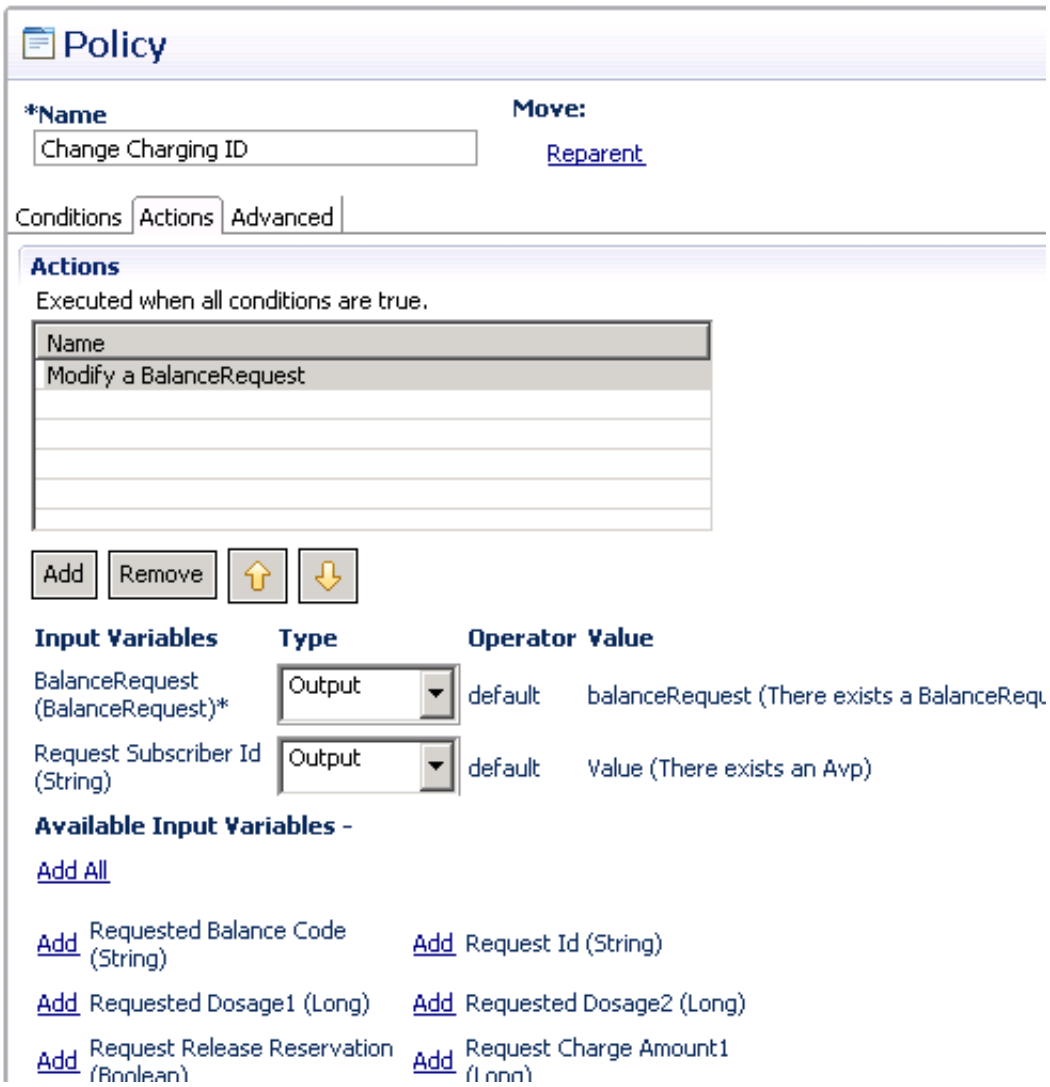

What Happens Next

Now, set your Cisco MsBM configuration against some test data on some test machines.

- Let Cisco MsBMrun for a time
- Use specific data so that you can predict and check the results.
- Check that your rates, tariff times, and policies all perform as required

Tips and Best Practices

Revised: July 10, 2015

This appendix covers the following sections:

Best Practices

- Save your work to the client repository frequently as you work. Use Ctrl-S, File > Save, or click the disk icon located above the tree.
- Publish your work to the Runtime Environment only when you are satisfied that it is correct.
- Check for updates frequently, or follow instructions from your Cisco technical representative to see if you have software updates to install.
- When logging in to Cisco Policy Builder for the first time, do not save your password to the local disk, it weakens the security of your system.
- To find out the version number of your Cisco Policy Buildersoftware components, click Tools > About Policy Builder.
- To see what features are installed, click Tools > About Policy Builder.

Session Information

Within CPS, these steps show how to find out what information is on the session. Knowing what is on the session, you can then create policy conditions and actions that take advantage of that data.

Have the CPS set up and taking traffic with subscribers, even if it is a test scenario.

Step 1 Log in to the Policy Builder interface and click Monitoring tab > Session Information folder.

Step 2 Open the Session Search Filters area in the main window and enter the criteria you want to search on.

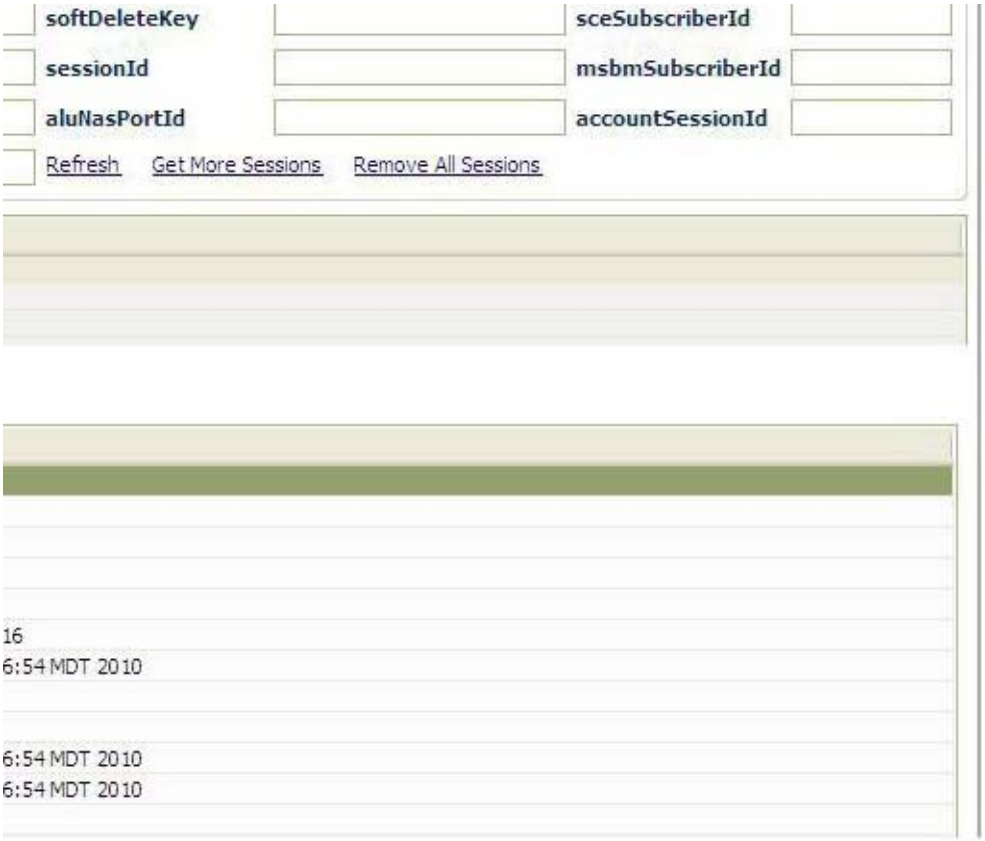

Typical Tasks for Everyday

This section shows you the techniques and shortcuts available in the Cisco Policy Builder interface.

Actions Menu and Copying

On many screens, in the pane in the right side, the Actions area displays a link for copying the information or object that you are looking at.

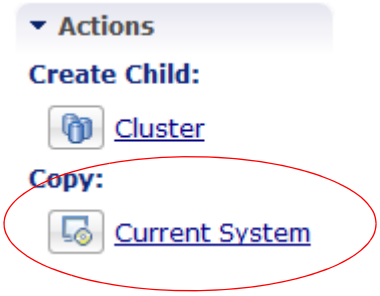

In this example, we can copy the System displayed on the screen, and then rename the new copy, changing the fields as necessary.

Objects that can be copied, and then changed, are:

- Systems, clusters, and instances
- Domains
- Policies, services and service options
- Use case templates
- Policy reporting records
- Notifications
- Customer reference data tables

Always look for the **Copy:** option in the Actions area if you want to make a copy of something.

At the least, you have to provide a new name for the object you copied.

Actions Menu and Create Child

On many screens, in the pane in the right side, the Actions area displays a link for creating a dependent of the object or screen that you are looking at.

The most useful type of Create Child Dependency is the Plugin configurations under systems, cluster, and instance. Plug-ins are really a Dependant to the system or cluster.

Under Reference Data > Account Balance Templates you might have a general account balance template for data. Under this, you can create the dependencies for these special types of data:

- One time quota
- Recurring
- Rollover

Note that not all information objects have the Create Child option. This option is only presented where it makes sense.

Menus

Menus and their drop down items provide general capabilities across all of the Cisco Policy Builder interface.

File Menu

The File menu has items that concern using the data repositories.

- Publish to Runtime Environment —use this when you are satisfied with your test setups and want to move the test repository into production.
- Save to Client Repository—save the repository you are working on.
- Switch Repository—change from any of the repositories you may have developed.

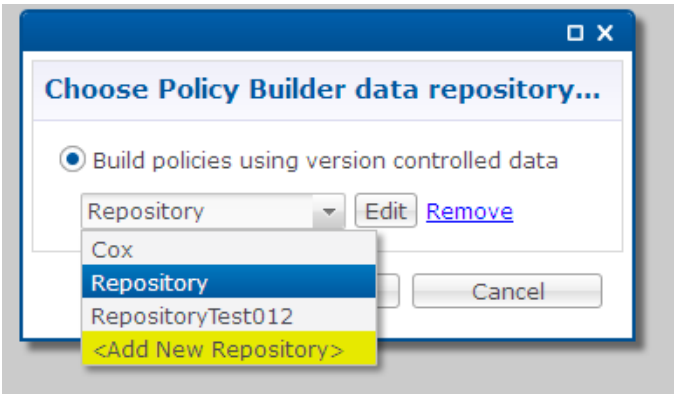

Tools Menu

The Tools menu shows informational subitems.

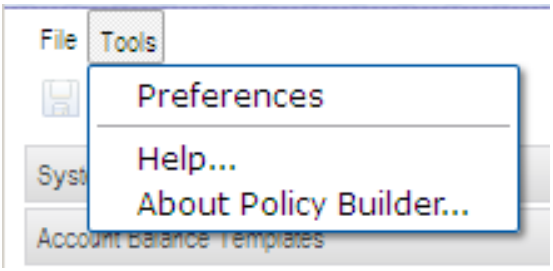

- Preferences—Not available to end users.
- Help—This window lets you select several documents in PDF format.
- About Policy Builder—This menu item displays the build numbers of policy builder and all its features.

Screen Legend

The policy builder screen has several iconic conventions.

- The ellipses ... means a screen has example data to start with.
- The labels at the top of the main window are referred to as tabs.
- The main headings in the screen are called nodes.
- In the tree, the page icon denotes documentation or a report.
- A folder denotes sub items in the tree.
- A right arrow means there are sub items available.
- A slanted arrow means the item is open.
- A red circle with an X in it means you have and error.
- A label with an asterisk *| means the field is required.

Interface Icons

Required fields have an asterisk.

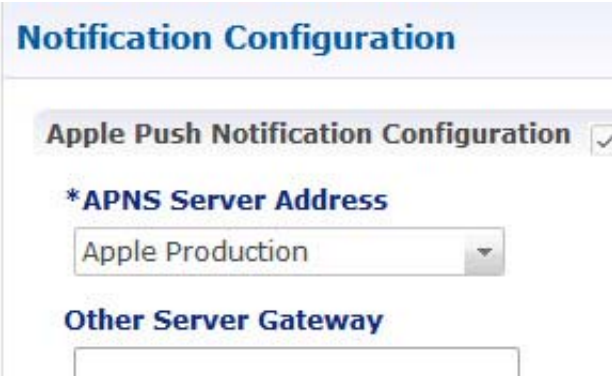

At the top of the tree on the left, helpful icons assist you to navigate the Cisco Policy Builder screens.

Save Your Work

Click the disk icon to save your current work to the client repository. When you are satisfied, you may publish the client repository to the server repository.

Undo and Redo

Use the yellow undo and redo icons as you edit and change screens. It can save typing time.

Delete a Node

Select a node in the tree, then click the red X to delete it. All the children in the node are deleted without any notification.

In the detailed pane on the right, use a Remove button or a delete link to remove specific data.

Page Forward Page Backward

Use the yellow paging icons to page back and forward through the screens you have visited. This can save you from having to go through many menus to compare screens.

Errors

If you see a red or error messages, click down through the tree on the left until you get to the node that has the error.

In the example, the system named test has an error in it. Click into the nodes until you find the screen with the error.

You can still work through your configuration and save it, but the errors still alert you to items that must be corrected.