C H A P T E R **10**

Using the Service Configuration Editor: Additional Options

This chapter explains how to use additional, advanced functionality available in the Service Configuration Editor.

The Service Security Dashboard

The Service Security Dashboard allows you to view and control all SCA BB security functionality.

The Dashboard is a gateway to a set of features that help you protect your network from security threats such as worms, DDoS attacks, and spam zombies. It allows configuration of the detection mechanisms (for example, attack thresholds) and of the actions to be taken when an attack is detected.

The Dashboard also allows you to access malicious traffic reports in the Reporter tool.

⚠ **Caution**   If anomaly-based detection of malicious traffic is enabled, any access control list (ACL) that is configured on the Service Control Engine (SCE) platform but is not applied to anything (for example, an interface, an access map, or anSNMP community string) might be deleted when a service configuration is applied to the platform.
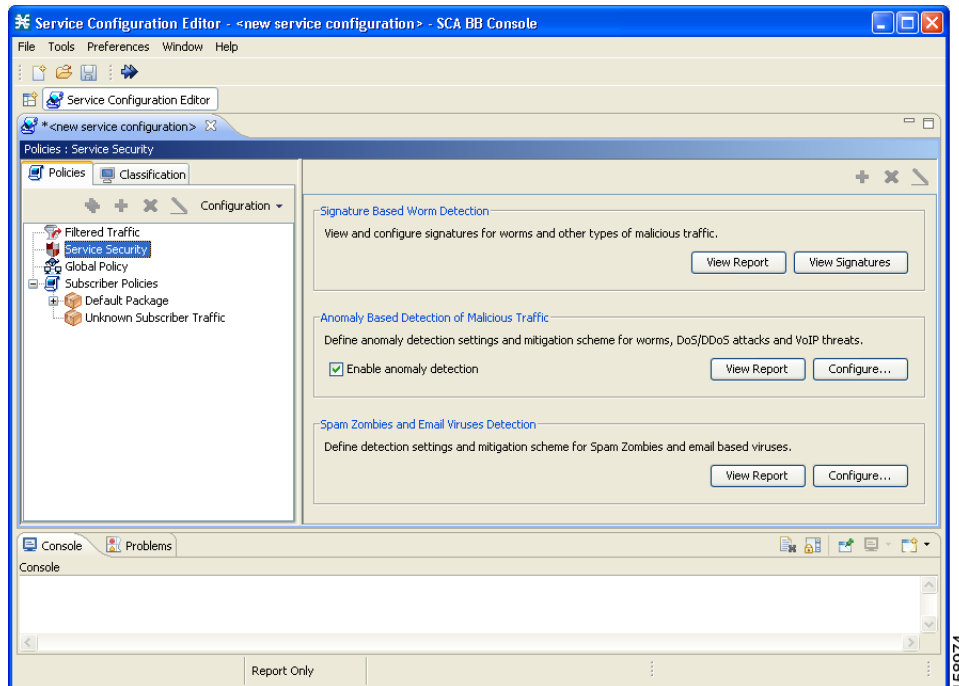
Workaround:

Disable anomaly-based detection of malicious traffic. (Clear the **Enable anomaly detection** check box.)

How to View the Service Security Dashboard

Step 1   In the Network Traffic tab, select **Service Security.**

Step 2   The Service Security Dashboard is displayed in the right pane.

Managing Worm Detection

SCA BB uses three mechanisms for detecting worms:

- Signature based detection—The SCE platform's stateful Layer 7 capabilities can detect malicious activity that is not easily detectable by other mechanisms. You can add signatures for new worms.

- Anomaly based detection—Overall traffic analysis can detect anomalies that might indicate worm activity. See Managing Anomaly Detection, page 10-3.

- Mass-mailing based detection—E-mail traffic analysis can detect anomalies that might indicate e-mail-based worms. See How to Configure Spam Detection Settings, page 10-15.

How to View Supported Worm Signatures

Step 1   In the Service Security Dashboard, click **View Signatures.**

The Signatures Settings dialog box appears, with Worm Signatures selected in the Signature Type drop-down list.

All supported worm signatures are listed.

Step 2   Click **Close.**

The Signatures Settings dialog box closes.

How to Add New Worm Signatures to a Service Configuration

Step 1 Either import the latest DSS or SPQI file provided by Cisco or create a DSS file containing any worm signatures that you wish to add to the service configuration.

Related Info

For more information, see Managing Protocol Signatures, page 7-33.

Managing Anomaly Detection

The most comprehensive threat detection method is anomaly detection.

Anomaly Detection

The basic principle of anomaly detection is monitoring successful (correctly established for TCP, bi-directional for other protocols) and unsuccessful (not properly established for TCP, unidirectional for other protocols) connection rates both to and from any IP address viewed by the system, and triggering an anomaly detection condition based of one of the following criteria:

- The total connection rate exceeds a predefined threshold.
- The suspicious connection rate exceeds a predefined threshold *and* the ratio of suspicious to unsuspicious connections exceeds a predefined threshold.

The ratio metric is a particularly robust indicator of malicious activity, and together with a rate qualifier serves as a reliable identifier for malicious activity.

Anomaly detection is divided into three categories based on the directional nature of the detected anomaly condition. The concepts used for the three categories are identical, but the nature of the detected malicious activity is different for each category.

- Scan/Sweep detector—Detects malicious activity based on an anomaly in connection rates *from* an IP address.
- DoS detector—Detects an anomaly in the connection rate between a pair of IP addresses: one of them is attacking the other. This can be either an isolated attack or part of a larger scale DDoS attack.

- DDoS detector—Detects an anomaly in the connection rate coming *to* an IP address, which means that it is being attacked. The attack can be by either a single IP address (DoS) or multiple IP addresses.

For all kinds of anomaly detection conditions, maximum flexibility is provided by the ability to define detection thresholds and the trigger actions to be taken for each:

- Flow direction
- Flow protocol
- (Optional) Port uniqueness for TCP and UDP

Note The GUI configuration described here replaces the CLI command set for configuring the Attack Filtering Module of the SCE platform, which was available in previous releases.

Anomaly Detection Parameters

For each anomaly detector category (Scan/Sweep, DoS, DDoS) there is one default detector. You can add additional detectors of each category. Detectors in each category are checked in order; the first match (according to the detector's threshold settings) triggers detection. You set the order in which detectors are checked; the default detector is checked last.

Anomaly detectors can contain up to 12 anomaly types associated with malicious traffic:

- Network initiated—Malicious traffic initiated from the network side:
    - TCP—Aggregate TCP traffic on all ports
    - TCP Specific Ports—TCP traffic on any single port
    - UDP—Aggregate UDP traffic on all ports
    - UDP Specific Ports—UDP traffic on any single port
    - ICMP—Aggregate ICMP traffic on all ports
    - Other—Aggregate traffic using other protocol types on all ports
- Subscriber initiated—Malicious traffic initiated from the subscriber side:
    - TCP
    - TCP Specific Ports
    - UDP
    - UDP Specific Ports
    - ICMP
    - Other

Note ICMP and Other anomaly types are not available for DoS attack detectors.

Each anomaly type on a detector has the following attributes associated with it:

- Detection thresholds—There are two thresholds, crossing either of them means that an attack is defined to be in progress:
    - Session Rate threshold—The number of sessions (per second) over specified ports for a single IP address that trigger the anomaly detection condition.

- Suspected sessions threshold— Suspected sessions are sessions that are not properly established (for TCP), or that are unidirectional sessions (for other protocols). Exceeding both the Suspected Session Rate *and* the Suspected Session Ratio will trigger the anomaly detection condition. (A relatively high session rate with a low response rate typically indicates malicious activity.)

    - Suspected Session Rate—The number of suspected sessions (per second) over specified ports for a single IP address.

    - Suspected Session Ratio—The ratio (as a percentage) between the suspected session rate and the total session rate. A high ratio indicates that many sessions received no response, an indication of malicious activity.

- Actions—Zero or more of the following actions may be taken when an anomaly detection condition is triggered (by default, no action is enabled):

> **Note** Logging of the anomaly to an on-device log file and generation of RDRs is not configurable per anomaly type.

    - Alert User—Generate an SNMP trap (see the "SCA BB Proprietary MIB Reference" chapter of the *Cisco Service Control Application for Broadband Reference Guide* for information about the Cisco proprietary MIB) indicating the beginning and end of an anomaly.

    - Notify Subscriber—Notify the relevant subscriber of the malicious activity, by redirecting his browsing sessions to a captive portal. To configure network attack subscriber notification, see Managing Subscriber Notifications, page 10-27.

    - Block Attack—Block the relevant sessions. Blocking is performed based on the specification of the malicious traffic that triggered the anomaly detection condition. If subscriber notification is enabled for the anomaly type, blocking is not applied to the port relevant for browsing (by default, this is TCP port 80; see Managing Advanced Service Configuration Options, page 10-39).

User defined detectors can also have one or more of the following attributes:

- IP address list—Limit detection to the listed IP address ranges. This applies to the source IP when detecting IP sweeps and port scans. It applies to the destination IP when detecting DoS and DDoS attacks.

- TCP port list—Limit detection to the listed destination TCP ports. This list is applied to TCP Specific Ports anomaly types only.

- UDP port list—Limit detection to the listed destination UDP ports. This list is applied to UDP Specific Ports anomaly types only.

How to View Anomaly Detection Settings

You can view a list of all anomaly detectors. The anomaly detectors are displayed in a tree, grouped according to detector category (Scan/Sweep, DoS, or DDoS).

For each anomaly detector you can view its associated parameters and see a list of all anomaly types included in the detector, together with their parameters.

Step 1 In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**
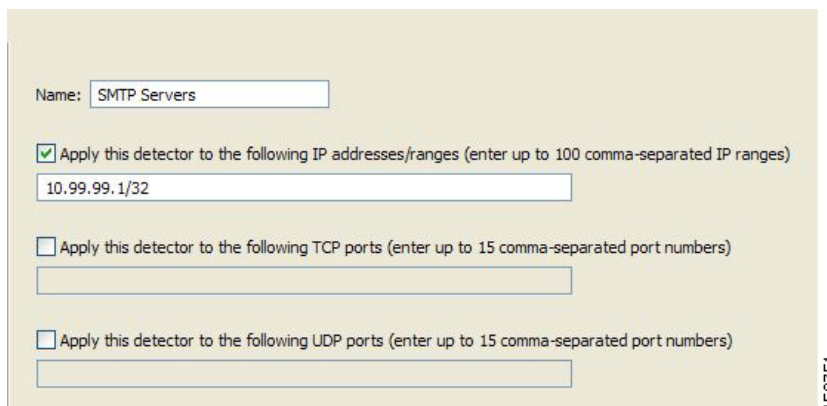
The Anomaly Detection Settings dialog box appears.

The detector tree is displayed in the left area of the dialog box; the right area is empty.
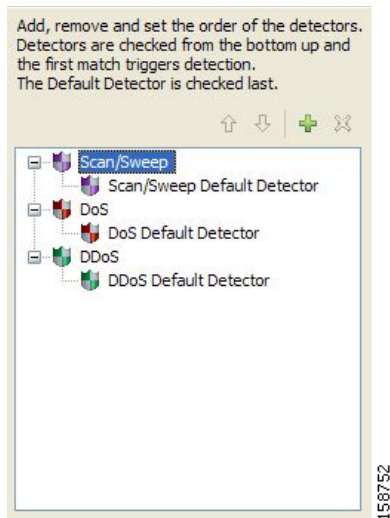
Step 2   In the detector tree, select a detector.

The detector parameters are displayed in the upper right area of the dialog box.

The detector's defined anomaly types are listed in the lower right area of the dialog box, together with the value of each parameter. The following screen capture shows the default parameter values for the Scan/Sweep default detector.

Initiating Side	Session Rate	Suspected Session Rate	Suspected Session Ratio	Alert User	Notify Subscriber	Block Attack
Network						
TCP	1000	500	50	Disable	Disable	Disable
TCP Specific Ports	1000	500	50	Disable	Disable	Disable
UDP	1000	500	50	Disable	Disable	Disable
UDP Specific Ports	1000	500	50	Disable	Disable	Disable
ICMP	500	250	50	Disable	Disable	Disable
Other	500	250	50	Disable	Disable	Disable
Subscriber						
TCP	1000	500	50	Disable	Disable	Disable
TCP Specific Ports	1000	500	50	Disable	Disable	Disable
UDP	1000	500	50	Disable	Disable	Disable
UDP Specific Ports	1000	500	50	Disable	Disable	Disable

Revert   Accept

210799

If unidirectional classification is enabled, the Suspected Session Rate is set equal to the Session Rate, which effectively disables anomaly detection by the suspected session trigger.

Initiating Side	Session Rate	Suspected Session Rate
Network		
TCP	1000	1000
TCP Specific Ports	1000	1000
UDP	1000	1000
UDP Specific Ports	1000	1000
ICMP	500	500
Other	500	500
Subscriber		
TCP	1000	1000
TCP Specific Ports	1000	1000
UDP	1000	1000
UDP Specific Ports	1000	1000

210815

Step 3   Click **OK.**

The Anomaly Detection Settings dialog box closes.

How to Add Anomaly Detectors

You can add new anomaly detectors. A service configuration can contain up to 100 anomaly detectors.

You define IP address ranges and TCP and UDP ports for the new detector, and one anomaly type.

After you have defined the detector, you can add other anomaly types (see Editing Anomaly Detectors, page 10-10).

Step 1   In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

**Step 2**   In the detector tree, select a detector category.

Step 3   Click ![plus icon] .

The Anomaly Detector Creation wizard appears, open to the Malicious Traffic Detector page.
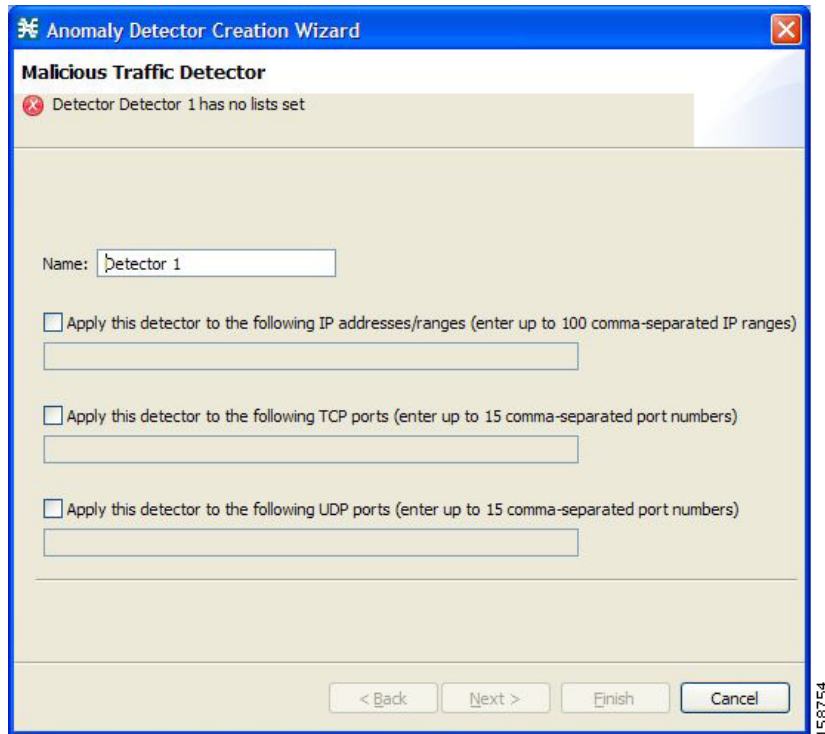
Step 4 In the Name field, enter a meaningful name for the detector.

Step 5 Check one or more of the check boxes to limit the scope of the detector.

The relevant fields are enabled.

Step 6 Enter lists of IP addresses or ports in the relevant fields.

Step 7 Click **Next.**

The Malicious Traffic Characteristics for a WORM attack page of the Anomaly Detector Creation wizard opens.
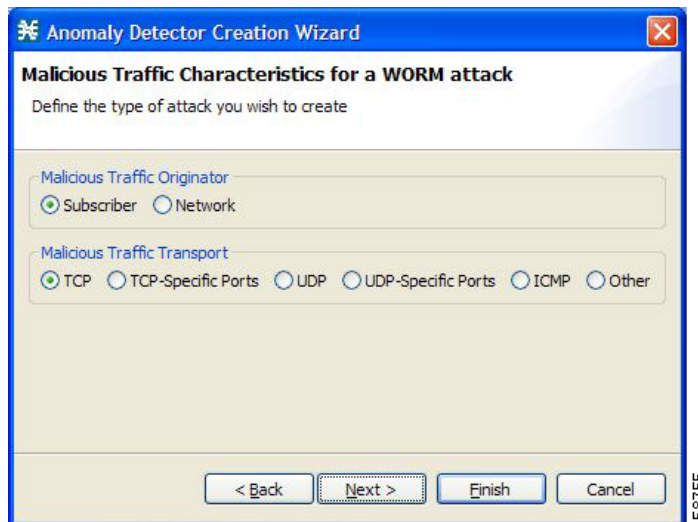
Step 8 Depending on the detector type that you are defining, select the originating side or the target side.

- If you are defining a Scan/Sweep detector or a DoS detector, select the originating side for the anomaly type you are defining.

- If you are defining a DDoS detector, select the target side for the anomaly type you are defining.

Step 9   Select a transport type for the anomaly type that you are defining.

Step 10   Click **Next.**

The Anomaly Detection Thresholds page of the Anomaly Detector Creation wizard opens.


Step 11   Set the detector settings for this anomaly type.

Do one of the following:

- To use the default detector's settings, check the **Use the Default Detector's settings** check box.

- Enter values in the Flow Open Rate, Suspected Flows Rate, and Ratio of Suspected Flow Rate fields.

Step 12   Click **Next.**

The Anomaly Detection Action Settings page of the Anomaly Detector Creation wizard opens.
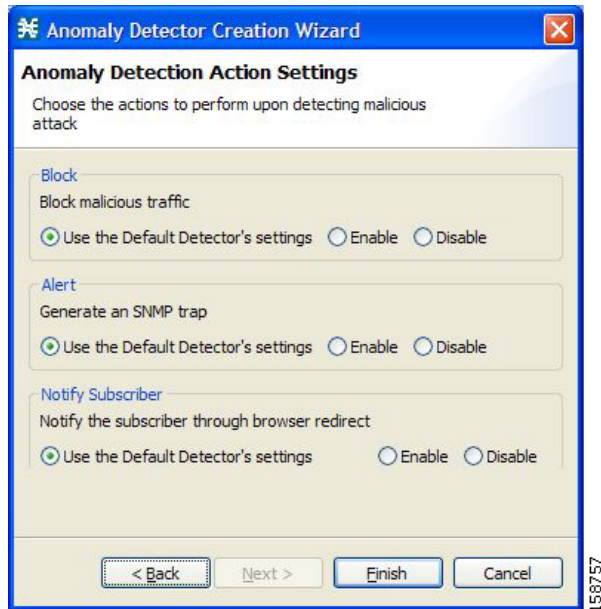
Step 13   Select Block, Alert, and Notify Subscriber actions.

Step 14   Click **Finish.**

The Anomaly Detector Creation wizard closes.

The new detector is added to the detector tree.

What to Do Next

You can now add additional anomaly types to the detector. (See Editing Anomaly Detectors, page 10-10.)

Editing Anomaly Detectors

You can perform the following actions on a user-defined anomaly detector:

- Edit detector parameters.
- Edit anomaly types.
- Add anomaly types.
- Delete anomaly types.
- Change the order of the detectors in the detector tree.

    For each detector category, detectors are checked, *bottom-up,* in the order that they are listed in the detector tree; the default detector is checked last.

You can edit the anomaly types of the three default detectors.

How to Edit Detector Parameters

Step 1   In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

Step 2   In the detector tree, select a detector.

The detector parameters are displayed in the upper right area of the dialog box.

Step 3   In the Name field, enter a new name for the detector.

Step 4   Check or uncheck the IP address range and ports check boxes.

Step 5   Enter or modify lists of IP addresses or ports in the relevant fields.

Step 6   Click **OK.**

The Anomaly Detection Settings dialog box closes.

Your changes are saved.

How to Edit Anomaly Types

Step 1   In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

Step 2   In the detector tree, select a detector.

Information about the anomaly types is displayed in the lower right area of the dialog box.

Step 3   Double-click an anomaly type.

The Anomaly Detector Creation wizard appears, open to the Anomaly Detection Thresholds page (see How to Add an Anomaly Type, page 10-12).

Step 4   Set the detector settings for this anomaly type.

Do one of the following:

- To use the default detector's settings, check the **Use the Default Detector's settings** check box.
- Change the values in the Flow Open Rate, Suspected Flows Rate, and Ratio of Suspected Flow Rate fields.

Step 5   Click **Next.**

The Anomaly Detection Action Settings page of the Anomaly Detector Creation wizard opens.

Step 6   Change Block, Alert, and Notify Subscriber actions.

Step 7   Click **Finish.**

The Anomaly Detector Creation wizard closes.

The anomaly type is updated with your changes.

Step 8   Repeat Steps 3 to 7 (or Steps 2 to 7) for other anomaly types.

Step 9   Click **OK.**

The Anomaly Detection Settings dialog box closes.

How to Add an Anomaly Type

Step 1 In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

Step 2 In the detector tree, select a detector.

The anomaly types are listed in the lower right area of the dialog box.

Step 3 Click (**Create New Detector Item Under Detector Items Feature**).

The Anomaly Detector Creation wizard appears, open to the Malicious Traffic Characteristics for a WORM attack page (see How to Add Anomaly Detectors, page 10-7).

Step 4 Select an origin for the anomaly type you are defining.

Step 5 Select a transport type for the anomaly type you are defining.

Step 6 Click **Next.**

The Anomaly Detection Thresholds page of the Anomaly Detector Creation wizard opens.

Step 7 Set the detector settings for this anomaly type.

Do one of the following:

- To use the default detector's settings, check the **Use the Default Detector's settings** check box.
- Enter values in the Flow Open Rate, Suspected Flows Rate, and Ratio of Suspected Flow Rate fields.

Step 8 Click **Next.**

The Anomaly Detection Action Settings page of the Anomaly Detector Creation wizard opens.

Step 9 Select Block, Alert, and Notify Subscriber actions.

Step 10 Click **Finish.**

The Anomaly Detector Creation wizard closes.

The new anomaly type is added to the anomaly type list.

Step 11 Repeat Steps 3 to 10 (or Steps 2 to 10) for other anomaly types.

Step 12 Click **OK.**
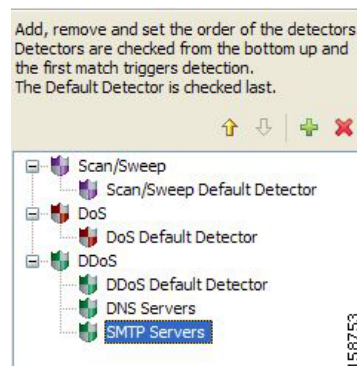
The Anomaly Detection Settings dialog box closes.

How to Delete an Anomaly Type

Step 1 In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

Step 2 In the detector tree, select a detector.

The anomaly types are listed in the lower right area of the dialog box.

Step 3 In the anomaly type list, select an anomaly type.

Step 4   Click ✖.

The selected anomaly type is deleted from the anomaly type list.

Step 5   Repeat Steps 3 and 4 (or Steps 2 to 4) for other anomaly types.

Step 6   Click **OK.**

The Anomaly Detection Settings dialog box closes.

How to Change the Order in which Detectors are Checked

Step 1   In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

Step 2   In the detector tree, select a detector.

The move up arrow, the move down arrow, or both are enabled, depending on the detectors location in the tree.

Step 3   Using these navigation arrows, move the detector to its desired location.

Step 4   Repeat Steps 2 and 3 for other detectors.

Step 5   Click **OK.**

The Anomaly Detection Settings dialog box closes.

Your changes are saved.

How to Delete Anomaly Detectors

You can delete any or all user-defined detectors.

You cannot delete the three default detectors.

Step 1   In the Service Security Dashboard, in the Anomaly Based Detection of Malicious Traffic pane, click **Configure.**

The Anomaly Detection Settings dialog box appears.

Step 2   In the detector tree, select one or more user-defined detectors.

Step 3   Click ![X].

A Confirm Delete message appears.
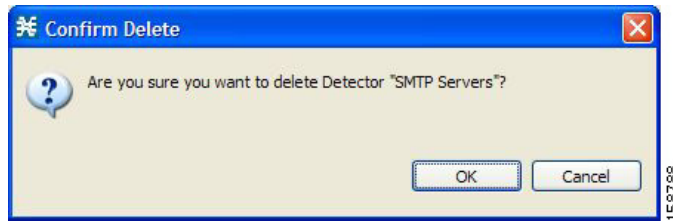
Step 4   Click **OK.**

The selected detectors are deleted and are no longer displayed in the detector tree.

Step 5   Click **OK.**

The Anomaly Detection Settings dialog box closes.

Managing Spam Detection

The anomalous e-mail detection method monitors SMTP session rates for individual subscribers. A high rate of SMTP sessions from an individual subscriber is usually an indicator of malicious activity that involves sending e-mail (either mail-based viruses or spam-zombie activity).

This method will work only if the system is configured in subscriber-aware or anonymous subscriber mode. This allows the SCE to accurately account the number of SMTP sessions generated per subscriber.

The detection method is based on the following:

- Typical broadband subscribers generate a small number of SMTP sessions (at most a single session each time they send an e-mail message).
- Typical broadband subscribers normally use the ISP's SMTP server (as configured in their mail client) as their only mail relay, and do not communicate with off-net SMTP servers.
- Spam zombies create many SMTP sessions, mainly to off-net servers (the mail servers of the destined recipient of the messages).

When configuring spam detection, you select an appropriate service to monitor. By default, this is the built-in SMTP service. To improve detection sensitivity, you can create a more specific service to narrow the scope of detection. Two possible services are:

- "Outbound SMTP"—SMTP sessions generated by the subscriber.
- "OffNet SMTP"—SMTP sessions that are *not* targeted to the SMTP server of the subscriber's ISP. Limiting the service to OffNet can avoid accounting for legitimate sessions.

Note   Prominent non-ISP e-mail providers (for example, Google and Yahoo!) now provide SMTP-based service, so OffNet is no longer a very good differentiator between legitimate and illegitimate activity. To refine the OffNet service, you must include an SMTP server list in the "OnNet SMTP" service definition; all other SMTP servers will be OffNet.

How to Configure Spam Detection Settings

Step 1   In the Service Security Dashboard, in the Spam Zombies and Email Viruses Detection pane, click **Configure.**

The Spam Setting dialog box appears.

Step 2   (Optional) To disable spam detection, uncheck the **Enable Spam detection** check box.

All other fields are disabled.

Continue at Step 7.

Step 3   From the Service to monitor for Spam drop-down list, select a service.

✎
Note   Leave the default value for the monitored service (SMTP), unless you have defined a more specific service, such as "Outbound SMTP" or "OffNet SMTP".

Step 4   Define the threshold e-mail session rate for anomalous behavior.

Step 5   From the Action upon detection drop-down list, select the action to be taken when malicious activity is detected.

Available actions are:

- **Ignore**
- **Block**
- **Notify**
- **Block and notify**

If you select Notify or Block and notify, the Subscriber Notification drop-down list is enabled.

Step 6   If you selected Notify or Block and notify, select a subscriber notification.

Note To define an appropriate subscriber notification, see Managing Subscriber Notifications, page 10-27.

Step 7 Click **Finish.**

The Spam Setting dialog box closes.

Viewing Malicious Traffic Reports

Information about detected traffic anomalies is stored in the Collection Manager database. You can use this information for network trending, detection of new threats, and tracking of malicious hosts or subscribers.

- Malicious Traffic Reports, page 10-16
- How to View a Service Security Report, page 10-16

Malicious Traffic Reports

A number of reports dealing with malicious traffic can be displayed in the Reporter tool:

- Global reports:
  - Global Scan or Attack Rate
  - Global DoS Rate
  - Infected Subscribers
  - DoS Attacked Subscribers
  - Top Scanned or Attacked ports
- Individual subscriber or hosts reports:
  - Top Scanning or Attacking hosts
  - Top DoS Attacked hosts
  - Top DoS Attacked Subscribers
  - Top Scanning or Attacking Subscribers

How to View a Service Security Report

Step 1 In the Service Security Dashboard, in the relevant pane, click **View Report.**

A Choose a report dialog box appears, displaying a tree of relevant reports.

Step 2 Select a report from the report tree.

Step 3 Click **OK.**

The Choose a report dialog box closes.

The Reporter tool opens in the Console, and displays the requested report.

Step 4   For information about manipulating and saving the report, see the "Working with Reports" chapter of the *Cisco Service Control Application Reporter User Guide.*

Filtering the Traffic Flows

Filter rules are part of service configurations. They allow you to instruct the Service Control Engine (SCE) platform, based on a flow's Layer 3 and Layer 4 properties, to:

- Bypass—Ignore the flow and transmit it unchanged.
- Quick forward—Duplicate the flow and send one copy directly to the transmit queue to ensure minimal delay. The second copy goes through the normal packet path. The flow is not controlled, but can be classified and reported.

When a traffic flow enters the SCE platform, the platform checks whether a filter rule applies to this flow.

If a filter rule applies to this traffic flow, the SCE platform passes the traffic flow to its transmit queues. No RDR generation or service configuration enforcement is performed; these flows will not appear in any records generated for analysis purposes and will not be controlled by any rule belonging to the active service configuration.

It is recommended that you add filter rules for OSS protocols (such as DHCP) and routing protocols (such as BGP) that might traverse the SCE platform. These protocols usually should not be affected by policy enforcement, and their low volume makes them insignificant for reporting.

A number of predefined filter rules are included in every new service configuration.

Note   By default, some, but not all, of the predefined filter rules are active.

Flows of certain protocols can also be filtered according to the flow's Layer 7 characteristics. (See Managing Advanced Service Configuration Options, page 10-39.) Like other filtered flows, Layer 7 filtered flows are not controlled, but can be classified and reported. The flows of the protocols that can be filtered are typically short and their overall volume is negligible, so filtering these protocols has little effect on network bandwidth and on the accuracy of the SCA BB reports.

- Information About Traffic Filtering, page 10-17
- How to View Filter Rules for a Package, page 10-19
- How to Add Filter Rules, page 10-20
- How to Edit Filter Rules, page 10-25
- How to Delete Filter Rules, page 10-26
- How to Activate and Deactivate Filter Rules, page 10-26

Information About Traffic Filtering

For certain types of traffic, service providers may need to reduce the latency and jitter introduced by the SCE platform or even to bypass the SCE platform to avoid traffic control as well. Typically, such decisions are made for a portion of the traffic, to reduce latency for delay sensitive applications, such as voice, and to bypass mission-critical traffic, such as routing protocols. The SCA BB Filtered Traffic mechanism is used to address this need.

Note   Most voice traffic is handled automatically by the SCE platform to reduce latency (see Automatic Quick Forwarding of Media Flows, page 10-19).

- The SCA BB Filtered Traffic Mechanism, page 10-18
- Filter Rule Actions, page 10-18
- Filter Rules and Service Rules, page 10-19
- Automatic Quick Forwarding of Media Flows, page 10-19

The SCA BB Filtered Traffic Mechanism

The SCA BB Filtered Traffic mechanism reduces latency or completely bypasses portions of the traffic by defining *filter rules* that match relevant flows and assign the correct action to them. A filter rule matches a packet according to its Layer 3 and Layer 4 properties, such as IP address, port number, and DSCP ToS, as well as the SCE platform interface (subscriber or network) from which the packet arrived. For packets that match a filter rule, the following actions can be applied:

- Bypass the current packet (to reduce latency and avoid traffic control).

  When this action is applied, the current packet is directly transmitted from the SCE platform without going through any service configuration processing or reporting. You must map the bypassed packet to a Class of Service (CoS) to assign it to one of the transmit queues of the SCE platform.

  Possible values for CoS are BE, AF1, AF2, AF3, AF4, and EF; where EF implies high processing priority and the other classes imply normal processing priority.

- Quick forward the flow (to reduce latency).

  When this action is applied, the current packet and all subsequent packets belonging to the same flow are duplicated and sent through two different paths: the original packet goes directly to the transmit queue, and thus has only a minimal delay, while a copy of the packet goes through the normal service configuration processing path for classification and reporting, and is then discarded.

- Assign the flow to the high priority processing input queue (to reduce latency).

Note   Not all platforms support this option.

- When this action is applied, the current packet and all subsequent packets belonging to the same flow enter the high priority processing input queue. They go through the normal service configuration processing path ahead of other packets that arrive simultaneously. You should map the flow to the EF CoS to assign it to the high processing priority transmit queue of the SCE platform.

A filter rule can perform DSCP ToS marking (by changing the DSCP ToS field of the packet) of the matched traffic in conjunction with any of the above actions.

Note   DSCP ToS marking and the assignment to CoS only take place when the operational mode of the system is Full Functionality (see System Operational Mode, page 10-34).

Filter Rule Actions

The Bypass and Quick forward actions apply to different scopes of traffic:

- The Bypass action only bypasses the current packet; every subsequent packet of the same flow goes through the Filtered Traffic mechanism. This means, for example, that when traffic is to be bypassed based on its destination port number, two rules should be created in order to match packets from both sides of a bidirectional flow.

  For example, to bypass all traffic to destination port 23, two filter rules are needed, one for packets arriving from the subscriber side addressed to network side port 23, and another for packets arriving from the network side addressed to subscriber side port 23.

- The Quick forward action is applied to the entire flow; once identified, all subsequent packets do not go through the filter rule mechanism, instead going through normal service configuration processing.

A packet may match more than one filter rule. If these filter rules have conflicting actions, Bypass takes priority over Quick forward.

Filter Rules and Service Rules

Filter rule actions to reduce latency allow the flow to be controlled by the SCE platform. This means that the flow can be blocked or given limited bandwidth if it matches a service rule. For example, if a filter rule is applied to reduce latency, but a service configuration rule is applied to block the same traffic, the traffic will be blocked.

The Bypass action is designed to avoid service configuration processing; bypassed traffic is not affected by service rules.

Automatic Quick Forwarding of Media Flows

The SCE platform reduces the latency of delay-sensitive voice and video media flows by applying the quick-forwarding action to SIP, MGCP, H323, Skinny, and RTSP media flows during classification. That is, when a media flow is classified as being of one of these types, it will be subject to quick forwarding immediately. The SCE platform does this automatically, regardless of filter rule configuration. These media flows might still be blocked or given limited bandwidth if they match a service rule.

How to View Filter Rules for a Package

You can view a list of the filter rules included in a service configuration.

The listing for each filter rule includes the name, the status, and a brief description (generated by the system) of the rule.

To see more information about a filter rule, open the Edit Filter Rule dialog box (see How to Edit Filter Rules, page 10-25).

Step 1   In the Network Traffic tab, select the **Filtered Traffic** node.

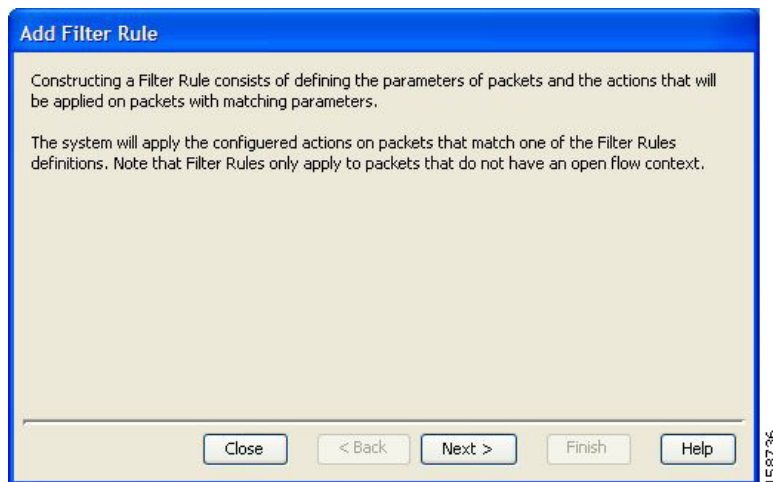A list of all filter rules is displayed in the right (Rule) pane.
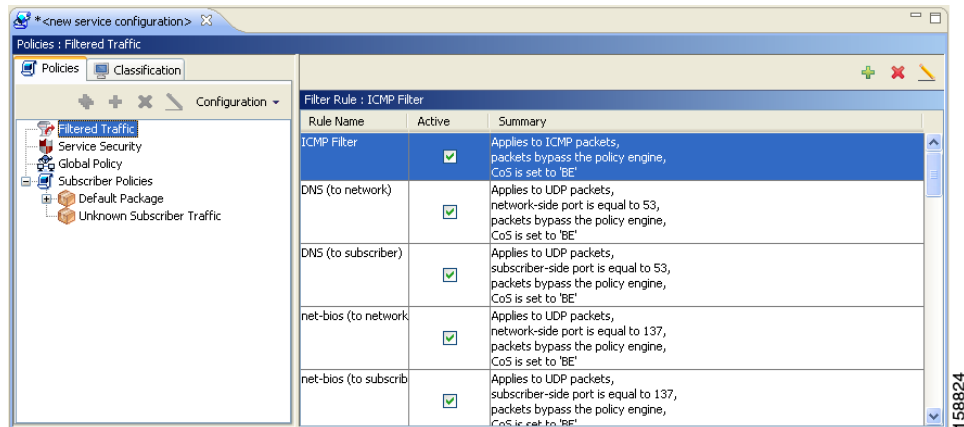
How to Add Filter Rules

The Add Filter Rule wizard guides you through the process of adding a filter rule.

Step 1   In the Network Traffic tab, select the **Filtered Traffic** node.

Step 2   Click (**Add Rule)** in the right (Rule) pane.

The Add Filter Rule wizard appears.

Step 3   Click **Next.**

The Transport Type and Direction page of the Add Filter Rule wizard opens.

Step 4   Select the transport type and initiating side and click **Next.**

The Subscriber-Side IP Address page of the Add Filter Rule wizard opens.

Step 5   Define the subscriber-side IP address and click **Next.**

The Network-Side IP Address page of the Add Filter Rule wizard opens.

Step 6 Define the network-side IP address and click **Next.**

If the transport type selected in Step 4 was *not* TCP or UDP, the ToS page of the Add Filter Rule wizard opens. Go to Step 9.

If the transport type selected in Step 4 was TCP or UDP, the Subscriber-Side Port page of the Add Filter Rule wizard opens.

Step 7 Define the subscriber-side port and click **Next.**

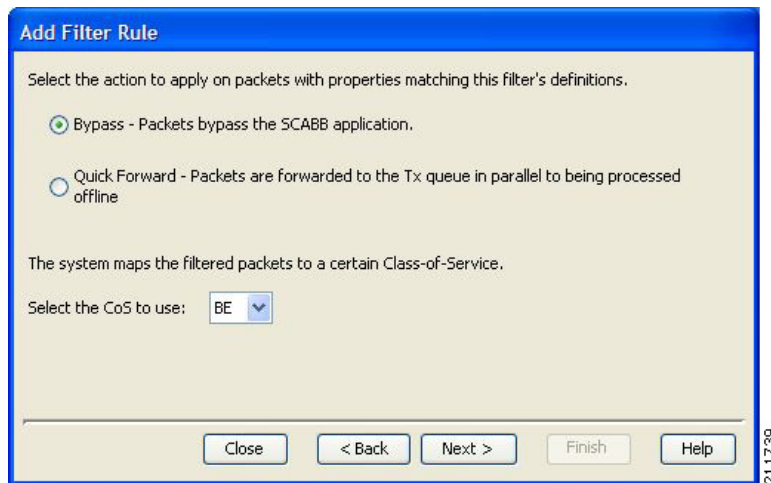The Network-Side Port page of the Add Filter Rule wizard opens.

Step 8   Define the network-side port and click **Next.**

The ToS page of the Add Filter Rule wizard opens.

Step 9   Define the ToS and click **Next.**

Note   Acceptable values for ToS are 0 to 63.

The Action and Class-of-Service page of the Add Filter Rule wizard opens.

Step 10   Select the radio button for the required action.

- **Bypass** —Packets that match this filter rule are not passed to SCA BB.

- **Quick Forward** —The SCE platform ensures low latency for packets that match this filter rule (use for delay sensitive flows). Packets are duplicated and passed to SCA BB for processing.

Step 11   Select a Class-of-Service value, and click **Next.**

The ToS Marking page of the Add Filter Rule wizard opens.
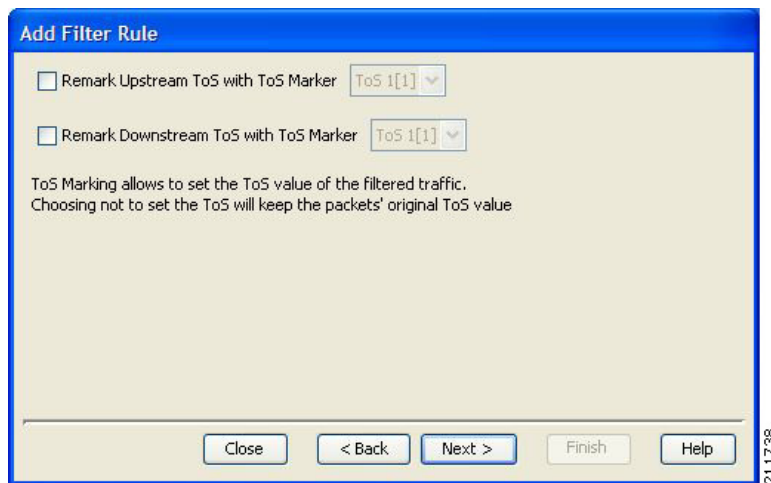
Step 12   (Optional) To change the DSCP ToS marker of packets in the filtered traffic, check the **Remark Upstream ToS with ToS Marker** and **Remark Downstream ToS with ToS Marker** check boxes, as required, select the required ToS marker from the drop-down list, and click **Next.**

- Disabling directional DSCP ToS marking in the ToS Marking Settings dialog box (see How to Manage DSCP ToS Marker Values, page 9-28) overrides DSCP ToS marking in that direction by a filter (that is, the DSCP ToS value will not be changed). In this case, the Problems View will display a Warning.

- If you filter for a flow in one direction in Step 4 but select ToS marking in the other direction in this Step, the filter rule will be created, but no DSCP ToS remarking will occur. In this case, the Problems View will display a Warning.

- If you select Quick Forward in the previous Step, SCA BB receives the *original* package and processes it. That is, the original DSCP ToS value is seen by the application regardless of the ToS marking action selected in the filter rule.

The Finish page of the Add Filter Rule wizard opens.
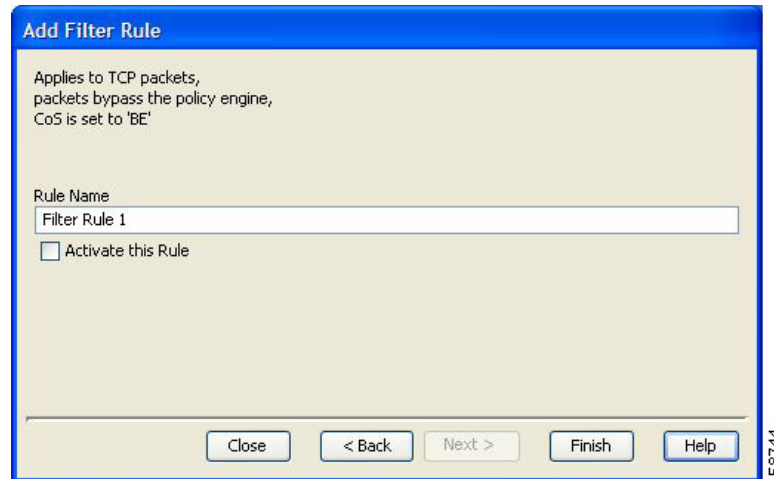
Step 13   In the Rule Name field, enter a unique name for the new filter rule.

Note   You can use the default name for the filter rule. It is recommended that you enter a meaningful name.

Step 14   (Optional) To activate the filter rule, check the **Activate this rule** check box. Traffic is filtered according to the rule only when it is activated.

Step 15   Click **Finish.**

The Add Filter Rule wizard closes.

The filter rule is added and is displayed in the Filter Rule table.

How to Edit Filter Rules

You can view and edit the parameters of a filter rule.

Step 1   In the Network Traffic tab, select the Filtered Traffic node.

A list of all filter rules is displayed in the right (Rule) pane.

Step 2   Select a rule in the Filter Rule table.

Step 3   Click ⬛ (**Edit Rule**).

The Introduction page of the Edit Filter Rule wizard appears.

The Edit Filter Rule wizard is the same as the Add Filter Rule wizard.

**Step 4**   Follow the instructions in the section How to Add Filter Rules, page 10-20, Steps 4 to 14.

Step 5   Click **Finish.**

The filter rule is changed and relevant changes appear in the Filter Rule table.

How to Delete Filter Rules

You can delete filter rules. This is useful, for example, when you want the system to resume handling the IP addresses and their attributes according to the individual rules that were previously defined for each subscriber IP address.

Step 1 In the Network Traffic tab, select the **Filtered Traffic** node.

A list of all filter rules is displayed in the right (Rule) pane.

Step 2 Select a rule in the Filter Rule table.

Step 3 Click ✖ (**Delete Rule**).

A Filter Rule Warning message appears.
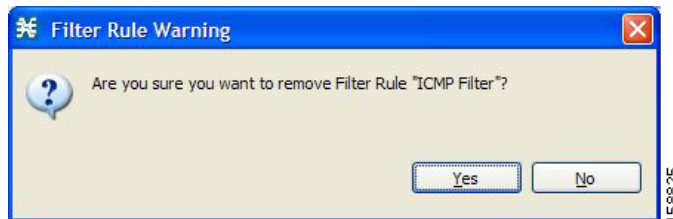
Step 4 Click **Yes.**

The filter rule is deleted and is no longer displayed in the Filter Rule table.

How to Activate and Deactivate Filter Rules

You can activate or deactivate filter rules at any time. Deactivating a filter rule has the same effect as deleting it, but the parameters are retained in the service configuration, and you can reactivate the filter rule at a later date.

Step 1 In the Network Traffic tab, select the **Filtered Traffic** node.

A list of all filter rules is displayed in the right (Rule) pane.

Step 2 Select a rule in the Filter Rule table.

Step 3 To activate the rule, check the **Active** check box.

Step 4 To deactivate the rule, uncheck the **Active** check box.

Step 5 Repeat Steps 3 and 4 for other rules.

Managing Subscriber Notifications

The subscriber notification feature pushes web-based messages to a subscriber by redirecting the subscriber HTTP traffic to relevant web pages. These web pages contain information relevant to the subscriber, such as notifications of quota depletion. HTTP redirection starts when the subscriber notification is activated and ceases when the notification is dismissed.

Note Subscriber notification is not supported when unidirectional classification is enabled.

The Cisco Service Control Application for Broadband (SCA BB) supports a maximum of 31 subscriber notifications, including the default notification and the Network Attack Notification.

- Subscriber Notification Parameters, page 10-27
- Network Attack Notification, page 10-29
- How to View Subscriber Notifications, page 10-30
- How to Add Subscriber Notifications, page 10-31
- How to Edit Subscriber Notifications, page 10-32
- How to Delete Subscriber Notifications, page 10-32

Subscriber Notification Parameters

A subscriber notification is defined by the following parameters:

- Name—Each subscriber notification must have a unique name.

Note You cannot change the name of the Default Notification or the Network Attack Notification.

- Destination URL—A configurable destination URL to which the subscriber's HTTP flows are redirected after redirection is activated. This web page usually contains the message that needs to be conveyed to the subscriber.
- Notification Parameters—The query part of the destination URL, which can be optionally added upon redirection.

  The format of the notification parameters to be added to the destination URL is:

  ```
  ?n=<notification-ID>&s=<subscriber-ID>
  ```
  where **<notification-ID>** is the ID of the notification that redirected the subscriber and **<subscriber-ID>** is the subscriber name.

Note There is a different format for the Network Attack Notification Parameters, page 10-29.

- The destination web server can use these parameters to carry a more purposeful message to the subscriber.
- Dismissal method—Indicates when to dismiss or deactivate the notification state. The dismissal method is one of the following:
  - Subscriber browses to destination URL (default)—As soon as the subscriber browses to the destination URL, they are considered as notified and the notification state is dismissed.

For example, if a quota was exceeded, the notification state is dismissed as soon as the subscriber browses to the destination URL that informs them of this fact (even though the subscriber still remains in a breach state).

– The condition that activated the notification no longer holds—The dismissal of the notification state is dependent on the resolution of the condition, rather than on the subscriber.

For example, if a quota was exceeded, the notification state is dismissed only when the subscriber completes the procedure to refresh their quota.

Note This option is *not* available for the Network Attack Notification. A subscriber must respond to the notification before the notification is dismissed.

– Subscriber browses to dismissal URL—The notification state is not dismissed until the subscriber proceeds from the destination URL to a different, final URL.

All HTTP flows are redirected until the notification is dismissed, which takes place when the subscriber accesses the dismissal URL. By default, the destination URL is also the dismissal URL and a notification is dismissed as soon as the first redirection takes place. However, you can define a different dismissal URL, so that the subscriber must acknowledge the notification.

For example, if a quota was exceeded, the web page at the destination URL may ask the subscriber to press an **Acknowledge** button after reading the message. The acknowledge URL would be defined as the dismissal URL and would deactivate further notifications.

The dismissal URL is composed of the URL hostname and the URL path, separated by a colon, in the following format:

`[*]<hostname>:<path>[*]`

– **<hostname>** may optionally be preceded by a wildcard (*), to match all hostnames with the same suffix.

– The path element must always start with "/".

– **<path>** may be followed by a wildcard (*), to match all paths with a common prefix.

For example, the entry ***.some-isp.net:/redirect/*** matches all the following URLs:

– www.some-isp.net/redirect/index.html

– support.some-isp.net/redirect/info/warning.asp

– noquota.some-isp.net/redirect/acknowledge.aspx?ie=UTF-8

• List of Allowed URLs—A list of URLs that will not be blocked and redirected even though redirection is activated.

After redirection is activated, all HTTP flows, except flows to the destination URL and to the dismissal URL, are blocked and redirected to the destination URL. However, subscribers can be permitted to access an additional set of URLs. This is useful, for example, to give subscribers access to additional support information.

Allowed URLs have the same format as the dismissal URL.

These parameters are defined when you add a new subscriber notification (see How to Add Subscriber Notifications, page 10-31). You can modify them at any time (see How to Edit Subscriber Notifications, page 10-32).

Network Attack Notification

Subscriber notification informs a subscriber in real-time about current attacks involving IP addresses mapped to that subscriber. (Enabling these notifications is described in The Service Security Dashboard, page 10-1.) SCA BB notifies the subscriber about the attack by redirecting HTTP flows originating from the subscriber to a server that supplies information about the attack.

One subscriber notification, Network Attack Notification, is dedicated to providing these notifications; it cannot be deleted. A Network Attack Notification is not dismissed at the end of an attack; subscribers *must* respond to it.

To allow redirection when blocking traffic, the system is configured to leave open one specified TCP port (by default, port 80). See Managing Advanced Service Configuration Options, page 10-39.

⚠
Caution   In earlier releases of SCA BB, configuring network attack notifications was performed using CLI commands. CLI commands should no longer be used for this purpose.

- Network Attack Notification Parameters, page 10-29
- Example of URL with Description Tail, page 10-30

Network Attack Notification Parameters

When a network attack is detected, HTTP flows of the subscriber are redirected to a configurable destination URL. This web page should display the warning that needs to be conveyed to the subscriber.

Optionally, the destination URL can include a query part containing notification parameters. The destination web server can use these parameters to create a more specific warning to the subscriber.

The query part of the URL has the following format:

```
?ip=<ip>&side=<side>&dir=<dir>&prot=<protocol>&no=<open-flows>&nd=<suspected-flows>&to=<open-flows-threshold>&td=<suspected-flows-threshold>&ac=<action>&nh=>handled-flows>
```
The meaning of each field in the tail is described in the following table:

Table 10-1   Description Tail Fields

| Field	Description	Possible Values
ip	Detected IP address	
side		• s—Subscriber • n—Network
dir		• s—Source • d—Destination
protocol		• TCP • UDP • ICMP • OTHER
open-flows	Number of open flows	
suspected flows	Number of attack-suspected flows	

Table 10-1   Description Tail Fields (continued)

Field	Description	Possible Values
open-flows-threshold	Threshold for open flows	
suspected-flows-threshold	Threshold for attack-suspected flows	
action		• R—Report • B—Block and report
handled-flows	Number of flows handled since the attack began (Non-zero only during and at the end of an attack)	

Example of URL with Description Tail

```
http://www.some-isp.net/warning?ip=80.178.113.222&side=s&proto=TCP&no=3
4&nd=4&to=34&td=10&ac=B&nh=100
```

How to View Subscriber Notifications

Step 1   From the Console main menu, choose **Configuration > Subscriber Notifications.**

The Subscriber Notifications Settings dialog box appears.
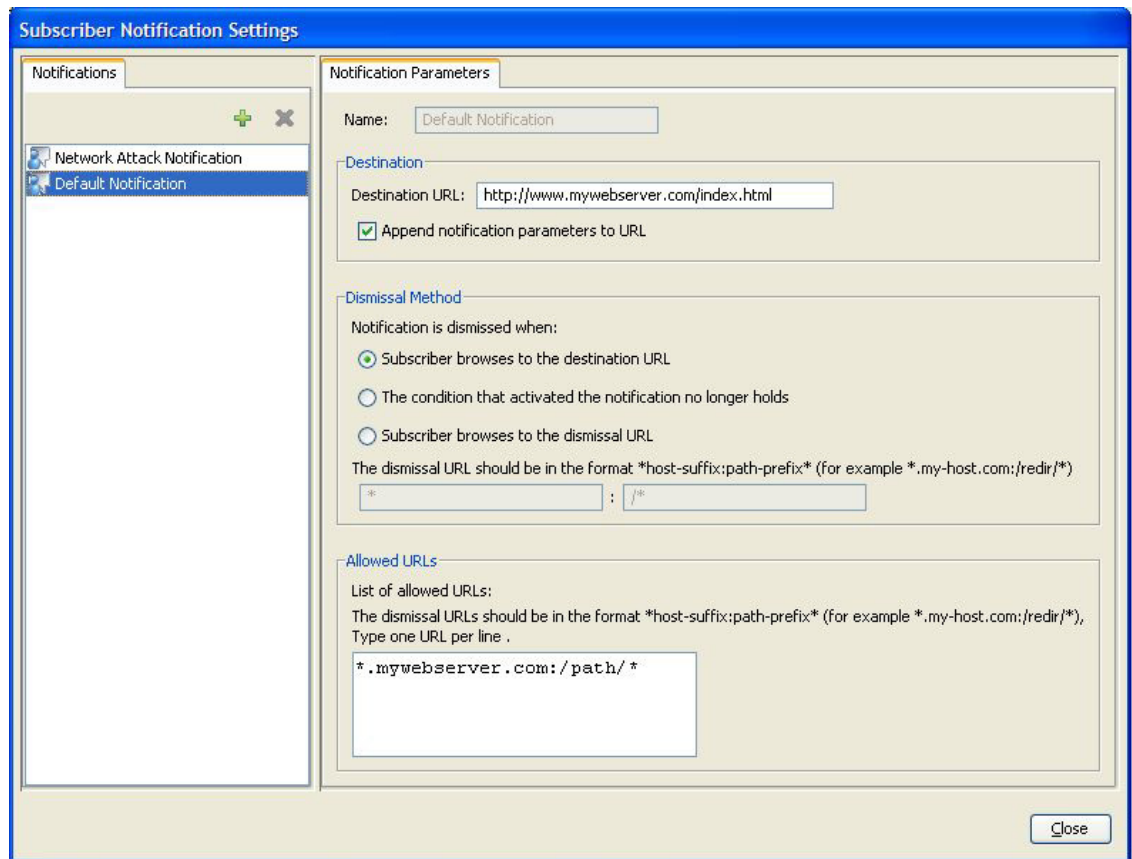
The Notifications tab displays a list of all subscriber notifications.

Step 2   Click a subscriber notification in the list to display its parameters.

The parameters of the subscriber notification are displayed in the Notification Parameters tab.

Step 3   Click **Close.**

The Subscriber Notifications Settings dialog box closes.

How to Add Subscriber Notifications

You can add up to 29 subscriber notifications to a service configuration.

> **Note**   Creating a subscriber notification does not activate the subscriber notification feature. After the subscriber notification is defined, it must be activated for a particular package. (See How to Edit Breach-Handling Parameters for a Rule, page 9-53.)

Step 1   From the Console main menu, choose **Configuration > Subscriber Notifications.**

The Subscriber Notifications Settings dialog box appears.

Step 2   Click [+] (**Add**).

Step 3 In the Name field, enter a unique name for the new subscriber notification.

Note You can use the default name for the subscriber notification. It is recommended that you enter a meaningful name

Step 4 In the Destination URL field, enter the destination URL.

Step 5 (Optional) If notification parameters will be appended to the destination URL, check the **Append notification parameters to URL** check box.

Step 6 Select one of the **Dismissal Method** radio buttons.

- **Subscriber browses to the destination URL**
- **The condition that activated the notification no longer holds**
- **Subscriber browses to the dismissal URL**

Step 7 If you selected Subscriber browses to the dismissal URL in Step 6, enter the dismissal URL host-suffix and path-prefix in the fields provided.

Step 8 Enter any allowed URLs, one per line, in the Allowed URLs text box.

Step 9 Click **Close.**

The Subscriber Notifications Settings dialog box closes.

How to Edit Subscriber Notifications

You can modify notification parameters at any time.

Step 1 From the Console main menu, choose **Configuration > Subscriber Notifications.**

The Subscriber Notifications Settings dialog box appears.

Step 2 Click a subscriber notification in the Notifications tab to display its parameters.

Step 3 Edit the parameters of the subscriber notification in the Notification Parameters tab.

Step 4 Click **Close.**

The Subscriber Notifications Settings dialog box closes.

How to Delete Subscriber Notifications

You can delete subscriber notifications at any time.

You cannot delete the default notification or the Network Attack Notification.

Step 1 From the Console main menu, choose **Configuration > Subscriber Notifications.**

The Subscriber Notifications Settings dialog box appears.

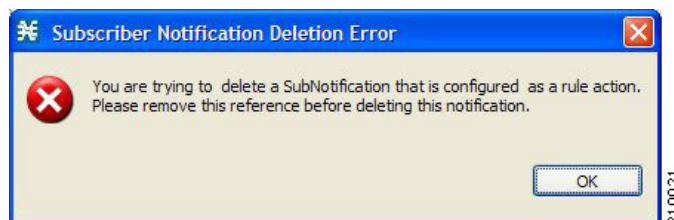**Step 2** Click a subscriber notification in the Notifications tab.

Step 3 Click ✖ (**Delete)**.

A Subscriber Notification Confirmation message appears.
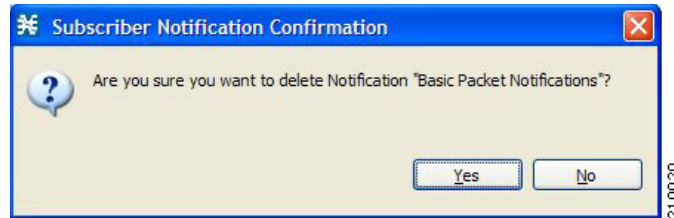
Step 4   Click **Yes.**

- If the specified subscriber notification is being used by a rule, a Subscriber Notification Deletion Error message is displayed.

Note   The subscriber notification cannot be deleted until you unassign it or deactivate it in all service rules. (See How to Edit Breach-Handling Parameters for a Rule, page 9-53.)

- The selected subscriber notification is deleted.

Step 5   Click **Close.**

The Subscriber Notifications Settings dialog box closes.

Managing the System Settings

The Console allows you to determine various system parameters that control:

- The operational state of the system
- Enabling and disabling asymmetric routing classification mode
- The redirection URLs for protocols that support redirection
- BW prioritization mode (see How to Set BW Management Prioritization Mode, page 9-43)
- Advanced service configuration options

Setting the System Modes

From the Console you can select:

- The operational mode of the system
- Asymmetric routing classification mode

Information About the System Modes

System Operational Mode

The operational mode of the system defines how the system handles network traffic.

Note   Each rule has its own operational mode (state). If this differs from the system mode, the "lower" of the two modes is used. For example, if a rule is enabled, but the system mode is report-only, the rule will only generate RDRs.

The three operational modes are:

- Full Functionality—The system enforces active rules on the network traffic and performs reporting functions (that is, generates RDRs).

- Report Only—The system generates RDRs only. No active rule enforcement is performed on the network traffic.

- Transparent—The system does not generate RDRs and does not enforce active rules on the network traffic.

Asymmetric Routing Classification Mode

Enabling unidirectional classification significantly improves classification accuracy when the SCE platform is deployed in an environment with a high rate of unidirectional flows.

Unsupported Features

The following SCA BB features are not supported when unidirectional classification is enabled:

- Flavors

- External quota provisioning

- Subscriber notification

- Redirection

- Flow Signaling RDRs

- Content filtering

- VAS traffic forwarding

- Anomaly detection (When you create a service configuration in asymmetric routing classification mode, the Suspected Session Rate is set equal to the Session Rate for all anomaly detectors (see How to View Anomaly Detection Settings, page 10-5). This effectively disables anomaly detection.)

When unidirectional classification is enabled, the service configuration editor indicates (in the Problems View) if the service configuration is consistent with the features that are supported in this mode.

The following features, which are not part of the service configuration, are also affected when unidirectional classification is enabled:

- Subscriber-Aware Mode (a mode in which subscriber information is dynamically bound to the IP address currently in use by the subscriber) is not supported.

- Enhanced flow open mode must be enabled.

The system gives no indication if the state of the above features is consistent with the state of the routing classification mode.

Protocol Classification

When unidirectional classification is enabled, protocol classification is performed in the normal way with the exception of unidirectional UDP flows. Because it is impossible to know the server side of a unidirectional UDP flow, SCA BB tries to classify the protocol using the destination port of the first packet; if no exact match is found, SCA BB tries to classify the protocol using the source port.

Switching to Asymmetric Routing Classification Mode

If you create a service configuration in symmetric mode and switch to asymmetric routing classification mode:

- Flavors are not used for classification.

- Periodic quota management mode is used.

- Data is not lost when you switch to asymmetric routing classification mode, but you cannot apply the service configuration to an SCE platform until all unsupported features are removed from the service configuration.
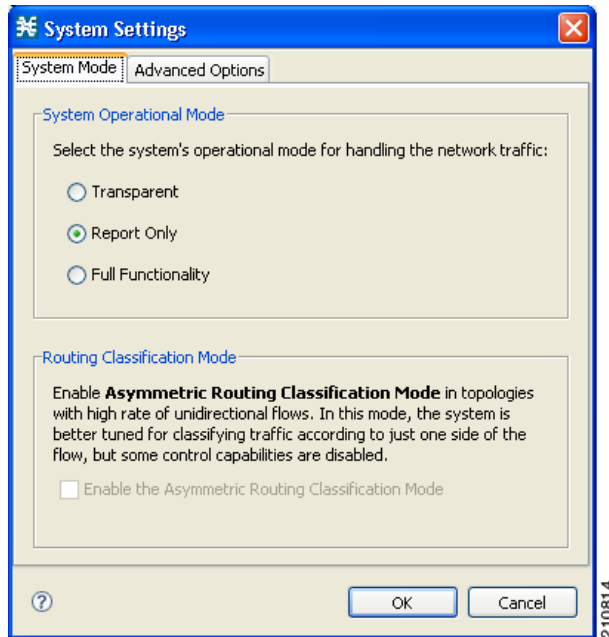
Switching from Asymmetric Routing Classification Mode

If you create a service configuration in asymmetric routing classification mode:

- The Suspected Session Rate is set equal to the Session Rate for all anomaly detectors.

- No flavors are created in the default service configuration, and no service elements have specified flavors.

- The quota management mode is periodic, with a daily aggregation period.

- Asymmetric routing classification mode limitations remain if you switch to symmetric mode. To change them, you must edit the service configuration.

How to Set the Operational and Topological Modes of the System

Step 1   From the Console main menu, choose **Configuration > System Settings.**

The System Settings dialog box appears.

Step 2   Select one of the **System Operational Mode** radio buttons:

- **Transparent**
- **Report Only**
- **Full Functionality**

Step 3   To change the routing classification mode, check or uncheck the **Enable the Asymmetric Routing Classification Mode** check box.

Step 4   Click **OK.**

The System Settings dialog box closes.

The new System Mode setting is saved.

Setting Redirection Parameters

The rules for a package may deny access to selected protocols. When a subscriber to the package tries to access a blocked protocol, the traffic flow can be redirected to a server where a posted web page explains the reason for the redirection (for example, a "Silver" subscriber trying to access a service available only to "Gold" subscribers). This web page can offer subscribers the opportunity to upgrade their packages. You configure which redirection set to use when defining rules (see How to Define Per-Flow Actions for a Rule, page 9-14).

![note pencil icon]

Note   Redirection is not supported when unidirectional classification is enabled.

The Console Redirection feature supports only three protocols:

- HTTP Browsing
- HTTP Streaming

  • RTSP Streaming

Each redirection set contains one redirection option for each of these three protocols. The system provides a default redirection set, which cannot be deleted. You can add up to 49 additional sets.

Each redirection URL includes the URL specified name, the Subscriber ID, and the Service ID in the following format:

```
<URL>?n=<subscriber-ID>&s=<service-ID>
```
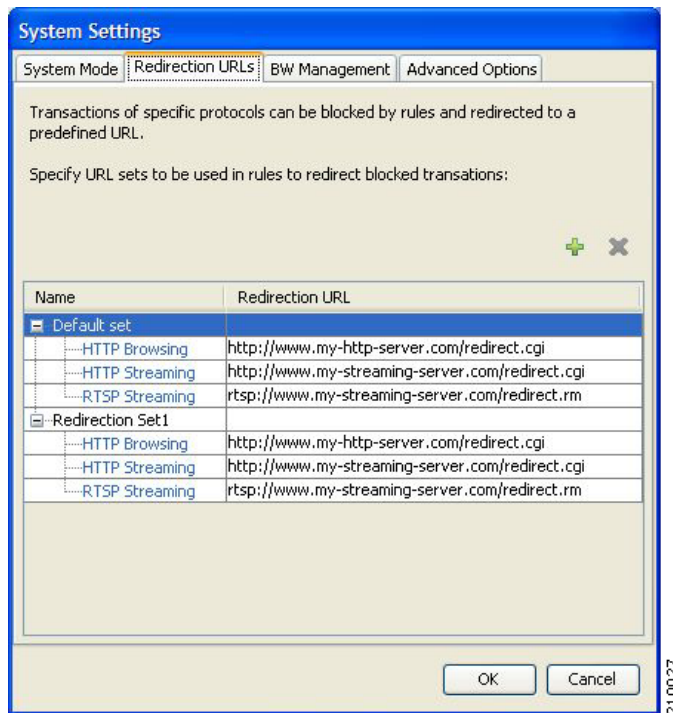
How to Add a Set of Redirection URLs

You can add up to 49 redirection sets.

Step 1   From the Console main menu, choose **Configuration > System Settings.**

The System Settings dialog box appears.

Step 2   Click the **Redirection URLs** tab.

The Redirection URLs tab opens.
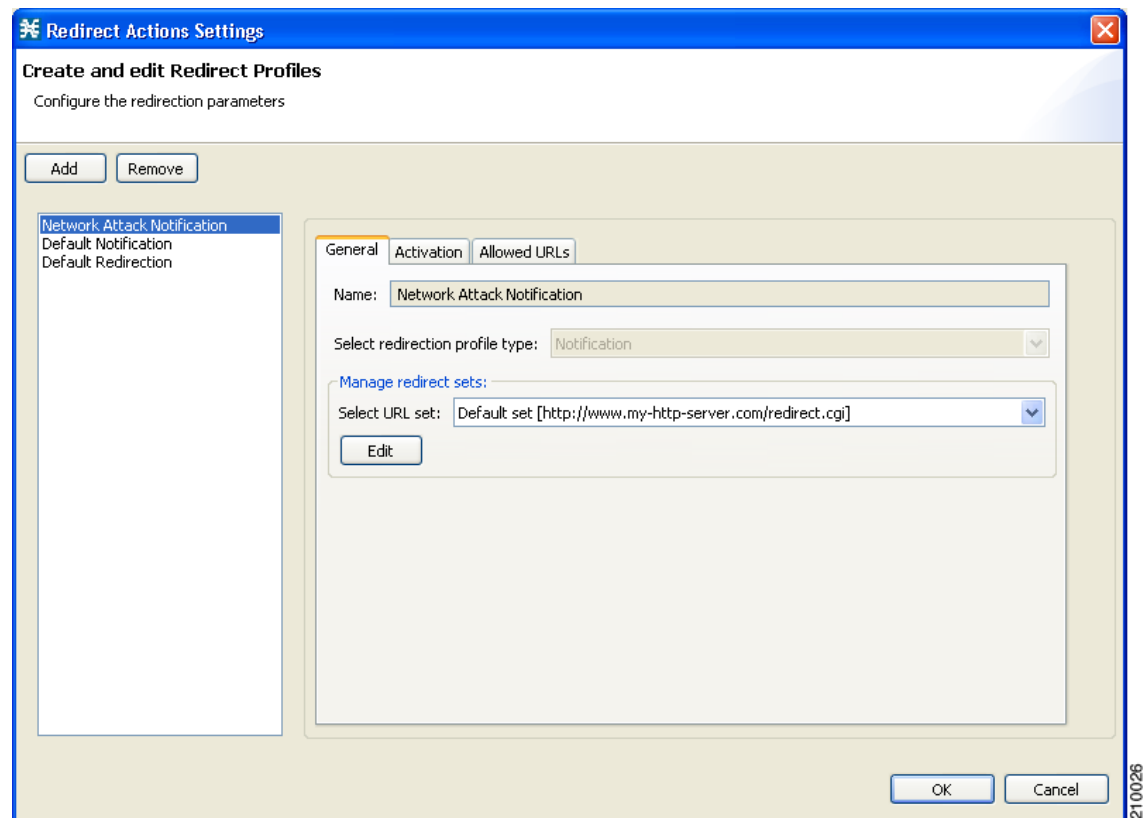
Step 3   Click (**Add**).

A new redirection set containing the default redirection URLs is added to the redirection set list.

Step 4 In the Name field, enter a unique name for the new redirection set.

✎

Note You can use the default name for the redirection set, but it is recommended that you enter a meaningful name.

Step 5 Enter new values in the Redirection URL cells of the new redirection set.

Step 6 Click **OK.**

The System Settings dialog box closes.

The Redirection group is added to the redirection set list.

How to Edit Redirection Parameters

Step 1 From the Console main menu, choose **Configuration > System Settings.**

The System Settings dialog box appears.

Step 2 Click the **Redirection URLs** tab.

The Redirection URLs tab opens.

Step 3 Click a URL in the **Redirection URL** column.

Step 4 Enter a new URL.

Step 5 Click **OK.**

The System Settings dialog box closes.

The Redirection settings are saved.

How to Delete a Set of Redirection URLs

Step 1   From the Console main menu, choose **Configuration > System Settings.**

The System Settings dialog box appears.

Step 2   Click the **Redirection URLs** tab.

The Redirection URLs tab opens.

Step 3   Click the name of a redirection set.

Step 4   Click ✖ (**Delete**).

A Redirection Warning message appears.
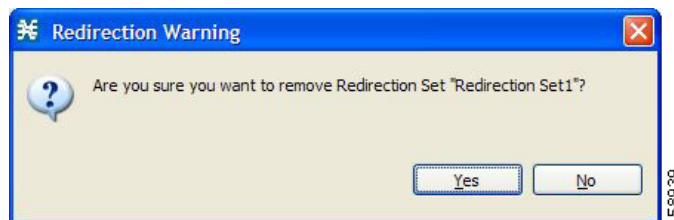
Step 5   Click **Yes.**

The redirection set is deleted.

Step 6   Click **OK.**

The System Settings dialog box closes.

The Redirection settings are saved.

Managing Advanced Service Configuration Options

Advanced service configuration options control the more sophisticated and less frequently changed attributes of the system. It is recommended that you do not change these options.

- The Advanced Service Configuration Properties, page 10-39
- How to Edit Advanced Service Configuration Options, page 10-44

The Advanced Service Configuration Properties

The following table lists the advanced service configuration properties:

Table 10-2   Advanced Service Configuration Properties

Property	Default Value	Description
Classification		
Guruguru detailed inspection mode enabled	FALSE	The Guruguru protocol is used by the Guruguru file-sharing application popular in Japan. SCA BB provides two inspection modes for classification of this protocol: • Default—Suitable for networks where little Guruguru traffic is expected. This is usual in all countries except Japan. • Detailed—Suitable for networks where Guruguru traffic is expected to be common. This should occur in Japanese networks only.
Kuro detailed inspection mode enabled	FALSE	The Kuro protocol is used by the Kuro file-sharing application popular in Japan. SCA BB provides two inspection modes for classification of this protocol: • Default—Suitable for networks where little Kuro traffic is expected. This is usual in all countries except Japan. • Detailed—Suitable for networks where Kuro traffic is expected to be common. This should occur in Japanese networks only.

Table 10-2   Advanced Service Configuration Properties (continued)

Property	Default Value	Description
Soribada detailed inspection mode enabled	FALSE	The Soribada protocol is used by the Soribada file-sharing application popular in Japan. SCA BB provides two inspection modes for classification of this protocol: • Default—Suitable for networks where little Soribada traffic is expected. This is usual in all countries except Japan. • Detailed—Suitable for networks where Soribada traffic is expected to be common. This should occur in Japanese networks only.
TCP destination port signatures	1720:H323	TCP destination port numbers for signatures that require a port hint for correct classification. Valid values are comma-separated items, each item in the form <port-number>:<signature-name>. Applicable signature names are: H323, Radius Access, Radius Accounting, and DHCP.
UDP destination port signatures	67:DHCP, 1812:Radius Access, 1645:Radius Access, 1813:Radius Accounting, 1646:Radius Accounting	UDP destination port numbers for signatures that require a port hint for correct classification. Valid values are comma-separated items, each item in the form <port-number>:<signature-name>. Applicable signature names are: H323, Radius Access, Radius Accounting, and DHCP.
UDP ports for which flow should be opened on first packet	5060, 5061, 67, 69, 1812, 1813, 1645, 1646, 2427, 2727, 9201, 9200, 123, 1900, 5190, 10000	Enhanced flow-open mode is disabled on the specified UDP ports, to allow classification according to the flow's first packet.

Table 10-2   Advanced Service Configuration Properties (continued)

Property	Default Value	Description
UDP source port signatures	1812:Radius Access, 1645:Radius Access, 1813:Radius Accounting, 1646:Radius Accounting	UDP source port numbers for signatures that require a port hint for correct classification. Valid values are comma-separated items, each item in the form <port-number>:<signature-name>. Applicable signature names are: H323, Radius Access, Radius Accounting, and DHCP.
V-Share detailed inspection mode enabled	FALSE	The V-Share protocol is used by the V-Share file-sharing application popular in Japan. SCA BB provides two inspection modes for classification of this protocol: • Default—Suitable for networks where little V-Share traffic is expected. This is usual in all countries except Japan. • Detailed—Suitable for networks where V-Share traffic is expected to be common. This should occur in Japanese networks only.
Winny detailed inspection mode enabled	FALSE	The Winny P2P protocol is used by the Winny file-sharing application popular in Japan. SCA BB provides two inspection modes for classification of this protocol: • Default—Suitable for networks where little Winny traffic is expected. This is usual in all countries except Japan. • Detailed—Suitable for networks where Winny traffic is expected to be common. This should occur in Japanese networks only.
Malicious Traffic		

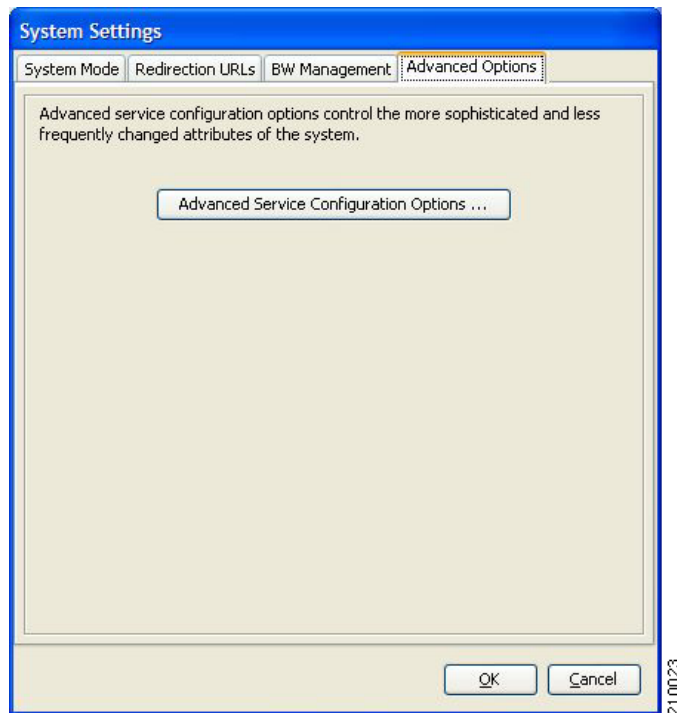*Table 10-2   Advanced Service Configuration Properties (continued)*

Property	Default Value	Description
Malicious Traffic RDRs enabled	TRUE	Specifies whether to generate Malicious Traffic RDRs.
Number of seconds between Malicious Traffic RDRs on the same attack	60	A Malicious Traffic RDR is generated when an attack is detected. Malicious Traffic RDRs are then generated periodically, at user-configured intervals, for the duration of the attack.
TCP port that should remain open for Subscriber Notification	80	You can choose to block flows that are part of any detected network attack, but this may hinder subscriber notification of the attack. The specified TCP port will not be blocked to allow notification of the attack to be sent to the subscriber.
Policy Check		
Ongoing policy check mode enabled	TRUE	Specifies whether policy changes affect flows that are already open.
Time to bypass between policy checks	30	Maximum time (in seconds) that may pass before policy changes affect flows that are already open.
Quota Management		
Grace period before first breach	2	The time (in seconds) to wait after a quota limit is breached before the breach action is performed. Policy servers should use this period to provision quota to a subscriber that just logged in.
Length of the time frame for quota replenish scatter (minutes)	0	The size of the window across which to randomly scatter the periodic quota replenishment.
Time to bypass between policy checks for quota limited flows	30	Maximum time (in seconds) that may pass before a quota breach affects flows that are already open.

Table 10-2   Advanced Service Configuration Properties (continued)

| Property	Default Value	Description
Volume to bypass between policy checks for quota limited flows	0	Maximum flow volume (in bytes) that may pass before a quota breach affects flows that are already open. A value of zero means that unlimited volume may pass.
Reporting		
Media Flow RDRs enabled	TRUE	Specifies whether to generate Media Flow RDRs.
Subscriber Accounting RDR enabled	FALSE	Specifies whether to generate Subscriber Accounting RDRs. The Subscriber Accounting RDR is used for SM-ISG integration. For more information, see the ISG documentation in the "Managing the SCMP" chapter of the *Cisco Service Control Engine (SCE) Software Configuration Guide.*

How to Edit Advanced Service Configuration Options

Step 1   From the Console main menu, choose **Configuration > System Settings.**

The System Settings dialog box appears.

Step 2   Click the **Advanced Options** tab.

The Advanced Options tab opens.

Step 3   Click **Advanced Service Configuration Options.**

The Advanced Service Configuration Options dialog box opens.
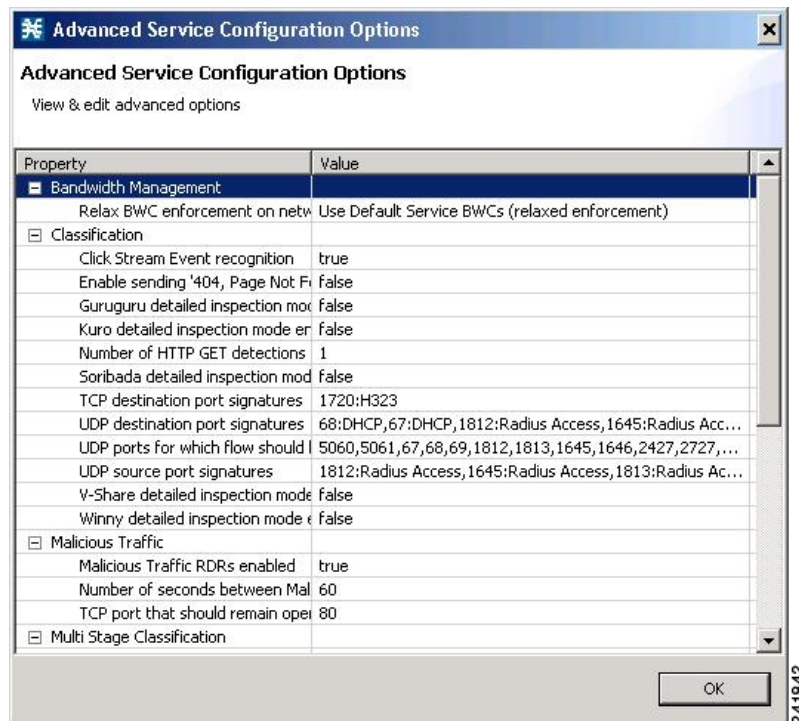
Step 4   Make your changes to the configuration options.

Step 5   Click **OK.**

The Advanced Service Configuration Options dialog box closes.

The changes to the advanced options are saved.

Step 6   Click **OK.**

The System Settings dialog box closes.

Managing VAS Traffic-Forwarding Settings

Traffic forwarding to Value Added Services (VAS) servers allows you to use an external expert system (VAS server) for additional traffic processing, such as intrusion detection and content filtering to subscribers. After processing, flows are sent back to the SCE platform, which then sends them to their original destinations.

The flows to be forwarded are selected based on the subscriber package and the flow type (IP protocol type and destination port number).

VAS traffic forwarding has the following limitations:

- Only the SCE 2000 4xGBE platform supports VAS traffic forwarding.
- A single SCE platform can support up to eight VAS servers.
- A service configuration can contain up to 64 traffic-forwarding tables.
- A traffic-forwarding table can contain up to 64 table parameters.
- VAS traffic forwarding is not supported when unidirectional classification is enabled.

Note   Because of the complexity of the VAS traffic-forwarding feature, VAS flows are not subject to global bandwidth control.

To use VAS traffic forwarding, you must also configure VAS services on the SCE platform. Additional information is available in the "Value Added Services (VAS) Traffic Forwarding" chapter of the *Cisco Service Control Engine (SCE) Software Configuration Guide.*

- How to Enable VAS Traffic Forwarding, page 10-46
- How to Disable VAS Traffic Forwarding, page 10-48
- How to Rename VAS Server Groups, page 10-48
- How to View VAS Traffic-Forwarding Tables, page 10-49
- How to Delete VAS Traffic-Forwarding Tables, page 10-50
- How to Add VAS Traffic-Forwarding Tables, page 10-51
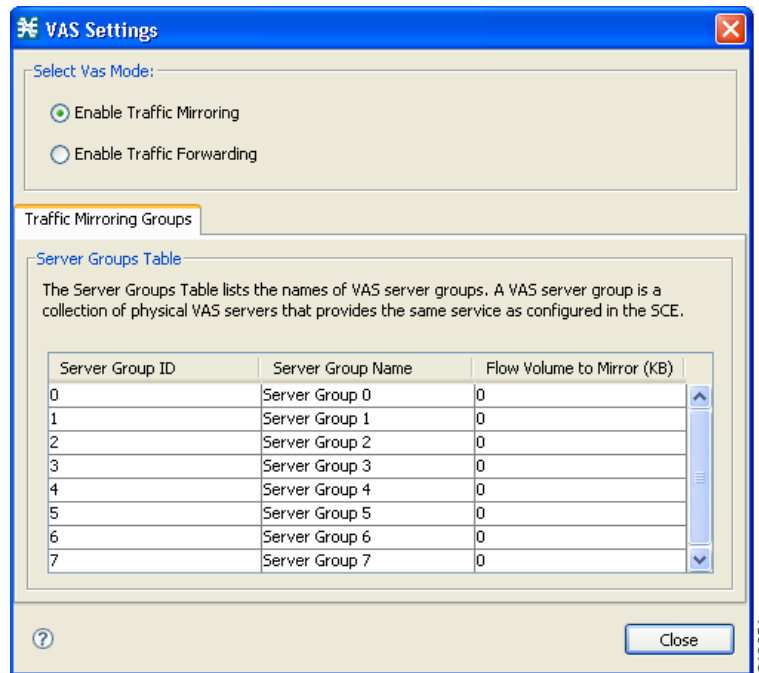- Managing VAS Table Parameters, page 10-51

How to Enable VAS Traffic Forwarding

By default, VAS traffic forwarding is disabled. You can enable it at any time.

Note   VAS traffic forwarding is not supported when unidirectional classification is enabled.

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   Check the **Enable Traffic Forwarding to VAS Servers** check box.

Note   VAS traffic forwarding is not supported in asymmetric routing classification mode. If you try to check the Enable Traffic Forwarding to VAS Servers check box when asymmetric routing classification mode is enabled, a VAS Error message appears.

Click **OK,** and continue at Step 3.

The VAS Traffic Forwarding Table drop-down list in the Advanced tab of the Package Settings dialog box is enabled (see How to Set Advanced Package Options, page 9-7).

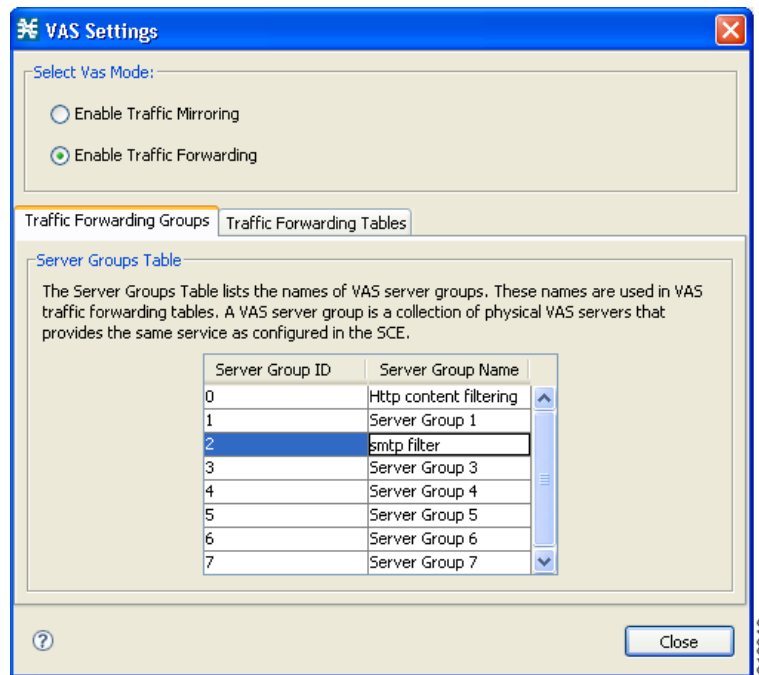**Step 3**   Click **Close.**

The VAS Settings dialog box closes.

How to Disable VAS Traffic Forwarding

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   Uncheck the **Enable Traffic Forwarding to VAS Servers** check box.

VAS traffic forwarding is disabled.

Step 3   Click **Close.**

The VAS Settings dialog box closes.

How to Rename VAS Server Groups

An SCE platform can forward flows to up to eight different VAS server groups. By default, the eight server groups are named "Server Group n", where n takes a value from 0 to 7. Give the server groups meaningful names; the names you give will appear in the drop-down list in the Advanced tab of the Package Settings dialog box (see How to Set Advanced Package Options, page 9-7) and in the Server Group field of the table parameters added to each traffic-forwarding table (see Managing VAS Table Parameters, page 10-51).

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   In the table in the Server Groups Table area, double-click in a cell containing a server group name.

Step 3   Enter a meaningful name in the cell.

Step 4   Repeat Steps 2 and 3 for other server groups you wish to rename.

Step 5   Click **Close.**

The VAS Settings dialog box closes.

How to View VAS Traffic-Forwarding Tables

SCA BB decides whether a flow passing through an SCE platform should be forwarded to a VAS server group based on a traffic-forwarding table. Each entry (table parameter) in a traffic-forwarding table defines to which VAS server group the specified flows should be forwarded.

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**
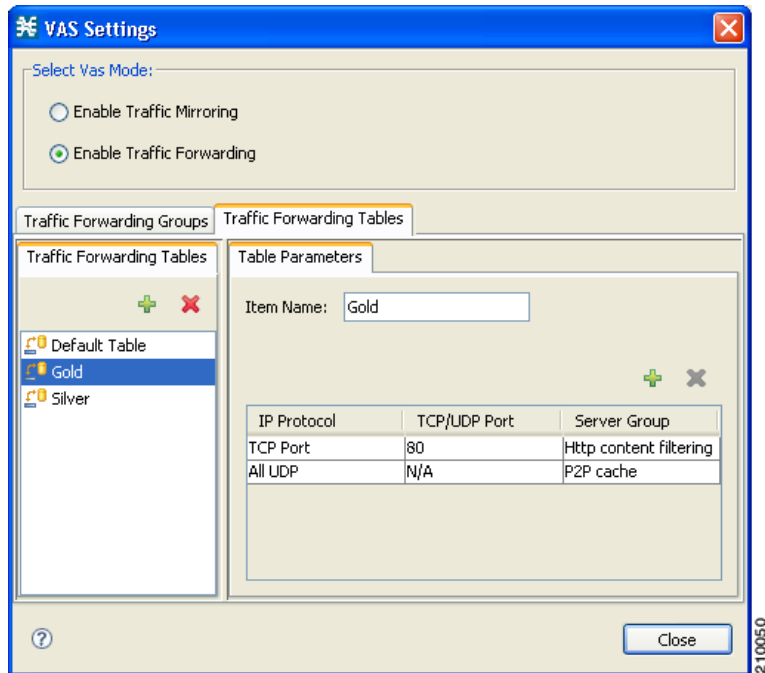
The VAS Settings dialog box appears.

Step 2   Click the **Traffic Forwarding Tables** tab.

The Traffic Forwarding Tables tab opens.

A list of all traffic-forwarding tables is displayed in the Traffic Forwarding Tables area.

Step 3   Click a table in the list of traffic-forwarding tables to display its table parameters.

A list of all table parameters defined for this traffic-forwarding table opens in the Table Parameters tab.

Step 4   Click **Close.**

The VAS Settings dialog box closes.

How to Delete VAS Traffic-Forwarding Tables

You can delete all user-created traffic-forwarding tables. The default traffic-forwarding table cannot be deleted.

Note   A traffic-forwarding table cannot be deleted while it is associated with a package.

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**
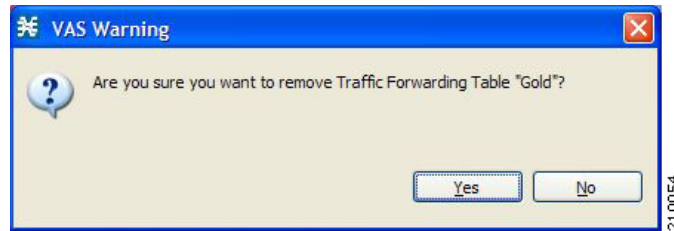
The VAS Settings dialog box appears.

Step 2   Click the **Traffic Forwarding Tables** tab.

The Traffic Forwarding Tables tab opens.

Step 3   From the list of traffic-forwarding tables in the Traffic Forwarding Tables area, select a table.

Step 4   Click (**Delete**).

A VAS Warning message appears.

Step 5   Click **Yes.**

The selected table is deleted and is no longer displayed in the list of traffic-forwarding tables.

Step 6   Click **Close.**

The VAS Settings dialog box closes.

How to Add VAS Traffic-Forwarding Tables

A default traffic-forwarding table is included in the service configuration. You can add up to 63 more traffic-forwarding tables, and then assign different traffic-forwarding tables to different packages.

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   Click the **Traffic Forwarding Tables** tab.

The Traffic Forwarding Tables tab opens.

Step 3   In the Traffic Forwarding Tables area, click (**Add**).

A new table named Table (n), where n is a value between 1 and 63, is added to the list of traffic-forwarding tables in the Traffic Forwarding Tables area.

The table name is also displayed in the Item Name box in the Table Parameters tab.

Step 4   In the Item Name field, enter a unique and relevant name for the traffic-forwarding table.

You can now add table parameters to the new traffic-forwarding table, see How to Add VAS Table Parameters, page 10-52.

Managing VAS Table Parameters

A table parameter is an IP protocol type, an associated TCP/UDP port (where applicable), and a VAS server group or a range of IP addresses.

A traffic-forwarding table is a collection of related table parameters.

A traffic-forwarding table can contain up to 64 table parameters.

- How to Add VAS Table Parameters, page 10-52
- How to Edit VAS Table Parameters, page 10-52
- How to Delete VAS Table Parameters, page 10-54

How to Add VAS Table Parameters

You can add up to 64 table parameters to a traffic-forwarding table.

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   Click the **Traffic Forwarding Tables** tab.

The Traffic Forwarding Tables tab opens.

Step 3   From the list of traffic-forwarding tables in the Traffic Forwarding Tables area, select a table.

Step 4   In the Traffic Parameters tab, click (**Add**).

A new table parameter is added to the list of table parameters in the Table Parameters tab.

✎ **Note**   Each new table parameter has the following default values:

Parameter	Default value
IP Protocol	TCP Port
TCP/UDP Port	80
Server Group	Server Group 0

You can now edit the new table parameter, as described in the following section.

Step 5   Click **Close.**

The VAS Settings dialog box closes.

How to Edit VAS Table Parameters

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   Click the **Traffic Forwarding Tables** tab.

The Traffic Forwarding Tables tab opens.

Step 3   From the list of traffic-forwarding tables in the Traffic Forwarding Tables area, select a table.

Step 4   In the table in the Table Parameters tab select a protocol, port, and server group.

    a.   Click in a cell in the IP Protocol column, and, from the drop-down list that opens, select an IP protocol type.
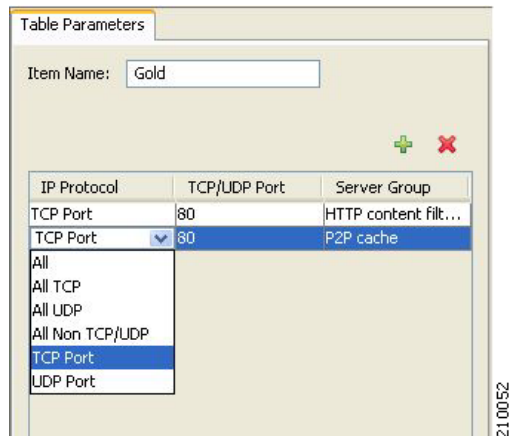
If you select All, All TCP, All UDP, or All Non TCP/UDP, "N/A" will appear in the TCP/UDP Port cell when you move to another cell in the table.
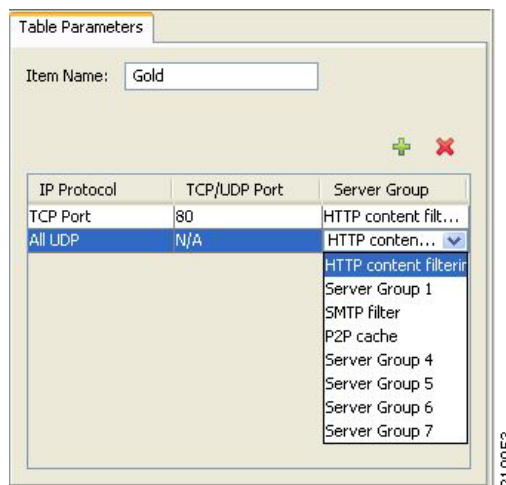
b. If you selected TCP Port or UDP Port, double-click in the cell in the TCP/UDP Port column, and enter the port number.

Note You cannot enter a range of ports in the TCP/UDP Port cell; you must add a separate table parameter for each port.

c. Click in the cell in the Server Group column, and, from the drop-down list that opens, select a server group.

Step 5 Click **Close.**

The VAS Settings dialog box closes.

How to Delete VAS Table Parameters

Step 1   From the Console main menu, choose **Configuration > VAS Settings.**

The VAS Settings dialog box appears.

Step 2   Click the **Traffic Forwarding Tables** tab.

The Traffic Forwarding Tables tab opens.

Step 3   From the list of traffic-forwarding tables in the Traffic Forwarding Tables area, select a table.

Step 4   From the list of table parameters in the Table Parameters tab, select a table parameter.

Step 5   Click ✖ (**Delete)**.

The selected table parameter is deleted and is no longer displayed in the list of table parameters.

Step 6   Click **Close.**

The VAS Settings dialog box closes.