Service Security Using the Cisco SCE Platform Application Note

Release 3.1
May 2007

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Customer Order Number:
Text Part Number: OL-10610-02

Service Security Using the Cisco SCE Platform Application Note

CONTENTS

CHAPTER 1

Introduction and Scope

This module provides an introduction to Service Security using the Cisco SCE.

Introduction and Scope

The need for protection from various types of attacks and malicious traffic that originate from the Internet has gained focus in recent years. DoS and DDoS attacks, worms, viruses, malicious HTTP content, and various types of intrusions are becoming increasingly popular.

SCE platforms are deployed inline and are stateful and programmable. These features position the SCE platform to detect and mitigate the effect of malicious traffic on service providers and their customers.

SCA BB includes Service Security functionality comprising anomaly detection, spam/mass-mailing detection, and signature detection. This functionality allows the SCE platform to address many of the threats that exist in current networks.

The SCA BB solution is not an airtight solution against network threats. It is effective in providing an insight into malicious activity in an operator network, and in mitigating large scale eruptions of malicious activity that might compromise overall network performance and degrade user experience.

This guide contains practical sections describing the specifics of configuration and reporting. The purpose is to explain and demonstrate the Service Security concepts. For a full operational description of the relevant management modules, refer to the SCA BB user guides.

CHAPTER 2

Functionality Overview

This module provides an overview of the service security functionality of the SCE platform.

Functionality Overview

The Cisco SCE platform uses three approaches for threat detection:

- Anomaly Detection—This set of mechanisms monitors the rate of connections (successful and unsuccessful) to and from each host IP address. It detects malicious activity based on exceeding "normal" connection rates and on the ratio between successful and unsuccessful connections. Anomaly detection characteristics can indicate the following categories of malicious activity:
  - Scan/Sweep/Attack—Based on an indication that a host is generating an anomalous rate of connections.
  - DoS/DDoS—Based on an indication that a host is a target for an anomalous rate of connections.
  - DoS—Based on an indication that a pair of hosts are involved in an activity where one is generating, and the other one is a target, for an anomalous rate of connections.
- The anomaly detection mechanism is effective in addressing zero-day threats—addressing threats as they appear without the need for preliminary knowledge about their exact nature and L7 signatures, but rather based on the characteristics of their network activity.

  For further details, see "Anomaly Based Detection".
- Mass-Mailing activity detection—This mechanism is based on monitoring SMTP session rates for individual subscribers. It uses the SCE platform's subscriber-awareness and can work in subscriber-aware or anonymous subscribers mode. SMTP is a protocol used for sending email; an excess rate of such sessions originating from an individual subscriber is usually indicative of malicious activity involving sending email: either mail-based viruses or spam-zombie activity.
- Signature based detection—The SCE platform provides stateful L7 capabilities that can be used to detect malicious activity that is not easily detectable by the other mechanisms. A user can independently configure signatures for such threats, thus achieving a fast turnaround time in addressing threats (details on this are not covered in this document).

All three detection approaches provide operators with several possible courses of action to be implemented based on their business needs.

Monitor—Inspect the network for malicious activity detected by each of these methods. This can be done using reports that are based on information collected for malicious activity that was detected, or using SNMP traps that can detect malicious activity using the anomaly detection module.

Block—Automatically block malicious activity that has been detected by the SCE platform to avoid threat propagation and adverse effects to the network.

Notify—Notify subscribers that they have been detected as being involved in malicious activity by redirecting their web sessions to a captive portal.

Operators have a high level of flexibility in tuning the detection methods and actions to be taken based on their specific needs. The SCA BB Security Dashboard is a GUI application that provides a simple front end for configuring and monitoring security functionality.

Figure 2-1 SCA BB Security Dashboard

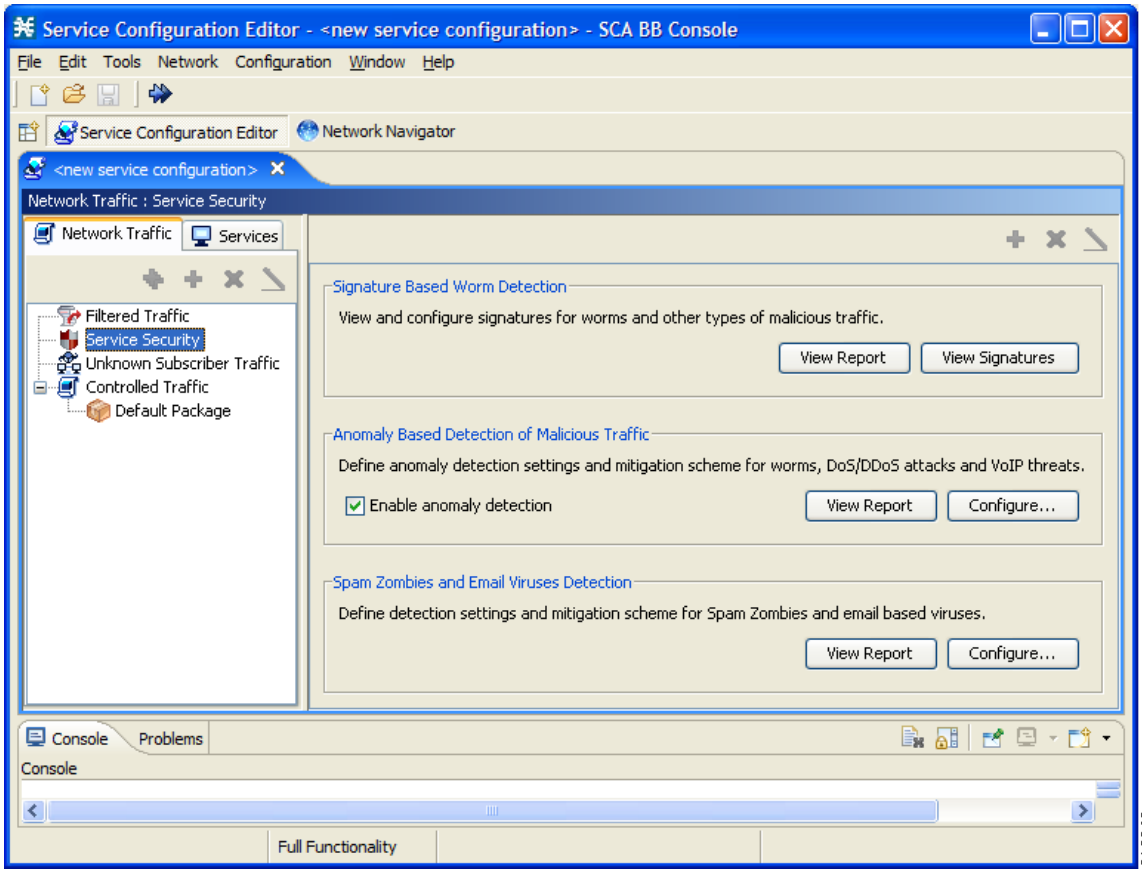

CHAPTER 3

Anomaly Based Detection

This module describes anomaly based detection using the Cisco SCE platform.

Overview

The most comprehensive threat detection module is the anomaly detection module. It monitors successful (correctly established for TCP, bi-directional otherwise) and unsuccessful (uni-directional, also termed "suspected") connection rates *to* and *from* any IP address viewed by the system. It triggers the detection of an anomaly based of one of the following criteria:

- Total connection rate exceeded a predefined threshold.

  OR

- Suspected connection rate exceeded a predefined threshold, and suspected-to-unsuspected connections ratio exceeded a predefined threshold.

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

Anomaly detection is split into categories based on the directional nature of anomalies as described in the following subsections. The concepts used by the three methods are identical, but they differ in the role of the endpoint that is monitored for the anomaly.

Scan/Sweep/Attack is a category of malicious activity that is based on detecting an anomaly in the connection rate from an IP address (the module ignores the destination IP addresses involved). The anomaly is detected based on the criteria specified above and can indicate one of the following:

- An attack—A host is participating in an attack on another host.
- A sweep—A host is sweeping the network in search of a vulnerable host (this is typical activity for network worms).
- A scan—A host is scanning the ports of other hosts to find out what services they are using and which ports are potentially vulnerable.

A denial of service (DoS) attack is detected based on an anomaly in the connection rate between a pair of hosts—one host is attacking the other. This can be either an isolated attack, or part of a larger scale DDoS attack.

A distributed denial of service (DDoS) attack is detected based on an anomaly in the connection rate to an IP address (the module ignores the source IP addresses involved), which indicates that it is being attacked. The attack can be from either a single IP address (DoS) or multiple IP addresses (DDoS).

For many types of anomalies, flexibility is reflected in the ability to define detection thresholds and the action to be taken for each:

- Anomaly direction (subscriber/network)
- Protocol (TCP/UDP/ICMP/Other)
- Port uniqueness for TCP/UDP—Whether the anomaly threshold applicable to a single port or the aggregate of ports.

The malicious traffic reports family in the reporter contains reports that allow monitoring different aspects of malicious activity.

Configuring Anomaly Detection

The configuration screen for anomaly detection shows the "tree" of various detectors. The tree comprises the three categories of anomalies.

The three categories are separate from a configuration perspective: Scan/Sweep, DoS, and DDoS.

Figure 3-1 Configuring Anomaly Detection

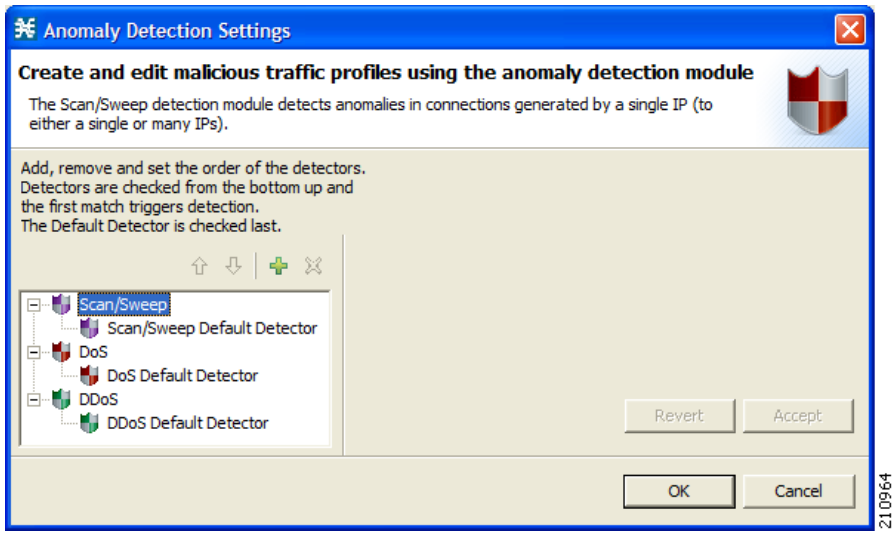

General Concepts in Anomaly Detection Configuration

The Detector Structure

The terms used in anomaly detection are described below. The term “related to” is used extensively; however, its meaning depends on the semantics of each anomaly and will be mentioned in the description given for each category of anomalies (scan/sweep/attack, DoS, and DDoS).

The first level of hierarchy within the detector contains two subcategories: Subscriber and Network. The Subscriber subcategory refers to malicious traffic related to IP addresses that are on the subscriber port. The Network subcategory refers to malicious traffic related to IP addresses that are on the network port.

These subcategories allow an operator to act on anomalies that are related to either the subscriber or network side.

Under each subcategory, six further granular subcategories exist:

- TCP—Refers to anomalies in connection rate related to a host over any TCP destination port. The aggregate rate of connections on all TCP ports is applicable for matching this rule.
- TCP Specific—Refers to anomalies in connection rate related to a host over any specific TCP destination port. For example, the rate of connections related to a specific host on port 80 is applicable for matching this rule, as well as the rate of connections related to a host over each of the TCP ports 23, 25, 110 (and any specific TCP port).
- UDP—Refers to anomalies in connection rate related to a host over any UDP destination port. The aggregate rate of connections on all TCP ports is applicable for matching this rule.
- UDP Specific—Refers to anomalies in connection rate related to a host over any specific UDP destination port. For example, the rate of connections related to a host on port 53 is applicable for matching this rule, as well as the rate of connections generated by a host over each UDP port.
- ICMP—Refers to anomalies in the connection rate related to a host using ICMP.
- Other—Refers to anomalies in the connection rate related to a host using non TCP/UDP/ICMP protocols.

Figure 3-2 Detector Structure

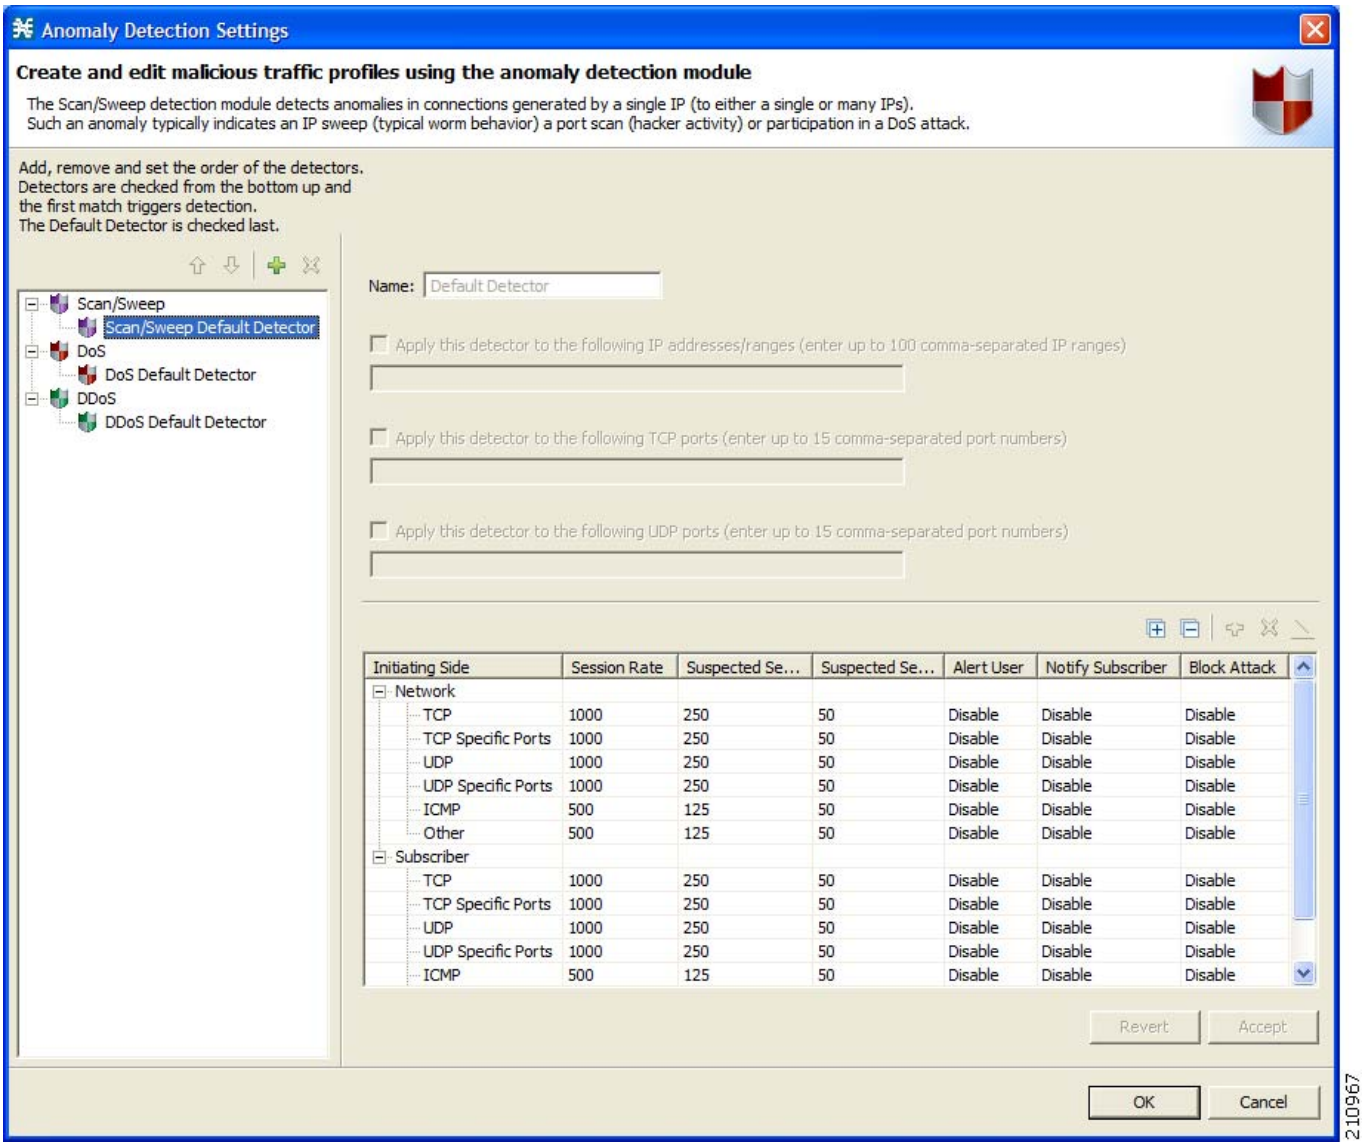

Configuring Detection Thresholds

The configuration of specific detection parameters and related actions is carried out on a per subcategory basis.

The configurable detection parameters for each anomaly subcategory are:

- Session Rate—Refers to the threshold of session-rate (sessions/second for an IP address) that for itself would trigger the anomaly of this type.

  For example, configuring the value 1000 on the session-rate for the TCP anomaly means that if a rate of 1000 TCP sessions/second (on any port) was detected from a host, the anomaly would be triggered.

- Suspected Session Rate—Refers to the threshold of suspected-sessions-rate (suspected-sessions/second) that for itself would trigger the anomaly line item.

  Suspected sessions are TCP sessions that were not properly established or unidirectional sessions for other protocols (UDP/ICMP/Other).

  This parameter is used in conjunction with the suspected-sessions ratio as an indicator of a traffic anomaly—a relatively high session-rate for which a large number do not respond typically indicates malicious activity.

  For example, configuring the value 1000 suspected-sessions/second for the TCP anomaly means that if a rate of 1000 TCP suspected-sessions/second (on any port) was detected from a host, *and* the suspected connection-ratio crossed the predefined threshold, the anomaly would be triggered.

- Suspected Sessions Ratio—The ratio between the suspected-sessions rate and the total session-rate. A high ratio indicates a high amount of "unresponded" sessions, which may indicate malicious activity.

Figure 3-3 Configuring Detection Thresholds

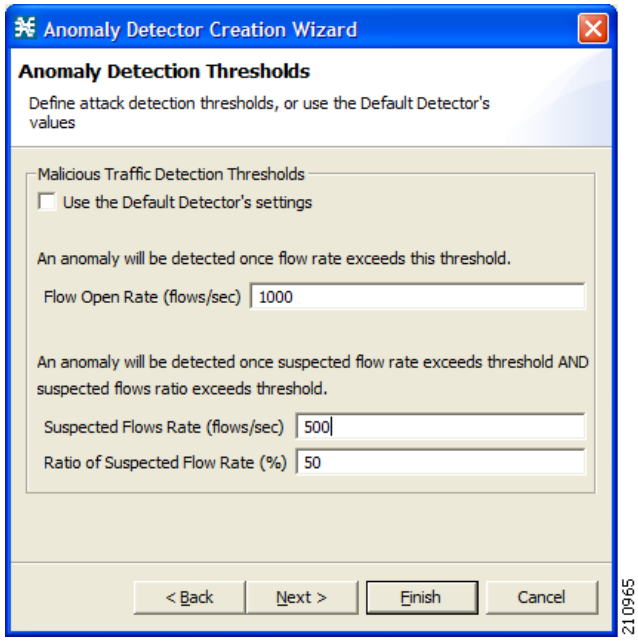

Configuring Actions

Each anomaly subcategory also includes an option to define the action to be taken upon detection.

It is not possible to configure, per anomaly subcategory, logging to an on-device log file and to the Collection Manager Database through RDRs (SCE Raw Data Records). To enable and disable RDRs for malicious traffic, from the Configuration menu, select RDR settings.

There are three independent and not mutually exclusive actions that are configurable per item:

- Alert User—Generate an SNMP trap (see Pcube proprietary MIB for details) indicating the beginning and end of such an attack.

- Notify Subscriber—Notify a subscriber of the malicious activity through redirecting their browsing sessions to a captive portal. To configure a subscriber notification, from the Configuration menu select Subscriber Notifications, then select Network Attack Notification. See the *Cisco Service Control Application for Broadband User Guide* for subscriber notification options.
- Block—Block the relevant session. Blocking is performed based on the specification of the malicious traffic that triggered the anomaly. For instance, if the anomaly detected is a port agnostic TCP scan from the subscriber side, all TCP sessions originating from the subscriber side will be blocked. Blocking will persist until the anomaly can no longer be seen (blocking will be removed intermittently to check for this). Note that if subscriber notification is also enabled for the anomaly, blocking is not applied to the port relevant for browsing (by default this is TCP port 80).

Figure 3-4 Anomaly Detection Action Settings

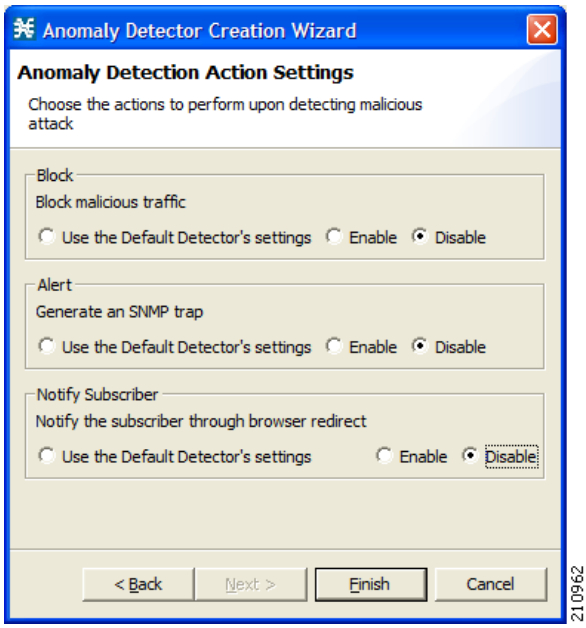

The Default Detector

For each category of malicious traffic a default detector exists that cannot be removed. The default detector is preconfigured with factory defaults for thresholds and actions.

Additional Detectors

For each category of malicious traffic (scan/sweep, DoS, DDoS), additional detectors, up to a total of 100 detectors in the three categories combined, can be defined under the default detector.

Such a detector applies to an IP-Address List, a UDP port list, a TCP port list, or a combination of the three. The detector should contain different thresholds and/or different actions to be applied for anomalies that match this list.

As an example, it is possible to have a DDoS detector, which corresponds to an IP address with no specific port, or to have another detector which monitors specific port attacks (e.g. DNS or SMTP detectors). Another example can be detectors that are targeted at a specific worm, and therefore include a specific list of ports.

When creating an additional detector, it is possible to apply it to one or more of the subcategories of the anomaly. For example, a user can create a DDoS detector for a specific IP list, and choose to apply it only for specific port attacks from the subscriber side. The new detector need not cover all "potential" subcategories of DDoS.

When creating a new anomaly subcategory, a user is required to define the actions to take, and can choose to use custom detection thresholds, or to inherit detection thresholds from the default detector.

Figure 3-5 Anomaly Detector Wizard

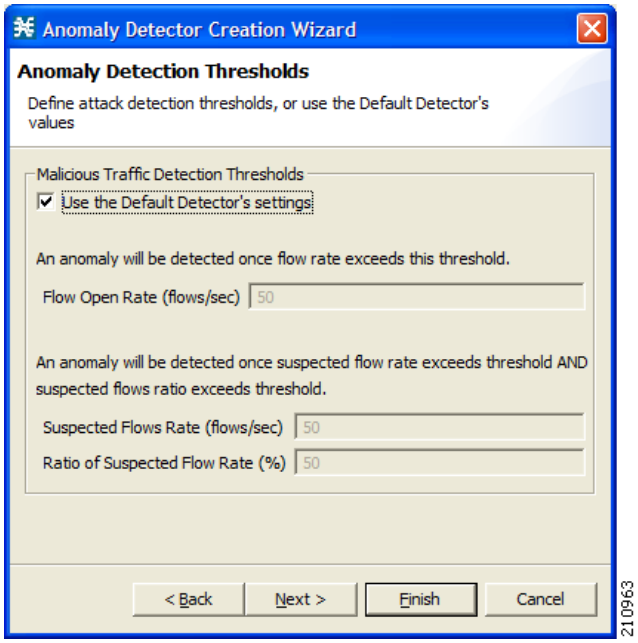

An anomaly can be triggered by *exactly* one detector. The order of detection is from the bottom of the list and up—the first detector that “matches” the anomaly threshold, together with the IP/Port list specification, sets the decision on the actions to take.

A user should use the up and down arrows above the anomaly tree window to arrange the detector in each category into the desired order.

Figure 3-6 Anomaly Tree Window

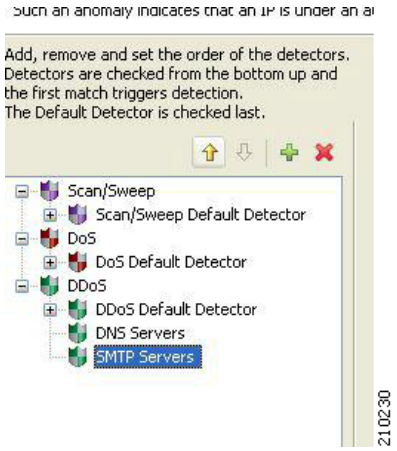

The Scan/Sweep Detector

The scan/sweep detector detects anomalies in connections generated by a host. These anomalies include total-connection-rate, suspected-connections-rate and suspected-connections-ratio, as described in the overview section.

A high ratio of suspected connections generated by a host, beyond some minimal rate, typically indicates the presence of a worm, which tries to propagate by sweeping the network for vulnerable hosts. The presence of a worm is even more obvious when an anomaly is detected on a specific port over which the vulnerability exists—this would typically trigger the TCP/UDP specific anomalies.

In addition, port scans and attacks performed by a host are also included under this category. Port scans are typically characterized by a high ratio of suspected connections, while host attacks can be either normal or suspected connections. Both types of malicious traffic could be detected by the non-port-specific anomaly detectors.

The first level of hierarchy within the detector contains two subcategories: Subscriber and Network. The Subscriber subcategory refers to scans, sweeps, and attacks that were detected from IP addresses that are on the subscriber port side. The Network subcategory refers to scans, sweeps, and attacks that were detected from IP addresses that are on the network port side.

These subcategories allow an operator to act on Scan/Sweep/Attack anomalies that originate from either the subscriber or network side.

Under each subcategory, six further granular subcategories exist:

- TCP—Refers to anomalies in the aggregate connection rate generated by a host over all TCP destination ports. For example, the aggregate rate of connections generated by a host to ports 23, 25, 80, 110, and any other TCP port is applicable for matching this rule.
- TCP Specific—Refers to anomalies in the connection rate generated by a host to any specific TCP destination port. For example, the rate of connections generated by a host to port 80 is applicable for matching this rule, as well as the rate of connections generated by a host to any other individual TCP port.
- UDP—Refers to anomalies in the aggregate connection rate generated by a host to all UDP destination ports. For example, the aggregate rate of connections generated by a host to ports 53, 445, and any other UDP port is applicable for matching this rule.

- UDP Specific—Refers to anomalies in the connection rate generated by a host to any specific UDP destination port. For example, the rate of connections generated by a host to port 53 is applicable for matching this rule, as well as the rate of connections generated by a host to any other individual UDP port.
- ICMP—Refers to anomalies in the connection rate generated by a host using ICMP.
- Other—Refers to anomalies in the connection rate generated by a host using non TCP/UDP/ICMP protocols.

Figure 3-7 Scan/Sweep Default Detector

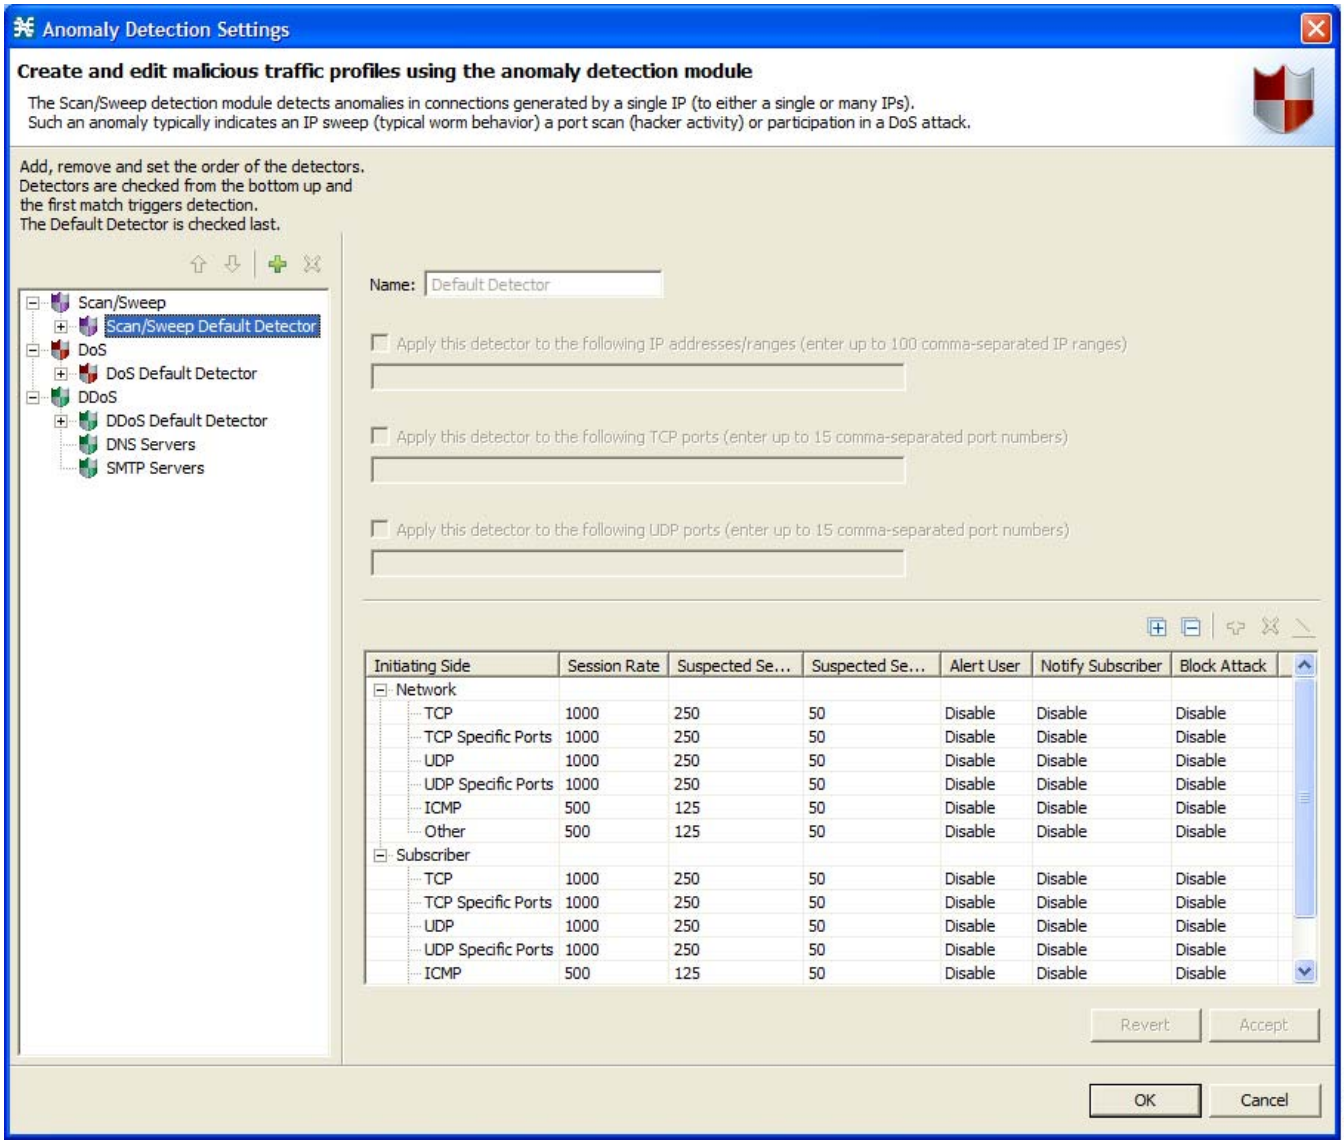

The DDoS Detector

The DDoS detector detects anomalies in connections for which a host is a destination. These anomalies include total-connection-rate, suspected-connections-rate and suspected-connections-ratio, as described in the overview section.

A high rate of connections to an IP address indicates that the IP address is being attacked.

A better indicator that an IP is being attacked is when there is a high suspected-sessions rate, in addition to a high suspected-sessions ratio.

The DDoS detector module detects DoS and DDoS attacks without differentiating between them because the IP addresses involved in generating the connections are not tracked by the module.

The first level of hierarchy within the detector contains two subcategories: Subscriber and Network. The Subscriber subcategory refers to attacks that were detected on IP addresses that are on the subscriber port side. The Network subcategory refers to attacks that were detected on IP addresses that are on the Network port side.

These subcategories allow an operator to act on anomalies that are towards either the subscriber or network side.

An operator might consider protecting subscribers from suspected attacks, while not doing this for network side activity, in order to not affect legitimate activity by subscribers who are falsely detected as participating in an attack.

Under each subcategory, six further granular subcategories exist:

- TCP—Refers to anomalies in the aggregate connection rate to a host over all TCP destination ports. For example, the aggregate rate of connections to ports 23, 25, 80, 110, and any TCP port is applicable for matching this rule.
- TCP Specific—Refers to anomalies in the connection rate to a host over any specific TCP destination port. For example, the rate of connections to a host on port 80 is applicable for matching this rule, as well as the rate of connections to a host on any other individual TCP port.
- UDP—Refers to anomalies in the connection rate to a host over all UDP destination ports. For example, the aggregate rate of connections to ports 53, 445, and any UDP port is applicable for matching this rule.
- UDP Specific—Refers to anomalies in the connection rate to a host over any specific UDP destination port. For example, the rate of connections to a host on port 53 is applicable for matching this rule, as well as the rate of connections to a host over any other individual UDP port.
- ICMP—Refers to anomalies in the connection rate to a host using ICMP.
- Other—Refers to anomalies in the connection rate to a host using non TCP/UDP/ICMP protocols.

Figure 3-8 DDoS Default Detector

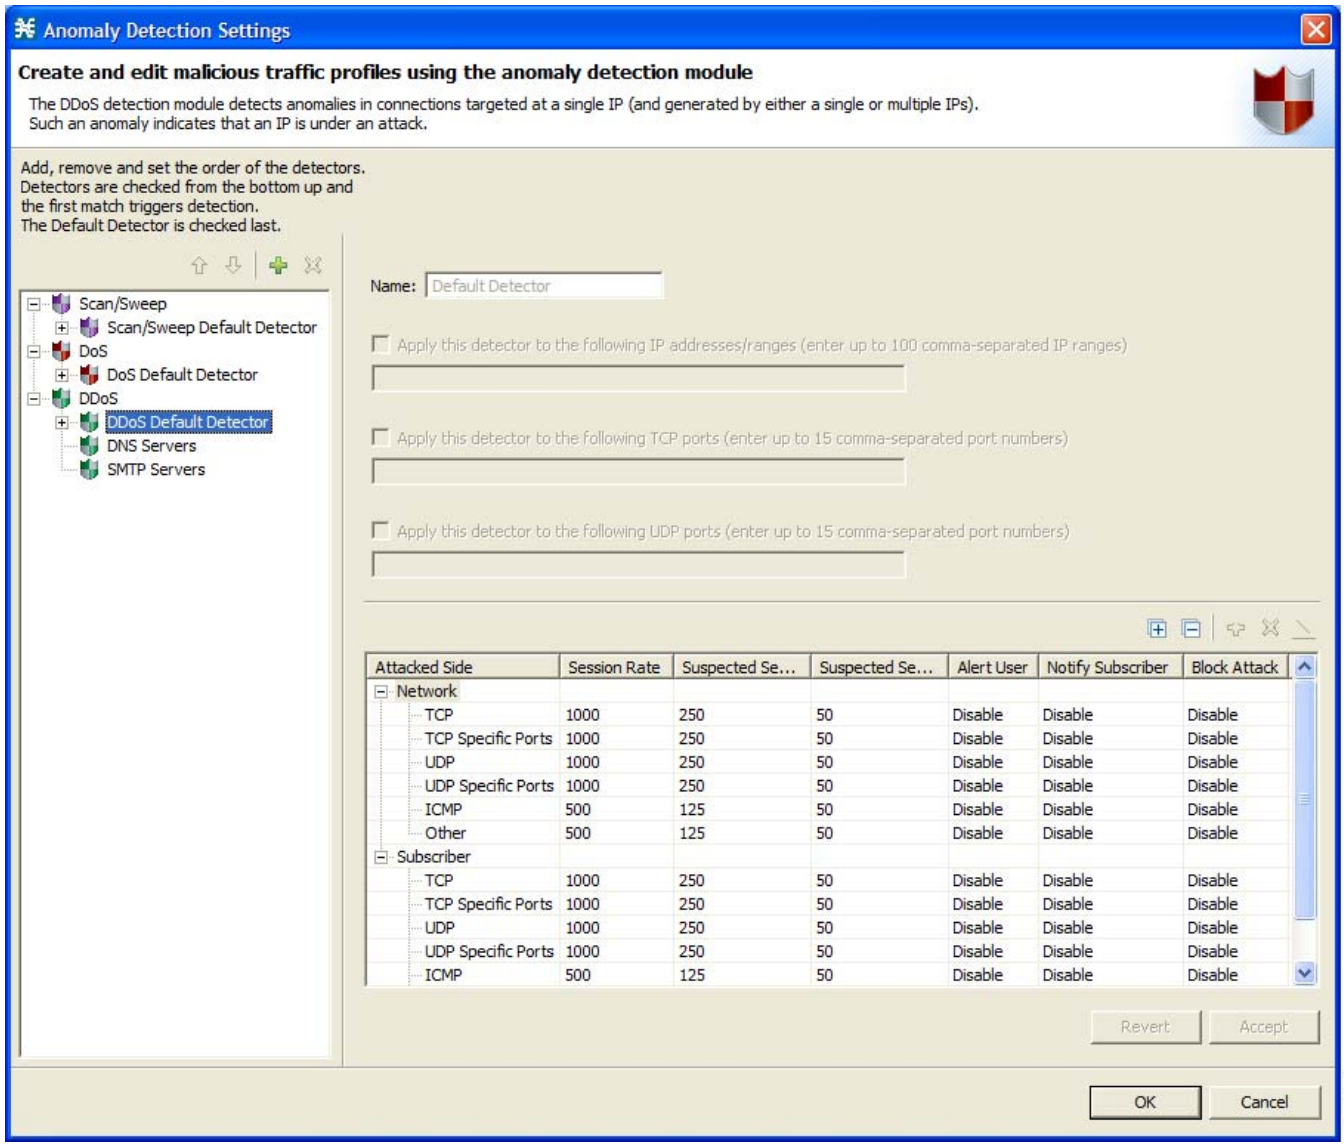

The DoS Detector

The DDoS detector detects anomalies in connections between a pair of hosts. These anomalies include total-connection-rate, suspected-connections-rate and suspected-connections-ratio, as described in the overview section.

A high rate of connections between a pair of hosts may be an indicator that the source host is attacking the destination host.

A better indicator that one host is attacking the other is when there is a high suspected-sessions rate, in conjunction with a high suspected-connections ratio.

This module monitors the rate of connections between a pair of hosts. This can include DoS and DDoS attacks without differentiating between them because there can be multiple hosts attacking the destination host, which will not be explicitly detected by this module.

The first level of hierarchy within the detector contains two subcategories: Subscriber and Network. The Subscriber subcategory refers to attacks that were detected from IP addresses that are on the subscriber port side. The Network subcategory refers to attacks that were detected from IP addresses that are on the Network port side.

These subcategories allow an operator to act on anomalies that are coming from either the subscriber or network side.

An operator might consider protecting subscribers from suspected attacks directed at them, but not doing this for network side activity, in order to not affect legitimate activity by subscribers who are falsely detected as participating in an attack.

Under each subcategory, four further granular subcategories exist:

- TCP—Refers to anomalies in the aggregate connection rate between a pair of hosts over all TCP destination ports. The aggregate rate of connections between a pair of hosts on ports 23, 25, 80, 110, and all other TCP ports is applicable for matching this rule.
- TCP Specific—Refers to anomalies in the connection rate between a pair of hosts over any specific TCP destination port. For example, the rate of connections between a pair of hosts on port 80 is applicable for matching this rule, as well as the rate of connections between a pair of hosts on any other individual TCP port.
- UDP—Refers to anomalies in the aggregate connection rate between a pair of hosts over all UDP destination ports. For example, the aggregate rate of connections between a pair of hosts on ports 53, 445, and all other UDP ports is applicable for matching this rule.
- UDP Specific—Refers to anomalies in the connection rate between a pair of hosts over any specific UDP destination port. For example, the rate of connections between a pair of hosts on port 53 is applicable for matching this rule, as well as the rate of connections between a pair of hosts over any other individual UDP port.

The ICMP category and the Other category do not exist for DoS detection since connections between a pair of hosts using these protocols cannot be distinguished from one another.

Figure 3-9 DoS Default Detector

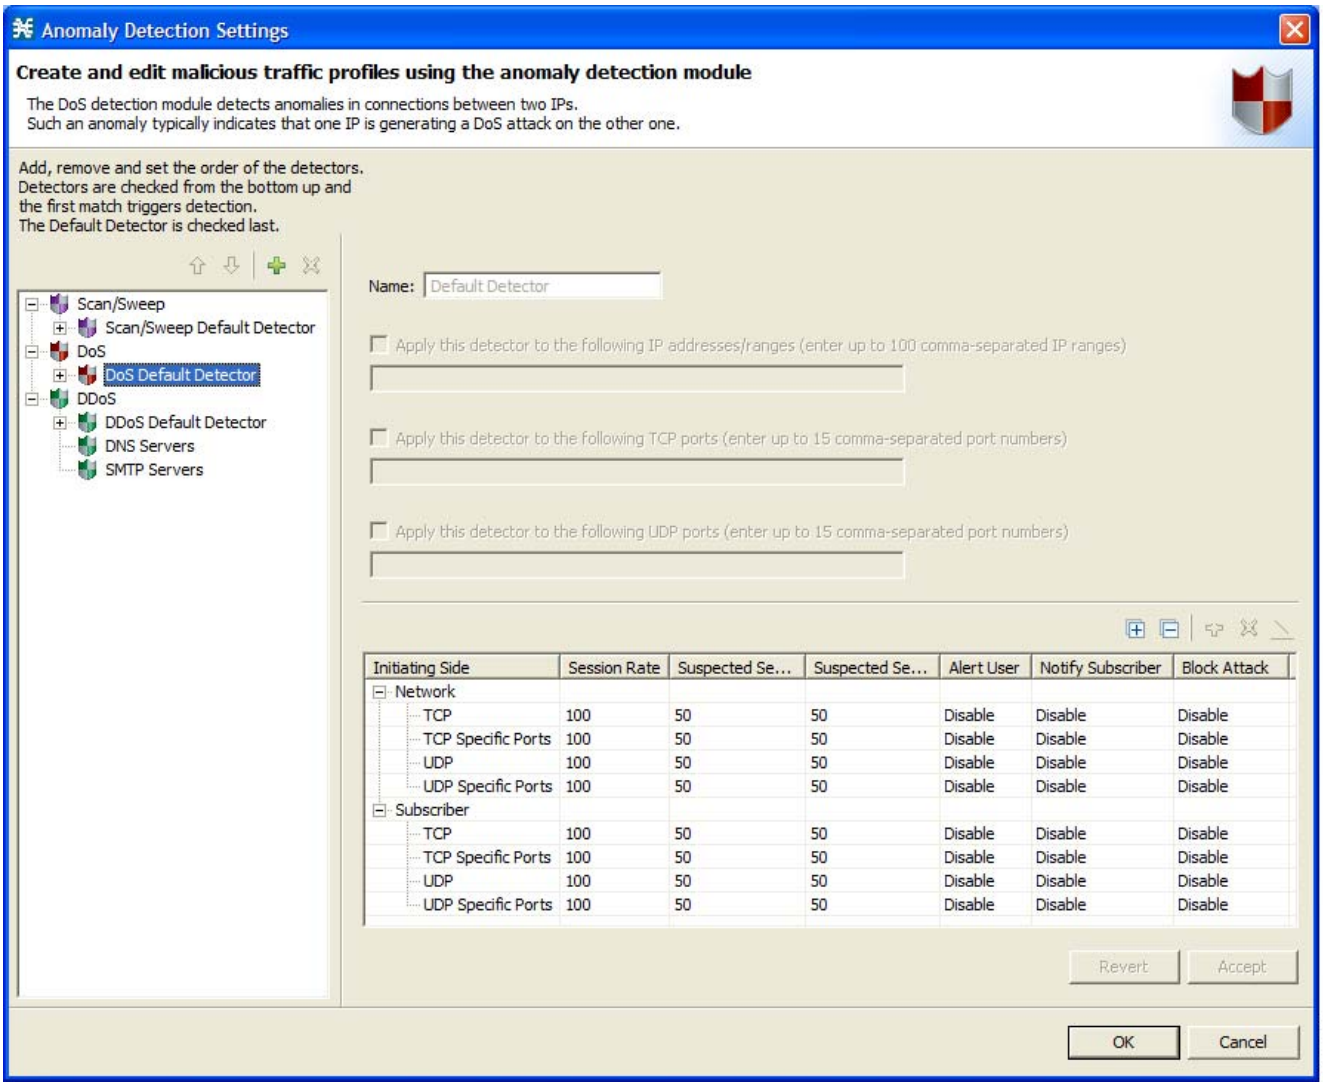

Monitoring Malicious Traffic

Information about traffic anomalies detected using the scan/sweep and the DDoS modules is sent over RDRs (SCE Raw Data Records) and is stored in the Collection Manager Database. This information can be used to find network trends, detect new threats, and track malicious hosts or subscribers.

Since anomaly detection is based on session rate threshold breaches, the actual information stored in the database can vary, depending on the detection thresholds set.

For example, it is possible that scan/sweep reports generated for a system with thresholds set at 100 sessions/second would look notably different (much more loaded with events) than reports generated for a system with thresholds set at 1000 sessions/second (assuming traffic patterns are identical for both systems).

There are a number of reports dealing with malicious traffic.

Global or "trending" reports include:

- Top scanned or attacked ports
- Global scan/attack rate
- Global DoS rate
- Infected subscribers
- DoS attacked subscribers

Individual subscriber or host reports:

- Top Scanning/attacking hosts
- Top Scanning/attacking subscribers
- Top attacked hosts
- Top attacked subscribers

Top Scanned or Attacked Ports

The Top Scanned or Attacked ports report is based on the records of sweep/attack activity detected by the SCE platform on a specific port.

The report presents the top ports over which such activity was detected. This provides a good indication of the current “vulnerable” ports being searched by worms, bots, and hackers.

The introduction of a new network worm is typically characterized by a rise in the amount of sweeps on a specific port. Ongoing monitoring of the network using this report would allow an operator to detect an outbreak of a new network threat based on a rise in the amount of malicious activity over some port.

Figure 3-10 Top Scanned or Attacked Ports Report

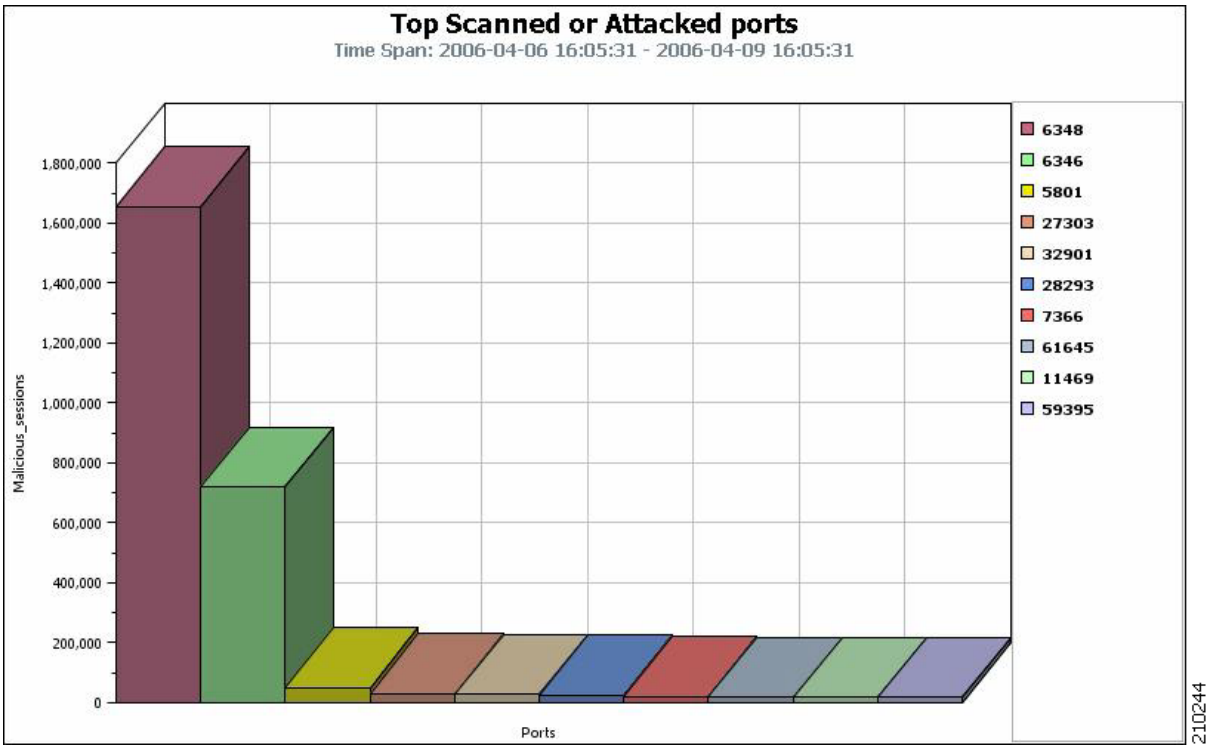

Global Scan or Attack Rate

The Global Scan or Attack Rate report is based on the records of sweep/attack activity detected by the SCE platform, but not limited to a specific port.

The report presents the global scan/attack rate over time, broken down by protocol. The report can be filtered by the scan direction: from the subscriber or from the network.

An outbreak of a network worm is typically reflected in a surge of scan activity, which this report can help detect.

Figure 3-11 Global Scan or Attack Rate Report

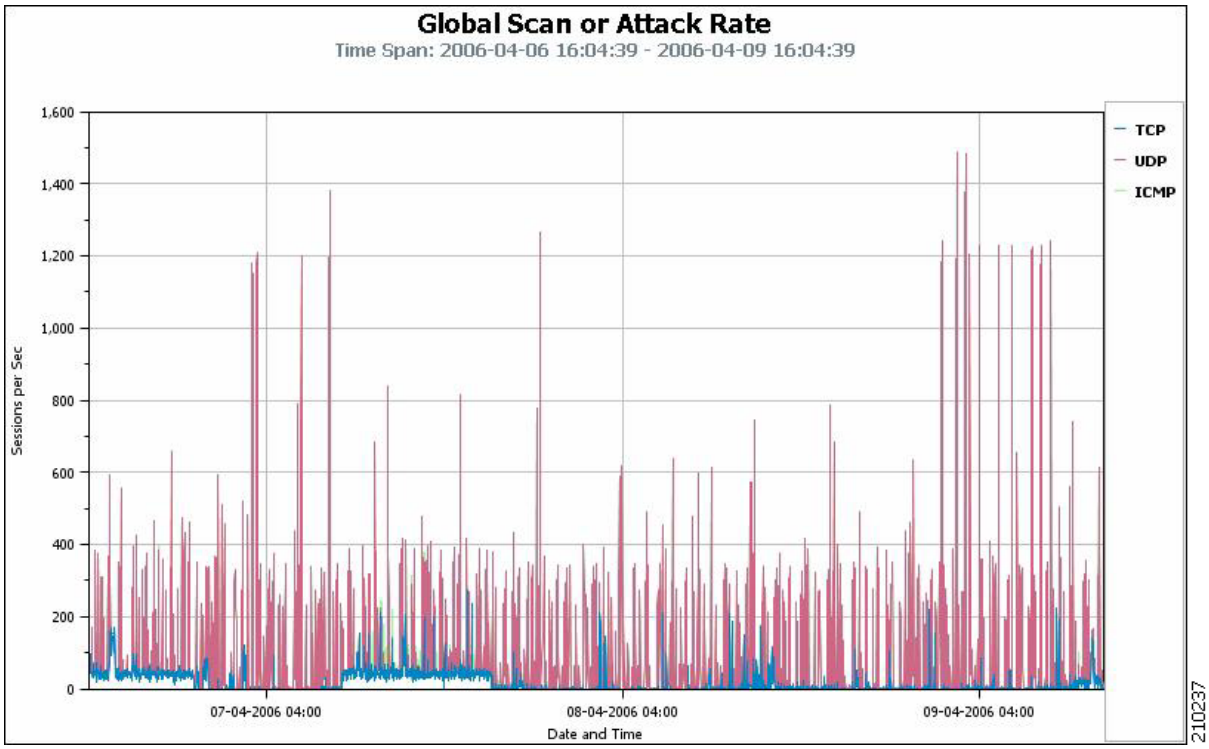

Global DoS Rate

The Global DoS Rate report is based on the records of DDoS activity detected by the SCE platform, but not limited to a specific port.

This report presents the global DoS rate over time broken down by protocol. The report can be filtered by the attack direction: to the subscriber or to the network.

Figure 3-12 Global DoS Rate Report

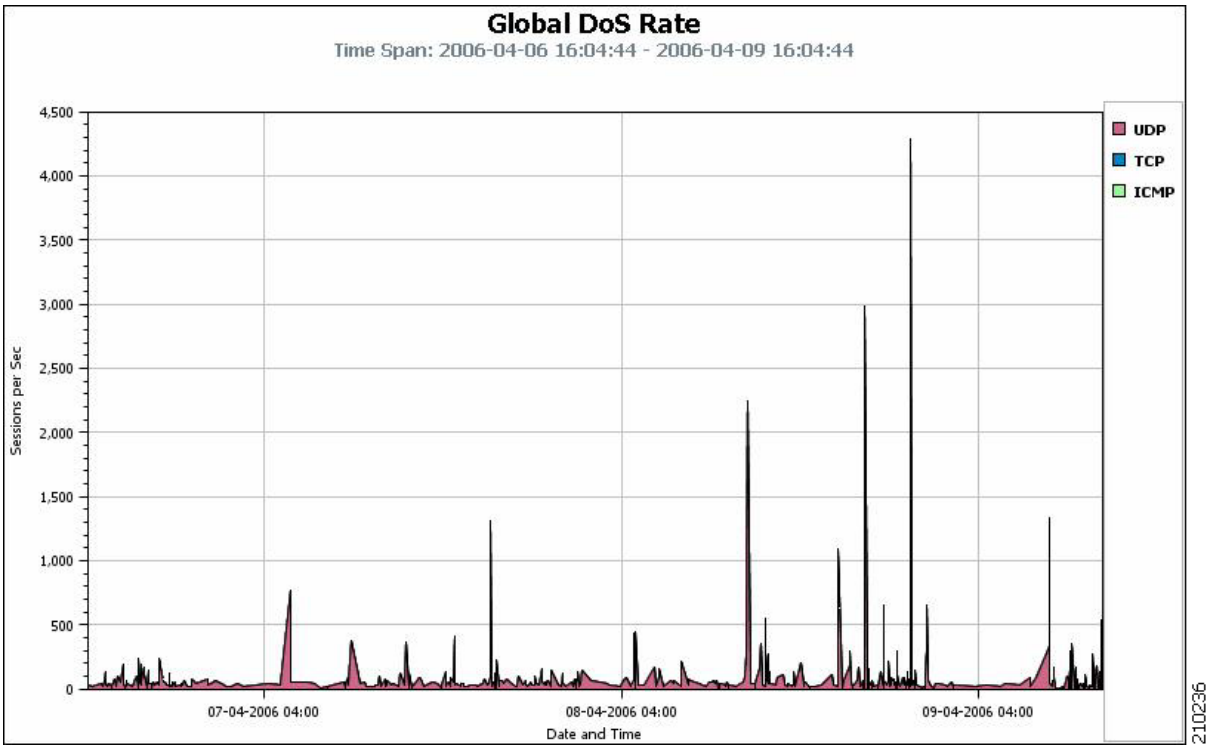

Infected Subscribers

The Infected Subscribers report is based on the records of scan/attack activity detected by the SCE platform, but not limited to a specific port.

The report estimates the number of infected subscribers over time (hourly time frames); this number represents the number of subscribers that were identified as generating malicious traffic during the hour in question. “Infected” reflects an assumption that traffic is generated by some malicious agent on the subscriber host.

The introduction of a new worm would typically trigger a considerable increase in the number of infected subscribers; therefore, this report is a good method of monitoring this activity.

Figure 3-13 Infected Subscribers Report

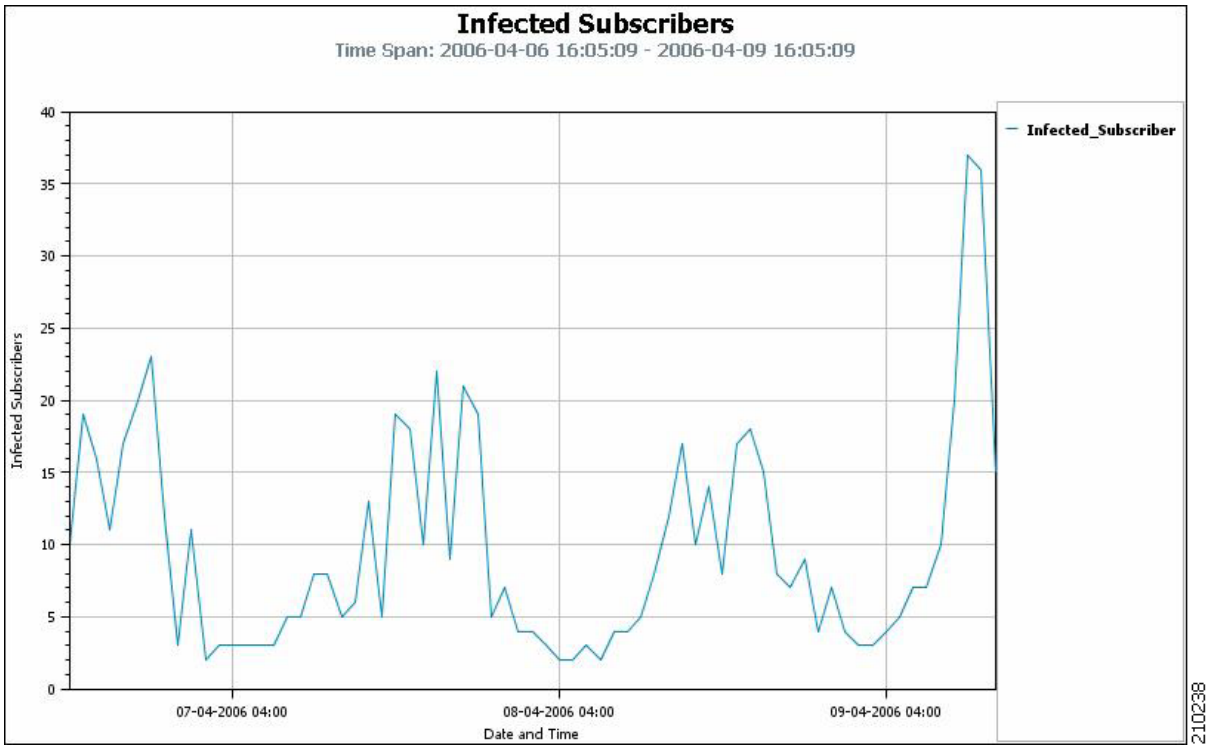

Top Scanning or Attacking Hosts

The Top Scanning or Attacking hosts report is based on the records of scan/attack activity detected by the SCE platform, but not limited to a specific port.

The report presents the top N scanning or attacking hosts during the time frame specified and can be filtered by subscriber or network, and by protocol.

Figure 3-14 Top Scanning or Attacking Hosts Report

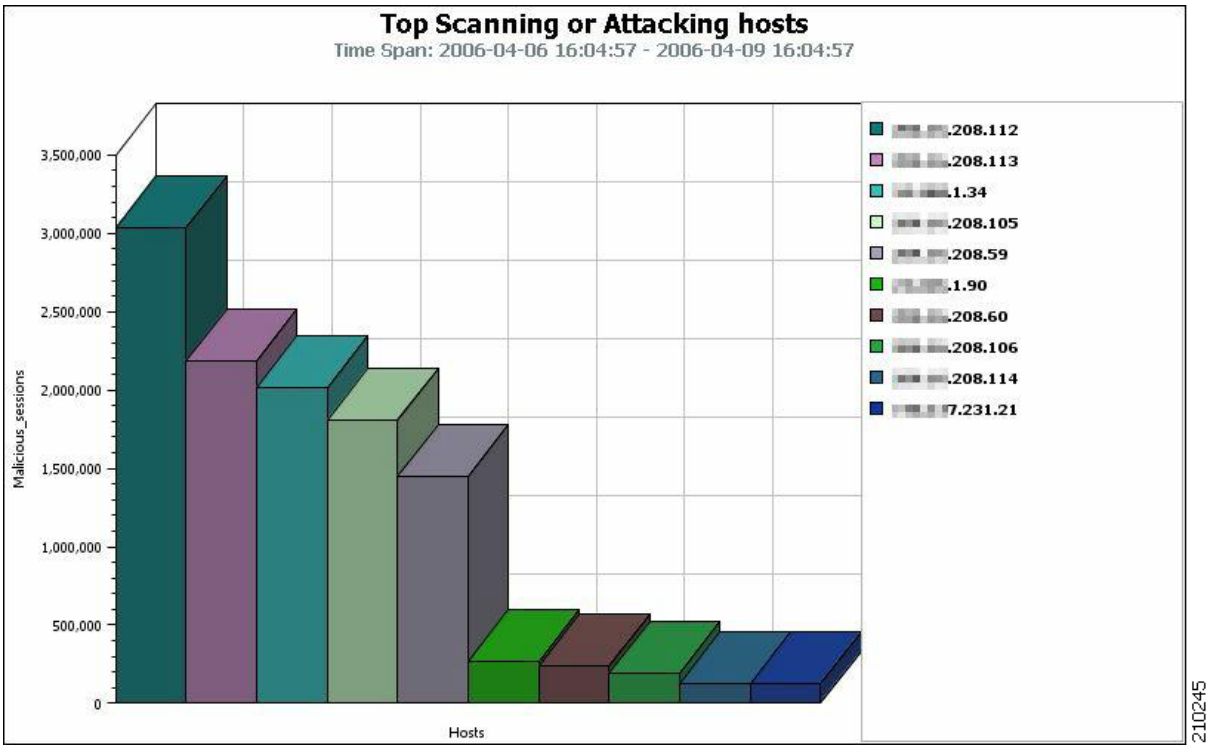

Top Scanning or Attacking Subscribers

The Top Scanning or Attacking Subscribers report is based on the records of scan/attack activity detected by the SCE platform, but not limited to a specific port.

This report presents the top N scanning or attacking subscribers during the time frame specified and can be filtered by protocol.

The N/A subscriber that typically exists in such a report is an aggregate of scan/attack traffic that could not be ascribed to a “named” subscriber—typically due to spoofing of the source IP address.

Figure 3-15 Top Scanning or Attacking Subscribers Report

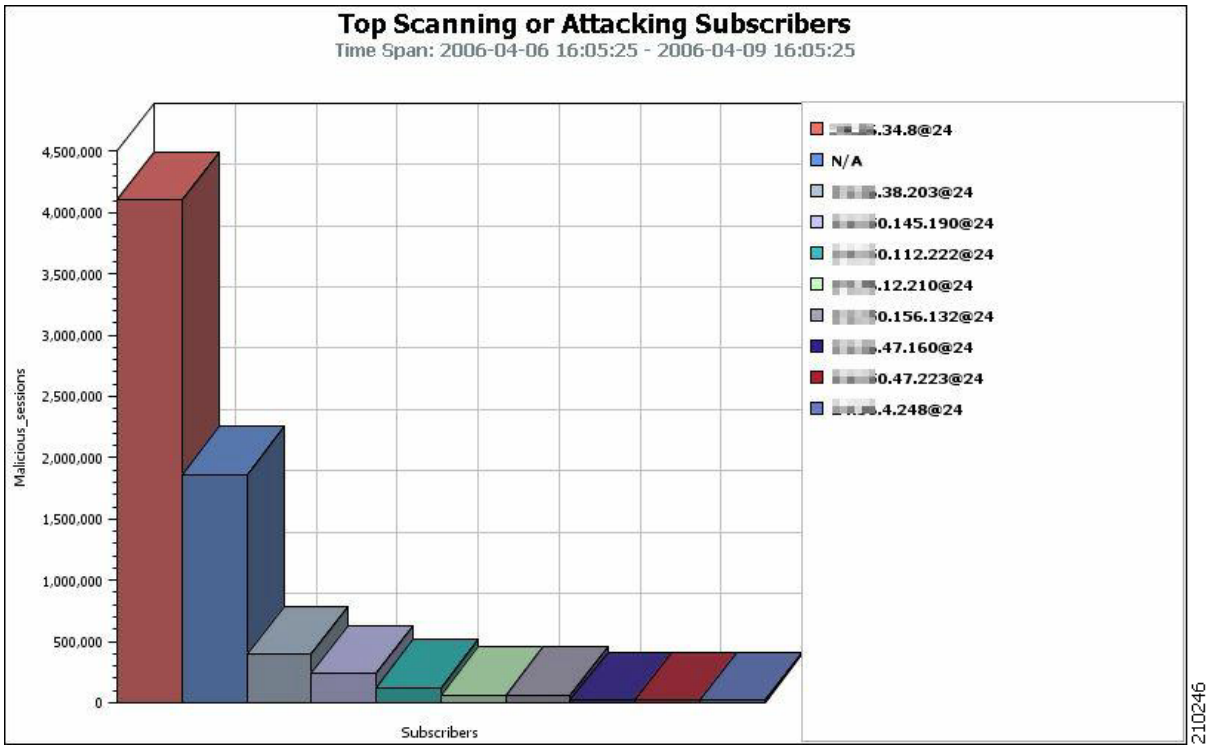

Top DoS Attacked Hosts

The Top DoS Attacked hosts report is based on the records of DDoS activity detected by the SCE platform, but not limited to a specific port.

This report presents the top N attacked hosts during the time frame specified and can be filtered by subscriber or network, and by protocol.

Figure 3-16 Top DoS Attacked Hosts Report

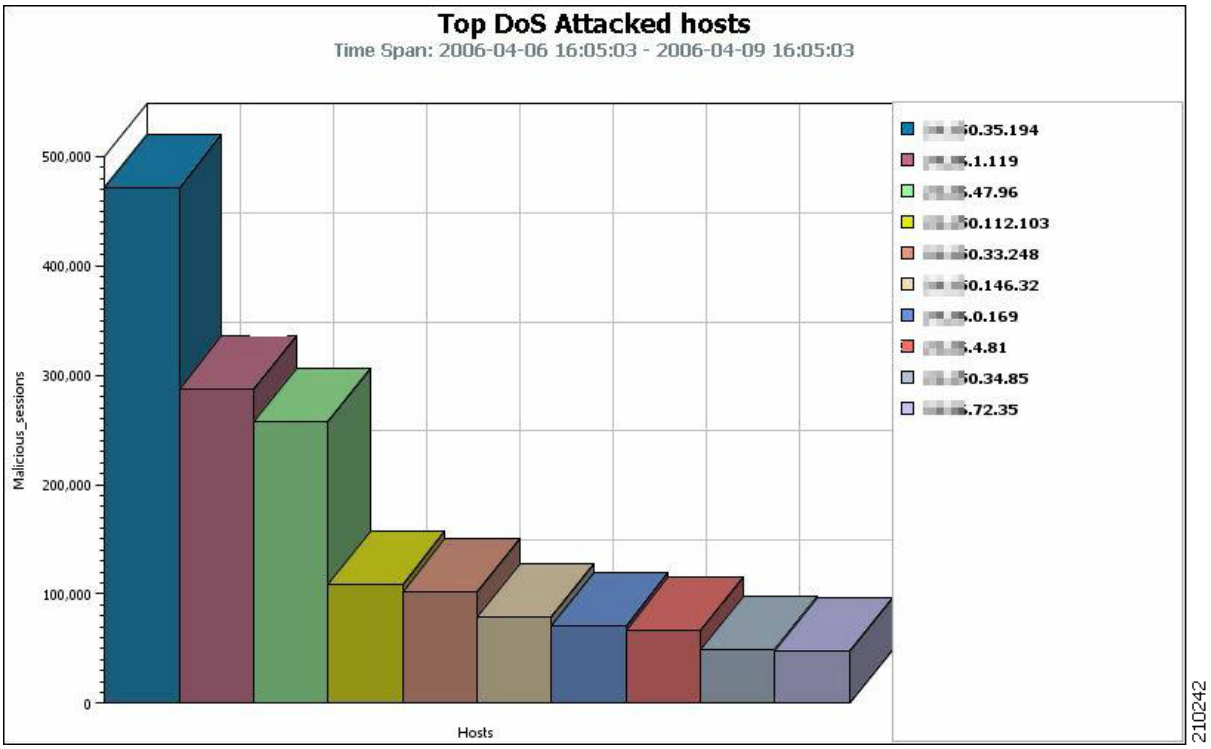

Top DoS Attacked Subscribers

The Top DoS Attacked Subscribers report is based on the records of DDoS activity detected by the SCE platform, but not limited to a specific port.

This report presents the number of attacked subscribers over time (hourly time frames); this number represents the number of subscribers that were identified as being attacked during the hour in question. The report can also be filtered by protocol.

Figure 3-17 Top DoS Attacked Subscribers Report

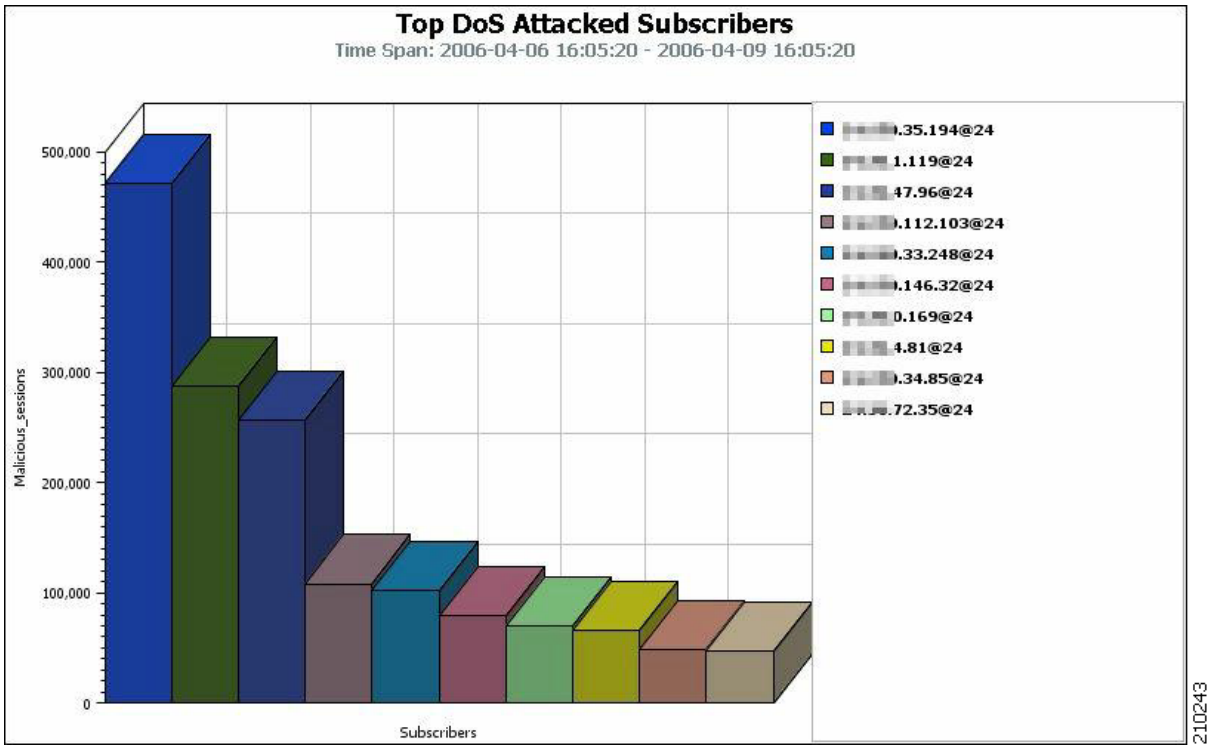

CHAPTER 4

Mass-Mailing Based Threats

This module introduces the concept of mass-mailing based threats and how to protect against them using the SCE.

Mass-Mailing Based Threats

The Mass-Mailing based threat detection module is based on monitoring SMTP session rates for individual subscribers. It uses the SCE platform's subscriber-awareness and can work in subscriber-aware or anonymous subscribers mode.

SMTP is a protocol used for sending email; an excess rate of such sessions from an individual subscriber is usually indicative of malicious activity involving sending email: either mail-based viruses or spam-zombie activity.

Configuration of Mass-Mailing Detection

Mass mailing detection is based on a subscriber breaching a predefined SMTP session quota.

In order for the functionality to work, the system must be configured to subscriber-aware or anonymous subscribers mode. This allows the SCE platform to accurately count the number of SMTP sessions generated by each subscriber.

Configuration is based on the following stages:

- Configuring the service for detection—The user should configure the appropriate service, which should have been built before this stage, for mass-mailing detection. It is common to use a service that includes only the SMTP protocol. Refinements can be made to narrow down the scope of detection and to potentially reduce the detection threshold.
  - "Outbound SMTP"—To account for only SMTP sessions generated by a subscriber. SMTP should not normally be seen as an inbound protocol because a subscriber is not expected to run an SMTP server on their own premises. Inbound SMTP connections may represent other kinds of malicious activity. To build such a service, a user should include the "outbound" attribute in the service definition.
  - "OffNet SMTP"—SMTP that is not targeted to a subscriber's "home SMTP server". Normal email clients send email through a home SMTP server, which later relays the email to wherever needed. Limiting the service to offNet can avoid accounting for "legitimate" sessions; i.e.

sessions that subscribers conduct with the SMTP server of their ISP. One caveat is that prominent non-ISP email providers such as Google and Yahoo! etc. have started providing an SMTP based service either for a fee, or free of charge, OffNet is no longer a suitable differentiator between "legitimate" and "non-legitimate" activity. To build such a service, a user should include an SMTP server list in the "onNet" service definition, which turns the rest of SMTP into "offNet".

  - A combination of the two.

- Define the quota to be used for indicating anomalous email activity. The quota is defined as a number of sessions for a given period—number of sessions and period length are both configurable. It is suggested that the user should base the values for these fields on some baseline monitoring of subscriber activity. See Monitoring Mass Mailing Activity, page 4-2.
- Define the action to be taken upon detecting Mass-Mailing activity. The action to be taken can be:
  - Block—SMTP or the more granular service over which detection is done is blocked once the quota is breached. The blocking is removed once the quota monitoring period ends. For example, for a 10 sessions / 60 seconds limit, blocking will be applied after 10 sessions occur within the quota period (quota period is started at an arbitrary point), and be removed once the quota period ends.
  - Notify—Redirect the subscriber browsing sessions to a captive portal presenting a message from the operator. This is done using "subscriber notification".
  - A combination of the two.

Figure 4-1 Spam Setting Window

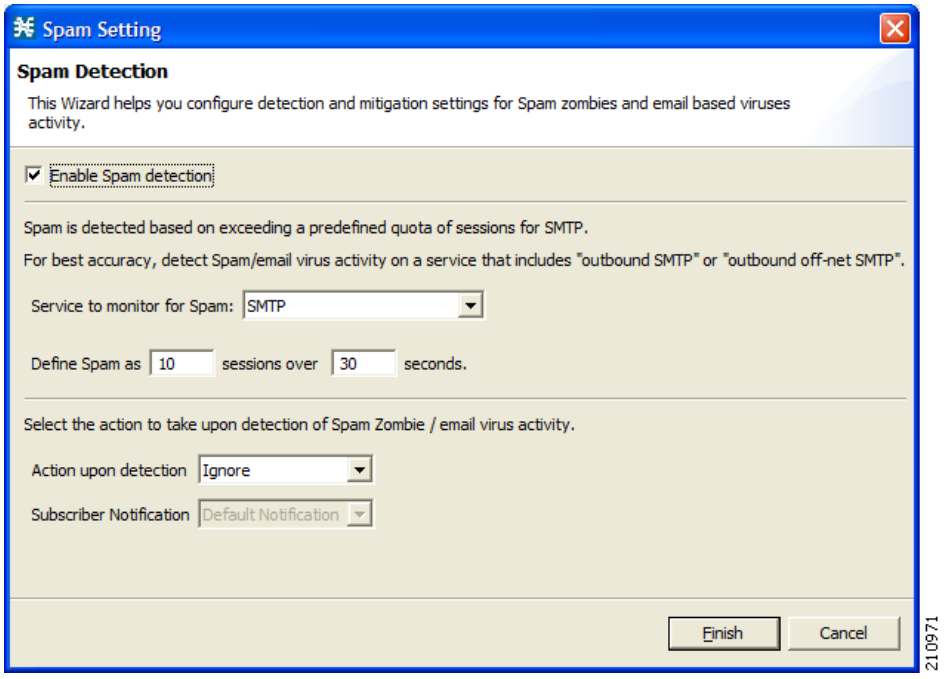

Monitoring Mass Mailing Activity

Mass mailing activity can be monitored based on information processed and stored in the CM database.

The most suitable report for detecting mass mailing activity by subscribers is the "top subscribers by sessions" report. This report is generated for the service and is used for mass-email detection (SMTP or a more granular service if it was defined). The report would highlight the IDs of subscribers most likely to be involved in mass mailing activity.

Figure 4-2 Top Subscribers Report