CISCO

CISCO SERVICE CONTROL SOLUTION GUIDE

Cisco Service Control Online Advertising Solution Guide:

Behavioral Profile Creation Using Traffic Mirroring

1 Overview

Online Behavioral Targeting is an online advertising approach that is based on presenting users with ads based on their interests, as deduced through monitoring their web browsing. The SCE platform can enable this based on an analysis of subscriber online usage patterns.

Such behavioral targeting does not require the analysis of each and every HTTP request on the line, as this would result in a lot of excess information. The SCE platform performs the first level of analysis in the behavioral targeting chain by inspecting the user browsing sessions, detecting the particular requests that are triggered by the actual user browsing (these events are termed ClickStream), and mirroring the traffic pertaining to these events. (Mirroring criteria may be different, depending on actual need.) The mirrored traffic is typically received by an entity that analyzes the nature of usage and creates a profile of the subscriber to be used later for targeting. The way the greater solution works is outside the scope of this document.

The mirroring capability on its own is useful also for a number of other solutions using the SCE platform. While this solution focuses on the behavioral targeting use case, the description of the mirroring capability and related configuration is also applicable for such solutions.

Figure 1 High Level Overview of an Mirroring-based Behavioral Targeting Solution

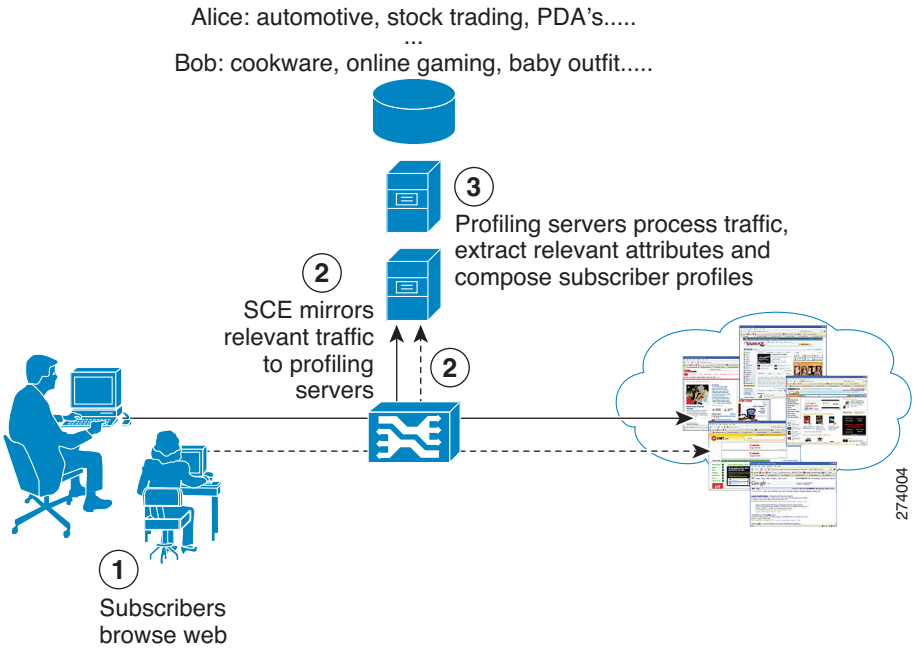

As mentioned above, the mirroring decision can be taken based on a number of criteria. In fact, the mirroring decision can be triggered based on each of the criteria that are used by SCA BB for classification of traffic.

One such particular example is traffic mirroring of HTTP traffic that is based on ClickStream. ClickStream detection is a fundamental capability of the solution, as it can detect which specific requests, out of the enormous number of HTTP requests generated throughout the subscriber web activity, were triggered by the subscriber. When a subscriber clicks a link, or types a URL to the browser address bar, an http request is generated for this URL. Typically, an html page is returned, which constitutes the outline of the contents requested. For the browser to be able to render this page, it must download multiple objects (tens or sometimes around a hundred for a single page viewed), which in turn results in multiple http requests for obtaining these objects.

To be able to conduct behavioral targeting, it is typically sufficient to understand what the user was trying to do (represented by the initial request, such as biz.publisher.com/ap/081120/world_markets.html ' "global markets"), rather than looking at each and every object downloaded as a secondary result of such a request (such as. http://ads.adnetwork.com/a/a/in/interbroke/300x250_yah.jpg --> broker ad).

ClickStream detection makes exactly this distinction, allowing the number of requests to be analyzed to be greatly reduced, which is necessary in order to enable a scalable analysis solution. At same time no data is provided concerning what the subscriber is actually doing.

Traffic that has been designated to be mirrored is replicated by the SCE platform and sent over a designated vlan and a designated pair of ports towards the listening servers.

The SCE platform supports multiple logical destinations for mirroring, each of which can be represented by one or more vlan, which will be load-shared by the SCE platform. Load sharing makes sure that all the traffic of a given subscriber belonging to a particular server group is handled by the same vlan.

Mirroring of a flow can continue indefinitely (until the flow is terminated) or be limited to a predefined volume passed over the flow, after which the mirroring will be stopped.

The impact of traffic mirroring on overall system performance depends on the actual percentage of traffic that is mirrored. It is recommended to monitor SCE performance when enabling traffic mirroring.

2 Configuring Traffic Mirroring Support: Highlights

This section provides the highlights of configuring the main components of traffic mirroring on the SCE platform. For complete configuration directions, see Step by Step Guide: Configuring an SCE Platform for Traffic Mirroring, page 9

Defining the Mirroring Server Groups

The mirrored traffic can be sent to one of eight possible server groups. These are server groups rather than individual servers, as the underlying infrastructure allows load-sharing the traffic destined to a serve group across multiple vlans.

These server groups are defined in the Policies tab of the Service Configuration Editor. Click **Configuration** and select **VAS Settings**.

Click the top radio button for traffic mirroring, and then define the names of the server-groups you will be using. You will use the server group IDs to define the transport setting for the solution later on.

For each server group, you can specify the flow volume (in L3 Kbytes) to mirror to the server. If left at 0 (the default), the entire flow will be mirrored. Otherwise, mirroring will be stopped once the specified volume has been mirrored.

Figure 2 VAS Settings

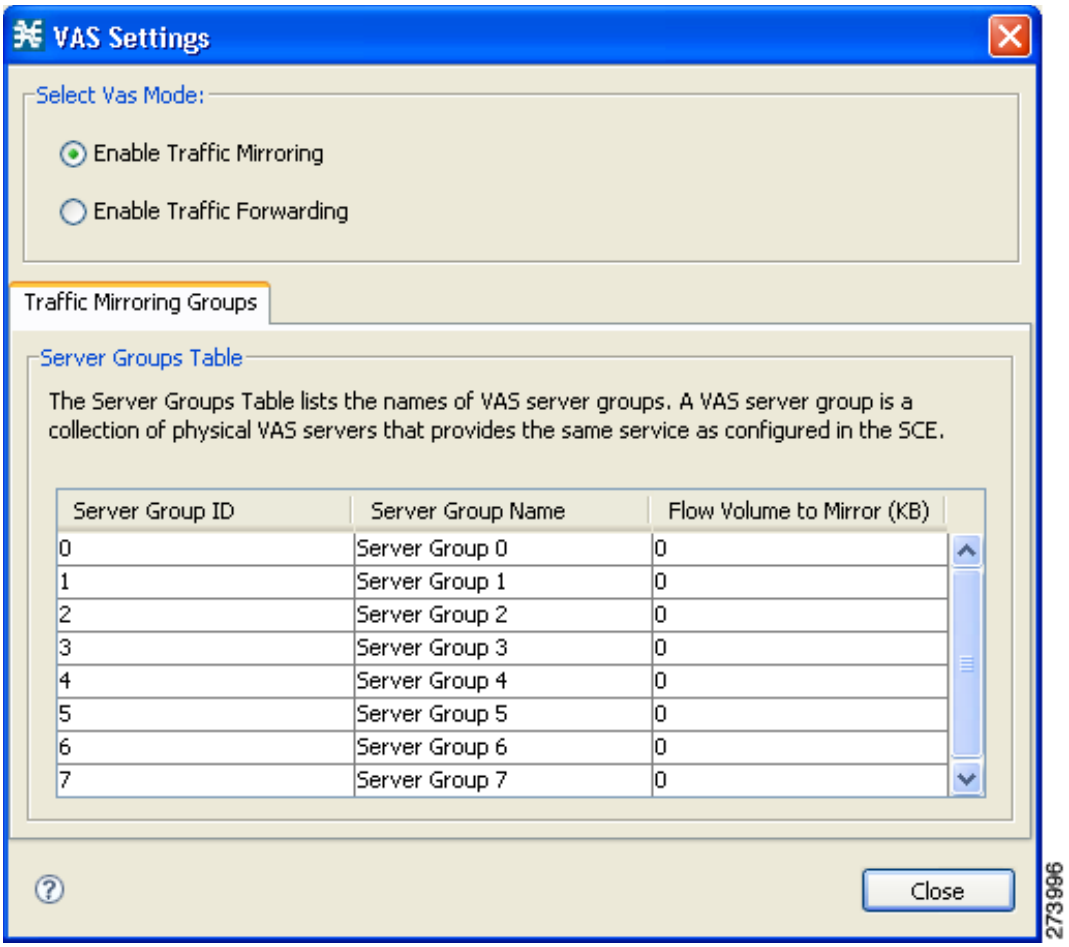

Creating a ClickStream Service

Note Identifying traffic as 'ClickStream' is one way of identifying traffic to be mirrored. Other approaches may involve classification based on other attributes, such as the URL matching a certain prefix or a user agent. This section is relevant for a case where ClickStream is used as the criteria for the traffic mirroring.

ClickStream signatures are mapped by default to the HTTP Browsing protocol and consequently to the browsing service. In order to be able to act on them separately, you will first have to move them to a protocol of their own, and assign this protocol to a service of its own.

Figure 3 Configuring the ClickStream Protocol

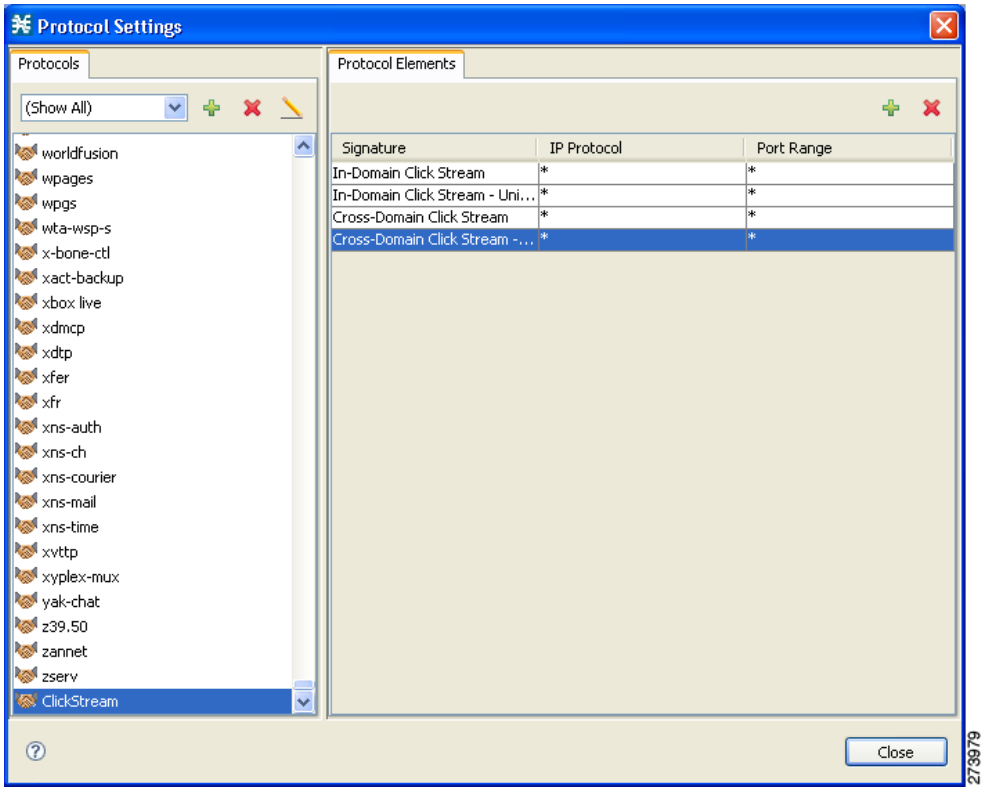

Figure 4 Configuring the ClickStream Service

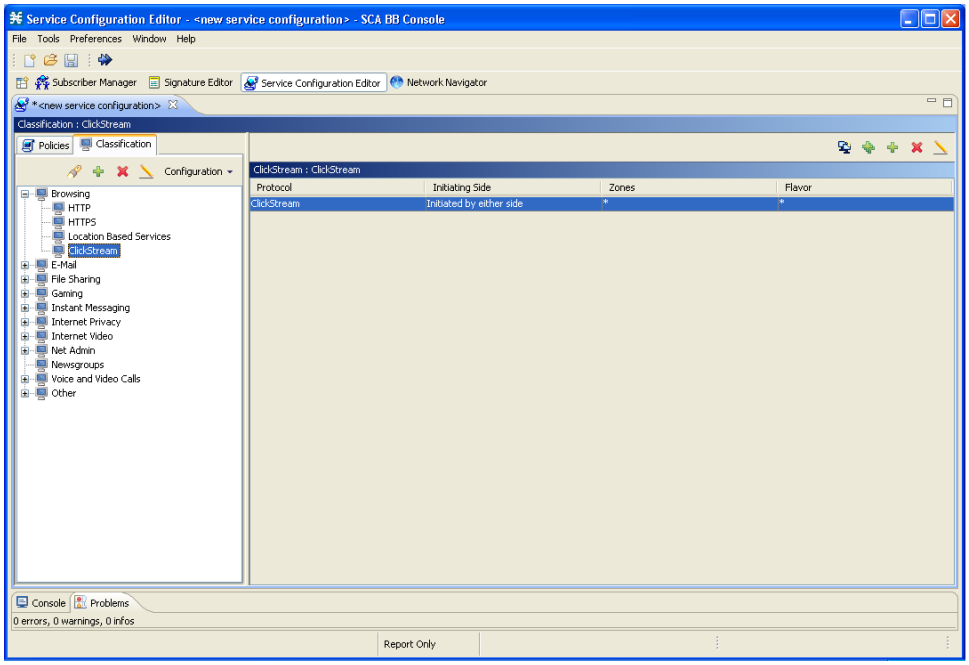

Enabling Deep HTTP Inspection

In order to enable comprehensive detection of the ClickStream events in the traffic stream, it is important to enable deep inspection of HTTP, which configures the SCE platform to analyze and classify all HTTP requests within a single flow.

Some browsers, in conjunction with some web server implementations, will use the same TCP flow to carry multiple requests triggered by clicks that are targeting the same host. Such events will not be detected if the classification is only done at the beginning of the flow (which is the default for SCA BB).

To enable deep HTTP inspection, in the SCA BB Console Service Configuration Editor, go to:

Configuration>System Settings>Advanced Options tab**>Advanced Service Configuration Options...**

Note Enabling deep HTTP inspection is expected to impact the SCE performance due to the excessive processing associated with it; the actual figure depending on the amount and on the nature of HTTP traffic. It is recommended that you monitor SCE platform performance when enabling this capability.

Creating Traffic Mirroring Rules

The traffic to be mirrored is defined by creating traffic rules that specify the mirroring action for the relevant traffic.

As a prerequisite, you must create a service that includes the type of traffic to be mirrored. This can be either the ClickStream service described above, or any other service defined through the SCA BB service configuration.

For each package with traffic to be mirrored, select the relevant service and activate mirroring to the proper server (that you have already configured using the VAS Settings dialog, see Defining the Mirroring Server Groups, page 3). The mirroring action is not exclusive, and you can configure it in parallel with other actions that need to be applied to the same service.

Note Leveraging subscriber awareness with traffic mirroring: Subscriber awareness is key to behavioral targeting using traffic mirroring, as it enables a network level opt-in or opt-out; a feature that is considered important to subscriber privacy. This is implemented using the SCE platform native subscriber awareness, by creating packages that allow or deny traffic mirroring, and assigning subscribers to these packages based on their opted-in or opted-out nature.

SCE Connectivity

Traffic mirroring is implemented by sending the mirrored packets over a designated vlan through a predefined link of the SCE platform. The link that has been defined for traffic mirroring can be either used exclusively for this purpose, or it can be one of the traffic ports, in which case the Tx capacity of the link will be shared between the original egress traffic and the mirrored traffic.

Traffic that is received on the subscriber interface on either link is sent over a vlan on the network interface over this predefined link. Traffic that is received on the network interface on either link is sent over a vlan on the subscriber interface over this predefined link.

The following illustration shows an SCE 2000 platform using a dedicated link for mirroring. The same topology is applicable using SCE8000 platform.

Figure 5 Traffic Mirroring on a Dedicated Link

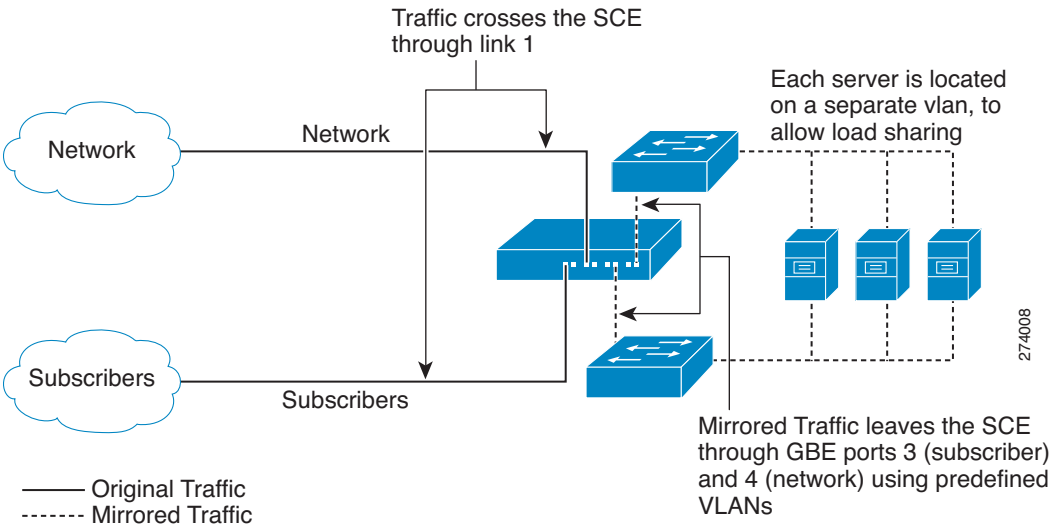

The following illustration shows an SCE 2000 platform using traffic ports for mirroring. The same topology is applicable using SCE8000 platform.

Figure 6 Traffic Mirroring over Traffic Ports

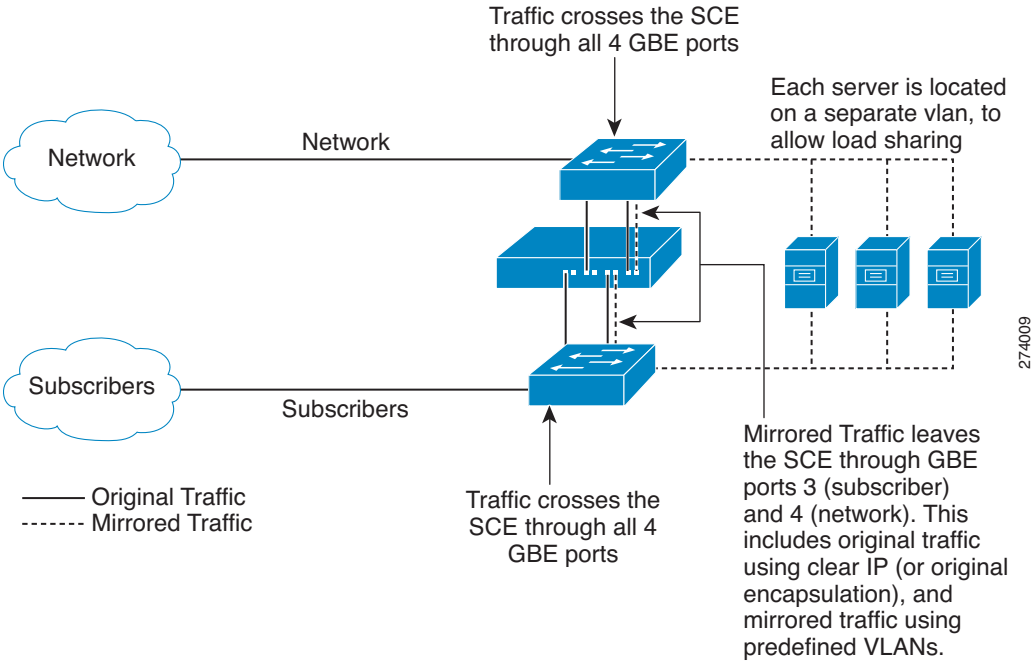

Configuring Traffic Mirroring Transport

Traffic mirroring transport is configured using the SCE platform CLI, and connects between the logical mapping to server groups, as defined through the SCA BB console, and the actual transmission of mirrored traffic, which is done over a VLAN.

You do this by defining physical servers that are mapped to VLANs, and associating these servers to server groups (which have been defined through the SCA BB console).

To configure the link over which traffic will be mirrored, use the following CLI command.

SCE(config if)# VAS-traffic-forwarding traffic-link {link-0|link-1}

To view the link over which traffic will be mirrored, use the following CLI command.

SCE#show interface linecard 0 VAS-traffic-forwarding

The server assigned to this traffic by the policy selects the vlan to send the traffic over. One or more vlans can be associated to each server, and the SCE platform load-shares the traffic destined to each server between these vlans. Load sharing is done at subscriber level (all traffic belonging to a specific subscriber is transmitted on the same vlan). Up to 64 distinct vlans can be supported by an SCE8000 platform, and up to 8 distinct vlans by an SCE 2000 platform.

To configure a vlan to be used for a particular server, use the following CLI command (linecard interface configuration mode:

SCE(config if)# VAS-traffic-forwarding VAS server-id *number* **VLAN** *vlan-id*

To view vlans that are used for a particular server, use the following CLI command:

SCE#show interface linecard 0 VAS-traffic-forwarding VAS server-id *id-number*

To remove vlan to from a particular server, use the following CLI command (linecard interface configuration mode:

SCE(config if)#no VAS-traffic-forwarding VAS server-id *number* **VLAN** *vlan-id*

To associate a server with a server group use the following CLI command (linecard interface configuration mode):

SCE(config if)#VAS-traffic-forwarding VAS server-group *group-number* **server-id** *id-number*

Mirrored Traffic - The Server Side

There are a few assumptions regarding mirrored traffic that the listening server should be aware of. Here are the highlights:

Start mirroring

Mirroring starts once the flow has been classified and matched to a service by the SCE platform. For TCP flows, this typically (but not always) happens on the first payload packet. As a result, the entire TCP handshake is not mirrored.

Mirroring of ACK only packets

Such packets (or more generically, packets with no payload at all) are not mirrored. While this should not affect the ability of a server to process the traffic, packets that were on the original data flow may be missing. RST and FIN packets are exceptions to this rule - see more below.

Mirroring of connection termination

- For connections that have been terminated in an orderly fashion - only the last FIN and ACK packets are mirrored.
- For connections that have been terminated using RST - only the RST packet is mirrored.
- For connections that for some reason have not been terminated - no connection termination indication will be sent.

Stop mirroring indication

When the SCE platform stops mirroring a flow because the specified volume has been already mirrored, it generates an RST packet over the mirrored VLAN, to indicate that mirroring has stopped for this flow.

Traffic encapsulation

Mirrored traffic will be encapsulated in a VLAN based on the VLAN number that has been assigned to that particular subscriber by the SCE platform.

In the case of traffic that has been originally encapsulated in a VLAN, an SCE8000 will remove the original VLAN and insert the mirroring VLAN instead. In such cases, the SCE2000 will add the mirroring VLAN on top of the original VLAN.

For all other types of encapsulation, the original packet will be encapsulated in a VLAN as-is.

3 Step by Step Guide: Configuring an SCE Platform for Traffic Mirroring

This section explains in detail how to configure a system for traffic mirroring.

- To configure a solution that mirrors ClickStream traffic, complete all the following steps.
- To configure a solution that does not mirror ClickStream traffic, skip to Step 22. (The beginning steps, define the ClickStream traffic, which is not relevant in this case.)

Step 1 In the SCA BB Policy Editor, select the **Classification** tab (left pane), click **Configuration,** and select **Protocols**.

Step 2 In the Protocol Settings window, select the **HTTP Browsing** service.

Step 3 In the **Protocol Elements** tab, remove the ClickStream related protocol elements:

- In-Domain Click Stream
- In-Domain Click Stream - Unidirectional Client Request
- Cross-Domain Click Stream
- Cross-Domain Click Stream - Unidirectional Client Request

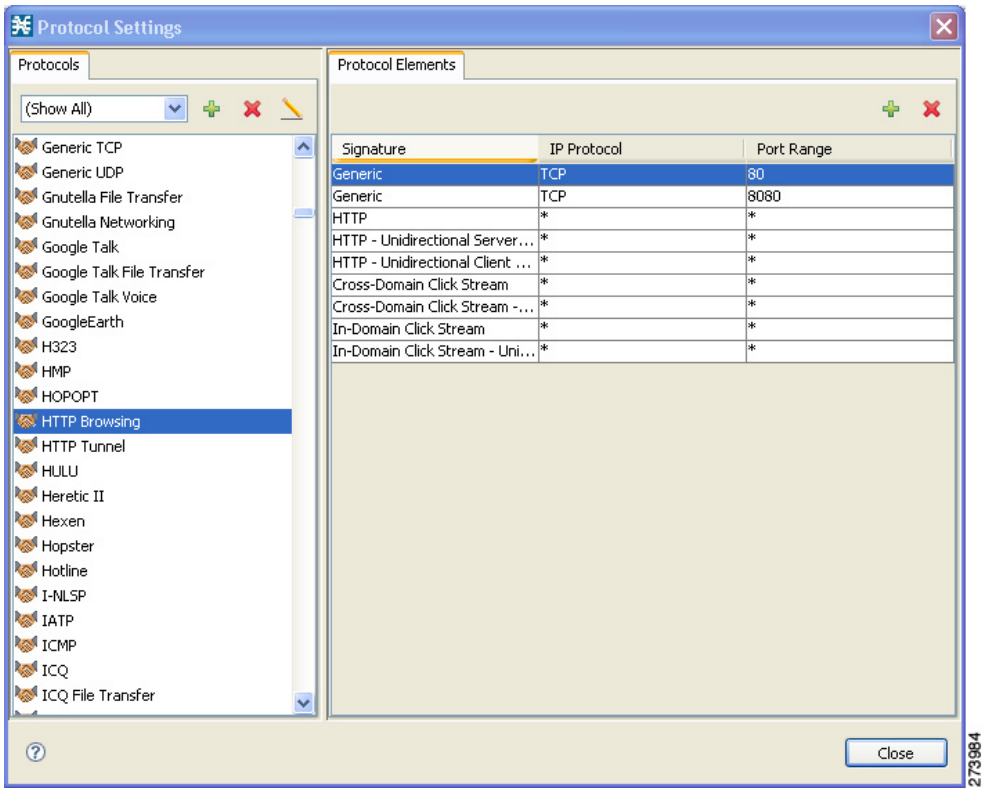

Step 4 In the Protocol Settings window, on the Protocols tab, click the '**+**' to add a new protocol.

Step 5 Name the new protocol **ClickStream Event** and click **OK**.

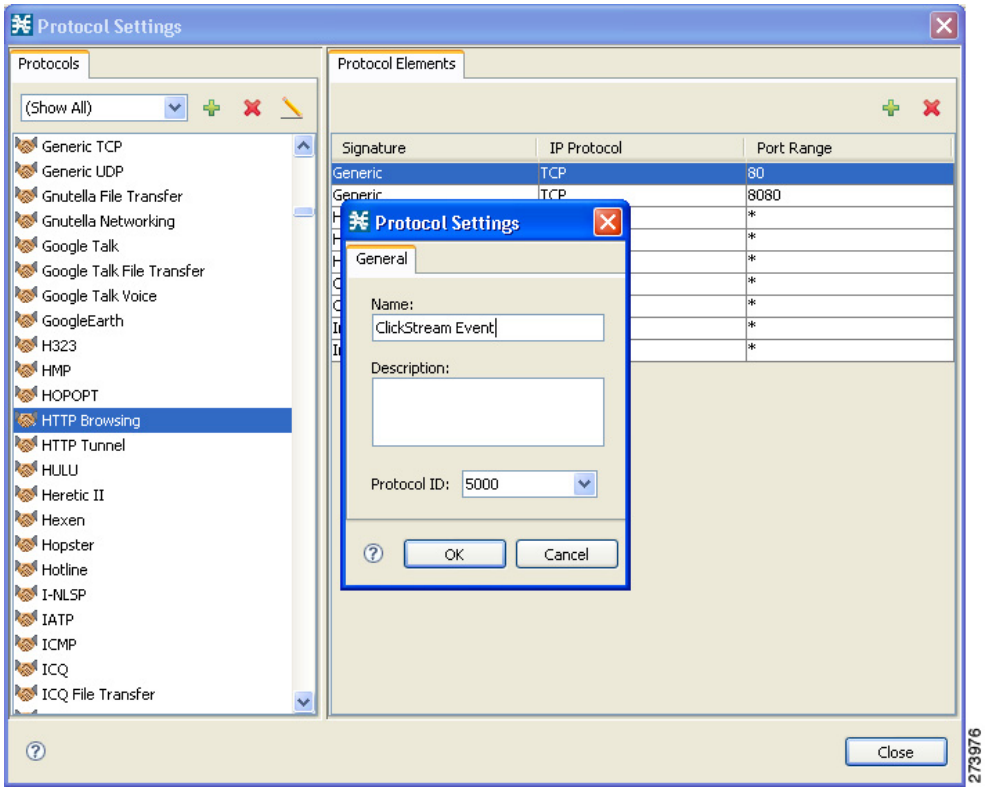

Step 6 In the Protocol Elements tab, click the '**+**' to add protocol elements to the ClickStream Protocol.

Step 7 For the new protocol element created., click in the Signature column on the '**...**' button.

Step 8 On the Select a Signature screen, add the **In-Domain Click Stream** signature and click **OK**.

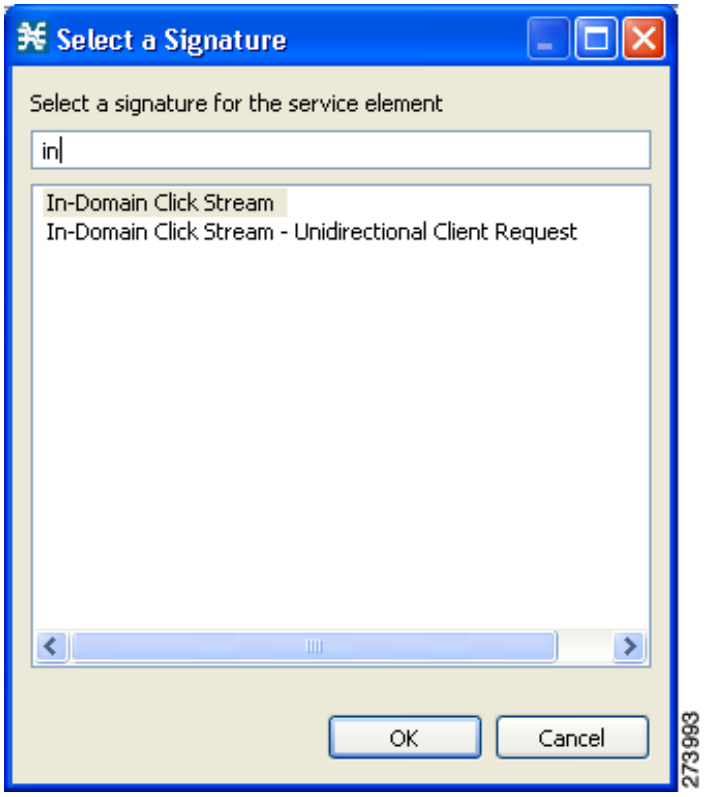

Step 9 Repeat Step 6 through Step 8 for the rest of the ClickStream signatures:

- In-Domain Click Stream - Unidirectional Client Request
- Cross-Domain Click Stream
- Cross-Domain Click Stream - Unidirectional Client Request

Step 10 In the SCA BB Policy Editor, select the Classification tab (left pane), and highlight the Browsing service

Step 11 Click the '**+**' to add a new service under the Browsing service.

Step 12 Name the service **ClickStream** (or any other name you choose).

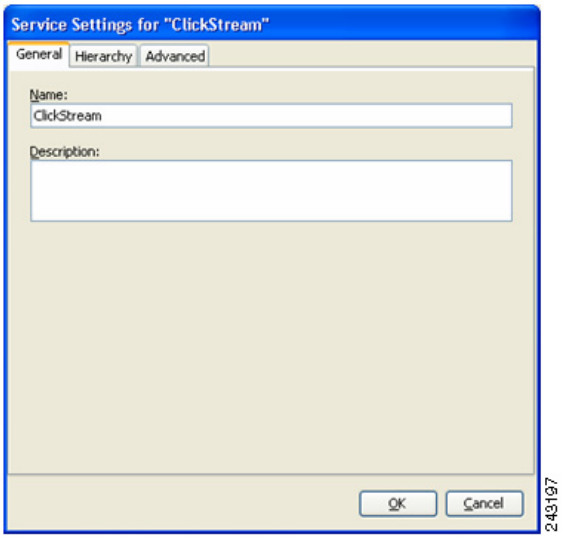

Step 13 Click the **Hierarchy** tab and check the two check boxes to add a dedicated service counter to the ClickStream Service.

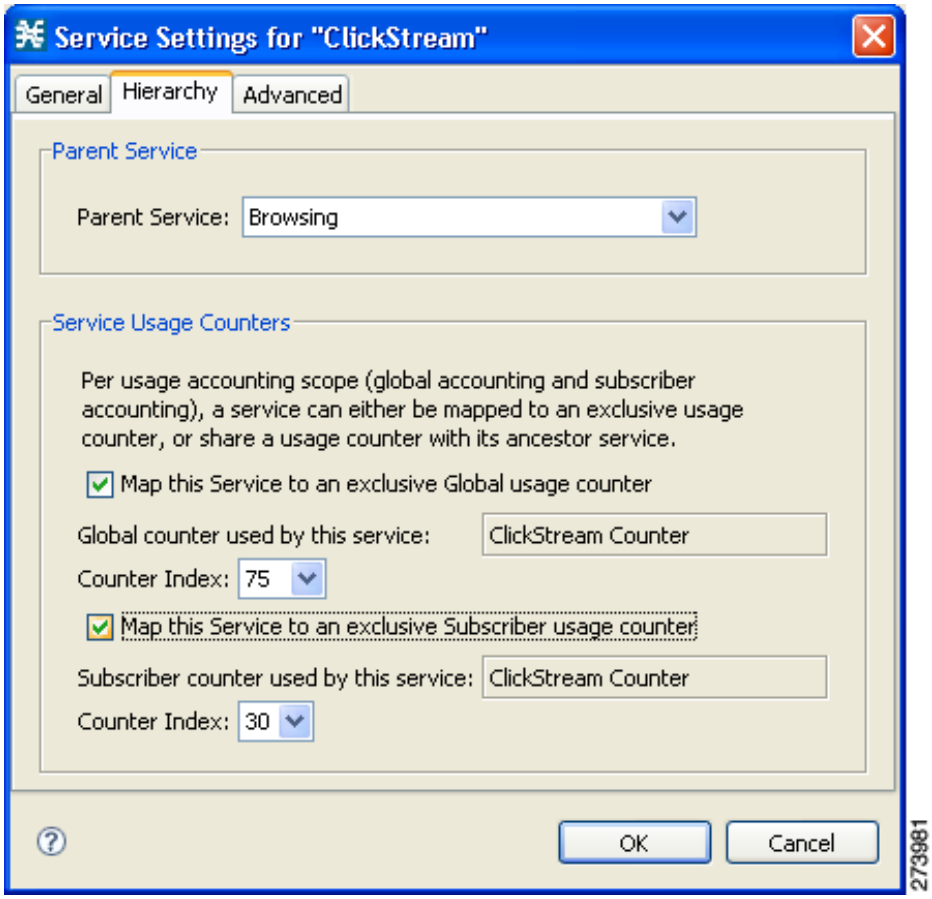

Step 14 Click **OK**.

Step 15 In the right pane, click the '**+**' icon to add a service element.

Step 16 In the dialog that opens, click **Select** next to the Protocol field and select the **ClickStream Event** protocol (or whatever you named your ClickStream protocol) from the list.

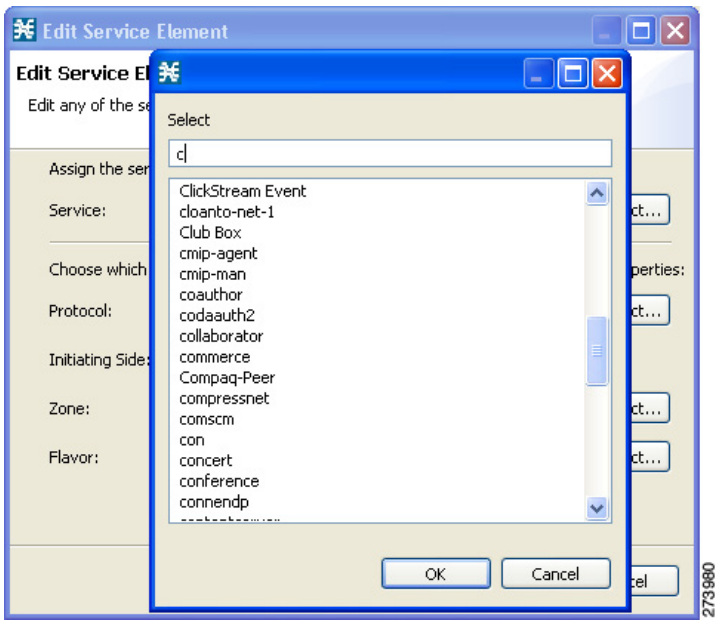

Step 17 Click **OK**.

Step 18 In the Policies tab of the Service Configuration Editor, select:

Configuration>VAS settings

Step 19 Click the **Enable Traffic Mirroring** radio button.

Step 20 In the lower part of the window, define a name for each of the server groups you will be using.

Step 21 For each server group, define the per-flow volume (in KB) to be mirrored to this group (for flows matching the criteria). Leaving the value '0' will allow the entire flow mirrored.

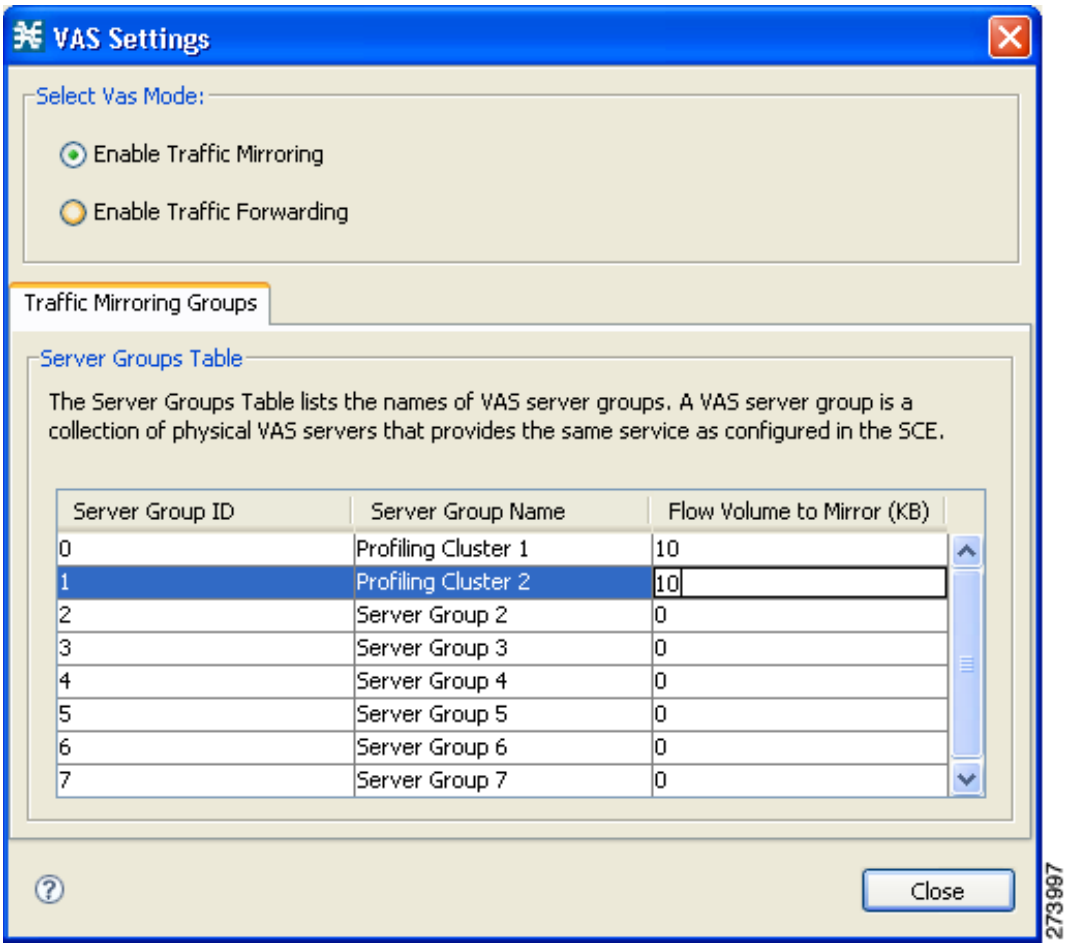

Step 22 In the SCA BB Policy Editor, select the **Policies** tab (left pane), and then select the package for which to mirror the traffic.

Step 23 In the right pane, click the '**+**' icon to add the ClickStream service (or any other service whose traffic mirrored) .

Step 24 In the dialog that opens, select **ClickStream** (or any other service) from the drop-down selection.

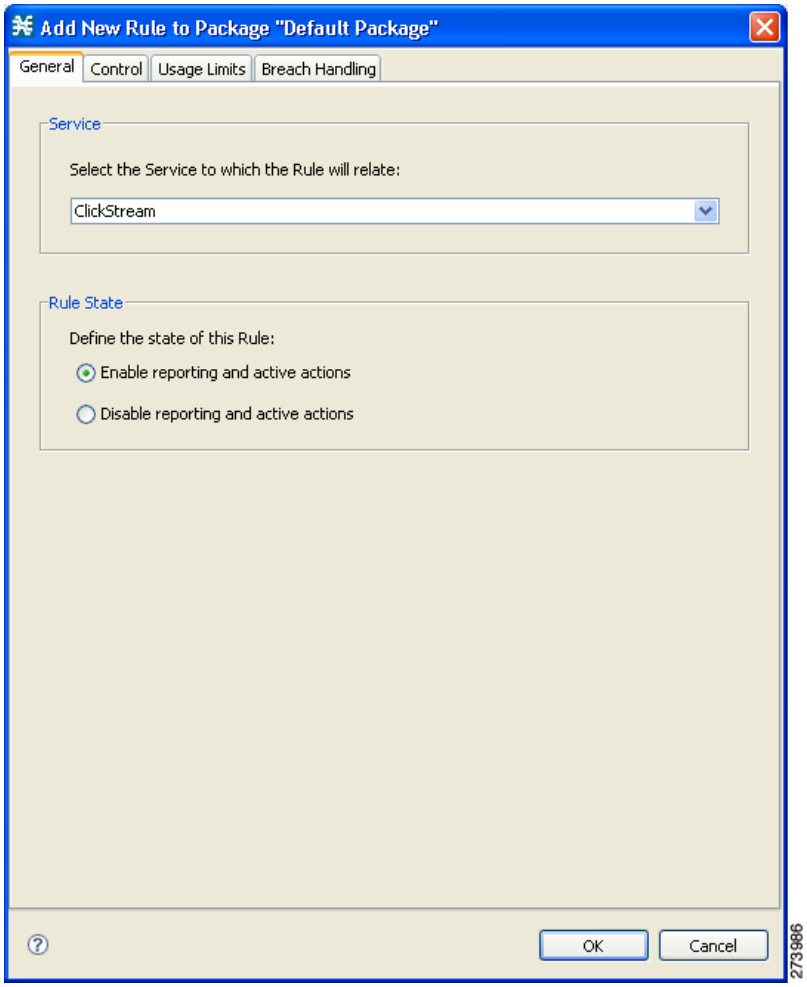

Step 25 Click the Control tab and check the **Mirror Traffic to Server Group** check box. From the associated drop down menu, select the server group to mirror the traffic to.

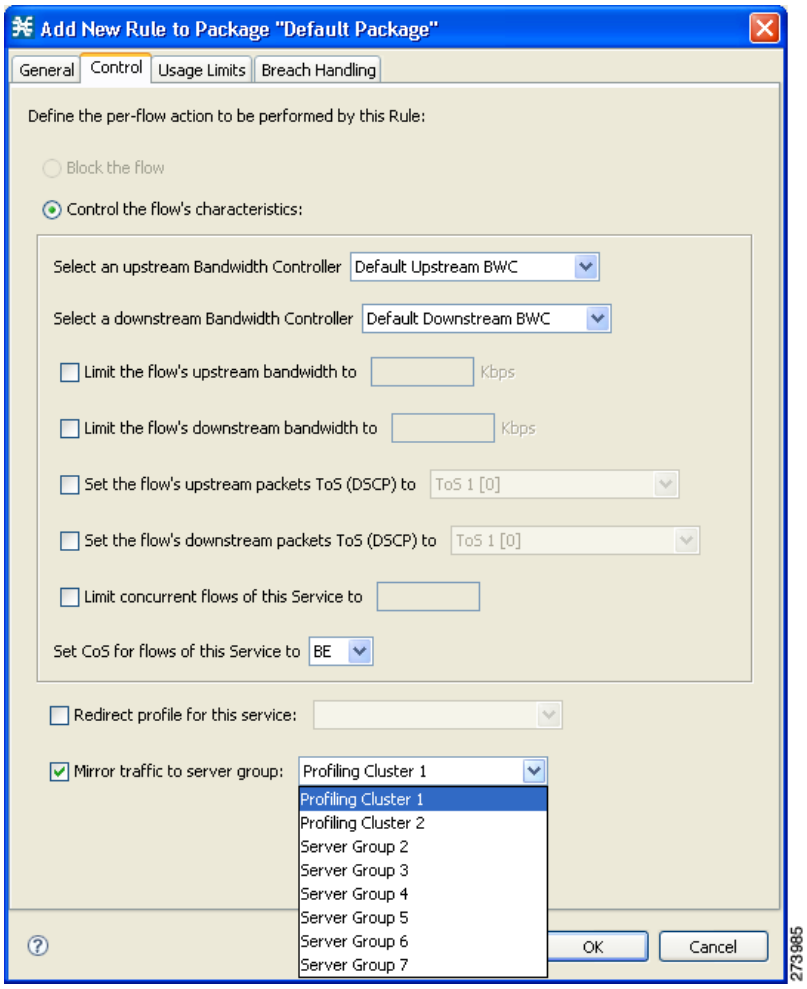

Step 26 Click **OK**.

Step 27 Repeat Step 23 through Step 26 for all services in the selected package that require traffic mirroring.

Step 28 Repeat Step 22 through Step 27 for all packages that require traffic mirroring.

Step 29 (Optional) Enable deep HTTP inspection. This allows the mirroring decision to be taken for each http request within a flow separately.

a. Select **Policies > Configuration >System Settings.**

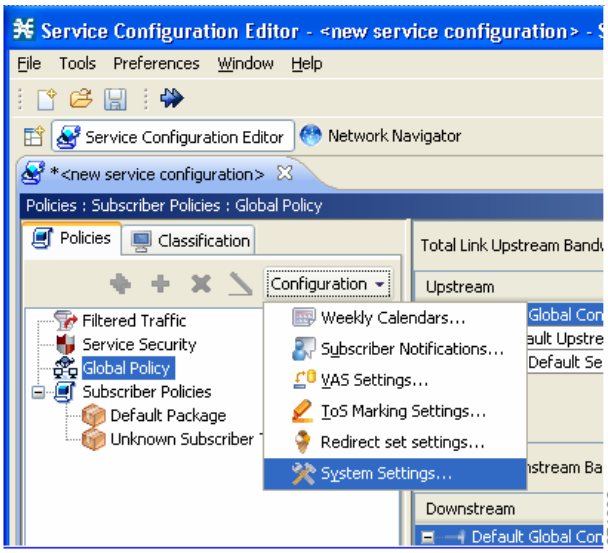

b. In the Advanced Options tab, click **Advanced Service Configuration Options** to enable deep inspection of http flows by setting the highlighted value to **64000.** This enables the analysis of multiple transactions within a single HTTP flow, which is important for comprehensive detection of ClickStream events.

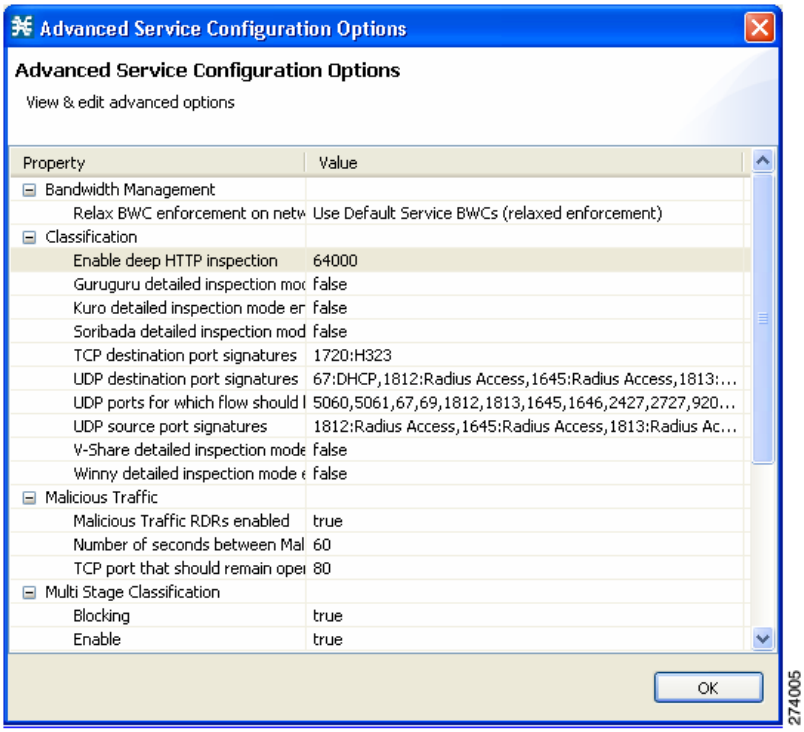

Step 30 This concludes the policy editing part of the configuration. Apply the Service Configuration to the SCE platform.

Step 31 Configure the link to be used for traffic mirroring on the SCE platform.

```
SCE(config if)#VAS-traffic-forwarding traffic-link {link-0 |link-1}
```

Step 32 Configure a VLAN tag for each physical VAS server.

```
SCE(config)#VAS-traffic-forwarding VAS server-id <number> VLAN <number>
```

Step 33 Assign each server to a server group:

```
SCE(config)#VAS-traffic-forwarding VAS server-group <number> server-id <number>
```

Step 34 Save the configuration:

```
SCE#copy running-config-all startup-config-all
```

4 Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly What's New in Cisco Product Documentation, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

Subscribe to the What's New in Cisco Product Documentation as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS Version 2.0.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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**.