corero
NETWORK SECURITY

IPS 5500

Configuration and Management Guide

Corero Network Security, Inc.
990-0188-19

Legal Information

Contents

Figures

Tables

Preface

This guide contains information about using the Corero IPS 5500 platform, which include the Corero IPS 5100-series hardware and the Corero IPS 5200-series hardware. Throughout this guide, the Corero IPS 5500 products are referred to as IPS Units.

Because there are several models of the IPS 5500 platform, not all features described in this document are applicable to all models. Features that apply to particular models are clearly identified.

The release notes that are shipped with the product may contain more recent information that was not available when this guide was published. For the latest information, please refer to the release notes.

Audience

This guide is intended for use by network and/or security administrators who are responsible for installing, configuring, and using network security equipment. It assumes the reader has a high level understanding of network operations.

Revision Information

This is an updated book.

Related Books

This guide is part of the Corero IPS 5500 documentation set. The IPS 5500 documentation set includes the following:

Documentation	Description
IPS 5500 Release Notes	Information on known problems, bug fixes, and technical tips.
5100-Series Hardware Installation Guide	Detailed information on the IPS 5100-series platform hardware features, including installation and initial configuration.
5200-Series Hardware Installation Guide	Detailed information on the PS 5200-series platform hardware features, including installation and initial configuration.
IPS 5500 Configuration and Management Guide	Conceptual and Procedural information for configuring and managing the integration of the IPS 5500 product into your network. The guide describes network and port role settings, bridging, system setup, and configuration, management, and monitoring of traffic security features.
IPS 5500 and DDS 5500 Online Help	Available through the IPS management application, the online help system provides detailed descriptions of configuration parameters, procedures, and notes regarding use of the IPS 5500 product features. Note that the IPS and DDS products share a common help system, and where the features in those products differ, the online help clearly indicates this.

Accessing the IPS 5500 Documentation

You can access the release notes and the guides from the documentation CD-ROM that ships with your product. The documentation can be viewed with Adobe Acrobat Reader.

You access the IPS 5500 online help system from within the IPS management application. You can enter the online help system in two ways:

- To access the table of contents for the online help system, click the Help button in the upper right corner of the IPS management application.
- To view context-sensitive help which pertains to the IPS management application dialog box you are currently viewing, click the Help button on the dialog box.

Accessing Information on the CD-ROM

Each IPS product includes a documentation CD-ROM. To view the documentation on the CD-ROM:

1. Insert the documentation CD-ROM into your CD-ROM drive.

   If you have autoplay enabled on your computer, the documentation menu displays.

   If you do not have autoplay enabled on your computer, click the `autorun.exe` program, located in the root directory on the CD-ROM. The documentation menu displays.

2. Click the documentation you wish to view.

Conventions

This book uses the following notation conventions.

- The notation Menu > Choice indicates that you should choose an item from a menu. For example, the following notation means, "Choose the Exit item from the File menu."

  Choose File > Exit.

- `Monospace` represents text that would appear on your display screen (such as commands, functions, code examples, and names of files and directories).
- *`Monospace italic`* represents terms that are to be replaced by literal values. The user must replace the monospace-italic term with a literal value.
- **`Monospace bold`** represents user input in examples and figures that contain both user input and system output (which appears in monospace).
- *Italics* is used to emphasize text.

Customer Support and Services

Corero Network Security offers two options for contacting Customer Services and Customer Support.

- Contact the Customer Services Center **by phone at + 1 978-212-1500**
  - Support is available for all customers with a Hardware or Software Warranty from 8:00 AM to 5:00 PM (Eastern US Time).
  - If you have purchased the Software Subscription Service, you can obtain service 7x24 by calling the support phone number and pressing Option 2. If the issue is critical, press Option 2 then Option 7.

NOTE

If, for any reason, the primary support phone number does not work, call Corero's answering service at +1.888.324.1246 (US) or +1.603.645.4145 (International) and a support representative will return your call.

- **On the web through the Customer Support Portal: https://support.corero.com. The Web Portal is the most effective way to log and track support issues.**

  This Portal provides:

  - Web-based incident management and customer support tracking system
  - Service request communications
  - Access to downloadable files including software and product documentation
  - An extensive knowledge base.

When you contact the Customer Services Center for assistance, have the following information ready:

- The case number, if you are calling about a previous problem
- Your name, and if someone else will be the contact person for the problem, the contact person's name.
- Your company name and location (city, state or province, and country)
- The telephone number (including area code) at which you or the contact person can be reached.
- The email address at which you or the contact person can be reached.
- The product name, model number, and serial number.
- A list of system hardware and software, including revision levels.
- A detailed problem description:

  Describe the symptom and the activities that preceded it.

  Include details about any recent configuration changes, if applicable.

  Be as specific as possible.

  Briefly describe your trouble-shooting steps and observations.

NOTE

When requesting support, problem resolution can go more quickly if you have access to the Documentation CD-ROM that accompanied your product.

For more information on Corero customer service and support programs, see Corero Services Overview (page 1-8)

How to Comment on This Book

At Corero Network Security, our goal is to provide the highest quality products and services to our customers. We value customer feedback and encourage users of Corero's systems to send their comments on the product, service, and documentation, so that we can continue to improve our products.

Please send your comments and suggestions, including features you would like to see in future releases, to the following address:

Customer Support and Services Center
Corero Network Security, Inc.

1 Cabot Road
Hudson, Massachusetts 01749 USA

Chapter 1
Corero's Three Dimensional Protection

Network security is a complex issue in our modern, resource-critical, multi-function, network environment. In order to truly protect your network, you must prevent improper use or overuse of your valuable network resources. Corero products guard your network and computer resources against resource misuse, helping to ensure your business can operate at peak performance.

This chapter introduces the concept of network intrusion prevention, describes Corero's three dimensional protection, and provides an overview of Corero products and their primary features. Finally, it provides examples of Corero product deployment.

This chapter contains the following sections:

What is Network Intrusion Prevention?

A Network Intrusion Prevention System is an in-line security appliance that inspects network traffic, identifying malicious, harmful, and/or unwanted network activity and blocking it. The traffic inspection performed by Corero Network Devices is done in real-time to ensure that good network traffic is able to pass through the device without noticeable delay.

There can be some overlap of functionality between Network Intrusion Prevention Systems and traditional firewalls, but it is clear that a firewall is not sufficient to protect against today's cyber threats. While each class of devices can block certain types of network transactions, how they affect networking configuration, how they perform traffic inspection, and how they approach system security are fundamentally different.

As a networking component, unlike most firewalls that also act as routers, a Network Intrusion Prevention System is a transparent device on the network that does not have a visible IP address, and requires no network reconfiguration to deploy. While a firewall's basic task is to regulate the type of network “conversations” that are allowed between computer systems of differing trust levels, a Network Intrusion Prevention System’s job is to inspect protocol and application content on the network to ensure that it does not contain harmful, malicious, or unwanted content. Rate-based algorithms protect against traffic floods, built-in stateful firewall filtering blocks unauthorized access to specific network assets, and finally, with the IPS rule sets and acceptable application use policies, users can define what types of traffic can pass to specific applications.

One of the unique features of Corero Network Devices is the protection they offer against Distributed Denial of Service (DDoS) attacks.

What are Distributed Denial of Service Attacks?

Corero Network Devices are designed to protect against Distributed Denial of Service (DDoS) attacks. These devices provide connection limiting and SYN flood limiting, and also offer targeted rules specifically designed to block HTTP and DNS attacks.

A Distributed Denial of Service (DDoS) attack is a cyber attack in which many, usually compromised, computers send a series of packets, data, or transactions over the network to the intended attack victim(s) in an attempt to make one or more of the victim’s computer-based services (such as a web application) unavailable to its intended users. DDoS attacks generally result from the concerted efforts of one or more malicious agents to stop an Internet site from functioning efficiently or at all.

Corero Network Devices mitigate the affect of both network-layer and application-layer DDoS attacks.

A DDoS attack is said to be a network-layer DDoS attack when it involves sending a flood of packets over the network at high volume to disrupt or overload the network infrastructure to the point where the infrastructure cannot transmit requests or responses, essentially making complete service transactions impossible. Network-layer DDoS attacks typically affect ISP links, routers, switches, firewalls, and servers, causing one or more of them to become bottlenecks, restricting or eliminating the ability of the server to deliver its service.

Application-layer DDoS attacks are a newer variant of this attack type. These attacks not only send network packets, but they actually complete TCP connections from the attacker to the victim service. Once the TCP connection is made, the attacking computers make repeated requests to the application in an attempt to exhaust the resources of the application, rendering it unable to respond to all of its other requests. These intelligent attacks are harder to defend against because they create denial of service conditions without causing the consumption of available network bandwidth, or overloading routers, firewalls, and switches. A repetitive HTTP GET request or DNS request is a common example of a transaction associated with application-layer DDoS attacks.

What is Three Dimensional Protection?

Corero Network Devices mitigate attacks in three major threat categories:

- Stops malicious content in network traffic, including exploits of Microsoft vulnerabilities, worms, Spyware, and other malware.
- Prevents undesired access to networks or systems, including unauthorized or illegal access.
- Defends against rate-based attacks on the infrastructure, such as SYN floods, and other Denial of Service attacks. These are attacks whose network traffic seems legitimate on the surface, but is not

In a fully protected network, all three attack approaches must be covered. A security gateway must act as a firewall (stop undesired access), an intrusion protection system (stop malicious content), and a rate based controller (stop flooding attacks).

Figure 1-1: Three Dimensional Protection

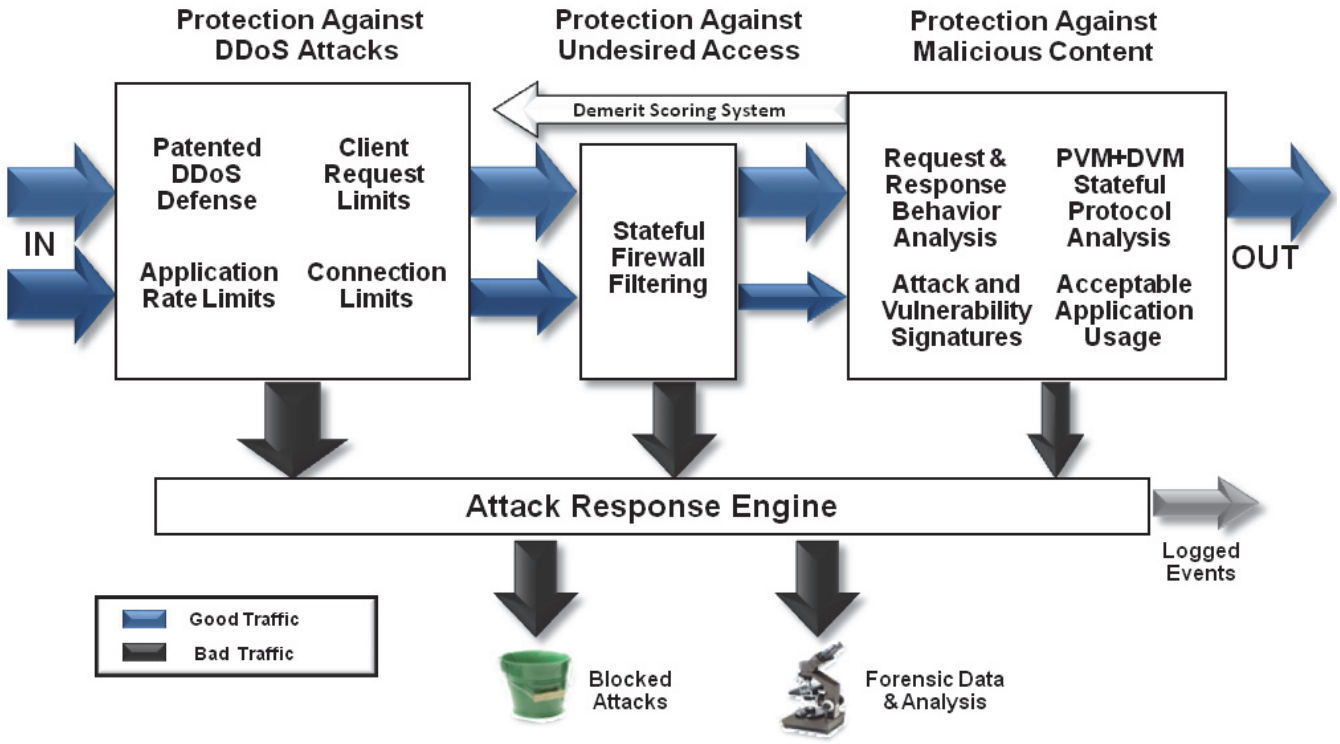

Corero Product Overview

Corero Network Security offers two Corero Network Device product families that provide Network Intrusion Prevention: IPS 5500 Units and DDS 5500 Units. These devices can be individually managed using a device-specific management application or, in order to manage multiple devices, you can purchase Corero's IPS Controller software.

The Corero **IPS product family** provides broad coverage for network resource threats. These products focus on content based protection for clients and servers, including firewall and rate-based protection, including DDoS-specific detection and mitigation. This product family also addresses client/browser exploits, malware protection, support for deep packet inspection, and awareness of a large number of protocols.

The Corero **DDS product family** provides focused coverage for network resource threats. Server protection is the primary focus, including DDoS-specific detection and mitigation. The available rules are defined and targeted for server-specific protection. The Corero DDS product family is more cost-effective for sites with DDoS-only requirements.

In order to simplify and streamline monitoring and management of multiple IPS and DDS Units, Corero offers the **IPS Controller**. The Corero IPS Controller is designed to provide centralized management for Corero Network Devices from both the IPS 5500 and DDS 5500 product families. The IPS Controller simplifies Corero Network Device management by creating policy groups, which are groups of devices that you can manage as a single entity. You can also create Corero ProtectionClusters out of two or more identical model units, providing high availability and high throughput processing. The IPS Controller also allows you to acquire and deliver Corero protection packs to ensure the latest protection for your network.

SecureCommand: a Centralized Management Solution

SecureCommand is the centralized management solution used to manage Corero Network Devices (IPS and DDS Units). It provides essential real-time security intelligence to help assess hacker/virus behavior, combat security threats, and meet regulatory compliance requirements across the IT infrastructure. SecureCommand provides convenient device management, event correlation, and robust scalable reporting.

The SecureCommand solution is comprised of several components:

- IPS Controller (page 1-5)
- Network Security Analyzer (page 1-5)
- TopResponse Updates (page 1-7)

IPS Controller

The IPS Controller tracks and manages software updates, TopResponse™ updates, and policy/configuration of multiple IPS and DDS units. Features include the following:

- Real-time status display for all Corero Network Devices, high-availability clusters, and groups
- Distribution of configuration and policies to one or more Corero Network Devices.
- User-defined groupings of Corero Network Devices simplify management tasks
- Easy management of high-availability clusters
- Corero Network Device software upgrade management
- Off-line editing of Corero Network Device configurations and policy using the GUI
- Off-line validation of configuration and policies
- Fully detailed audit trail for configuration and policy changes.
- High-level “dashboard” summary display
- TopResponse research and automated update service
  - Keeps protection and management elements up-to-date
  - Policy and signatures
  - Internet Topology information
  - Spyware site information

Network Security Analyzer

The Network Security Analyzer (NSA) provides security professionals with the essential real-time security intelligence to help identify and understand hacker, virus, and SPAM/spyware behavior, security breaches, denial-of-service, and unauthorized access to sensitive information.

NSA helps minimize incident response time by automatically collecting and correlating event data from a variety of multi-vendor network devices - routers, switches, firewalls, VPNs, IPS systems, and proxy servers, as well as anti-spyware, antivirus, SPAM management, and content filtering web security appliances. Reported information helps eliminate false positives, identify security breaches and corporate violations, improve security operations, and deliver the necessary tools to meet Sarbanes-Oxley, PCI, GLBA, HIPAA, and FISMA compliance.

In today’s environment, one of the primary key features for a security management solution is the ability to scale to large networked environments. Network Security Analyzer provides a distributed architecture for small to medium enterprises that scales to thousands of network devices. The architecture supports both a standalone deployment for smaller networks and a distributed deployment for medium enterprise installations. The flexibility of the NSA

architecture allows for the creation of a security information and event management solution that can adapt to any environment

The architecture allows MSSPs to take advantage of out-of-the-box reporting and monitoring portals to offer new value-added revenue generating services or expand their current remote monitoring services to include comprehensive on-demand reporting and compliance audit log management. The built-in XML based API allows MSSPs and enterprise customers to integrate NSA's reporting, alerting, and monitoring data with other third-party portals.

NSA real-time monitoring and alerting features include:

- Heterogeneous Real-time Monitoring: Monitors security event data across the entire network of security devices in real-time.
- Real-time Correlated Alerting: Template driven Alert Manager allows creation and definition of any number of alerts to reduce false positives and identify blended attacks.
- Real-time Event Manager: View security event data from thousands of heterogeneous and multi-vendor network devices and prioritize the actions based on business impact of each event, allowing for corrective actions before an incident occurs.
- Event Drill-down: Advanced on-the-fly event correlation and analysis of significant security events.
- Monitoring Dashboard: Monitoring dashboard provides a quick, consolidated view of the environment. Create and view any number of user specific monitoring views and toggle between the different views.

NSA security reporting features include:

- Reporting Portal with Powerful Drill-down: Reporting portal gives access to over 600 reports. Powerful drill-down feature displays 2nd and 3rd level details with a single click.
- Correlated Reporting: Get a holistic view and understanding of hacker and virus activity by correlating data across all network devices instead of looking at each device's data separately.
- Intrusion and Rule based Reporting: Through over 50 attack and rule based reports, NSA provides essential information to help security administrators get a comprehensive understanding of the intrusions and rule violations.
- Protocol and Web Usage Reporting: Get a firm handle on protocol and web usage patterns by user, department and/or device.
- SPAM and Spyware Reporting: Generates over 30 SPAM and spyware activity related reports.
- Antivirus Reporting: Generates over 100 anti-virus activity related reports that identify the presence of viruses across small and medium enterprise networks.
- Vulnerability Reporting: Integrates and reports on vulnerability data derived from NESSUS vulnerability scans.
- Content Categorization Reporting: Generates content categorization related reports to help understand employee web usage patterns.
- Automated Report Generation/Distribution: Generates more than 600 reports. E-mail reports automatically to multiple recipients in HTML, MHTML, PDF, Word, Excel and text formats.

NSA compliance audit lifecycle management (CALM) features include:

- Automated Log Archiving for Compliance: Automatically compresses, encrypts and archives log for investigative analysis and regulatory compliance.
- Compliance Monitoring: Centralized monitoring and alert correlation for real-time investigation of security incidents with regulatory compliance implications.

- Compliance Reports: Detailed reports to Sarbanes Oxley (SOX), Health Insurance Portability and Accountability Act (HIPAA), Gramm-Leach-Bliley Act (GLBA), and the Federal Information Security Management Act (FISMA).
- Scalable Search: An easy-to-use mechanism to search hundreds of GB of log data across multiple devices based on user search criteria to aid in investigative/forensics analysis.
- Activity Investigation: Identify anomalies and employee corporate policy violations.

TopResponse Updates

The IPS Controller also includes access to TopResponse updates. TopResponse is an Automated Protection Update program that provides Corero Network Device customers with advanced security services to maximize security, availability, and performance of their network.

It offers proactive protection from zero-day threats and resolution to security issues. Specifically, TopResponse provides automated updates, technical support, security advisory and software subscription services, along with access to Corero's Knowledge Base.

NOTE

For customers who do not use the IPS Controller central management system, there is a standalone TopResponse update software utility designed for use with IPS Units called the TopResponse Update Manager. This application enables downloading and activating security updates to each of your deployed Corero Network Devices.

Corero Services Overview

Table 1-1 describes the provisions available to you if your Corero equipment is covered under a one or more Corero service agreements or warranties.

NOTE

For full Corero product licensing and warranty information, see http://www.corero.com/en/support/end_user_agreements.

All customer service agreements provide access to Corero's web-based customer request tracking and ticketing system.

NOTE

When you purchase a Corero product, Hardware Warranty support (12 months from the date of shipment) and Software Warranty support (90 days from the date of shipment) are included.

Table 1-1: Corero Service Features

Corero Service	Service Agreement Provisions
Hardware Warranty or Advanced Hardware Replacement Service	• Telephone support from Monday through Friday, 8AM to 5PM Eastern (US) Time. For information on how to contact Corero Customer Service, see the section titled Corero Services and Support in the preface of this guide. • Hardware unit repair or replacement. Replacement includes door-to-door delivery. • A Hardware Warranty qualifies you for same business day shipment for product replacement ahead of damaged unit return if the product is delivered in a damaged or inoperative state. The Advanced Hardware Replacement Service provides the same replacement service during the lifetime of the product.
Software Warranty	Telephone support from Monday through Friday, 8AM to 5PM Eastern (US) Time. For information on how to contact Corero for services and support, see the section titled Corero Services and Support in the preface of this guide.
Software Subscription Service	• Telephone support 24x7. For information on how to contact Corero for support and services, see the section titled Corero Services and Support in the preface of this guide. • Notification of software releases. • Entitlement to all major, minor, and maintenance releases and downloads. • Access to the Corero Support Knowledge Base.
Threat Update Service	• Protection Packs that include updated signatures, filters, configuration files, rules, and malicious IP addresses. • Attack advisories delivered by signed Email.

Table 1-1: Corero Service Features *(Continued)*

Corero Service	Service Agreement Provisions
SecureWatch Service or SecureWatch PLUS Service	In order to support these services, the customer typically purchases Corero SecureCommand (IPS Controller, Network Security Analyzer, and Threat Update Service), the Software Subscription Service, and the Advanced Hardware Replacement service. When a customer purchases the **SecureWatch Service** or the **SecureWatch PLUS Service**, Corero customer support will: • Test to verify connectivity to Corero products. • Work with customer to create a change management process for deliver and installation of Corero software updates and security updates. • Provide the customer with notification of Threat Update Service advisories and Corero software updates • Monitor Corero device operation and automatically initiate the Advanced Hardware Replacement Service if a problem is detected. • Verify that NSA reports are successfully generated, and report to the customer if they are not. • Provide a central means of contact and trouble reporting. • Provide weekly configuration, performance, fault, and security activity reports via email • Provide automatic backup of modified configuration files. • Apply software updates within 2 days of customer approval using the customer-specific change management process. • Apply Threat Update Service protection packs within 1 business day of customer approval using the customer-specific change management process. • Apply Threat Update Service security advisory implementations (including rule modifications) within 2 business days of customer approval using the customer-specific change management process.

Table 1-1: Corero Service Features *(Continued)*

Corero Service	Service Agreement Provisions
SecureWatch PLUS Service	In addition to all of the features available in the SecureWatch Service, SecureWatch PLUS provides the following: • Corero assigns a named Technical Account Manager to the customer with overall responsibility for service delivery and customer communications. The account manager will visit the customer's location twice a year to meet with the customer and discuss all aspects of Corero services. • Around-the-clock monitoring and support by the state-of-the art Corero Security Operations Center (SOC). and 8 AM to 8 PM access to Corero SOC staff. • Ongoing tuning and optimization to defend against changing attack vectors. • Apply configuration updates within one business day of customer approval. • Network Security Analyzer report and alert generation and verification • Generate an audit report on the customer IT environment including topology, protocols, traffic types, average traffic flows, and network usage. Customized configuration to conform to the customer's policies and requirements. • Create and deploy a customized and complete defense configuration for all customer equipment based on the customer's security policy, business objectives, and security best practices • Monthly communications between the technical account manager and the customer about network environment, threat awareness, defense configuration maintenance, and the attack response plan. • List critical monitored conditions that would signal the onset of a DdoS attack at the customer location • 24x7 availability of Corero defense expertise in the event of attack, providing and coordinating support and according to the attack response plan. Corero engagement continues until the attack is mitigated. Initial response to the attack will occur in less than 1 hour, and reports will be given every two hours. • Creation of post-incident report with attack assessment, impact, and recommended measures to improve prep & response in the future • Formulation of a joint customer-Corero incident response plan. • Immediate and continuous engagement through the duration of an attack. • Post-incident analysis and recommended follow-up action after an attack.

Corero Network Device Rate-Based Protection

Rate-based protection is applied in a specific order, and this is the first step in the overall detection process:

1. When a Corero Network Device receives traffic, it looks up the IP address and gathers the information associated with that address.
2. When performing rate-based protection, the Corero Network Device simultaneously assesses whether any rate-based policies have had rules triggered in the three assessment areas: SYN Flood Mitigation, Client Request Limiting, and Connection Limiting.
3. If the traffic passes the rate-based protection check, it is processed by the applicable Firewall policy.

   NOTE

   Note that Application Rate Limiting is performed as part of the Firewall inspection process. Application Rate Limiting is completely separate from Packet Rate-Based detection.

4. If the traffic passes the Firewall policy check, it proceeds to the IPS Policy check.
5. Only once it has progressed successfully through these mitigation processes is the traffic permitted to pass beyond the Corero Network Device.

Figure 1-2 shows the order in which protection is applied

Figure 1-2: Traffic Processing Order

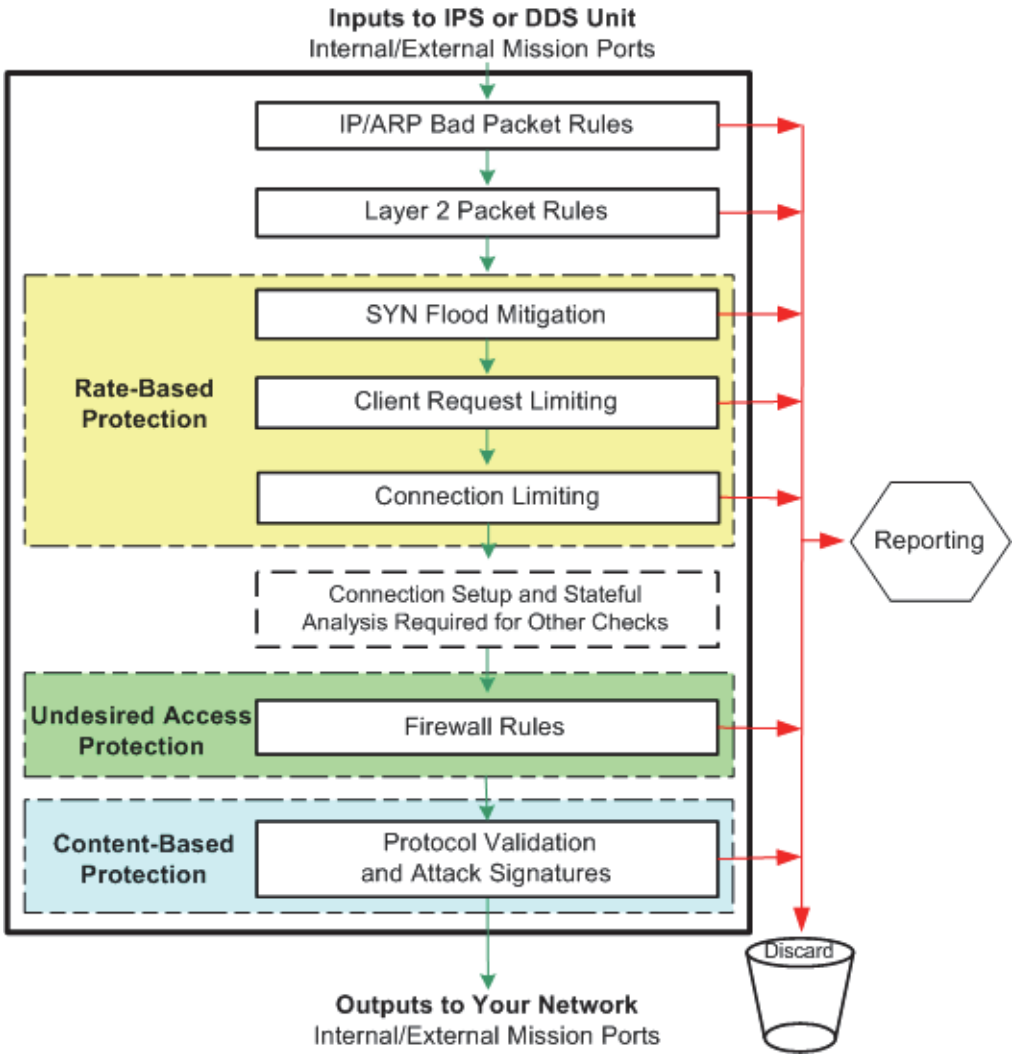

Corero Network Device Packet-Based Stateful Analysis and Connection Setup

Each Corero Network Device records critical information about each packet in a flow record stored in its Flow Table (connection table), an internal memory structure. Recording this information enables the Unit to statefully inspect each packet and to reorder packets for proper analysis. At the start of each transaction, the IPS or DDS Unit creates the appropriate Flow Table entries. It then checks these entries when it receives subsequent packets for the same transaction.

Corero Network Devices provide best-effort service in the face of resource exhaustion or overload conditions. The main resource loss that causes connectivity problems is Connection Table (also called the Flow Table) exhaustion. The Unit avoids this resource problem for critical applications by allowing you to control which applications the system will inspect and mitigate.

By reserving Flow Table space during periods of high resource usage, a Corero Network Device maintains the ability to record critical information about key applications under these conditions.

You can configure the device not to create flows for any non-mission-critical applications, including TCP-based applications that tend to create large numbers of connections.

NOTE

Non-IP packets do not create flows and do not have to be restricted.

Note that Corero Network Devices can only perform stateful analysis and deep packet inspections for a given transaction if they can create a flow. Corero Network Devices also have hardware support for fast reclamation of properly terminated TCP flows and all aged out flows. Fast reclamation allows the flow resources to recover quickly for use by new connections.

High Availability: The ProtectionCluster™

Corero Network Security products are designed with High Availability (HA) in mind. High-MTBF hardware design with no rotating media, redundant hot-swappable power supplies, and a hot-swappable N+1 fan tray ensures non-stop operation. Port bypass on all internal and external network ports ensures network availability even in the unlikely event of an internal failure.

Multiple Corero Network Devices can be combined into a ProtectionCluster, offering protection from a single-point of failure, and supporting continuous state sharing between devices to ensure continued network operation, even in the event of a fail-over.

5200 Series Corero Network Devices offer two 10 Gb HA interfaces, and 5100 series Corero Network Devices offer four dedicated gigabit-speed HA interfaces. These ports allow automatic traffic balancing and continued communication with other devices in the ProtectionCluster, even if one of the HA links is down. ProtectionClusters can be configured to support High Availability / Redundancy, High Throughput, or both.

Suggested Corero Network Device Deployment Locations

Before you can deploy a Corero Network Device in your network as an inline device, you need to decide on a deployment mode. Depending on your particular configuration, you may want to place one or more IPS or DDS Units in the locations described in Table 1-2.

Network Intrusion Prevention systems such as Corero Network Devices are suitable for both perimeter and core deployments. Perimeter deployments typically place the device behind the firewall, allowing the firewall to apply its access controls first, and then the device further inspects traffic that the firewall has allowed through. Corero Network Devices have advanced DDoS protection capabilities that make them well suited to deployment in front of the firewall, preventing the firewall from becoming a single point of failure in the event of a botnet attack.

The figures in the following sections provide general guidelines for placing your Corero Network Device. For more detailed help, contact your Corero vendor.

NOTE

Corero Network Devices are deployed inline versus offline or in passive mode whereas IDS configurations using SPANs or Taps are always passive.

Table 1-2: Deployment Locations

Configuration	Protection	Placement
Critical Online Asset Protection	Protects network segments from threats and provides containment of infected segments.	Place your Corero Network Device in front of your Internal network. For more information, see Protect Critical Online Assets (page 1-16).
High Throughput	Provides additional, shared processing for high volume environments.	Configure two Corero Network Devices in a Single Inline With Peer configuration. For more information, see High Volume Configuration (Single Inline with Peer) (page 1-17).
ProtectionCluster	Provides active redundancy to your current configuration.	Add multiple, redundant, Corero Network Devices. For more information, see ProtectionCluster Configuration (page 1-18).
Network Perimeter	Increases protection against targeted DDoS attacks and application-level threats.	Place your Corero Network Device in front of the Firewall. For more information, see Protect Your Network Perimeter (page 1-19).
Network Perimeter	Protects the network from cyber-threats that may traverse the VPN link. Place your IPS Unit behind the VPN concentrator.	For more information, see Protect Critical Online Assets (page 1-16).
Critical Online Asset Placement (dedicated server protection)	(Protects assets from network and application level threats regardless of whether they originate from inside or outside.	Place your IPS Unit in front of a server farm, Intranet, or Extranet. For more information, see Protect Your Hosting Center (page 1-20) and Protect Servers in Your Enterprise (page 1-21).

Protect Critical Online Assets

Figure 1-3 shows three Corero Network Devices deployed to protect critical online assets.

Figure 1-3: Protecting Critical Online Assets

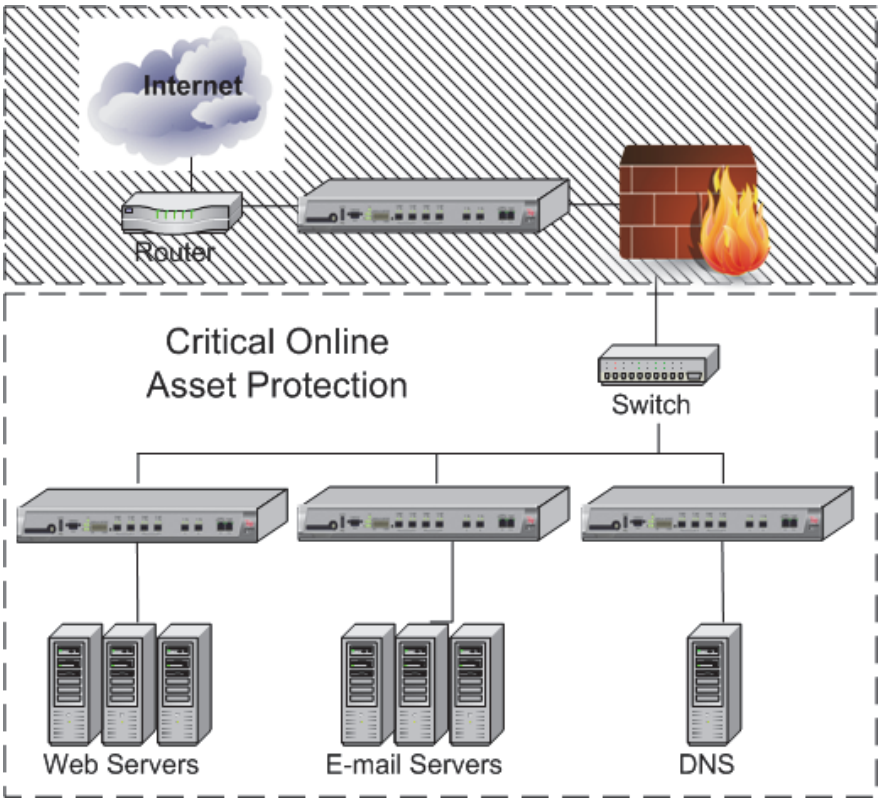

High Volume Configuration (Single Inline with Peer)

Figure 1-4 shows a Corero ProtectionCluster deployed for high volume environments. In the Single Inline with Peer (Leaf Node) configuration, only one Corero Network Device passes network traffic, but the second device assists in detection processing and flow setup operations, dramatically increasing the traffic load that the devices can handle and almost doubling the number of connections that can be created and analyzed.

Figure 1-4: Inspecting High Volume Traffic

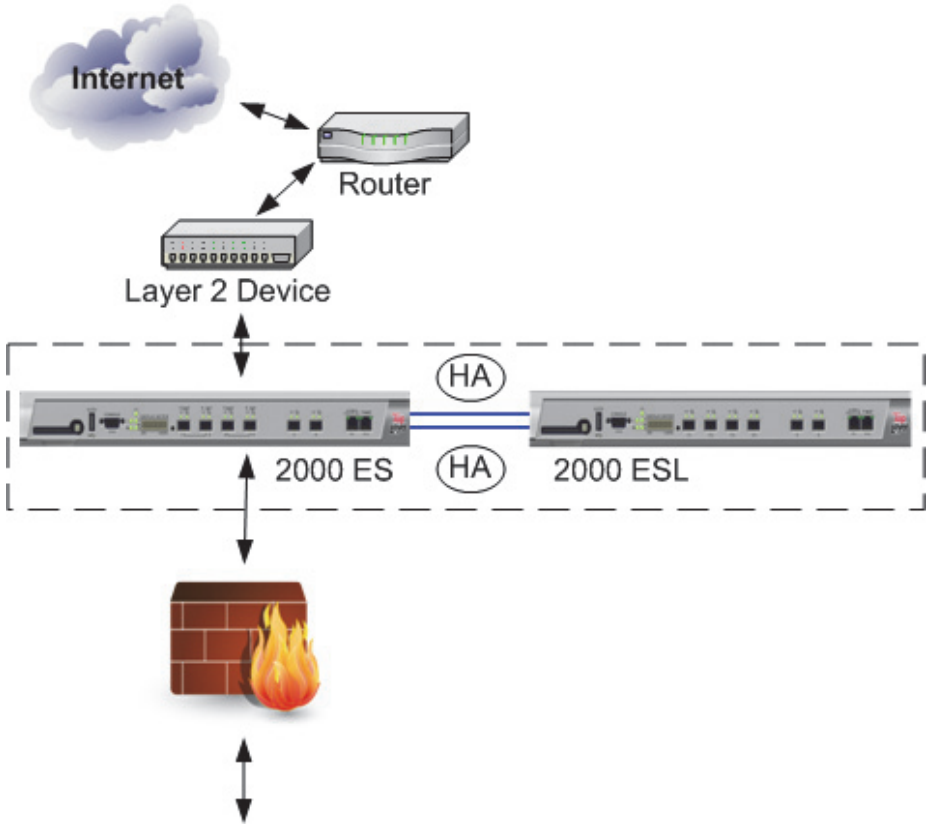

ProtectionCluster Configuration

Figure 1-5 shows a two-unit ProtectionCluster. More Corero Network Devices can be added to a ProtectionCluster.

Refer to the Release Notes for your product for information about the maximum number of Units supported in a ProtectionCluster.

Figure 1-5: Two Unit ProtectionCluster Configuration - 5100 Series

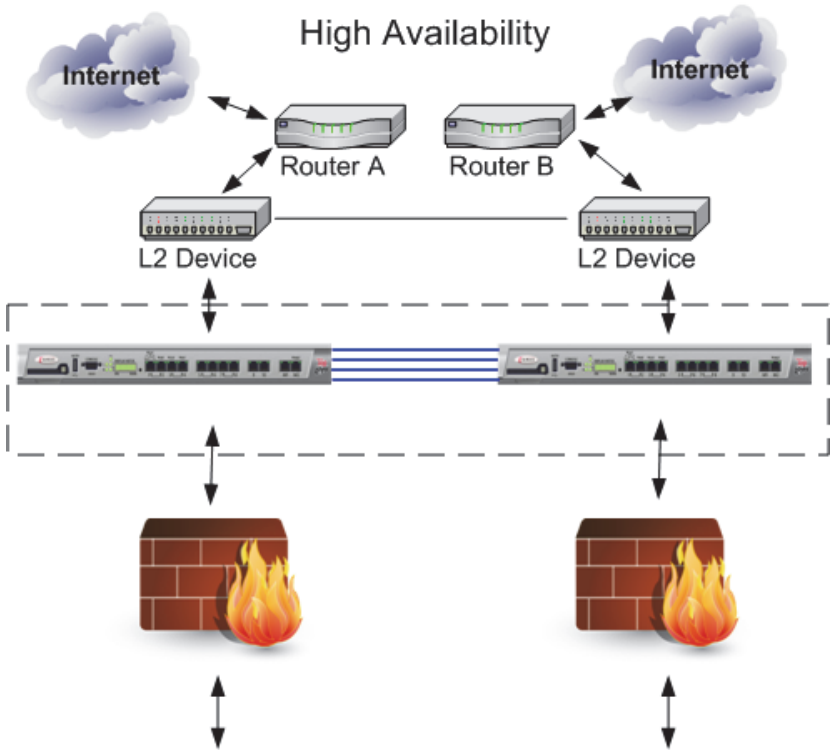

A ProtectionCluster refers to a network configuration option that provides higher bandwidth and redundancy. This configuration connects multiple Corero Network Devices together. On the 5100-Series Corero Network Devices, this is done by using up to four of the 10/100/1000 ports that can be configured for this purpose, HA1 through HA4. On the 5200-Series Model 2000 Units there are two 10,000 ports available. On the 5200-Series Model 2400 Units up to eight 10,000 ports available. A ProtectionCluster configuration also enables the Units to share the intense processing required for deep and stateful protocol analysis necessary to detect attempted exploits of application-level vulnerabilities.

In this configuration, both sides of the configuration receive and pass your network traffic, unless there is a failure. This solution, using a combination of Corero Network Devices, protects up to two full duplex Gigabit input ports: stopping the bad traffic, while permitting the “good” traffic to pass to its destination

If Corero Network Devices will not be co-located, consider the following maximum link distances for fiberoptic cable:

Fiber Core Diameter	Fiber Bandwidth	Maximum Link Distance
62.5um	160 MHz*Km	220 Meters
62.5um	200 MHz*Km	275 Meters
60um	400 MHz*Km	500 Meters
50um	500 MHz*Km	550 Meters

Protect Your Network Perimeter

Figure 1-6 shows a Unit deployed on the perimeter of the network that uses a firewall.

Figure 1-6: Protecting the Network Perimeter

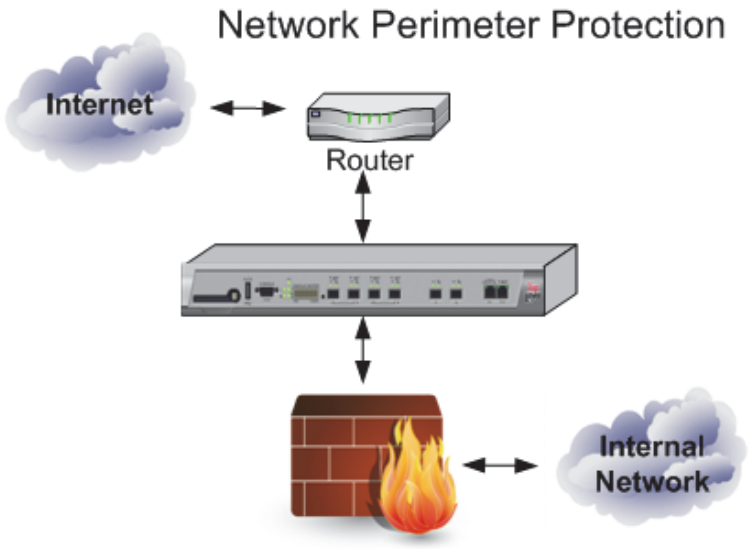

Protect Your Hosting Center

Figure 1-7 shows a Unit deployed as protection for a hosting center.

Figure 1-7: Protecting a Hosting Center

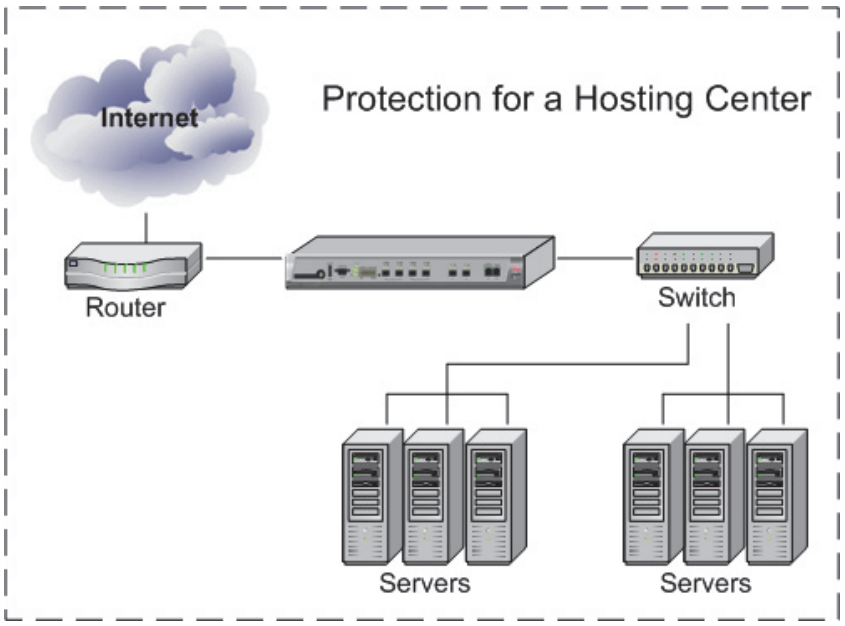

Protect Servers in Your Enterprise

Figure 1-8 shows a Unit deployed both before and after the firewall to protect servers within your enterprise.

Figure 1-8: Protecting Enterprise Servers

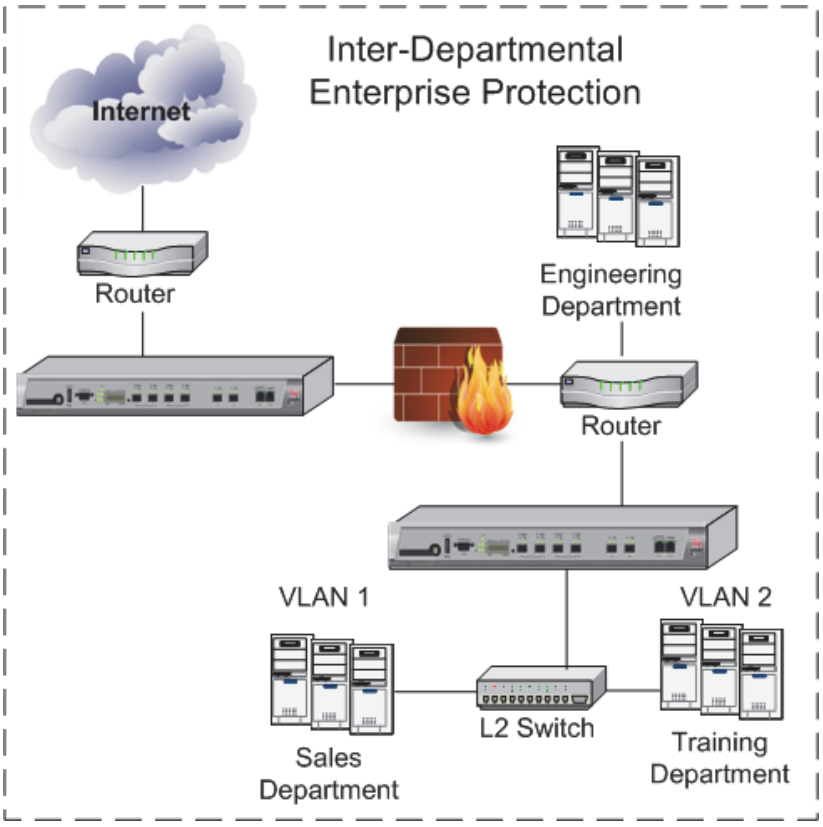

Chapter 2
IPS Overview

The IPS Unit's unique intrusion prevention architecture is the first high-performance inline security device that provides non-stop protection against both network level and application level cyber threats.

Leveraging its patented and award-winning DDoS defense technology, The Corero IPS Unit is specifically focused and highly optimized, designed to identify and deflect attacks while allowing genuine traffic to pass with minimum disruption. The IPS Unit provides maximum protection for critical IT assets while allowing full access to legitimate users and applications.

The IPS 5100 and IPS 5200 Series hardware platforms use advanced multi-core processor technology, with patented algorithms integrated with proven stateful analysis techniques, advanced deep packet inspection and industry-leading DoS (Denial of Service) attack protection to provide comprehensive protection from Internet-based and internal threats.

This chapter introduces the following IPS 5500 features:

Protocol Anomaly Detection

An Intrusion Prevention System must be able to determine whether the packets violate protocol standards, as this may be indicative of malware. In addition to determining whether the packets violate the standards, it must also be able to determine whether the data within the protocol adheres to expected usage. This expected usage could be industry-wide or at the enterprise level. For example, if peer-to-peer (P2P) applications were disallowed by an enterprise by policy, legitimate P2P traffic would traverse the firewall but should be blocked by the IPS. In contrast, a corporate policy may allow P2P, but disallow file sharing or other attachments. In this case the IPS must be able to identify any attachments associated with the protocol and strip out the attachments to be discarded. Corero products apply stateful protocol inspection, which enables it to make more intelligent decisions than intrusion prevention systems that rely primarily on signatures.

Data File Inspection

A significant proportion of attacks seen today results from malware contained in data that are used by applications, even though the transport protocol may adhere to the appropriate RFCs. For example, many attackers take advantage of vulnerabilities in Microsoft Office applications to launch their attack once the application runs the data with the embedded malware. For this reason, Corero products inspect the data files.

Acceptable Application Usage

It is important that an IPS can restrict what an application is able to process thereby preventing unauthorized operations. The ability to combine access control and approved usage checks on application layer traffic is important. For example, a web server is able to process far more commands than a typical user would use in practice. By only permitting traffic to the web server that utilizes the allowed commands you would eliminate complete classes of potential attacks. When applied by the IPS, this type of protection can be effective at blocking zero-day exploits.

Signature Matching

There are several techniques that have been created over the years for applying signatures to network traffic to determine whether the packets contain malware. The earliest and most simple version was referred to as simple pattern matching. A more efficient form of pattern matching referred to as regular expression defines complex search patterns that increase the accuracy of malware detection. In order to minimize latency, a significant amount of hardware acceleration has been built in to the IPS device.

Real-time Shunning

The Corero IPS has an effective protection capability called shunning that can quickly block traffic from IP addresses, temporarily or permanently, that are suspected of originating or being related to an attack. The advanced protection capabilities from shunning can be summarized as follows:

- Attack Source Identification - The Security Event Viewer enables users to identify a set of attacker IP addresses associated with blocked and detected attacks.
- Malicious IP Address Shunning - isolate events of interest and automatically shun all IP addresses associated with a particular attack event. Users can set time periods for how long each address should be shunned, as well as manually unshun addresses that are determined safe.
- Attack Defense Dashboards – The user interface allows Security Operations Center personnel to switch between routine monitoring and incident response.
- Additional Router Protection - Administrators can export a list of IP addresses being shunned

Robust Protection

Corero's purpose-built Tilera multicore processor architecture features Gigabit speed TopInspect™ deep packet inspection algorithms. The IPS Unit's robust High Availability (HA) configurations, high-MTBF hardware design, redundant capabilities, hot-swappable power supplies and swappable fan-tray, secure custom operating system, and flexible port-bypass capabilities provide non-stop reliability. The IPS 5500 family consists of products with the performance and capacity to handle throughputs from 100Mbit/sec to 8Gbit/sec, with transaction rates up to 40,000 sessions/sec.

Attack Mitigation

The IPS 5500 provides stateful matching of attack signatures for IP, UDP, and reassembled TCP session payloads. In addition to the factory provided signatures, you can add and edit your own signatures.

The IPS 5500 provides acceptable application use policies, including:

- Deep packet inspection for HTTP, FTP, DNS, SMTP, Telnet, SSH, MS-RPC, MS-CIFS, and other application protocols.
- Critical vulnerability protection against injection attacks, access attacks, DDoS attacks, unauthorized servers, back doors, and the like.
- Transaction and data protection rules for application-level checking of HTTP, FTP, DNS, SMTP, Telnet, SSH, MS-RPC, MS-CIFS, and other application protocols.
- Configurable data validation modules that inspect the content and format of known and unknown file types when carried as payloads of supported L3, L4, and L5 protocols.

The IPS 5500 provides protocol and file validation, including:

- Configurable transport layer protection rules for TCP and UDP including flexible enforcement criteria.
- Protocol normalization for reordering and coalescing IP fragments, and reordering TCP segments.
- Configurable file-format protection rules for files carried in protocol payloads.
- File format usage policies

The IPS 5500 provides stateful firewall filtering, including:

- Policy-based undesired access protection through stateful firewall filtering with no performance degradation
- Configurable data link protection against illegal or ill-formed MAC and data link headers, IEEE 802.1Q VLAN filters, and MAC address filters.
- Configurable protection against attempts to use TCP retransmissions and segment overlap as evasion mechanisms.
- Configurable network protocol protection rules for IPv4, ICMP header fields, IP address filters

The IPS 5500 ensures the availability of applications and services, even when under botnet-initiated attacks. This protection includes:

- Denial of Service & DDoS Protection:
  Patented algorithms for protection against SYN floods, ICMP floods, UDP floods, and application overload attacks.
- Policy-Based Rate Limits:
  Policy based rules that limit traffic rates.
- Connection Limits:
  Configurable rules that protect your network resources (such as servers and routers) from being overwhelmed by too many active connections

- Client Request Limits
  Configurable rules that limit the rate at which individual clients or groups of clients can initiate transactions.

Deep Packet Inspection

The IPS Unit is installed inline as a Layer 2 network forwarding element. It inspects all traffic to prevent undesired access, filters illegal packets and illegal headers, stops network attacks and DoS attacks, prevents exploits of critical vulnerabilities, mitigates service overload attacks, and thwarts application level attacks.

In addition to its stateful analysis, firewall, and anti-DDoS features, the IPS Unit protects critical online assets with the TopInspect deep packet inspection technology by:

- Focusing on protecting against critical remotely exploitable vulnerabilities.
- Deep, thorough analysis of network and application transactions to prevent harmful and/or malicious activity.
- Protocol Validation Module (PVM) architecture verifies that the protocol in use is the one expected and that it is being used correctly. In addition, PVMs validate protocol rules and check for known vulnerabilities.
- Advanced RFC-validation to protect against zero-day and short-notification-window exploits.
- Data Validation Module (DVM) architecture focuses on current and future vulnerabilities carried in attachments to HTTP and Email traffic

High Availability: The ProtectionCluster™

The IPS Unit is designed with High Availability (HA) in mind. High-MTBF hardware design with no rotating media, redundant hot-swappable power supplies, and a hot-swappable N+1 fan tray ensures non-stop operation. Port bypass on all internal and external network ports ensures network availability even in the unlikely event of an internal failure.

Multiple IPS Units can be combined into a ProtectionCluster, offering protection from a single-point of failure, and supporting continuous state sharing between devices to ensure continued network operation, even in the event of a fail-over.

Each IPS Unit has either two (5200 Series) or four (5100 Series) dedicated Gigabit-speed HA interfaces to allow automatic traffic balancing and continued communication with other devices in the ProtectionCluster, even if one of the HA links is down. Corero's experience in deploying in-line intrusion prevention has led to IPS Unit security solutions that network managers readily accept in their networks.

For more information about ProtectionClusters, see Chapter 10, "ProtectionCluster Configuration".

NOTE

HA ports are not available on some older models of the IPS 5500 product.

Chapter 3
Getting Started with the IPS Management Application

The IPS Unit management application is a Java Web Start™ Graphical User Interface (GUI) application that runs as a stand alone application. The GUI is the primary management application for IPS Unit configuration and monitoring.

This chapter introduces you to the management application, its main features, and how to use it.

For other management methods, refer to Management Services (page 8-3).

This chapter contains the following information:

Accessing the Management Application

To launch the management application:

1. The IPS Unit's GUI is a Java Web Start™ application and requires that you have the proper version of the Java Runtime Environment installed. Refer to the IPS 5500 Release Notes for JRE version and availability.
2. To display the IPS Unit's Welcome window, point your browser at the IP address you assigned to the IPS Unit during the Setup procedure.
3. From the Welcome window, select the Graphical User Interface. If this is the first time you have accessed this application, the IPS Unit downloads the Java Web Start Application to your computer.
4. The login window displays. Enter admin, with no password, as the initial login.

CAUTION

For security purposes, be sure to change this login and password as described in Configuring Management Port Access (page 8-6).

5. If you are the first user to launch a version of the software that is accompanied by an updated End User License Agreement (EULA), you will be required to accept the EULA before you can access the management application. For more information on viewing EULA text, see Viewing "About..." Information (page 18-5).

NOTE

If you are unable to change your IPS Unit from Bypass to Inline mode, this indicates your trial system license has expired. Contact Corero for assistance.

6. The main window displays (Figure 3-1).

Figure 3-1: Management Application

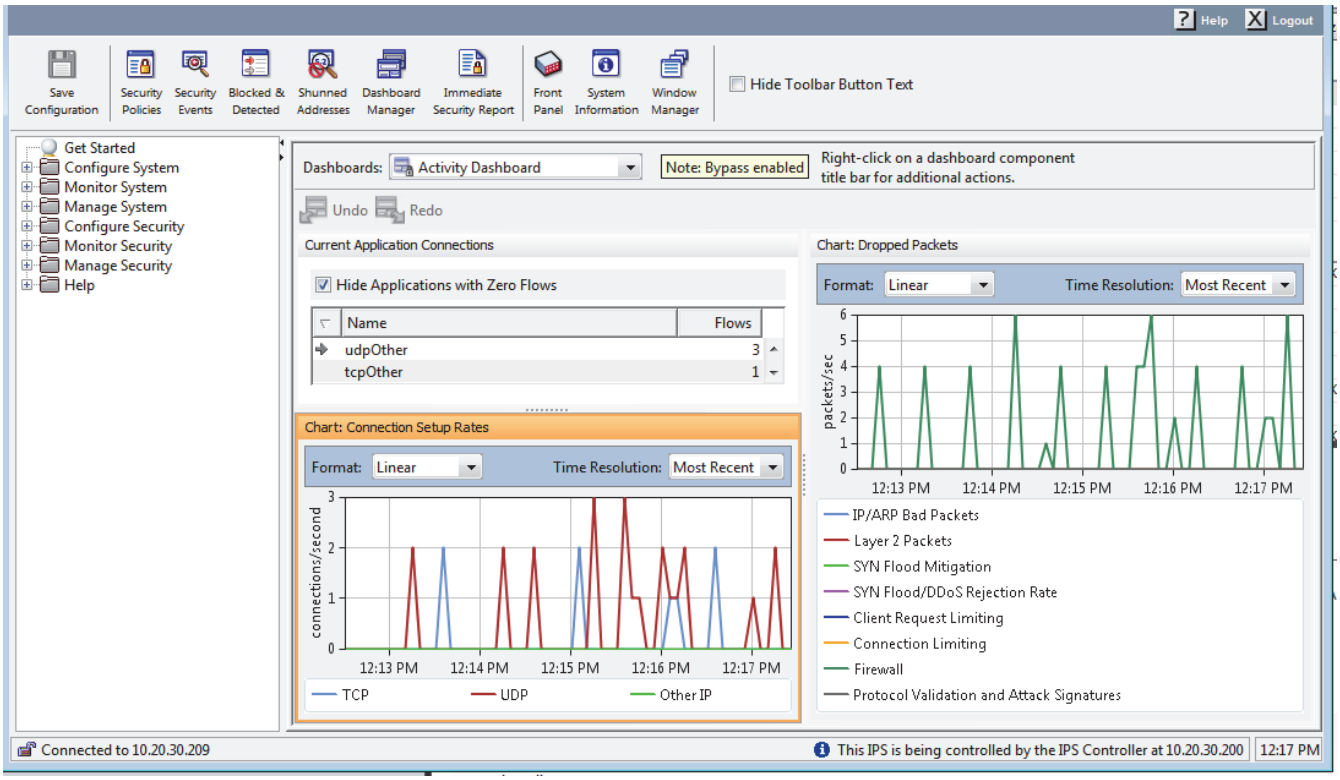

The main window of your Corero Network Device management application contains the following areas:

- The top of the screen provides access to main features through toolbar buttons. For a description of these buttons, see Using the Toolbar Buttons (page 3-4).
- The left side of the screen provides access to all features through a navigation tree. For a description of navigation tree options, see Using the Navigation Tree (page 3-5).
- The bottom of the screen displays a status bar. For more information on the status bar, see Using the Status Bar (page 3-6).
- The remainder of the screen provides the work area, where you can view, configure, and manage network traffic and policy information.

Using the Toolbar Buttons

Located at the top of the main window, these buttons provide quick access to the commonly used functions described in Table 3-1.

Table 3-1: Toolbar Buttons

Icon	Label	Corresponding Navigation Tree Option	Description	For More Information, See...
	Save Configuration	(None)	Save policy and system configuration change.	Saving Changes (page 3-10)
	Security Policies	Configure Security > Security Policies *and* Manage Security > Security Policies	Display the Firewall + Intrusion Protection System policy configuration window.	Chapter 11, "About Security Policies" Chapter 12, "Managing FW+IPS Security Policies"
	Security Events	Monitor Security > Security Event Viewer	Display the Security Event Viewer to examine and manage network security events.	About the Security Event Viewer (page 19-19)
	Blocked & Detected	Monitor Security > Blocked and Detected Attacks	Display the blocked and detected window to examine network security events.	Viewing Blocked and Detected Attacks (page 19-16)
	Shunned Address	Monitor Security > Shunned Address Viewer	Display the Shunned Address Viewer to view, add, modify, and delete Shunned Address information.	About Using IP Address Shunning to Stop an Attack (page 19-4)
	Dashboard Manager	Monitor Security > Dashboards > Dashboard Manager	Manage the GUI dashboards.	Using the Dashboard Displays (page 3-8)
	Immediate Security Report	Monitor Security > Immediate Report	Run an Immediate security report.	Generating an Immediate Security Report (page 16-10)
	Front Panel	Monitor System > Front Panel	View the interactive Front Panel view of the IPS Unit's configuration.	Using the Front Panel View (page 18-2)
	System Information	Monitor System > System Info	View the System Information window.	Viewing System Information (page 18-6)
	Window Manager	(None)	Display the Window Manager tool. For more information.	Managing Application Windows (page 3-10)

Using the Navigation Tree

The Navigation Tree, located at the left of the main window, provides access to IPS Unit port, network, and other system configuration and management windows as well as security configuration and management.

The Navigation Tree groups IPS Unit features into a series of top-level choices, which are described in Table 3-2.

Table 3-2: Navigation Tree Top-Level Options

Navigation Tree Option	Description
Get Started	Launches the Getting Started Wizard, which walks you through initial configuration of the IPS Unit.
Configure System	Provides access to features associated with initial configuration of the IPS Unit, including ports, users, time settings, management access, and whether or not the IPS Unit is managed by an IPS Controller.
Monitor System	Provides access to features that enable you to monitor system operation, statistics, and events.
Manage System	Enables you to manage system operation including reboot, software upgrade, and resetting to factory defaults.
Configure Security	Provides access to all aspects of security policy configuration, as well as access to security logs and reports.
Monitor Security	Provides access to IPS Unit views that display system events, attack information, statistics, and graphs. You can also query an IP address from this selection.
Manage Security	Launches the Security Policies window.
Help	Provides access to information about the current software version Provides access to the End User License Agreement text. Enables you to launch the online help system, which provides detailed descriptions of configuration parameters, procedures, and notes regarding use of the IPS 5500 product features. Note that the IPS and DDS products share a common help system, and where the features in those products differ, the online help clearly indicates this.

Using the Status Bar

The gray status bar displays across the bottom of the main window.

- The information on the left side indicates the IP address of the Corero Network Device currently being managed.
- The right side displays the current time.
- If the Corero Network Device is currently being managed by an IPS Controller, this information will display.
- If the Corero Network Device is using a trial system license status bar will display trial license information including:
  - When the device is running in Trial mode (before the license expiration date).
  - When the license is approaching its expiration date.
  - When the trial license has expired (which locks the device in bypass mode until another license is entered).

NOTE

For detailed information about viewing license information and entering licenses, see System License Management Key (page A-2).

Using the Online Help

The Graphical User Interface (GUI) provides access to detailed conceptual and procedural configuration and operation information. You can access the help in two ways:

- If you want to view or find a topic on a subject of interest, do one of the following:
  - Click the Help button at the upper right corner of the main window.
  - Choose Help > Help Topics from the navigation tree.

  This will launch the online help system. From here you can peruse the table of contents and the index, or search for particular terms.
- If you want information on using the specific dialog box you are currently displaying, click the Help button on that dialog box. This will display context-sensitive help, showing information that pertains specifically to the dialog box whose help button you clicked.

Using the Dashboard Displays

The dashboard consists of one or more individual graphs and charts that are displayed simultaneously, providing snapshots of the IPS Unit's operation and attack mitigation results. The dashboard is located in the work area to the right of the navigation tree in the main window. You can choose from among several default dashboards for different views, or you can create your own.

The individual components in a dashboard can be moved around within that dashboard to best fit the user's screen space and desired look. Also, components can be overlapped with each other and accessed via tabs by dragging one component to the title bar of another component. A component that is tabbed can be un-tabbed by dragging the tab to a new location in the dashboard.

You can select a dashboard at the top of the work area. The dashboard you select dictates which graphs or charts display on the main window. There are several default dashboards for different views, or you can create your own using the Dashboard Manager, which you can access by clicking the Dashboard Manager toolbar button.

Available default dashboards are described in Table 3-3.

Table 3-3: Default Dashboards

Dashboard	Contents
Activity	Displays information on connection setup rates, dropped packets, and current application connections.
Chart	Displays information on connection setup rates, dropped packets, IP threat levels, and SYN flood statistics.
General Attack	Displays information on blocked and detected attacks, dropped packets, and IP address information, as well as displaying the security event viewer.
Health/Monitoring	Displays information on blocked and detected attacks, CPU activity, dropped packets, and current application connections.

The charts and components you can view in default dashboards, and use to make a custom dashboard, are described in Table 3-4.

Table 3-4: Available Dashboard Components

Chart	Description
Blocked and Detected Attacks	The Blocked and Detected Attacks window dynamically displays information about current attacks.
Dropped Packet Statistics	Lists the number of packets dropped in each of the following categories: received unicast, received data link errors, transmit unicast, transmit data link errors, malformed, layer 2 bridge filtered, firewall blocked, IPS blocked, connection limited, application priority, link outbound congestion, load shedding, DDoS rejection, FTP load shedding, fragment limiting, malformed fragments, malformed TCP segments, ICMP rate limiting, and client rate limiting.
Chart: Connection Setup Rates	Displays the current rate of connection setups per second for the TCP, UDP, or other IP connection types.
Chart: Connection Usage	Displays graphs representing the traffic going through Corero Network Devices including TCP, UDP, Other IP, and Aged connections.
IP Address Query	Displays device, group, cluster, threat, SYN, and connection information for a specified IP address.

Table 3-4: Available Dashboard Components *(Continued)*

Chart	Description
Chart: CPU Activity	Represents the percentage of CPU activity in the following categories: utilization, maintenance, TCP setup, UDP setup, and IP connection.
Chart: Dropped Packets	Indicates the number of packets dropped due to IP/ARP bad packets, layer-2 filtered packets, SYN flood mitigation, SYN flood / DDoS rejection, client request limiting, connection limiting, firewall, and protocol validation and attack signatures.
Security Event Viewer	Enables you to easily examine and react to the traffic that triggers a Corero Network Device's security rules. Using the viewer you can examine details for every event triggered by a rule.
Chart: IP Threat Levels	The number of addresses that fall into each of the available address threat levels. Threat levels include unknown, trusted, suspicious, malicious, and DDoS rejection IP address categories.
Chart: SYN Flood Statistics	Provides information on the rate at which malicious SYN flood packets are handled based on the packets per second dropped for the following packet types: malicious SYN packets, SYN flood / DDoS rejection, client proxy fail, server proxy fail, proxy resource drop.
View Rule	Displays the name, description, and confidence category for the selected rule.
Current Application Connections	Provides the number of current connections for a network service, listed by network protocol/port combination.
Front Panel View	Displays an interactive graphical view of the physical front panel of the selected Corero Network Device (IPS or DDS Unit). This view shows the device's ports, and, using special icons, shows the roles currently assigned to the ports and other port-related feature settings.
Port Statistics	The Port Statistics table reports information for the transmitted and received packets for each Corero Network Device port.
Shunned Address Viewer	The Shunned Address Viewer is a security monitoring and management tool that enables you to easily examine and update the IP addresses that are currently being shunned by a Corero Network Device
Chart: Custom Chart	This graph enables you to select from a number of statistics over a user-specified period of time. For each statistic you select, the system provides default settings for the low, medium, and high thresholds for display. When you select a statistic, you can modify these thresholds as needed.

Helpful Hints

The Graphical User Interface (GUI) enables you to easily add and manage elements required by your IPS Unit's system and security functions. Typically, you access needed configuration or management functions from the Navigation Tree.

Adding Items

When you define an element using the Add window, you can choose to add a single element or multiple elements as follows:

- If you want to add a single element, once you have entered the desired information, click the Done button.
- If you want to add more than one element:
  a. Once you have entered the desired information, click Add. This will save the element you added, and the Add dialog box will display with empty fields.
  b. Add as many elements as you wish in this fashion.
  c. When you have finished entering information for the last element, click Done.

Choosing Multiple Items

When selecting rows in a table, you may use standard Windows selection keys to select multiple elements (Ctrl+mouse click or arrow) or a range of elements (Shift + mouse click or arrow).

If you want to select all items in the currently selected table, click Ctrl+A.

Managing Application Windows

The Graphical User Interface includes a Window Manager (shown in Figure 17) which enables you to view a list of the currently open management windows (including graphs), and switch to a given window.

1. To access the window manager, double click the icon in the top right section of the main window, or press the F11 function key.
2. To jump to a particular window, select the window from the list, then click Show.

Saving Changes

When you modify IPS Unit settings, the toolbar indicates when you have specified changes, but have not yet saved them permanently to the flash drive. If you save your changes using the Save Configuration toolbar button, your changes will be preserved across IPS Unit restarts.

To save your changes, click the "Save Configuration" Toolbar button ().

Note that configuring security policies differs from other configuration changes in that you can choose either to Undo your changes or Apply them before you Save them.

Reserved Words

The following words are used in special ways by the IPS management application. They cannot be used, in any form, to begin the names of host groups, IP address ranges, port pairs, or other items.

- Any
- IP
- TCP

- UCP
- ICMP
- Inbound
- Outbound
- Internal
- External

Chapter 4
Initial IPS Unit Configuration Tasks

After you have completed installation and cabling of the IPS Unit hardware, you must perform several initial configuration tasks prior to initial use.

This chapter describes the following initial configuration tasks:

When you are configuring your Corero Network Device, in this case an IPS Unit, there are certain settings that apply to system operation of the unit itself. These settings involve system time and system log servers.

NOTE

For information on Corero Network Device management access related features, refer to Chapter 8, ‘‘Management Access”.

Managing Syslog Servers

Your Corero Network Device is designed to log device-specific information. The Event Logging System (ELS) stores this information in a log file, but due to space limitations, Corero strongly recommends you use Syslog servers to ensure capacity for proper event storage. You can also configure the event logging system to send log information to a Syslog server that you specify.

To view, add, or modify Syslog Servers:

1. From the Navigation Tree, choose Configure System > Syslog Servers. The Syslog Servers dialog box displays.

   The Syslog Servers dialog box displays the following information:

   - IP Address of the Syslog server.
   - The UDP Port, which is, by default, the well known UDP port 514.
   - The Mode, specifying whether messages to the Syslog server are enabled or disabled.

2. To Add a server to which the Event Logging System (ELS) should send messages, click Add. The Add Syslog Server dialog box displays. Enter the parameters listed above.

   When you have finished, do one of the following:

   - If you only wish to add a single server, click Done. The dialog box closes.
   - If you want to add additional servers, click Add. The dialog box remains, so you can add another server. When finished, click Done.

3. To modify a Syslog server, select the server you wish to modify, then click Edit. The Edit Syslog Server dialog box displays. You can only enable and disable messages from the event logging system to the Syslog Server.

   Select the desired mode, then click OK.

4. To delete a Syslog server, select the desired server, then click Delete. A confirmation dialog box displays, asked you to verify your choice. To proceed, click Yes.
5. Save your changes by clicking the Save Configuration toolbar button.

Managing Audit Logs

Your Corero Network Device has an audit function which logs every change to the unit's configuration. These log items are kept in a log file, and, if Syslog server(s) have been setup, can also be configured to send them to the Syslog server(s).

All configuration changes, save operations, boot ups, and failed and successful authentications are logged. Audit logging is disabled by default.

NOTE

Enabling the audit function may cause some slowdown in the use of the management application.

Audit messages are stored in the audit log file in the following format:

Table 4-1: Audit Log Information

Date	Time	Unit_IP	Unit_Model	ID	Device	User	Audit_Fields
Aug 31	10:52:41	10.25.36.102	IPS5500-1000EC	id=rr-nn	pt=TLN-TQ	user=peterz	xx= xy= yz=

Where:

- ID is a unique number identifying the audit entry.
  rr is the number of times that the Corero Network Device has been rebooted.
  nn is a sequentially increasing number since the last reboot.
- User is the name of the user logged in performing the action, or "Not Known" if it cannot be identified (such as when the unit is powered on - this event is audited but there is no user logged in)
- Audit_fields contain different values depending upon the operation being audited.

NOTE

Audit fields contain different values depending on the operation being audited.

Audit messages are kept in an audit log file stored on the Corero Network Device's compact flash. Up to 10 audit files are maintained, the oldest being deleted to make available space for a new one when required. Audit log files can be viewed by the System Log Viewer.

To enable audit logging on your Corero Network Device:

1. Select Configure System > Advanced System Config > Audit Logging. The Audit Logging dialog box displays.
2. Select the Enable Audit logging check box.
3. To additionally send audit log messages to Syslog servers, once you have enabled audit logging, select the Send Audit Data to Syslog Server(s) check box.

   When you select this option, in addition to being stored in the audit log file, data is also sent to the local Syslog server. These messages are assigned a facility of 13, and a priority/severity of 6 (information).

NOTE

For information on how to configure Syslog servers, see Managing Syslog Servers (page 4-2).

4. To view audit log information, choose Monitor System > System Log Viewer. The View Log File dialog box displays.
5. In the Log Type drop-down, choose Audit Log.

   You can select a numbered Audit Log file from the drop-down list, or you can select the Current File to view the most recent information.

   If you choose to view contents of the Current File, you can click Refresh to display any more current information, if available.

6. In the File Name drop-down, select the desired log file, then click OK.

   The View Audit Log File dialog box displays, with the most recent data displayed first.

   You can scroll through the audit log data, or search for specific terms. You can also sort the data based on a column's contents by clicking the column's heading.

Managing Network Time Protocol (NTP) Servers

You can use Network Time Protocol (NTP), which is documented in RFC 1305, to automatically update time settings for the Corero Network Device. NTP synchronizes time among distributed time servers and clients. Synchronization enables time-specific events, such as system logs, to be correlated. All NTP servers and clients use Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT). If the device loses NTP sync, this information will be logged to the system log file and Syslog server(s) (if configured).

With NTP enabled, the Corero Network Device determines the system time by receiving NTP broadcast or multicast messages or by querying an NTP server at the time interval you configure. The device then chooses the NTP server with the lowest stratum number (as defined by the NTP algorithm) and updates the system clock. The stratum number describes how many NTP hops away the device is from an authoritative time source, with stratum 1 being the time source itself.

For example, a stratum number of 1 means a radio or atomic clock is directly attached. A stratum number of 2 means the Corero Network Device receives its time through NTP from a stratum 1 time server, and so on.The device supports connection to an NTP server with a stratum number of two or lower.

An NTP server automatically chooses, as its time source, the machine with the lowest stratum number.

CAUTION

The Corero Network Device does not accept a response from an NTP server with a time that is more than 45 minutes away from the device's current time. After three such responses, the query is no longer sent to that NTP server. To avoid this, use the Time Settings window to set the device's time as close to the real time as possible.

NOTE

Because the Corero Network Device's time is updated in small increments, for NTP updating to work quickly, you should manually configure the device's current time to be as close as possible to your local time. For information on setting the current time, see Configuring the Current Time (page 4-7).

1. From the Navigation Tree, select Configure System > Time > Network Time Protocol. The Network Time Protocol dialog box displays (Figure 4-1).

Figure 4-1: Network Time Protocol Dialog Box

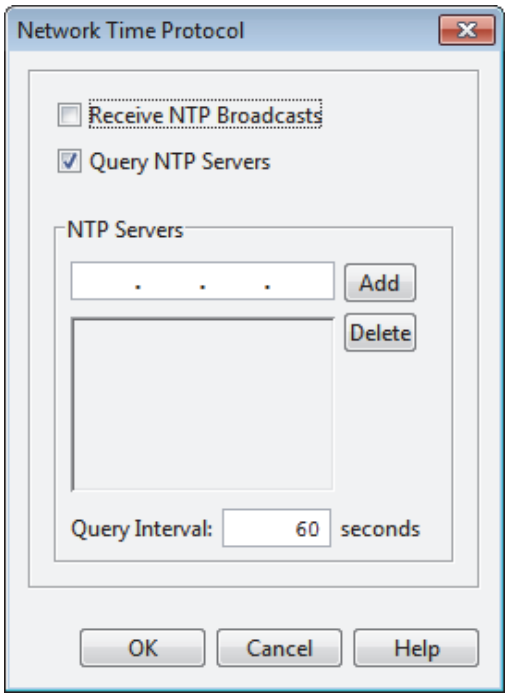

The Network Time Protocol window displays a list of current NTP servers.

2. To add an NTP Server:
   a. Specify whether or not the Corero Network Device should Receive NTP Broadcasts. When this check box is selected, the device is configured to accept time update broadcast messages from the NTP servers you identify.
   b. Specify whether or not the device should query this NTP Server for an updated time.
   c. Of you have selected Query NTP Servers, you can specify the Query Interval, which indicates how frequently the device requests the current time from the NTP server.
3. To delete an NTP server, select the server in the list and click Delete. A dialog box displays asking you to confirm your selection.
4. Save your changes by clicking the Save Configuration toolbar button.

Configuring the Current Time

If needed, you can configure the current date and time on the internal clock of the Corero Network Device. This is typically only done after powering up the system.

NOTE

If you intend to use an NTP server to update the Corero Network Device's clock, be sure to set the current time as accurately as possible. When system time synchronizes with an NTP server, it is gradually changed until the two are the same. Since the time is corrected gradually, if you enter an incorrect time, it will take longer for the NTP server's input to correct the time.

To configure the current time:

1. From the Navigation Tree, select Configure System > Time > Current Time. The Current Time window displays.
2. You can modify the following information:
   - The local date, month, and four-digit year (YYYY).
   - The local time, specified using a 24-hour format (HH:MM:SS).
3. When you have finished, click OK.
4. Save your changes by clicking the Save Configuration toolbar button.

Configuring the Time Zone

You can configure the time zone used by the internal clock of the Corero Network Device.

To configure the time zone:

1. From the Navigation Tree, choose Configure System > Time > Time Zone. The Time Zone dialog box displays.
2. Select the local time zone.
3. To automatically adjust the clock for daylight savings time (DST) changes, click the check box.

NOTE

If you choose not to enable this feature, and you are operating in a location that follows daylight savings time, you will need to manually change the time on the device during the transition to and from DST.

4. When you have finished, click OK.
5. Save your changes by clicking the Save Configuration toolbar button.

About The Getting Started Wizard

The Getting Started Wizard is the starting point for configuring the IPS Unit. The wizard makes it easy to establish port roles for your IPS Unit based on your network configuration. You define the basic port configuration based on your network, and the wizard chooses and configures the actual ports along with their roles and settings.

From the wizard you can establish settings for the following port characteristics:

- Basic port operation modes such as whether the IPS Unit treats ports as port pairs (two ports linked together as input and output ports) and how many ports of each type you need.
- Special port features such as port tracking.
- Whether the IPS Unit should initially monitor traffic but send all traffic through without any mitigation (Always Bypass — default setting), or act in some other manner (Security Bypass mode).
- Configuration for mirror, management, discard, and capture ports.

NOTES

1. For a discussion of the concepts and terminology used with the IPS Unit's ports, including port pairs and special port roles, refer to "Chapter 5, ''Understanding Ports".
2. For a detailed discussion of the Getting Started Wizard, see Running the Getting Started Wizard (page 4-10).

Traffic Bypass Considerations

Before you run the Getting Started Wizard, you should consider how you want traffic to flow through the IPS Unit when you initially turn it on. Although the factory default security setting for the IPS Unit suffices for many networks, you may want to initially see what types of traffic are being reported as suspect before actually allowing the IPS Unit to process and mitigate traffic. For a detailed discussion of the Bypass Settings feature, see Selecting the Bypass Settings Mode (page 6-8).

On initial configuration, you can set the Bypass Settings feature to Always Bypass, which is the default setting. When the IPS Unit is set to Always Bypass, the IPS Unit examines and reports on traffic issues, but passes all the traffic it receives, regardless of its findings.

There are two ways to access the Bypass Settings feature:

- By running the Getting Started Wizard in the Navigation Tree.
- By selecting the Configure System > Bypass Settings option in the Navigation Tree.

NOTE

If a Corero Network Device is using a trial (time-limited) System License, when the license expires, the system will automatically enter Bypass mode. When locked in Bypass mode, you will be able to modify the bypass settings, but your changes will not take affect until you acquire a new system license key from Corero and apply it to the device. Contact Corero for assistance.

Running the Getting Started Wizard

To run the Getting Started Wizard:

1. Launch the Getting Started Wizard by clicking Get Started, which is the first option in the Navigation Tree.
2. The introductory window explains the purpose of the wizard. click Next.

   NOTE

   You can navigate forward and backward through the wizard at any time using the Next and the Back buttons.

3. The Mission Port Pairs page displays.

   A mission port pair consists of two ports that only send traffic through one another. One is an internal port, and one is an external port. Mission ports always operate in port-pair forwarding mode. Bridge forwarding is not supported for Mission ports. For more information on port roles, see Port Role Overview (page 5-2).

   The table at the bottom of the window displays the number of currently selected ports and their current roles. It updates dynamically as you make your selections.

   Specify how many Mission Port pairs you want to configure on this IPS Unit, then click Next.

   NOTE

   After you have finished configuring Mission Port pairs, it is helpful to give each pair a meaningful name. For more information on naming port pairs, see Viewing and Naming Port Pairs (page 6-7)

4. The Mission Port Pair Settings page displays.

   Use this page to select Port Tracking and to specify the Bypass Settings state for all Mission Port pairs.

   Port Tracking tracks the port link state for each Mission Port pair. For more information on Port Tracking, see Port Tracking (page 5-8).

   Bypass Settings specify the hardware-based bypass feature between Mission port-pairs. The term bypass indicates that traffic is not being mitigated by the IPS Unit. You can choose whether traffic never bypasses the IPS Unit, whether traffic always bypasses the IPS Unit, or whether traffic only bypasses the IPS Unit when the system is down or being reset. For more information on bypass settings, see Bypass Settings (page 5-6).

   Specify your Mission Port Pair settings, then click Next.

5. The Maintenance Ports page displays.

   Use this page to specify additional Management, Mirror, Discard (forensic), or Capture ports. The types of ports that are available for configuration vary depending on the IPS Unit model. For more information on port roles, . For information on specifying a port for capture, see Modifying Traffic Capture Settings (page 6-9)

   Specify the number of Management ports you want to configure for this IPS Unit. Specify which other port roles you want to implement. Then click Next.

6. The Summary window displays a summary of your specified settings. Review your selections, then do one of the following:
   - If you want to change your configuration settings, click Back.
   - To implement your configuration settings, click Finish.

- If you do not want to make your configuration changes, click Cancel.

NOTE

If you have changed the port configuration, ensure you reconfigure the physical cable connections to match the port configuration you specified.

Additional System Configuration Tasks

Table 4-2 provides references to other sections of this guide to enable you to perform additional configuration tasks with your IPS Unit.

Table 4-2: Additional System Configuration Tasks

To accomplish this task....	Refer to...
Understanding and modifying port roles	Chapter 5, "Understanding Ports" Chapter 6, "Viewing and Configuring Ports"
Port Tracking feature	Port Tracking (page 5-8)
Capture Port	Modifying Traffic Capture Settings (page 6-9)
Dynamic view of port role configuration	Using the Front Panel View (page 18-2)
User groups, user accounts, and global user security settings	Chapter 7, "Managing Users"
Management access controls	Chapter 8, "Management Access"
Advanced port topics, including VLAN support and static MAC addresses	Chapter 9, "Advanced Port Configuration"
IPS 5500 system-related tasks such as system reboot, reset to factory defaults, configuration files, and software upgrades	Appendix A, "IPS Unit System Management"

Chapter 5
Understanding Ports

Most of the ports on Corero Network Devices, such as IPS Units and DDS Units, can have several possible roles. Each port offers specific capabilities and can operate in one of several roles at any given time, and some ports have fixed roles. This chapter describes port roles, and which roles apply to specific ports on the device.

For information on viewing and configuring port settings, see Chapter 6, ‘‘Viewing and Configuring Ports”.

This chapter contains the following sections:

Port Role Overview

The ports on a Corero Network Device can be configured for different port roles. Each port role provides specific capabilities and operations for a given port. Note that some ports on the device are assigned permanent roles during manufacturing, while other ports can be assigned one of several roles at the customer site. Note that a port can only be configured for one role at any given time.

Setting Port Roles

To set the role of a specific port, use the Getting Started wizard available from the Graphical User Interface. The wizard queries the user to specify the number of ports needed for a specific role. The wizard then offers the appropriate ports for these roles. Using this wizard allows the Corero Network Device to enforce a port role's requirements and to locate the ports and group them according to their roles (for example, specifying adjacent ports as a bypass port pair). Note that you can run the Getting Started wizard at any time.

Mission, Management, and Maintenance Ports

In addition to port roles, Corero Network Devices support the concept of Mission, Management, and Maintenance ports to further classify ports based on their role type.

- Mission ports - These ports process internal and external network traffic. Two matched ports, one that handles internal traffic, and one that handles external traffic, are known as a Mission port-pair. The Corero Network Device supports a minimum of one, and a maximum of four, Mission port-pairs on any one device. Ports specified as a Mission port-pair are adjacent to one another on the device. A Mission port-pair handles traffic using port-pair forwarding, where external traffic always enters and leaves through the external port, and internal traffic always enters and leaves through the internal port. The device keeps track of the link state for Mission port-pairs using its LAN port tracking feature. For more information, see Port Tracking (page 5-8).
- Management ports - These ports are used to manage the Corero Network Device itself. Management traffic is completely isolated from Mission traffic.
- Maintenance ports - These ports on the Corero Network Device are used to manage events and mirror traffic on the device. These ports can have the role of Capture, Mirror, or Discard. For more information, see Table 5-1.
  - On 5000 series hardware, maintenance ports are Fast Ethernet ports, except for Capture ports which are either Gig ports or Fast Ethernet ports.
  - On 5100 and 5200 series hardware, maintenance ports are 10/100/1000 Gigabit ports.
  - On 5200 Series hardware, port #7 (M1) is the only port used for mirror or discard. No other ports are available for these functions.

The predefined and available port roles for the ports on a specific Corero Network Device model varies between product models. For more information on product-specific port role information, see the product-specific sections at the end of this chapter.

Port Role Types

Table 5-1 lists all of the port role types available for Corero Network Devices.

Table 5-1: Port Role Types

Role	Description
Capture	Use a Capture port as a single, port-based, mirroring output port. You can specify that one of the Mission ports has all of its received and transmitted packets sent to this port. For configuration information, see Modifying Traffic Capture Settings (page 6-9).
Discard	You can specify that a Corero Network Device send blocked and discarded traffic to this port. You can configure policies to specify which packets go to the Discard port. This port is typically connected to an analysis tool.
External (Outside)	An External port is used to connect to the external network, such as a network outside your corporation or organization. The External port does not allow management access. The External port receives packets and forwards them (subject to policy checks) to its paired Internal port. You can specify an External and an Internal port to carry traffic using port-pair forwarding mode.
High Availability (HA)	High Availability (HA) ports are directly connected to a redundant Corero Network Device. The HA port is used to balance traffic between redundant devices and guarantee that all packets of a given flow go through the same device. Note: High Availability (HA) ports are not available on Model 75EC IPS Units. On the 5100 Model Corero Network Devices, ports 5-8 are dedicated HA ports when HA is enabled. Corero recommends that you use all four HA ports on this model for maximum throughput and performance.
HA Interconnect Switch	A Switch port enables you to connect multiple Corero Network Devices in a ProtectionCluster. When you connect more than two devices in a ProtectionCluster, you must use a switch to connect the HA links. Ports S1 through S4 on the Model 2000 ESL are dedicated to this interconnection, so this model can be used in place of a switch in HA configurations.
Internal (Inside)	An Internal port is used to connect to an internal network. The Internal port does not allow management access. The Internal port receives packets and forwards them (subject to policy checks) to its paired External port. You can specify an External and an Internal port to carry traffic using port-pair forwarding mode.
Management	A Management port enables you to manage the Corero Network Device. It can also be used as an output port for reporting traffic (using standard Syslog and SNMP traps).
Mirror	Identify one or more Mirror ports to create a mirror (copy) group. When you specify a Mirror port, the Corero Network Device copies all traffic that meets the conditions of a particular policy entry to the Mirror port(s). If there is more than one Mirror Port, the device uses a Round Robin algorithm to balance traffic among the Mirror ports. Note that all ports in the mirror group must be set to the same speed.
Unused	An Unused port is a port that is not configured with another role. The Unused port does not accept traffic nor does it send any traffic. The Corero Network Device will not recognize a link to this port.

Port Role Features

Table 5-2 displays a summary of port role features.

Table 5-2: Port Role Features

Port Role	Accept Management Traffic?	Generate Management Traffic?	Mirror Traffic from the Port	Mirror Traffic to the Port	Receive Dropped Traffic or a Copy of Monitored Traffic
Capture	No	No	No	No	Yes
Discard	No	No	No	No	Yes
External	No	No	Yes	No	No
High Availability	No	No	No	No	No
Internal	No	No	Yes	No	No
Management	Yes Port 2616 is the IPS Controller Management Port.	Yes (Syslog, SNMP, Traps)	No	No	No
Mirror	No	No	No	Yes	Yes
Unused	No	No	No	No	No

Port Pair Forwarding

Mission ports *always* operate in port pair forwarding mode. In port pair forwarding, each individual External port is paired with a single Internal port. The Corero Network Device forwards all packets received on either of the ports in the pair to the other port in the pair, subject to the defined security policy filtering.

In port pair forwarding, the Corero Network Device:

- Does not perform any MAC address learning; therefore, there is not any natural bridge filtering (that is, the device does not drop any packets whose destination is the same as the incoming port).
- Does not perform any flooding to multiple ports for multicasts, broadcasts, or unknown destination addresses. This traffic is forwarded to the other port in the pair.
- Tracks the link state of each port in the pair (if you enable the Port Tracking feature). If either port has a down link state, the device takes both ports down. In addition, both ports must be connected and maintain an up link state before the device will begin sending traffic through the ports. This feature enables a cable/port failure to be propagated from one side of the device to the other, which enables outside redundancy mechanisms to detect the loss of the ports.

NOTE

A port pair is called a segment when configuring security policies. It is important that you provide meaningful names for your port pairs so you can easily identify the segments when you create security policies. For more information, refer to Port Role Overview (page 5-2).

Bypass Settings

Corero Network Devices provide a software-based bypass feature between Mission port-pairs. When the device is operating in bypass mode it is not mitigating the traffic passing through it in any way. Bypass settings affect all mission ports.

If there is a software failure, the device stops inspecting traffic and recording the results, but the device continues to pass all traffic, with the following exceptions:

- For 5100/5200 ES-series hardware the device only passes traffic when the unit is powered on.
- For 5100 EC-series hardware, the bypass setting is preserved when there is no power to the Corero Network Device. This way, if a 5100 series hardware unit is in bypass mode and power is lost, mission traffic on ports will still be passed via the hardware bypass capability. But management traffic is never passed on management ports when there is no power.

NOTE

If a Corero Network Device is using a trial (time-limited) System License, when the license expires, the system will automatically enter Bypass mode. When locked in Bypass mode, you will be able to modify the bypass settings, but your changes will not take affect until you acquire a new system license key from Corero and apply it to the device.

There are three modes of bypass control as described in Table 5-3. Choose the mode that reflects the combination of normal and failure operation that you want to occur.

Table 5-3: Bypass Control Modes

Mode	Normal Operation	Failure Behavior	Description
Never Bypass	Inspect, Mitigate, Record	Stop all Traffic	The Corero Network Device inspects all traffic, mitigates problem traffic, and records all information. All traffic flows through the device's functions. If the software fails, the device acts as an open wire and does not forward any traffic. This mode provides the most protection and ensures that, in case of failure, unchecked traffic will not pass. This mode is also known as Fail Close in firewall terminology, because the system closes the door to all traffic if a system failure occurs.
Always Bypass (default)	Inspect, Record, Never Mitigate	Pass all Traffic	The Corero Network Device inspects all traffic and records traffic statistics, but does not mitigate. All traffic always passes through the device. This mode is useful when you are testing the device. Traffic is never blocked, even if there is a software failure. You can examine the information produced by the device to see what traffic would have been blocked if the unit were performing mitigation. This mode is also known as Fail Open in firewall terminology, because the system (opens the door and allows traffic if a failure occurs.

Table 5-3: Bypass Control Modes *(Continued)*

Mode	Normal Operation	Failure Behavior	Description
Bypass During System Reset	Inspect, Mitigate, Record	Pass all Traffic	The Corero Network Device initially operates in Never Bypass mode, checking and mitigating all traffic. The unit transitions to Always Bypass mode (passes all traffic) if there is a software failure and the device needs to reboot. Once the device resumes normal operation, it returns to full mitigation behavior. Bypass During System Reset mode is useful once you have completed testing the Corero Network Device and you want to mitigate traffic, but you want to pass unchecked traffic during a software failure rather than block unchecked traffic.

NOTE

In a ProtectionCluster environment where asymmetric network traffic is possible, all Corero Network Devices should be in either Always Bypass or Never Bypass mode. This is to ensure that when one device is rebooting, the other device(s) in the ProtectionCluster will not see partial flows. The exception to this is when all mission ports reside on a single device in the ProtectionCluster. In this configuration the bypass mode can be safely set to Bypass During System Reset.

Port Tracking

In port pair forwarding, each external mission port is paired with a single internal mission port, creating a mission port pair. The Corero Network Device forwards all packets received on either of the ports in the pair to the other port in the pair, subject to the defined security policy filtering. The primary purpose of port tracking is to track the link state of mission port-pairs and ensure that both ports in a pair reflect any change in the link status of either port.

The Corero Network Device is an in-line device and is often deployed in a redundant network. Link state tracking between peer ports on the device is essential for those failover mechanisms that do not use health checks and, instead, rely on link state. To support failover operation in devices such as firewalls and routers, the device propagates the end-to-end link state.

If Port Tracking is enabled on a Corero Network Device, and both ends of the link report different link status for two consecutive time periods (one second, by default), the device changes the link states so that they match (if one link was down, it takes the second link down also). When a failed link recovers, the device waits according to the Recover Wait Time, then reevaluates the status of both sides of the link and adjusts the link states to match the new condition.

Using the Getting Started wizard, you can enable or disable the Port Tracking feature for Mission port-pairs. For more information, see Port Role Overview (page 5-2).

Port Roles for 5100 Series Units

Port Roles on Corero Network Devices vary in their use, and their flexibility. Some port roles are preassigned during manufacturing, and cannot be modified. Some ports can be assigned one of several port roles. Note that a port can only be assigned one role at a time. In addition, each model has restrictions on how many of a given port role type can be implemented at any given time.

For a description of port roles, see Table 5-1.

Table 5-4 explains port roles for 5100 Series Corero Network Devices.

Table 5-4: 5100 Series Port Roles

Port Role	Number of Ports that Can Concurrently Share This Role	Operating Speed
Capture	0 or 1	10/100/1000
Discard	0 or 1	10/100/1000
External (connecting to an outside network)	1, 2, 3, or 4	10/100/1000
High Availability (HA) IPS 5100 75EC models do not support high availability.	0, 2, or 4	10/100/1000
Internal (connecting to an inside network)	1, 2, 3, or 4	10/100/1000
Management	1 or 2	10/100/1000
Mirror	0, 1, 2, or 3	10/100/1000
Unused These ports are not configured with a specified role	(Not Applicable)	10/100/1000

5100-Series Preconfigured and Configurable Port Role Assignments

Table 5-5 contains information the 5100-Series Corero Network Device's default port role configuration, as well as the possible roles available for each port. The ports are listed in the order in which they appear on the device (from left to right).

For information about the types of port roles, see Table 5-1.

Note that the default port roles are those supplied in manufacturing. These roles are in effect prior to running the Getting Started wizard.

Table 5-5: 5100-Series Port Role Assignments

Port Number	Speed	Peer Port Number In Bypass Mode	Possible Roles	Default Role 5500-150EC/ES 5500-500EC/ES 5500-1000EC/ES	Default Role 5500-75EC (Only 4 Ports, No HA Ports)
1	10/100/1000	2	External	External	External
2	10/100/1000	1	Internal	Internal	Internal
3	10/100/1000	4	External	External	External

Table 5-5: 5100-Series Port Role Assignments *(Continued)*

Port Number	Speed	Peer Port Number In Bypass Mode	Possible Roles	Default Role 5500-150EC/ES 5500-500EC/ES 5500-1000EC/ES	Default Role 5500-75EC (Only 4 Ports, No HA Ports)
4	10/100/1000	3	Internal	Internal	Internal
5	10/100/1000	6	External, HA	External	N/A
6	10/100/1000	5	External, HA	Internal	N/A
7	10/100/1000	8	External, HA	External	N/A
8	10/100/1000	7	External, HA	Internal	N/A
9	10/100/1000	N/A	Management, Mirror, Capture, Discard, Unused When requested, the device will automatically select a single capture or discard port from the ports available.	Unused	N/A
10	10/100/1000	N/A	Management, Mirror, Capture, Discard, Unused When requested, the device will automatically select a single capture or discard port from the ports available.	Unused	N/A
M1	10/100/1000	N/A	Management, Mirror, Capture, Discard, Unused When requested, the device will automatically select a single capture or discard port from the ports available.	Management	Management
M2	10/100/1000	N/A	Management	Management	Management

Port Roles for 5200 Series Model 2000ES Units

Port Roles on Corero Network Devices vary in their use, and their flexibility. Some port roles are preassigned during manufacturing, and cannot be modified. Some ports can be assigned one of several port roles. Note that a port can only be assigned one role at a time. In addition, each model has restrictions on how many of a given port role type can be implemented at any given time.

For a description of port roles, see Table 5-1.

Table 5-6 explains port roles for 5200 Series Model 2000 ES devices.

Table 5-6: 5200 Series Model 2000 ES Port Roles

Port Role	Number of Ports that can Concurrently Share This Role	Operating Speed
Capture	0	10/100/1000
Discard	0 or 1	10/100/1000
External (connecting to an outside network)	1 or 2	10 Gigabit Ethernet
High Availability (HA)	0 or 2	10 Gigabit Ethernet
Internal (connecting to an inside network)	1 or 2	10 Gigabit Ethernet
Management	1	10/100/1000
Mirror	0 or 1	10/100/1000
Unused These ports are not configured with a specified role	(Not Applicable)	10 Gigabit Ethernet

5200-Series Model 2000ES Preconfigured and Configurable Port Role Assignments

Table 5-7 contains information the 5200-Series Model 2000 ES device's default port role configuration, as well as the possible roles available for each port. The ports are listed in the order in which they appear on the device (from left to right).

For information about the types of port roles, see Table 5-1.

Note that the default port roles are those supplied in manufacturing. These roles are in effect prior to running the Getting Started wizard.

Table 5-7: 5200-Series Model 2000 ES Port Role Assignments

Port Number	Speed	Peer Port Number In Bypass Mode	Possible Roles	Default Role
1	10 Gigabit Ethernet	2	External	External
2	10 Gigabit Ethernet	1	Internal	Internal
3	10 Gigabit Ethernet	4	External	External
4	10 Gigabit Ethernet	3	Internal	Internal
5	10 Gigabit Ethernet	N/A	HA	Unused

Table 5-7: 5200-Series Model 2000 ES Port Role Assignments *(Continued)*

Port Number	Speed	Peer Port Number In Bypass Mode	Possible Roles	Default Role
6	10 Gigabit Ethernet	N/A	HA	Unused
M1	10/100/1000	N/A	Mirror, Discard, Unused	Unused
M2	10/100/1000	N/A	Management	Management

Port Roles for 5200 Series Model 2000ESL Units

Port Roles on Corero Network Devices vary in their use, and their flexibility. Some port roles are preassigned during manufacturing, and cannot be modified. Some ports can be assigned one of several port roles. Note that a port can only be assigned one role at a time. In addition, each model has restrictions on how many of a given port role type can be implemented at any given time.

For a description of port roles, see Table 5-1.

Table 5-8 explains port roles for 5200 Series Model 2000 ESL Corero Network Devices.

Table 5-8: 5200 Series Model 2000 ESL Port Roles

Port Role	Number of Ports that can Concurrently Share This Role	Operating Speed
Discard	0 or 1	10/100/1000
High Availability (HA)	0, 1, or 2	10 Gigabit Ethernet
High Availability (HA) Interconnect Switch	0, 2, or 4	10 Gigabit Ethernet
Management	1	10/100/1000
Mirror	0 or 1	10/100/1000
Unused These ports are not configured with a specified role	(Not Applicable)	10 Gigabit Ethernet

5200-Series Model 2000ESL Preconfigured and Configurable Port Role Assignments

Table 5-9 contains information the 5200-Series Model 2000 ESL device's default port role configuration, as well as the possible roles available for each port. The ports are listed in the order in which they appear on the device (from left to right).

For information about the types of port roles, see Table 5-1.

Note that the default port roles are those supplied in manufacturing. These roles are in effect prior to running the Getting Started wizard.

Table 5-9: 5200-Series Model 2000 ESL Port Role Assignments

Port Number	Speed	Possible Roles	Default Role
S1	10 Gigabit Ethernet	HA Interconnect Switch	HA Interconnect Switch
S2	10 Gigabit Ethernet	HA Interconnect Switch	HA Interconnect Switch
S3	10 Gigabit Ethernet	HA Interconnect Switch	HA Interconnect Switch
S4	10 Gigabit Ethernet	HA Interconnect Switch	HA Interconnect Switch
5	10 Gigabit Ethernet	HA	HA
6	10 Gigabit Ethernet	HA	HA
M1	10/100/1000	Mirror, Discard, Unused	Unused
M2	10/100/1000	Management	Management

Port Roles for 5200 Series Model 2400ES Units

From a management and operations perspective, the Model 2400 ES is comprised of two distinct subsystems:

- The Upper Subsystem of the Model 2400 ES is equivalent to a Model 2000 ES. In the user interface, you will see this subsystem referred to as 5500-2000ES (2400ES-Upper). For information on the ports available on the Model 2000 ES unit, see Port Roles for 5200 Series Model 2000ES Units (page 5-11).
- The Lower Subsystem of the Model 2400 ES is equivalent to a Model 2000 ESL. In the user interface, you will see this subsystem referred to as 5500-2000ESL (2400ES-Lower). For information on the ports available on the Model 2000 ESL unit, see Port Roles for 5200 Series Model 2000ESL Units (page 5-13).

Note that the user interface treats the Model 2400 ES as two separate entities: a Model 2000 ES above, and a Model 2000 ESL below.

Whenever you use the Management application to interact with the Model 2400 ES, you will either interact with the 2000 ES subsystem, or the 2000 ESL subsystem. This includes operations such as using the Getting Started Wizard, performing configuration functions, and viewing status.

Chapter 6
Viewing and Configuring Ports

Most of the ports on a Corero Network Device have fixed roles, while you can assign a few of the ports to one of several possible roles. Each port role entails specific capabilities and operations. Some ports can be assigned more than one role, and some ports have fixed roles.

If you are unfamiliar with the available port roles, and which port roles are available on which devices, see Chapter 5, “Understanding Ports”.

This chapter contains the following sections:

Viewing Port Status

You can use the management application to display a dynamically changing view of port status. This display is called the Front Panel view. When you run the Getting Started wizard, or modify port settings, the device updates the Front Panel view to reflect your configuration choices.

The Front Panel View enables you to view port role, state, and statistical information. It also enables you to view system information, and modify port information.

To display the Front Panel View on a Corero Network Device:

1. Do one of the following:
   - Choose Monitor System > Front Panel from the Navigation Tree.
   - Select the Front Panel icon from the toolbar at the top of the main window. 

   The IPS Front Panel View displays (Figure 6-1).

For a complete description of the Front Panel view and its features, see Using the Front Panel View (page 18-2).

Figure 6-1 shows the Front Panel View for an IPS Unit.

Figure 6-1: IPS Front Panel View

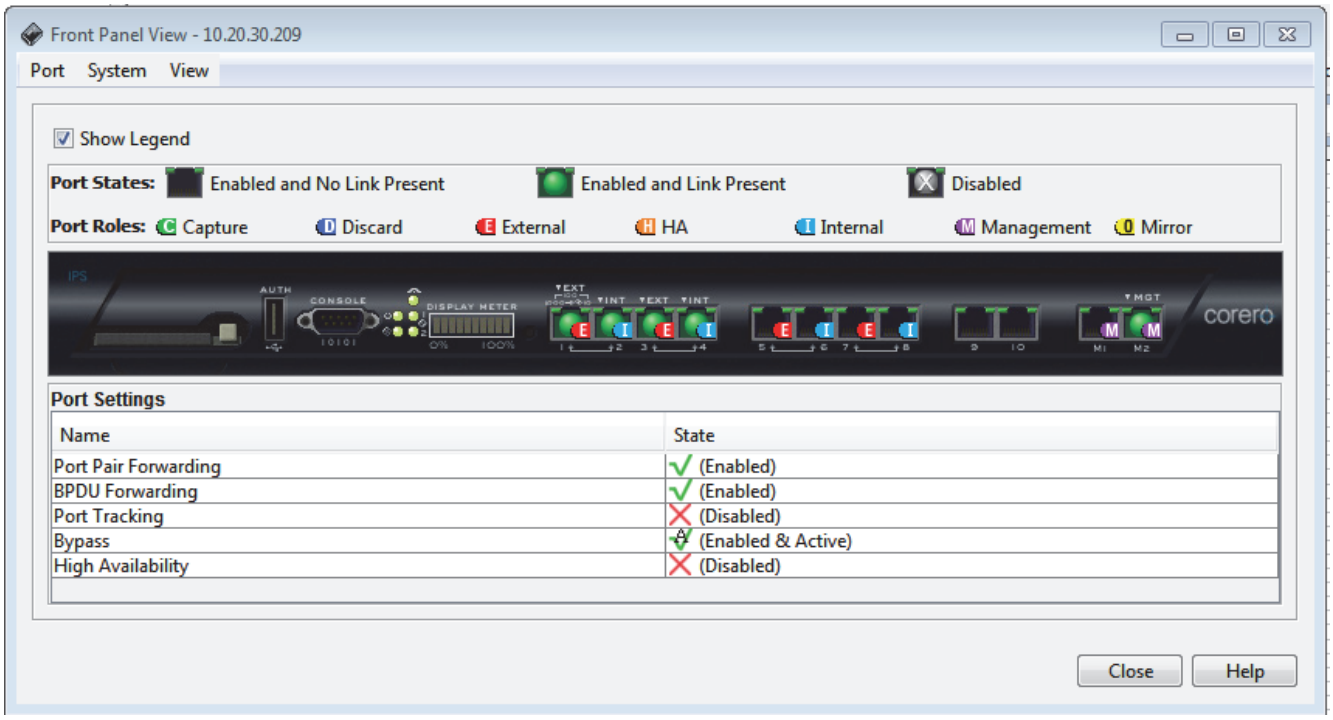

Configuring Corero Network Device Ports With the Getting Started Wizard

Each port has a set of acceptable roles along with a default port role assignment. For detailed information on port roles, see Chapter 5, ‘‘Understanding Ports”.

To set the role of a specific port, use the Getting Started wizard. The wizard queries the user to specify the number of ports needed for a specific role. The wizard then chooses the appropriate ports for these roles. Using the wizard allows the Corero Network Device to enforce a port role's requirements and to locate the ports and group them according to their roles (for example, port pairs are placed next to each other).

You can change the port role assignments of the configurable ports by re-running the Getting Started wizard at any time. Only those items whose settings you change are affected.

For detailed instructions on running the Getting Started Wizard for your IPS Unit, refer to Running the Getting Started Wizard (page 4-10)

Viewing and Modifying Port Settings

The Ports window summarizes the settings for each port.

Once you have configured a port using the Getting Started Wizard, you can edit a subset of the port settings at any time.

Note that there are occasions when the system's operational mode will force a particular port mode for a period of time. For example, the ports are forced to full duplex in bypass mode, because when the Corero Network Device is in bypass mode, mission ports may power up in a half duplex mode. This can result in the device reporting collisions and dropped packets. To avoid this potential problem, the ports are forced into full duplex mode when the device powers up in bypass mode. For a detailed description of Security Bypass, refer to Bypass Settings (page 5-6).

To view port information using the Corero Network Device management application:

1. Choose Configure System > Ports > Ports from the Navigation Bar. The Ports dialog box displays (Figure 6-2) .

Figure 6-2: Ports Dialog Box

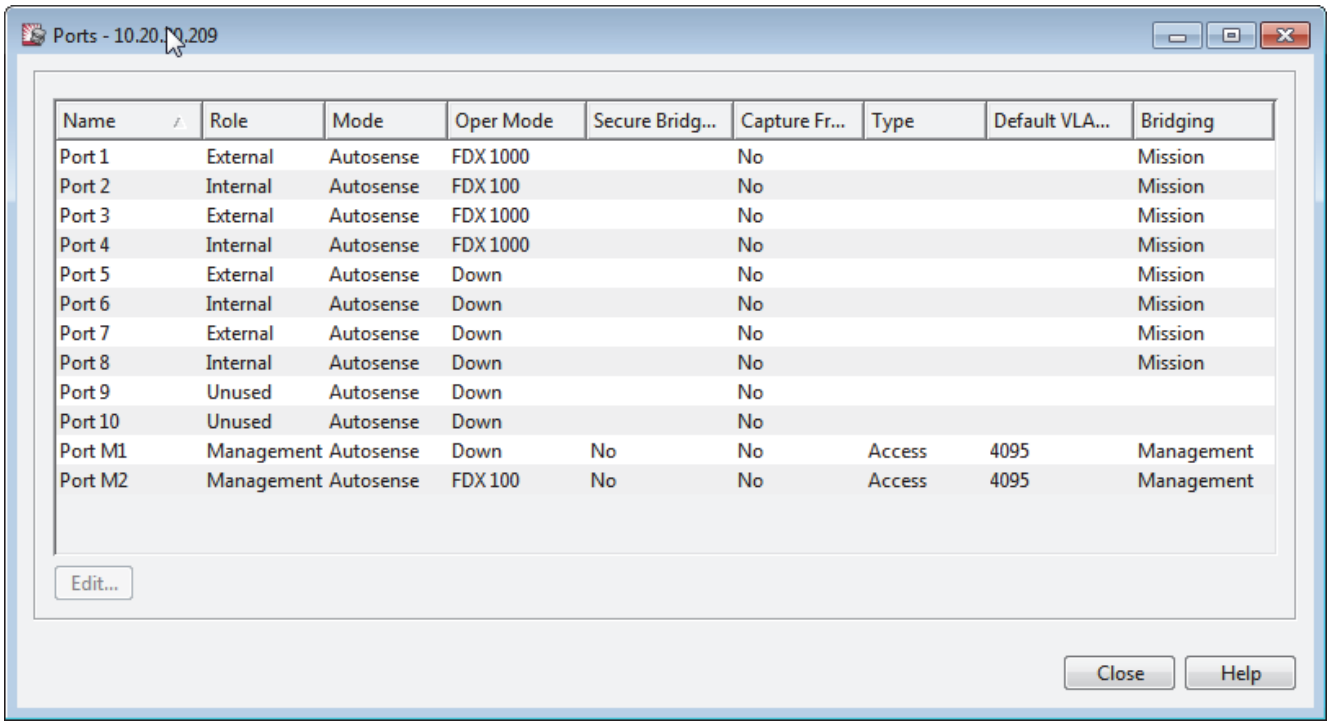

Table 6-1 summarizes the information in the Ports dialog box.

Table 6-1: Ports Dialog Box

Column	Description
Name	Corresponds to one of the port numbers on the front of the Corero Network Device. On 5100-Series hardware ports 5, 6, 7 and 8 are used by the device to communicate with a second device in a redundant configuration, or a single inline with peer configuration. On 5200-Series hardware, ports 5 and 6 are used.

Table 6-1: Ports Dialog Box *(Continued)*

Column	Description
Role	Indicates the assigned function for this port. For more information about port roles, see Port Role Overview (page 5-2).
Mode	Indicates the port's transmit/receive speed/mode setting. For more information, see Viewing and Modifying Port Settings (page 6-4).
Oper(ating) Mode	Indicates whether this port is connected and operational. If operational, the value indicates the port's actual speed.
Secure Bridging	Limits secure port traffic to user-defined host MAC addresses. There is no dynamic learning performed on Secure ports. Allowed hosts must be configured as a static MAC address with a fixed port.
Capture From	Indicates whether the traffic from this port is copied to the Capture port.
Type	For ports that are part of a VLAN, indicates the port's VLAN type: Access or Trunk.
Default VLAN ID	For ports that are part of a VLAN, indicates the default VLAN ID for this port. **Note:** For detailed information on this setting's use and how to modify it, see Chapter 9, "Advanced Port Configuration".
Bridging	Indicates the bridging domain that this port is assigned to: • **Mission** ports handle your network traffic. • **Management** ports handles management and maintenance traffic. Ports with the roles Management, Capture, Discard and Mirror are assigned to management.

2. In the Ports dialog box, select the port you want to edit, then click Edit. The Edit Port Settings dialog box displays.
3. For all port roles, you can modify the Mode. You can set a port to Autosense or to a specific setting for port speed and operation. The following options are available:
   - Autosense — Port speed is automatically set based on observed traffic.
   - Disabled — The port is not available for use.
   - HDX10 — Half duplex, 10 Mbps
   - HDX 100 - Half duplex, 100 Mbps
   - FDX10 — Full duplex, 10 Mbps
   - FDX100 — Full duplex, 100 Mbps
   - FDX1000 — Full duplex, 1Gbps (available on all 5100-Series hardware)
   - FDX10GbE — Full duplex, 10Gbps (only available on 5200-Series 10GbE ports)

   The list of options presented depends on the type of port you selected.

   NOTE

   Typically, during deployment, if you are using port-pair forwarding, choose selections to specifically match-up the settings for each external-internal port pair, rather than using Autosense to detect speed and operation.

4. If you are editing the port settings for a non-mission port, you can also choose whether or not to select Secure Bridging. Secure Bridging limits port traffic to user-defined host MAC addresses. There is no dynamic learning

performed on secure ports. Allowed hosts must be configured as a static MAC address with a fixed port as described in VLAN Overview (page 9-3).

When Secure Bridging has been selected, the port only receives packets whose source MAC address/VLAN ID pair has been configured as a static entry. The Corero Network Device only forwards packets to a Manual Mode port if the packet's destination MAC/VLAN ID pair has been configured as Static.

If you are editing the port settings for a non-mission port, you can also specify the Default VLAN ID. The Default VLAN ID specifies the ID that is applied to untagged packets received on a Trunk port, and is the ID that is expected if a tagged packet is received on an Access port. It is also the ID used to tag packets on a Trunk port when they are originated by the Corero Network Device (for example, management packets). If the VLAN ID setting is enabled (that is, the VLAN ID is nonzero) and the port is an Access port, then tagged packets received on this port must have this ID; otherwise, they are dropped. For detailed information on setting the Default VLAN ID, see Chapter 9, "Advanced Port Configuration".

5. When you are finished, click OK.

Viewing and Naming Port Pairs

A port pair is a dedicated pair of ports. Traffic that travels through the Corero Network Device will always enter through one port in the pair, and exit through the other. Port pairs are also called segments. You specify segments when you are configuring a FW+IPS policy.

When you run the Getting Started wizard, it creates port pairs (called segments). Depending on how you configure and use your device, you will use one or more of these port pairs to handle your internal and external mission traffic.

You can configure port pairs when you run the Getting Started wizard. You can then use the Port Pairs window to provide meaningful names to each port pair that you intend to use for internal/external traffic. The names you assign are the names you will see when you define the segments that apply to each firewall (FW+IPS) security policy.

NOTES

1. A port pair name cannot begin with a reserved word such as Internal, External, Inbound, Outbound, IP, TCP, UDP, ICMP, or Any. These dedicated words are solely used to define security policies.
2. Port pair members are automatically assigned by the Corero Network Device, and cannot be modified by the user.

To modify port pairs using the Corero Network Device management application:

1. Choose Configure System > Ports > Port Pairs from the Navigation Bar. The Port Pairs dialog box displays.

   The Port Pairs dialog box lists all of the mission port pairs configured for the selected devices, shows which ports belong to each port pair, and indicates the status of the links for the port pair (Up or Down).

2. To rename a port pair:
   a. Select a port pair and click Edit. The Edit Port Pair dialog box displays.
   b. Enter the port pair name in the Name field, then click OK.

3. In the Corero Network Device management application, save your changes by clicking the “Save Configuration” Toolbar button.

Selecting the Bypass Settings Mode

When you initially configure your Corero Network Device, you can specify the bypass setting using the Getting Started Wizard. Thereafter, you can modify the bypass setting using the getting started wizard, or you can modify it using the Bypass Settings dialog box, as described below.

NOTE

The bypass mode setting applies to all mission port pairs.

For a detailed description of Bypass Settings, refer to Bypass Settings (page 5-6).

To modify bypass settings using the Corero Network Device management application:

1. Choose Configure System > Bypass Settings from the Navigation Bar. The Bypass Settings dialog box displays.
2. Select the desired Bypass Setting, then click OK. Options include:
   - Never Bypass — Perform mitigation on all traffic. Should a software failure occur, pass no traffic.
   - Always Bypass — (Default) Inspect all traffic and report on attacks, but pass all traffic.
   - Bypass During System Reset — Perform mitigation on all traffic. Should a software failure occur, pass all traffic. Continue with full mitigation as soon as operation is restored.

   NOTE

   If a Corero Network Device is using a trial (time-limited) System License, when the license expires, the system will automatically enter Bypass mode. When locked in Bypass mode, you will be able to modify the bypass settings, but your changes will not take affect until you acquire a new system license key from Corero and apply it to the device.

3. Save your changes by clicking the “Save Configuration” Toolbar button.

Modifying Traffic Capture Settings

If you added a Capture port when you ran the Getting Started wizard, use the Select Port to Capture window to choose which Internal or External port's traffic the Corero Network Device will send to the capture port. You may also want to enable or disable the capture function.

To modify traffic capture settings using the Corero Network Device management application:

1. Choose Configure System > Ports > Capture Port from the Navigation Bar. The Select Port to Capture dialog box displays.
2. On the Select Port to Capture dialog box, use the check box to specify whether to enable (selected) or disable (deselected or cleared) capturing from the port.
3. Use the drop-down list to choose the port from which you want to capture data.
4. When finished, click OK.
5. Save your changes by clicking the "Save Configuration" Toolbar button.

Chapter 7
Managing Users

Corero Network Devices maintain user accounts that uniquely identify each user and the user's allowed management privileges.

In addition to the user's name and password, each account stores information on whether the user account is active or inactive, and what privileges are available to that user.

This chapter describes how to create and manage user accounts and user groups. It contains the following sections:

User Account Passwords

The management application authenticates a user through a login name and password. For security, the device's authentication process requires that passwords meet the criteria listed in Table 7-1.

Table 7-1: User Account Password Parameters

Parameter	Description
Password Length	The absolute minimum length for a password is 8 characters. The maximum length for a password is 64 characters. You can modify the minimum password length required.
Allowed Characters	A password may include any printable character.
Allowed Passwords	There are several criteria to ensure the password cannot easily be guessed: • The password cannot be the same as the user name. • When the user changes the password, they cannot reuse previous passwords. You can modify how many previous passwords are checked for comparison. • The password must contain one uppercase letter, one lowercase letter, and one digit.
Expiration	Users must change their password within a specific time frame. You can modify how frequently users must change their password. You can modify how far in advance the Corero Network Device warns users that their password is going to expire. If you set this value to 0 (zero), the password will never expire.

NOTE

You cannot change a user's login name. If a user's name changes, you must delete the account and create a new account with the new name.

Some of the above authentication parameters are set for an individual user, as described in Managing Users (page 7-3), and some are global, as described in Configuring Global User Security Settings (page 7-7).

User Account Lockouts

There are certain conditions under which a user's account status is changed to locked and the user will not be permitted to log into the management application. These cases include:

- When the user's account has expired. If a user's account is locked due to password expiration, an administrator with sufficient privileges can unlock the account. Once the account is unlocked, the user is required to change the password on their next login.
- When the user or another person has attempted to log in, but has failed a predetermined number of times. The number of allowed attempts is configurable as described in Configuring Global User Security Settings (page 7-7).
- When a security administrator has changed the status of the user's account to inactive.

Managing Users

When you create or modify a user account, you must specify the user's name and password.

You must also specify the account status. Account status options are listed in Table 7-2.

Table 7-2: User Account Status Options

Status	Description
Active	The user can log in and perform the activities authorized by the user's privilege setting.
Inactive	The user is not permitted to perform any activities.
Locked	The user cannot log in, either because they entered the incorrect password too many times, or because the user's password has expired. The user must contact an administrator in order to have the account activated. If the account was locked because the password expired, the administrator will be required to change the password when unlocking the account.

In addition, you must specify the Privilege level. Privilege levels are listed in Table 7-3.

Table 7-3: User Account Privilege Options

Status	Description
Monitor	The user can view all Corero Network Device information, but cannot add, delete, or modify any settings.
Administrator	The user can perform all device configuration tasks and view all operational information, statistics, and reports.

NOTE

User account settings can be modified or superceded by global user security settings. For example, you can specify that administrators cannot set the user status when creating the user. For information on these settings, see Configuring Global User Security Settings (page 7-7).

To manage users:

1. From the Navigation Tree, choose Configure System > Management Access > Management Users > Manage Users.

   The Management Users dialog box displays (Figure 7-1).

Figure 7-1: Management Users Dialog Box

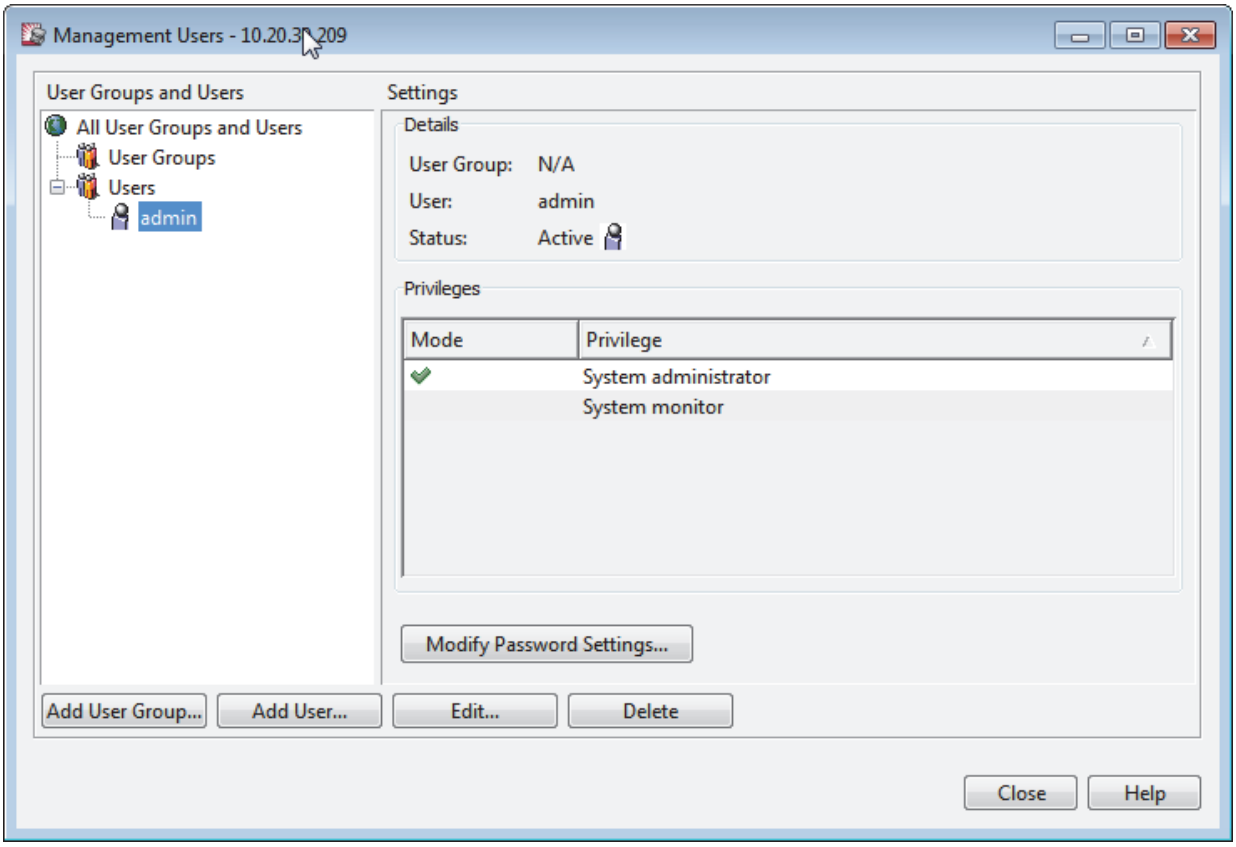

2. To add a user:

   a. Do one of the following:
   - To create a user in a user group, select the user group and click Add User.
   - To create a user who is not part of a user group, select All Users and User Groups and click Add User.

   The Add User dialog box displays.

   b. In the Add User dialog box, enter the user's name.

   c. If you have enabled the "Allow Setting a User's Status When a User Is Created" feature, you can set the user's status at this time. Status options are listed in Table 7-2. For more information on this feature, see Configuring Global User Security Settings (page 7-7).

   d. Select the user's privilege level. Privilege levels are described in Table 7-3.

   NOTE

   If the user is part of a user group, the privilege level you select must match the privilege level for the group.

   e. Do one of the following:
   - If you want to create the user in a user group, ensure the desired user group is selected.
   - If you do not want to create the user in a user group, do not select a user group.

   f. Enter the password for this account.Enter it again for verification.

NOTE

For more information about passwords and global security settings that affect them, see User Account Passwords (page 7-2).

g. Specify when the password will expire. When the password is nearing its expiration, the user will be informed that they must select a new password. If you do not want the password to expire, enter zero.

h. Do one of the following:
- If you want to add another user, click Add, then specify information for the next user.
- If you do not want to add another user, click Done.

3. To modify a user:

a. Select the user and click Edit. The Edit User dialog box displays.

b. If desired, modify the user's status and privilege level (described above)

c. If desired, specify a user group to which the user will belong.

d. When finished, click OK.

NOTE

You cannot modify the user account you are currently using (logged in to).

4. To modify a user's password

a. Select the user and click Modify Password Settings. The Modify Password Settings dialog box displays.

b. If desired, specify the number of days before the new password will expire. If you do not want the password to expire, enter zero.

c. To change the password, select the Change Password check box.

d. Enter the old password.

e. Enter the new password twice, for verification.

f. When finished, click OK.

5. To delete a user

a. Select the user.

b. Click Delete.

c. You are asked to confirm your selection. Click Yes to delete the user, or click No to retain it.

6. Save your changes by clicking the Save Configuration toolbar button.

Managing User Groups

For ease of management, you can create groups of users who have the same level of privileges. For more information on user privileges, see Managing Users (page 7-3).

At any given time, a user can belong to only one group.

NOTE

You cannot modify an existing User Group. If you want to make changes to an existing group, delete the existing user group and add a new group with the changes you desire.

To manage user groups:

1. From the Navigation Tree, choose Configure System > Management Access > Management Users > Manage Users.

   The Management Users dialog box displays (Figure 7-1).

2. To view the available user groups, click User Groups in the left pane.
3. To add a user group:
   a. Click Add User Group. The Add User Group dialog box displays.
   b. Enter the name of the user group, and its privilege level. For more information on privilege levels, see Managing Users (page 7-3).
   c. Do one of the following:
      - If you want to add another user, click Add, then specify information for the next user.
      - If you do not want to add another user, click Done.
4. To delete a user group:
   a. Select the user group you want to modify, then click Delete.
   b. You are asked to confirm your request.
5. Save your changes by clicking the Save Configuration toolbar button.

NOTE

When you delete a user group, all users in that group are transferred to the default group.

Configuring Global User Security Settings

The management application enables you to modify security settings that apply when you are creating or modifying a user account.

NOTE

These settings only affect user accounts you create or modify in the future, not those for existing users.

To modify global user security settings:

1. From the Navigation Tree, choose Configure System > Management Access > Management Users > User Security Settings.

   The User Security Settings dialog box displays.

2. Modify the settings as desired, then click OK. The settings are described in Table 7-4.

3. Save your changes by clicking the Save Configuration toolbar button.

Table 7-4: Global User Security Settings

Setting	Description
Allow setting a user's status when a user is created	This allows a user with administrator privileges to both create and activate a user in one operation. **Note:** If this box is not selected, the administrator must create the user with an inactive status, exit the new account, then edit the account separately to activate the user.
Change a user's status to "Locked" after failed attempts	You can specify how many times a user can try to enter their password. If this number of attempts is exceeded, the user account status is automatically locked, requiring intervention from an administrator to reactivate it.
(Password) Minimum Length	The minimum number of characters allowed for a password (from 8 through 64).
(Password) Default Expiration	The number of days that a new password remains valid (from 1 through 999). You can also specify zero, which indicates that the password does not expire. Note that you can modify this setting when creating a user account.
(Password) Expiration Warning	The number of days before the password expires that the Corero Network Device will warn the user during login that the password will soon expire (from 1 through 49,710 days).
(Password)| History Depth	The number of previous passwords that the Corero Network Device should remember and compare against the new password when the user attempts to change the password (from 3 through 8 per user).

Chapter 8
Management Access

The Corero Network Device permits secure management access in several ways:

- Secure Architecture — The device's architecture separates management traffic (which is traffic entering the device's on one of the management ports) from mission traffic (which is traffic entering the device's external and internal ports). The device can bridge management traffic from one management port to another management port, but does not forward management traffic to any other port, nor does it forward non-management traffic to any management port.
- User Account Administration — Create and manage user accounts, providing different sets of privileges to individual users and groups of users.
- User Authentication — Establish a user's identity and privilege level.
- Management Service Access Control — Provide overriding controls for each form of management access.
- Mission and Management Traffic Isolation - Mission and Management traffic are never combined in a Corero Network Device.

This chapter describes the various forms of device management access and their controls.

The chapter includes the following sections:

Management Session Overview

Your Corero Network Device creates a management session whenever a successfully authenticated user connects to the device. The session continues until the user ends the session or the session times out.

NOTE

Any user can have multiple simultaneous management sessions over any management service. A maximum of eight simultaneous management sessions (total for all users) is allowed. Note that the Console port supports only one management session.

Each management service has a time-out period that applies to all management sessions of that type. You can modify this setting as described in Configuring Management Port Access (page 8-6).

Information is sent to the Event log and also any configured Syslog servers. The device reports all relevant information for each management session including:

- User name
- Session privileges
- Login time (start of session)
- Management access method (HTTP, HTTPS, Telnet, Console, etc.)
- IP address of user
- MAC address of user
- Unique session ID
- Logout time (end of session)

Management Services

Corero Network Devices support a number of management interfaces. Table 8-1 summarizes the capabilities of each interface.

Table 8-1: Corero Network Device Management Interfaces

Interface	Description	For More Information, See...
Serial Console	The CONSOLE port on the front of the Corero Network Device provides access to a limited Command Line Interface that you can use to perform basic setup.	Serial Console Access and the Command Line Interface (page 8-4)
Command Line Interface Over Telnet	You can access a Command Line Interface (CLI) using a Telnet session. The CLI provides access to nearly all of the device's configuration and management commands.	Understanding SNMP Management (page 8-13)
Management Application using HTTP and HTTPS	The Management application runs as a Java Web Start™ program over the World Wide Web. The Graphical User Interface (GUI) provides interface windows for configuring and managing the device and provides access to context sensitive online help. The device provides persistent management sessions using HTTP (port 80) or HTTPS (port 443). These ports allow full management access through your firewall. The Graphical User Interface for the Management application is not dependant on a Web browser. Instead, it uses the Java Web Start technology to provide a self-updating application that operates over the Web.	Chapter 3, "Getting Started With the DDS Management Application" Managing SSL Certificate and Key Information for HTTPS Access (page 8-8)
SNMP	The device's Simple Network Management Protocol interface provides read-only access to many of the device's settings. You can use SNMP to read the values of your device's proprietary MIB objects. You cannot change these values using SNMP.	Understanding SNMP Management (page 8-13) Managing SNMP Parameters (page 8-15)
IPS Controller	The IPS Controller is used to centrally manage multiple devices. It runs on a separate system.	IPS Controller Interface Settings (page 8-16)

NOTE

When changing the IP address of the Corero Network Device via the console, it is necessary to reboot the device to enable the change to take effect.

Serial Console Access and the Command Line Interface

The basic Command Line Interface (CLI) supports Corero Network Device setup and diagnostic functions.

The Setup command, which is accessed from the CONSOLE port on the front of the physical device, enables you to set the following device-specific parameters:

- IP address
- Subnet mask
- Time and time zone parameters
- Administrator information
- Default gateway (Optional)

Serial Console Port Authentication

The serial console port supports an authentication feature, which provides an extra layer of protection where the Corero Network Device is deployed in an environment without physical security. In these cases, if authentication on the serial console port is enabled, only authorized users are allowed to access the CLI.

By default, serial console port authentication is disabled.

To enable serial console port authentication:

1. Connect a computer to the serial console port.
2. Enter the following CLI command:

   set auth[enticate]

3. Thereafter, when the device starts, or when the serial console cable is connected and authentication is enabled, the console port user must supply the device's "admin" username and password. Once these are entered, the user assumes admin privileges with full access to the Command Line Interface (CLI).

NOTES

1. You cannot enter a blank line for the password. If you are using the default admin account and password, you must enter a single asterisk "*" followed by the Enter key.
2. Only one user with admin privileges is allowed access to the CLI at a time.
3. Disconnecting the console cable (loss of DSR) forces a logout if authentication is enabled and someone is logged in.

The Serial Console Port CLI commands are listed in Table 8-2.

Table 8-2: Serial Console Port CLI Commands

Command	Description
set auth[enticate] <on|off>	Enables or disables console authentication. The default setting is off (disabled). • When authentication is set to off (disabled), the console user automatically assumes admin privileges when either the Corero Network Device is rebooted or the serial port cable is connected. • When authentication is set to on (enabled), the console user retains admin privileges until they issue the logout command, the console cable is disconnected, or the console port is automatically logged out due to console inactivity.

Table 8-2: Serial Console Port CLI Commands *(Continued)*

Command	Description
show auth[enticate]	Displays the console authentication status.
login	Prompts the user for the correct username and password in order to log into the CLI with admin privileges. **Note:** You cannot enter a blank line for the password. If you are using the default admin account and password, you must enter a single asterisk "*" followed by the Enter Key.
logout	Allows the user to stay connected to the CLI, but with read-only access. **Note:** For security purposes and when using authentication, you should always log out of the CLI when it is not in use.
show status	If authentication is enabled, displays the current authenticated user name.
set autol[ogout] <off|default|30...3600>	When set to off (disabled), this command disables automatic logout due to console inactivity after the specified period of time. When set to default, autologout is enabled, when authentication is on (enabled). If authentication is off (disabled), there is no autologout event. The default inactivity period is 300 seconds (5 minutes), but it can be set to any value from 30 seconds to 3600 seconds (1 hour).
show autol[ogout]	Displays the autologout state. • If autologout is disabled, this information is displayed. • If autologout is set to the default (enabled), the amount of time before autologout occurs is displayed.

Configuring Management Port Access

For added security, your Corero Network Device provides restrictions for using remote access. You can choose whether to allow or deny management port access through HTTP, HTTPS, SNMP, and Telnet. And for the Telnet, HTTPS, and HTTP services, you can specify permitted IP address ranges. Locking down specific IP address ranges that are permitted to manage the device enables the unit to block scanning traffic received on the management port, because the source address does not fall within the approved IP address range.

CAUTION

Be careful when restricting management access to the Corero Network Device that you do not unintentionally disconnect your current management session. For example, if you are accessing the device through a Telnet session and you restrict management access to an IP address range that does not include the IP address of the current Telnet session, the current session will be disconnected.

To configure access to management ports:

1. From the Navigation Tree, choose Configure System > Management Access > Mgmt Port Control/SNMP. The Management Port Control dialog box displays.
2. To modify the access to a service (HTTP, HTTPS, SNMP, or Telnet), do the following:
   a. Select the service, then click Edit. The Edit Access Settings dialog box displays.
   b. Select the desired access (allow or deny), then click OK.
3. To modify the client IP address ranges that are permitted to access the management port:
   a. Select the desired service (HTTP, HTTPS, SNMP, or Telnet), then click Client IP Ranges. The Access Settings - Client IP Ranges setting dialog box displays, showing all permitted IP ranges for that service type.
   b. To add a client IP range, click Add. The Add Client IP Range dialog box displays. You can add IP addresses in four ways:
      - As an IP address/Prefix (for example 192.0.8.31/24, where the first 24 bits would have to match exactly.)
      - As an IP address/Mask (for example 192.0.8.31/255.255.255.0).
      - As a range specified by a first and last IP address (for example 192.0.128.0-192.01.128.255).
      - As a single IP address (for example 192.0.8.31).
   c. To remove a client range, select the range, then click Delete.
4. Save your changes by clicking the Save Configuration toolbar button.

NOTE

Management IP addresses can only be accessed via the dedicated management port(s).

Telnet CLI Commands

You can use a Telnet session to examine most Corero Network Device management objects and set the values of most configuration variables. The commands listed in are available from the Telnet accessible command line interface.

NOTE

Before attempting to use a Telnet session, verify that the device's Management Port Control restrictions for Telnet are set to Allow, as described in Configuring Management Port Access (page 8-6).

Table 8-3: Telnet CLI Commands

Command	Description
ping	Used to send an ICMP Echo packet to a specified IP Address. `ping <your-ip-address>`
set	Used to set configuration parameters on the device. To obtain a listing of all set commands, enter the following command: `set ?`
show	Used to show configuration parameters and status information on the device. To obtain a listing of all show commands, enter the following command: `show ?`
mib	Used to display all MIB information. `mib`
save	Used to save all configuration information to the Compact Flash card. `save`

reboot	• To reboot the device, and retain current settings, enter:
	`reboot`
	(Executing this command is equivalent to choosing Manage System > Reboot from the Navigation Tree)
	• To reboot the device and reset it to factory defaults, enter:
	`reboot factory`
	(Executing this command is equivalent to choosing Manage System > Reset to Factory Defaults from the Navigation Tree.

quit	Used to end this command session. `quit`

Managing SSL Certificate and Key Information for HTTPS Access

If you are going to use HTTPS to access your Corero Network Device, you must restrict these management services to a port and/or an IP address range as described in Configuring Management Port Access (page 8-6).

In addition to this access, the device must have a security certificate. A default SSL certificate is provided and activated during manufacturing. The currently active certificate must be retained after the device reboots. The device also supports an optional SSL Certificate Chain file.

NOTE

Each Corero Network Device supports a maximum of three SSL certificates.

Before you upload an SSL Certificate or Key, consider the following:

- The Certificate and Key file names cannot be longer than eight characters (excluding the extension).
- The Certificate and Key file names must match, with the exception of the file extensions.
- The Certificate file name must end with the .cer extension.
- The Key file name must end with the .key extension.
- The (optional) Certificate Chain file must:
  - Not have a filename longer than eight characters (excluding the extension).
  - End with the .pem extension.
  - Be smaller than 16 KB in size.
  - Contain no more than eight certificates
  - Contain certificates of no more than 8 KB in length, when decoded in CER format.

To manage SSL Certificate files:

1. From the Navigation Tree, select Configure System > Management Access > SSL Certificates. The SSL Certificate Management dialog box displays (Figure 8-1).

Figure 8-1: SSL Certificate Management Dialog Box

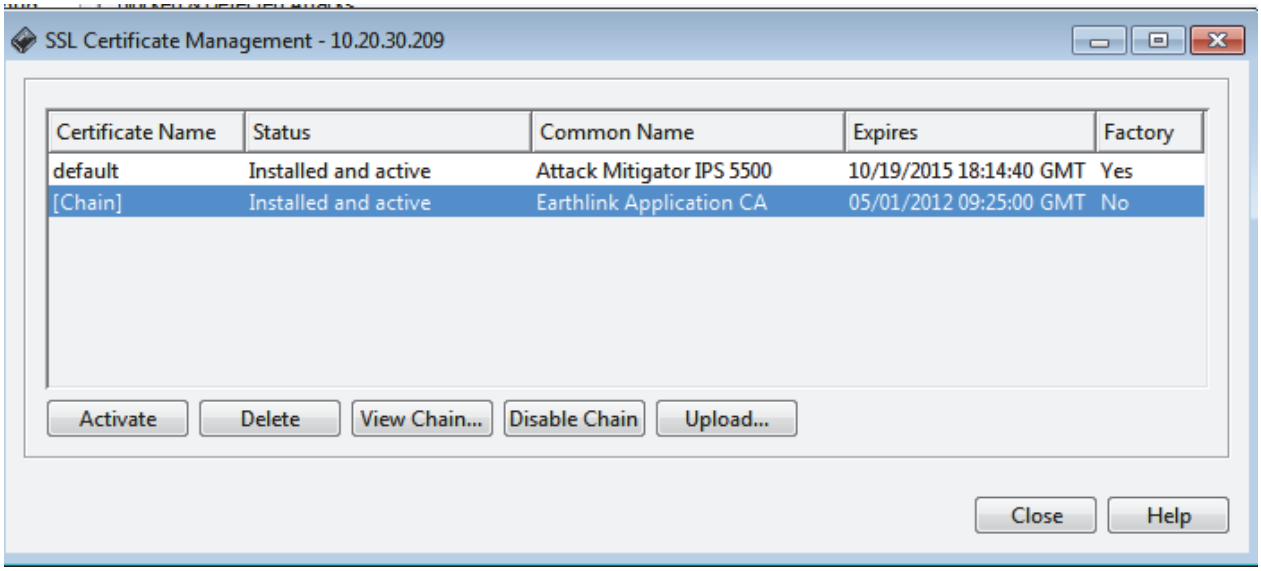

2. View the certificate listings, status, and expiration date of each certificate. Valid status options include:
   - Installed - this certificate has been installed but is not currently active.
   - Installed and Active - this certificate has been installed and activated.
   - Uploading - this certificate is in the process of being uploaded.
   - Invalid - The installation of this certificate failed.
3. To upload a certificate:
   a. Click Upload. The Upload SSL Certificate and Key window displays.
   b. To specify the Certificate file name, select the Upload SSL Certificate check box, then browse to and select the Certificate File Name (ending in .cer).
   c. To specify the Key file name, browse to and select the Key File Name (ending in .key).
   d. Click the Upload button on the Upload SSL Certificate and Key window. The system responds with the certificate status (see above).
4. To upload an SSL Certificate Chain,
   a. Click Upload. The Upload SSL Certificate and window displays.
   b. To specify the SSL Certificate chain, select the Upload SSL Certificate Chain checkbox, then browse to and select the SSL Certificate Chain file name (ending in .pem).
   c. Click the Upload button on the Upload SSL Certificate window. The system responds with the certificate status (see above).
5. To activate a certificate or chain, select the certificate or chain then click Activate.
6. To delete a certificate or chain, select the certificate or chain and then click Delete.
7. To view the certificates inside the certificate chain, select the certificate chain and then click View Chain.
8. To disable the certificate chain, select the certificate chain and then click Disable Chain.
9. When you are finished, save your changes by clicking the Save Configuration toolbar button.

User Authentication Settings

Corero Network Devices enable you to specify the authentication method to use to allow users access to the device's management functions. The device supports three methods for authenticating users to allow access to management functions, as described in Table 8-4.

Table 8-4: User Authentication Methods

Authentication Method	Description
Local Only	Users are authenticated against a locally maintained database on the Corero Network Device and Radius is not used.
Radius Only	Radius (Remote Authentication Dial-In User Service) is used to authenticate users. If authentication fails, that user is denied access. The device allows management access to be authenticated using Radius. Radius support includes authentication of user name/password combinations and administrator/monitor authorization. No accounting is supported. Radius is supported as defined in RFC 2865. Radius is used to provide a centralized repository of users, along with their passwords, authorization, and possibly other information. A client, in this case the device, may request a Radius server to authenticate a user's name and password, and return additional information about the user. The device uses Radius only for authentication and administrator/monitor authorization, which is required; any additional information returned by the server is ignored.
Try Radius, and if no response, use Local	Try Radius, and if no response, Use Local - Radius is first used to authenticate users. If Radius fails, local authentication is attempted. **Note:** Only a failure to reach a server will result in a local authentication attempt. If Radius authentication succeeds, but denies the user access, no attempt is made to override this with local authentication.

Radius is a user authentication protocol, in which a Radius server allows or denies the user access to a given resource. If you choose to use one or more Radius servers for authentication, you can also select the manner in which the Corero Network Device searches for the Radius server. For information on how to define Radius servers, see Managing Radius Servers (page 8-11).

You can identify up to eight Radius servers for each device. At least one server profile must be enabled for Radius to operate.

When a Radius request is generated, the Radius server is selected in the manner you specify in your authentication settings. Alternate Radius server search methods are described in Table 8-5.

Table 8-5: Alternate Radius Server Search Methods

Search Method	Description
Priority	If this option is selected, all new requests are sent (initially) to the server with the lowest configured priority value. If a request times out, the request is rebuilt and delivered to the next server (determined by it's priority value). However many Radius servers you specify, if all configured servers time out, the user is denied access. Or, if Try Radius, And If No Response, Use Local is selected, local authentication will be attempted. **Note:** When selecting servers by priority, a time out causes all new requests to be sent to the next server in the list. There is no mechanism for the Corero Network Device to return automatically to the highest-priority server (unless successive failures eventually lead back to it).

Table 8-5: Alternate Radius Server Search Methods *(Continued)*

Search Method	Description
Round-Robin	If this option is selected, the Radius Server is selected using a standard round-robin algorithm, with each new request going to the next Radius server on the rotational list. In this case, the server priority is ignored. Each new request is sent to the next server in the list (if multiple servers are configured). If a request times out, the request is rebuilt and delivered to the next server (determined by round robin rotation). However many Radius servers you specify, if all configured servers time out, the user is denied access. Or, if "Try Radius, And If No Response, Use Local" is selected, local authentication will be attempted. **Note:** By default, the device is configured to select servers using Round Robin.

To specify user authentication settings:

1. Verify that you have added any Radius servers you want to use for authentication. This procedure is described in Managing Radius Servers (page 8-11).
2. Choose Configure System > Management Access > Authentication Settings from the navigation tree. The Authentication Settings dialog box displays.
3. Select the authentication method you wish to use to authenticate users. These authentication methods are described in Table 8-4.
4. If you are using Radius servers, and have more than one of them available for user authentication, you must specify the alternate radius server search method. These search methods are described in Table 8-5.
5. When finished, click OK.

Managing Radius Servers

If you will be using Radius to provide user authentication, you will need to configure one or more Radius servers on for use by your Corero Network Device.

To configure Radius servers:

1. Select Configure System > Management Access > Management Users > Radius Servers. The Radius Servers dialog box displays.
2. Do one of the following:
   - To add a Radius server, click Add.
   - To modify an existing Radius server, select the server, then click Edit.
   - To remove an existing Radius server, select the server, then click Delete.
3. When you choose to add or edit a Radius server, a dialog box appears with the fields listed in Table 8-6.
4. When finished specifying Radius server parameters, click OK.
5. Save your changes by clicking the Save Configuration toolbar button.

Table 8-6: Radius Server Settings

Field	Description
IP Address	The address of the Radius server. This value may not be 0.0.0.0, and must be unique across all Radius servers.
Description	Enter text describing this Radius server. This text is not used during the authentication process, it is there to help the user distinguish among several defined Radius servers.

Table 8-6: Radius Server Settings *(Continued)*

Field	Description
Mode	Indicates whether this server profile is enabled or disabled. A Corero Network Device will only attempt to authenticate a user through this server if the profile is enabled. The device will not attempt to authenticate users through disabled Radius servers.
UDP Port	The UDP port used by this Radius server. Valid values are from 0 to 65535. UDP port 1812 is typically the default port used for Radius authentication, though some older Radius implementations use port 1645. Refer to the documentation supplied with your Radius server for information about its UDP port number.
Timeout	The seconds that the Corero Network Device should wait before timing out a user authentication request to this Radius server. If you have multiple Radius servers defined, this is the amount of time that the device will wait before requesting authentication from the next authentication server. You can specify any value between 1 and 255 seconds, though a timeout value of 2-3 seconds is usually sufficient.
Retries	The number of times the device may retry this Radius server if there is no response to the initial request after the timeout value is exceeded. Once the retries are exhausted, the device tries the next Radius server (if one is available). You can specify any value between 0 and 255 retries. If this value is set to zero, only one attempt will be made to authenticate the user with this Radius server. If multiple Radius servers are available, you should set this value to 1 or 2; otherwise, it should be set to 3 or 4.
Priority	When you specify more than one Radius server, you can assign a different priority to each one. This is called the Alternative Radius Server Search Method Priority. You can specify any value from 0 through 8. The server with the lowest configured priority value is always used, unless a failure is detected. The default priority is zero. **Note:** This priority value does not apply to Radius servers configured to use the Round Robin search algorithm. For more information on Radius server search methods, see User Authentication Settings (page 8-10)
Secret	Enter the Secret value, which is the Shared Secret, as defined by the Radius specification. The value entered here must exactly match (case-sensitive) the value configured on the Radius server itself. The text for the shared secret can be between eight and sixty-four characters long.

Understanding SNMP Management

Corero Network Devices can generate SNMP traps when they detect alarms or other system events. The device comes with an SNMP agent as part of its firmware. When enabled, SNMP traps can be sent to a designated trap host. A network administrator can send trap messages to up to eight trap hosts (servers).

NOTE

In line with CERT advisories, SNMP is disabled by default on Corero Network Devices.

The device supports the SNMP Get function. The device supports standard SNMP management read operations (GET, GET-NEXT), and a proprietary SNMP MIB. Port 8 is used as the management port for SNMP.

You can modify the default Get Access community string (the default is private).

NOTE

The SNMP network management system (NMS) and trap servers must be located on the management LAN connected to the management port (MGT, Port 8).

The Corero Network Device's software supports industry standard SNMP MIBs including:

- RFC1213 MIB-II (excluding at, ipRouteTable, ipForwardTable, and egp)
- RFC1493 Bridge MIB (excluding dot1dStatic)
- RFC1643 Ethernet MIB (excluding dot3CollTable, dot3Tests, and dot3Errors)
- RFC2096 IP Forwarding Table MIB

Proprietary SNMP MIB

A proprietary SNMP Management Information Base (MIB) is provided with this product. You can access the MIB from the Technical Information Menu (readme_first.html) located on the CD-ROM and also from the Corero Customer Support Portal at `https://support.corero.com`.

SNMP Get Operations Supported

Corero Network Devices enable you to GET the following standard SNMP management MIB objects on port roles:

- **System Table**s
  - System
  - Interface
  - IP
  - ICMP
  - TCP
  - UDP
  - SNMP
- **Bridge Tables**
  - dot1dBase
  - dot1dTP

Supported SNMP Traps

When you enable SNMP on your Corero Network Device, it can send SNMP trap messages to up to eight IP addresses. When configured to do so, SNMP Trap messages are sent for the following events:

- Attack Detection
- Configuration Change
- Redundant Link Up/Down
- Authentication Failure
- Link Up/Down
- Mirror Link Up/Down
- Cold Start
- Login Failure
- Other event logging system Event

Managing SNMP Parameters

The SNMP selection on the Management Port Control window provides access to SNMP settings. You can specify whether traffic from the SNMP service is allowed or denied. You can also specify SNMP-related information for the Corero Network Device itself.

In the Simple Network Management Protocol (SNMP), a managed node (agent function) can be configured to send a message (called a trap) to a management station (called a trap server) to report an exception condition. The trap servers should be located on the network connected to Port 8, the device's management port. You can configure the device to send SNMP traps to up to eight trap servers.

To manage SNMP parameters for a Corero Network Device:

1. From the Navigation Tree, choose Configure System > Management Access > Mgmt Port Control/SNMP. The Management Port Control dialog box displays.
2. Select the SNMP service and click the Edit button. The Edit Access / Community Strings / SNMP Trap Addresses dialog box displays.
3. To view or modify access settings, click the Edit Access Settings tab. The current access displays.

   To modify access settings, use the drop-down to select the desired access (Allow or Deny).

4. To view or modify the community string, click the Community Strings tab. The current SNMP Get community string displays.

   To modify the community string, enter the new community string (up to 32 characters).

   NOTE

   Corero strongly recommends that you change the default community string.

5. To view or modify the IP addresses where the Corero Network Device sends SNMP traps, click the SNMP Trap Addresses tab. The tab displays a list of the current SNMP host IP addresses. An address of 0.0.0.0 indicates that no trap server is configured for this server ID.

   Do one of the following:

   - To add an IP address, click Add. In the Add Trap Address dialog box, specify the IP address of the system, and the associated community string.
   - To modify the community string for an IP address, select the address, then click Edit. In the Edit Trap Address dialog box, modify the community string as needed.
   - To modify the IP address for a trap server, delete the existing address and add the correct address.
   - To delete an IP address, select the address, then click Delete.

6. Save your changes by clicking the Save Configuration toolbar button.

IPS Controller Interface Settings

The IPS Controller is used to manage multiple Corero Network Devices, and runs on a separate system. The IPS Controller communicates with IPS or DDS Units via a secure communications tunnel. Further security is provided by use of a shared key, which you enter on both the Corero Network Device and the IPS Controller. Each device under IPS Controller management has its own shared key. You can use the same key for multiple devices, or you can use unique keys for each device.

To communicate with an IPS Controller (IPSC), the following Corero Network Device criteria must be met:

- The device must be configured to permit IPS Controller management, as described below.
- The device and the IPS Controller must be given the same shared key. This text-based key is used during the authentication process between the two devices. This shared key is defined, both on the device and the IPS Controller, by the user.
- The device must have port 2616 open for communication with the IPS Controller.

CAUTION

If the network traffic from the device on which the IPS Controller software is running passes through the mission ports before connecting to the Corero Network Device's management port, then it is important to ensure that port 2616 is not blocked by a firewall policy. On the device, Port 2616 is not blocked by default. If Port 2616 is blocked, then the device and the IPS Controller will be unable to communicate with each other.

There are three modes of communication you can select for communication between a Corero Network Device and the IPS Controller, as described in Table 8-7.

Table 8-7: IPS Controller Management Options

Option	Description
Block IPS Controller	Any communication attempts by the IPS Controller to manage a Corero Network Device are ignored. Use this mode when you want to fully manage the device locally using the device's own management application.
Block IPS Controller until Logout	Any communication attempts by the IPS Controller to manage the device are ignored until the management session is completed. Use this mode when you want to temporarily manage the device locally.
Allow IPS Controller	The IPS Controller is allowed to communicate with (and manage) the device. Note that the communication attempt will remain unsuccessful if the shared key is not set correctly.

In addition, for secure intercommunication, each Corero Network Device uses a shared key to communicate with the IPS Controller.

To modify the settings that control whether or not the device can be managed by an IPS Controller:

1. To modify the IPS Controller settings, from the Navigation Tree, choose Configure System > Management Access > IPS Controller Management > Settings. Select the desired setting (settings are listed in Table 8-7), then click OK.
2. To modify the Corero Network Device-to-IPS Controller shared key:
   a. From the Navigation Tree, choose Configure System > Management Access > IPS Controller Management > Shared Key. The IPS Controller Shared Key dialog box displays.
   b. Enter the old shared key.

c. For verification, enter the new shared key twice.

d. When finished, click OK.

3. Save your changes by clicking the Save Configuration toolbar button.

Chapter 9
Advanced Port Configuration

The following sections discuss concepts and configuration issues related to Corero Network Device network interfaces. Concepts include how the device separates regular mission traffic from management traffic and how it handles VLANs.

This chapter describes:

Mission Traffic and Management Traffic Isolation

Corero Network Devices enforce the concept of isolated management traffic. The device isolates mission traffic from management traffic, providing separate, dedicated transmission for each type separately.

Mission Ports can only operate in Port Pair Forwarding mode, where the device only forwards the traffic from one port of a Mission Port Pair to the other port in the pair. Mission and management traffic remain separated.

With 5100-Series devices, the device can bridge management traffic from one management port to another management port, but does not forward management traffic to any other device port, nor does it forward non-management traffic to any management port.

This isolation enables you to create a separate management subnet or VLAN that only trusted devices can access, thus providing more secure management access to the device.

Corero Network Devices always have one dedicated management port. Depending on the model, you can also configure additional management ports on your device.

With 5100-Series devices, all configured management ports are bridged together and the device sends flooded management traffic to all of them.

5200-Series devices do not provide bridging between management ports.

VLAN Overview

A VLAN is a virtual LAN created within a physical local area network by tagging packets in a way that distinguishes the VLAN traffic from other packets on the network. Based on its VLAN ID, a packet is forwarded to the specific port (or ports) associated with that VLAN ID. This method provides the separation and security of individual, virtual networks within the same physical network.

The VLAN tag on a tagged packet includes three pieces of information:

- VLAN ID
- User priority
- Canonical format indicator (indicates whether or not the VLAN ID is in the accepted standard format)

Corero Network Devices are only concerned with the VLAN ID.

On Mission Port Pairs, the device is transparent to 802.1Q VLANs.

A Management port may be assigned to a VLAN. When this is configured the device will expect management traffic to belong to the 802.1Q VLAN ID configured. The default of 4095 is configured to “No VLAN” for management traffic.

VLAN Port Types

Corero Network Devices support two types of VLAN ports:

- Access Ports
  Each access port supports a single VLAN which you identify by assigning the port a default VLAN ID. Access ports generally receive untagged traffic and the device transmits untagged packets to an access port.
- Trunk Ports
  Trunk ports can support more than one VLAN. When using trunk ports, you must define a default VLAN ID that applies to all the ports in the Mission Bridge or the Management Bridge domain. The device associates the default VLAN ID with any untagged packets it receives on any Trunk port in the bridge.

When a Corero Network Device floods a packet to multiple ports within a bridge, it always transmits the same format (tagged or untagged) to all the ports. It accomplishes this by applying the following rules to ports within a bridge:

- All the ports for a given bridge (Management Bridge or Mission Bridge) are of the same type (access or trunk).
- All trunk ports within a bridge use the same default VLAN ID.

VLAN Forwarding Algorithm

When deciding how to handle VLAN enforcement, Corero Network Devices follow the process shown in Figure 9-1.

Figure 9-1: VLAN Forwarding Algorithm

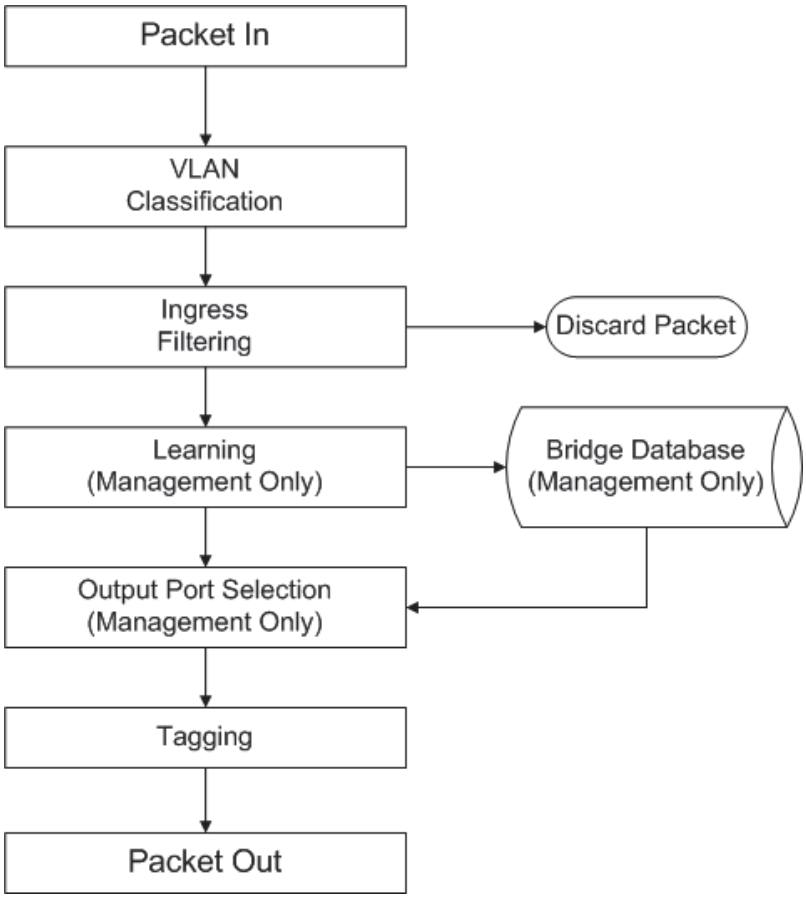

The forwarding algorithm steps shown in Figure 9-1 are described in Table 9-1.

Table 9-1: VLAN Forwarding Algorithm

Step	Description
VLAN Classification	Corero Network Devices assign every arriving packet to a VLAN using the following rules: • Untagged packets receive the default VLAN ID for the port on which they arrived. • Packets tagged with an ID of zero (null VLAN packets) indicate that no VLAN was specified for the packet. The device gives these packets the default VLAN ID for the port on which they arrived. • Tagged packets are assigned the VLAN corresponding to their VLAN ID.

Table 9-1: VLAN Forwarding Algorithm *(Continued)*

Step	Description
Ingress Filtering	Corero Network Devices use a packet's input port and VLAN ID, along with the following ingress filtering rules, to determine whether it should accept the packet for the learning process or discard the packet: • Discard packets with a VLAN ID that has not been configured on the device (invalid VLAN). • Accept packets if the VLAN ID is included in the port's member set. • Discard packets if the VLAN ID is not in the port's member set (unless VLAN enforcement is turned off).
Learning (Management Only)	During the learning process, a Corero Network Device observes the source MAC address of packets received on a given port and updates the bridge MAC database to associate the MAC address with the port.
Output Port Selection (Management Only)	Next, the Corero Network Device determines the list of ports to which the packet may be forwarded. It uses the following port selection rules: • The device never forwards a packet to the port on which the packet arrived. • It never forwards a packet to a port in a different bridge domain. Packets arriving at a port in the Mission Port Pair domain stay in the Mission Port Pair domain and the same for Management Bridge packets. • For a packet with a unicast MAC destination address, the device uses the associated port from the bridge MAC database. If the MAC address is in the database, the device forwards the packet to its associated port under the following conditions: • The output port is not the same as the input port. • The VLAN assigned to the packet is in the member set of VLANs for the output port, or the feature to enforce VLANs for the output port is turned off. If the MAC address is not in the database, the device floods the packet to a broadcast or multicast destination MAC address according to the following rules: • The output port must be in the same domain as the input port. • The output port must be different from the input port. • The VLAN ID assigned to the packet must be in the member set of the output port (unless VLAN enforcement is turned off).
Counter Incremented	Whenever the Corero Network Device discards the packet, it increments the VLAN Destination Filter counter for the packet's input port.
Tagging	The Corero Network Device transmits the packet in VLAN tagged format under the following conditions: • The output port is a Trunk port. • The VLAN ID is not the default VLAN ID for the output port. The device tags the packet with the VLAN ID determined during the classification phase of packet processing. However, it never transmits a null VLAN ID. If the device received the packet in tagged format, it retains the User Priority and CFI fields as they were received. If the packet originally arrived untagged, the device sets the User Priority and CFI fields to zero.

VLAN Handling for Ports with Special Roles

This section explains VLAN treatment for special ports you may configure for your Corero Network Device.

For a discussion of port roles, refer to Chapter 5, ‘‘Understanding Ports”.

VLAN Handling for Discard and Capture Ports

You can assign these ports to the Management Bridge domain. When a Discard or Capture port is part of the Management Bridge domain, management packets received on these ports are subject to all the rules that apply to the VLAN forwarding algorithm.

VLAN Handling for Mirror Ports

The port type, default VLAN, and VLAN member set properties do not apply to mirror ports. VLAN forwarding behavior depends on the traffic type as follows:

- The Corero Network Device transmits mirror operation packets without regard to VLAN egress filtering rules, and in the same format that the corresponding packet would have on the destination port of the Mission Bridge domain.
- You may choose to bridge TCP reset packets that are received on a mirror port.
- Any packet received on a mirror port that is not a TCP reset packet is dropped.

VLAN Handling of Management Entity Traffic

The management entity receives management traffic from the management ports. It also generates traffic, such as management messages, to send to the management ports.

The management entity operates using a configured VLAN ID:

Only management traffic directed to this VLAN will be received by the management entity.

Traffic can be initiated by either the client or the management entity.

- Client Initiated Traffic

  When an external client initiates communication with the management entity, a Corero Network Device determines the VLAN ID using the classification rules described in VLAN Forwarding Algorithm (page 9-4). The management entity responds using the same VLAN ID, and the tagging behavior follows the normal rules for transmitting packets, which are also described in VLAN Forwarding Algorithm (page 9-4).

- Management Entity Initiated Traffic

  When the management entity initiates communication with a client (for example, to send a Syslog report), the management entity uses the VLAN ID you assigned to the management entity. Output port selection and tagging behavior follow the normal rules for transmitting packets as described in VLAN Forwarding Algorithm (page 9-4).

Changing Management Entity VLAN ID

Corero Network Devices come with the following default VLAN settings for the Management Entity and Management Bridge:

On an IPS 5000, the default is 1.

Model 5100 and 5200 units use 4095, which is both the management entity VLAN ID and the default VLAN ID for the Management Bridge.

- Management Entity VLAN ID: **4095**
- Default VLAN ID for Management Bridge: **4095**

When you are using the management application for a specific Corero Network Device, you can also modify the VLAN port type.

If needed, you can modify these settings to a non-standard configuration.

CAUTION

If you need to change these VLAN settings, you must do so carefully so you do not block connectivity between the Corero Network Device and your management host.

NOTE

For a discussion of how the Corero Network Device's Management Entity handles VLAN traffic, refer to VLAN Handling of Management Entity Traffic (page 9-7).

If you need to change these default settings, you must follow the procedure below carefully so that you do not unintentionally cut the Corero Network Device off from your management host. If you do cut yourself off, refer to Recovering Connectivity When You Accidentally Lose Management Access (page C-2).

To view and modify the VLAN ID for a device-specific management port from the Corero Network Device management application:

1. From the Navigation Tree select: Configure System > Ports > Ports. The Ports dialog box displays.
2. Select a Management Port and click Edit.
3. Modify the default VLAN ID as desired.

   NOTE

   For information on how to modify other port settings on the Edit Port Settings dialog box, see Chapter 6, ‘‘Viewing and Configuring Ports”.

4. When finished, click OK.

Managing One-Arm Routing

One-arm routing refers to a configuration where a router forwards traffic between multiple subnets/VLANs that are received over a single physical LAN segment. The router is referred to as a one-armed router or a stub router. This type of configuration is common in Service Message Block (SMB) networks that are partitioned into subnets/VLANs where 80/20 traffic rule applies. This means that 80% of the traffic is intra-subnet/VLAN and the remaining 20% is the inter-subnet/VLAN traffic.

The primary difference between one-arm routing and other types of routing is that, with one-arm routing, each packet sent in either direction traverse the same Corero Network Device *twice*.

Layer 2 switches handle the intra-subnet traffic and the smaller (20%) inter-subnet traffic is forwarded to the one-armed router. The router then updates the Ethernet headers and forwards the packets back out the same physical LAN segment they were received on but to a different VLAN/subnet.

It is not uncommon to see a Corero Network Device deployed in a one-arm routing configuration. In this configuration, the device is generally placed between a Layer 2 switch, that provides the fan-out, and a router that provides Internet connectivity, as well as inter-subnet connectivity for systems on the inside of the enterprise.

The Client and Server are on different subnets and on the inside of an enterprise Intranet. The device is placed between the one-armed router (External Link of the device) and the Layer 2 switch (Internal Link of the device) to provide connectivity to the Clients and Servers on the inside.

When one-arm routing is enabled, the Corero Network Device identifies when packets for a given flow are traversing the device more than once. If this occurs, this setting enables the device to function properly in this configuration.

This is accomplished by assigning locations to the Client and Server. When a new flow is established, the Client location is noted (Internal versus External). The device assigns the opposite location for the Server (External, if client was internal and vice versa). The device inspects traffic and applies its policies only when the packet is received on a port that matches the location of the sending system.

The traffic flow for the configuration in the One-Arm Routing graphic is shown below. In this case, the Client is assigned to the Internal location, which means the Server is assigned to the External location. The instance when packet inspection occurs (that is, Internal versus External) is indicated in bold.

In the example configuration shown in Figure 9-2, Client to Server traffic travels as follows:

Client > Layer 2 Switch > Corero Network Device (Internal) > Router > Corero Network Device (External) > Layer 2 Switch > Server

Server to Client traffic travels as follows:

Server > Layer 2 Switch > Corero Network Device (Internal) > Corero Network Device External > Layer 2 Switch > Client

NOTE

This one-arm routing implementation requires that the Corero Network Device be configured for Port-Pair forwarding.

Figure 9-2: One-Arm Routing Configuration

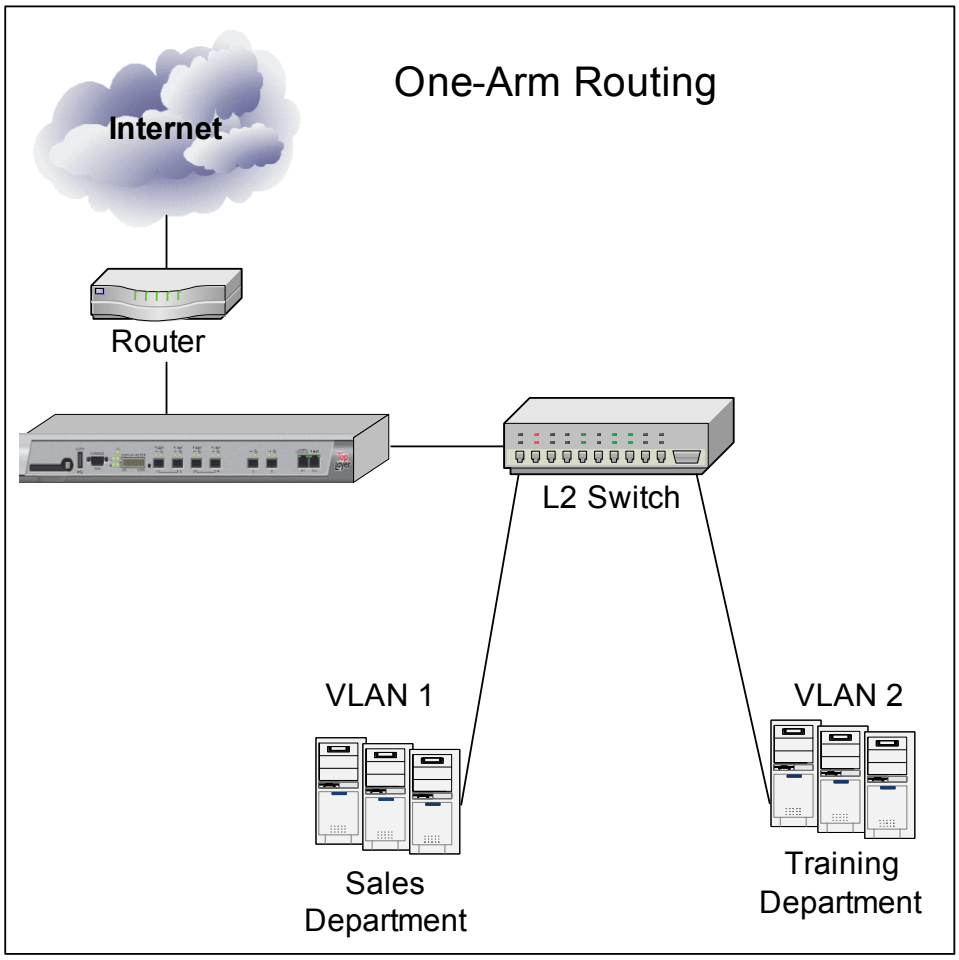

One-Arm Routing Considerations

If you are implementing one-arm routing, consider the following:

- Spoof Checks (page 9-10)
- Firewall Policies (page 9-11)
- Delayed Packet Inspection (page 9-11)

Spoof Checks

Spoof checks (rule: tln-001001 Connection from Spoofed IP Address) are performed as follows:

1. At flow set up time to ensure that the mission port type (Internal versus External) for the received packet matches the location configured in the IP Range table.
2. Once the flow is established to ensure that the Client and Server packets are received on ports that match their configured location.

When one-arm routing is enabled, the first part of spoof checking (item 1 above) is still performed, but the second part (item 2) is no longer performed as it conflicts with the behavior of one-arm routed packets.

Firewall Policies

Even though packets traverse the Corero Network Device more than once in a one-arm routed configuration, there is only one flow setup event that occurs. Therefore, if there are Firewall policies entries that pertain to specific direction (for example, Inbound versus Outbound), these policy entries are applied when the packet arrives the first time.

In the case of a SYN packet, the Firewall policy is applied the first time and if the same SYN packet arrives a second time (because it was one-arm routed), these policies are not applied again. This statement holds true even when one-arm routing is not enabled.

Delayed Packet Inspection

The implementation of one-arm routing support may delay the inspection of packets on the Corero Network Device until the packet appears the second time. In the example configuration, all Server-to-Client packets are inspected only on their return from the router and not the first time they appear from the Layer 2 switch. In some cases, it may appear that packets that would have failed some checks on the device are being forwarded incorrectly. The checks will be applied when the packets return from the router (hence, delayed) and acted on as configured by the user.

Configuring One-Arm Routing

To enable or disable one-arm routing:

1. Choose Configure System > Advanced System Config > One-Arm Routing from the navigation tree. The One-Arm Routing dialog box displays.
2. Do one of the following:
   - To enable one-arm routing, select (check) the Enable One-Arm Routing check box.
   - To disable one-arm routing, clear (uncheck) the Enable One-Arm Routing check box.
3. When you enable one-arm routing, you need to disable spoof checks for your IP address definitions. To do this:
   a. Select Manage Security > Security Policies from the navigation tree, then click the Host Groups tab.
   b. Select the Host Group that contains the relevant IP addresses.
   c. Under the Host Group Membership area, click Add.
   d. Add the address range associated with the one-arm router.
   e. Under Spoof Checks, select Disable.
   f. When finished, click OK.
4. Save your changes by clicking the Save Configuration toolbar button.

Chapter 10
ProtectionCluster Configuration

For those Corero Network Device models that provide High Availability (HA) ports, multiple Corero Network Devices (IPS or DDS Units) can be configured to work together providing redundancy. When two or more devices are configured this way, their combination is called a ProtectionCluster.

NOTE

IPS 5500 Model 75 EC Corero Network Devices do not support high availability.

In addition to providing redundancy, A ProtectionCluster can also be configured to increase processing power, improve throughput of inspected packets, and flow balancing. ProtectionClusters also provide additional capacity which enables the Corero Network Devices to share the intense processing required for the deep and stateful protocol analysis necessary to detect attempted exploits of vulnerabilities.

In order to ensure proper traffic flow, ProtectionCluster operation is designed to pass all packets in a given flow in and out of the same Corero Network Device.

This chapter contains the following sections:

ProtectionCluster Overview

You can configure two or more Corero Network Devices to work together to provide network hardware redundancy and higher bandwidth throughput. This type of configuration is called a ProtectionCluster.

A ProtectionCluster configuration connects multiple Corero Network Devices together using selective 10/100/1000 ports on 5100-Series hardware devices or 10GbE ports on 5200-Series hardware devices. An HA configuration also enables Corero Network Devices to provide the intense processing necessary for deep and stateful protocol analysis of high-speed and high-volume traffic.

A ProtectionCluster must be comprised of two or more of the exact same Corero Network Device type and model, although whether the model is an EC or ES can vary within a ProtectionCluster. For example, you could create a ProtectionCluster out of three DDS 5500 1000ES Units, or you could create a ProtectionCluster out of two IPS 5500 500ES Units, but you could not combine the two different device types in a single cluster.

You can configure a two-device ProtectionCluster without an IPS Controller. An IPS Controller is required for ProtectionClusters with 3 or more members, and suggested for a two-member cluster because the IPS Controller management interface simplifies the process of synchronizing settings on ProtectionCluster members.

NOTE

Refer to the Configuration and Management Guide for your device for additional information on ProtectionCluster configurations, The location of HA ports, and installation diagrams.

Dual-Device High Availability ProtectionClusters

When two Corero Network Devices are connected in a highly available ProtectionCluster configuration it is known as a high availability (HA) configuration. In addition to providing automatic failover, these links provide higher bandwidth for flow rebalancing in redundant configurations. During normal operation, all the Corero Network Devices in the configuration contribute to analyzing and passing network traffic. During a failure, the operational devices within the HA configuration handle the traffic previously handled by the failed device. A 2-unit High Availability ProtectionCluster uses ports 5 and 6 (5200-Series hardware), ports 5-8 (5100-Series or 5000-Series hardware) to connect multiple devices together.

High Availability Ports

Figure 10-1 shows the four ports in the 5100-Series hardware that can be used to carry the communication needed for one of the two Corero Network Devices to instantly take over if there is a failure of the other unit. Devices also use these ports to ensure that a single device sees all packets associated with a given network connection. 5100-Series units require that you interconnect four HA ports between ProtectionCluster members.

NOTE

High Availability (HA) ports are not available on Model 75EC IPS Units.

Figure 10-1: 5100-Series High Availability Ports: 5 Through 8

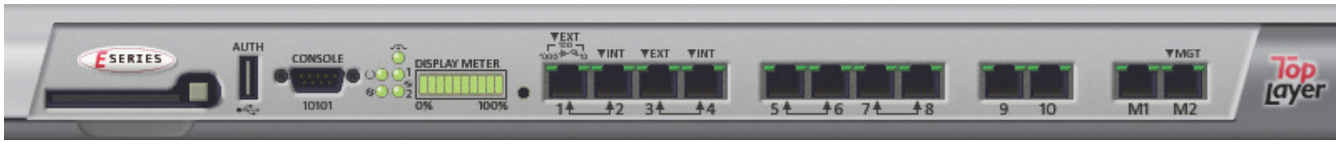

5200-Series hardware requires that you interconnect two HA ports between ProtectionCluster members. Ports 5 and 6 are dedicated to high availability, as shown in Figure 10-2.

Figure 10-2: 5200-Series High Availability Ports: 5 and 6

High Availability ProtectionCluster Configurations

Corero Network Devices are designed for maximum compatibility with your existing network configurations. If you have currently configured your network for redundant operations, you can insert Devices into this configuration.

Figure 10-3 shows a typical customer's existing high availability, fully redundant configuration. In this configuration, all network traffic goes through Router A or Router B. If either router fails, all network traffic goes through the other router.

Figure 10-3: Existing Customer High Availability Configuration

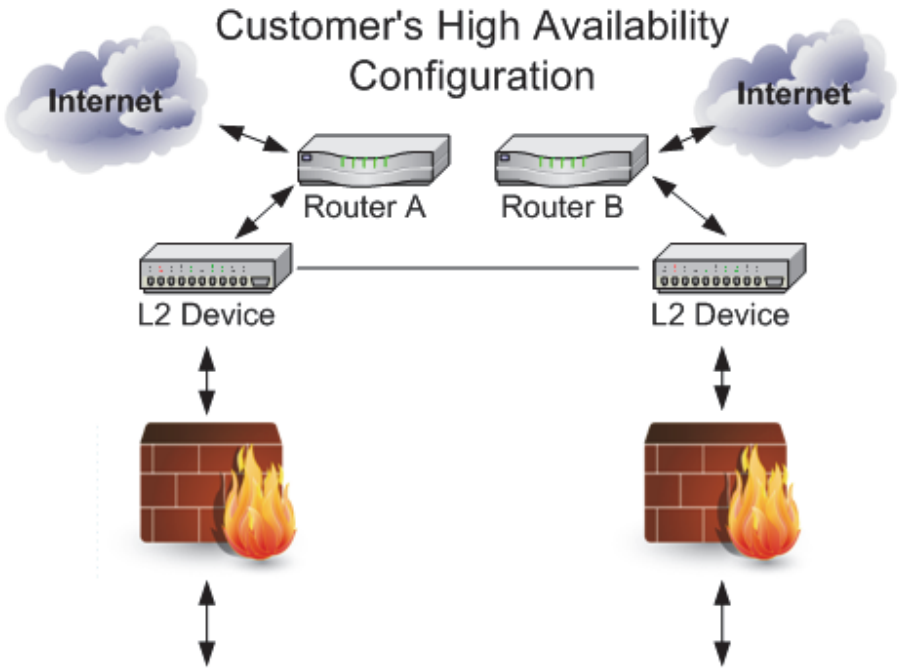

Figure 10-4 depicts two Corero Network Devices inserted into an existing redundant network configuration. The devices are deployed inline in the network.

When there is no failure, both units share the network load, thereby increasing bandwidth. If there is a failure in one Corero Network Device, the second device takes the entire load, ensuring continued network operation. If another network device in the configuration fails, the remaining workload can continue to be shared across both operational Devices.

The lines between the Corero Network Devices in Figure 10-4 represent the redundant, high availability links that the devices use to communicate their operational status, share state information, and share the task of performing packet inspection on large volumes of network traffic.

Figure 10-4: Dual Inline High Availability Configuration

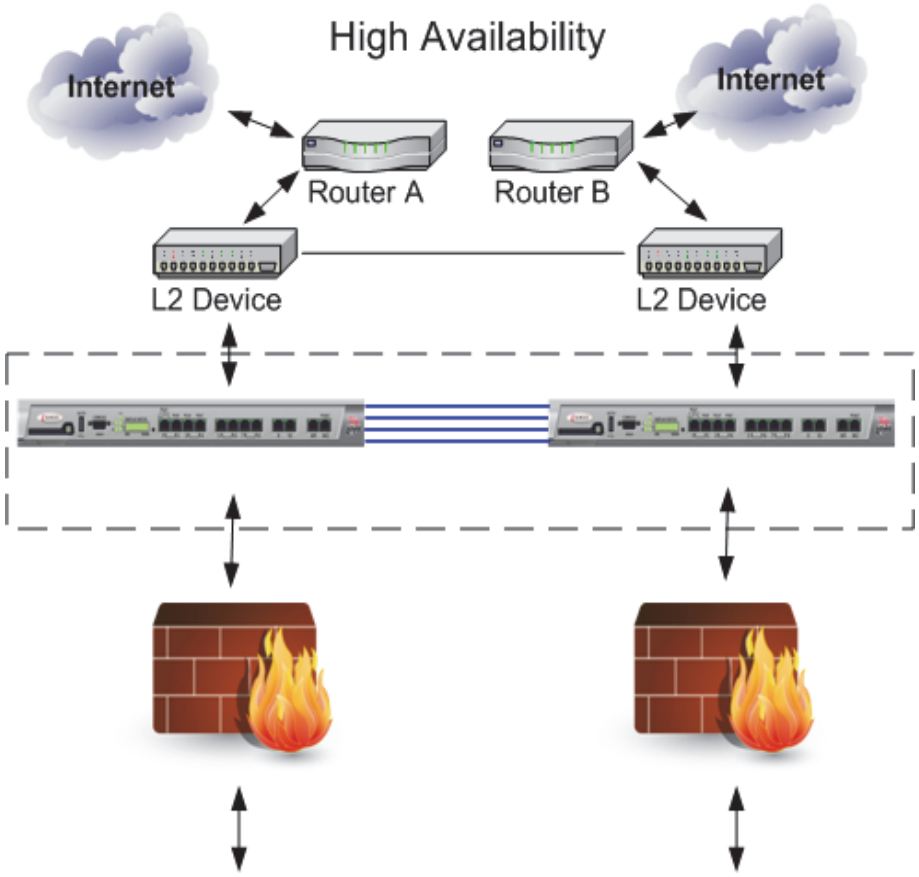

High Capacity ProtectionCluster Configurations

You can also use a ProtectionCluster to increase processing power and inspection throughput on your network. In the example shown in Figure 10-5, a Model 2000 ES and a Model 2000 ESL are configured as inline peers. The 2000 ES, called the unit in this configuration, provides the mission ports where traffic enters and exits the ProtectionCluster. The 2000 ESL, called the leaf unit, provides additional processing capacity.

Note that you can connect up to 8 Model 2000 ES units in a ProtectionCluster configuration.

The 2000 ES and the 2000 ESL have an identical number of ports, but on each model the ports have different purposes.

- The 2000 ES has four mission ports (numbered 1-4).
- The 2000 ESL has four dedicated Corero Network Device HA Interconnect switch ports (numbered S1-S4).

This high capacity configuration allows traffic entering the main unit to be processed on the inline peer unit (also called a leaf node). Traffic processing is distributed across both units, while continuing to use port pair forwarding, which is required on ProtectionCluster units.

The increased capacity provided by the leaf node can be used to:

- Build in additional capacity for planned near-term use.
- Accommodate higher current levels of traffic than a single Corero Network Device can inspect by itself.
- Provide additional capacity to increase network resiliency against specific types of attacks, such as Distributed Denial of Service (DDoS) attacks.

Figure 10-5: High Capacity 2000 ES and ESL Configuration

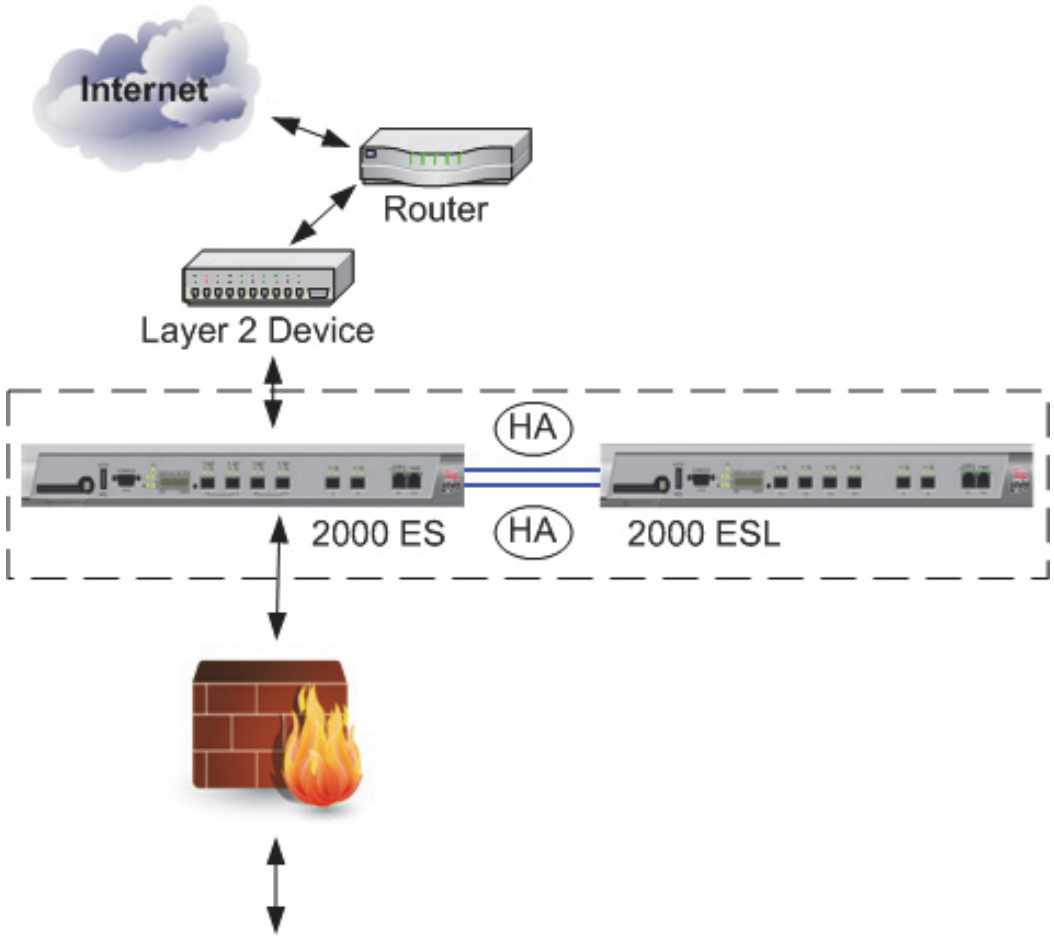

If you want to implement a dual-inline configuration for high capacity, and you have an existing Corero Network Device 5200 Model 2000 ES, you can either add a Model 2000 ES or a Model 2000 ESL to your existing unit. However, if you are installing 5200 units for the first time, you can purchase the 5200 Model 2400 ES, which is a single-box solution comprised of a 2000 ES and a 2000ESL contained in the same chassis.

You can interconnect a Model 2000 ES and a Model 2000 ESL. When you do this:

- The four HA Interconnect ports on the Model 2000 ESL. These switch ports are for dedicated use interconnecting the HA ports on multiple 5500 Model 2000 ES units.
- How a 5200 Model 2000 ES and a Model 2000 ESL would be interconnected. If you are using more than one Model 2000 ES in a ProtectionCluster, you would use the HA Interconnect ports on the 2000 ESL to connect the HA ports of the additional units.

Figure 10-6 shows the connections required between a Model 2000 ES IPS Unit and a Model 2000 ESL IPS Unit.

Figure 10-6: IPS 5500 Model 2000 ES and Model 2000 ESL Interconnections

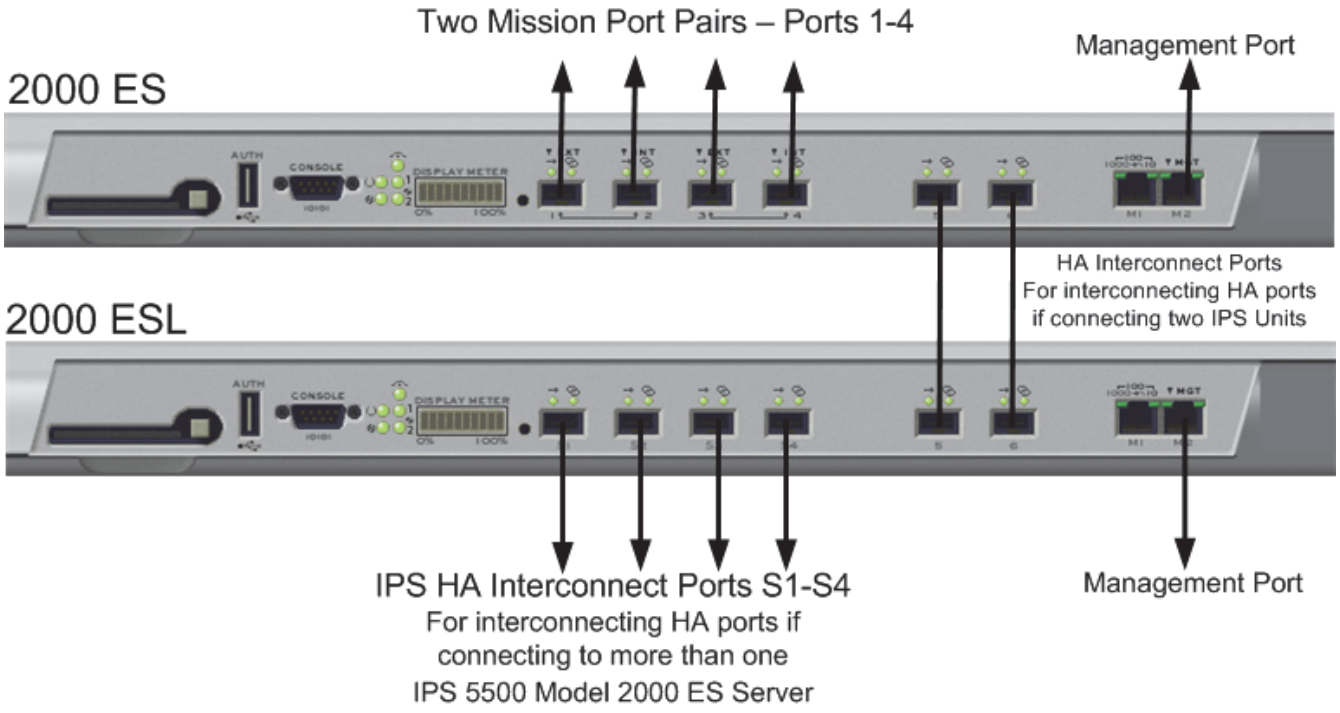

Figure 10-7 shows how two IPS 5500 Model 2400 ES units could be interconnected in a ProtectionCluster configuration. You could also connect a single Model 2400 ES unit an additional 2000 ES unit.

Figure 10-7: IPS 5500 2400 ES to 2400 ES Interconnections

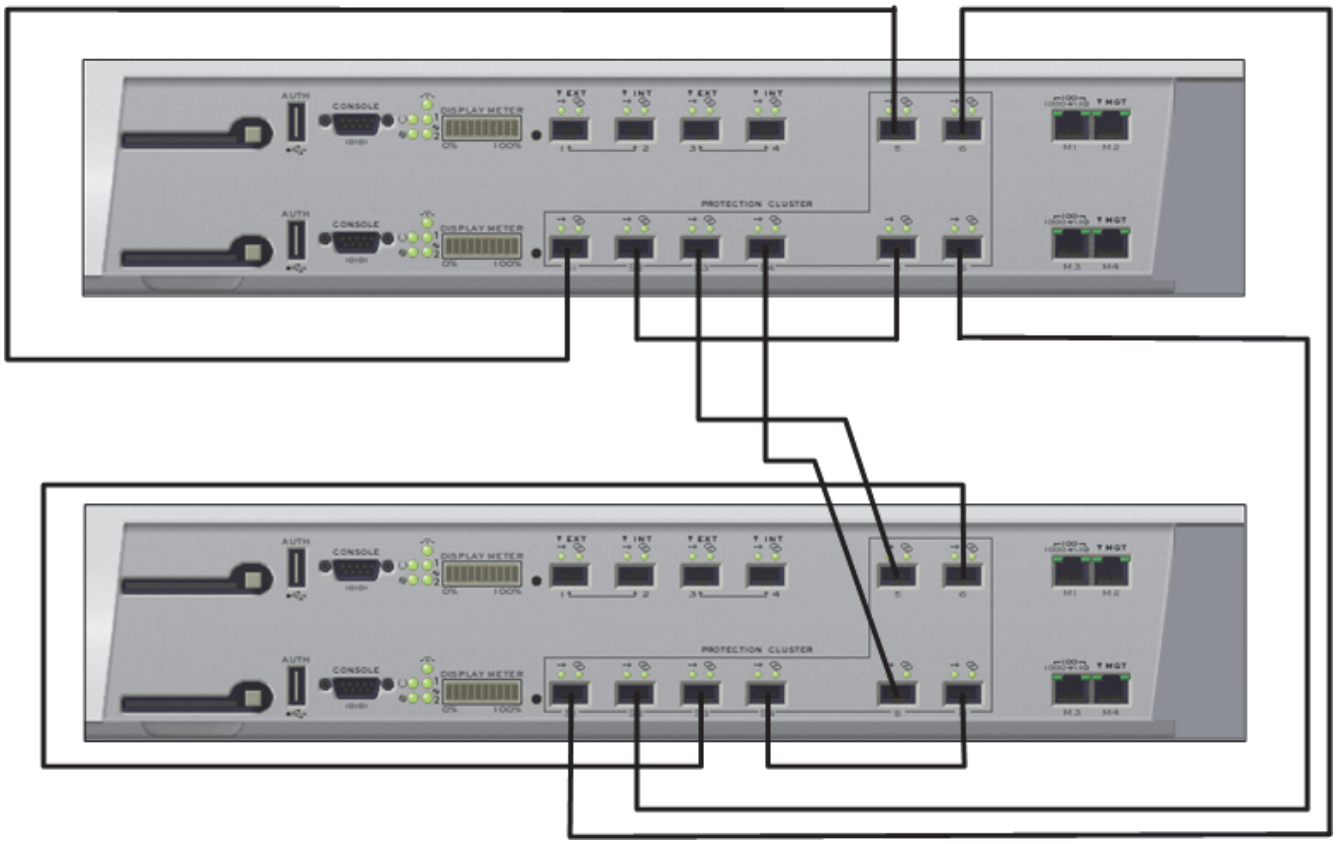

Figure 10-8 shows how a high-throughput ProtectionCluster comprised of two older IPS 5500 Model 2400 ES units could be used for perimeter defense in a customer network. You could use the same configuration for two DDS 5500 Model 2400 ES units.

Figure 10-8: High-Throughput Perimeter Defense ProtectionCluster Configuration

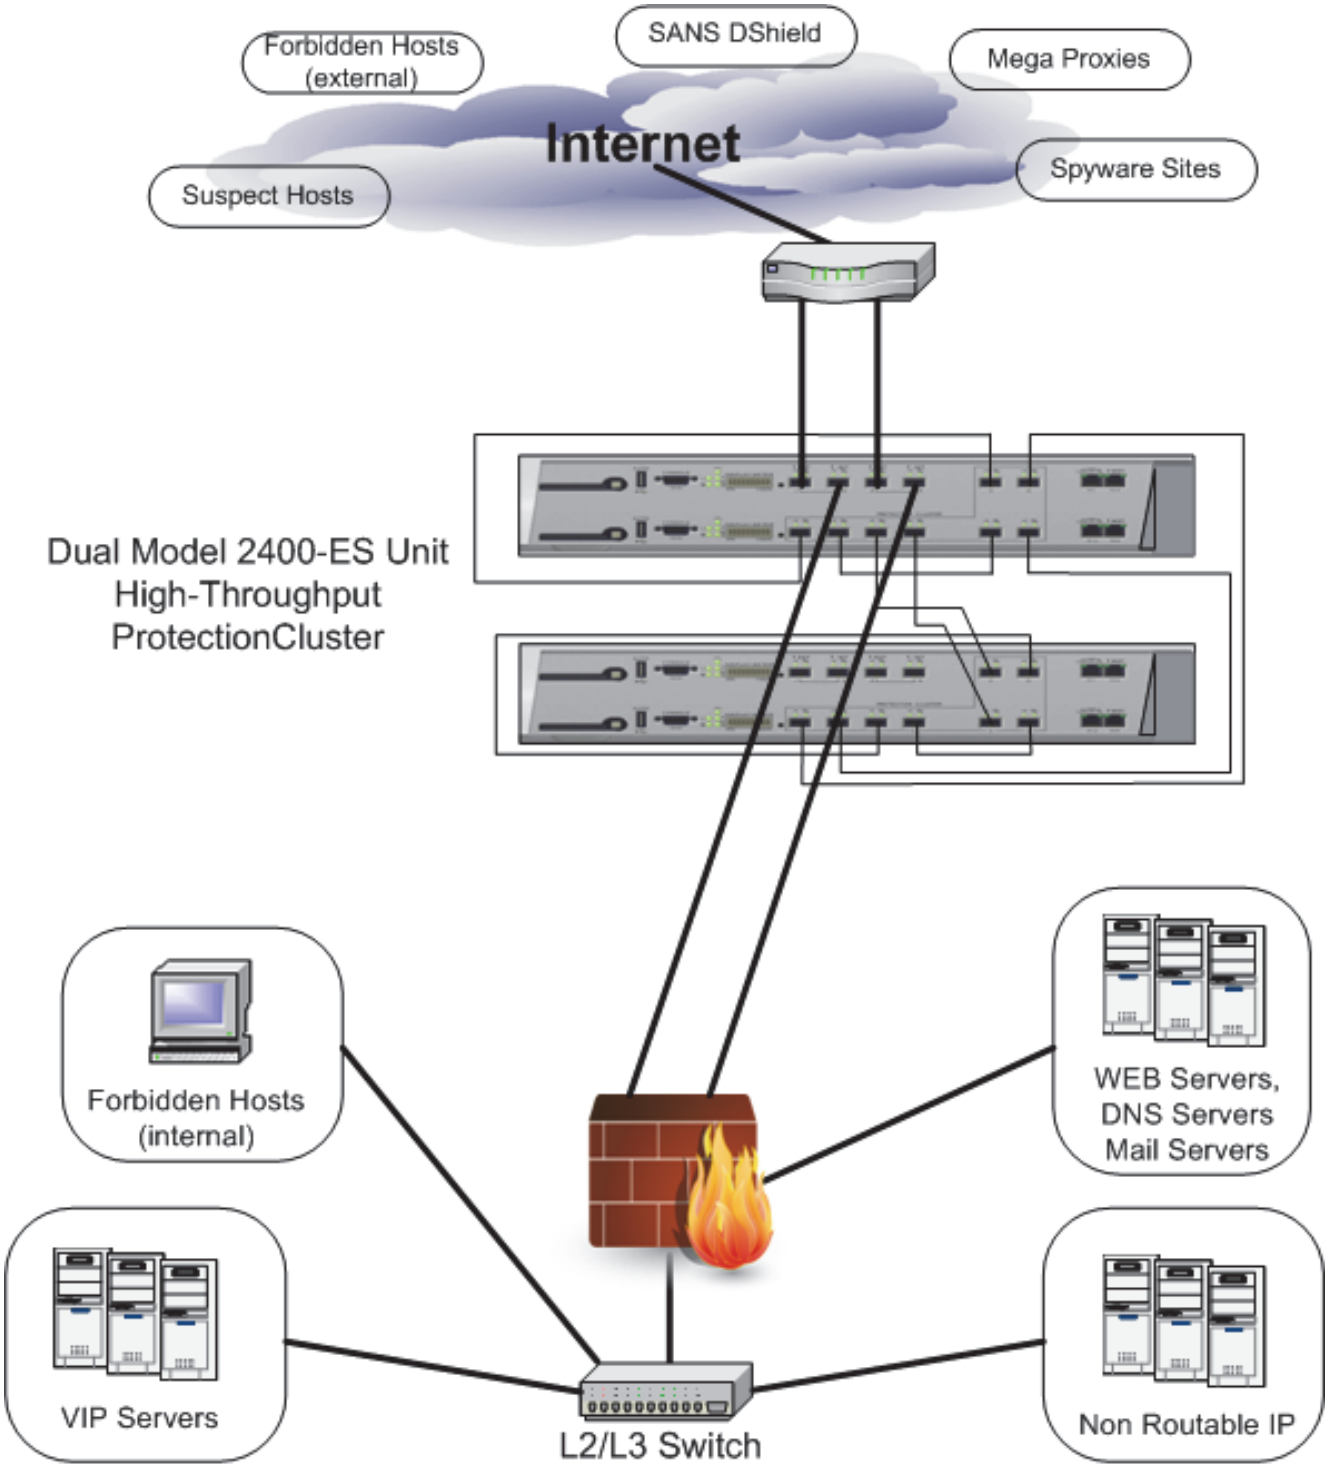

ProtectionCluster Planning and Preparation

When preparing to configure a ProtectionCluster, consider the following:

- All members of a ProtectionCluster must be the same Corero product (IPS or DDS) and model (5200, 5100, 0r 5000), running the same software version, and be members of the same Policy Group.
- ProtectionCluster configurations of between two and eight Corero Network Devices are supported when using the IPS Controller.
- You cannot modify the members comprising a ProtectionCluster. If you want to change the constituency of a ProtectionCluster, you must delete and recreate it with the updated membership.
- If the ProtectionCluster contains more than two Corero Network Devices, then a dedicated switched interconnect must be provided between the HA ports of all the cluster members. The four HA Interconnect ports built into the Model 2000 ESL are designed specifically for this purpose, meeting the specified switching requirements.

  This switch must provide a jumbo-frame-tolerant L2 path between all Corero Network Devices in the ProtectionCluster, and must provide non-blocking interconnect performance, adequately buffered for large bursts.
- You must use Port Pair Forwarding traffic mode on all Corero Network Device ProtectionCluster members.
- If the network attached to the Mission ports requires a Port Tracking feature to enable high availability, you must also enable Port Tracking mode from the Getting Started Wizard.
- If you use a Mirror, Discard, or Capture port on more than one member of a ProtectionCluster, designate the same port on each of the cluster members for consistency.
- Both the forward and return packets for all traffic that you wish to protect must be seen by one of the Corero Network Devices. To ensure this happens, your network configuration must not contain loops that allow packets to bypass the devices completely.
- The L2 switch port and the firewall port on a Corero Network Device must have matching settings or entering and exiting Bypass Mode can cause link problems. The ports between the L2 switch and the firewall should be configured and tested as a direct connection to one another prior to insertion of the Corero Network Device between them. This way, the Corero Network Devices will match the settings of the adjacent devices.
- In a ProtectionCluster environment where asymmetric network traffic is possible, all Corero Network Devices should be in either Always Bypass or Never Bypass mode. This is to ensure that when one device is rebooting, the other device(s) in the ProtectionCluster will not see partial flows. The exception to this is when all mission ports are on one device and the other device(s) in the Protection Cluster are used as leaf nodes. In this configuration, the bypass mode can be safely set to Bypass During System Reset.
- There is a maximum link distance between Corero Network Devices at different locations, as shown in Table 10-1

Table 10-1: Maximum Fiber Link Distance Between Corero Network Devices

Fiber Core Diameter	Fiber Bandwidth	Link Distance
62.5 um	160 MHz*Km	220 Meters
62.5 um	200 MHz*Km	275 Meters
50 um	400 MHz*Km	500 Meters
50 um	500 MHz*Km	550 Meters

Creating a Dual-Device ProtectionCluster Without an IPS Controller

Corero strongly recommends that you configure and manage dual-device ProtectionClusters using an IPS Controller. In fact, an IPS Controller is required to create and manage ProtectionClusters with three or more members.

However, the management application for your Corero Network Devices does enable you to create a two-member ProtectionCluster without the use of an IPS Controller.

NOTE

High Availability (HA) ports are not available on Model 75EC IPS Units.

To create a ProtectionCluster containing two Corero Network Devices without the IPS Controller.

1. Verify that the Corero Network Devices you want to include in the ProtectionCluster meet the following prerequisite criteria:
   - Ensure all of the devices you want to include in the ProtectionCluster are the same type (IPS or DDS Unit) and the same model.
   - Ensure all devices you want to include in the ProtectionCluster have been installed and fully configured. They must be up and available for proper cluster configuration.
   - Ensure all devices to be clustered are running the same version of software. You cannot create a ProtectionCluster from devices running different software versions.
   - Physically connect the high-availability (HA) ports between the two devices. For instructions on how to do this, refer to the Hardware Installation Guide for your cluster devices.

   NOTE

   For the rest of this procedure will refer to the members of the ProtectionCluster as Unit A and Unit B.

2. For Unit B, launch the management application and click the System Information toolbar button. While viewing the System Information dialog box, write down the unique MAC address for Unit B.
3. For Unit A, launch the management application, then choose Configure System > High Availability from the Navigation Tree. The High Availability dialog box displays.
4. On the High Availability dialog box, select the Enable High Availability check box, then and enter the MAC address for Unit B, which you noted earlier.

   NOTE

   You do not need to configure the High Availability and MAC Address information for Unit A from Unit B, as you will, instead, be transferring this information using Unit A's configuration files as described in the following steps.

5. Run the Getting Started wizard on Unit A. As you complete the wizard, configure Unit A with the following settings on the Mission Port Pair Settings window:
   - Ensure that you have selected Port Pair Forwarding. This setting is required.
   - If devices connected to the mission ports require port tracking for high availability failover operation, select Port Tracking.
   - Ensure the bypass feature is set to Never Bypass.
6. When finished, click Save to save your changes to the configuration files on Unit A.

7. On Unit A, from the management application, choose Manage System > Configuration Files from the Navigation Tree. The Configuration Files dialog box displays.
8. On the Configuration Files dialog box, select the HA configuration file, then click Download. The file's contents display in a browser window.
9. Choose File > Save As from the browser window's menu bar and save the contents to a file on your management station. Do not modify the file name.
10. On Unit B, from the management application, choose Manage System > Configuration Files from the Navigation Tree. The Configuration Files dialog box displays.
11. On the Configuration Files dialog box, select Upload. The Upload Configuration File dialog box displays.
12. Browse to the location of the HA file you saved, then click Upload.

    NOTE

    You may receive a Security Advisory message generated by the Java system asking if you will allow this action. If this message displays, allow it.

13. When upload is complete, the file's contents are automatically validated. Once validation is complete, the device automatically activates the uploaded file and begins using it.
14. After you change the High Availability configuration on a Corero Network Device, you must reboot the unit.
15. On either member of the ProtectionCluster, verify that the High Availability is properly configured by clicking the Front Panel toolbar button to launch the Front Panel view. High Availability status information is displayed at the bottom of the Front Panel display

Chapter 11
About Security Policies

This chapter provides an overview of concepts and terminology used to create individual security policies for Corero Network Devices.

Security policies enable you to allow or block traffic depending on its characteristics. You can craft security policies that permit necessary traffic, and block unnecessary traffic, not only to reduce traffic volume, but also to prevent attacks. You can block specific traffic types in different ways. For example, if you wanted to block Instant Messaging traffic, you could block the services associated with that type of traffic, you could block traffic of a particular type using a specific port, or you could block traffic based on a known host name from which the traffic comes.

This chapter includes the following topics:

Overview of Security Policies

Security policies are the logical constructions that guide a Corero Network Device's security decisions as it examines traffic flows for the device subsystems to meet the needs of your network.

NOTE

In addition to setting up a security policy to handle network traffic, there is an additional feature which allows for shunning IP addresses that are deemed suspect. See Chapter 19, ''Security Management and Monitoring" for details of this feature.

The three types of policies for subsystems guide the types of security checks listed in Table 11-1.

A security policy is comprised of one or more firewall policies, one or more IPS policies, and one or more rate-based policies

Table 11-1: Policy Types

Policy Type	Description
Firewall Policy	Provides classic firewall blocking for traffic, based on IP addresses, Layer 4 ports, and segments (port pairs).
IPS Policy	Provides the following: • Protocol validation • File attachment content validation (only applicable to IPS E-Series models) • Attack Signatures • Acceptable application-usage policies
Rate-Based Policy	Protects resources from overuse by legitimate users, as well as abusive denial-of-service attackers. Provides limits for: • Client requests • Connections • SYN Flood controls • Application rate limiting

Figure 11-1 shows how policies are used to mitigate traffic.

Figure 11-1: Three Dimensional Protection

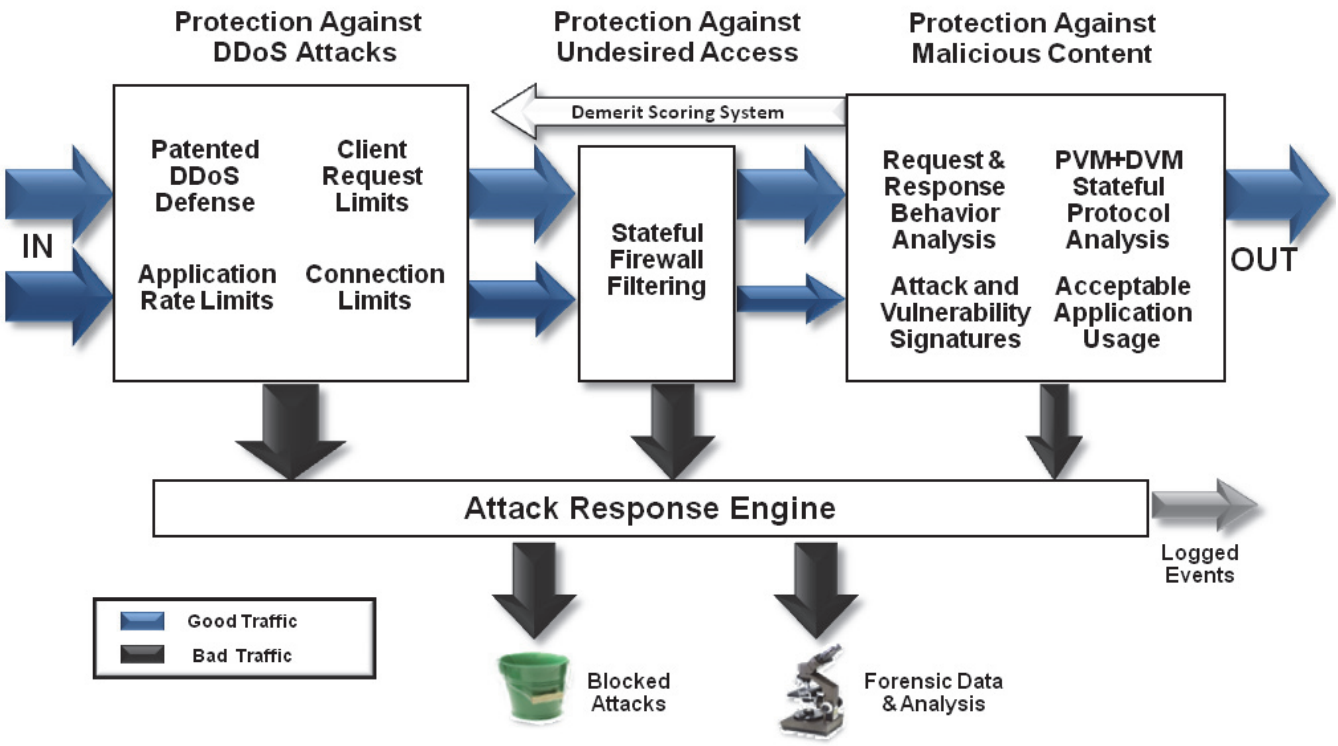

Elements of a Security Policy

A Corero Network Device security policy is a logical definition that says: if a given traffic flow meets the following conditions, treat it in the specified manner.

Table 11-2 describes the parts that comprise a security policy.

Table 11-2: Security Policy Components

Component	Description
Traffic Flow Condition	The traffic flow is specified using one or more of the following: • The segment (port) or segments (ports) on which the traffic arrived. • The server or servers (host group or groups) from which the traffic came. • The protocol or protocols the traffic uses. • The client or clients (host group or groups) to which the traffic is targeted.
Firewall Action	The firewall action is the response to the traffic: Allow, Drop, or Reject.
System Response	The system response specifies how the system responds to the traffic. This could include logging an event, or copying the traffic to a discard port.

When you view the FW+IPS Policies tab on the Configure Security Policies dialog box (Figure 11-2), each column represents a different component of the policy. These components are grouped according to Conditions, Firewall Treatment, and IPS Treatment (rule sets and rule dispositions).

Figure 11-2: FW+IPS Tab Columns Indicate Policy Components

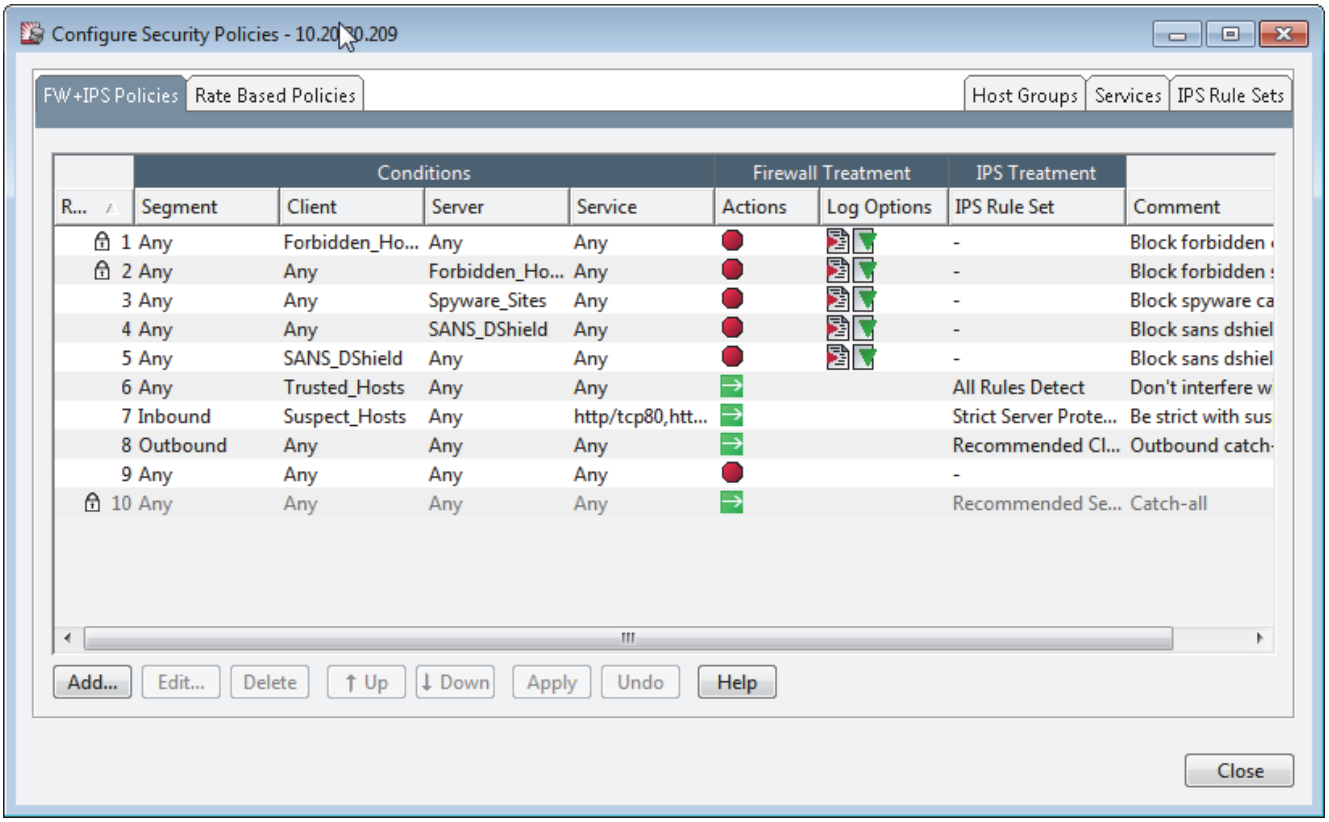

The following sections describe each component of the Corero Network Device's security policies:

Segments

Corero Network Devices are designed to run in port pair forwarding mode. This operating mode associates two mission ports together as a port pair. When traffic is received by the device through one port, the traffic will be sent out through the other port in the pair (assuming the traffic is deemed to be safe).

A port pair is called a segment. You can define a security policy (for example, a firewall policy) that only applies to a single segment or to a set of segments.

When you run the Getting Started wizard, the device takes your defined set of port requirements and selects ports that match those requirements.

The Select Segments window (available when you are defining Firewall + security policies), displays the available sets of port pairs, or segments, as shown in Figure 11-3.

Figure 11-3: Select Segments Dialog Box

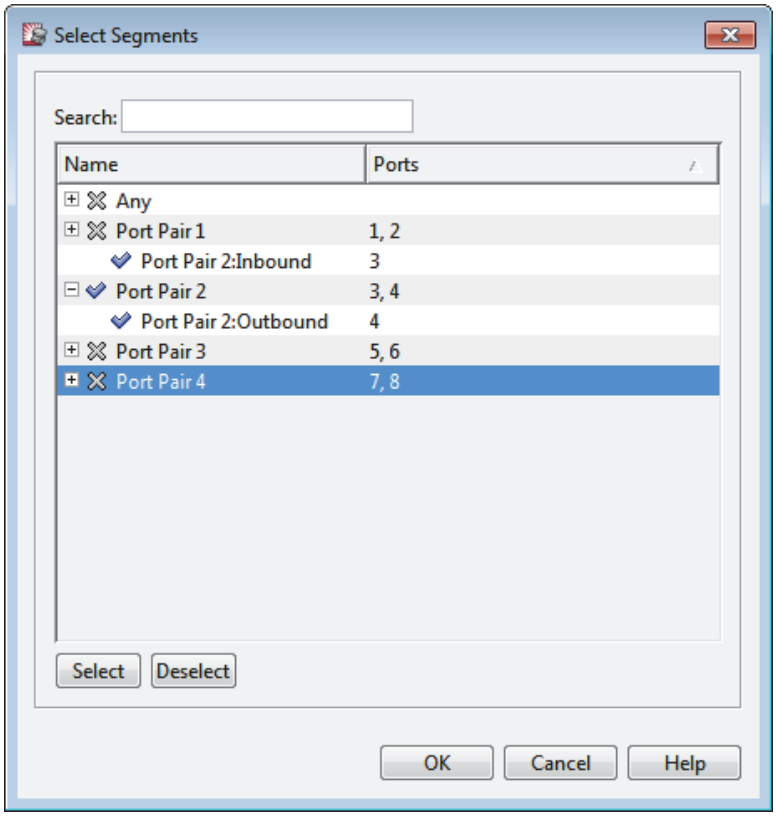

Host Groups

A host group is a named set of IP address ranges. A given host group may act as both clients and servers. A client host group is a host group whose members are initiating connections. A server host group is a host group whose members receive connections.

If you want host group traffic to be treated differently (depending on whether traffic is going to or coming from host group members), you will need to create more than one security policy entry for the same host group and apply specific conditions and treatment based on whether the host group is initiating or responding to traffic requests. Any given IP address range (including a range containing a single address) may only be in one host group.

NOTE

For a discussion of security policy entries refer to Elements of a Security Policy (page 11-4).

This product comes with a set of predefined host groups which you can modify. You can also add your own groups.

Figure 11-4 shows how some of the default host groups could be used with an IPS installed for Internet perimeter defense. Figure 11-9 (Default FW + IPS Policy) shows the default policy entries that define these groups.

Figure 11-4: Host Groups Used in a Perimeter Defense Implementation

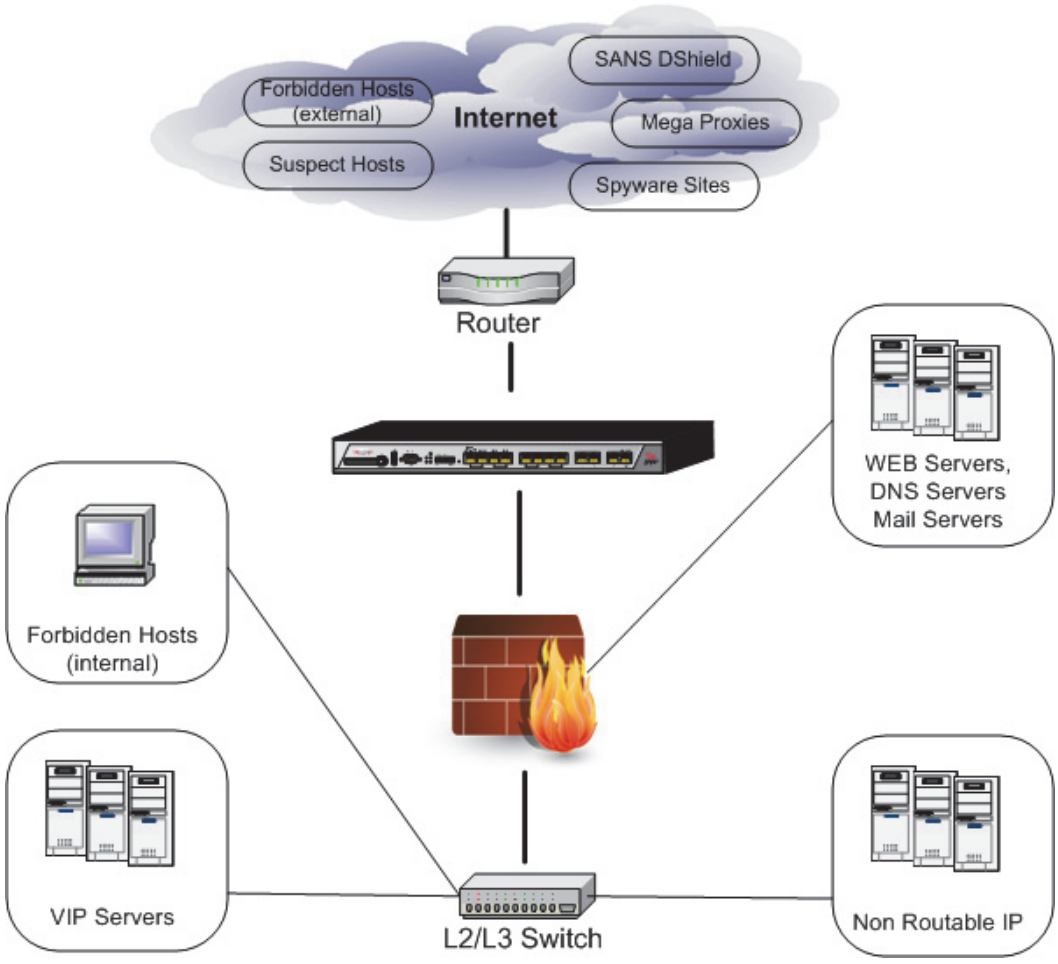

Named IP Address Ranges

Named IP address ranges are used to simplify host group definition. A named IP address range can be a single IP address or a set of addresses. You can create a named IP address range and specify the following attributes for it:

- Associated host group
- IP address or range specification (individual IP address, IP address range, subnet)
- Type of spoof check treatment
- Source and/or destination filters
- Range attributes such as subnet or broadcast

Spoof checking is used to identify attacks where hosts modify the IP address to imitate a different (internal or external) IP address. You can instruct the Corero Network Device to check the type of port (internal or external) on which traffic with this IP address arrives. You can disable spoof checking, or you can specify whether you want to allow traffic from an IP address in this range only if it appears on the internal port in a port pair, or only if it appears on the external port.

IP Address Specification Considerations:

When specifying IP address ranges, *it is important that you consider the following*:

- An IP address can belong to one, and only one, host group.
- You can, intentionally or unintentionally, specify IP addresses or address ranges for multiple host groups that match a single IP address.
- When you create, modify or delete a host group, this can result in one or more IP addresses being moved from that host group to a different host group.
- The management application does not inform you when creating, editing, or deleting a host group results in the reassignment of an IP address from one host group to another.

NOTE

To determine the current host group associated with an IP address, perform an IP Query on that address. For more information, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).

- If an IP address matches the specifications for more than one host group, there is an order of precedence the system follows in order to assign the address. The order of precedence goes from most specific to least specific, as described in Table 11-3.
- If an IP address matches multiple specifications that have the same order of precedence (for example, more than on specified address range), it will be assigned to the host group associated with the most recently defined IP address range.

Table 11-3: IP Address Host Group Assignment Order of Precedence

Order	Description
First	The address is specified as an individual IP address in the host group definition.
Second	The address falls into a specified address range.
Third	The address falls into a specified subnet.

A named IP address range can also be a subset of another range, as long as it falls completely within the larger range. However, by creating a subset address range, you remove those addresses from the larger range (creating a hole in the larger address range). (A singleton IP address is considered a range of one address.)

For example:

1. If you had an initial range (Range A) that contained the addresses from 10.20.30.40 to 10.20.30.255.
2. Then you created a second range (Range B) that contained the addresses from 10.20.30.100 to 10.20.30.125.
3. Range A would automatically be modified to include 10.20.30.40 to 10.20.30.99, and 10.20.30.126 to 10.20.30.255.

Services

A service is a policy element that identifies a protocol or set of protocols to a Corero Network Device, for example, HTTP, DNS, or FTP. The device includes many definitions of valid and invalid services, and you can add custom definitions. You can specify handling properties (such as connection time-out and discard priority) for services. All default services are associated with Any server. When you define a Firewall + IPS security policy, you can identify a set of services and an associate them with a specific server group or groups.

You can specify detailed information for each service. For example, you can protect a server by only opening those inbound ports that relate to the application the server is providing (for example, port 80 for HTTP). Some protocols have both fixed control ports and additional auxiliary ports. For this type of server, only allow traffic from those fixed ports in the required range, and block all unused ports. For additional protection, you can also block the auxiliary port range on other servers that are not running the application that requires them.

You can also define new services based on new protocols, and you can also define services for protocols that run on non-standard ports or port ranges. You can use these definitions to restrict or limit access to specific network services. Refer to the online help for more information.

Rules

Each Corero Network Device contains hundreds of rules that it uses to check whether a given flow of traffic is acceptable or not. A rule may be based on packet checks (for example, Illegal ICMP Header) or protocol checks (for example, HTTP Unknown Method Name).

The device contains rules for the following categories:

- Packet-based rules (global rules apply to all traffic)
- Firewall rules
- Intrusion protection system rules (IPS Rules)
- Rate based rules

Rules also have treatments associated with them that tell the Corero Network Device what you want to happen when a rule is triggered. Treatments include the action the device should apply to the traffic that triggers the rule (allow, drop, reject), and the type of logging you want performed. There is one set of packet-based rules and one set of rate based rules. You can modify the treatments for each individual rule, as described in Chapter 15, “Managing Rules and Rule Sets”.

Rule Sets

For rules, there are multiple copies of the rules called rule sets, from which you can create your own copies. You can configure the treatment for the same rule differently in different rule sets.

A named set of rules, including the treatment you configure for each rule in the set, is called a Rule Set. The device comes with several pre-defined rules sets such as RecommendedServerProtection and StrictServerProtection.

For more information about rule sets, see Chapter 15, “Managing Rules and Rule Sets

Elements of a Firewall + IPS Security Policy

There are two primary types of security policy:

- Firewall Security Policy (Figure 11-5)
  Firewall policies provide classic firewall blocking for traffic. These policies base blocking decisions on IP address, layer 4 ports, and segments (port pairs). Each policy includes a set of firewall conditions (defining a specific type of traffic), and the treatment for that traffic.
- IPS Security Policy (Figure 11-6)
  Includes traffic conditions and a selected rule set to describe actions to take when the specified traffic triggers the rule. Security policies validate protocol and file attachment content, and also check for attack signatures and acceptable application-usage policies.

Figure 11-5: Elements of a Firewall Policy

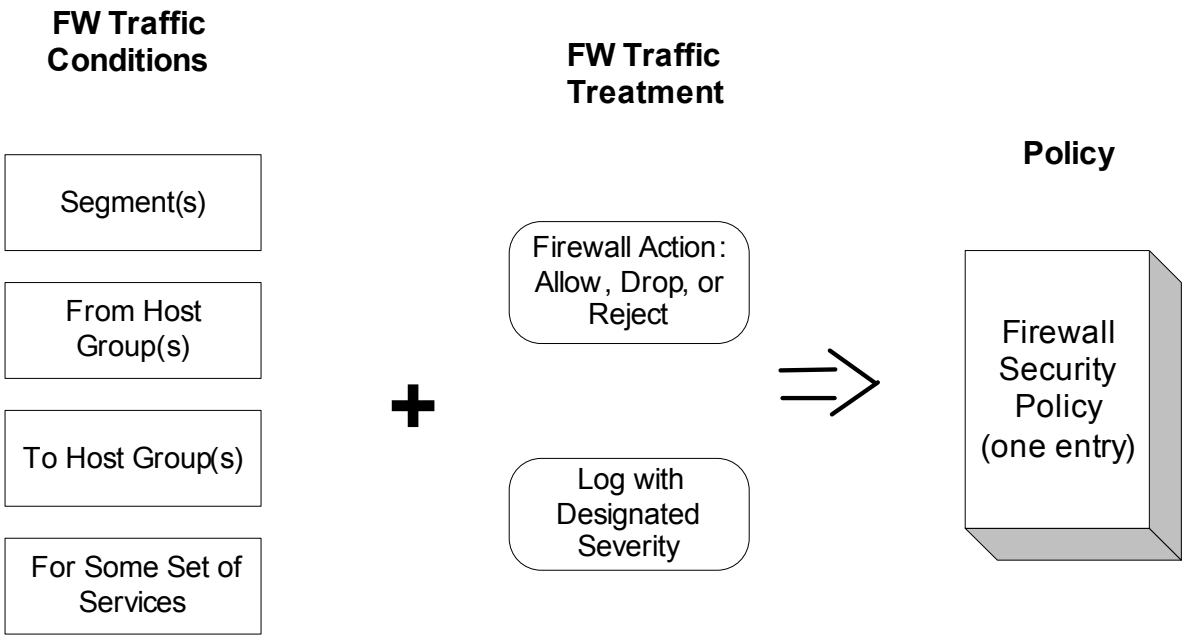

Figure 11-6 shows the elements used in a firewall policy.

Figure 11-6: Elements of an IPS Policy

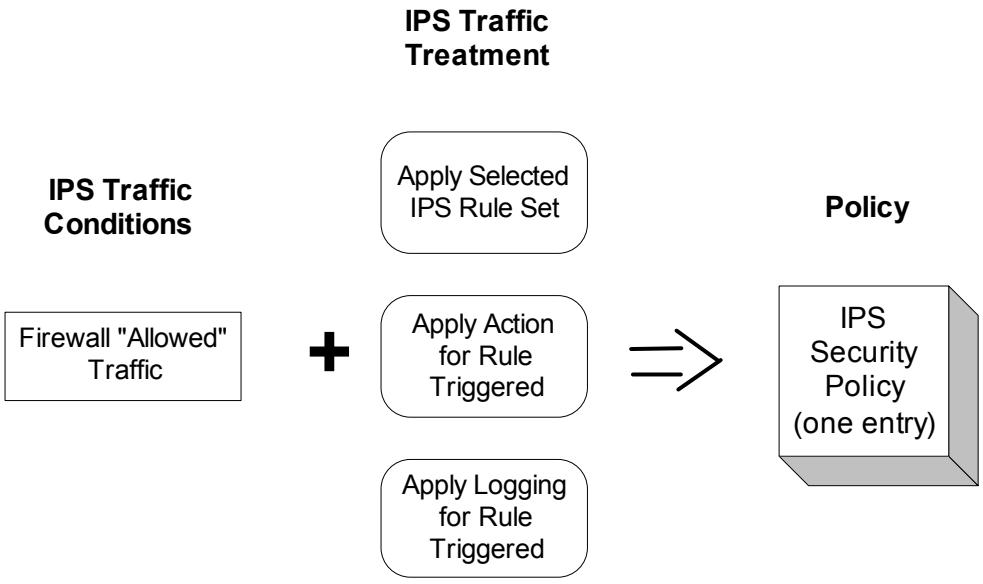

These two policy specifications are combined in a Firewall + IPS Policy. Initially, the Firewall policy allows or denies different types of traffic. Then, the allowed traffic in a firewall policy has an IPS policy applied to it. These two forms of security policy are combined into a FW+IPS policy as shown in Figure 11-7.

Figure 11-7: Elements in a Combined FW+IPS Policy

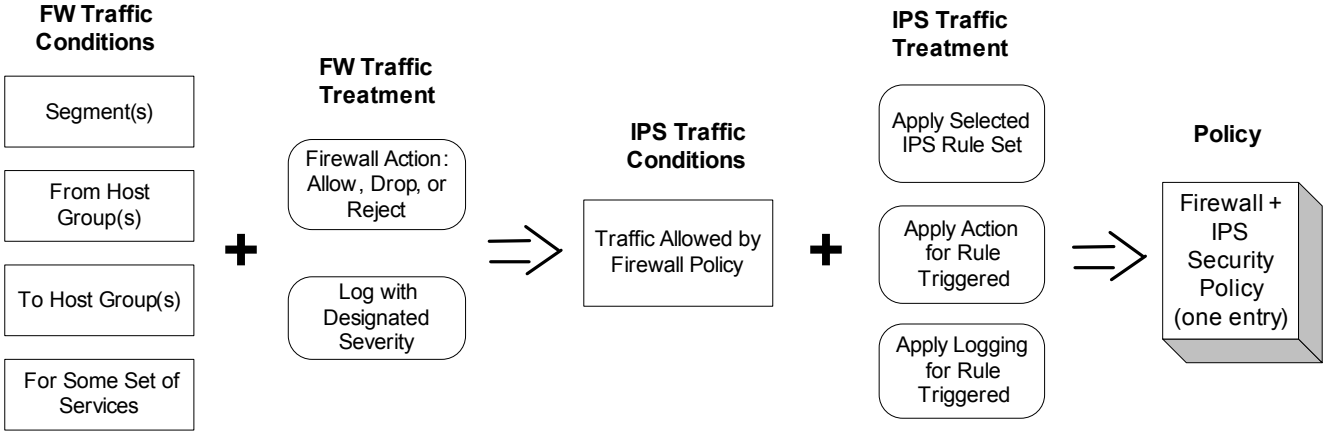

When you view the FW+IPS Policies tab on the Configure Security Policies window of the management application, you can see that each combined FW+IPS policy entry (one line in the table) is comprised of the components listed in Table 11-4.

Each row of the FW+IPS Policies table provides one complete policy entry. Each entry includes the conditions that identify the traffic, the treatment for that traffic, and, for traffic that passes the firewall, its IPS treatment.

Typically, when you need to define policies for your Corero Network Device, you will add them in pairs. You will add a row to the policy table defining the specific traffic type you want to permit, and specify the client and host server

groups from which you will permit it. Then you will specify a second row in the policy, immediately below the first row (indicating a lower priority), that specifies that the rest of the traffic of that type will not be permitted to or from any Any client host group or Any server host group.

For example, the following components comprise the highlighted default policy entry shown in Figure 11-8.

Figure 11-8: Default Strict Server Protection Policy

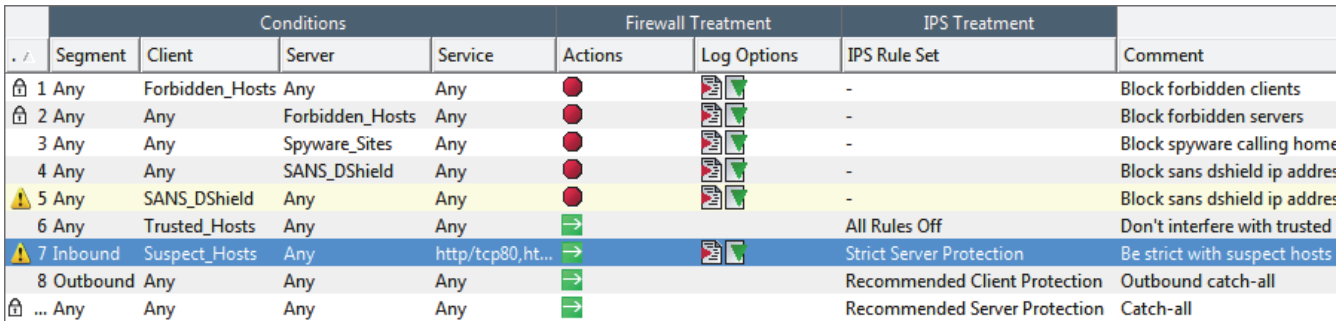

		Conditions			Firewall Treatment		IPS Treatment	
	Segment	Client	Server	Service	Actions	Log Options	IPS Rule Set	Comment
1	Any	Forbidden_Hosts	Any	Any			-	Block forbidden clients
2	Any	Any	Forbidden_Hosts	Any			-	Block forbidden servers
3	Any	Any	Spyware_Sites	Any			-	Block spyware calling home
4	Any	Any	SANS_DShield	Any			-	Block sans dshield ip addres
5	Any	SANS_DShield	Any	Any			-	Block sans dshield ip addres
6	Any	Trusted_Hosts	Any	Any			All Rules Off	Don't interfere with trusted
7	Inbound	Suspect_Hosts	Any	http/tcp80,ht...			Strict Server Protection	Be strict with suspect hosts
8	Outbound	Any	Any	Any			Recommended Client Protection	Outbound catch-all
...	Any	Any	Any	Any			Recommended Server Protection	Catch-all

Table 11-4 describes the settings for the policy shown in Figure 11-8, and how they apply to traffic.

Table 11-4: Description of a Sample Default FW+IPS Policy

Column	Selection	How This Policy Affects Incoming Traffic
Priority		
Row	7	If no other policy before Row 7 has yet applied to the current traffic...
Conditions		
Segment:	Inbound	If the traffic comes in on an inbound port...
Client:	Suspect_Hosts	If the traffic is coming from any IP address in the Suspect_Hosts group...
Server:	Any	If the traffic is going to any IP address in any server host group...
Service:	http/tcp80,httpSsl/tcp443,ftp/tcp21,ftpSsl/tcp90,smtp/tcp25, smtpSsl/tcp465,dns/tcp53,dns/udp53	If the traffic is using one of the specified services...
Firewall Treatment		
Firewall Treatment - Action:	Allow	Then let the traffic pass through the firewall policy, so the IPS Rule Set can be applied.
Firewall Treatment - Log Options:	Low severity	Log any traffic this policy identifies with a severity of Low.
IPS Treatment		
IPS Treatment - IPS Rule Set:	Strict Server Protection	Then apply the Strict Server Protection set of IPS rules, following the treatment specified in the rule set for any IPS rule that the traffic triggers.

Default FW+IPS Policy Operation

Corero Network Devices apply policy entries on the FW+IPS Policies tab in order, from top to bottom, until the traffic being examined matches a policy entry. It then applies the conditions defined by that entry.

The IPS Unit ships with a default firewall and intrusion protection policy containing the entries shown in Figure 11-9.

Figure 11-9: IPS Unit Default FW+IPS Policies

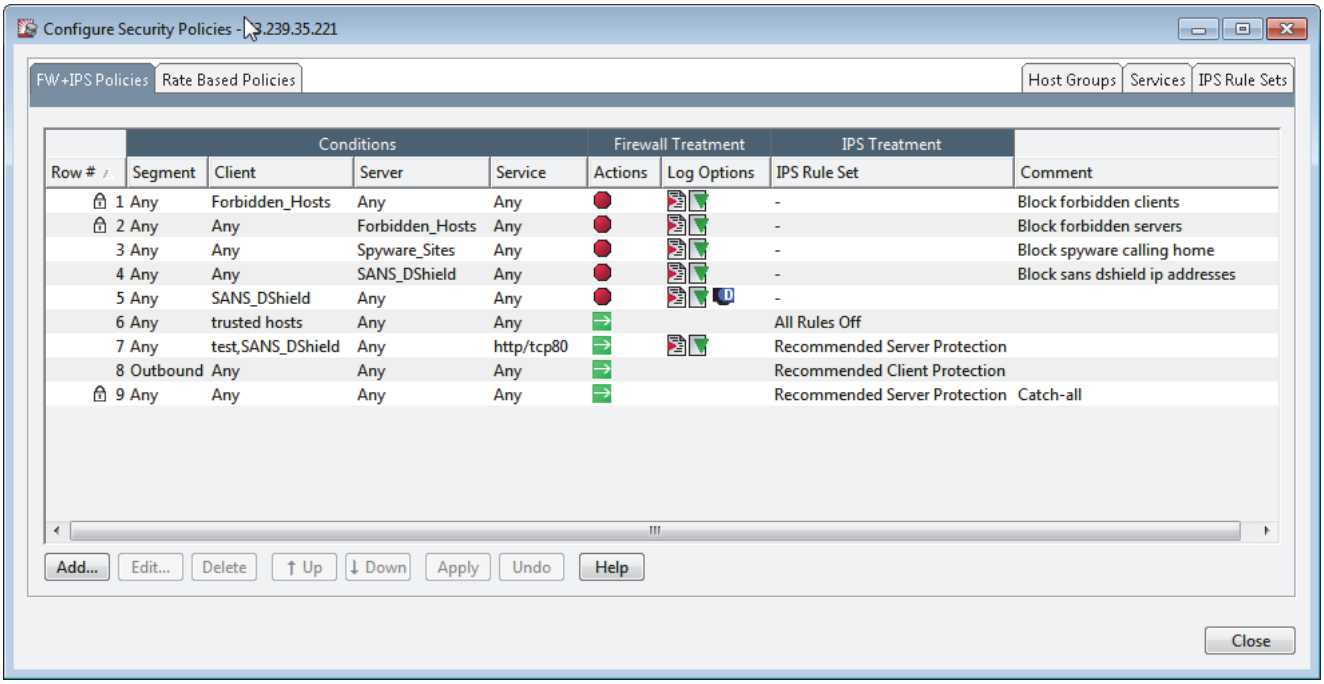

Each row in the FW+IPS tab table is a separate policy entry. Together, all of the entries make up the complete Firewall +IPS security policy. Policies are processed in row order, with Row 1 processed first, and so forth.

NOTE

You may only add policy entries after the initial default policies (Forbidden_Hosts), and before the final Catch-All policy that covers all traffic to which no current policy applies.

The default entries use the default host groups listed in Table 11-5, and block or restrict much of the obviously questionable traffic. The table below describes each entry.

Table 11-5: Default FW+IPS Policies

Row	Description	Summary
1	From any port in any port pair, any IP address in Forbidden_Hosts trying to talk to any server connected to an internal port on the Corero Network Device, using any service, drop and log the traffic. This is a locked default policy that cannot be moved or deleted.	Forbidden_Hosts can't get in or out. Place attackers here for immediate, total, blocking. Also stops internal hosts in this group from initiating connections with the Internet.

Table 11-5: Default FW+IPS Policies *(Continued)*

Row	Description	Summary
2	From any port in any port pair, any IP address trying to initiate a connection to any server in Forbidden_Hosts, using any service, drop and log the traffic. This is a locked, default policy that cannot be move or deleted.	Forbidden_Hosts can't be servers. Stop internal hosts from connecting to infected web sites.
3	From any port in any port pair, any internal IP address trying to talk to any IP address in Spyware_Sites, using any service, drop and log the traffic. Does not stop a spyware infection, but stops any of your infected hosts from sending data back to the spyware site. The IP addressed in this host group are updated via Top Response Protection Pack updates to the Corero Network Device.	Breaks the spyware feedback loop.
4	From any port in any port pair, any internal IP address trying to talk to any IP address in SANS_DShield, using any service, drop and log the traffic. These are updated via Top Response Protection Pack updates to the Corero Network Device.	Uses the DShield list provided by the SANS Institute. These are IP addresses that SANS deems offensive.
5	From any port in any port pair, any host in Trusted_Hosts trying to talk to any IP address using any service, allow the traffic and apply the Null rule set (all rules off).	Allows authorized hosts to avoid being blocked.
6	From any inbound port in any port pair, any IP address in Suspect_Hosts trying to talk to any server host group, allow only a specific set of services, and then apply the Strict Server Protection rule set to that traffic.	Allow only limited traffic, carefully screened.
7	From any outbound port in any port pair, any IP address, trying to talk to any IP address using any service, allow the traffic and apply the Recommended Client Protection Rule Set.	Allows any outbound traffic not specifically forbidden by earlier entries.
8	From any port in any port pair, any IP address, trying to talk to any IP address, using any service, allow the traffic and apply the Recommended Server Protection Rule set. A gateway device allows traffic that is not specifically identified as bad; whereas, a firewall device stops any traffic not specifically allowed. Policy item #8 indicates that this device is, by default, set up as a security gateway, not a firewall. You can easily change this line in the policy to drop all other traffic (making the device act more like a firewall), or, if desired, you can apply a stricter rule set to the allowed traffic. This is a locked default policy that cannot be moved or deleted.	Allow any leftover traffic, but apply recommended rules.

Elements of a Rate Based Security Policy

Rate based policies establish traffic limits . These policies provide DDoS protection and connection limits to help prevent your network resources from becoming overwhelmed.

For more information on rate-based policies, see Chapter 20, ‘‘SYN Flood and Connection Limiting Security”, Chapter 21, ‘‘Client Rate Limiting” and Chapter 22, ‘‘Advanced Client Rate Limiting”.

The elements for an IPS rate based security policy are shown in Figure 11-10.

Figure 11-10: IPS Rate-Based Security Policy Elements

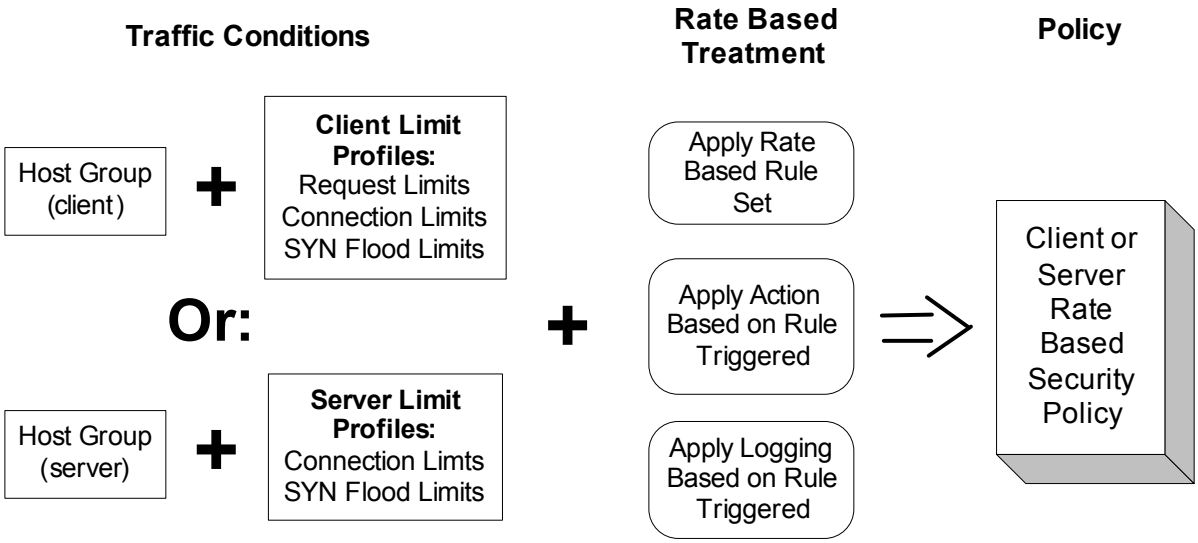

Chapter 12
Managing FW+IPS Security Policies

This chapter describes how to customize your Firewall + IPS (FW+IPS) security policy. These policies enable you to craft specific ways that different types of traffic are treated by your Corero Network Devices. You can specify which types of traffic to drop and permit based on firewall settings. For traffic that successfully passes through the firewall, you can specify which Corero IPS rules are applied to the traffic, and if the traffic matches a rule, how that traffic will be treated by the device.

For detailed information on security policies and how they work, see Chapter 11, "About Security Policies".

This chapter includes the following topics:

NOTE

For information on managing rate-based policies, see Chapter 21, "Client Rate Limiting".

Viewing FW+IPS Policies

The management application makes it easy for you to establish security policies for the device's subsystems. Because Firewall (FW) and Intrusion Protection System (IPS) policies share comment elements, the GUI uses a common screen to configure a combined FW+IPS. Traffic that is not dropped based on the Firewall policy is then examined and handled based on the associated IPS policy.

Figure 12-1 shows the default FW+IPS policy with an additional line added by the IT department to limit company users from playing "World of Warcraft" over the network. It is assumed that the IT department has created a host group called Internal_Networks which contains the IP addresses of systems on the internal networks of the organization.

Figure 12-1: FW+IPS Policy Example

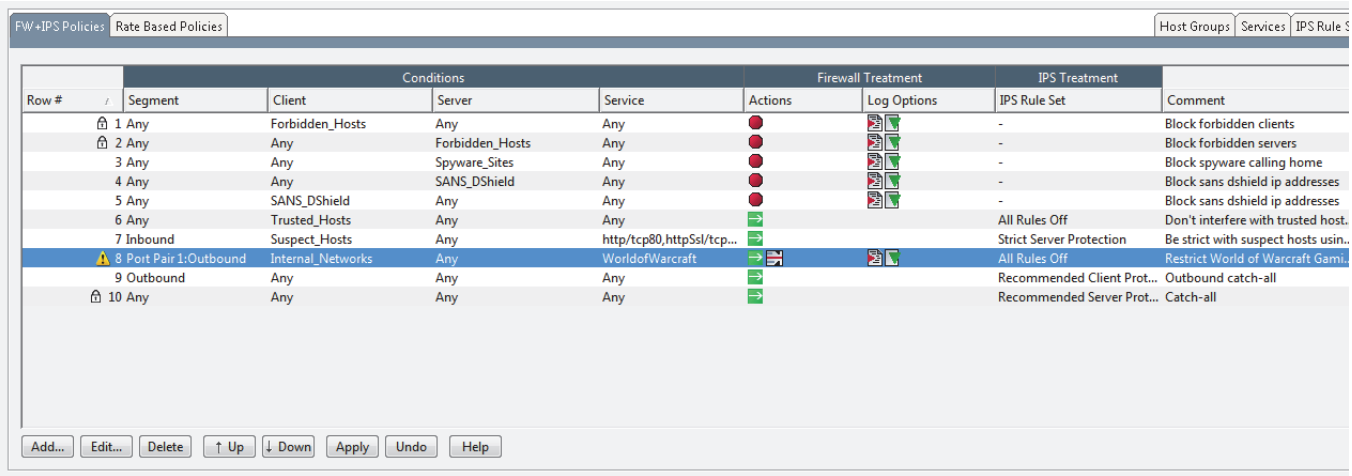

Each row in the FW+IPS tab table is a separate policy entry. Together, all of the entries make up the complete Firewall +IPS security policy.

Policies are processed in row order, with Row 1 processed first, and so forth.

NOTE

You may only add policy entries after the initial default policies (Forbidden_Hosts), and before the final Catch-All policy that covers all traffic to which no current policy applies.

To view FW+IPS Policy information using the Corero Network Device management application:

1. Do one of the following:
   - Click the Security Policies toolbar button.
   - From the navigation tree, choose Configure Security > Security Policies.
   - From the navigation tree, choose Manage Security > Security Policies.

   The Configure Security Policies dialog box displays.
2. To view FW+IPS Policies, click the FW+IPS Policies tab (Figure 12-1).

The columns displayed for each policy are listed in Table 12-1.

Table 12-1: FW+IPS Policy Table Columns

Column	Description
Status Icons and Row Background (Highlight) Colors	• Default policy entries are marked by a lock icon. They can be modified but not deleted. Policy entries that have been modified are highlighted in yellow. • To their left, these entries display a yellow triangle icon with an exclamation point inside. This indicates that you have made changes to the policy, but have not applied them. To apply your changes, click Apply. Alternatively, if you do not want the changes applied, click Undo. **Note:** Your changes will not be preserved across a system restart until you Save them.
	A row that is highlighted in gray is **not** processed. There are three reasons that a policy row may be light gray, each indicated by a different icon: • **!** - If the row is preceded by an exclamation point, an element used to define the policy entry has been deleted, rendering the policy invalid. • **X** - If the row is preceded by a gray X, the policy entry has been disabled by a user. • **?** - If the row is gray and preceded by a question mark, it indicates that the policy entries prior to this entry already handle the traffic that it defines. Because a higher priority policy is processing the traffic, this policy entry is not used.
Row	Policies are applied in the order in which they are listed, with row 1 being applied first. When traffic is processed, policies are applied in numeric order (from top to bottom), until it finds a policy that matches the traffic under scrutiny. Once it finds a policy that applies to the current traffic, it applies the Firewall policy. Then, if the Firewall treatment is Allow, it applies the policy's IPS rule set and its treatment.
Segment	Indicates the port pair(s) to which this policy entry applies. Once you specify a pair, you can indicate whether you want the device to evaluate only inbound or outbound traffic.
Client	The Client column specifies that the policy only applies to traffic *coming from* the selected Host Group (set of IP addresses).
Server	The Server column specifies that the policy only applies to traffic *going to* the selected Host Group (set of IP addresses).
Service	Indicates which set of network applications are included in this policy.
Firewall Treatment (Action)	Specifies what firewall action the Corero Network Device should take if the current traffic meets the conditions defined by this policy. Firewall actions are: • Allow - Enables the traffic to pass the firewall, then be processed by the IPS policy. • Drop - Drops the traffic without informing the source. No further processing occurs. • Reject - Drops the traffic and informs the source. No further processing occurs. If Apply Rate Limiting is selected, the specified limit is displayed.

Table 12-1: FW+IPS Policy Table Columns *(Continued)*

Column	Description
Firewall Treatment (Log Options)	Specifies the following information: • Whether the firewall action should be logged. • If the firewall action is logged, specifies the severity of the event. • If the packet was dropped or rejected, specifies whether the packet will be copied to the Discard Port. **Note:** If you have configured a Discard Port, you can log all traffic that triggers a rule even if the Action for that rule is set to Allow. This feature enables you to record usage information about applications you allow as well as those you disallow.
Rule Set	Indicates what set of IPS (protocol validation and content inspection) rules the device should apply to any traffic that has been allowed to pass by the Firewall policy.
Comments	A user-specified description of the policy.

Understanding the Difference Between Making and Activating Policy Changes

The management application for your Corero Network Device enables you to view policy changes before they are applied to traffic.

These types of changes include:

- Adding or Modifying a policy row
- Deleting or moving a policy row
- Enabling or disabling a policy row

If you make changes to the policy table, the changes appear in the FW+IPS Policy tab, but are not used by the device until you click the Apply button. If you do not want to apply your changes, click Undo. The Undo button restores the policy table to the condition it was in immediately after the last time you clicked the Apply button. If you want your changes preserved across system restarts, ensure you click the Save Configuration toolbar button.

Modifying a Policy's Priority

Policies are processed in row order. Row 1 is the first policy applied to incoming traffic, row 2 is next, and so forth. If the traffic matches the specifications in a particular row, then the traffic is processed based on the treatment specified by the policy row. If the traffic does not match the specifications in a particular row, then it is tested to see if it meets the specifications in the next policy row, based on row priority.

You can use the order in which policies are applied to craft a sequence of policies that meet your site's requirements. Typically, more specific policies are placed in lower row numbers and given first priority. More general policies are listed in later rows.

CAUTION

Carefully consider where to place policies in the policy table to ensure specific types of traffic are properly treated.

Several default policies are given specific places in the policy table, and cannot be moved. The first two rows always specify policies for Forbidden_Hosts clients and server systems. The last row is always a default row, called the catch-all row, that applies to traffic that does not match any higher priority policy.

In order to modify a policy's priority, you must change its order (placement) in the policy table. To do this:

1. View the current policy rows as described in Viewing FW+IPS Policies (page 12-2).
2. Select the row whose order you want to change.
3. To move the policy to a lower row number (higher priority), click Up.
4. To move the policy to a higher row number (lower priority), click Down.
5. Do one of the following:
   - If you are satisfied with your changes, click Apply.
   - If you do not want to keep your changes, click Undo.
6. If you applied your changes, save your changes by clicking the Save Configuration toolbar button.

Configuring FW+IPS Policies

To configure an FW+IPS Policy:

1. Ensure the elements you wish to use in your FW+IPS policy are already available in the management application. These components can include:
   - Host Groups - for detailed instructions, see Chapter 13, "Managing Host Groups".
   - Services - for detailed instructions, see Chapter 14, "Managing Services".
   - IPS Rule Sets - for detailed instructions, see Chapter 15, "Managing Rules and Rule Sets".

   These items comprise the building blocks of your policy. If the desired policy elements are not available, create them before you proceed.
2. Go to the FW+IPS Policies tab as described in Viewing FW+IPS Policies (page 12-2).
3. Do one of the following:
   - To create a new FW+IPS Policy, click Add. The Add FW+IPS Policy dialog box displays.
   - To modify an existing FW+IPS Policy, select the policy, then click Edit. The Edit FW+IPS Policy dialog box displays.

   Figure 12-2 shows the Add FW+IPS Policy dialog box. The Edit FW+IPS Policy dialog box contains the same tabs and options.

Figure 12-2: Configuring a Policy: Conditions Tab

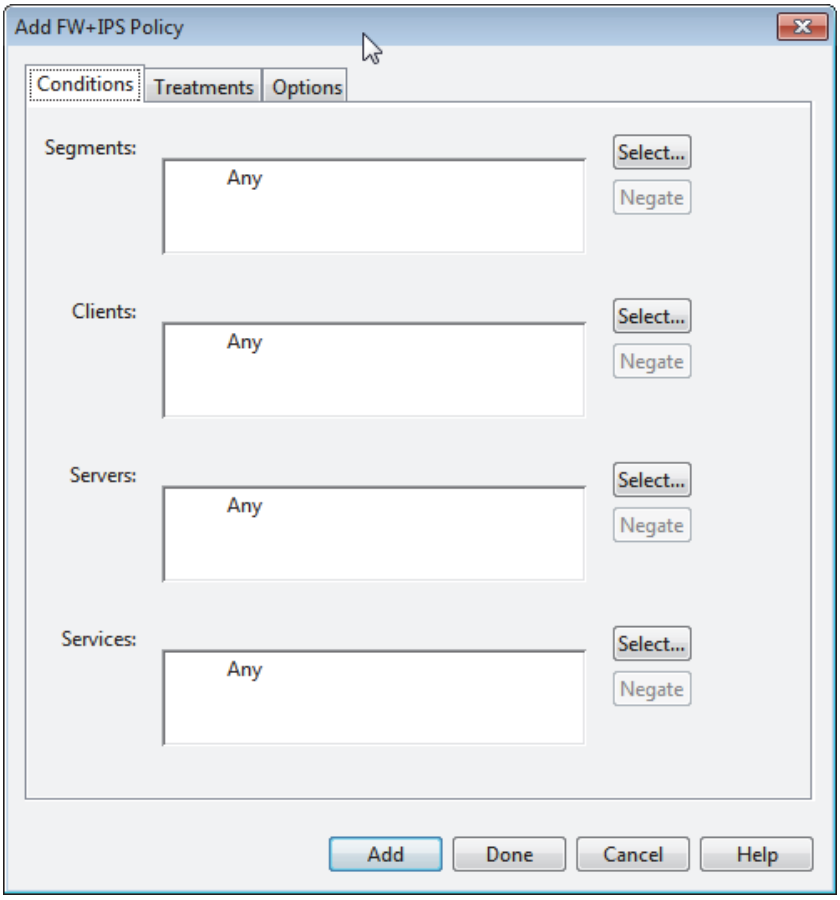

4. On the Conditions tab (Figure 12-2), specify the segments (inbound and outbound port pairs) whose traffic you want affected by the policy.

   To specify one or more segments:

   a. Adjacent to the Segments list, click Select. The Select Segments dialog box displays. Selected segments are marked with a check. Deselected segments are marked with an X. You can choose individual segments, or select the option for Any. You can also search for segments by entering text in the Search field.

   b. To specify segments, do one of the following:
   - Click an individual segment.
   - Click a segment, then Shift+Click another segment to select those two and all segments between them.
   - Click a segment, then Ctrl+Click additional segments to select specific segments.

   c. When you have selected the desired segments, do one of the following:
   - Click Select to include these segments.
   - Click Deselect to exclude these segments.

   d. When you are finished specifying segments, click OK.

   e. If you want the policy applied to these segments, you are finished. If you *do not want* the policy applied to these segments (that is, you want the policy applied to all segments *except* for these) click Negate. A red X displays adjacent to your selection in the Segments list.

5. On the Conditions tab (Figure 12-2), specify the Clients whose traffic you want affected by the policy. You do this by selecting one or more Host Groups whose IP Addresses you want to choose as clients.

   NOTE

   After you select clients, you can specify whether you *do* want the policy applied to them, or whether you *do not* want the policy applied to them.

   To specify one or more clients:

   a. Adjacent to the Clients list, click Select. The Select Clients dialog box displays. Selected client host groups are marked with a check. Deselected client host groups are marked with an X. You can choose individual client host groups, or select the option for Any. You can also search for host groups by entering text in the Search field

   b. To specify client host groups, do one of the following:
   - Click an individual host group.
   - Click a host group, then Shift+Click another host group to select those two and all host groups between them.
   - Click a host group, then Ctrl+Click additional host groups to select specific groups.

   c. When you have selected the desired client host groups, do one of the following:
   - Click Select to include these host groups.
   - Click Deselect to exclude these host groups.

   d. When you are finished specifying client host groups, click OK.

   e. If you want the policy applied to these client host groups, you are finished. If you *do not want* the policy applied to these client host groups (that is, you want the policy applied to all host groups *except* for these) click Negate. A red X displays adjacent to your selection in the Clients list.

6. On the Conditions tab (Figure 12-2), specify the Servers whose traffic you want affected by the policy. You do this by selecting one or more Host Groups whose IP Addresses you want to choose as servers.

NOTE

After you select servers, you can specify whether you *do* want the policy applied to them, or whether you *do not* want the policy applied to them.

To specify one or more servers:

a. Adjacent to the Servers list, click Select. The Select Servers dialog box displays. Selected servers host groups are marked with a check. Deselected servers host groups are marked with an X. You can choose individual server host groups, or select the option for Any. You can also search for host groups by entering text in the Search field

b. To specify server host groups, do one of the following:
- Click an individual host group.
- Click a host group, then Shift+Click another host group to select those two and all host groups between them.
- Click a host group, then Ctrl+Click additional host groups to select specific groups.

c. When you have selected the desired server host groups, do one of the following:
- Click Select to include these host groups.
- Click Deselect to exclude these host groups.

d. When you are finished specifying server host groups, click OK.

e. If you want the policy applied to these server host groups, you are finished. If you *do not want* the policy applied to these server host groups (that is, you want the policy applied to all host groups *except* for these) click Negate. A red X displays adjacent to your selection in the Servers list.

7. On the Conditions tab (Figure 12-2), specify the Services whose traffic you want affected by the policy.

NOTE

After you select services, you can specify whether you *do* want the policy applied to their traffic, or whether you *do not* want the policy applied to their traffic.

To specify one or more services:

a. Adjacent to the Services list, click Select. The Select Services dialog box displays. Selected services are marked with a check. Deselected services are marked with an X. You can choose individual services, or select the option for Any. You can also search for services by entering text in the Search field

b. To specify services, do one of the following:
- Click an individual service.
- Click a service, then Shift+Click another service to select those two and all services between them.
- Click a service, then Ctrl+Click additional services to select specific services.

c. When you have selected the desired services, do one of the following:
- Click Select to include these services.
- Click Deselect to exclude these services.

d. When you are finished specifying services, click OK.

e. If you want the policy applied to these services, you are finished. If you *do not want* the policy applied to these services (that is, you want the policy applied to all services *except* for these) click Negate. A red X displays adjacent to your selection in the Services list.

8. Click the Treatments tab (Figure 12-3). This tab enables you to specify the firewall and IPS treatments you want applied to traffic that matches all of the parameters on the Conditions tab.

Figure 12-3: Configuring a Policy: Treatments Tab

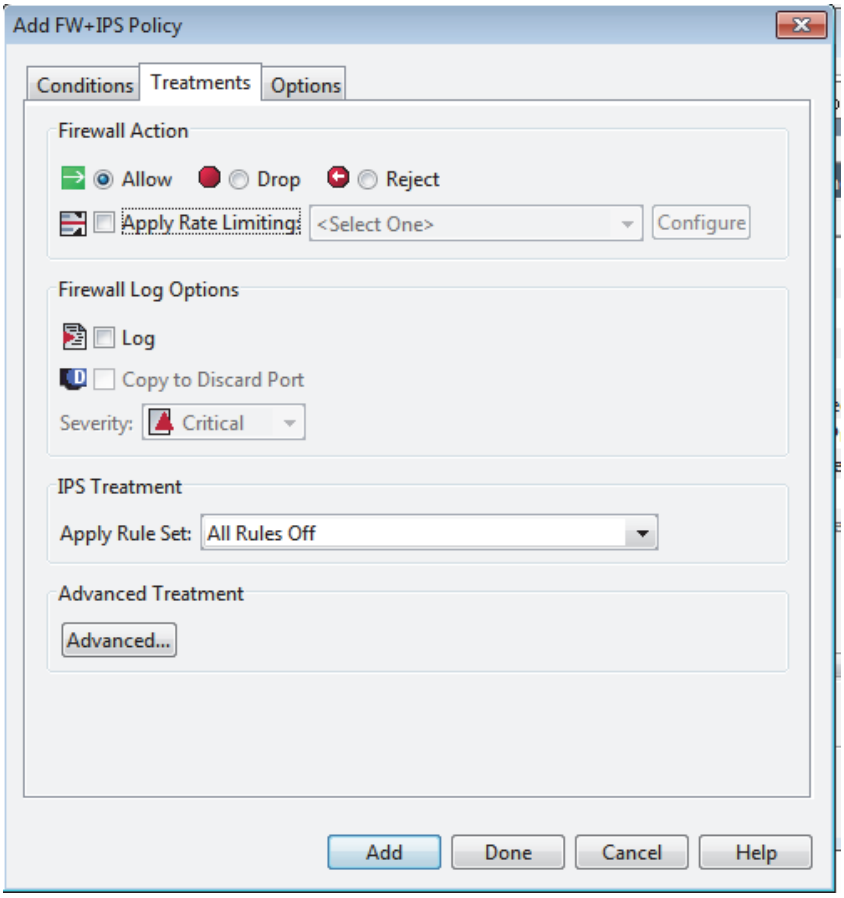

9. On the Treatments tab (Figure 12-3), specify a Firewall Action. When traffic meets the conditions you configured earlier in this procedure, the Firewall Action defines how you want the traffic treated. To specify an action:
   a. Specify how you want the traffic treated:
      - Selecting Allow lets the traffic through to the next step in the processing flow.
      - Selecting Drop causes the device to silently drop the traffic.
      - Selecting Reject causes the device to drop the traffic and send a TCP reset to actively block the sender.
   b. If you specify that you want the traffic Allowed, you can also specify a rate limit. To do so, click the Apply Rate Limiting check box.
   c. If you have selected the Rate Limiting check box, select the desired application rate limit (in kbps) from the drop-down.

      NOTE

      If the rate limit you want to specify does not appear in the drop-down list, click Configure to add, edit, or delete a limit.

10. On the Treatments tab (Figure 12-3), specify Firewall Log Options. These log options specify how and whether you want traffic that meets the conditions for this policy logged. You can also specify whether you want this traffic copied to the discard port.
    a. If you want the event to be logged every time this policy applies to traffic, click the Log check box.
    b. In the Severity drop-down select the severity you want associated with the log entry.

c. To copy the packet to a designated discard port, click the Copy to Discard Port check box.

NOTE

If you have configured a Discard Port, you can log all traffic that triggers a rule, even if the Action for that rule is set to Allow. This feature enables you to record usage information about applications you allow as well as those you disallow.

11. On the Treatments tab (Figure 12-3), specify an IPS Rule Set by selecting the desired rule set from the drop-down list. For more information on IPS Rule Sets, see Chapter 15, "Managing Rules and Rule Sets".
12. On the Treatments tab (Figure 12-3), specify Advanced Treatment settings by clicking the Advanced button. Select the desired check boxes, then click OK. Options for the Advanced Treatment Settings dialog box are described in Table 12-2

Table 12-2: Security Policy Advanced Treatment Settings

Setting	Description
Block IP Fragments	When an incoming network packet is too large for the network equipment to handle, the packet can be fragmented and sent on to its destination, where the fragments will be reassembled into the original packet. If you have an application that must use fragments, you can configure it as an exception, but since fragments are often used as a source of attacks, consider blocking fragments for all applications. When you block fragments, an ICMP error indication (MTU Exceeded) will be returned to the source, and the source can then send the information as smaller packets to begin with.
Check for TCP Connections with Missed Setups (Mid-Flows)	Many types of attacks use common scanning techniques which can be identified as incomplete TCP connections. This setting checks to ensure the 3-way TCP handshake is complete for a particular connection. If the 3-way handshake is not completed, we drop the connection and may trigger rule tln-001017 (TCP Connection with Missed Setup) or rule tln-001025 (TCP Connections with Missed Setup - RST [Reset] Packet Only). Unless operation of your network or business requires the use of mid-flows, consider dropping them. At worst, this can result in a lost connection which is then retried. You can configure policies that: • Drop all TCP mid-session traffic at all times. • Only allow TCP mid-session traffic to pass for specified, critical, internal or external host groups. • Handle all HTTP mid-session traffic differently than other TCP mid-session traffic.
Mirror Flow	If you have specified a mirror port for your configuration, you can specify whether allowed traffic that matches this policy is sent to the mirror port. If there are multiple mirror ports defined, traffic is balanced among them using a round robin algorithm.

13. Click the Options tab. To specify policy options:
    a. Optionally, you can enter a meaningful description of the policy in the Comment area. This type of information can help you identify why you created or modified a policy, or how this policy differs from other, similar policies.
    b. Specify whether the policy is Enabled or Disabled by clicking the appropriate radio button. A disabled policy will not affect traffic.

14. When you have finished specifying policy settings in the Corero Network Device management application, click Done at the bottom of the Configure Security Policies dialog box. If you want to save the policy and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
15. To apply your changes to the device, click Apply.
16. Save your changes by clicking the Save Configuration toolbar button.

Chapter 13
Managing Host Groups

A host group is a named set of IP addresses. You specify a host group when defining security policies.

A host group can define a group of clients or a group of servers. If the members of a given host group will act as both clients and servers, you may need to create more than one security policy entry for the same host group and apply specific conditions and treatments based on whether the host group is initiating or responding to traffic requests.

You can specify IP address members of a host group in several ways: as a single IP address, as a small group of IP addresses, as an IP address range, as a network block of IP addresses, or as any combination of these.

Corero Network Devices come with a set of predefined host groups that will meet most users' requirements. You can modify predefined host groups, or define your own.

NOTE

Any given IP address, whether specified alone or in a range, can only be in one host group.

This chapter includes the following topics:

Defining Host Groups

When you define a host group, you do so by specifying named IP address ranges. This helps to simplify host group definition. A named IP address range can be a single IP address or a set of addresses. You can create a named IP address range and specify the following attributes for it:

- Associated host group
- IP address or range specification (individual IP address, IP address range, subnet)
- Type of spoof check treatment
- Source and/or destination filters
- Range attributes such as subnet or broadcast

Spoof checking is used to identify attacks where hosts modify the IP address to imitate a different (internal or external) IP address. You can instruct the Corero Network Device to check the type of port (internal or external) on which traffic with this IP address arrives. You can disable spoof checking, or you can specify whether you want to allow traffic from an IP address in this range only if it appears on the internal port in a port pair, or only if it appears on the external port.

IP Address Specification Considerations:

When specifying IP address ranges, *it is important that you consider the following*:

- An IP address can belong to one, and only one, host group.
- You can, intentionally or unintentionally, specify IP addresses or address ranges for multiple host groups that match a single IP address.
- When you create, modify or delete a host group, this can result in one or more IP addresses being moved from that host group to a different host group.
- The management application does not inform you when creating, editing, or deleting a host group results in the reassignment of an IP address from one host group to another.

NOTE

To determine the current host group associated with an IP address, perform an IP Query on that address. For more information, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).

- If an IP address matches the specifications for more than one host group, there is an order of precedence the system follows in order to assign the address. The order of precedence goes from most specific to least specific, as described in Table 13-1.
- If an IP address matches multiple specifications that have the same order of precedence (for example, more than on specified address range), it will be assigned to the host group associated with the most recently defined IP address range.

Table 13-1: IP Address Host Group Assignment Order of Precedence

Order	Description
First	The address is specified as an individual IP address in the host group definition.
Second	The address falls into a specified address range.
Third	The address falls into a specified subnet.

A named IP address range can also be a subset of another range, as long as it falls completely within the larger range. However, by creating a subset address range, you remove those addresses from the larger range (creating a hole in the larger address range). (A singleton IP address is considered a range of one address.)

For example:

1. If you had an initial range (Range A) that contained the addresses from 10.20.30.40 to 10.20.30.255.
2. Then you created a second range (Range B) that contained the addresses from 10.20.30.100 to 10.20.30.125.
3. Range A would automatically be modified to include 10.20.30.40 to 10.20.30.99, and 10.20.30.126 to 10.20.30.255.

Default Host Groups

The Corero Network Device provides a set of predefined host groups which you can modify. You can also add your own groups.

By default, most of the predefined host groups do not have IP address ranges assigned to them. You must add IP addresses to the default groups you decide to use.

The default host groups are defined in Table 13-2. Note that the IP addresses 0.0.0.0 and 255.255.255.255 can be moved from the "other_Hosts" group to a different host group but cannot be deleted.

Table 13-2: Default Host Groups

Host Group Name	Pre-Populated or Empty	Default FW Policy	Default IPS Policy	Other Policy Notes
AOL_IM_Servers	Populated with clients such as AOL IM Server IP Addresses	None	None	Use to when you want to block AOL Instant Messaging (IM) Servers or when you would like to rate limit AOL Instant Messaging in your network.
DNS_Servers	Empty	None	None	Use to create specific policy settings for DNS servers.
Forbidden_Hosts	Empty Add Known Bad Hosts, such as the loop back network address range (127.0.0.0).	FW blocks these hosts	N/A Dropped before IPS subsystem	By default, the device quickly and efficiently blocks all traffic from members of this host group using a built-in Firewall rule. There are two firewall entries for Forbidden Hosts; as clients and as servers.
Mail_Servers	Empty	None	None	Use to create specific policy settings for mail servers.
Mega_Proxies	Populated with clients such as AOL proxies	None	None	Because these IP addresses can be expected to generate much larger volumes of traffic than other host groups, you may want to configure a different rate-based policy for this group.
Non_Routable_IP	Populated with private and reserved IP addresses	None	Apply the Recommended Client or Server Protection Rule Set	This host group contains default IP addresses that are either private or reserved. The device follows the policy as defined in the Recommended Client or Server Protection Rule Set.
other_hosts	Populated with the entire IP address space.	None	None	This host group contains all IP addresses. As addresses are assigned to other host groups, they are automatically removed from this one.

Table 13-2: Default Host Groups *(Continued)*

Host Group Name	Pre-Populated or Empty	Default FW Policy	Default IPS Policy	Other Policy Notes
SANS_DShield	Empty This group is populated during a TopResponse update.	Drop traffic from these servers	N/A Dropped before IPS subsystem	There is a default firewall rule that blocks all traffic from this server host group. Corero recommends you apply a TopResponse update as soon as installation is complete.
Spyware_Sites	Populate. Add other known sites.	Drop traffic from these servers	N/A Dropped before IPS subsystem	There is a default firewall rule that blocks all traffic from this server host group.
Suspect_Hosts	Empty	Allow only HTTP port 80 traffic	Strict Server Protection	Could be used for hosts that may be generating attacks. In addition to the strict IPS policy, you could apply a more restrictive rate-based policy or rule set to these hosts.
Trusted_Hosts	Empty	None	None	Add hosts that you want to give access to sensitive resources or hosts that you want to apply a less strict security policy to.
User_Mega_Proxies	Empty	None	None	Use to establish your own mega proxy group.
VIP_Services	Empty	None	None	Use to create specific policy settings for VIP services.
WEB_Servers	Empty	None	None	Use to create specific policy settings for web servers.

Viewing Host Groups

To view host groups:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.

2. Click the Host Groups tab for the IPS Unit (Figure 13-1).

Figure 13-1: IPS Host Groups Tab

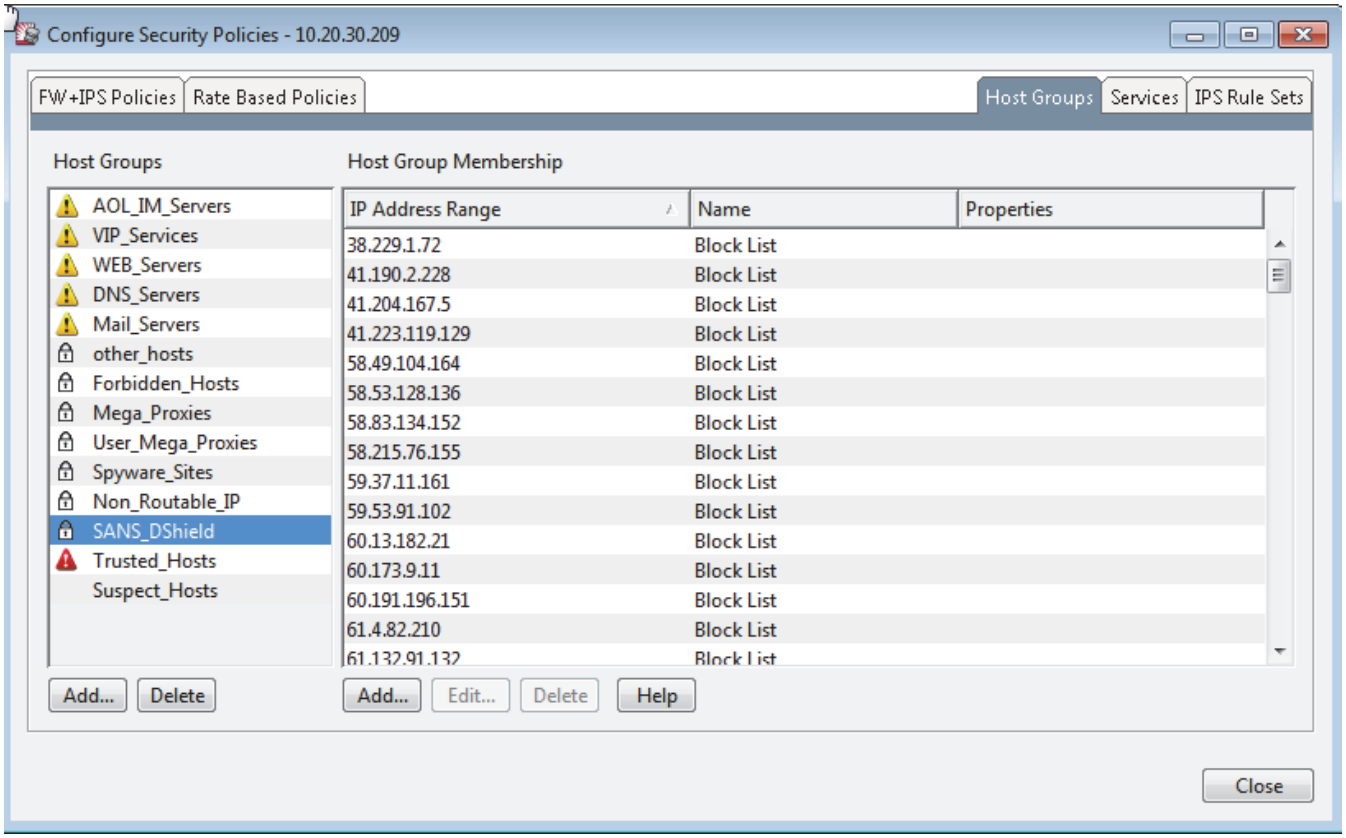

Host group members display in the right pane.

Each default host group has an icon next to it:

- A host group has a lock next to it if this group is part of the default configuration. You can modify the IP addresses contained in this host group, but you cannot delete the group itself.
- A host group has a red warning triangle next to it if no rate based protection has been applied to it.
- A host group has a yellow information triangle next to it if it is not used in a firewall policy.

NOTE

If a host group has no rate-based protection and is not used in a firewall policy it will only have the red warning triangle.

Adding or Editing Host Groups

When you add or modify a host group, you can do one or more of the following:

- Add a host group to the list of host groups.
- Add an IP address or range to the selected host group.
- Edit the selected IP address or range in the selected host group.
- Delete the selected IP address or range from the selected host group.

For information on how to delete a host group, see Deleting Host Groups (page 13-10)

To add or modify a host group:

1. Go to the Host Groups tab as described in Viewing Host Groups (page 13-6).
2. To add a host group:
   a. Under the Host Groups list, click Add. The Add Host Group dialog box displays.
   b. Enter the name of the new host group.
   c. If you want to create another host group, click Add; otherwise, click Done.
3. To add an IP address range to a host group:
   a. Select the host group to which you want to add the address range.
   b. Under the host group membership area, click Add. The Add IP Address Range dialog box displays.
   c. Optionally, specify a name for this address range. Specifying a meaningful name for a range makes it easier to recognize and work with later on.
   d. If, for some reason, you did not select the correct host group initially, select the host group with which you want this IP address range associated.
   e. Specify IP address information. You can add IP addresses in four ways:
   - As an IP address/Prefix (for example 192.0.8.31/24, where the first 24 bits would have to match exactly.)
   - As an IP address/Mask (for example 192.0.8.31/255.255.255.0).
   - As a range specified by a first and last IP address (for example 192.0.128.0-192.01.128.255).
   - As a single IP address (for example 192.0.8.31).
   f. Specify spoof check settings. You can choose from the following settings:
   - Disable
   - Allow from internal ports only
   - Allow from external ports only
   g. To identify the first IP address as a subnet address, click the Advanced button, select the Identify First IP Address As A Subnet Address check box, then click OK.
   h. To identify the last IP address as a broadcast address, click the Advanced button, select the Identify Last IP Address As A Broadcast Address check box, then click OK.
4. To edit an IP address range in a host group
   a. Select the host group in which you want to edit the address range.
   b. Under the host group membership area, select the desired address range, then click Edit. The Edit IP Address Range Settings dialog box displays.
   c. Optionally, specify a name for this address range.
   d. If, for some reason, you did not select the correct host group initially, select the host group with which you want this IP address range associated.

e. Specify spoof check settings. You can choose from the following settings:
   - Disable
   - Allow from internal ports only
   - Allow from external ports only

f. To identify the first IP address as a subnet address, click the Advanced button, select the appropriate check box, then click OK.

g. To identify the last IP address as a broadcast address, click the Advanced button, select the appropriate check box, then click OK.

5. To delete an IP address range from a host group

   a. Select the host group from which you want to delete the address range.

   b. Under the host group membership area, select the desired address range, then click Delete.

   c. You are prompted to confirm your selection.

6. When you have finished specifying host group settings in the Corero Network Device management application, click Done.

7. Save your changes by clicking the Save Configuration toolbar button.

Deleting Host Groups

At some time, you may want to delete a host group. Many users prefer to delete host groups that are empty and will not be used again, because the interface does not display information about whether or not a host group contains IP addresses.

When you delete a host group, existing flows will continue to use host groups until they are finished. For best results, if you plan to delete a host group that contains IP addresses, Corero recommends that you delete all addresses from the host group, and then delete the host group itself.

To delete a host group:

1. Go to the Host Groups tab as described in Viewing Host Groups (page 13-6).
2. Select the host group you want to remove, then click Delete.
3. When you have finished deleting the host group in the Corero Network Device management application, click Done.
4. Save your changes by clicking the Save Configuration toolbar button.

Chapter 14
Managing Services

When you specify a policy, you include information about the services to which that policy applies. You can specify services such as HTTP or DNS, as well as the associated port for the traffic you want addressed by the policy.

You can, for example, block services to host groups that do not require access by those clients, limiting undesired access to your network. If you know a server is dedicated to a specific service (HTTP, for example), you can block all other services for that server.

The most common applications (services) are predefined on your Corero Network Device. You can modify predefined services, or define your own. Service settings are commonly modified to extend the time-out for a specified service.

You may also want to define services so you can tailor client rate limiting settings for some or all servers that communicate with your network.

This chapter includes the following topics:

Viewing Services

The Services window displays a list of services and their attributes. A service is the combination of a protocol, port number, and a server group. The management application provides a large number of default services. You can also define services that run on non-standard ports.

To view services:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.

2. Click the Services tab (Figure 14-1).

Figure 14-1: Services Tab

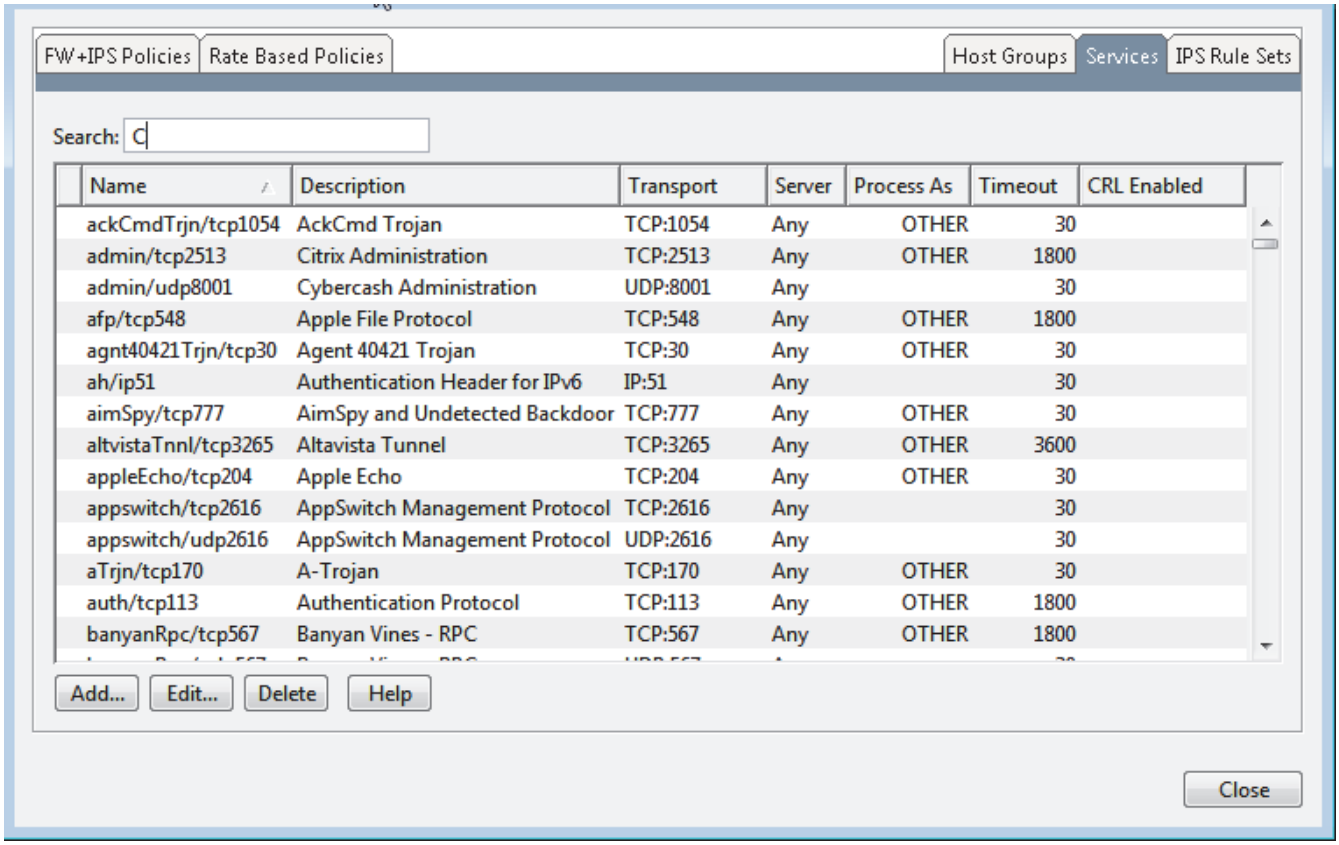

NOTE

When defining a policy, you can select one or more services that you want to include in the policy.

3. To search for a particular service, enter a search string (in the Search box) to display specific services based on the content of the various fields in the table.
4. You can view the settings available for each service. Settings are described in Table 14-1

Table 14-1: Services Tab Settings

Setting	Description
Name	A name that uniquely identifies the service.
Description	A description that further identifies the service. This field is examined by the Search function, so it is important to add a description that will help you pick out the new service from the list of services.
Transport	The protocol (either IP, ICMP, TCP, or UDP) and the port number for the service.
Server	The server or group of servers associated with this service. Possible values are: • Any — Includes the entire IP address range. This group is in effect unless superseded by custom address specifications you define and then applied to a specific application. • Server Group — The servers for this service were selected from the drop-down list of defined server groups. • Address — Indicates that the server was defined as a single IP address, range of IP addresses, IP address range and port range, or combinations of these.
Process As	Identifies a specific protocol or method to use for deep packet inspections.
Timeout	Number of seconds without traffic for this service that the Corero Network Device should wait before timing out the connection.
CRL Enabled	Specifies whether Client Rate Limiting is enabled for this service. If Client Rate Limiting is enabled, if any rate-based policies are configured for this service, they will be used to detect and affect service-specific traffic.

Adding or Editing a Service

There may be times when you want to add a service that was not included in the list of default services, or you want to modify a service to specify a different IP address range or port.

To add or edit a service:

1. Access the Services tab as described in Viewing Services (page 14-2).
2. Do one of the following:
   - To add a service, click Add.
   - To edit a service, select the service, then click Edit.
3. Enter a name and a description for this service.

   NOTE

   If you do not specify one, the Corero Network Device automatically gives the service a name based on its protocol and other information you provide. However, since this name will be automatically generated, it may not be very usable. It is important that you enter a meaningful name and description for this service so you can easily identify it later.

4. In the Connection Timeout field, enter the number of seconds that the device should wait before timing out a connection for this service.
5. If you are adding a new service, you must specify the Transport Settings. To do so, choose the type of service from the Transport drop-down box, then enter the port assigned to this service.
6. If you are adding a new service, you must define the server or group of servers to associate with this service:
   - Any — Include the entire IP address range. This group is in effect unless superseded by custom address specification you define and then applied to a specific application.
   - Host Group — From the drop-down list, select a host group to associate with this application.
   - Address Spec — Enter a single IP address, range of IP addresses, IP address range and port range, or combinations of these. Separate entries by commas. For example, the following assigns a single IP address and a range of IP addresses with specifically defined ports:
     10.20.30.40,10.20.30.45-10.20.30.47:8000-8008
7. Modify advanced service settings as described in Specifying Advanced Service Settings (page 14-5).
8. When you have finished specifying service settings in the Corero Network Device management application, click Done at the bottom of the Configure Security Policies dialog box. If you want to save the service and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
9. Save your changes by clicking the Save Configuration toolbar button.

Specifying Advanced Service Settings

When you are adding or modifying a service, you will want to specify advanced service settings. You use the Advanced Service Settings dialog box to set handling of this service under extreme traffic conditions, and to indicate other special handling for packets associated with this service.

To specify advanced service settings:

1. Add a service or edit service settings as described in Adding or Editing a Service (page 14-4).
2. When finished specifying service settings on the Add Service or Edit Service dialog box, click the Advanced button. The Advanced Service Settings dialog box displays (Figure 14-2).

Figure 14-2: Advanced Service Settings Dialog Box

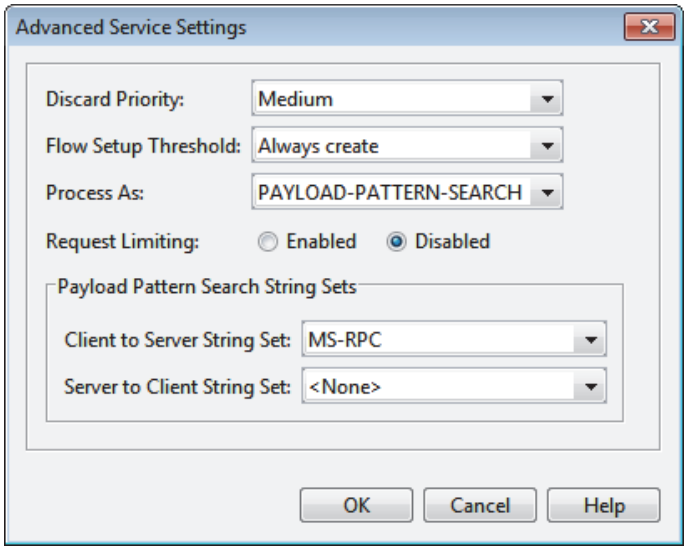

3. Specify the Discard Priority.

   This value indicates how likely it is that the Corero Network Device will discard packets for this service during periods of extremely heavy traffic. Selecting a higher priority specifies that you want the traffic associated with this service to be *more* likely to be discarded.

4. Specify the Flow Setup Threshold.

   This parameter indicates how critical it is to record connection information for this service. The Corero Network Device uses a Flow Table to record certain state information about each traffic connection. Under extreme traffic conditions, the Flow Table could become full.

   This parameter enables you to specify how critical it is to record connection information for a specific service. By not creating an entry in the Flow Table, you preserve the table for more critical services during periods of heavy traffic. For each service, you can choose whether you never want to create an entry in the Flow Table, or whether you always do. Alternatively, you can specify that the Corero Network Device create an entry in the table only if the table is less than a specific percentage full. This way, as the table fills, fewer flows are set up for lower priority traffic.

CAUTION

Improperly modifying this setting can adversely affect system operation. Before changing this setting, contact Corero Network Security.

5. Specify the Process As protocol settings.

   For this parameter, select the type of protocol expected for this service. The Corero Network Device uses the checks for this protocol to perform Deep Packet Inspection. If the service uses one of the protocols listed in the Process As drop-down list, select the protocol to identify the packet's expected contents to the device; otherwise select None to disable Deep Packet Inspections.

   You can choose from the following options

   - None — The protocol is not one of those listed. Do not perform Deep Packet Inspections for this service.
   - For an IPS Unit, you can select from the following protocol types:

AOL-IM	CIFS	DNS	ECHO
FTP	HTTP	HTTP-IM	LDAP
MSN-IM	MSRPC	NETBIOS	OTHER
PAYLOAD - PATTERN - SEARCH	PPS	SIP	SNMP
SMTP	SSH	SSL	Telnet
TFTP	YAHOO-IM		

6. Select either the Enabled or Disabled radio button for Request Limiting.
7. If you selected Payload Pattern Search from the Process As drop-down menu, specify the Payload Pattern Search String Sets.

   These are sets of search strings that the Corero Network Device uses when examining traffic. You can choose a different set of search strings for this service depending on whether the traffic is going from client to server, or server to client. You can select separate search string sets for client-to-server traffic and server-to-client traffic.

   NOTE

   You can create user-identified string sets (payload signatures) from the IPS Rules Customization dialog box. For more information, refer to Attack Signatures Overview (page 15-19).

8. When finished, click OK.
9. When you have finished specifying service settings in the Corero Network Device management application, click Done at the bottom of the Configure Security Policies dialog box. If you want to save the service and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
10. Save your changes by clicking the Save Configuration toolbar button.

Deleting a Service

At some point, you may want to delete a service. You may want to do this, for example, if you have created a service and associated it with a subset of IP addresses for policy use, but you now want to modify the policy to apply across the whole Corero Network Device. In this case, you would modify the policy to apply to the system-wide service, then delete the custom service you were previously using.

To delete a service:

1. Access the Services tab as described in Viewing Services (page 14-2).
2. Select the service you want to remove, then click Delete.
3. Save your changes by clicking the Save Configuration toolbar button.

Chapter 15
Managing Rules and Rule Sets

Corero Network Devices perform a majority of their security policy operations based on intrusion protection system rules (IPS rules). These rules govern how the Corero Network Device examines and treats traffic that is allowed by the Firewall subsystem.

Corero provides predefined rule sets that will be applicable to most users' requirements, but you can also modify existing rule sets, or create your own.

Not all of these rules apply to every traffic situation, and the treatment you should apply to traffic that triggers one of these rules may not always be the same, so carefully consider the policy you define.

This chapter contains the following sections:

About Rules

Rules are designed to sense particular conditions that may be malicious in origin. Some rules represent very strict conditions that you may not want to apply to every host, while other rules represent known definite issues that should be universally applied.

CAUTION

Improperly modifying rule settings can adversely affect system operation. Consider consulting with Corero Network Security before modifying these settings.

Rules are a primary building block of security policies. You can create a named set of rules, modify the treatment applied to individual rules within the set, and then use the named rule set as a building block when creating Firewall + IPS Rule policies.

There are two basic types of rules: packet-based rules and rate-based rules. Packet-based rules are applied based on Layer 2 network traffic. The settings for these rules have been optimized by Corero, and they should not generally be changed under normal system operation. Rate-based rules are applied based on the rate at which requests, connections, or network communications are sent to the Corero Network Device. You may choose to disable some of these rate-based rules which are applied system-wide.

You can view detailed information about each rule including its identifier (which starts with tln-) and a description of the rule's purpose.

Security Event Category

Every rule begins with a six character security event category prefix that allows for sorting by rule prefix anywhere a list of rules displays. Rule prefixes are listed in Table 15-1.

Table 15-1: Security Event Categories (Rule Prefixes)

Prefix	Description
AAUPV:	Acceptable application usage policy violation or condition match
DDOSA:	Rate-based attack
EXPLT:	Attempt to exploit a known vulnerability
FWALL:	Firewall policy violation or condition match
NETWK:	Network behavior issue - an IP, UDP, or TCP issue
OTHER:	Another issue or interest
PROTO:	Protocol anomaly or violation
RATEV:	Rate-based policy usage violation concerning connection limits or client rate limits
RECON:	Reconnaissance in the form of port scans or sweeps
RRBDx:	Request response behavior - DNS
RRBHx:	Request response behavior - HTTP
SPYWR:	Spyware was found in the inspected body
TROJN:	Trojan or backdoor program
VIRUS:	Virus and/or worm in executable file

In addition to a security event category, each event has a confidence category, from one to three stars, that designates how likely it is that the rule will yield false positives.

Rules are organized into default rule sets based on their predetermined confidence level. The confidence level reflects how likely the system is to trigger a false positive, which might misidentify traffic as malicious or troublesome when it is not.

Confidence Levels

Rules and rule sets have a predetermined confidence level. The confidence level reflects how likely the rule is to trigger a false positive, which might misidentify traffic as malicious or troublesome when it is not. Confidence levels are described in Table 15-2.

Table 15-2: Confidence Levels

Level	Definition
Three Stars	Rules with three stars are considered **safe.** These rules should never cause a false positive. Three star rules that trigger events should always be considered malicious.
Two Stars	Rules with two stars are **recommended**. These rules seldom trigger false positives. Two star rules that trigger events should always be considered suspect and probably malicious.
One Star	Rules with one star should be used for **strict** enforcement. One star rules may provide false positives, but they are extremely useful in situations where you are leery of a particular host group.
No Stars	Rules with no stars are **extremely targeted**. They are to be used in special cases, under certain conditions, or for specific, restricted situations. When used indiscriminately, these rules can easily trigger false positives.

User-Modifiable Rule Settings

In addition to the specific issue it is designed to detect, every rule has user-modifiable treatment settings that affect whether or how that rule is used by a security policy to detect malicious traffic. These settings include:

Rules are designed to be very granular in their detection abilities. For example, not all rules identify known bad traffic. Some identify when too much traffic is coming in, whether good or bad. And, in some cases, what might be bad traffic from one client is acceptable traffic from another client.

For information on how to edit the settings for a given rule, refer to Modifying Rule Settings (page 15-14).

NOTE

Some rules cannot be modified by the user.

For both the default rule sets and the rule sets you create, you can modify the following characteristics on a rule-by-rule basis:

- Status: Whether the rule is enabled or disabled.
- Actions: The action that the Corero Network Device should take if the rule is triggered.
- Logging options: Including whether the traffic that triggered the rule should be copied to the discard port.

NOTE

Your treatment modifications apply only to the instance of that rule in the named rule set in which you edited it.

Status

Every rule set always contains every rule. The way to control which rules are applied in a specified rule set and which are not is by modifying the status of the rule for that particular rule set. You can either enable the rule, which means it will be applied for that policy, or you can disable it.

Actions

An action is a response by the Corero Network Device when traffic triggers a security rule. It is part of the treatment portion of a security policy. When traffic through the device triggers a rule, the device can take one of the following actions

- Allow— Pass the traffic.
- Drop— Passively block the traffic by dropping the traffic with no indication to the offending client.
- Reject— Actively block the traffic by dropping the traffic and taking specific action, for example, by resetting the connection.

When you create the IPS portion of a Firewall + IPS policy, you apply a rule set to the IPS part of the policy. Each rule in the rule set has an action associated with it, which you can modify.

In the case of the firewall portion of the Firewall + IPS policy, the same action applies to all of the firewall rules. For more information, see Elements of a Firewall + IPS Security Policy (page 11-9).

Logging Options

Logging options are also part of the treatment portion of a security policy. They specify the reporting actions the Corero Network Device should take when traffic triggers a rule. Logging options include sending information to a log file based on its severity rating, and copying the associated traffic to the discard port.

You can specify logging options for all traffic that triggers an IPS or rate-based rule, even if the action you choose for that rule is “allow”. Discovering when traffic triggers a rule enables to you record usage information about applications you allow as well as those you block.

Limit Profiles for Rate Based Policies

Rate based security policies contain host groups and rules, but each policy also contains limit profiles.

The following types of rate limit profiles are part of the policy:

- Connection Limits: This feature limits the number of simultaneous connections allowed for a host group, and for individual members of the group.
- Client Request Limits: This feature limits the traffic the system will accept from each client in the group. Client request limiting is performed based on the client, server, and service associated with a particular packet.
- SYN Flood Limits
  Provides limits for the number of incomplete SYN requests for servers and for various categories of clients (trusted, suspicious, malicious, and so forth).

About Rule Sets

The Corero Network Device uses IPS rule sets to define its IPS policy operation for various categories of traffic. For example, you may want to apply a recommended set of rules to one group of servers and a strict set of rules to another group.

A named set of rules, including the treatment you configure for each rule in the set, is called a Rule Set. The device comes with several pre-defined rules sets such as RecommendedServerProtection and StrictServerProtection.

Although each rule set contains *all* of the IPS rules, rule sets vary in the following important ways:

- Some rules in the rule set may be disabled.
- The traffic control action (allow, drop, or reject) that the device takes when a specific rule triggers may vary by rule set.
- Logging options may vary for the same rule, based on the rule set.
- Each rule set can have an overriding set of parameters that implement all of its rules differently, based on the confidence level assigned to each rule in the rule set. For example, you could modify a rule set to only apply rules that have a confidence level of safe (three stars).

Although the Corero Network Device contains default rules sets, you can modify those rule sets. Or, if you prefer, you can copy a rule set and use the copy to create a custom rule set to use in your IPS policies.

IPS Units provide both Client and Server Protection rule sets.

Default Rule Sets

The Corero Network Device provides several default rule sets that you can apply to different situations when creating security policies. Table 15-3 describes the default rule sets.

Table 15-3: Default Rule Sets

Rule Set	Description
All Rules Block	Used as a benchmark check. Applies all rules to a host group's traffic and blocks any traffic that triggers the rule, regardless of the confidence level of that rule.
All Rules Detect	Used as a benchmark check. Enables you to see what rules are being triggered by a particular host group. The device logs the results, but passes the traffic.
All Rules Off	This could be used for a host group in which you have very high confidence, or for which you only want to apply firewall and DDoS protection. Traffic processing for this group will be somewhat faster, since the traffic is diverted past some of the device's subsystems.
Recommended Server Protection	Includes server-oriented rules from the three-star (Safe) and two-star (Recommended) confidence categories.
Strict Server Protection	Applies server-oriented rules from the three-star (Safe), two-star (Recommended) and one-star (Strict) confidence categories.
Recommended Client Protection	Includes client-oriented rules from the three-star (Safe) and two-star (Recommended) confidence categories.
Strict Client Protection	Applies client-oriented rules from the three-star (Safe), two-star (Recommended) and one-star (Strict) confidence categories.

Viewing Rule Sets

Corero Network Devices include a large number of intrusion protection system rules that govern how it examines and treats traffic that is allowed by the Firewall subsystem. Not all of these rules apply to every traffic situation. In addition, the treatment you should apply to traffic that triggers one of these rules may not always be the same.

For this reason, the device uses rule sets to define its policy operation for various categories of traffic. For example, you may want to apply a recommended set of rules to one group of servers and a strict set of rules to another group.

When you view a rule set, you are viewing a scrollable, searchable list of all IPS rules. Since each rule set contains all rules, the list of rules is the same for each rule set. However, the settings for individual rules will differ between rule sets.

To view a Rule Set:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.

2. Click the IPS Rule Sets tab. The Rule Sets tab displays.

   NOTE

   The Rule Set Membership table on the IPS Rule Sets tab provides a scrolled, searchable list of all the rules. The list is the same for each rule set, however, the parameters for individual rules will be different in different rule sets.

Figure 15-1 shows the IPS Rule Sets tab.

Figure 15-1: IPS Rule Sets Tab

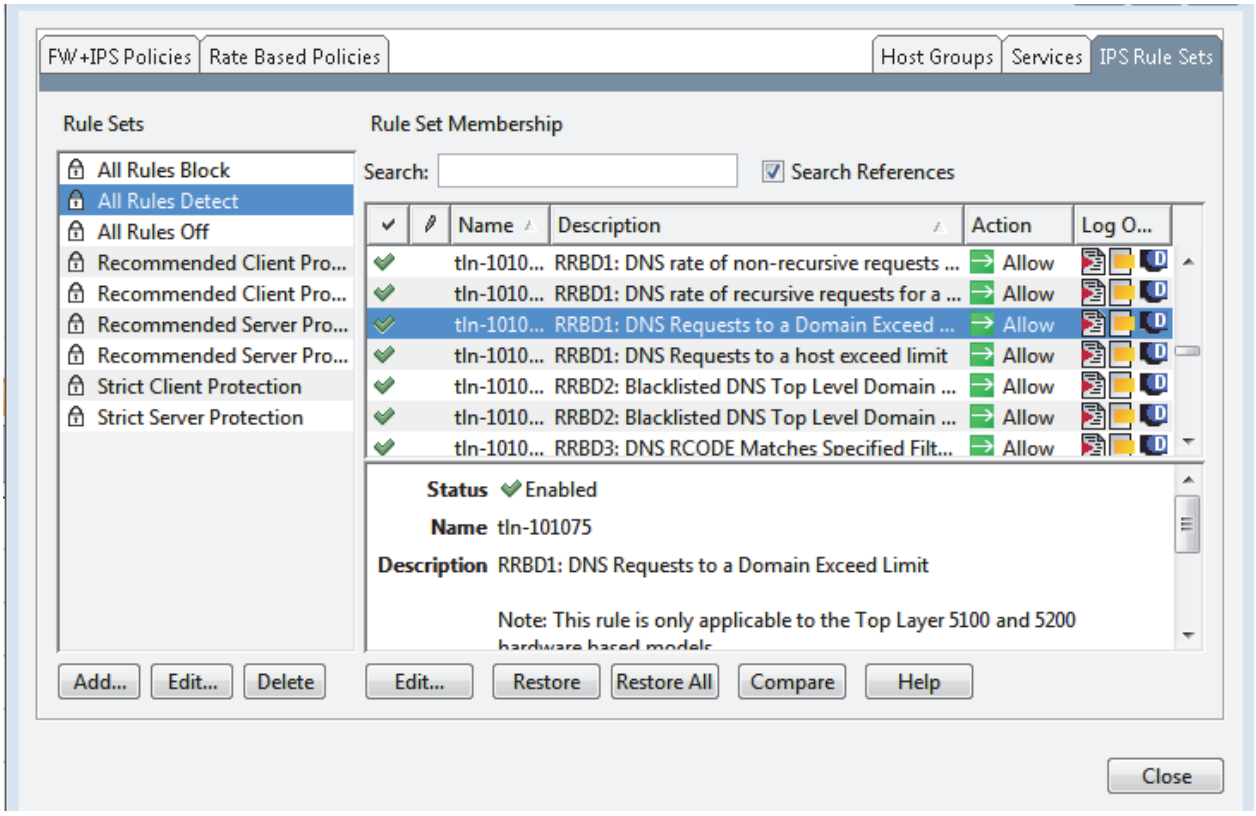

3. To view information about a rule set, select the set in the Rule Sets list. The rules comprising that rule set display in the right pane under Rule Set Membership.

   The Rule Set Membership table on the IPS Rule Sets tab provides a scrolled, searchable list of all the rules. The list is the same for each rule set, however, the parameters for individual rules will be different in different rule sets.

   The Rule Set Membership table displays the information listed in Table 15-4

Table 15-4: Rule Set Membership Table Contents

Column	Description
Status	• A green check mark indicates the rule is enabled for this rule set. • A grey X indicates the rule set is disabled, and is not applied for this rule set. When a rule in a rule set is disabled, traffic that triggers this rule is allowed through and is not logged.
Edited	A pencil icon in this column indicates a rule has been modified from its default configuration.
Name	The system-defined name (number) for this rule.
Description	Short description for the rule. Every rule begins with a security category prefix. See Table 15-1 for a listing of the prefixes.

Table 15-4: Rule Set Membership Table Contents *(Continued)*

Column	Description
Action	Action icons include: • Allow— Pass the traffic. • Drop— Passively block the traffic by dropping the traffic with no indication to the offending client. • Reject— Actively block the traffic by dropping the traffic and taking specific action, for example, by resetting the connection.
Log Options	The following icons indicate the logging options selected for this rule: • Log with Severity (color): Send information to a log file based on its severity rating: - Green indicates Low - Yellow indicates Moderate - Red indicates Critical • Copy to Discard Port —Copy the associated traffic to the Discard port.

4. The search window at the top of the table enables you to display a select set of rules from the entire list.

   To search for test in the name or brief description displayed in the Rule Set Membership list, enter the information in the Search text box.

   After you enter a search string, you can choose more thorough searches by selecting one or both of the following check boxes:

   - Search References —Searches the rule title.
   - Search Full Description —Searches the longer descriptions found in the material displayed in the bottom pane.

     NOTE

     If you choose to search the full description, be aware that you may need to scroll through the bottom pane to see why a particular rule was included in your search results.

5. To view information for a specific rule in the rule set, click the rule, and the information displays below the Rule Set Membership table (Figure 15-4). For more information on managing rules, see Rules Customization (page 15-22).

Managing Rule Sets

The Corero Network Device includes a large number of intrusion protection system rules that govern how it examines and treats traffic that is allowed by the Firewall subsystem. Not all of these rules apply to every traffic situation and the treatment you should apply to traffic that triggers one of these rules may not always be the same.

You can create a new rule set by adding all of the desired rules manually, or you can create a new rule set that is a copy of an existing rule set which you can then modify.

In addition, although multiple rule set can contain the same set of rules, they can vary in the following important ways:

- Some rules in the rule set may be disabled.
- The traffic control action (allow, drop, or reject) that the device takes when a specific rule triggers may vary by rule set.
- Logging options may vary for the same rule, based on the rule set.
- Finally, each rule set can have an overriding set of parameters that implement all of its rules differently, based on the confidence level assigned to each rule in the rule set. For example, you could modify a rule set to only apply rules that have a confidence level of safe (three stars).

You can manage rule sets by adding, modifying, or deleting them.

To manage Rule Sets:

1. Access the Rule Sets tab as described in Viewing Rule Sets (page 15-6).

   NOTE

   The Rule Set Membership table on the IPS Rule Sets tab provides a scrolled, searchable list of all the rules. The list is the same for each rule set, however, the parameters for individual rules will be different in different rule sets.

2. To add a rule set:
   a. Under the Rule Sets list, click Add. The Add Rule Set dialog box displays.
   b. Specify a name for the new rule set.
   c. If you want to base the new rule set on an existing rule set, specify the rule set that you want to copy in the Copy of drop-down list.
   d. If you want to create additional rule sets, click Add. Alternatively, if you just want to create this one new rule set, click Done.
3. To modify a rule set:
   a. Select the desired rule set in the Rule Sets list, then click Edit. The Edit Rule Set dialog box displays (Figure 15-2).

Figure 15-2: Edit Rule Set Dialog Box

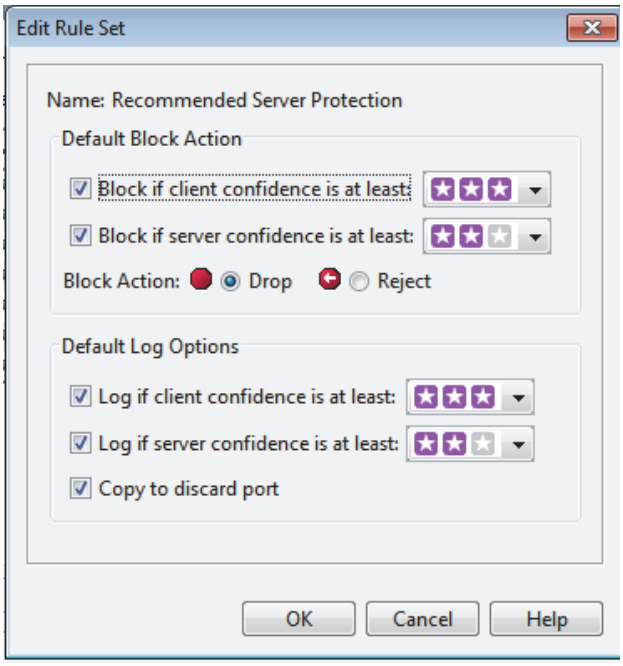

b. Specify the desired settings for the rule set. These settings are described in Table 15-5.

Table 15-5: Rule Set Parameters

Setting	Description
Block if Client Confidence is ...	This allows you to specify that all rules at or above a certain confidence rating result in blocked client traffic.
Block if Server Confidence Is ...	This allows you to specify that all rules at or above a certain confidence rating result in blocked server traffic.
Block Action	When traffic is blocked, you can choose how any packet that triggers the rule is treated. • Drop - This selection drops the packet, and, in many cases, blocks the flow so no more packets that are part of that connection can pass. • Reject - This selection ends a TCP RST packet to the client and server to attempt to kill the connection state on both the client and server.
Log if Client Confidence is ...	Enables you to specify that all rules at or above a certain confidence rating are logged when they are triggered by client traffic.
Log if Server Confidence is ...	Enables you to specify that all rules at or above a certain confidence rating are logged when they are triggered by server traffic.
Copy to Discard Port	Enables you to specify whether you want any packet that triggers the rule to be copied to the discard port, if one is configured.

4. To delete a rule set, select the desired rule set in the Rule Sets list, then click Delete. You are prompted to confirm your selection.
5. When you have finished specifying IPS rule settings in the Corero Network Device management application, click Done.

6. Save your changes by clicking the Save Configuration toolbar button.

Viewing Packet-Based and Rate-Based Rules

Corero Network Devices perform packet-based security checks that are designed to detect and eliminate traffic that is obviously malformed (either deliberately or through transmission problems). They also perform rate-based checks related to SYN Flood settings and connection limiting.

NOTE

For detailed information on SYN Flood and Connection limiting, see Chapter 20, "SYN Flood and Connection Limiting Security".

Packet-based rules are executed first, followed by Rate-Based Protection rules.

Figure 15-3 shows where these checks fit into the device's security subsystems.

Figure 15-3: Packet-Based Checks

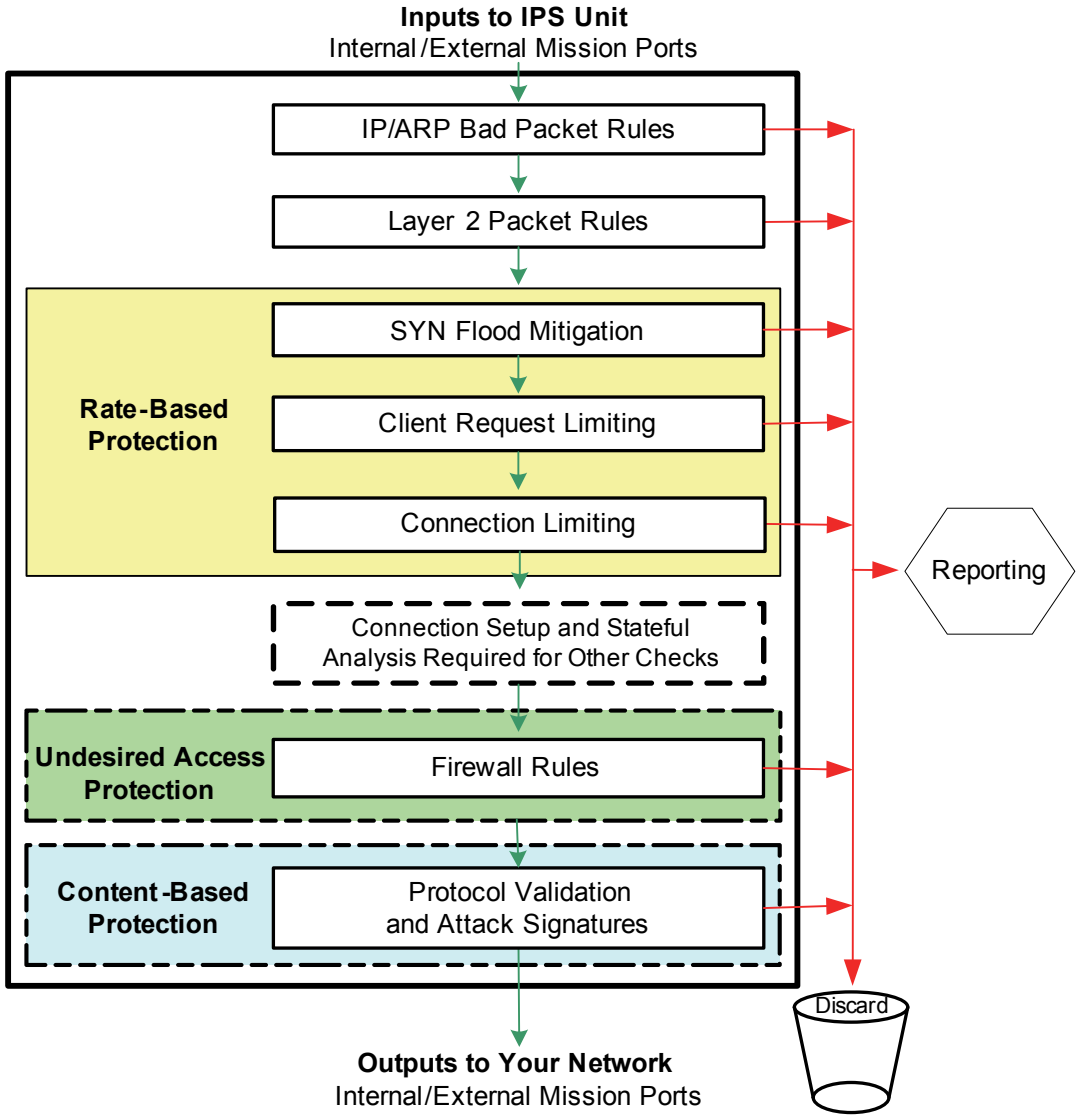

Since the packet-based checks typically detect mangled and defective traffic that should not be passed on to your network, most users will never want to disable the rules that control these checks. However, it is possible to change the setting for these checks on a rule-by-rule basis. For example, you may want to alter the logging treatment for a given rule.

To view rules:

1. Do one of the following:
   - From the Navigation Tree, choose Configure Security > Advanced Security Config > Packet Based Rules.
   - From the Navigation Tree, choose Configure Security > Advanced Security Config > Rate Based Rules.

   The appropriate rule dialog box displays. Figure 15-4 shows the Packet Based Rules dialog box, but the Rate Based Rules dialog box displays similar information.

Figure 15-4: Packet Based Rules Dialog Box

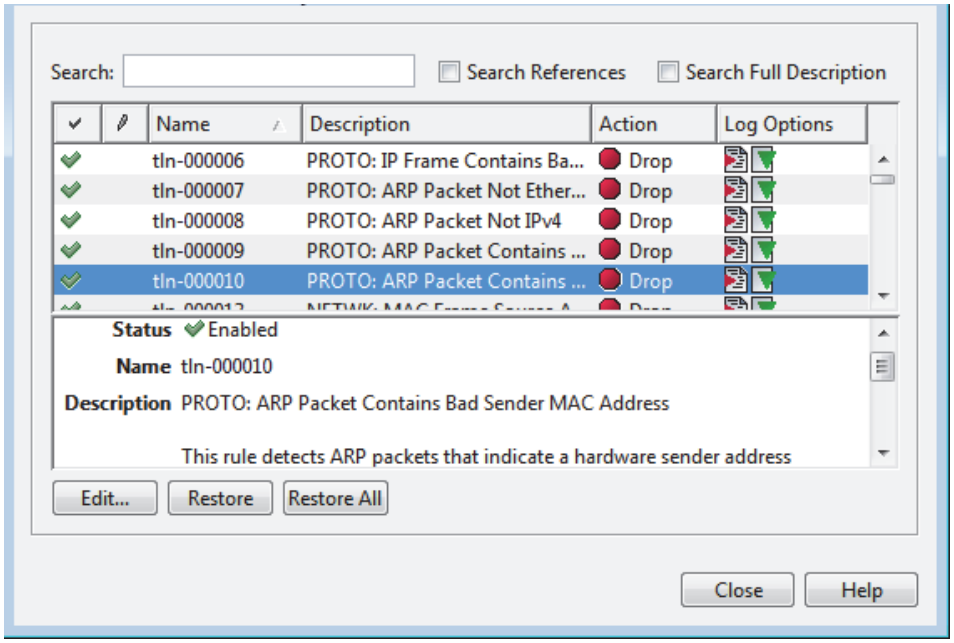

2. To view information on a specific rule, select the rule in the list. Rule information displays in the lower portion of the dialog box.

Modifying Rule Settings

To modify rule settings:

1. From the management application for a Corero Network Device, do one of the following:

In order to...	You must...
Modify packet-based rule settings	from the Navigation Tree, choose Configure Security > Advanced Security Config > Packet Based Rules. The Packet Based Rules dialog box displays.
Modify rate-based rule settings	From the Navigation Tree, choose Configure Security > Advanced Security Config > Rate Based Rules. The Rate Based Rules dialog box displays.
Modify IPS rule settings	1. Click the Security Policies button on the toolbar. The Configure Security Policies dialog box displays. 2. Click the IPS Rule Sets tab. 3. Select a rule set containing the desired instance of the rule.

2. To modify the settings for a rule, select the rule, then click Edit. The Edit Rule Settings dialog box displays (Figure 15-5).

Figure 15-5: Edit Rule Settings Dialog Box

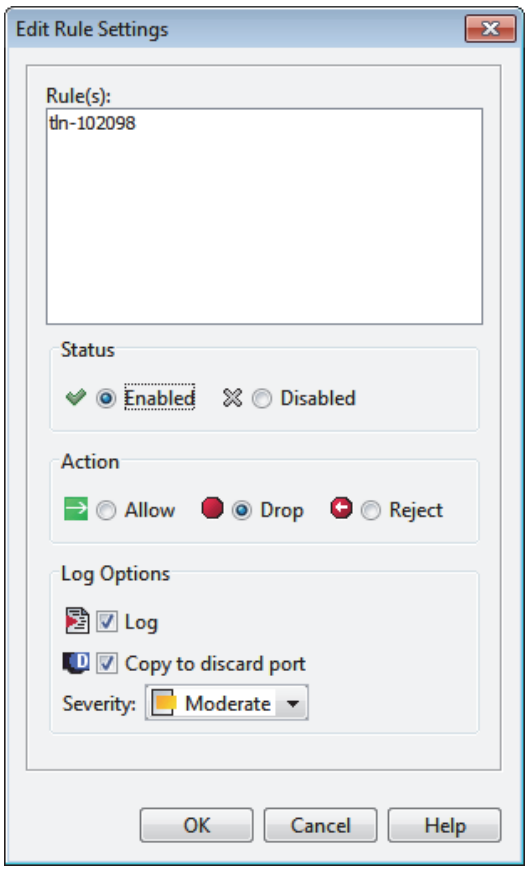

3. You can enable or disable individual rules.

   This setting applies only to the selected rule in the chosen rule set, and it overrides the settings established using the Edit a Rule Set window.

   If you disable a rule, packets that would normally trigger the rule are passed, and no logging is performed.

   If you enable a rule, you can set its Action. Possible actions are:

   - Allow— Pass the traffic.
   - Drop— Passively block the traffic by dropping the traffic with no indication to the offending client.
   - Reject— Actively block the traffic by dropping the traffic and taking specific action, for example, by resetting the connection.

4. Specify the Log options for this rule as follows:
   - Log - Send information to a log file based on its severity rating.
   - Copy to Discard Port - Copy the associated traffic to the Discard port based on its severity rating.

     NOTE

     If you have configured a Discard Port, you can log all traffic that triggers a rule, even if the Action for that rule is set to Allow. This feature enables you to record usage information about applications you allow as well as those you disallow.

   - Severity - The log severity determines whether packets for this rule are processed by the device's Event Logging System. To be processed, the priority you set for this rule must be equal to, or above, the overall priority set for event messages, as described in Logging Options (page 15-4). Severity levels include:
     - Low (green)
     - Moderate (yellow)
     - Critical (red)
5. When finished, click OK. When the rule appears in the list of rules, an icon displays indicating that this rule has been modified from its default settings.
6. If you modify an individual rule's settings, you can restore the settings to the factory default values for Status, Action, and Log Option. To do so:
   a. Select the rule.
   b. Click Restore.
   c. You are prompted to confirm your selection.
7. If you have modified multiple rules in the rule set, you can reset all rules in the rule set to the factory default values for Status, Action, and Log Option. To do so:
   a. Click Restore All.
   b. You are prompted to confirm your selection.
8. When you have finished specifying rule settings in the management application for a Corero Network Device, click OK.
9. Save your changes by clicking the Save Configuration toolbar button.

   NOTE

   You can also specify a limit to how many events can be sent per rule per minute to the logs and the Security Event Viewer. This helps ensure these event listings are not

overwhelmed by frequent triggering of a single rule. The default limit is 300 events per rule every minute. You can only access this setting from the Blocked and Detected Attacks page. For more information on modifying this setting, see Viewing Blocked and Detected Attacks (page 19-16).

Comparing Two Rule Sets

If you want to view the differences between settings on two rule sets:

1. Using the management application for a Corero Network Device, click the Security Policies button on the toolbar. The Configure Security Policies dialog box displays.
2. Click the IPS Rule Sets tab.
3. Select a rule set in the Rule Sets list.
4. Click Compare. The Compare IPS Rule Sets dialog box displays.
5. From the drop-down list, select the IPS rule set to which you want to compare the rule set you selected. The table lists all rules, and indicates any differences between the disposition for an individual rule between the two rule sets.

Restoring Rules to Default Settings

At some point, you may want to revert a rule in a particular rule set to its factory settings. To do so:

1. Do one of the following:
   - From the Navigation Tree, choose Configure Security > Advanced Security Config > Packet Based Rules.
   - From the Navigation Tree, choose Configure Security > Advanced Security Config > Rate Based Rules.
2. Click the IPS Rule Sets tab.
3. Select the desired rule set in the Rule Sets list.

   NOTE

   You can only restore rule settings for one rule list at a time.

4. Do one of the following:
   - To restore a single rule to its factory default settings, select the rule in the Rule Set Membership list, then click Restore.
   - To restore all rules in the Rule Set Membership list to their factory default settings, click Restore All.
5. You are prompted to confirm your selection.
6. When you have finished specifying rule settings in the management application for a Corero Network Device, click Done.
7. Save your changes by clicking the Save Configuration toolbar button.

Attack Signatures Overview

This feature provides the ability to search the payloads of network protocols that are not natively parsed and decoded by the Corero Network Device. The device's rules that correspond to these content patterns are also referred to as “signatures”. These patterns can be either case sensitive or case insensitive ASCII printable character strings or a sequence of binary bytes.

The rules that correspond to these content patterns are also referred to as signatures. You can also define your own payload signature patterns, which provides another way to define the various limits and parameters. These signatures enable you to search the payloads of network protocols that are not natively parsed and analyzed by your Corero Network Device. These patterns can be case sensitive or case insensitive ASCII printable character strings, or a sequence of binary bytes.

The IPS Unit provides 32 string set patterns, of which 9 can be user-defined.

The pattern set that the device uses to monitor a flow is based on the association of the network application to one of the defined string sets. For example, the string set used for the Finger protocol is different than the string set for the Microsoft RPC protocol.

Pattern Formats

The patterns that are supported by the engine can be either case sensitive or insensitive ASCII printable character strings or they can be a sequence of binary bytes. Binary bytes are entered in hex surrounded by the | character. For example, the following are valid patterns. The first is ASCII text, the second is entered as binary bytes, and the third is a combination of the two:

select/**/

|61 6E 69 68 A8|

select|61 6E 69 68 A8|/**/

Number of Strings Supported Depends on Total Length of All Strings

The number of strings supported is a function of the number of overall search bytes defined. Generally, the String Search Engine (SSE) will support up to 512 patterns that are each 32 bytes in length. The maximum pattern that can be specified is 64 bytes.

String Search Engine Pattern Matching

The SSE will start the search for patterns at the first byte of TCP or UDP payloads. The search will continue across the “stream” of bytes associated with the flow. The searching function is stateful and can stop/resume at any arbitrary byte in the stream. Packet, fragment and segment boundaries will not affect the searching operation. The SSE does not provide the ability to specify the search operation at a starting offset or stream depth in this release.

Actions for Matched Strings

Each string has an individual IPS rule identifier located in the Protocol Checks subsystem. As such, the disposition for each string matched will be under full policy control in this subsystem. If an SSE pattern is matched, the policy action can be to ignore, detect or block the traffic associated with the flow. If the action is to block the flow, the packet is dropped, as are all remaining packets for the flow. If the policy is to detect, the rule identifier is reported in the event and the SSE resumes pattern matching on the stream. It is possible that the SSE will report a subsequent pattern match for the same flow. In all cases, only one block event will be supported per flow.

The presence of a signature in this table will be treated as a filter specification. Each payload signature must be associated with an existing string set name selected from the Payload Signatures Set table.

Managing Attack Payload Patterns

The management application enables you to view or modify attack signature payload patterns. A pattern is comprised of a signature name, a specified string set, and an ID label.

For information about modifying string sets, see Payload Signature Sets (page 15-21).

To manage attack signature payload patterns:

1. Using the management application for a Corero Network Device, choose Configure Security > Advanced Security Config > IPS Rules Customization from the Navigation tree.

   The IPS Rules Customization dialog box displays (Figure 15-6).

2. View available patterns by choosing Attack Signatures > Patterns in the Filters area.
3. To modify a pattern, select the Signature name, then click Edit.
4. In the Signature field, specify the signature to be used as a filter.
5. Select the String Set you want to use for this payload pattern. If the String Set you want is not available, you can modify an existing String Set as described in Payload Signature Sets (page 15-21).
6. Specify an ID label, which is a text string that describes the signature pattern.
7. Do one of the following:
   - If you are finished adding patterns, click Done.
   - If you wish to add another pattern, click Add. The pattern you created will be saved, and the dialog box will remain so you can enter another pattern.
8. Click Apply to save your pattern changes.
9. When you have finished specifying patterns in the Corero Network Device management application, click Done. If you want to save the service and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
10. Save your changes by clicking the Save Configuration toolbar button.

WARNING

When you add patterns (signatures) to a string set, you must click the Apply button on the IPS Rules Customization dialog box if you make any changes; otherwise, these changes will not be saved.

Payload Signature Sets

When performing pattern matching tasks, the strings that the Corero Network Device searches are gathered into sets of related searches. For example, there is a pattern set named FINGER for use with the FINGER service. For the string sets used in pattern matching, you can edit the name of the signature sets or change the case sensitivity. Note that you cannot add string sets, because the system has a predetermined number of them, but you can modify their contents.

To edit the name of the signature set or change the case sensitivity:

1. Using the management application for a Corero Network Device, choose Configure Security > Advanced Security Config > IPS Rules Customization from the Navigation tree.

   The IPS Rules Customization dialog box displays (Figure 15-6).

2. View available string sets by choosing Attack Signatures > Sets in the Filters area.
3. To modify a signature set, select the Signature Set name, then click Edit.
4. If desired, enter a new name for the signature set.
5. Use the check box to indicate whether the signatures in this set are case sensitive.
6. Click OK.
7. When you have finished specifying rule settings in the management application for a Corero Network Device, click Close.
8. Save your changes by clicking the Save Configuration toolbar button.

Rules Customization

The IPS Rules Customization selection from the Navigation Tree provides access to a large number of windows that enable an advanced user to customize protocol-related and signature-related security settings. These settings are automatically managed when you download new configuration settings from Corero (assuming you have subscribed to the TopResponse™ Service), but if needed, you can modify parameters yourself. By default, these options are set to satisfy most network requirements.

CAUTION

Corero recommends that you contact the Customer Services Center if you want to modify customization settings. Improper settings can negatively affect system operation and any associated network traffic.

You can customize rules in several areas:

- Network Protocols:
  You can specify information such as maximum ping and ICMP lengths, TCP midflow blocking settings, and configure permissions for IP options.
- Protocol Validation Modules:
  You can specify parameters for application protocols, including DNS, FTP, HTTP, SSH, SMTP, MSNET (RPC and CIFS), and Telnet.
- Attack Signatures:
  This feature provides the ability to search the payloads of network protocols that are not natively parsed and decoded by the device.

To view or modify IPS rule parameters:

1. Using the management application for a Corero Network Device, choose Configure Security > Advanced Security Config > IPS Rules Customization from the Navigation tree.

   The IPS Rules Customization dialog box displays (Figure 15-6).

Figure 15-6: IPS Rules Customization Dialog Box

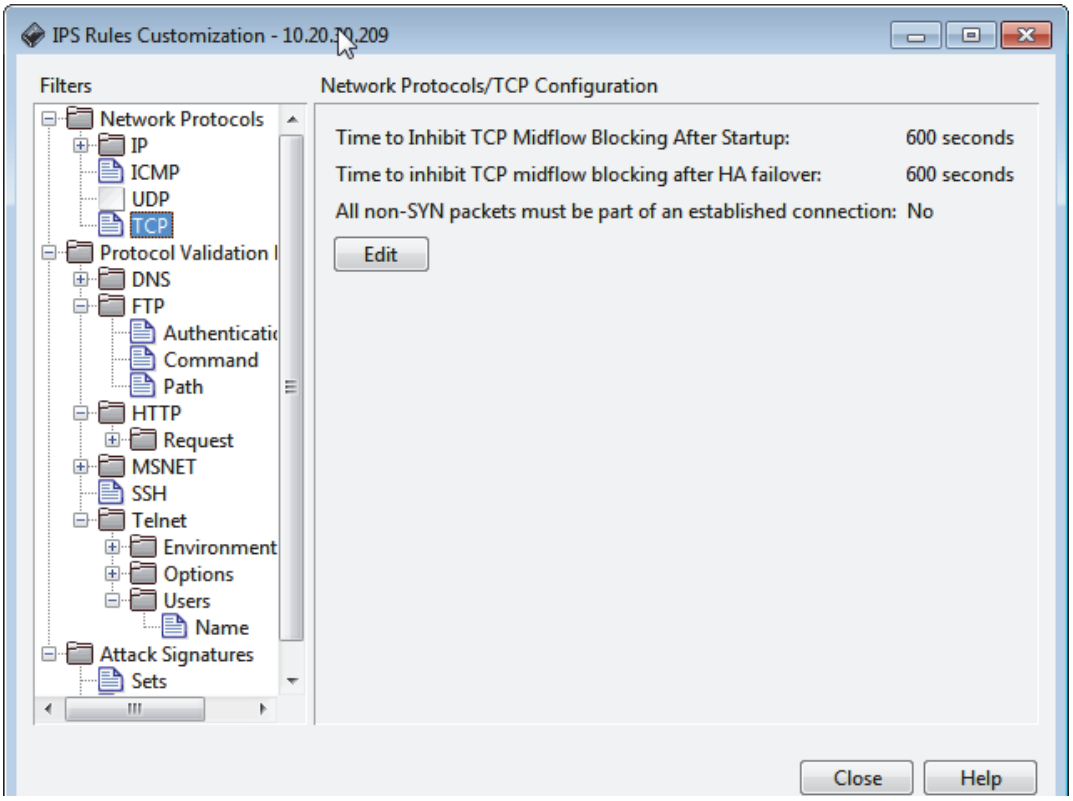

2. Select the parameter in the left pane and view its settings in the right pane.
3. If desired, click Edit, and modify the parameters.
4. When finished, click OK.
5. When you have finished specifying rule settings in the management application for a Corero Network Device, click Close.
6. Save your changes by clicking the Save Configuration toolbar button.

Chapter 16
Generating and Viewing Security Reports

You can generate reports for a Corero Network Device that help you understand system operation and the security-related decisions the unit makes while processing network traffic according to your security policies.

This chapter contains the following topics:

About Security Reports

You can generate preconfigured security reports that provide summary and detailed information about a device's security and general operations.You can configure reports on a scheduled basis (called a Periodic Report), or on-demand (called an Immediate Report).

Your Corero Network Device automatically allocates a specific volume of memory to store generated security reports. Once the allocated memory is fill, the device deletes old reports when space is needed to make room for new reports.

NOTE

All stored reports are lost when the Corero Network Device reboots.

You can configure the following items for security reports:

- The time and frequency when the report is generated, including enabling or disabling periodic report generation. Note that, even if periodic reporting is disabled, you can still generate an immediate version of the report.
- The level of detail for the report (based on the report template used to generate the report)

Table 16-1 lists the available report templates:

Table 16-1: Standard Report Templates

Report Name	Description
Complete Report	Provides summary and detailed information about security events, packets blocked, processor utilization, active traffic flows, port utilization, packet analysis, and system diagnostic information which includes SYN flood mitigation, CPU overload protection, and any device resources that had to be limited.
Standard Report	Includes all of the security information sections contained in the Complete Report, but does not include the system diagnostic information.
Security Overview Report	Provides summary information about security events, blocked packets, and Top 10 attackers.
PCI Compliance Report	Contains the results of the device's self-assessment for PCI DSS (Payment Card Industry Data Security Standards) compliance, and recommendations for PCI security remediation. **Note:** The PCI Compliance report indicates the state of the device‘s compliance with the applicable PCI DSS sections.

Understanding the Data Collection Periods for Security Reports

To obtain values for its various reports, your Corero Network Device records many different types of traffic and mitigation events using counters and gauges. For a given report period, the device compares the counter values at the beginning and end of the report period and uses the data to calculate total, current, average, and peak values. In some cases, the device also calculates rates. See About Security Reports (page 16-2) for more information.

Each generated report has two data collection points:

- At the report's start time, which is the beginning of the reporting interval.
- At the report's end time, when the report is generated.

The start and end times for each report type differ as follows:

- Periodic Report|
  The device compares the counter data beginning at the last time it took a data snapshot (the last time it generated a Periodic report), with the data at the time it generates the report. If the device has rebooted during that period, data collectors are reset and it uses those reset values as the start counter values.
- Immediate Report
  The data collection period begins at the last time the device generated a Periodic report (or the device rebooted if that occurred after the report), and ends at the point when you request the Immediate security report.

When you generate reports, consider the following:

- If you generate an Immediate report, it does not change the start point data used for the next Periodic report.
- If you disable generation of the Periodic report, the device continues collecting data based on the currently defined data collection interval.
- If you disable Periodic security report generation (the report settings specify generation times, but generation is disabled), the device continues to take counter snapshots at the configured intervals, and advance the start and end times, but does not generate the report. When you enable report generation again, it uses the data from the latest start interval and the next scheduled end interval.

Report Generation Schedule Example

In the following example:

- The user set the Periodic security report to be generated twice a day, at 12 PM and 6 PM.
- The user also requested two Immediate reports that covered portions of two Periodic reports' time intervals.

  Note that these requests did not reset the interval for the Periodic reports.

Table 16-2: Report Generation Schedule Example

Time	Action
6:00 AM	System Boot. The Corero Network Device takes a snapshot of its counters at system boot.
7:05 AM	The user specifies the times for Periodic reporting as 12:00 PM and 6:00 PM.
9:20 AM	The user generates an immediate report, which is generated using the following range: Start Time: system boot time End Time: 9:20 AM (the requested time)

Table 16-2: Report Generation Schedule Example *(Continued)*

Time	Action
12:00 PM	The system generates its first periodic report, which is generated using the following range: Start Time: system boot time End Time: 12:00 PM (the scheduled time)
5:00 PM	The user generates another immediate report, which is generated using the following range: Start Time: 12:00 PM (when the last periodic report was generated) End Time: 5:00 PM (the requested time)
6:00 PM	The system generates its second periodic report, which is generated using the following range: Start Time: 12:00 PM (when the last periodic report was generated End Time: 6:00 PM (the scheduled time)
7:15 PM	The user decides they would rather generate periodic reports at 10:00 AM and 4:00 PM, and changes the report settings accordingly.
10:00 AM	The system generates its third periodic report, which is generated using the following range: Start Time: 6:00 PM (when the last periodic report was generated) End Time: 10:00 AM (the scheduled time)

Security Report Contents

You can generate several different types of reports for Corero Network Devices. Not only can you choose from among several templates, you can also choose whether you want to generate a periodic (scheduled) report, or an immediate (on-demand) report.

Figure 16-1 shows the first page of a sample report.

Figure 16-1: Sample Periodic Security Report

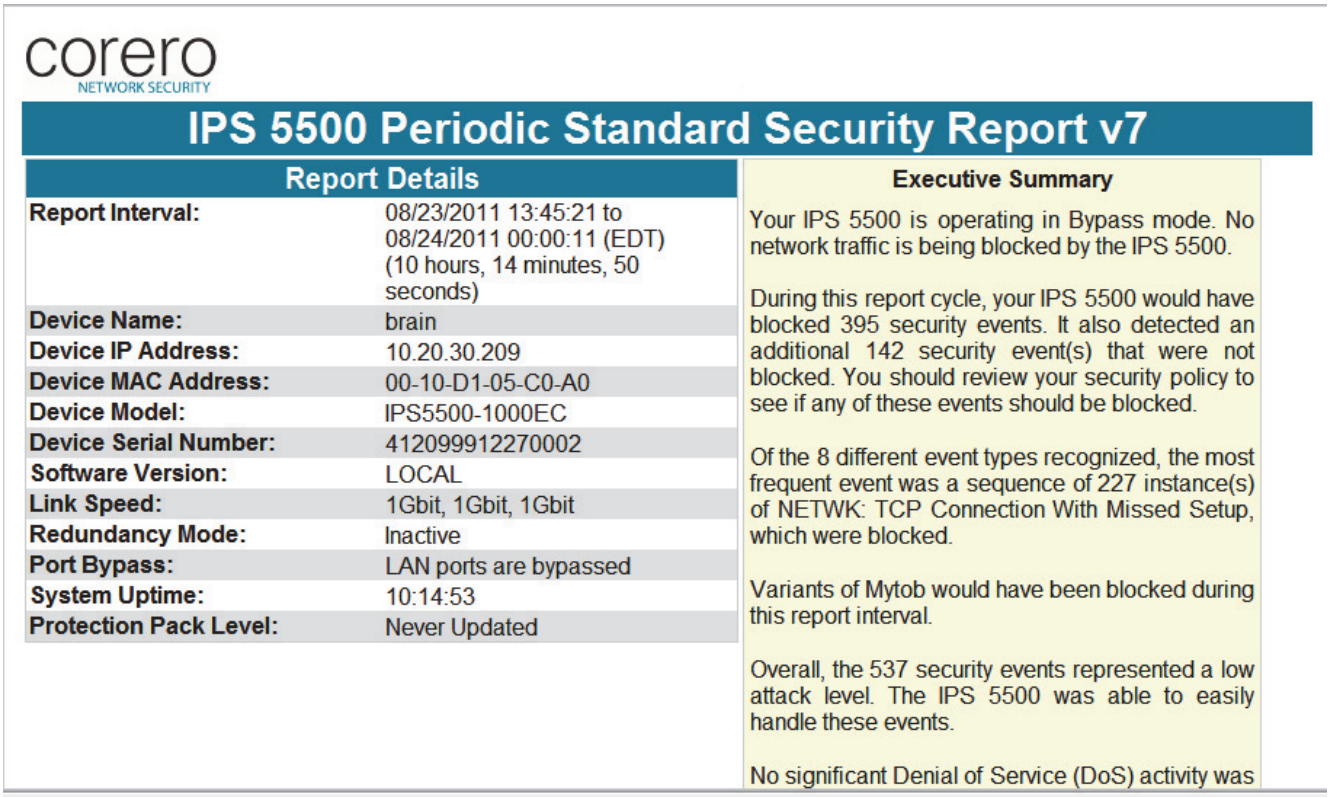

corero
NETWORK SECURITY

IPS 5500 Periodic Standard Security Report v7

Report Details	
Report Interval:	08/23/2011 13:45:21 to 08/24/2011 00:00:11 (EDT) (10 hours, 14 minutes, 50 seconds)
Device Name:	brain
Device IP Address:	10.20.30.209
Device MAC Address:	00-10-D1-05-C0-A0
Device Model:	IPS5500-1000EC
Device Serial Number:	412099912270002
Software Version:	LOCAL
Link Speed:	1Gbit, 1Gbit, 1Gbit
Redundancy Mode:	Inactive
Port Bypass:	LAN ports are bypassed
System Uptime:	10:14:53
Protection Pack Level:	Never Updated

Executive Summary

Your IPS 5500 is operating in Bypass mode. No network traffic is being blocked by the IPS 5500.

During this report cycle, your IPS 5500 would have blocked 395 security events. It also detected an additional 142 security event(s) that were not blocked. You should review your security policy to see if any of these events should be blocked.

Of the 8 different event types recognized, the most frequent event was a sequence of 227 instance(s) of NETWK: TCP Connection With Missed Setup, which were blocked.

Variants of Mytob would have been blocked during this report interval.

Overall, the 537 security events represented a low attack level. The IPS 5500 was able to easily handle these events.

No significant Denial of Service (DoS) activity was

Table 16-3 provides a description of the contents of each section of the available security report templates.

NOTE

The PCI Compliance report contains different information that is specific to its function.

Table 16-3: Security Report Contents

Section	Description	Security Overview Report	Standard Report	Complete Report
Report Details	Provides device, report interval, software, and other pertinent report background details.	X	X	X
Executive Summary	Provides a high-level overview of the device's performance and security events it encountered during the reporting period, including: • Total number of events detected • Total number of events blocked • Overall security event level based on average events per minute: - Low: Less than 10 - Moderate: Between 10 and 100 - High: Greater than 100	X	X	X
Security Events Blocked Summary	Events blocked, listed by name. Each blocked event is counted once per flow (connection). Later packets in the flow are also blocked and are included in the Blocked Packet Details section of the report. Once the flow is blocked, the time-out value for that flow's application is reduced to 30 seconds.	X	X	X
Security Events Detected Summary	Events detected but not blocked, listed by name. Events can be detected more than once for a given flow.	X	X	X
Top 10 Attackers	Lists IP addresses that appear most often as a source of attacks.	X	X	X
Blocked Packet Summary	Provides a diagram of the various security subsystems and the total number of packets blocked by each subsystem. It also shows the total received and transmitted packets. This section helps you quickly zero in on the most common areas of concern and how serious the attacks are.	X	X	X
System Processor Utilization	Provides usage statistics for the device's main traffic handling processors. Breaks CPU utilization into categories and provides the following numbers for each category. **Note:** You should be concerned if peak CPU usage in the Total CPU Usage category is near 100%. This indicates either an attack or a very high network load.		X	X
System Session Table Usage	The System Session Table holds state information that the device uses to analyze the packets in a flow. The report gives the maximum number of flows for which the table can store state details. That total depends on the model you have installed. Flow statistics are summarized by the type of flow information the table holds: Total flows, TCP, UDP, IP, and Reserved flows. Flows are not created for ICMP traffic.		X	X
System Session Active Flows	Displays all services that currently have active flows in the System Session table.		X	X
System Session Setup Rate	Displays the rate of flow set up per second for various types of flows. If the total of all flows set up per second nears 50,000 flows, the device is considered very busy.		X	X

Table 16-3: Security Report Contents *(Continued)*

Section	Description	Security Overview Report	Standard Report	Complete Report
System IP Address Summary	The report gives the maximum number of hosts for which a Corero Network Device can assess the threat level. That total depends on the model you have installed. The report provides current, average, and peak statistics for IP addresses in each of the threat level categories: • Unknown — The host IP address is known, but the device has not yet determined its threat level. Depending on your settings, the device may or may not proxy requests from hosts in this category. • Trusted — Behavior of these hosts is within acceptable boundaries. The device sends requests from these hosts on to their destination servers to handle. • Suspicious — These hosts have a suspicious level of incorrect behavior. The device proxies their requests on behalf of the intended server. • Malicious — The device has determined that these hosts are behaving maliciously. It drops requests from these hosts. A second table tracks new hosts seen during a period of distributed denial of service (DDoS) attacks. During an attack, the device places these new hosts into a separate table and requires them to “prove themselves” by demonstrating reasonable back off retry times before it adds them to the Unknown hosts of the regular IP address table. A non-zero entry in this table would indicate that the device entered DDoS rejection mode and was dealing with a DDoS attack or attacks.		X	X
LAN Port Utilization	Provides the percent utilization for each of the mission ports you have defined. Also provides actual transmit and receive packet counts for each of the mission ports.		X	X
Packet Analysis Details	Number of packets received and transmitted on each Mission port. An important number to watch is “Total packets dropped due to resource depletion” which is normally zero. A non-zero value here indicates an issue which could be memory, table space, or simply an extremely high volume of traffic on Gigabit Mission ports. **Note:** The total packets received on all mission ports will not exactly match the total of the blocked, dropped, and transmitted packets because the data on each port is collected at a slightly different time. Also, in Bridging mode (Port Pair Forwarding disabled), unknown unicast packets can be transmitted on more than one port.		X	X

Table 16-3: Security Report Contents *(Continued)*

Section	Description	Security Overview Report	Standard Report	Complete Report
SYN Flood Mitigation Details	Provides details on the rate at which packets were dropped due to SYN Flood mitigation activities. Dropped packets fall into the following categories: • Malicious client blocked — Packets dropped per second because the device determined that the client is acting in a malicious manner. • Client TCP handshake failed — Client did not complete the TCP handshake process within thirty seconds. • Server TCP handshake failed — Client completed the handshake process but the server did not. The server could be overloaded, or could be receiving SYN requests for an application that it does not handle. • No proxy queues available — A SYN flood attack or system overload caused the device to temporarily run out of proxy queues and it was unable to proxy some server requests, and dropped those requests packets. • DDoS rejection — During a DDoS attack, the device requires new clients to pass a “well-behaved” test before it will process their packets. This count indicates how many packets per second the device dropped because it entered DDoS Rejection mode and was applying this test to new clients.			X
CPU Overload Protection	The Corero Network Device invokes a CPU overload protection mode when the Forwarding Engine cannot keep up with the packets arriving. The report indicates the number of times that the device entered into each level of CPU protection. Protection levels are: • Level 1: New Session Setup Suspended — The device briefly stopped setting up new sessions. The device continues to process packets that match existing sessions. • Level 2: Packet Forwarding of Existing Session Suspended — The device briefly stopped setting up new sessions, and briefly stopped forwarding packets for existing sessions. • Level 3: Packet Reception on Gigabit Ports Suspended — In addition to Level 1 and Level 2 behavior, the device briefly shut off one or more Gigabit ports. If the device did not need to use any of the protection levels, it displays the following message: During the report interval, the device did not invoke any CPU Overload Protection mechanisms. **Note:** The number of packets dropped due to CPU overload protection is listed in the System Resource Limits Exceeded section of the report.			X

Table 16-3: Security Report Contents *(Continued)*

Section	Description	Security Overview Report	Standard Report	Complete Report
System Resource Limits Exceeded	Lists the number of packets dropped due to limited resources within the device. The report presents counts for the following resource limit events: • Link outbound congestion — Output queue for the port is overrun. • No flood descriptor for multicast packet — Flooding resources exceeded. • CPU Overload Protection — Packets dropped while in CPU Overload Protection. • SYN Flood: No proxy queue available — The Corero Network Device has temporarily used up all available proxy queues. Additional requests cannot be proxied and are dropped. **Note:** Setting the device to proxy hosts in the Unknown threat level state uses up some of the available proxy queues. If there were no dropped packets due to limited resources, the report contains the following statement: During this report interval, no resource limits were exceeded.			X

Generating an Immediate Security Report

To generate an immediate security report:

1. Click the Immediate Security Report toolbar button.

   The View Immediate Security Report dialog box displays.

2. Select the desired Report Template. Report templates are listed in Table 16-1.
3. Click View. The Immediate Report displays in a web browser.

Specifying Periodic Security Report Settings for a Corero Network Device

You can instruct the management application to automatically generate a report based on a specific report template at a scheduled time each day. When specifying periodic security report settings, consider the following:

- Even if you disable generation of the periodic report, data collection continues, based on the data collection interval you have configured.
- If you generate an immediate report while periodic report generation is disabled, the Corero Network Device uses the current portion of the collected data to generate the Immediate report. For more information, see Understanding the Data Collection Periods for Security Reports (page 16-3).
- A Corero Network Device holds up to two immediate reports for each report type (report template), and up to seven periodic security reports for each report type (report template). When the number of a particular type of report has reached its limit, old reports of that type are deleted as new reports of that type are generated.

To specify settings for a periodic report using the management application for a Corero Network Device:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Report Settings.

   The Security Report Settings dialog box displays.

2. Modify the settings as desired. You can choose from three different methods of specifying the reporting interval:
   - Once a day at a time specified in hours and minutes.
   - Twice a day at times specified in hours and minutes.
   - Every 1, 2, 3, 4, 6, 8, or 12 hours, starting at a time specified in hours and minutes.

   NOTE

   The specified times are local times based on the Corero Network Device's system clock.

3. Once you have specified your settings, click OK.
4. In order to specify the report template for the report you want to generate, choose Configure Security > Security Logs and Reports > Manage Report Templates from the Navigation Tree. The Manage Security Report Templates dialog box displays. The available security report templates are listed in Table 16-1.
5. Click the Settings button. The Edit Report Template dialog box displays.
6. If desired, modify the following report options.
   - Select the desired (Security Events) List option. You can select whether the report will list the Top 10 security events, the Top 20 security events, All security events, or All Non-Zero security events (all security events that have a blocked count greater than zero).
   - To enable periodic generation, select the Generate Periodic Report check box. To disable generation, clear (uncheck) the check box.

   When you are finished, click OK.
7. Click the Save Configuration Toolbar button to save your changes.

Viewing Saved Security Reports

By default, all generated reports, whether Immediate or Periodic, are saved.

To view either an immediate or a periodic security report:

1. From the Navigation Tree, choose Monitor Security > Reports. The View Security Reports dialog box displays.

   On a Corero Network Device, previously generated reports are listed in chronological order. Each row in the table displays the report file name and template, whether the report was an Immediate report, and the date and time it was generated.

2. Select the desired report.
3. Click View. The selected report displays in a web browser.

Deleting Saved Security Reports

At regular intervals, you should delete saved security reports from your system. Note that all generated reports are automatically saved by the system.

When identifying which security reports to delete, consider the following:

- A Corero Network Device holds up to two immediate reports for each report type (report template), and up to seven periodic security reports for each report type (report template). When the number of a particular type of report has reached its limit, old reports of that type are deleted as new reports of that type are generated.

To delete a saved security report:

1. From the Navigation Tree, choose Monitor Security > Reports. The View Security Reports dialog box displays.
2. Select the desired report. You can select multiple reports using Ctrl-Click and Shift-Click.
3. Click Delete. You are prompted to confirm your selection.

Managing Security Report Templates

The content and form of the periodic security reports are determined by predefined templates.

To manage security report templates:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Manage Report Templates. The Manage Security Report Templates dialog box displays.
2. To view the PHP code that is used to generate a specific security report, select the desired template and click View Template. The PHP report template code displays in a browser window.
3. To create and upload a modified security report template:
   a. View the desired template and save it to a new file name.
   b. Modify the template file as needed and save it to a new file name.
   c. On the Manage Security Report Templates dialog box, click Upload. The Upload Report Template dialog box displays.
   d. Browse to the new report template file, then click Upload.
4. Save your changes by clicking the Save Configuration toolbar button.

Chapter 17
Managing Security Logs

Corero Network Devices enable you to manage the type and severity of events that are logged. This chapter describes how to manage and view security logs.

This chapter contains the following topics:

Understanding Event Logging

During the process of receiving and transmitting traffic, Corero Network Devices perform many checks and other operations. All of these operations, and all of the system events and user-related management interface tasks produce event messages.

You can use these messages to understand the state of the components in the device and the quality of the traffic flowing through your network. Based on your analysis of the message you can take actions such as modifying the configuration of the device to deal in specific ways with certain servers that need to be controlled, or with traffic sources that appear to be malicious.

The rest of this section describes the following:

What is an Event?

Corero Network Devices generate a large variety of messages based operational and security events that occur, from purely system-related events such as ports going up or down to very specific traffic checks that the device performs.

Some of the reasons that the device generates messages include:

- Different categories of traffic filtering
- Various operational and filtering thresholds that are crossed
- Configuration changes
- Engine, processing, and hardware events
- Management changes
- Device reboot or restart
- Port setting changes
- Component failure
- Successfully setting up or tearing down a connection
- An instance where the device acts as a proxy for one of your servers
- An device configuration change
- Failure of a packet to successfully pass a specific integrity test

NOTE

The Event Logging System online help, available on the documentation CD-ROM, provides detailed information about the messages and the subsystems that generate them.

In addition to generating messages for each event, the device also tracks the total number of messages in each major category by incrementing counters that you can examine.

For example, it tracks the total number of connections it makes per second for the traffic flowing through the part of the network it is watching.

Event Logging System Outputs

The Event Logging System receives message input from the various subsystems and sends the messages to a wide variety of user destinations based on the controls that you supply through the Graphical User Interface.

Based on your input, the event logging system determines whether a message it receives from a subsystem should be sent to one or more of the following destinations:

- Console — CONSOLE port on the front of the device.
- Memory — Store a selected set of messages on the device. Based on your configuration, divide the messages into two broad groups:
  - Event Messages — Non-critical, non-security-related events such as port status and connection creation.
  - Alert Messages — Critical and security-related events such as equipment failure, and network attacks.
- Syslog — One or more user-defined Syslog servers.

You have complete control over which destinations receive which messages and whether the message should actually be sent.

Message Control Hierarchy

The event logging system provides the user with a hierarchical form of message control. This control can be as coarse as turning the entire logging system off or on, and as fine as setting parameters for a specific event message.

The event logging system supports the following levels of message control:

- Global Level — At the highest level, turn all message generation on or off and set a minimum severity threshold that a message must meet to be transmitted by the event logging system.
- Subsystem Level — Control message output from each subsystem individually. You can turn messages for a subsystem on or off, and set the priority that all messages for that subsystem should have.
- Message Groups — Place messages into predefined and user defined groups, and then control the processing of all message within a group. You can turn all messages in the group on or off, add and delete messages in a group, set the group's priority, or set the destination for all messages in the group.
- Message Level — At the lowest level, control all the settings for individual messages. You can enable and disable a given message, set its priority, indicate the destinations for that message.

NOTE

Even if you turn off event logging, the Corero Network Device continues to properly maintain all of its event counters.

Viewing the Events Log

Your Corero Network Device stores important system-related events such as boot events, and management application access login and logout in its Events Log.

To view the Events Log:

1. From the Navigation Tree, do one of the following:
   - Click Monitor System > Events Log.
   - Click Monitor System > System Log Viewer, and choose Events Log from the drop-down list.

   When you make this selection, you may receive a download confirmation message.
2. The Events Log provides information such as startup details, and when management sessions start and stop.

   Whenever a management session is opened or closed, the Events Log provides the following information:

   - user — The name under which the user logged in, or attempted to log in.
   - session type — The type of management access, for example HTTP or SNMP.
   - cip — The client IP address from which the management session was initiated.
   - cprt — The client port used to initiate this management session.
   - msg — Message regarding status of this session (Logon, Logged out, Unknown user, etc.)

Viewing the Alerts Log

The Alerts log provides a record of the attacks detected and blocked. The Alerts log records the first few events from any given attack; then, aggregates the remaining attacks into summary messages.

NOTE

The Alerts log is primarily used when troubleshooting issues with the help of Corero Services and Support.

You can access the Alerts log directly from the Graphical User Interface as follows:

1. From the Navigation Tree, do one of the following:
   - Click Monitor Security > Alerts Log.
   - Click Monitor System > System Log Viewer, and choose Alerts Log from the drop-down list.
2. A browser window opens, displaying the contents of the Alerts log.

Table 17-1 lists the information that can display in the Alerts log.

Table 17-1: Alerts Log Information

Item	Description
Date and Time	Date and time the message was generated.
IP Address	The IP address of the device that received the attack.
Model Number	The model number of the device that received the attack.
id	The Event Logging System ID for this message.
prot	The network protocol used in the attack traffic.
cip	The source (client) IP address for the attack traffic.
cprt	The network port number (client port) associated with the client originating the attack.
sip	The destination (server) IP address for the attack traffic.
sprt	The destination (server) port for the attack traffic.
atck	The TLN rule that was triggered by this traffic.
disp	Disposition indicates the final determination made by a forwarding process with respect to a network traffic security event. Types of disposition include unknown, blocked, detected, and so forth.
ckt	The physical device port that received this traffic.
src	Indicates whether the source of the traffic was internal or external to your network.
msg	A brief message describing how the device handled the traffic (for example Blocked By Firewall).
num	The sequential number uniquely identifying this log entry.

-

Viewing Audit Logs

Your Corero management application has an audit function which logs every change made through the user interface. These items are kept in a log file, and, if Syslog server(s) have been setup, can also be configured to send them to the Syslog server(s).

All configuration changes, save operations, boot ups, and failed and successful authentications are logged. Note that audit logging is disabled by default.

NOTE

For information on configuring audit logging for your IPS Unit, see Managing Audit Logs (page 4-3).

Audit messages are stored in the audit log file in the following format:

Table 17-2: Audit Log Information

Day	Time	Unit_IP	Event ID	Device	User	Details
Aug 31	10:52:41	10.25.36.102	rr-nn	TLN-TQ	peterz	Text Description

Where:

- ID is a unique number identifying the audit entry.
  rr is the number of times that the device has been rebooted.
  nn is a sequentially increasing number since the last reboot.
- User is the name of the user logged in performing the action, or “Not Known” if it cannot be identified (such as when the unit is powered on - this event is audited but there is no user logged in)
- Details contains information about the operation being audited.

NOTE

Audit fields contain different values depending on the operation being audited.

Audit messages are kept in an audit log file stored in compact flash memory. Up to 10 audit files are maintained, the oldest being deleted to make available space for a new one when required.

To view audit logs:

1. To view audit log information, choose Monitor System > System Log Viewer from the menu bar.

   The View Log File dialog box displays.

2. In the Log Type drop-down, choose Audit Log. The View Log File dialog box displays.
3. Select the desired Audit Log file from the drop-down list, then click OK. Select a past audit log file or choose the current file if you want to view the most recent information.

   If you choose to view contents of the Current File, you can click Refresh to display any more current information, if available.

4. The View Audit Log File dialog box displays the date and time, event ID, and event details, with the most recent data displayed first.

You can scroll through the audit log data, or search for specific terms. You can also sort the data based on a column's contents by clicking the column's heading.

Managing Event Groups

You can assign an event message to one event group, either one of the default groups, or a group that you define. You can then enable or disable all the messages in a group.

- When a group is enabled, the device processes (logs and sends) messages in that group.
- When a group is disabled the device does not process (log or send) messages in that group.

> NOTE
>
> Disabling the messages for a group does not disable counters associated with those messages.

To manage event groups:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Event Logging > Event Groups. The Event Groups dialog box displays.
2. To create an event group:
   a. Click Add. The Add Event Group dialog box displays.
   b. Enter the name of the new group. Provide a meaningful name for the type of message you intend to associate with this group.
   c. Specify whether messages in the group are Enabled or Disabled.
   d. When finished, click Add.
3. To Enable or Disable the messages in a group:
   a. Select the desired group, then click Edit.
   b. Set the Mode to Enable or Disable.
   c. When finished, click OK.

> NOTE
>
> You cannot modify the name of a group. To change the name of a group, delete the existing group, then create a new group using the correct name and the desired settings.

4. To Delete a group:
   a. Select the group, then click Delete.
   b. You are asked to confirm your selection.

Setting Global Event Logging Controls

To set global logging controls:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Event Log Settings. The Event Log Settings dialog box displays.
2. To enable the Corero Network Device to process all messages (subject to other controls, such as individual message enabled, group enabled, and so forth), set the Mode to Enabled.
3. To disable message processing, set the Mode to Disabled. The device does not process any event messages, regardless of other message settings.
4. Specify the maximum log level for reporting events. The log level provides an overall threshold that messages must meet or exceed before the device processes them.

   Log level settings (Table 17-3) follow Syslog severity conventions with zero indicating the most severe events, and 7 indicating the least severe.

Table 17-3: Event Log Levels

Level	Description
0	Emergency (the highest setting)
1	Alert
2	Critical
3	Error
4	Warning
5	Notice
6	Information
7	Debug

5. Save your changes by clicking the Save Configuration toolbar button.

Setting Message Controls by Event Subsystem

Messages are assigned to subsystems by the event logging system. You cannot add or delete subsystems, nor can you change the subsystem to which a message is assigned. However, the event logging system allows control of messaging at the subsystem level.

You can enable or disable all the messages for a given subsystem.

- When a subsystem is enabled, the device processes (logs and sends) messages associated with that subsystem.
- When a subsystem is disabled the device does not process (log or send) messages associated with that subsystem.

You can also specify the summary frequency. This value controls the interval (in seconds) between times when the device generates summary messages for this subsystem. A summary message provides the total number of messages (including the summary messages) generated for that subsystem since the device started. By comparing the total in the summary message with the actual messages of that type stored in your Syslog file, you can determine whether you have received all messages, or whether you may have a problem such as an overloaded Syslog server.

A summary message provides the total number of messages (including the summary messages) generated for that subsystem since the device started. By comparing the total in the summary message with the actual messages of that type that you have in your Syslog file, you can determine whether you have received all of the messages or whether you may have a problem such as an overloaded Syslog server.

Different message types are associated.with each subsystem. The subsystems and their message types are listed in Table 17-4.

Table 17-4: Event Subsystems

Subsystem	Message Types
AM Filters	• Attack mitigation filtering of TCP, UDP, ICMP and IP traffic • Aggregation of event messages • DDoS rejection • Threshold crossed (fragment load shedding) • MIB object changed
Bridge Forwarding	• High and low thresholds crossed • MIB object changed
Classification	• Thresholds crossed • MIB object changed
Configuration	• X.509 certificate update • NTP config update • Threshold crossed • MIB object changed
Flow	• TCP, UDP, IP and ICMP flow events • Threshold crossed (connection rates and flow resource usage) • MIB object changed

Table 17-4: Event Subsystems *(Continued)*

Subsystem	Message Types
Interface	• Interface events • Threshold crossed • MIB object changed
IP Forwarding	• IP traffic forwarding • Threshold crossed • MIB object changed
Policy	• Threshold crossed • MIB object changed
Re-Order Engine	• Reorder engine IP fragment configuration error • IP fragment error • IP fragment aged error • IP fragment IP options error • Threshold crossed • MIB object changed
Remote Access	• Session started, ended, or terminated • Authentication failure • Threshold crossed • MIB object changed
Statistics	• Threshold crossed • MIB object changed
System	• System, engine and hardware status • Software upgrade • Thresholds crossed (CPU heavy load, CPU idle time) • MIB object changed

To modify event subsystem settings:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Event Logging > Event Subsystems. The Event Subsystems dialog box displays.
2. Select the desired subsystem, then click Edit. The Edit Event Subsystem dialog box displays.
3. Set the mode to Enabled or Disabled (as described above).
4. Specify the Summary Frequency (as described above).
5. When finished, click OK.

Configuring Event Thresholds

For certain operational events, you can set thresholds to determine when the device should generate an event message. When the device crosses the threshold, it sends a message to the Event Logging System.

Thresholds can mark when a particular system activity crosses out of the normal range of operation, and can also specify when they cross back into the normal range.

You can configure the thresholds that the device uses to determine when to generate many event messages. Event-specific threshold values are listed in Table 17-5.

You can also specify whether the threshold is enabled or disabled:

- When a threshold is enabled, and that threshold is crossed during system operation, a message will be generated (unless the group or subsystem associated with that message is disabled).
- When a threshold is disabled, crossing the threshold will not trigger an event message.

Figure 17-1 shows the Event Thresholds dialog box.

Figure 17-1: Event Thresholds Dialog Box

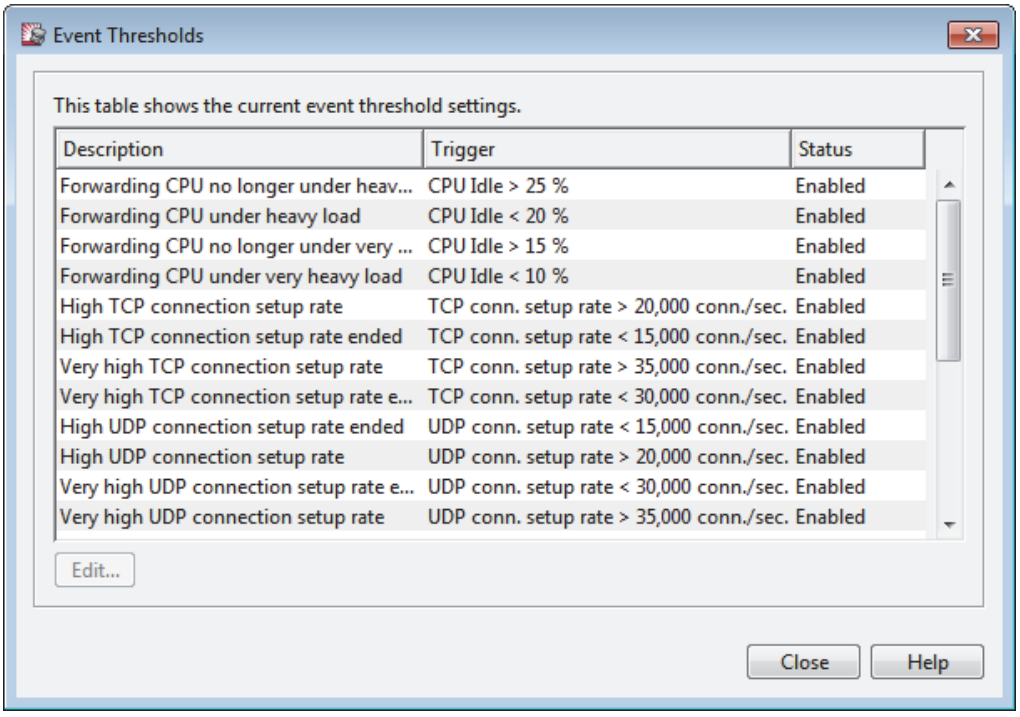

To modify an event threshold:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Event Logging > Event Thresholds. The Event Thresholds dialog box displays (Figure 17-1).
2. Select the desired event threshold, then click Edit. the Edit Event Threshold dialog box displays.
3. Select whether the threshold is Enabled or Disabled (as described above).
4. Modify the trigger as needed.
5. When finished, click OK.

6. Save your changes by clicking the Save Configuration toolbar button.

Table 17-5 describes the user-configurable threshold triggers for each event:

Table 17-5: User-Configurable Event Threshold Triggers

Event	Threshold Value
Forwarding CPU no longer under heavy load	CPU idle time percentage goes above your configured value.
Forwarding CPU under heavy load	CPU idle time percentage goes under your configured value.
Forwarding CPU no longer under very heavy load	CPU idle time percentage goes above your configured value.
Forwarding CPU under very heavy load	CPU idle time percentage goes under your configured value.
High TCP connection setup rate	TCP connection setup rate goes above your configured number of connections per second.
High TCP connection setup rate ended	TCP connection setup rate goes below your configured number of connections per second.
Very high TCP connection setup rate	TCP connection setup rate goes above your configured number of connections per second.
Very high TCP connection setup rate ended	TCP connection setup rate goes below your configured number of connections per second.
High UDP connection setup rate ended	UDP connection setup rate goes below your configured number of connections per second.
High UDP connection setup rate	UDP connection setup rate goes above your configured number of connections per second.
Very high UDP connection setup rate ended	UDP connection setup rate goes below your configured number of connections per second.
Very high UDP connection setup rate	UDP connection setup rate goes above your configured number of connections per second.
High flow resource usage ended	Available positions in the flow table (free flows available) goes above your configured number of flows.
High flow resource usage	Available positions in the flow table (free flows available) goes below your configured number of flows.
Very high flow resource usage ended	Available positions in the flow table (free flows available) goes above your configured number of flows.
Very high flow resource usage	Available positions in the flow table (free flows available) goes below your configured number of flows.
High IP address resource usage	Available positions in the IP address table (free IP addresses available) goes below your configured number of addresses.
High IP address resource usage ended	Available positions in the IP address table (free IP addresses available) goes above your configured number of addresses.
Very high IP address resource usage	Available positions in the IP address table (free IP addresses available) goes below your configured number of addresses.
Very high IP address resource usage ended	Available positions in the IP address table (free IP addresses available) goes above your configured number of addresses.

Table 17-5: User-Configurable Event Threshold Triggers *(Continued)*

Event	Threshold Value
Receive packet load shedding triggered	Load shedding starts or stops (not configurable).
Transmit packet load shedding triggered	Load shedding starts or stops (not configurable).
Receive IP fragment load shedding triggered	Load shedding of fragments starts or stops (not configurable).

Modifying Individual Message Settings

You can modify settings associated with individual event messages, including their severity, where the message will be sent, and the group with which the message is associated.

Event message settings include whether the message is enabled or disabled, the message severity, the message priority, and whether the message will be sent to specified servers or logs.

In addition, you can view information on how many times a particular message was triggered. You can view three categories of message counts:

- Sent
  The number of messages actually sent to the enabled destinations (Console, Syslog, and nonvolatile memory log file).
- Failed
  The number of messages that, for some internal reason, were triggered but not sent.
- Filtered
  The number of messages that, because of your settings, were not sent. Settings that affect filtering include message, group, and subsystem mode settings, message severity levels, as well as master logging settings.

The sum of the Send, Failed, and Filtered message counts equals the total number of instances generated for a given message.

Figure 17-2 shows the Event Messages dialog box.

Figure 17-2: Event Messages Dialog Box

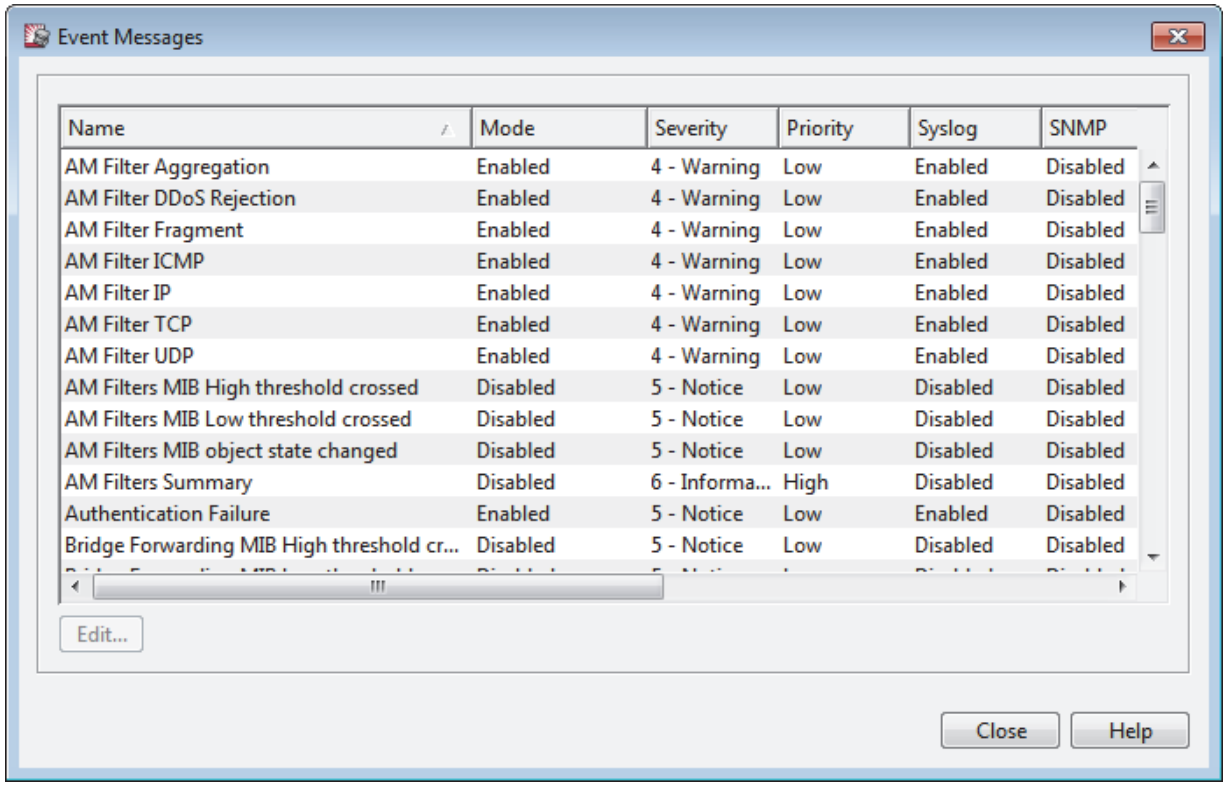

To modify the settings for an event message:

1. From the Navigation Tree, choose Configure Security > Security Logs and Reports > Event Logging > Event Messages. The Event Messages dialog box displays (Figure 17-2).
2. Select the desired message, then click Edit. The Edit Event Message dialog box displays.
3. Modify the settings as desired. The settings are described in Table 17-6.
4. When finished, click OK.
5. Save your changes by clicking the Save Configuration toolbar button.

Table 17-6: Event Message Settings

Setting	Description
Name	The name of this event message. Event message names cannot be modified.
Mode	Specify whether this event message is enabled or disabled: • When an event message is enabled, if a triggering event occurs, a message will be generated (unless the group or subsystem associated with that message is disabled). • When an event message is disabled, a triggering event will not result in an event message.
Severity	The severity you want to associate with this message. Severity levels are listed in Table 17-3. This setting must be less than or equal to the global logging severity setting or the message will not be processed.
Priority	The priority associated with the event. You cannot modify the priority for an event message.
Syslog	• When enabled, the event logging system will send this message to any configured Syslog servers. • When disabled, the event logging system will not this message to any configured Syslog servers.
SNMP	• When enabled, the event logging system will send this message to any configured SNMP trap servers. • When disabled, the event logging system will not this message to any configured SNMP trap servers.
CF Log	• When set to None, the event logging system will not send this message either to the Event log file or to the Alert log file on the device's nonvolatile (compact flash) memory module. • When set to Alert, the event logging system will send this message to the Alert log file on the device's nonvolatile memory module. • When set to Event, the event logging system will send this message to the Event log file on the device's nonvolatile memory module.
Group	To assign this message to a group, choose a group from the drop-down list. You can then control this message by enabling or disabling its group.

Chapter 18
System Monitoring

Corero Network Devices provide a number of ways you can view information on their status and operation. You can view information on device and component status, statistics, application connections, and IP addresses.

This chapter contains the following sections:

Using the Front Panel View

The Front Panel display is a dynamically changing view of port status. You can click various areas to display port roles, port states, port settings, and other port and system information. You can also access port configuration windows for individual ports. After you run the Getting Started wizard, the device automatically updates the Front Panel display to reflect your configuration choices.

A legend displays above the front panel, showing the Port State and Port Role icons available for each port. At the bottom is information on current port settings.

To display the front panel using the management application for a Corero Network Device:

1. Do one of the following:
   - Click the Front Panel toolbar button.
   - From the navigation tree, choose Monitor System > Front Panel.

Figure 18-1 shows the Front Panel View for an IPS Unit.

Figure 18-1: IPS Front Panel View

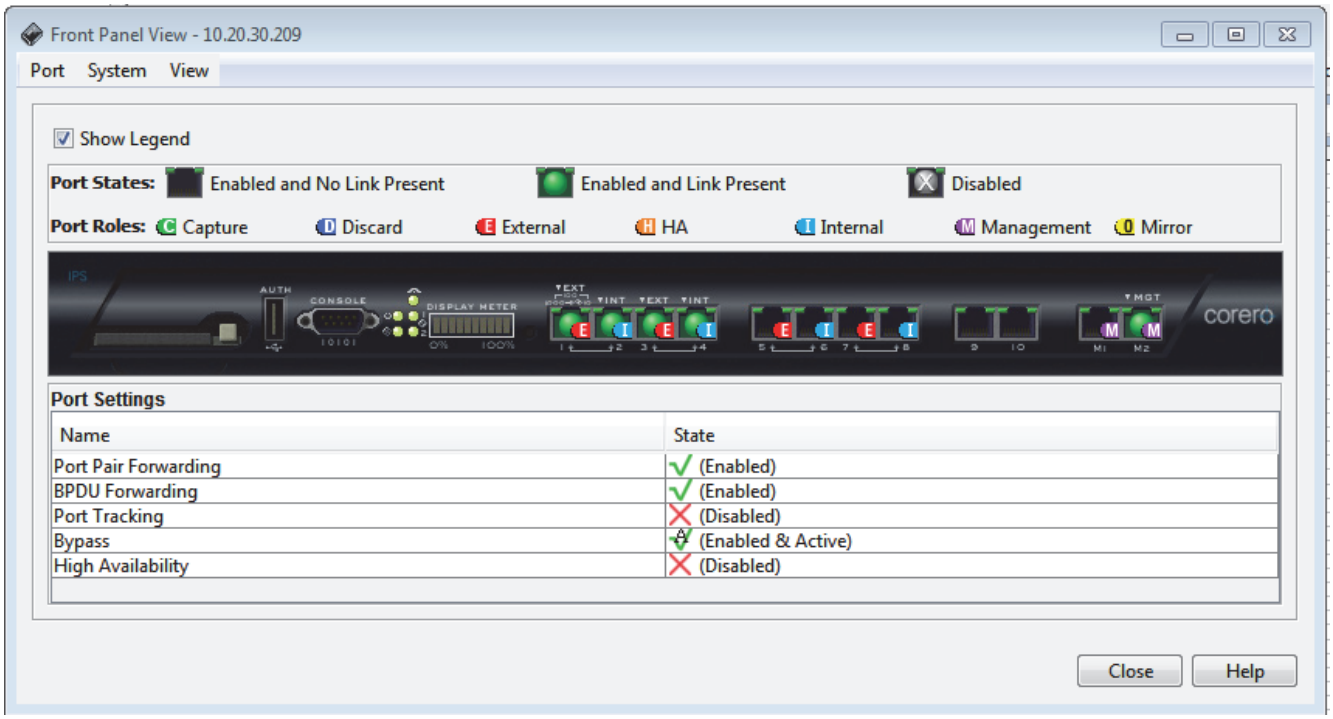

The Front Panel displays the information listed in Table 18-1.

Table 18-1: Port Icons on the Front Panel View

Port Icon	Description
Port State	Indicates two things: • Whether a port is enabled or disabled • Whether the device senses a link (connection) for that port. In the previous figure, Port M1 shows icons for Port Enabled and Link Present.
Port Role	Indicates by letters and colors each port's assigned role. A check box on the Front Panel View enables you to show or hide a legend describing port roles. • Management ports are indicated by a purple M. • External mission ports are indicated by a red E. • Internal mission ports are indicated by a light blue I. • A capture port is indicated by a green C. • A discard port is indicated by a dark blue D. • A mirror port is indicated by a yellow O.
Display Meter	The Front Panel view contains a horizontal Display Meter that indicates the current traffic load, in connections per second, that the device is handling. This meter is located immediately to the left of the available ports. The meter contains ten LED segments. Each segment represents 5000 connections per second.
Port Settings	The Front Panel view also provides visual indication of the status of the following major features: • Port Pair Forwarding — If enabled, the device forwards traffic between two matched input and output Mission ports. • Port Tracking — If enabled, the device tracks the state of Mission port pairs. If one port of the pair changes state, the device changes the state of the other port to match it. • Bypass Settings Indicator — Indicates which of the three bypass modes you have currently selected. Bypass can be Enabled or Disabled, or you can choose to have the system Bypass During System Reset. For more information on bypass settings, see Selecting the Bypass Settings Mode (page 6-8). In addition, for Always Bypass mode (default), the display provides an indication of whether the port is Active or Inactive. An Active indication means that the device is currently sending traffic through without performing any mitigation, actively bypassing all security functions. **IMPORTANT:** Be sure to change this setting after you finish configuring the device. • High Availability — If enabled, indicates that this device is part of a ProtectionCluster.

To view and manage features from the Front Panel View:

1. To display the Front Panel view, do one of the following:
   - Click the Front Panel toolbar button.
   - From the navigation tree, choose Monitor System > Front Panel.

The Front Panel View displays.

Figure 18-1 shows the IPS Front Panel View.

2. To view or modify settings for a particular port, do one of the following:
   - Select the port in the Front Panel View, then choose Port > Settings from the menu bar.
   - Right-click the port in the Front Panel View, then choose Settings from the pop-up menu.

   The Edit Port Settings dialog box displays.

3. To view or clear statistics for a particular port, do one of the following:
   - Select the port in the Front Panel View, then choose Port > Statistics from the menu bar.
   - Right-click the port in the Front Panel View, then choose Statistics from the pop-up menu.

   The Port Statistics dialog box displays . For additional details on viewing port statistics, see Viewing Port Statistics (page 18-8).

4. To view information specific to this Corero Network Device, do one of the following:
   - Choose System > Information from the menu bar.
   - Right click anywhere except over a port and choose Information from the pop-up menu.

   For additional details on viewing system information, see Viewing System Information (page 18-6)

5. To show or hide role-specific information on the Front Panel View, select View > Role. Hiding role information on the display can help you focus on port states.

Viewing “About...” Information

Using the management application, you can view product-specific “About...” information at any time.

To view “About...” information:

1. In the management application for a Corero Network Device, from the Navigation Tree, choose Help > About.

   The “About...” dialog box displays the information fields listed in Table 18-2.

Table 18-2: “About...” Information

Field	Description
License Information	Abbreviated end user license agreement information for the current software.
View License Agreement button	Provides access to the full text of the End User License Agreement that was accepted for the current software version.
Hardware Type	Identifies the model of this device.
Software Version	The currently-running version of the software.
Serial Number	The serial number of this device.
Java Runtime Environment Version	The currently-installed Java Runtime Environment client version.

2. If desired, click the View License Agreement button. The full text of the End User License Agreement that was accepted for the current software version displays.

NOTE

The license text dialog box includes a link to the Corero web site so you can view the most recent license agreement information.

Viewing System Information

Using the management application for a Corero Network Device, you can view detailed system information at any time.

To view device-specific information:

1. In the management application for a Corero Network Device, do one of the following:
   - Click the System Information Toolbar button.
   - From the Navigation Tree, choose Monitor System > System Info.
   - From the Front Panel View, choose System > Information from the menu bar.
   - From the Front Panel View, right click anywhere except over a port and choose Information from the pop-up menu.

   NOTE

   If you only need to view the hardware type, software version, serial number, and Java Runtime Environment Version, choose Help > About from the navigation tree.

The System Information dialog box displays the information fields listed in Table 18-3.

Table 18-3: System Information

Field	Description
Name	The name (a text string with a maximum of 20 characters) used to identify this device.
Location	The location (a text string with a maximum of 20 characters) used to indicate the physical location of this device.
Contact	The contact person (maximum of 20 characters) who manages this device.
Hardware Type	Identifies the model of this device.
Software Version	The currently-running version of the software.
Protection Pack Level	The most recently installed TopResponse™ protection pack.
Serial Number	The serial number of this device.
MAC Address	The MAC address of this device.
License Key Status	The status of the current license key. Status messages include: • Permanently Unlocked - The Corero Network Device is operating with a permanent system license, with no expiration date or restrictions. • Valid Until (expiration date) - The Corero Network Device is operating with a trial system license, and the expiration date is listed. Note that the expiration date is specified in GMT (UMT) time. • Locked In Bypass Mode - The Corero Network Device was operating with a trial system license, but the license has expired. The device will remain locked in bypass mode until a new license is applied. Contact Corero for assistance. • Unknown - The Corero Network Device is currently unable to obtain information on the license status. If this problem persists, contact Corero.
Enter System License Key button	This button enables you to enter a System License Key.

NOTE

For information describing the operation and entry of a System License Key on an IPS Unit, see System License Management Key (page A-2).

Viewing Port Statistics

The Port Statistics window displays a list of the ports, with statistical information about port activity. You can view statistics for a single port, or for all ports. The statistics displayed are standard Ethernet Statistics Group counters (defined by RFC 1757, RMON MIB) or Interfaces table counters (defined by RFC 1213, MIB II).

1. To view port statistics on the management application for a Corero Network Device, do one of the following:
   - To view statistics for all ports, from the Navigation Tree, select Monitor System > Port Statistics.
   - To view statistics for a specific port, on the Front Panel View, right click any port and, from the pop-up menu, select Statistics.

   The Port Statistics dialog box displays. For details on the information displayed in the Port Statistics dialog box, see Table 18-4.

2. To clear (zero) all statistics for a specific port, select the port and click Clear Statistics.

Table 18-4: LAN Port Information

Column	Description
Name	The physical port number.
Receive Link Util	Receive link utilization for this port, expressed as a percentage of available bandwidth.
Transmit Link Util	Transmit link utilization for this port, expressed as a percentage of available bandwidth.
Total Packets	The total number of packets received, including bad packets, broadcast packets, multicast packets. and 1518 octets (excluding framing bits but including FCS octets) but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
Total Octets	Total number of octets of data received on the network, including those in bad packets but excluding frame check sequence (FCS) octets. This object provides a reasonable estimate of Ethernet utilization. Refer to RFC 1757 for information if more precision is desired.
Broadcast Packets	Total number of good packets received that were directed to the broadcast address (this does not include multicast packets).
Multicast Packets	Total number of octets of data received on the network, including those in bad packets but excluding frame check sequence (FCS) octets. This object provides a reasonable estimate of Ethernet utilization. Refer to RFC 1757 for information if more precision is desired.
Bad CRC	Total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error). A high number of bad CRCs can indicate a port speed mismatch.
Collisions	Best estimate of the total number of collisions on this segment. Refer to RFC 1757 for more information about this counter. A high number of collisions can indicate a port speed mismatch.
Receive Unicast Packets	Number of unicast packets delivered to a higher-layer protocol. Refer to RFC 1757 for more information about this counter.
Receive Non-Unicast Packets	Number of non-unicast packets (broadcast or multicast packets) delivered to a higher-layer protocol. Refer to RFC 1757 for more information about this counter.
Receive Octets	Total number of packets received that were between 64 and 1518 octets in length (including bad packets), excluding framing bits but including FCS octets.

Table 18-4: LAN Port Information *(Continued)*

Column	Description
Transmit Unicast Packets	Total number of packets that higher-level protocols requested be transmitted to a unicast address, including those that were discarded or not sent.
Transmit Non-Unicast Packets	Total number of packets that higher-level protocols requested be transmitted to a non-unicast (broadcast or multicast packets) address, including those that were discarded or not sent.
Transmit Octets	Total number of octets transmitted out of the interface including framing characters.
Transmit Collisions	Total number of packets that experienced a collision during transmission.
Fragment	Total number of packets that were fragmented during transmission.
Undersized	Total number of packets received that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed.
Oversized	Total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
Jabbers	Total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).

Viewing Current Application Connections

The Current Application Connections dialog box displays information about successfully established connections for each defined network service (application). An application specifies a name for a specific network protocol/port combination.

NOTE

You can also view application connection usage and connection setup rates from the dashboard.

To view current application connections:

1. From the Navigation Tree, choose Monitor System > Current Application Connections.

   The Current Application Connections dialog box displays. For each application, the table displays the number of current connections established by the device for this application.

2. By default, the device only displays applications that have one current connection. If you would like to view all applications, even those with no current connections, clear (deselect) the Hide Zero Counters check box.

Viewing the Bridge MAC Address Table

To display the bridging details about each MAC address that the Corero Network Device has handled:

1. From the Navigation Tree select Monitor System > Bridge MAC Address Table.
2. The Bridge MAC Address Table dialog box displays. For each MAC address the device has seen, it displays the following information:
   - The MAC address
   - The VLAN most recently associated with this MAC address.
   - The port on which traffic arrived
   - Whether this MAC address is part of a user-specified set of predefined (static) MAC addresses, or whether the MAC address was learned by the device.
3. To delete the current port information for all known MAC addresses, click Flush Entries.

Viewing the Management Port ARP Table

The Address Resolution Protocol (ARP) table provides a list of discovered IP address/physical address pairs. To display the ARP information for all known IP addresses:

1. From the Navigation Tree, choose Monitor System > Management Port ARP Table.
2. The Management Port ARP Table dialog box displays. For each IP address the device has seen as a source or destination address for internal or external traffic, the table displays the corresponding MAC (physical hardware) address.

Chapter 19
Security Management and Monitoring

Security management and monitoring are the most frequently performed tasks in the management application. The management application provides several primary views that enable you to quickly detect issues, drill down to identify additional information, and apply both short term and long term solutions to resolve them.

This chapter also describes the shunning feature, which is used tor quickly block traffic initiated by a suspect IP addresses. You can easily enable (shun) and disable (unshun) this treatment for a specific IP address.

This chapter contains the following sections:

Security Monitoring Overview

The task of monitoring security includes identifying security issues, then researching additional information to help identify the cause.

Table 19-1 lists the primary views used to identify security issues.

Table 19-1: Security Monitoring Tools

Tool	Description	For more information, see...
Blocked and Detected Attacks	This view displays a table of all attacks detected or blocked. This display is updated in real-time. From the Blocked and Detected Attacks view, you can easily navigate from a particular attack to the Security Event Viewer for additional analysis and management.	• Viewing Blocked and Detected Attacks (page 19-16)
Security Event Viewer	A powerful event reporting tool that enables you to view, sort, and filter security events. From the Security Event Viewer, you can easily navigate from a particular event to any associated information including policies, rules, and detailed IP address information.	• About the Security Event Viewer (page 19-19) • Viewing Security Events and Security Event Details (page 19-23) • Security Event Viewer Filter Tool (page 19-25)

NOTE

If you discover a serious security issue, one of your first actions can be to shun, or completely block, the offending IP address. For more information see About Using IP Address Shunning to Stop an Attack (page 19-4).

Table 19-2 summarizes the tools you can use to research issues you discover using the Blocked and Detected Attacks page or the Security Event Viewer.

Table 19-2: Security Issue Research Tools

Tool	Description	For more information, see...
Reports	Choose from among the available security report templates to display detailed security information. You can generate scheduled reports (periodic reports), or generate reports on demand (immediate reports).	• Chapter 16, “Generating and Viewing Security Reports”
Alert Logs	Displays a record of the events associated with attacks detected and blocked. Entries for multiple attacks of the same kind are aggregated. Log information provides details not found in the Blocked and Detected Attacks window.	• Chapter 17, “Managing Security Logs”
IP Address Query	Provides details about the behavior of a host as it is requesting and completing connections. Also provides the ability to reset the SYN flood and/or connection counters for the individual IP address identified.	• Reset (Clear) SYN Flood and Connection Counters (page 19-31)

Table 19-2: Security Issue Research Tools *(Continued)*

Tool	Description	For more information, see...
Statistics	You can view statistics on port activity or dropped packets. The dropped packet statistics view summarizes information on both received and dropped packets. Provides an instant view of the volume of issues since startup.	• Viewing Dropped Packet Statistics (page 19-33)
Graphs	Set of visual representations showing overall traffic activity, security-related activity, and other system-level information. You can view graphs displaying information about dropped packets, SYN flood statistics, IP threat levels, connection usage, connection rates or CPU activity. You can also design your own custom graph.	• Viewing Charts and Graphs (page 19-37)

About Using IP Address Shunning to Stop an Attack

Corero Network Devices have an effective protection capability called shunning that can quickly block traffic from IP addresses, temporarily or permanently, that are suspected of originating or participating in an attack. Shunning an attacker's IP address at an ingress point to the network reduces the possibility of the attack expanding to other targets within the environment protected by the device.

When shunning, you group IP addresses into collections called Shun Labels. You can assign IP addresses to a shun label based on any criteria you desire, including common features of the IP source system, by date, or by attack type.

NOTE

Shunning is only available in E-series IPS Unit models.

The advanced protection capabilities provided by shunning are described in Table 19-3.

Table 19-3: Shunning Capabilities

Capability	Description
Attack Source Identification	The Security Event Viewer enables users to identify a set of attacker IP addresses associated with blocked and detected attacks.
Malicious IP Address Shunning	Isolate events of interest and automatically shun all IP addresses associated with a particular attack event. Users can set time periods for how long each address should be shunned. They can also manually unshun addresses that are later determined to be safe.
Attack Defense Dashboards	The user interface allows Security Operations Center personnel to switch between daily monitoring and under-siege incident response.
Additional Router Protection	Administrators can export a list of IP addresses being shunned so that they can be imported into a router for blocking by the router.

You can identify addresses to shun in three ways:

- Using the IP addresses from the events selected in the Security Event Viewer (assuming that events have been selected). See Viewing Security Events and Security Event Details (page 19-23) .
- Using the IP addresses from all the events matching the current filter in the Security Event Viewer. See Security Event Viewer Filter Tool (page 19-25).
- Manually entering IP addresses in the Shun Attackers dialog box.

Once you have identified one or more IP addresses associated with events of interest, you can shun one, some, or all of them. IP addresses you want shunned can also be entered manually. Shunning is a temporary activity, and various time periods can be specified for the addresses to be shunned.

Shunning Considerations

When using the shunning feature, consider the following:

- When a shunning action is performed, all IP addresses associated with that action are assigned a shun label, which is supplied by default as a timestamp but which can be changed by the user to any descriptive text desired.
- Multiple shunning actions can be performed, each one of which is allocated its own shun label.
- A maximum of 256 shun labels are supported.

- A shun label is an attribute of an IP address, and an IP address is only ever associated with one shun label. If the user associates an IP address with a shun label, and that IP address is already associated with an existing shun label, it will lose its association with the previous shun label.
- Each shun label is associated with a shun rule. When an IP address sends a packet and that packet is blocked as a result of the shunning function, a security event is generated. The rule that is triggered is one of the shunning rules. There are 256 shunning rules, they are numbered tln-033001 through tln-033256, with rule names of FWALL: IP Address Shunned with Label 001: unassigned. The label number refers to the shun label number, and the unassigned label is replaced with the name of the shun label.
- The time remaining for IP addresses associated with a shun label can be:
  - Increased by the user - this is useful if an attack continues for a longer than expected period of time.
  - Set to zero - as an alternative to unshunning an IP address.
- Because shunning is a temporary action, all shun information is transient and is lost after the device is rebooted. Any attackers that the user determines should be permanently blocked should be added to a host group and entered into an appropriate policy row in the policy table.
- A list of currently shunned IP addresses can be exported to a file.
- A maximum of 128K (131,072) IP addresses can be shunned at any one time.

Typical Scenarios for Using Shunning

Table 19-4 lists some common reasons to use the shunning function.

Table 19-4: Common Shunning Scenarios

Purpose	Description
Network Slowdown Determined	If the network security administrator is notified that there is a slowdown in the network, reviewing the blocked and detected attacks and using the security event viewer may enable the person to identify a set of attacker IP addresses. These addresses can be selected and shunned for a period of time. The administrator can verify that the traffic is being blocked from these IP addresses because the device generates security events for each shunned packet.
Initial Shunning Not Effective	If the network slowdown is not resolved by shunning the IP addresses, the administrator may chose to unshun them and determine a different approach to diagnosing the root cause.
Good IP Addresses Shunned by Mistake	If the network security administrator is notified that traffic from one or more IP addresses is being blocked and that it should not be blocked, the addresses being shunned can be reviewed to determine if this is the reason they are blocked, and if so, selectively unshun them.
Using a Router to Block Attackers	The administrator can export a list of IP addresses being shunned so that they can be imported into a router for blocking by the router.
Changes in the Threat	If the administrator suspects that a set of IP addresses being shunned pose either more or less of a threat than when they were initially shunned, the shun time can be extended or shortened. A subset of the IP addresses can also be given a different shun duration by moving them to a new shun label.

Shunning IP Addresses

During an attack, you would typically use the Security Event Viewer to identify several events of interest. Once you have selected these events, you can then create a shun label so the IP addresses associated with these events can be effectively blocked.

NOTE

Shunning is only available on E-series IPS Unit models.

To shun one or more IP addresses:

1. To shun only addresses selected from the Security Event Viewer:
   a. Use the Blocked and Detected Attacks dialog box or click the Security Events toolbar button to display the Security Event Viewer.
   b. On the Security Event Viewer, deselect (clear) the Active Mode check box.

      NOTE

      The Shun Attackers button remains unavailable (grayed out) until you clear the Active Mode check box.

   c. Select an event associated with each IP address you want to shun.
   d. Click Shun Attackers. The Shun Attackers dialog box displays (Figure 19-1).
2. To shun only manually entered IP addresses:
   a. Choose Monitor Security > Shunned Address Viewer from the Navigation Tree. The Shunned Address Viewer dialog box displays (Figure 19-2).
   b. To shun addresses using a new shun label, In the Labels area at the top of the Shunned Address Viewer dialog box, click New.
   c. To modify the addresses associated with an existing shun label, select the shun label, then click Edit.
   d. The Shun Attackers dialog box displays (Figure 19-1).

Figure 19-1: Shun Attackers Dialog Box

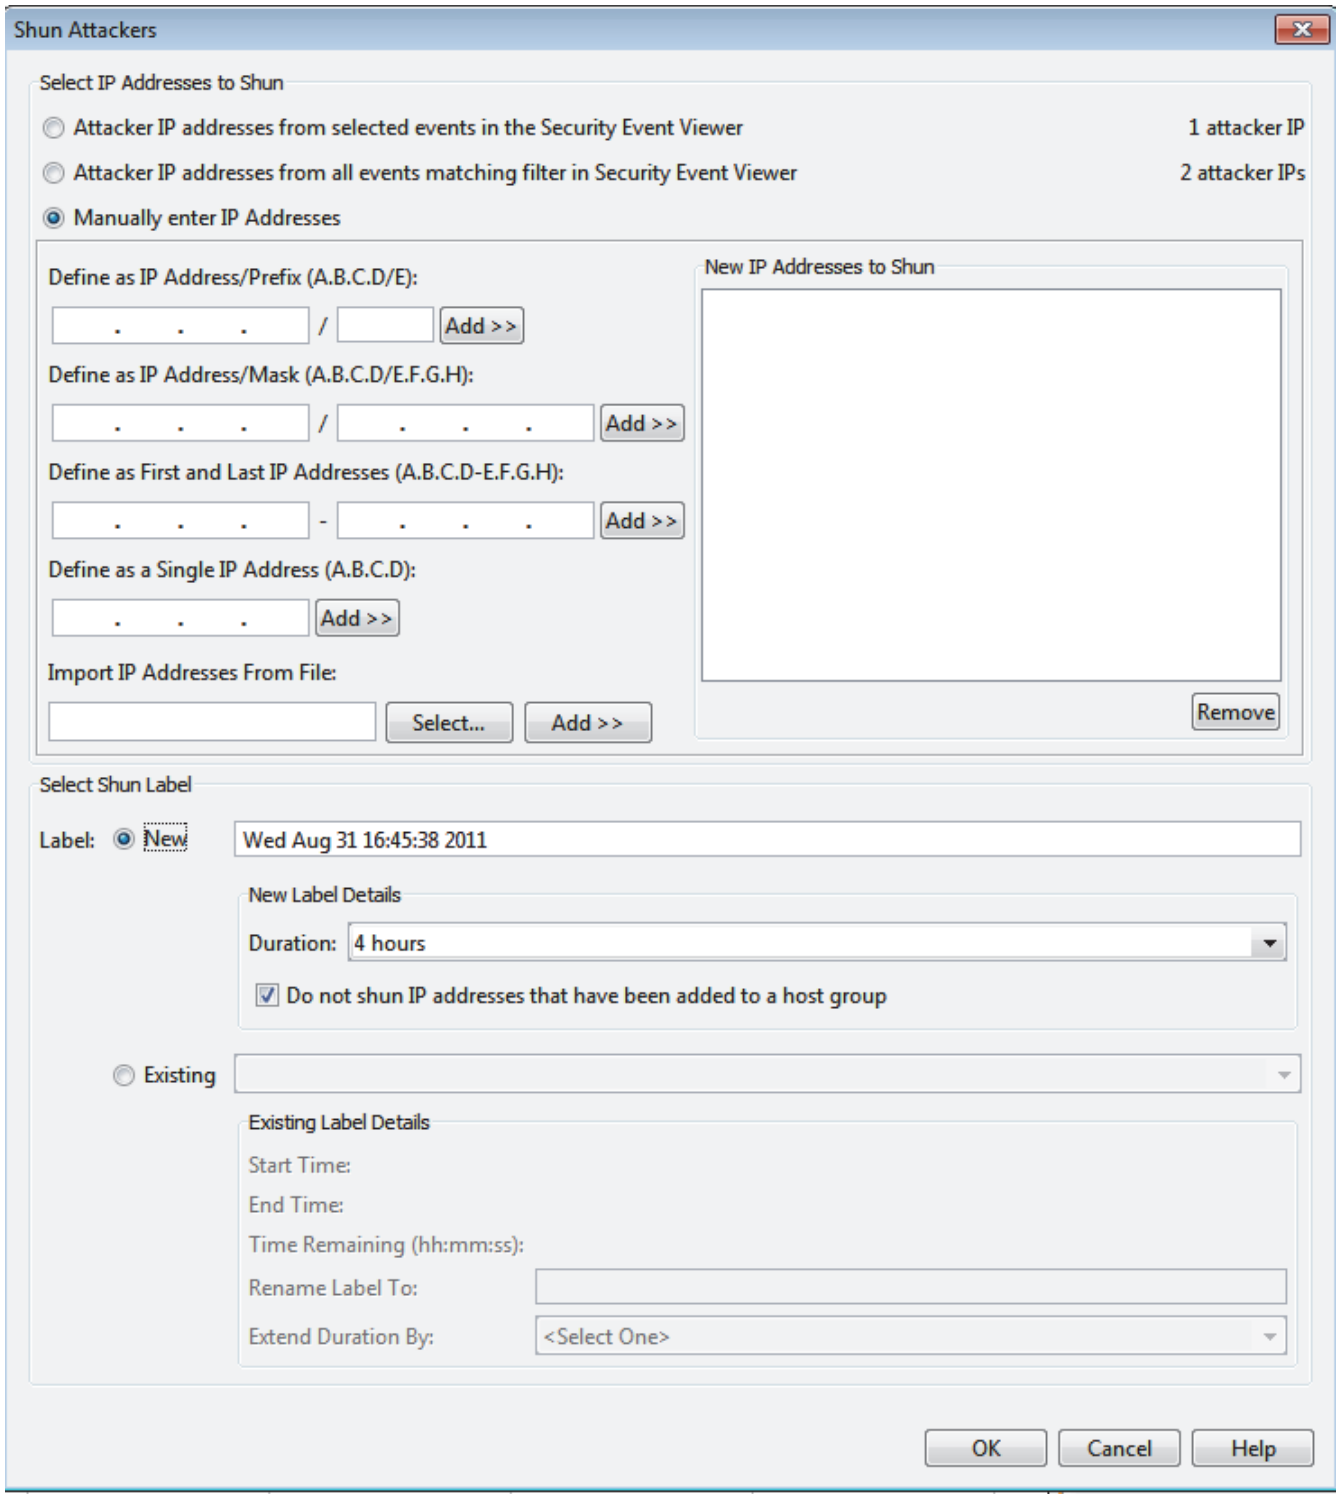

3. Choose how you would like to specify IP addresses to shun. You can select from the following options:
   - Shun attacker IP addresses from selected events in the Security Event Viewer. When you choose this option, the IP addresses are automatically selected.

- Shun attacker IP addresses from all events matching the specified filter in the Security Event Viewer. When you choose this option, the IP addresses are automatically selected.
- Manually enter (or import) IP addresses you would like to shun.

4. If you opted to manually enter (or import) IP addresses you would like to shun, you can specify addresses in one or more ways.
   - To **Define as IP Address/Prefix**, enter the IP address and prefix information in the appropriate locations, then click Add>>.
   - To **Define as IP Address/Mask**, enter the IP address and mask information in the appropriate locations, then click Add>>.
   - To **Define as First and Last IP Addresses** (specifying a range), enter the first and last addresses in the range in the appropriate locations, then click Add>>.
   - To **Define a Single IP Address**, specify the IP address in the appropriate location, then click Add.
   - To **Import IP Addresses From a File**, click Select to choose the file containing a list of the addresses, then click Add>>.

     IP addresses to be shunned can be imported from a csv file. In the file (which must have the .csv file extension), each line is either an IP address, or an IP address range defined by two IP addresses separated by a comma. Prefixes and masks for IP addresses cannot be used with the import function.

     NOTE

     If you accidentally add an incorrect address, select the address in the list area to the right and click Remove.

5. Once you have finished either automatically or manually selecting IP addresses, you need to select a Shun Label. There are two ways to do so:
   - If you want to specify a new shun label, click the New radio button. Enter a duration, specifying the period of time that you want the address shunned. You can select time periods as short as five minutes, or as long as an unlimited period of time (the address is shunned until you delete the shun label, or unshun the address).
   - If you want to specify an existing shun label, click the Existing radio button. From the drop-down list, select the desired shun label. The label details display. You can optionally rename the shun label, or extend the shunning duration by a specified period of time.
6. Selecting the Do Not Shun IP Addresses That Have Been Added to a Host Group check box enables you to specify that you want the system to block attackers, but you want to avoid blocking IP addresses that are trustworthy (those that are defined in an existing host group).
7. When finished, click OK.
8. Save your changes by clicking the Save Configuration toolbar button.

Viewing and Managing Shunned Addresses

The shunned address viewer selection from the Navigation Tree is a security monitoring and management tool that enables you to easily examine and update the IP addresses that are currently being shunned.

Shunned IP addresses are grouped by shun label. You can view information about each shun label, such as how many IP addresses are in it, shun time remaining, and the total number of packets dropped by all IP addresses within that shun label group. You can also view information about each shunned address, such as the shun label currently applied to it, and the shunning start and end time.

The viewer includes features that enable you to manipulate the data, and quickly modify the security management steps.

NOTE

Shunning is only available in E-series IPS Unit models.

Management features of the viewer include the following:

- Use powerful filter mechanisms that enable you to display particular IP addresses.
- Remove IP addresses from a shun label.
- Extend the shunning time for IP addresses in a shun label.
- Export a full list of IP addresses being shunned for archiving, additional analysis, and use in other programs and products.

NOTE

All references to time are to the time on the Corero Network Device, which may differ from the time on the user's workstation running the GUI if the user is in a different time zone to the device.

To view shunned addresses:

1. Do one of the following:
   - Click the Shunned Addresses Toolbar button.
   - From the Navigation Tree, choose Monitor Security > Shunned Address Viewer.

   The Shunned Address Viewer displays (Figure 19-2).

Figure 19-2: Shunned Address Viewer

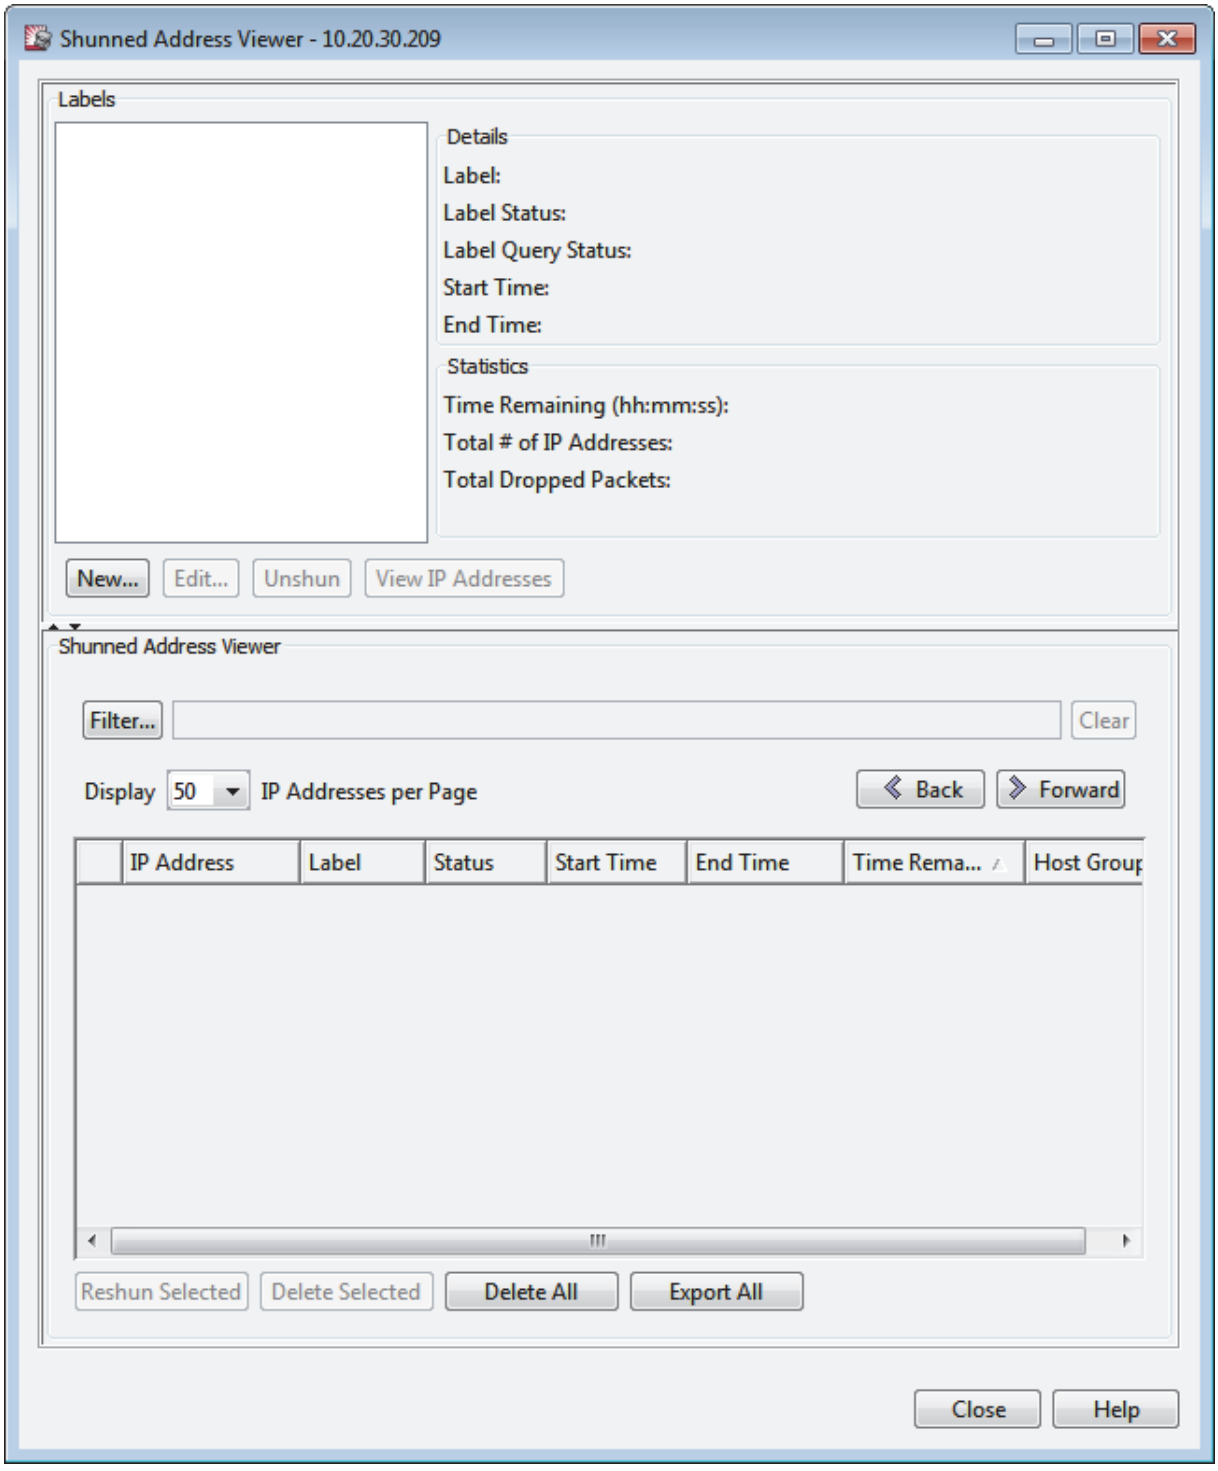

2. To view information about a shun label, select the shun label in the Labels area, then click View IP Addresses. The information described in Table 19-5 displays.

NOTE

The information displayed by the Shunned Address Viewer is not dynamically updated. To receive an updated display you must close and re-open the page.

Table 19-5: Shun Label Parameters

Parameter	Description
Label	The name of the shun label.
Label Status	Label status is one of the following: • Shunning — IP addresses are being actively shunned. • Deleting — The device is in the process of deleting IP addresses from the shun label so they will no longer be shunned. • Cancelled — The user initiated a cancellation of the shunning action, so the IP addresses associated with this label are not currently being shunned. • Expired — The timer for the shun label has expired, so the IP addresses associated with this label are not currently being shunned.
Label Query Status	The Status is one of the following: • Processing — IP addresses are being associated with this shun label so that they implement the shun label time period. • Done — All IP addresses associated with this shun label are now applying shun label actions.
Start Time	The start time for the shunning associated with this shun label.
End Time	The end time currently scheduled for the shunning associated with this shun label.
Time Remaining	The remaining amount of time left for IP addresses associated with this shun label to be blocked.
Total # of IP Addresses	The number of IP addresses currently associated with this shun label.
Total Dropped Packets	The total number of packets dropped (blocked) as a result of shunning by IP addresses associated with this shun label.

3. To add a shun label and associate it with one or more IP addresses, click New. The Shun Attackers dialog box displays with the New radio button selected. For information on how to shun attackers, see Shunning IP Addresses (page 19-6).
4. To modify a shun label:
   a. Select the shun label in the Labels area, then click Edit. The Shun Attackers dialog box displays with the Existing radio button selected.
   b. You can now modify shun label information including renaming the shun label, extending the duration of the shun label, and modifying the IP addresses associated with the label. The settings on the Shun Attackers dialog box are described in Shunning IP Addresses (page 19-6).
5. To unshun (discontinue shunning) all addresses in a shun label, select the shun label in the Labels area, then click Unshun.

NOTE

When you unshun a label, the IP addresses will remain internally associated with the shun label until the Corero Network Device needs to reuse the internal table space in

which the information is stored. This gives the user the ability to reshun the IP addresses at a later time. The duration in which the IP addresses remain associated with the shun label will vary depending upon how quickly additional IP addresses are being shunned.

6. To view information about the IP addresses associated with a shun label, select the shun label in the Labels area, then click View IP Addresses. The IP addresses then display in the list at the bottom of the page. Detailed information is displayed for each shunned IP address in the viewer, as listed in Table 19-6.

NOTE

The information displayed by the Shunned Address Viewer is not dynamically updated. To receive an updated display you must close and re-open the page.

Table 19-6: Shunned IP Address Parameters

Parameter	Description
IP Address	The IP address.
Label	The name of the shun label associated with this IP address.
Status	The status is one of the following: • Processing — The IP address is in the process of being associated with a shun label and its time period. It is not yet being blocked according to the shun label settings. o that they implement the shun label time period. • In-Sync — The displayed IP address is associated with the shun label specified by the user, and is currently being blocked according to the shun label time period. • Deleting - The IP address is in the process of being deleted from the shun label. This process can take some time, depending on the number of IP addresses being removed. Once the IP address has been deleted, it will no longer display in the IP address list associated with the selected shun label.
Start Time	The start time for the shunning associated with this shun label.
End Time	The end time currently scheduled for the shunning associated with this shun label.
Time Remaining	The remaining amount of time left for IP addresses associated with this shun label to be blocked.
Host Group	The host group with which this IP address is associated (if any).

7. You can filter the IP addresses listed at the bottom of the page. For more information, see Shunned Address Viewer Filtering (page 19-14)
8. If you want to reshun IP addresses in a shun label:
   a. Select the desired shun label in the Labels list.
   b. Ensure that the shun label has the Cancelled or Expired label status.
   c. If desired, specify filter in formation as described in Shunned Address Viewer Filtering (page 19-14).
   d. Select one, some, or all of the IP addresses.
   e. Click Reshun Selected.
9. If you want to remove IP addresses from a shun label:
   a. Select the desired shun label in the Labels list.

b. If desired, specify filter in formation as described in Shunned Address Viewer Filtering (page 19-14).

c. Select one, some, or all of the IP addresses.

d. Click Delete Selected.

10. If you want to export a list of all IP addresses associated with a shun label:

a. Select the desired shun label in the Labels list.

b. If desired, specify filter in formation as described in Shunned Address Viewer Filtering (page 19-14).

c. Select one, some, or all of the IP addresses associated with that shun label.

d. Click Export All. The addresses are sent to a CSV file, which you can save.

11. When finished, save your changes by clicking the Save Configuration toolbar button.

Shunned Address Viewer Filtering

The Shunned Address Viewer includes a powerful filtering tool that enables you to zero in on a specific set of IP addresses from the full list associated with a selected shun label.

NOTE

Shunning is only available in E-series IPS Unit models.

To access the Shunned Address Viewer Filter tool,

1. Do one of the following:
   - Click the Shunned Addresses Toolbar button.
   - From the Navigation Tree, choose Monitor Security > Shunned Address Viewer.

   The Shunned Address Viewer dialog box displays (Figure 19-2).
2. In order to populate the filter with information from a particular shun label, select the shun label in the Labels list, then choose View IP Addresses.
3. If you want to specify more detailed filtering information, click Filter. The Shunned Address Filter dialog box displays (Figure 19-3).

Figure 19-3: Shunned Address Filter Dialog Box

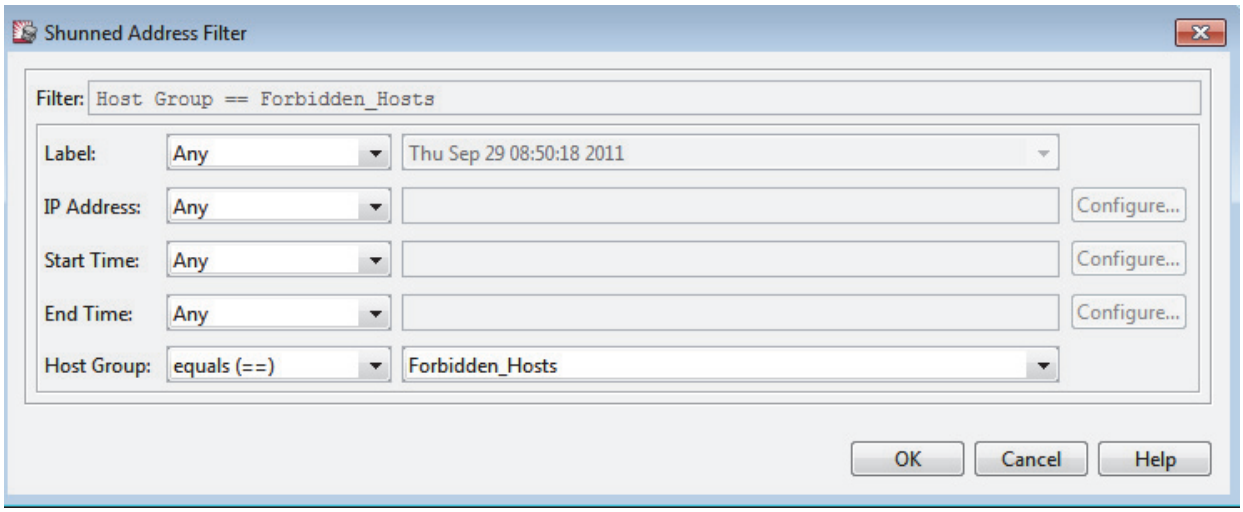

4. When specifying a filter, you can specify how you want the filter results to match. Note that not all filter options are available in each category.
   - Any - Include results matching any option in this category.
   - Equals - Include results matching only the option you specify in this category.
   - Does Not Equal - Only include results that do not match the option you specify in this category.
   - Less Than - Only include results that are less than (lower than or before) the specified value.
   - Greater Than - Only include results that are greater than (higher than or after) the specified value.

   The Security Event Filter tool enables you to filter based on the categories listed in (Table 19-7).

Table 19-7: Security Event Filter Options

Filter	Description
Label	Specifies the shun label associated with the IP addresses.
IP Address	Specifies the IP addresses
Start Time	Specifies the start time for the current shunning period.
End Time	Specifies the end time for the current shunning period.
Host Group	The host group associated with the IP addresses.

Viewing Blocked and Detected Attacks

The Blocked and Detected Attacks window dynamically displays information about current attacks, automatically sorted by rule. This page is the first place you should go to look for current security issues.

By default, the management application only displays the events to which it has reacted. If you want it to display all event types (all detection rules), even those with no recorded events, uncheck the Hide Rules with Zero Events check box.

To view blocked and detected attack information:

1. Do one of the following:
   - Select Monitor Security > Blocked and Detected Attacks from the navigation tree.
   - Click the Blocked & Detected toolbar button.

   The Blocked and Detected Attacks dialog box displays (Figure 19-4).

Figure 19-4: Blocked and Detected Attacks Dialog Box

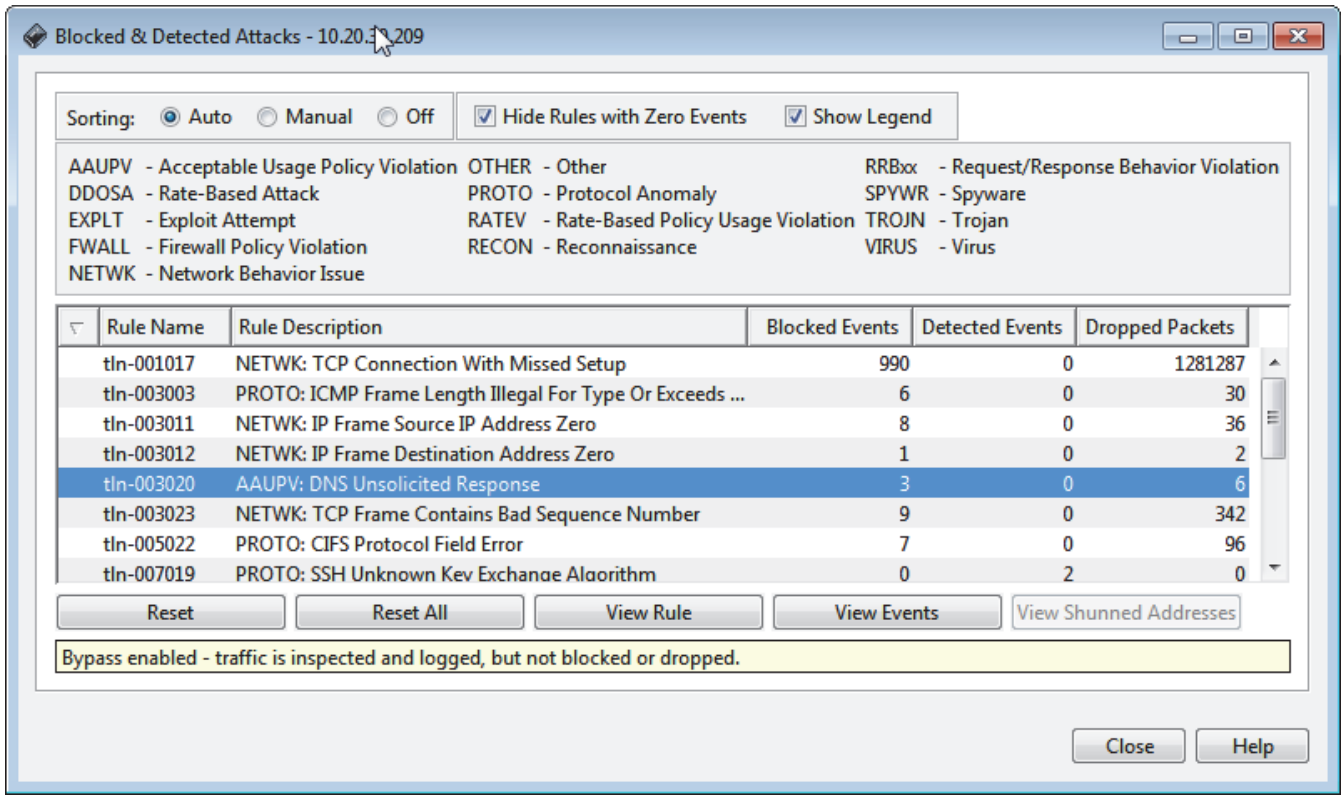

The Blocked and Detected Attacks page displays the information listed in Table 19-8.

Table 19-8: Blocked and Detected Attacks View

Information	Description
Arrow Indicator	An arrow indicator in the left-most column indicates how recent the attack was. • A red arrow indicates a very recent attack (within the past few seconds). • A grey arrow indicates a less recent attack.

Table 19-8: Blocked and Detected Attacks View *(Continued)*

Information	Description
Rule Name	The name of the triggered rule.
Rule Description	A brief description of the triggered rule, including the five-character rule category.
Blocked Events	The number of blocked events associated with this rule since the last time this item was reset. A Blocked Event is raised only once for a given flow. The packet that causes the Blocked Event is also counted as one Dropped Packet. Any further packets that arrive for that flow are also dropped, but it still only counts as one event. This means the Dropped Packet Count is usually a larger number than the Blocked Event Count
Detected Events	The number of packets detected associated with this rule since the last time the counter for this attack was reset.
Dropped Packets	The number of dropped packets associated with this event.

NOTE

At the bottom of the page, a status area indicates whether bypass is enabled.

2. By default, the information in the display is automatically sorted. If you would like additional details on a specific attack, you can manually sort the list by selecting the Manual radio button, then clicking the header of the column on which you want the list sorted.
3. If you would like the position of the rows in the list to remain fixed, you can turn off sorting by selecting the Off radio button.

NOTE

When you turn off sorting, the Management Application continues to update the Blocked and Detected values for each row. However, it stops dynamically repositioning the rows based on the number of detected events.

4. To view additional information about the rule associated with an attack, select the attack, then click View Rule. The Rule Details dialog box displays the following information:
   - The rule's name and description.
   - You can specify a limit (the Event rate Limit) to how many events can be sent per rule per minute to the logs and the Security Event Viewer. This helps ensure these event listings are not overwhelmed by frequent triggering of a single rule. The default limit is 300 events per rule every minute.

     You can modify this setting by clicking Edit.
   - Internet references associated with this rule.
   - The rule's confidence level.
5. To view detailed information about the events associated with an attack, select the attack, then click View Events. The events display in the Security Event Viewer. For more information, see Viewing Security Events and Security Event Details (page 19-23).
6. To view detailed information about shunned addresses associated with an attack, select the attack, then click View Shunned Addresses. Shunned address information displays in the Shunned Address Viewer. For more information, see Viewing and Managing Shunned Addresses (page 19-9).

7. To reset the Blocked, Detected, and Dropped Packet counters for an individual detection rule, select the attack and click the Reset button.
8. To reset the event counters for all the rules, click the Reset All button.

CAUTION

Counters provide historical data for generated reports. When you select Reset All and these counters are reset, the historical data you cleared will not be available for inclusion in generated reports for that time period.

About the Security Event Viewer

The Security Event Viewer is a security monitoring and management tool that enables you to easily examine and react to the traffic that triggers security rules. The viewer includes features that enable you to filter and focus on event details, then quickly take security management steps.

Using the Security Event Viewer, you can:

- View basic event information such as severity, action taken by the device, rule triggered, and date and time of the event.
- View traffic details such as protocol, IP address, and port information.
- Display real-time or historical event data.
- Display a detailed description of the rule that triggered the event.
- Jump to, and modify if desired, the policy line entry that identified the event.
- Filter the display to focus on a smaller set of events based on criteria you specify.
- Examine the actual attack packets associated with events.
- Download event information in PCAP and CSV formats for archiving and additional analysis.
- Block attackers associated with selected or all security events by shunning them. Shunning enables you to quickly and temporarily block all traffic initiated by IP addresses that are suspected of originating an attack or otherwise identified as requiring that their traffic be blocked.

NOTE

The Shunning feature is only available in the E-series IPS Unit models.)

When using the Security Event Viewer, consider the following:

- You should have no more than three Management User Interfaces running simultaneously with the Security Event Viewer in Active Mode. Otherwise, the Security Event Viewer may not update events in a timely fashion.
- When using the Security Event Viewer > Download PCAP button, the PCAP (packet capture) file contains all packets that match the current Security Event Filter criteria. The maximum file size is one Megabyte of data, and older packets are listed first. If there are too many packets that match the selected criteria, modify the criteria using the Filter button to obtain a smaller set of events.

Figure 19-5 shows the Security Event Viewer in the management application for a Corero Network Device.

Figure 19-5: Corero Network Device Security Event Viewer

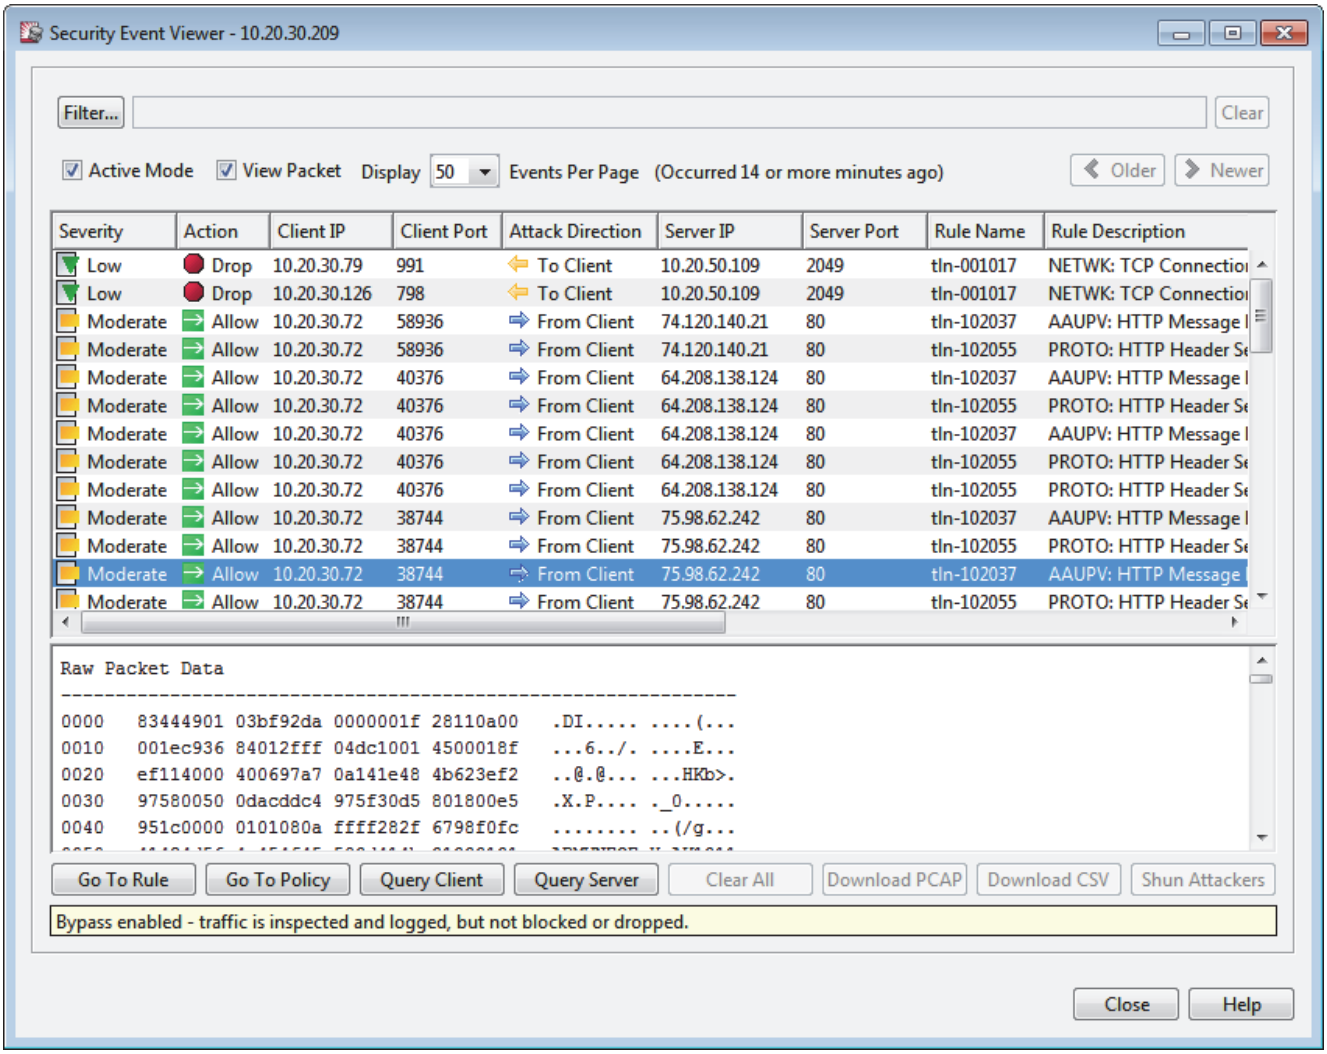

Figure 19-6 shows the Security Event Viewer in the IPS Controller management application.

Figure 19-6: IPS Controller Security Event Viewer

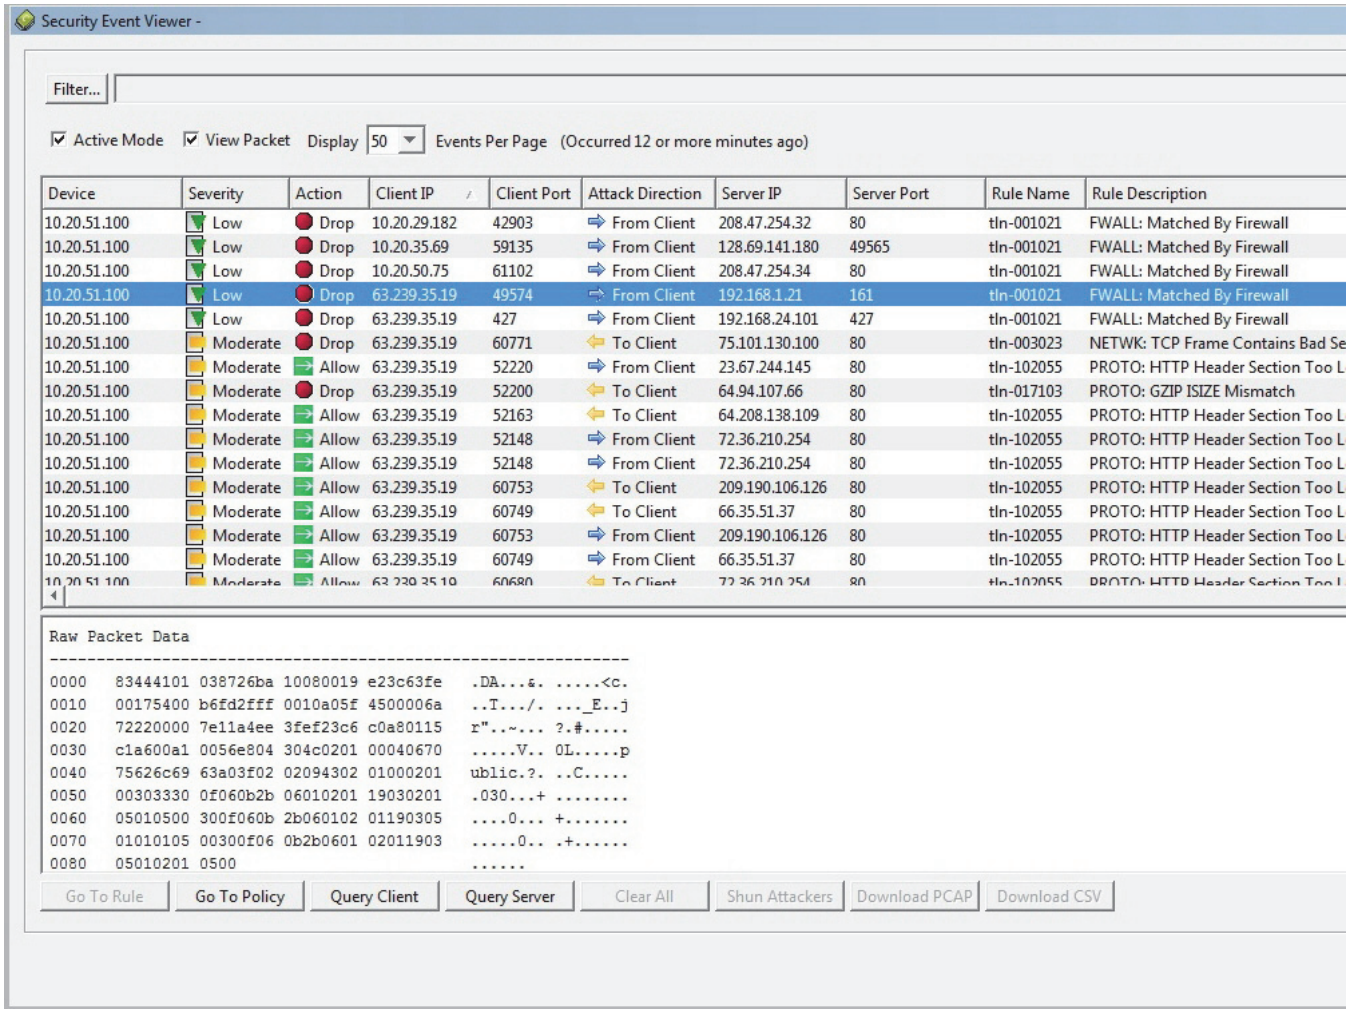

The information shown in the Security Event Viewer is described in Table 19-9.

Table 19-9: Security Event Viewer Information

Column	Description
Severity	The level of danger this type of event poses: low, moderate or critical. The value displayed is based on the value currently assigned to the rule that triggered this event. All rules come with a preconfigured severity level. You can modify the severity setting by editing the rule's parameters. For information on modifying rule parameters, see Modifying Rule Settings (page 15-14). NOTE: To view a description of the rule that triggered an event, select the event and click the View Rule button.

Table 19-9: Security Event Viewer Information *(Continued)*

Column	Description
Action	Indicates how the Corero Network Device handled the traffic that triggered the event. Actions may be any of the following: • Allow— Pass the traffic. • Drop— Passively block the traffic by dropping the traffic with no indication to the offending client. • Reject— Actively block the traffic by dropping the traffic and taking specific action, for example, by resetting the connection. **Note:** The action that the device takes when traffic triggers a rule is part of the settings for that rule in the currently selected rule set. You can modify a rule's action by clicking the Go To Rule button. For information on modifying rule parameters, see Modifying Rule Settings (page 15-14).
Client IP	The source IP address for the event's traffic.
Client Port	The logical source port associated with the network protocol for this event.
Attack Direction	Indicates whether the attack occurs for inbound or outbound traffic.
Server IP	The destination IP address for the event's traffic.
Server Port	The logical destination port associated with the network protocol for this event.
Rule Name	The internal alpha-numeric designation for the rule that triggered the event.
Rule Description	A short description of the rule that triggered this event. To view more detail, select the event and click the Go To Rule button. Every rule begins with a security category prefix, which enables you to sort the rules based on their security category. See Table 15-1 or a listing of the rule prefixes.
Timestamp	The date and time of the event.
Event Number	The Event Logging System (ELS) assigns a unique number to each event. The viewer can display the last 160,000 of these unique events. A copy of the Event Logging System online help is accessible on the Documentation CD-ROM supplied with your device.
Protocol	The network protocol, if applicable, of the traffic that caused this event.
Origin Port	The physical port that was the source for this traffic.
HA ID	If applicable, the ID of the high availability device associated with this event.

Viewing Security Events and Security Event Details

To view security events using the Security Event Viewer:

1. Do one of the following:
   - From the Navigation Tree, choose Monitor Security > Security Event Viewer.
   - From the Toolbar, click the Security Events button.

   The Security Event Viewer dialog box displays (Figure 19-5).

   It displays a multi-page table listing security event information.
2. To update the display as events occur, select the Active Mode check box. This box is selected by default.
3. If you want to view more or fewer events in the list at one time (page size), enter the desired number of Events Per Page.
4. You can change the time frame for displayed event.
   - View events before the ones currently displayed by clicking Older.
   - View more recent events by clicking Newer.
5. To sort the events in a particular order, select the heading of the column by which you want the table sorted.
6. If you want to view the raw data for the packet triggering an event, select the View Packet check box. When this box is selected, any time you click an event in the Security Event Viewer table, the packet data displays at the bottom of the window.
7. If you want to perform tasks using a current snapshot of event activity, and do not want newer events to display, deselect (clear) the Active Mode check box.

   NOTE

   The Active Mode check box must be cleared in order to access many Security Event Viewer features.
8. The Security Event Viewer includes a powerful filter tool that enables you to display a user-specified set of events so you can find those that are relevant to the current attack situation. For instructions on how to filter the list of displayed events, see Security Event Viewer Filter Tool (page 19-25).

   To clear any filter settings you have specified, click Clear. All events display again.
9. More advanced feature permit you to view more detailed information or perform a particular action related to a selected event. Table 19-10 lists advanced Security Event Viewer features, and how to access them. Note that features may be accessed from buttons, a pop-up menu, or both.

Table 19-10: Advanced Security Event Viewer Features

Select the...	And ensure the Active Mode check box is...	In order to...
Go To Rule button *or* Go To Rule option from the pop-up menu	Either selected (checked) or deselected (cleared)	Display the rule that was triggered by the selected event. **Note:** If you have changed any of the conditions that were established for the FW+IPS policy entry that triggered this event, the search for the rule will fail, and an error message will display. Moving a policy entry line up or down, or changing the treatment for an entry should not affect Go To Rule results.

Table 19-10: Advanced Security Event Viewer Features *(Continued)*

Select the...	And ensure the Active Mode check box is...	In order to...
Go To Policy button *or* Go To Policy option from the pop-up menu	Deselected (cleared)	Display the policy that triggered the selected event. **Note:** If you have changed any of the conditions (segment, client, server, service) that were established for the FW+IPS policy entry that triggered this event, the search for the policy will fail, and an error message will display. Moving a policy entry line up or down, or changing the treatment for an entry, should not affect Go To Policy results.
Query Client button Query Server button *or* Query Client or Query Server option from the pop-up menu	Either selected (checked) or deselected (cleared)	Performs an IP Query on the client or server associated with the selected event, and displays details about the behavior of a client or server host as it is requesting and completing connections. For more information on querying an IP address, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).
Clear All button *or* Go To Policy option from the pop-up menu	Deselected (cleared)	Clear all displayed events. **Note:** This applies only to the current management session, and only to the current management user.
Download PCAP button *or* Download PCAP option from the pop-up menu	Deselected (cleared)	Download the data associated with the event to your management station in PCAP (packet capture) form for further analysis. Once the data is downloaded, you can save it to a file.
Download CSV button *or* Download CSV option from the pop-up menu	Deselected (cleared)	Download the event data to your management station in CSV (comma separated value) format. Once the data is downloaded, you can save it to a file. Then you can import the data into a spreadsheet or other data analysis tool.
Shun Attackers button *or* Shun Attackers option from the pop-up menu	Deselected (cleared)	Create a shun label to block the attacking IP address associated with the selected event. For more information on shunning, see About Using IP Address Shunning to Stop an Attack (page 19-4). **Note:** This feature is only available for E-series IPS Unit models.
Quick Filter option from the pop-up menu	Deselected (cleared)	Filter the entire set of events based on the value in the currently selected event. You can filter based on Severity, Rule, Protocol, Client IP, Client Port, Server IP, Server Port, and Origin Port.
WHOIS Lookup option from the pop-up menu	Deselected (cleared)	Perform a WHOIS lookup on the IP address over which you have placed the cursor.
Reverse DNS Lookup option from the pop-up menu	Deselected (cleared)	Perform a reverse DNS lookup on the IP address over which you have placed the cursor.
Filter by Selected Row option from the pop-up menu	Deselected (cleared)	Launch the Filter tool prepopulated with any information associated with the selected row.

Security Event Viewer Filter Tool

The Security Event Viewer includes a powerful filtering tool that enables you to zero in on a specific set of events from the hundreds or thousands of events currently stored in memory.

The tool, shown in Table 19-11, provides several categories for filtering data. Some categories apply to both active data (events as they are occurring) and inactive data (fixed set of events), and some apply only when you are using inactive mode. For each category, you can choose one of the following three operators:

- Any — This category does not limit the set of displayed events.
- Equals — Only events that equal a chosen value in this category are displayed.
- Does not equal — Only events that do not equal a chosen value in this category are displayed.

Figure 19-7: Security Event Filter Tool

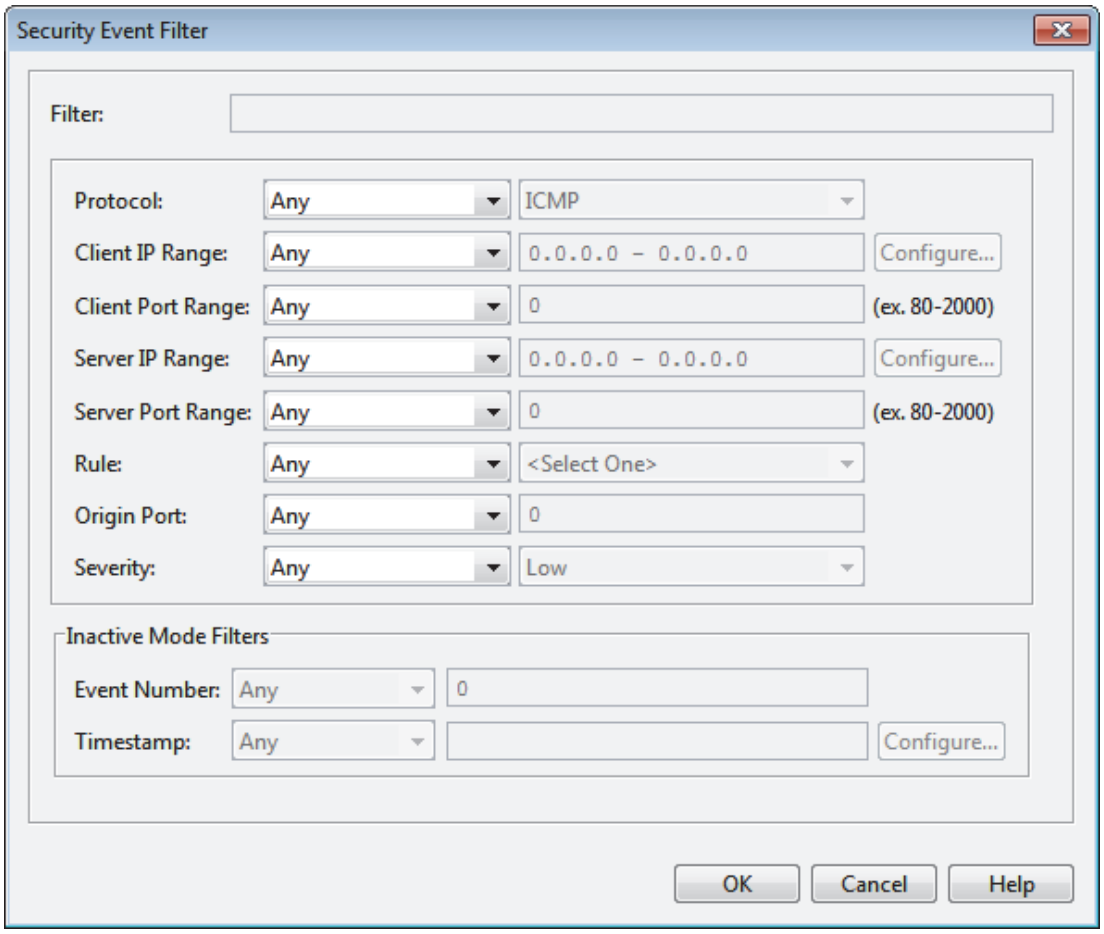

To access the Security Event Viewer Filter tool,

1. Do one of the following:
   - From the Navigation Tree, choose Monitor Security > Security Event Viewer.
   - From the Toolbar, click the Security Events button.

   The Security Event Viewer dialog box displays (Figure 19-5).

2. In order to populate the filter with information from a particular event, select the event, then choose Filter By Selected Row from the pop-up menu.
3. If you want to specify more detailed filtering information, click Filter. The Security Event Filter dialog box displays (Figure 19-7).

When specifying a filter, you can specify how you want the filter results to match. Note that not all filter options are available in each category.

- Any - Include results matching any option in this category.
- Equals - Include results matching only the option you specify in this category.
- Does Not Equal - Only include results that do not match the option you specify in this category.
- Less Than - Only include results that are less than (lower than or before) the specified value.
- Greater Than - Only include results that are greater than (higher than or after) the specified value.

The Security Event Filter tool enables you to filter based on the categories listed in (Table 19-11).

Table 19-11: Security Event Filter Options

Filter	Description
Protocol	Specify the network protocol associated with each packet (ICMP, TCP, UDP).
Client IP Range	Specify the client IP range whose events you want to view. To add a client IP range, click Add. The Add Client IP Range dialog box displays. You can add IP addresses in four ways: • As an IP address/Prefix (for example 192.0.8.31/24, where the first 24 bits would have to match exactly.) • As an IP address/Mask (for example 192.0.8.31/255.255.255.0). • As a range specified by a first and last IP address (for example 192.0.128.0-192.01.128.255). • As a single IP address (for example 192.0.8.31).
Client Port Range	Specify the client port range whose events you want to view.
Server IP Range	Specify the server IP range whose events you want to view. To add a client IP range, click Add. The Add Client IP Range dialog box displays. You can add IP addresses in four ways: • As an IP address/Prefix (for example 192.0.8.31/24, where the first 24 bits would have to match exactly.) • As an IP address/Mask (for example 192.0.8.31/255.255.255.0). • As a range specified by a first and last IP address (for example 192.0.128.0-192.01.128.255). • As a single IP address (for example 192.0.8.31).
Server Port Range	Specify the client port range whose events you want to view.
Rule	In order to view events associated with a particular rule, select that rule.
Origin Port	In order to view events associated with a particular port on the device, select that port.
Severity	Select the severity of the events you want to view.
Action	Specify the action associated with the rule.

Table 19-11: Security Event Filter Options *(Continued)*

Filter	Description
Event Number	Specify an event number, then select whether you want to view events above or below that number. **Note:** You can only use this filtering criteria if the Active Mode check box on the Security Event Viewer is cleared (deselected).
Time	Specify a date and time, then select whether you want to view events before or after that time. **Note:** You can only use this filtering criteria if the Active Mode check box on the Security Event Viewer is cleared (deselected).

Using IP Address Query to Learn About a Host and Clear Counters

You can use the IP Address Query feature to view details about the behavior of a selected host when requesting and completing connections. Once the device begins limiting an IP address that is producing a large number of incomplete connection requests or a higher than allowed number of completed requests, the device continues to block traffic from that address until enough time has passed to ensure that the IP address is producing normal activity.

NOTE

In addition, from the Security Event Viewer, if you clear (deselect) the Active Mode check box, you can place your cursor over an IP address and choose WHOIS Lookup or Reverse DNS Lookup for additional information on the address.

To query an IP address using the management application for a Corero Network Device:

1. Do one of the following:
   a. From the Navigation Tree, choose Monitor Security > Query IP Address.
   b. From the Security Event Viewer, with the Active Mode check box cleared (deselected), select an event then click Query Client or Query Server.

   The IP Address Query dialog box displays (Figure 19-8).

Figure 19-8: Corero Network Device IP Address Query Dialog Box

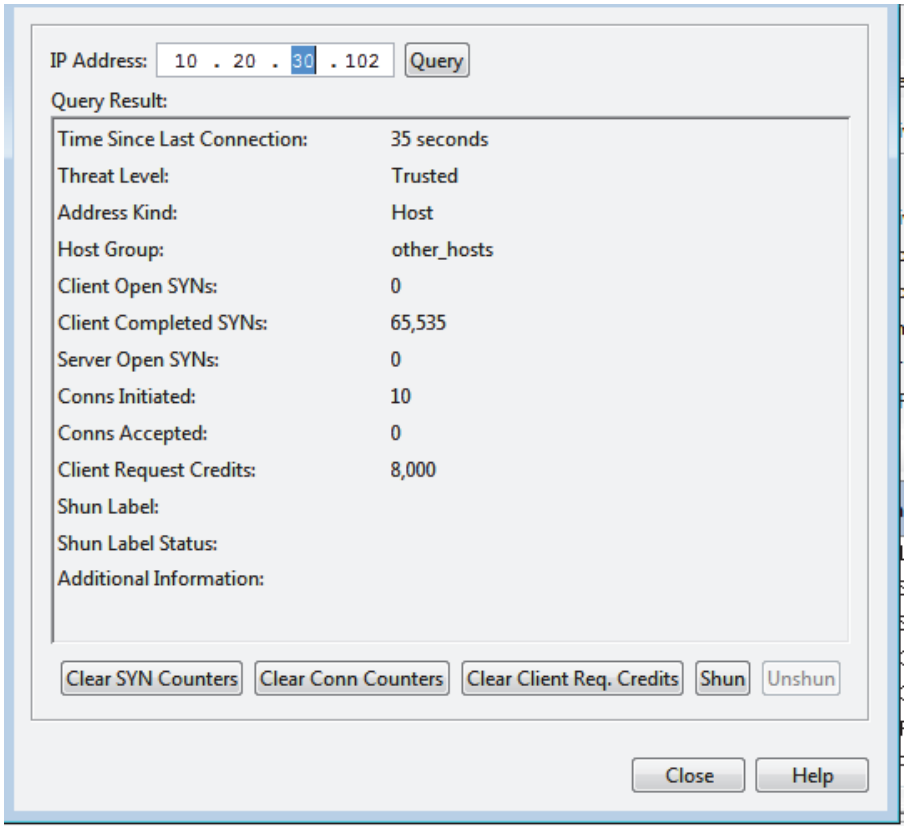

2. Enter the IP address of the host about which you want to view additional information, then click Query. The IP Address Query dialog box displays the query results.

   The IP Address Query dialog box displays the information listed in Table 19-12.

Table 19-12: Corero Network Device IP Address Query Information

Item	Description
Time Since Last Connection	The number of seconds since this IP address was last seen as a client in a completed connection.
Threat Level	Current SYN flood threat level that the device has assigned to this IP address: • Unknown — TCP connection patterns are not yet identified for this address. • Trusted — This IP address has made a user-specified number of completed connection requests and the number of initiated, but not completed requests is below the user-defined Suspicious threshold. The device forwards requests from this address to the destination address. • Suspicious — This IP address is initiating enough incomplete TCP connections to cause the device to proxy any connection request to any device from this source IP address. • Malicious — This address is creating a dangerous level of incomplete connections. The device blocks any connection requests from this address until the address stops initiating requests for a timeout period, or the number of incomplete requests decays to the Suspicious level.
Address Kind	Indicates that this is a host address.
Host Group	The host group, if applicable, for this IP address.
Client Open SYNs	If the queried address is that of a client, this number is the total number of uncompleted connection requests that this IP address has generated since the value was last reset. The Clear SYN Counters button will reset this value to zero. This counter decreases gradually over time.
Client Completed SYNs	If the queried address is that of a client, this number is the total number of completed connection requests that this IP address has generated since the value was last reset. The Clear SYN Counters button will reset this value to zero. This counter decreases gradually over time.
Server Open SYNs	If the queried address is that of a sever, this number is the total number of uncompleted connection requests directed at this server since the value was last reset. The Clear SYN Counters button will reset this value to zero. This counter decreases gradually over time.
Conns Initiated	Total number of currently active TCP connections. The Clear Conn Counters button will reset this value to zero. This counter decreases gradually over time.
Conns Accepted	Total number of completed, and currently active, TCP connections that this address has accepted since the value was last reset. The Clear Conn Counters button will reset this value to zero.
Client Request Credits	Current number of request credits remaining for this client. If the client is “overspent”, this value will be negative. Client periodically receives additional credits based on the Request Limit value associated with the Client Group for this client's IP address.
Shun Label	If the IP is currently being shunned, this identifies the shun label which is being used to shun the IP address.
Shun Label Status	The shun label status - either shunning, cancelled or expired.
Additional Information	Miscellaneous additional information about the IP address, if any.

3. If you know that a particular IP address was being used maliciously but is no longer a threat, or that its rate of completed connections is now within proper thresholds, you can reset its SYN flood, connection counters, or client request credits manually. To reset the counters for a specific IP address, click one of the following buttons:
   - Clear Syn Counters — Treats this address as newly seen with no uncompleted SYNs.

     Note that these counters are only meaningful if SYN Flood Limits are enabled on the Client Host Group to which the IP address belongs. If SYN Flood Limits are not enabled, these counters are always zero.
   - Clear Conn Counters — Treats this IP address as if has not yet initiated or accepted any connections. (Existing connections are not affected and are not part of the new count.)
   - Clear Client Req. Credits — Sets client request credits to the Request Limit. The device resets the threat level of the IP address to Unknown, completed connections to zero, and/or client request credits to the Request Limit.

   NOTE

   For more information on clearing counters, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).

4. You can also choose to Shun (block) or Unshun the selected address using the buttons at the bottom of the page. For more information on shunning IP addresses, see About Using IP Address Shunning to Stop an Attack (page 19-4).

Reset (Clear) SYN Flood and Connection Counters

Your Corero Network Device records information regarding completed and incomplete SYN requests and the current number of active connections for each IP address. Sometimes, due to an attack or special conditions, the device may be preventing certain IP addresses from attempting to send a SYN packet or create a new connection. This could occur, for example, if hosts are used by an attacker to perform a SYN flood attack.

Once you have stopped the attack, or are otherwise comfortable that the client should be allowed to resume normal operations, you can reset the SYN flood and/or connection counters using the Clear Counters dialog box (Figure 19-9).

NOTE

Alternatively, you can use the IP Address Query window to reset counters for a specific IP address. For more information, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).

To reset, or clear, SYN flood or connection counters:

1. From the Navigation Tree, choose Monitor Security > Clear Counters. The Clear Counters dialog box displays (Figure 19-9).

Figure 19-9: Clear Counters Dialog Box

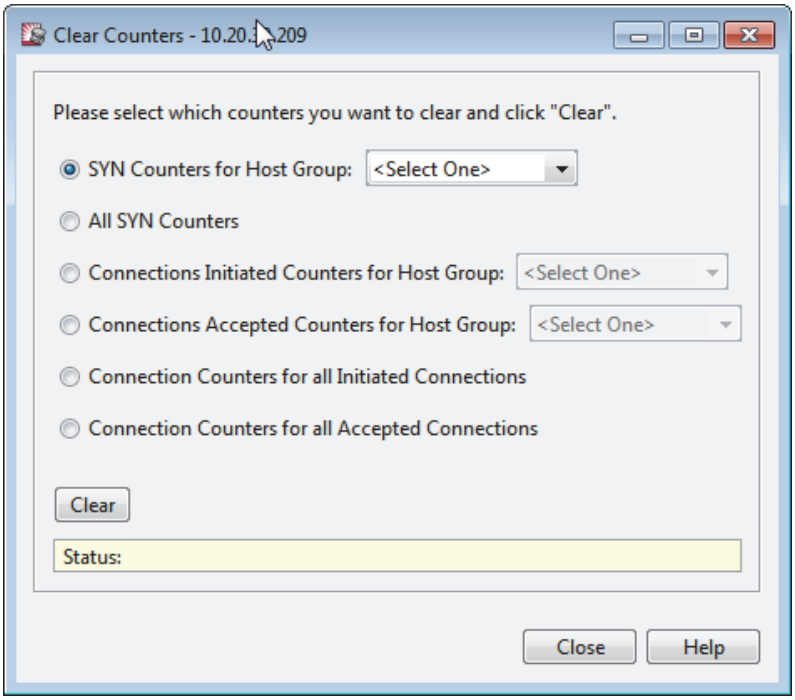

2. Select a counter to clear, as described below.

If you want to...	Then you should....
Return all clients in the selected group to their initial SYN flood state (typically the Unknown state), you must clear the SYN Flood counters for a specified host group.	Select SYN Flood Counters for Host Group, then select a host group from the list.

If you want to...	Then you should....
Return all clients in all groups to their initial SYN flood state (typically the Unknown state), you must clear all SYN Flood counters.	Select All SYN Counters.
Clear all client connection counters associated with clients in a specific host group, but not affect the counters for individual clients.	Select Connections Initiated Counters for Host Group, and select a host group from the list.
Clear all he server connection counters for the selected host group, but not affect the counters for individual servers.	Select Connections Accepted Counters for Host Group, and select a host group from the list.
Clear the connection count for all individual clients.	Select Connection Counters for all Initiated Connections.
Clear the connection count for all servers.	Select Connection Counters for all Accepted Connections.

3. When you have finished making your selection, click Clear. The Status area indicates the results of the operation.

Viewing Dropped Packet Statistics

You can view a detailed list of dropped packets, organized by the reason the packet was dropped.

To view dropped packet statistics:

1. From the Navigation Tree, select Monitor Security > Dropped Packets Statistics.

   The Dropped Packet Statistics dialog box displays.

2. This dialog box displays a table listing the current number of dropped packets and the reasons that the packets were dropped. Table 19-13 lists the reasons why packets might be dropped.

Table 19-13: Dropped Packet Reasons

Reason	Description
Load Shedding	Packets dropped when the device waited to process a new connection or packets in an existing connection to create load shedding during periods of very heavy traffic.
Received Data Link Errors	Ethernet data link errors. There are many reasons for data link errors, such as: Bad CRC on the datagram. Mismatched speeds on the Ethernet ports. For example: • One device set to 100 Mbps with the complementary device set to 10Mbps. • One device set to full duplex, with the complementary device set to half duplex.
Malformed Packets	The packet was improperly formed and could not be processed. Examples might include packets that: • Are too short • Use out-of-range addresses • Include improper options • Use the same source and destination addresses • Use invalid flags • Have invalid checksums
Malformed Fragments	Fragments of a packet were malformed and could not be processed. For example, the packet contained fragments with overlapping offset fields, or may contain an invalid fragment option.
Fragment Limiting	The maximum number of fragments for a packet was exceeded, so the packet was dropped.
Layer 2 Bridge Filtered	The packet was dropped due to layer 2 filtering. Layer 2 filtering enforces low level access control, enhances security, removes enforces layer 2 and layer 3 compliance, and mitigates layer 2 attacks.
SYN Flood	A SYN flood is a form of denial-of-service attack in which an attacker sends a succession of SYN requests to a target system in an attempt to consume enough server resources to make the system unresponsive to legitimate traffic. If a Corero Network Device receives an inappropriate number of SYN requests from one or more clients, these packets are dropped.
DDoS Rejection	During a DDoS attack, the IPS Unit requires new clients to pass a test to ensure appropriate network behavior before it will process their packets. This count indicates how many packets per second the IPS Unit dropped because it entered DDoS Rejection mode and was applying this test to new clients.

Table 19-13: Dropped Packet Reasons *(Continued)*

Reason	Description
ICMP Rate Limiting	The rate limit for ICMP traffic was exceeded, so these ICMP packets were dropped.
Client Request Limiting	The number of client requests exceeded the specified limit, so these client requests were dropped.
Connection Limited	Client or client groups exceeded their maximum number of connections, so packets for additional connections were dropped.
ALG Load Shedding	Packets were dropped due to performance-based load management activity. ALG load shedding is associated with older model Corero Network Devices.
Malformed TCP Segments	Packets containing malformed TCP segments were dropped.
Session Table Limit	The maximum number of concurrent sessions was reached, so additional session requests were dropped.
Firewall Blocked	These packets were blocked by the firewall portion of the FW+IPS policy.
IPS Blocked	These packets were blocked by the IPS Rules portion of the FW+IPS policy.
Link Outbound Congestion	The output queue capacity for one or more ports was exceeded, so these packets were dropped in order to manage system resources.
Transmit Data Link Errors	These packets were dropped due to Ethernet data link errors. Errors of this type include: • Bad CRC on the datagram. • Ethernet port speed mismatch. For example: - One device was set to 100 Mbps and the other was set to 10Mbps. - One device set to full duplex and the other was set to half duplex.

Viewing Port Statistics

The Port Statistics table reports information for the transmitted and received packets for each mission port.

To view port statistics:

1. From the Navigation Tree, select Monitor Security > Port Statistics.

   The Port Statistics dialog box displays the information listed in Table 19-14.

2. From this dialog box, you can clear (zero) all counters for one or more selected ports. To do so:
   a. Select one or more ports.
   b. Click Clear Statistics.

Table 19-14 lists the information displayed in the Port Statistics dialog box table.

Table 19-14: Port Statistics Information

Column	Description
Name	Port number.
Receive Link Util	Receive link utilization, for this port, expressed as a percentage of available bandwidth.
Transmit Link Util	Transmit link utilization, for this port, expressed as a percentage of available bandwidth.
Total Packets	Total number of packets received, including bad packets, broadcast packets, and multicast packets and 1518 octets (excluding framing bits but including FCS octets) but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
Total Octets	Total number of octets of data received on the network, including those in bad packets but excluding frame check sequence (FCS) octets. This object provides a reasonable estimate of Ethernet utilization. Refer to RFC 1757 for information if more precision is desired.
Broadcast Packets	Total number of good packets received that were directed to the broadcast address (this does not include multicast packets).
Multicast Packets	Total number of octets of data received on the network, including those in bad packets but excluding frame check sequence (FCS) octets. This object provides a reasonable estimate of Ethernet utilization. Refer to RFC 1757 for information if more precision is desired.
Bad CRC	Total number of packets received that were less than 64 octets in length (excluding framing bits but including FCS octets) and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
Collisions	Best estimate of the total number of collisions on this segment. Refer to RFC 1757 for more information about this counter.
Receive Unicast Packets	Number of unicast packets delivered to a higher-layer protocol. Refer to RFC 1757 for more information about this counter.
Receive Non-Unicast Packets	Number of non-unicast packets (broadcast or multicast packets) delivered to a higher-layer protocol. Refer to RFC 1757 for more information about this counter.
Receive Octets	Total number of packets received that were between 64 and 1518 octets in length (including bad packets), excluding framing bits but including FCS octets.
Transmit Unicast Packets	Total number of packets that higher-level protocols requested be transmitted to a unicast address, including those that were discarded or not sent.

Table 19-14: Port Statistics Information *(Continued)*

Column	Description
Transmit Non-Unicast Packets	Total number of packets that higher-level protocols requested be transmitted to a non-unicast (broadcast or multicast packets) address, including those that were discarded or not sent.
Transmit Octets	Total number of octets transmitted out of the interface, including framing characters.
Transmit Collisions	Total number of packets that experienced a collision during transmission.
Fragment	Total number of packets that were fragmented during transmission.
Undersized	Total number of packets received that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed.
Oversized	Total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and were otherwise well formed.
Jabbers	Total number of packets received that were longer than 1518 octets (excluding framing bits but including FCS octets) and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).

Viewing Charts and Graphs

The Graphical User Interface (GUI) provides several graphs or charts displaying current information for ongoing operations.

To view a graph:

1. Do one of the following:
   - To view a specific graph, from the Navigation Tree, choose Monitor Security > Graphs, then choose the name of the desired graph.
   - To view the graphs specified as part of the Chart Dashboard, from the Dashboard drop-down, choose Chart Dashboard.

   The specified chart information displays.
2. To increase the size of the graph (and the amount of information displayed) drag a side or a corner of the graph with your mouse.
3. If the data in a specific graph spans a wide range, consider changing the graph from a Linear display format to a Logarithmic one.
4. If you have multiple graphs open simultaneously, you can select a particular graphic to view by clicking the Window Manager toolbar button.
5. Using the Time Resolution drop-down, you can change the time frame covered by a graph to display a smaller or larger amount of time. Information about the current time is always at the right.

The available graph types are listed in Table 19-15.

Table 19-15: Graph Types

Graph	Description
Dropped Packets	Provides an indication of the number of packets dropped by the different subsystems and checks. The graph displays information for the following packet types: • IP/ARP Bad Packets — Various types of poorly formed packets dropped. • Layer-2 Bridge Filtered — Packets filtered out due to Layer-2 forwarding rules. • SYN Flood Mitigation — Packets dropped for clients characterized as malicious. • SYN Flood/DDos Rejection Rate — Packets dropped for IP addresses that the device characterized as malicious. • Client Request Limiting — Packets dropped because an IP address is generating traffic • Connection Limiting — Packets dropped because of limitations you established from servers or user groups based on a group's allowed number of connections. • Firewall — Packets filtered out by the Firewall rules you have established. above its configured request limit. • Protocol Validation and Attack Signatures — Packets filtered due to violations of protocol rules found during deep packet inspections.

Table 19-15: Graph Types *(Continued)*

Graph	Description
SYN Flood Statistics	This graph provides an indication of the handling of malicious SYN flood packets. It displays the rate (in packets per second) at which the following types of packet drops occur: • Malicious SYN packet rate — Device is receiving and dropping packets from malicious clients. • SYN Flood/DDoS Rejection Rate — The rate at which packets are dropped because the device designates their IP addresses as malicious. • Client Proxy Fail Rate — The rate at which packets are dropped because the client did not complete a proxied handshake process. • Server Proxy Fail Rate — The rate at which packets are dropped because the server did not complete a proxied handshake process. (This can occur if, for example, a client attacks a network by trying to connect to a server that does not exist, or tries to connect to a service that is not run on that server.) • Proxy Resource Drop Packet Rate — Packets dropped due to major SYN flood attack that severely limits available resources.
IP Threat Levels	This graph displays the number of IP addresses that fall into each of the address threat levels assigned by the device: • Unknown IP addresses — Addresses whose threat level is currently undetermined. • Trusted IP addresses — Addresses that have completed enough connections to be considered trusted. • Suspicious IP addresses — Addresses that currently have enough incomplete connections that they are considered suspicious. The device proxies requests from these addresses. • Malicious IP addresses — Addresses that have enough incomplete connections that the device believes that they are generating a DDoS attack and, therefore, is blocking their requests. • DDoS Rejection IP addresses — Addresses that were affected by a major DDoS attack that severely limited available resources.
Connection Usage	Displays the types of connections that represent the traffic going through the device. Includes the number of TCP, UDP, and Other IP connections, as well as number of Aged Connections during the selected time period. • TCP Connections — Number of TCP connections that have incomplete 3-way handshakes (Syn-SynAck-Ack). • UDP Connections — Number of established UDP connections. • Other IP Connections — Number of active connections for non-IP protocols and for IP protocols other than TCP. • Aged Connections — Number of TCP connections for which a Fin or Reset has been seen, but the connection has not yet aged out. (Age time for a closed connection is approximately two seconds.)
Connection Setup	Displays the device's current rate of setup for various types of connections, including TCP, UDP, and Other IP.

Table 19-15: Graph Types *(Continued)*

Graph	Description
CPU Activity	Represents the device's CPU activity by percentage for the following activities: • Utilization — Percentage of utilization for all CPU activities. • Maintenance — Percentage for maintenance activities. • TCP Setup — CPU activity percentage for TCP connection setups. • UDP Setup — CPU activity percentage for UDP connection setups. • IP connection — CPU activity percentage for other IP connection setups.
Custom Chart	This graph enables you to select from a number of statistics over a user-specified period of time. For each statistic you select, the system provides default settings for the low, medium, and high thresholds for display. When you select a statistic, you can modify these thresholds as needed. You can select one or more statistics, even those with different scales and metrics, and the graph will superimpose them on the display. The Custom chart displays data for each selected statistic in a different color, with a legend displaying the statistic associated with each color at the bottom of the chart. The system retains your most recent Custom Chart settings, so the next time you access the Custom Chart feature from the same system, using the same account, a graph using your most recent Custom Chart settings will automatically display.

Chapter 20
SYN Flood and Connection Limiting Security

You can add security policies for your Corero Network Devices that protect your network's resources from overuse and abuse. Rate-based policies protect resources from overuse by legitimate users, but primarily results from abusive denial-of-service attackers. You can modify the default rate-based limits on a per-host-group basis.

This chapter describes how to specify rate limits for SYN Flood and Connection Limiting.

NOTE

Both SYN Flood limiting and Connection limiting are turned on by default. The factory configuration specifies very high values for connection limiting.

This chapter contains the following sections:

NOTE

Corero Network Devices provide client request limiting security for client host groups and server host groups. For more information, see Chapter 21, "Client Rate Limiting".

Connection Limiting Overview

Sometimes attacks come in the form of an overwhelming number of connection requests to one or more targets. If permitted, this number of connections will consume resources on the server until no new connections can be made, resulting in a denial of service to legitimate traffic. Yu limit connection requests to both a specified Host and a specified Host Group.

To implement a connection limit rate-based policy, you must:

1. Enable client rate limiting for the service to which you want the specified limits applied.
2. If needed, create client host groups to which you want to apply rate-based policies.
3. Specify one or more connection limit profiles. A connection limit profile contains a per-host connection limit, and a per-host-group connection limit. For more details, see Table 20-1.
4. Apply the desired connection limit profiles to the specified host groups.

NOTE

When initially configuring rate limits, start with one rate-based policy, then monitor system operation with the policy enabled. This allows you to view how normal system operation is affected by the limit. You can then tune your setting until the responses to legitimate traffic and malicious traffic are as desired.

For each profile, you specify the parameters listed in Table 20-1.

Table 20-1: Connection Limit Profile Settings

Parameter	Description
Profile Name	Specify a unique name for your profile. Note that profile names must be unique across all profile types on the device.
Maximum Group Connections	The maximum connection limit for all members (IP addresses) in a host group. This value is associated with the following rules: • tln-002002: TCP Active Connections From Client Group Exceed Specified Limit • tln-002004: TCP Active Connections To Server Group Exceed Specified Limit
Maximum Member Connections	The maximum connection limit for a specific member (IP address) in a host group. This value is associated with the following rules: • tln-002001: TCP Active Connections From Single Client In Group Exceed Specified Limit • tln-002004: TCP Active Connections To Single Server Exceed Specified Limit

NOTE

Connection Limiting and Application Rate Limiting rules are the only rate-based rules you can set to Allow. This enables you to view information about rule trigger events while still allowing traffic to pass.

SYN Flood Rate Limiting Overview

A SYN flood is a form of denial-of-service attack in which an attacker sends a succession of SYN requests to a target system in an attempt to consume enough server resources to make the system unresponsive to legitimate traffic.

WARNING

SYN Flood Mitigation parameters are difficult to set properly. If you feel your site's settings need to be modified, contact the Corero Customer Services Center.

Normally, when a client attempts to start a TCP connection to a server, the client and server exchange a series of messages. This is called the TCP three-way handshake, and is the foundation for every connection established using the TCP protocol.

1. The client requests a connection by sending a SYN (synchronize) message to the server.
2. The server acknowledges this request by sending SYN-ACK (synchronize acknowledgement) back to the client.
3. The client responds with an ACK (acknowledgment), and the connection is established.

A SYN flood attack works by sending SYN messages requesting a connection, but then not responding to the server's SYN-ACK reply with the expected ACK code. The malicious client can either simply not send the expected ACK, or by spoofing the source IP address in the SYN, cause the server to send the SYN-ACK to a falsified IP address. The falsified IP address will not respond, because it did not initiate the SYN.

The server will wait for the acknowledgement for some time, as simple network congestion could also be the cause of the missing ACK. But during an attack increasingly large numbers of half-open connections will consume resources on the server until no new connections can be made, resulting in a denial of service to legitimate traffic.

You can configure several SYN-based thresholds in order to detect malicious behavior on the part of a source IP address. It also allows you to specify whether or not the device will proxy requests from unknown IP addresses. These settings are grouped into a SYN Flood profile, which you can select for use with a particular host group.

NOTE

When initially configuring rate limits, start with one rate-based policy, then monitor system operation with the policy enabled. This allows you to view how normal system operation is affected by the limit. You can then tune your setting until the responses to legitimate traffic and malicious traffic are as desired.

To implement a SYN Flood rate-based policy, you must:

1. If needed, create client host groups to which you want to apply rate-based policies.
2. Specify the SYN flood parameters you want to use in one or more SYN Flood profiles. These parameters are described in Table 20-2.
3. Apply the desired SYN Flood profiles to the specified host group.

When specifying SYN Flood protection profiles, you provide the information listed in Table 20-2.

Table 20-2: SYN Flood Protection Profile Settings

Setting	Description
Name	The profile name.

Table 20-2: SYN Flood Protection Profile Settings *(Continued)*

Setting	Description
Trusted Threshold	The number of successful and well-executed connections required to establish this IP address as a trusted IP address. Requests from trusted clients are forwarded to the destination address.
Suspicious Threshold	The source IP address has reached a user specified threshold (Suspicious Threshold) due to the number of open SYNs. The device will proxy all connection requests to any device from that source IP address. As a proxy, the device assumes the role of the destination device and responds to the connection set up request. If the source address successfully completes the connection setup, then the device creates the connection with the real destination device and begins to forward traffic to that connection. If this limit is exceeded, the following rule is triggered and the action defined for this rule is followed: tln-001018 'Connection From Client That Fails Proxy Handshake'
Suspicious Exit Threshold	The device maintains a count of the number of incomplete TCP connections destined to a server group and allows the user to set a value for this threshold (known as the Connection Threshold). When a source Client Group is in the Unknown or Trusted states, and the destination's Server Connection Threshold has not been reached, then SYN packets received from this source to the given destination are passed directly on to the server group (subject to other security policy checks). If the source has been deemed Suspicious or the destination's Server Group's Connection Threshold has been reached, the device will not directly pass the SYN packet, but will instead act as a proxy on the server's behalf. The device ensures that the source creates a legitimate TCP connection before passing this connection to the server; thereby, protecting the server from bogus open SYNs. The suspicious exit threshold is the number of completed SYNs that a client must acquire in order to exit proxy mode. **Note:** By default, this value should be set significantly lower than the Suspicious Threshold to ensure a client does not vacillate between proxy and no proxy. If this limit is exceeded, the following rule is triggered and the action as defined for this rule is followed: tln-001019 'Connection To Server That Fails Proxied Handshake'
Malicious Threshold	The source IP address has reached an even higher user specified threshold (Malicious Threshold) because of the number of open SYNs. The device blocks any connection request from this address until one of two things happens: The address stops initiating requests for a given period of time called the timeout period. The number of new, open SYNs drops enough over time (due to a decay mechanism) to allow the total open SYNs to decay to the suspicious level. If this limit is exceeded, the following rule is triggered and the action as defined for this rule is followed: tln-001007 'Connection From Malicious Source IP Address'

Table 20-2: SYN Flood Protection Profile Settings *(Continued)*

Setting	Description
Proxy Requests from Unknown IP	Used for TCP connection traffic patterns for this address that are not yet identified. For IP addresses with an unknown threat level, you can configure the device to either: • Forward connection requests from this address to the destination device. • Proxy connection requests until the IP address becomes trusted. Each connection request, until completed, is an instance of an open SYN and could cause the total open SYNs from all IP addresses to pass the acceptable threshold for that server. In this case, the device proxies the connection request regardless of the forwarding setting for unknown IP addresses. During a DDoS attack, the device takes additional steps to protect your resources. One of these steps is to initially deny a connection request by a new (potentially spoofed and dangerous) client during the DDoS attack. Once the device verifies that the client is a good client, the unit processes the client's connection requests If this limit is exceeded, the following rule is triggered and the action as defined for this rule is followed: tln-001020 'Connection From New Client During DDoS Attack'

Enabling SYN Flood, Connection, and Client Request Limiting

You use slightly different process to enable SYN Flood Limiting, Connection Limiting, and Client Request limiting. Table 20-3 describes the differences between these processes.

Table 20-3: Enabling SYN Flood, Connection, and Client Request Limiting

To enable...	Do this...	And consider the following...
SYN Flood Limiting	On the multi-tabbed Configure Security Policies dialog box: 1. Click the Rate-Based Policies tab. 2. Select the desired Client host group and specify a SYN Flood Limit value as described in Configuring a Connection or SYN Flood Rate Limit (page 20-8). 3. Select the desired Server host group and specify a SYN Flood Limit value as described in Configuring a Connection or SYN Flood Rate Limit (page 20-8). 4. Enable the relevant rules for the desired rule set.	If the SYN Flood Limit for *either* the client host group or the server host group is set to No Limit, SYN Flood Limiting is disabled for all members of that host group. For detailed instructions, see Configuring a Connection or SYN Flood Rate Limit (page 20-8).
Connection Limiting	On the multi-tabbed Configure Security Policies dialog box: 1. Click the Rate-Based Policies tab. 2. Select the desired Client host group and specify a Connection Limit value as described in Configuring a Connection or SYN Flood Rate Limit (page 20-8). 3. Select the desired Server host group and specify a Connection Limit value as described in Configuring a Connection or SYN Flood Rate Limit (page 20-8). 4. Enable the relevant rules for the desired rule set.	If the Connection Limit for either the client host group or the server host group is set to No Limit, Connection Limiting is disabled for all members of that host group. For detailed instructions, see Configuring a Connection or SYN Flood Rate Limit (page 20-8).
Client Rate Limiting for Client host groups.	On the multi-tabbed Configure Security Policies dialog box: 1. Click the Rate-Based Policies tab. 2. Specify a Client Request Limit value for the desired Client host group. 3. Click the Services tab. 4. Edit the desired service. 5. Click the Advanced button on the Edit Service dialog box. 6. Select the Enabled radio button to enable Request Limiting. 7. Enable the relevant rules for the desired rule set.	If the Client Request Limit for the client host group set to No Limit, or if Request Limiting is disabled for a particular service, client rate limiting is disabled. For detailed instructions, see Chapter 21, "Client Rate Limiting".

Table 20-3: Enabling SYN Flood, Connection, and Client Request Limiting *(Continued)*

To enable...	Do this...	And consider the following...
Client Rate Limiting for Server Host Groups	On the multi-tabbed Configure Security Policies dialog box: 1. Click the Rate-Based Policies tab. 2. Specify a Client Request Limit value for the desired Client host group. 3. Click the Services tab. 4. Edit the desired service. 5. Click the Advanced button on the Edit Service dialog box. 6. Select the Enabled radio button to enable Request Limiting. 7. Enable the relevant rules for the desired rule set. 8. Modify rate limiting profiles for the Request/Response Behavior (RRB) rules.	If the Client Request Limit for the client host group set to No Limit, or if Request Limiting is disabled for a particular service, client rate limiting is disabled. For detailed instructions, see Chapter 22, “Advanced Client Rate Limiting”.

Configuring a Connection or SYN Flood Rate Limit

When configuring a Connection or SYN Flood rate limit, consider the following:

- When setting the parameters for SYN Flood and Connection Limiting features for edge devices, try to consider the requirements of all the devices in the set of devices being protected.
- For protection of internal servers, be aware that the set of services being handled by the servers, and, therefore, the traffic requirements, may be quite different from a set of servers "on the edge".
- For both Corero Network Devices protecting edge devices (servers with external clients) and Corero Network Devices protecting internal servers, start with the default settings and monitor the device for blocked traffic. If the current settings are blocking valid traffic, increase these settings.

There are four steps to configuring a Connection or SYN Flood rate limit:

- Step 1: Preparing Host Groups (page 20-8)
- Step 2: Enabling Request Limiting for a Service (page 20-11)
- Step 4: Creating a Rate Based Policy for a Specific Host Group (page 20-13)
- Step 4: Creating a Rate Based Policy for a Specific Host Group (page 20-13)
- Step 5: Enabling the Relevant Rules for the Desired Rule Set (page 20-16)

Step 1: Preparing Host Groups

When you configure Connection and SYN FLood rate limiting, you specify limits for a specific host group. So the first step in configuring client rate limiting is preparing the host groups to which you will apply limits.

NOTE

For more information about host groups, see Chapter 13, ‘‘Managing Host Groups".

To prepare host groups:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.
2. Click the Host Groups tab (Figure 20-1). You may find that you can add client IP addresses to existing host groups (such as the Other Host Group, or the Suspicious Host Group). Alternatively, you may want to create new host groups, such as one for Public Web Servers.
3. To add a new host group, in the Host Groups area, click Add, then specify a name for the new host group.
4. To add or edit the IP address membership of a host group (Figure 20-1):
   a. Select the desired host group in the Host Groups area.
   b. To add more IP addresses to an existing host group, in the Host Group Membership area, click Add. The Add IP Address Range dialog box displays. Optionally, you can specify a name for this address group.
   You can add IP addresses in four ways:
   - As an IP address/Prefix (for example 192.0.8.31/24, where the first 24 bits would have to match exactly.)
   - As an IP address/Mask (for example 192.0.8.31/255.255.255.0).
   - As a range specified by a first and last IP address (for example 192.0.128.0-192.01.128.255).
   - As a single IP address (for example 192.0.8.31).

c. To edit the addresses in an existing host group, in the Host Group Membership area, click an IP Address range, then click Edit. The Edit IP Address Range dialog box displays, enabling you to modify the (optional) IP Address Range name and the associated host group.

d. Whether you have added or edited addresses, you can modify Spoof Check settings. Spoof Checks are used to identify attacks where hosts modify the IP address to imitate an internal (or external) IP address. Instruct the Corero Network Device whether or not to perform spoof checks, and if they will be performed, specify the type of port (internal or external) from which traffic with this IP address will be permitted.

e. Whether you have added or edited addresses, you can also modify advanced settings. To do so, on the Add or Edit IP Address Range dialog box, click the Advanced button. This enables you to specify whether to identify subnet or broadcast addresses associated with the IP address range you specified.

5. To Delete an IP Address Range from the Host Group, select the IP Address Range, then click Delete.
6. When you have finished specifying host group settings in the management application for a Corero Network Device, click Done.
7. Save your changes by clicking the Save Configuration toolbar button.

Figure 20-1: Modifying Host Group Membership

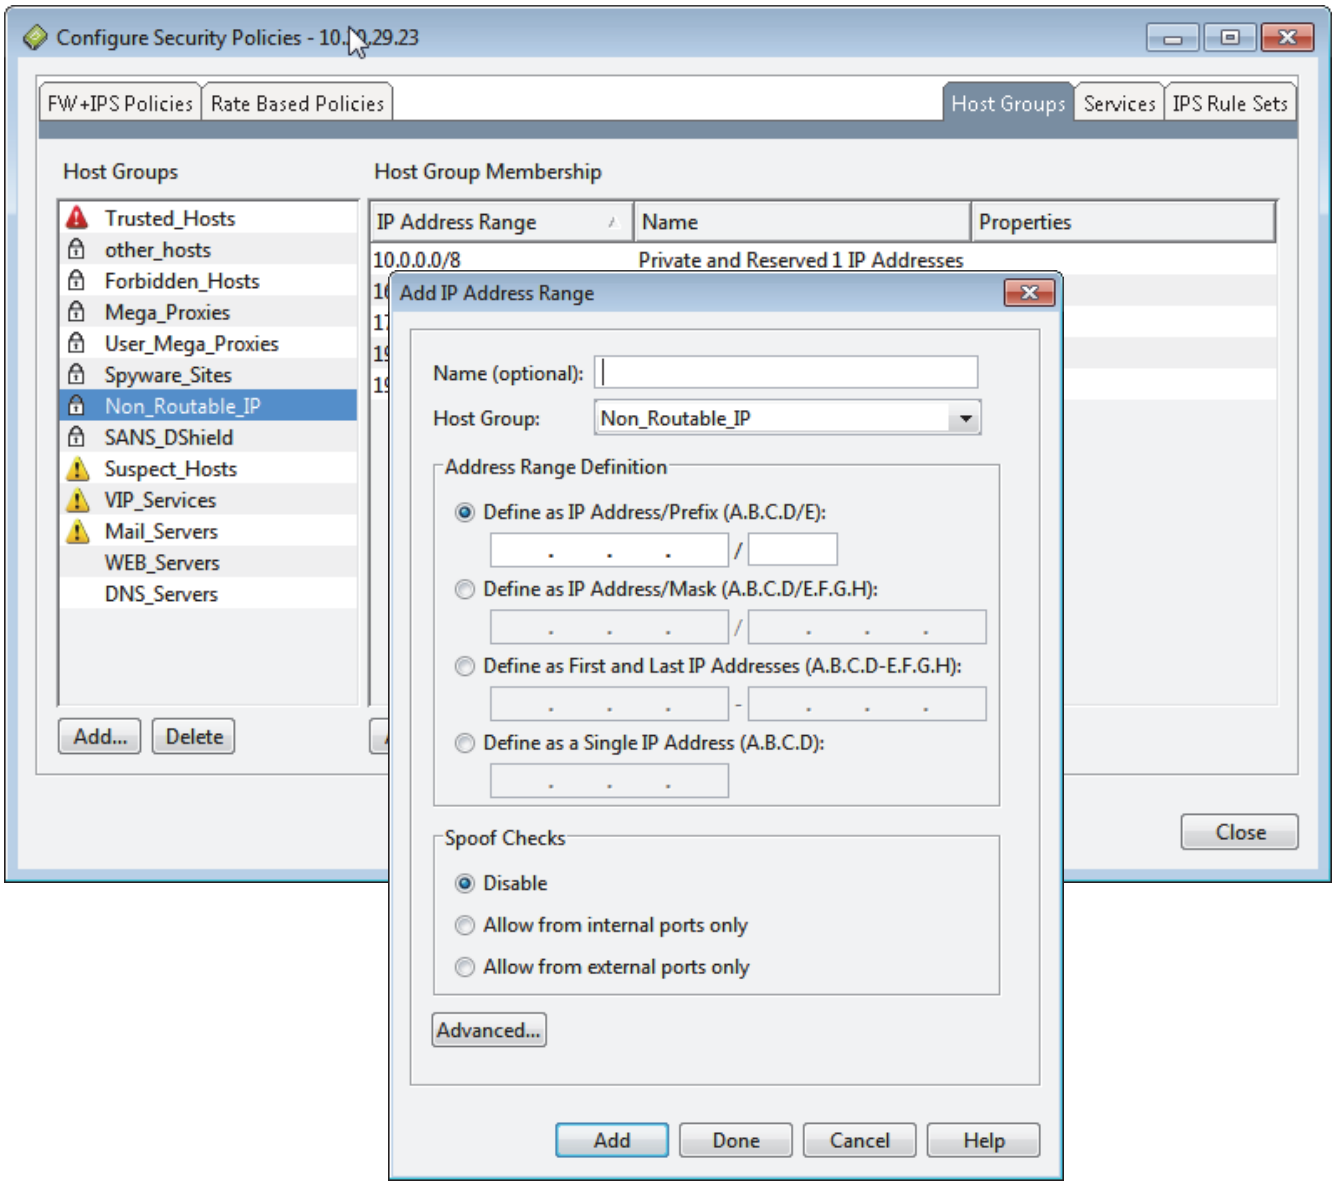

Step 2: Enabling Request Limiting for a Service

The request limiting feature is only applied to those services that have request limiting enabled. In the case of client rate limiting, you would typically enable request limiting on both the DNS and HTTP services.

NOTE

If you are performing initial configuration for client request limiting, Corero recommends that you create a service associated with just one or a few servers, then enable rate limiting on that service. You can then watch the progress of normal traffic and ensure that it passes properly.

NOTE

For more information about services, see Chapter 14, ‘‘Managing Services”.

To enable request limiting for a service:

enable request limiting for a service:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.
2. Click the Services tab (Figure 20-2). You can view or modify services from this page.

   NOTE

   If you do not see the service for which you want to enable request limiting, you will need to add it. For instructions on adding a new service, see Chapter 14, ‘‘Managing Services”.

3. Select the service to which you want to apply client rate limiting, then click Edit. The Edit Service dialog box displays (Figure 20-2).
4. Click Advanced. The Advanced Service Settings dialog box displays.
5. For request limiting, click the Enabled radio button for Request Limiting, then click OK.
6. When you have finished specifying service settings in the Corero Network Device management application, click Done at the bottom of the Configure Security Policies dialog box. If you want to save the service and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
7. Save your changes by clicking the Save Configuration toolbar button.

Figure 20-2: Modifying Services

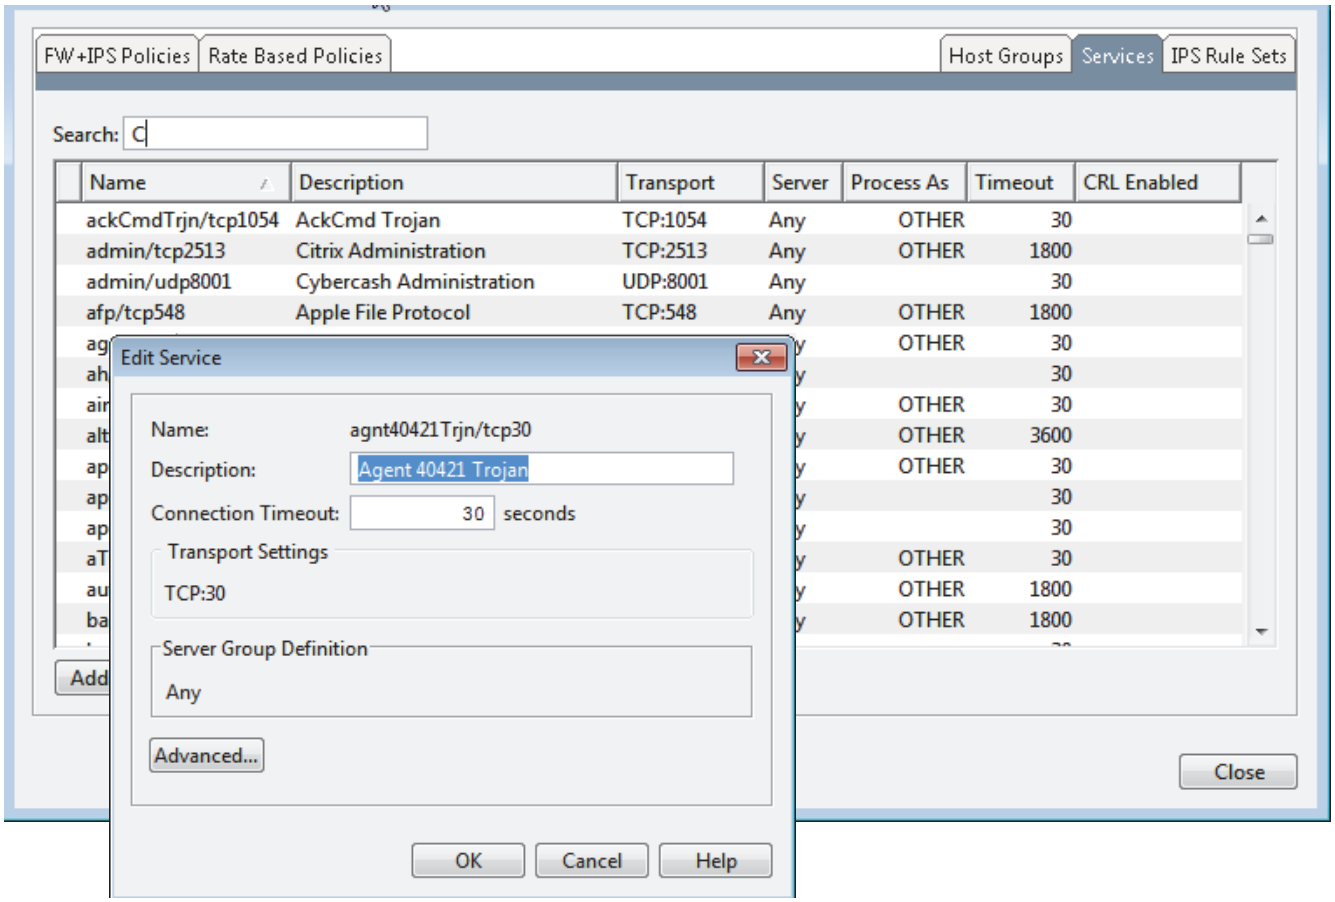

Step 4: Creating a Rate Based Policy for a Specific Host Group

You can add or edit Connection or SYN Flood limiting for a particular traffic type. To create a Connection or SYN Flood limiting rate-based policy for a specific host group:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.
2. Click the Rate Based Policies tab.
3. Click either the Clients or Servers category, depending on the host group to which you want the limits applied.
4. Select the host group whose clients you want to rate limit, then click Edit
   - If you selected Clients, the Edit Client Limits dialog box displays (Figure 20-3).
   - If you selected Servers, the Edit Server Limits dialog box displays. The SYN Flood and Connection Limit options are on the General tab.

Figure 20-3: Rate Based Policies Tab

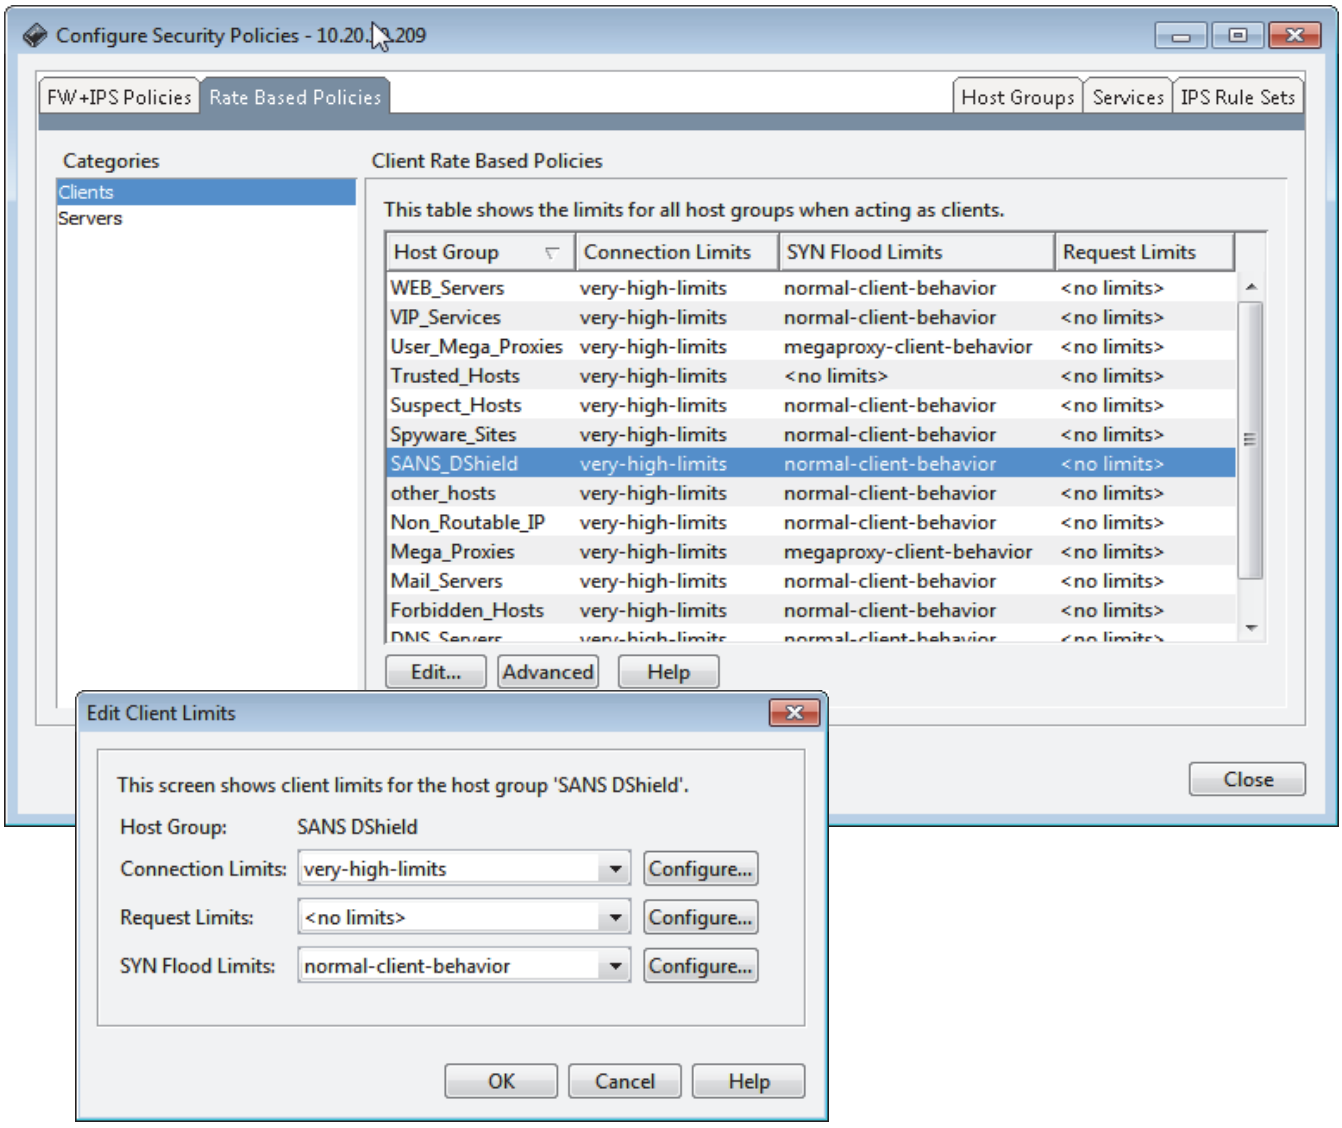

5. To modify Connection or SYN Flood rate limits for this host group.
   - Use the drop-down list to select the desired profile for SYN Flood Limits.
   - Use the drop-down list to select the desired profile for Connection Limits.

   NOTE

   You can view profile information by clocking the Configure button for either Connection Limits or SYN Flood Limits.

6. If the Connection Limit profile or settings you want are not available in the Connection Limits drop-down list:
   a. Click the Configure button adjacent to the Connection Limits drop-down.
   b. If you want to add a new profile, click Add. To modify an existing profile, select the profile and click Edit.

c. In the configuration dialog box that displays, you can add to or edit the available Connection limits. Information on these parameters is listed in Table 20-1.

d. When finished, click Close. Now you can select it in the drop-down list.

7. If the SYN Flood Limit profile or settings you want are not available in the SYN Flood Limits drop-down list:

   a. Click the Configure button adjacent to the SYN Flood Limits drop-down.

   b. If you want to add a new profile, click Add. To modify an existing profile, select the profile and click Edit.

   c. In the configuration dialog box that displays, you can add to or edit the available SYN Flood limits. Information on these parameters is listed in Table 20-2.

   d. When finished, click Close. Now you can select it in the drop-down list.

8. The Rate Limit dialog box also enables you to manually assign SYN Flood threat levels to individual IP addresses. To do this:

   a. Click Advanced. The Advanced SYN Flood Client Settings dialog box displays.

   b. To specify a new IP address and assign a threat level to it, click Add. Enter the IP address and select the desired threat level. Available threat levels include Unknown, Trusted, Suspicious, and Malicious.

   c. To modify the threat level associated with an existing IP address, select the address, then click Edit. You can now modify the threat level associated with that address.

   d. To remove the threat level from an IP address altogether, select the IP address, then click Delete.

9. When finished, click OK.

10. When you have finished specifying host group settings in the management application for a Corero Network Device, click Done.

11. Save your changes by clicking the Save Configuration toolbar button.

Step 5: Enabling the Relevant Rules for the Desired Rule Set

Once you have specified parameters for Connection and SYN Flood rate limiting, you need to ensure the associated rules are enabled, and that the settings have been modified to meet your current requirements.

To modify rule settings:

d limiting rate-based policy for a specific host group:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.

2. Click the IPS Rule Sets tab. Select the desired rule set.

   NOTE

   For more information on managing rule sets, see Chapter 15, "Managing Rules and Rule Sets".

3. To view the rules associated with SYN flood, connection, and client limiting that are not associated with HTTP or DNS on the management application for a Corero Network Device, choose Configure Security > Advanced Security Config > Rate Based Rules from the Navigation Tree.
4. To modify the settings for a rule, select the rule, then click Edit. The Edit Rule Settings dialog box displays (Figure 20-4).

Figure 20-4: Modifying Rule Settings

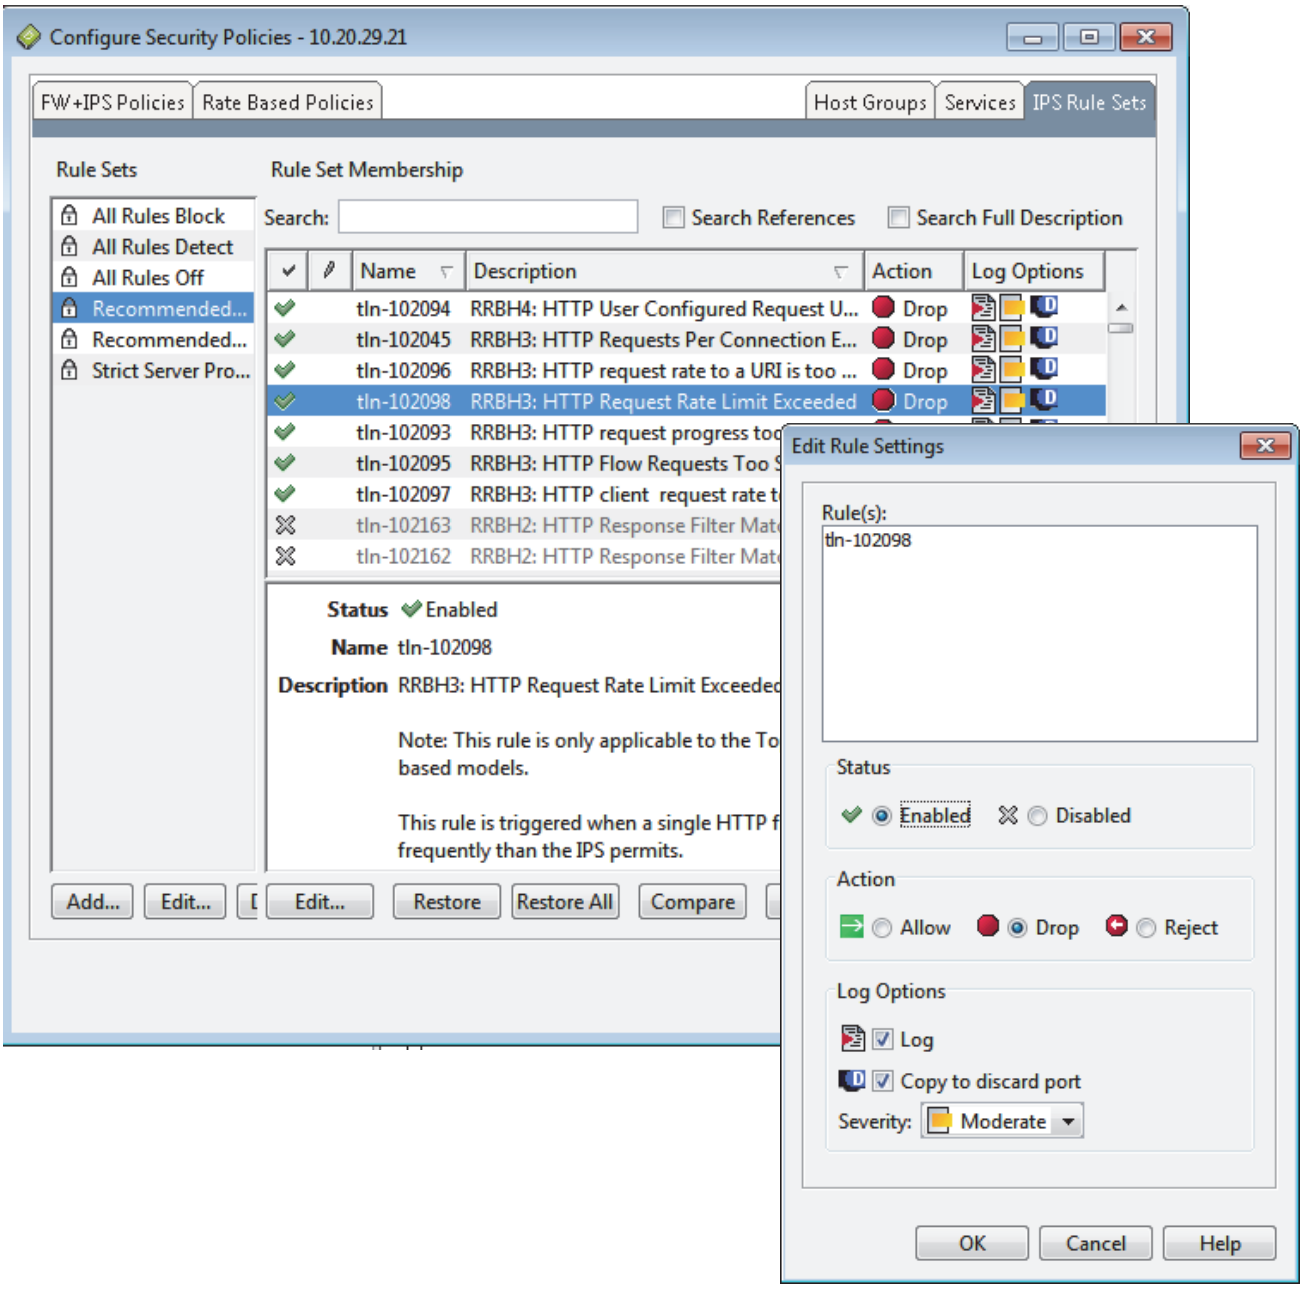

5. You can enable or disable individual rules.

NOTE

This setting applies only to this rule and it overrides the rule set settings established using the Edit a Rule Set window.

If you disable a rule, packets that would normally trigger the rule are passed, and no logging is performed.

If you enable a rule, you can set its Action. Possible actions are:

- Allow— Pass the traffic.

- Drop— Passively block the traffic by dropping the traffic with no indication to the offending client.
- Reject— Actively block the traffic by dropping the traffic and taking specific action, for example, by resetting the connection.

Specify the Log options for this rule as follows:

- Log - Send information to a log file based on its severity rating.
- Copy to Discard Port - Copy the associated traffic to the Discard port based on its severity rating.

NOTE

If you have configured a Discard Port, you can log all traffic that triggers a rule, even if the Action for that rule is set to Allow. This feature enables you to record usage information about applications you allow as well as those you disallow.

- Severity - The log severity determines whether packets for this rule are processed by the device's Event Logging System. To be processed, the priority you set for this rule must be equal to, or above, the overall priority set for event messages, as described in Logging Options (page 15-4). Severity levels include:
  - Low (green)
  - Moderate (yellow)
  - Critical (red)

6. When you have finished specifying rule settings in the management application for a Corero Network Device, click OK.
7. Save your changes by clicking the Save Configuration toolbar button.

   When the rule appears in the list of rules, a pencil icon displays indicating that this rule has been modified from its default settings.

Checking the Number of Open SYNs and Current Connections for an IP Address

You can use the IP Address Query feature to view current counter and client request credit information for a particular IP address. To query an IP address:

1. On the management application for a Corero Network Device, do one of the following:
   a. From the Navigation Tree, choose Monitor Security > Query IP Address.
   b. From the Security Event Viewer, with the Active Mode check box cleared (deselected), select an event then click Query Client or Query Server.

   The IP Address Query dialog box displays.
2. Enter the IP address of the host about which you want to view additional information.
3. Select the Corero Network Device you want to query. If you want to query all devices, choose <All>.
4. On the management application for a Corero Network Device, click Query.
5. If the IP address is known, from the IP Address Query dialog box, you can:
   - View the current number of Client Open SYNs, Client Completed SYNs, and Server Open SYNs.
   - Clear (zero) the all SYN Counters.
   - Reset other system counters.

   NOTE

   For additional information on the IP Address Query dialog box, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).

Chapter 21
Client Rate Limiting

The Corero Network Device management application enables you to add a security policy that protects your network's resources. Rate-based policies protect resources from overuse by legitimate users, as well as abusive denial-of-service attackers. You can modify the default rate-based limits on a per-host-group basis.

In order to perform Client Rate Limiting, the Corero Network Device assesses the packet source and its Client Host Group, the packet destination and its Server Host Group, the associated service, and current client credit information for the IP source address. The device makes client rate limiting decisions based on this information.

NOTE

All client rate limiting rules are turned off by default.

Note that client rate limits are specified per client group, but the costs associated with various types of traffic are specified per server group.

This chapter contains the following sections:

NOTE

In order to prevent DDoS and similar attacks, you can specify HTTP and DNS limits. For more information, see Chapter 22, "Advanced Client Rate Limiting".

Client Rate Limiting Overview

Your Corero Network Devices routinely capture detailed information about network traffic that flows through it. For example, it stores a list of the most recently seen IP addresses. For each address, it stores connection counts, including attempted, accepted, and completed connections. In addition, in order to stop certain rate-based types of attacks, The device stores a value for the current number of credits that reflect the balance of behavior of that IP address, known as the Client Request Credits.

This feature is designed to allow each and every client to have individual limits applied, ranging from simple packet or request per second, to complex protocol based request / response violations, thus preventing complex rate based attacks from succeeding.

A client is defined as the initiator of a connection or flow.

Each IP address is given a number of Client Request Credits every 30 seconds. The number of credits is configured by profile for each client group.

The device decreases the number of credits based on client behavior as follows:

- For each packet that corresponds to a CRL enabled service, configured globally in advanced security config for packet types.
- For Protocol based Request / Response behavior violation, configured by profile for each server group, and individually enabled or disabled by configuring RRBxx rules in IPS policies.

When a client IP address's credits are below zero, packets from that client are dropped, providing protection for infrastructure and services from complex rate-based attacks.

There are also some advanced settings:

- The maximum credit value per client is capped, this cap setting is known as the Burst Rate.
- When a negative credit value (known as the overdraft limit) is exceeded, all packets from the client will be discarded. This feature acts as a virtual shun of the address.

Client Rate Limiting Configuration Elements

You must enable rate limiting for the services whose traffic you wish to limit.

When you specify client rate limiting for a client host group, you specify the number of client credits you will permit per minute.

There are several basic steps to implement a client rate based policy,:

1. If needed, create client host groups to which you want to apply rate-based policies.
2. Enable client rate limiting for the service(s) to which you want the specified limits applied.
3. Ensure the relevant rules are enabled.
4. Add or edit limit values, called profiles, that specify the client limits you want to be able to select for a client host group.
5. Apply the desired limit profiles to the specified host groups.
6. Configure the Global Client per packet costs, in Advanced Security configuration.

NOTE

Customizing global client per-packet costs is only available on more recent E-Series IPS Units. For previous models, the per-packet cost is fixed at 1.

NOTE

When initially configuring client rate limits, it may be easier to start with one client rate-based policy, then monitor system operation with the policy enabled. Once you have seen how that policy affects normal traffic, you can tune it for optimal results.

Client Rate Limiting Calculation Example

The number of Client Request Credits available for a particular client IP address dynamically changes based on traffic patterns. Traffic is only passed when there is a positive number of credits available for that IP address.

When rate limit credits are calculated, there are several important factors:

- The current value of the number of Client Request Credits. When a traffic flow begins, the number of credits available is equal to the rate limit value. Thereafter, the current number of Client Request Credits is based on the traffic behavior in the flow.
- The rate at which new Client Requests are made. Every half minute, the device replenishes half of the rate limit value.
- The incoming traffic rate, and the configured cost for the packet types.
- Any RRBxx rule-related per-packet costs for Protocol Rate policy violations
- The burst rate, an advanced setting that controls the cap on the maximum number of Client Request Credits allowed to accumulate.
- The overdraft point, an advanced setting that specifies the negative credit limit. Once the number of Client Request Credits reaches this value, all packets from that client are dropped, regardless of the configured service.

When the credit limit goes below zero, packets from the specified IP address sent to the specified service will be dropped. When the credit limit reaches the (negative) overdraft level, all packets from that IP address will be dropped.

Imagine a system configuration in which the following occurs:

- The client request limit is set to 1000 packets per minute for a particular IP address.
- The inbound traffic rate for a particular traffic type is 100 packets per second.

Table 21-1 Shows how the credits available for that traffic would change over a period of time. After 10 seconds, traffic will be blocked. Note that no additional credits are added (burst), because that would happen halfway through the one minute sampling period.

NOTE

To view the current number of credits available for a given IP address, query that IP address as described in Reset (Clear) SYN Flood and Connection Counters (page 19-31)

Table 21-1: Example: Credit Rate Limit Values Over Time

Elapsed Time (seconds)	Credits Deducted	Credits Added	Credits Available
0	0	1000	1000
1	100	0	900
2	100	0	800
3	100	0	700
4	100	0	600
5	100	0	500

Table 21-1: Example: Credit Rate Limit Values Over Time *(Continued)*

Elapsed Time (seconds)	Credits Deducted	Credits Added	Credits Available
6	100	0	400
7	100	0	300
8	100	0	200
9	100	0	100
10	100	0	0

Configuring a Client Rate Limit

Requests are made from clients and received by servers.

The following sections describe the procedure used to configure a client rate limit:

Step 1: Preparing Host Groups

When you configure client host group-based client rate limiting, you specify limits for a specific client host group. The first step in configuring client rate limiting is preparing the host groups for which you will define limits.

NOTE

For more information about host groups, see Chapter 13, ‘‘Managing Host Groups”.

To prepare host groups:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.
2. Click the Host Groups tab. You may find that you can add client IP addresses to existing host groups (such as the Other Host Group, or the Suspicious Host Group). Alternatively, you may want to create new host groups, such as one for Public Web Servers.
3. To add a new host group, in the Host Groups area, click Add, then specify a name for the new host group.
4. To add or edit the IP address membership of a host group (Figure 21-1).
   a. Select the desired host group in the Host Groups area.
   b. To add more IP addresses to an existing host group, in the Host Group Membership area, click Add. The Add IP Address Range dialog box displays. Optionally, you can specify a name for this address group.
   You can add IP addresses in four ways:
   - As an IP address/Prefix (for example 192.0.8.31/24, where the first 24 bits would have to match exactly.)
   - As an IP address/Mask (for example 192.0.8.31/255.255.255.0).
   - As a range specified by a first and last IP address (for example 192.0.128.0-192.01.128.255).
   - As a single IP address (for example 192.0.8.31).
   c. To edit the addresses in an existing host group, in the Host Group Membership area, click an IP Address range, then click Edit. The Edit IP Address Range dialog box displays, enabling you to modify the (optional) IP Address Range name and the associated host group.
   d. Whether you have added or edited addresses, you can modify Spoof Check settings. Spoof Checks are used to identify attacks where hosts modify the IP address to imitate an internal (or external) IP address. Instruct the Corero Network Device whether or not to perform spoof checks, and if they will be performed, specify the type of port (internal or external) from which traffic with this IP address will be permitted.

e. Whether you have added or edited addresses, you can also modify advanced settings. To do so, on the Add or Edit IP Address Range dialog box, click the Advanced button. This enables you to specify whether to identify subnet or broadcast addresses associated with the IP address range you specified.

5. To Delete an IP Address Range from the Host Group, select the IP Address Range, then click Delete.
6. When you have finished specifying host group settings in the Corero Network Device management application, click Done at the bottom of the Configure Security Policies dialog box. If you want to save the host group and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
7. Save your changes by clicking the Save Configuration toolbar button.

Figure 21-1: Modifying Host Group Membership

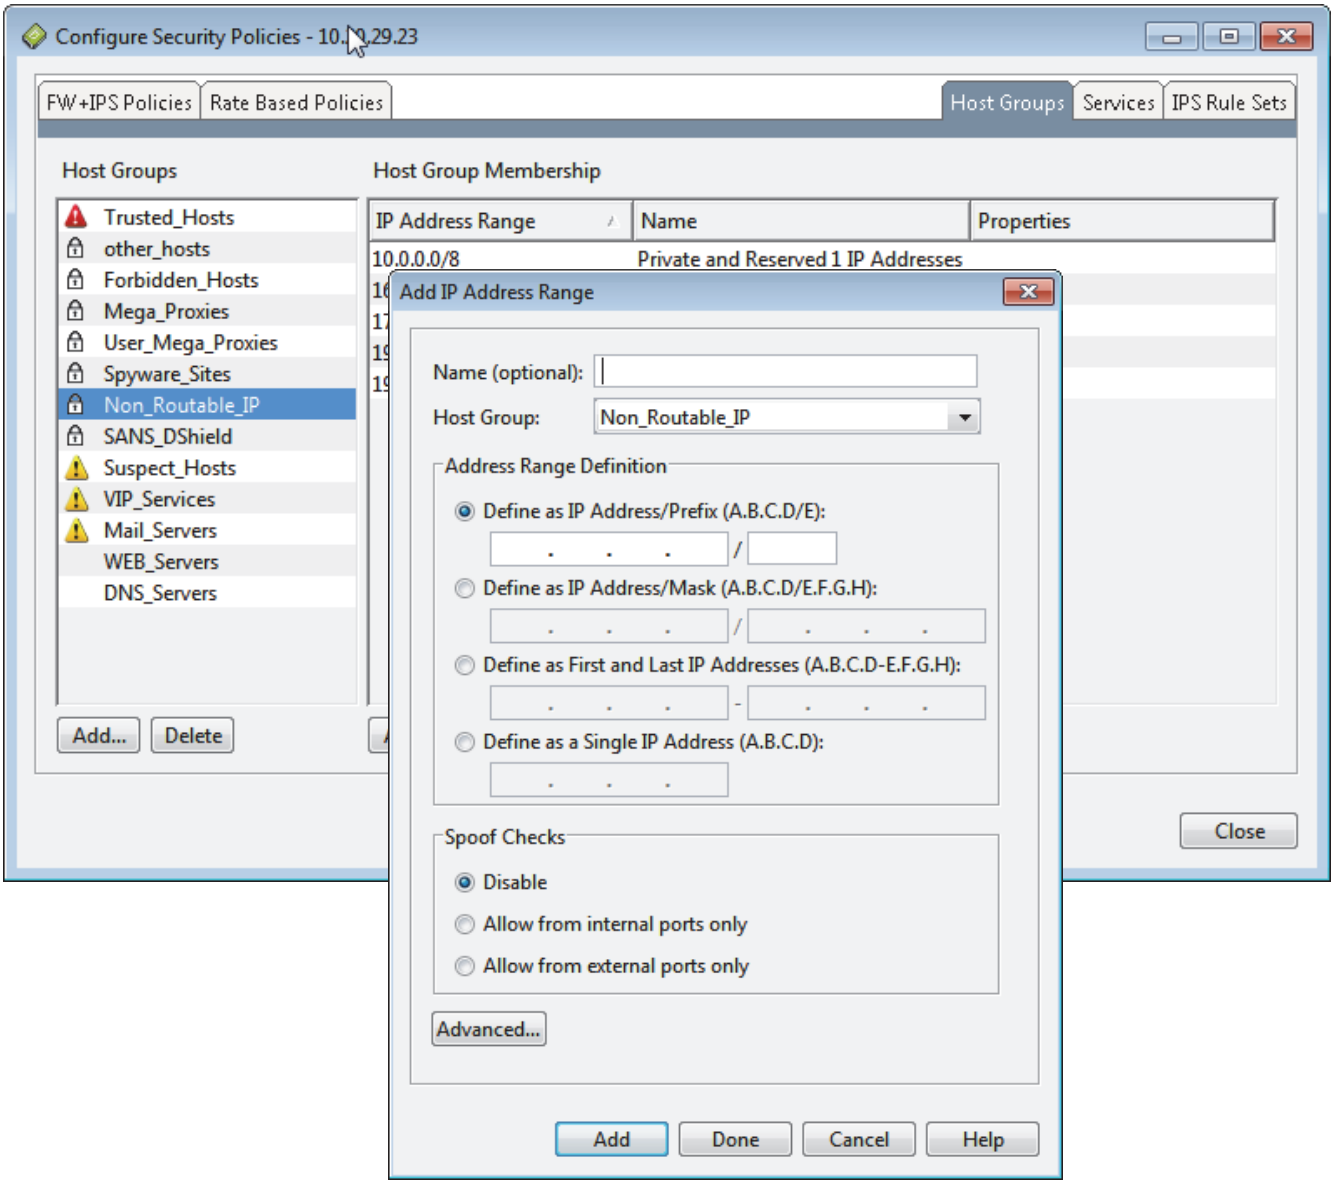

Step 2: Enabling Request Limiting for a Service

The client request limiting feature is only applied to those services that have request limiting enabled. Note that client rate limiting is turned off by default.

CAUTION

It is important to note that, if you specify more than one service for a policy, the limit will be applied to the total value of all specified services' requests, rather than to each service individually.

For more information about services, see Chapter 14, "Managing Services".

To enable request limiting for a service:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.
2. Click the Services tab. You can view or modify services from this page.

   NOTE

   If you do not see the service for which you want to enable request limiting, you will need to add it. For instructions on adding a new service, see Chapter 14, "Managing Services".
3. Select the service to which you want to apply client rate limiting, then click Edit. The Edit Service dialog box displays (Figure 21-2).
4. Click Advanced. The Advanced Service Settings dialog box displays.
5. For request limiting, click the Enabled radio button for Request Limiting, then click OK.
6. When you have finished specifying service settings in the Corero Network Device management application, click Done at the bottom of the Configure Security Policies dialog box. If you want to save the service and create another one, click Add. Alternatively, if you do not want the changes applied, click Cancel.
7. Save your changes by clicking the Save Configuration toolbar button.

Figure 21-2: Modifying Services

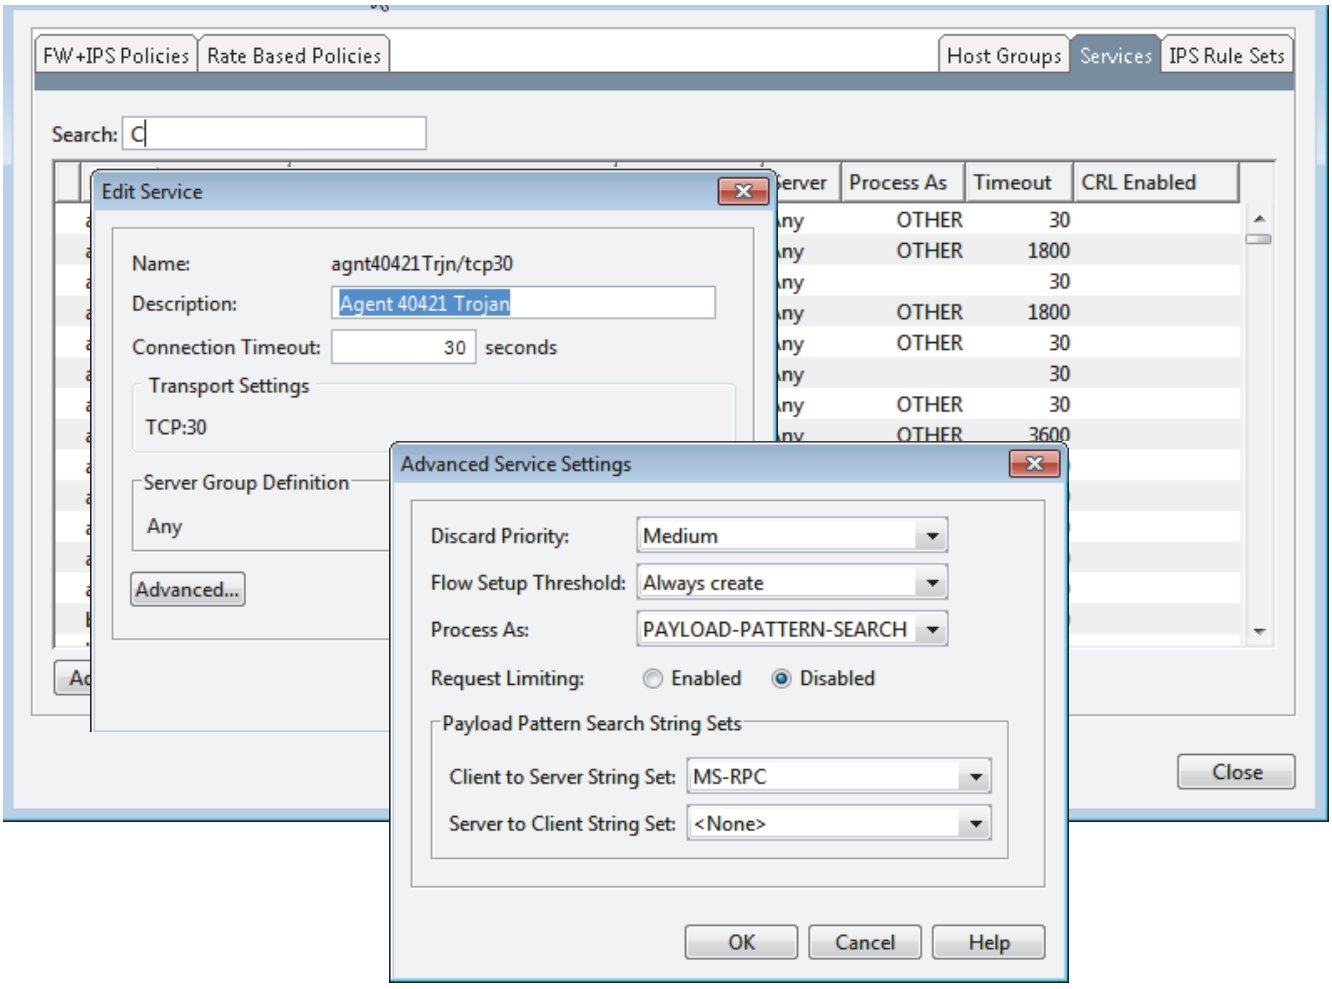

Step 3: Creating a Rate Based Client Limit Policy for a Client Host Group

You can manage traffic flowing into your network by specifying a client rate limit for a client host group.

To create a rate-based policy for a specific host group:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.
2. Click the Rate Based Policies tab.
3. Click Clients, then select the host group whose client requests you want to rate limit
4. Click Edit. The Edit Client Limits dialog box displays (Figure 21-3).
5. Select a Request Limit value from the drop-down list.

   If the limit you desire is not available, you can create or modify a limit to suit your needs. To do so:

   a. Click the Configure button adjacent to the Client Limits drop-down.
   b. If you want to add a new limit, click Add. To modify an existing limit, select the limit and click Edit.
   c. In the configuration dialog box that displays, you can add to or edit the request limit.
   d. When you are finished, click Close. Then you can select the limit from the drop-down list.
6. If desired, specify the SYN Flood and Connection Limits. For detailed information about specifying SYN Flood and Connection Limit profile settings on the General Limits tab, see Chapter 20, ‘‘SYN Flood and Connection Limiting Security.
7. When finished, click OK.
8. When you have finished specifying host group settings in the management application for a Corero Network Device, click Done.
9. Save your changes by clicking the Save Configuration toolbar button.

Figure 21-3: Rate Based Policies Tab

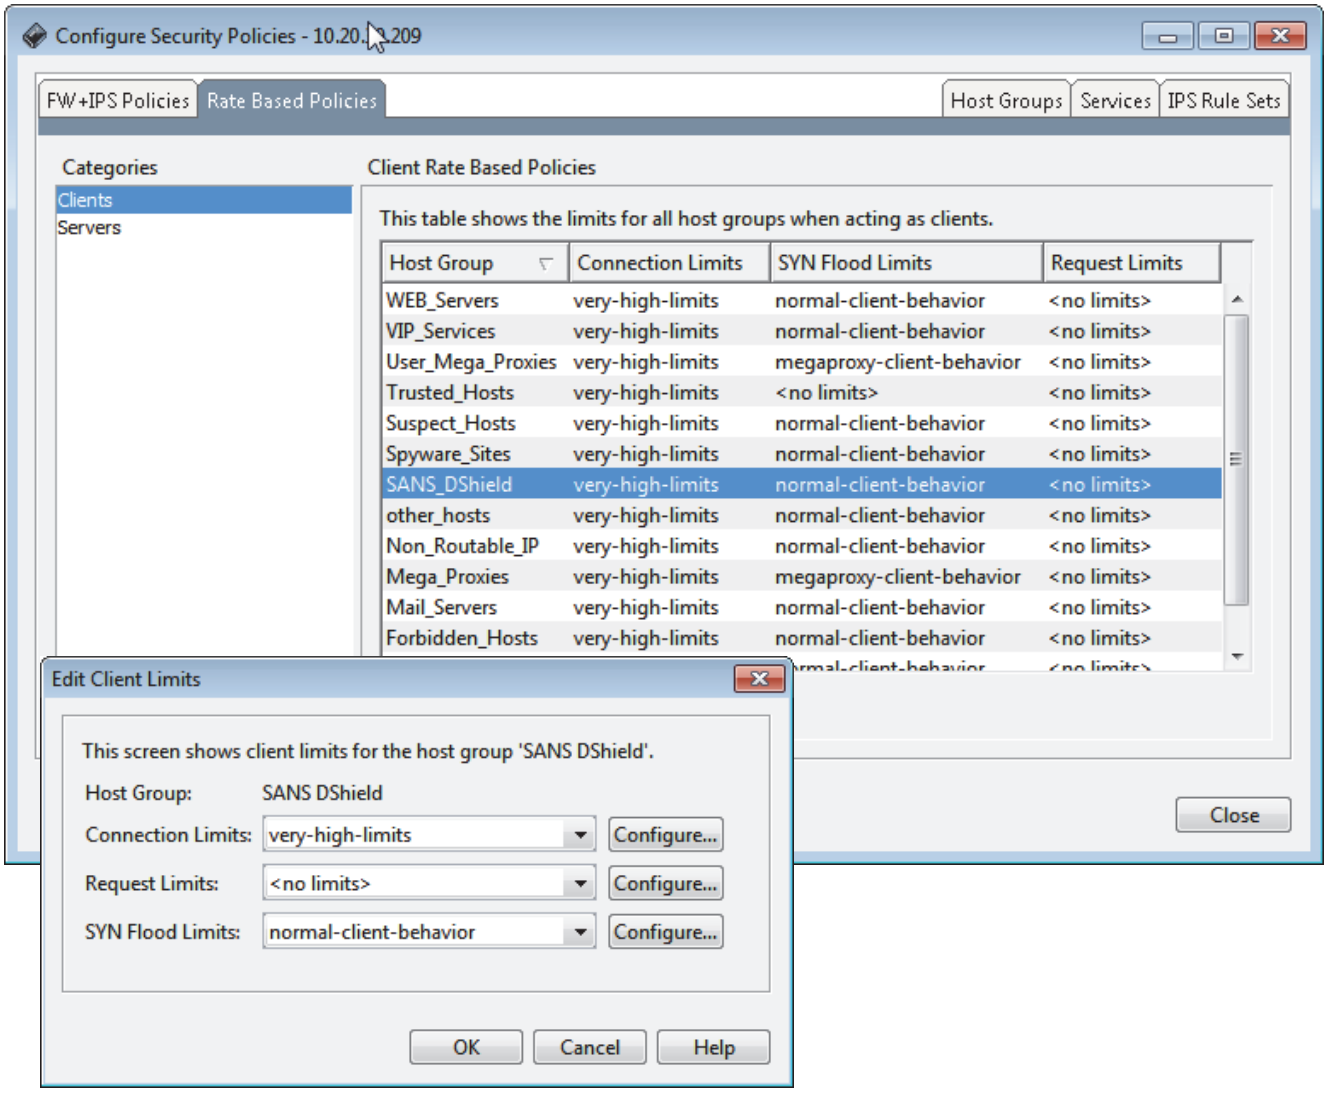

Step 4: Enabling the Relevant Rules for the Desired Rule Set

Once you have specified parameters for client request limiting, you need to ensure the desired rules are enabled, and that the settings have been modified to meet your current requirements. To modify rule settings:

d limiting rate-based policy for a specific host group:

1. Do one of the following:
   - Click Security Policies on the toolbar
   - Choose Manage Security > Security Policies from the navigation tree.

   The Configure Security Policies dialog box displays.

2. Click the IPS Rule Sets tab. Select the desired rule set.

   NOTE

   For more information on managing rule sets, see Chapter 15, "Managing Rules and Rule Sets".

3. To view the rules associated with SYN flood, connection, and client limiting that are not associated with HTTP or DNS on the management application for a Corero Network Device, choose Configure Security > Advanced Security Config > Rate Based Rules from the Navigation Tree.
4. To modify the settings for a rule, select the rule, then click Edit. The Edit Rule Settings dialog box displays (Figure 21-4).

Figure 21-4: Modifying Rule Settings

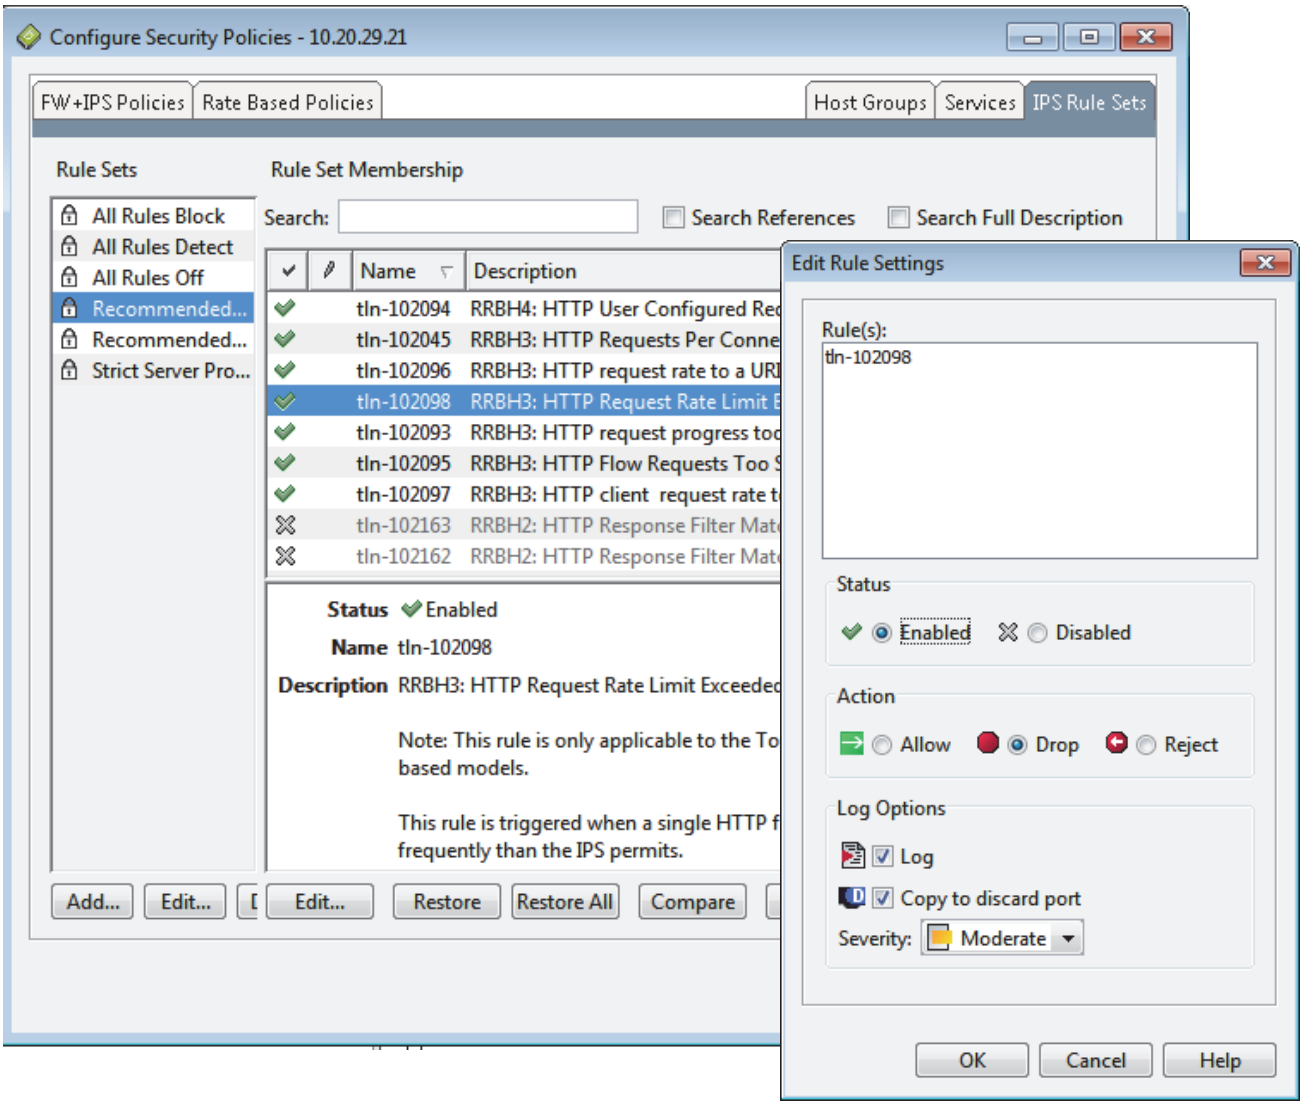

5. You can enable or disable individual rules. If you enable a rule, you can set its Action.

NOTES

1. This setting applies to this individual rule, and overrides the settings established using the Edit a Rule Set window.
2. While it is possible to set a rate-based rule to Block, the rate measurement features are only useful if the rule is set to Allow, enabling traffic to incur Client Request Credit costs that will engage rate-limiting penalties.

Possible actions are:

- Allow— Pass the traffic.
- Drop— Passively block the traffic with no indication to the offending client.

- Reject— Actively block the traffic by dropping the traffic and taking specific action, for example, by resetting the connection.

CAUTION

If you disable a rule, packets that would normally trigger the rule are passed, and no logging is performed.

6. Specify the Log options for this rule as follows:
   - Log - Send information to a log file based on its severity rating.
   - Copy to Discard Port - Copy the associated traffic to the Discard port based on its severity rating.

NOTE

If you have configured a Discard Port, you can log all traffic that triggers a rule, even if the Action for that rule is set to Allow. This feature enables you to record usage information about applications you allow as well as those you disallow.

   - Severity - The log severity determines whether packets for this rule are processed by the device's Event Logging System. To be processed, the priority you set for this rule must be equal to, or above, the overall priority set for event messages, as described in Logging Options (page 15-4). Severity levels include:
     - Low (green)
     - Moderate (yellow)
     - Critical (red)
7. When you have finished specifying rule settings in the management application for a Corero Network Device, click OK.
8. Save your changes by clicking the Save Configuration toolbar button.

   When the rule appears in the list of rules, a pencil icon displays indicating that this rule has been modified from its default settings.

NOTE

If Client Request Limiting is configured, and the rule is enabled, each RRBxx IPS rule deducts the cost configured in the appropriate Server Rate profile from the Client Request Credits for the IP address of the client whose traffic is triggering the rule.

Checking Client Request Credits for an IP Address

You can use the IP Address Query feature to view current counter and client request credit information for a particular IP address. For additional information on the IP Address Query dialog box, see Reset (Clear) SYN Flood and Connection Counters (page 19-31). To query an IP address:

1. On the management application for a Corero Network Device, do one of the following:
   a. From the Navigation Tree, choose Monitor Security > Query IP Address.
   b. From the Security Event Viewer, with the Active Mode check box cleared (deselected), select an event then click Query Client or Query Server.

   The IP Address Query dialog box displays (Figure 19-8).
2. Enter the IP address of the host about which you want to view additional information.
3. Select the Corero Network Device you want to query. If you want to query all devices, choose <All>.
4. On the management application for a Corero Network Device, click Query.
5. From the IP Address Query dialog box, you can:
   - View the current number of client request credits.
   - Clear (zero) the number of client request credits.
   - Reset other system counters.

NOTE

For additional information on the IP Address Query dialog box, see Reset (Clear) SYN Flood and Connection Counters (page 19-31).

Chapter 22
Advanced Client Rate Limiting

This chapter describes the more advanced features of Client Rate Limiting. As a prerequisite, you must familiarize yourself with the information in Chapter 21, ‘‘Client Rate Limiting”. In addition to the information in that chapter, a working knowledge of HTTP and DNS, and a clear understanding of the operating environment are also critical to proper understanding and configuration of these features.

The Corero Network Device contains detailed information about specific protocols, used for protocol validation. In the case of HTTP and DNS, this extends to monitoring for specific rate-based and behavior-based monitoring. This advanced monitoring can be used to identify different types of rate-based attacks.

NOTE

These features are useful for protecting servers and services against attack. If you are deploying the Corero Network Device in a mixed environment, you must take care to correctly identify clients and servers and apply the correct profiles and rules to relevant traffic. Failure to do so may result in network traffic interruption.

As described in Chapter 21, ‘‘Client Rate Limiting”, particular user-specified limits can trigger rules when suspicious or malicious behavior occurs. Advanced client rate limiting enables you to take advantage of additional HTTP and DNS profile settings related to specific rules called Request/Response Behavioral (RRB) rules. Note that these rules can also be configured to block malicious traffic independently of Client Rate Limiting features, which can enable you to block certain types of attacks, such as DNS attacks where the source IP address is random and spoofed.

The features and rules described in this chapter are only available on the IPS 5100 and 5200 hardware models.

NOTE

For assistance in configuring your HTTP and DNS advanced Client Request Limiting features, contact Corero Network Security.

This chapter contains the following sections:

How Do I Customize Client Rate Limit Credit Deductions?

During normal system operation, credits are deducted as traffic is processed. By default, one credit is deducted from the available credits for a particular IP address and packet type.

NOTE

If you are considering modifying the per-packet costs for your Corero Network Device, Corero recommends that you perform an initial test using a custom host group with only a few IP addresses. This way, you can test whether you are providing enough credits on a regular basis to ensure proper operation. You will likely need to tune your available credits and your per-packet costs.

To modify the per-packet cost of different *client-initiated* packet types:

1. From the navigation tree, choose Configure Security > Advanced Security Config > Client Packet Cost Customization. The Configure Per-Packet Costs dialog box displays.
2. Specify the per-packet cost for the traffic types described in Table 22-1.

NOTE

Specifying a per-packet cost of zero (0) is a valid setting, and may be useful for some configurations.

3. When finished, click OK.
4. Save your changes by clicking the “Save Configuration” Toolbar button.

Table 22-1: Configure Per-Packet Costs

This Protocol Type...	Includes Client-Initiated Packets Such as...
TCP SYN packets (But not SYN-ACK)	SYN **Note:** Does not include SYN-ACK, because this packet type comes from the server.
TCP packets containing data	TCP packets with a TCP payload, such as GET.
Other TCP packets	TCP packets with no TCP payload, such as ACK.
UDP packets	Any UDP packet
ICMP packets	Any ICMP packet
Other packets	Any other packet type, including those with a protocol field that is not UDP, TCP, or ICMP.

How Do I Customize Protocol Client Rate Limit Deductions?

There are two areas of configuration required, in addition to the basic Client Request Limiting, which is described in Chapter 21, "Client Rate Limiting", that must be configured for these rules to be effective.

RRBxx Rules, when applied to network traffic, cause Client Request Credit deductions to occur when triggered. In addition, these rules enable you to specify whether the Corero Network Device will Allow or Drop the specific connection. Note that Client Request Credit deductions will always be taken for those services with Client Request Limiting enabled. These deductions are taken whether the rules are designated to Allow or Drop the specified traffic.

The server Rate Based Policy Profiles define the parameters for these IPS rules, specifically the behavioral parameters of when the rules trigger, and the Client Request Credit cost associated with each triggering incident. These policies are applied to Server Host groups. HTTP and DNS IPS Rules apply Client Request Credit deductions differently, depending on how they are configured.

Rate based RRB rules will trigger blocked and dropped packet counts and apply Client Request Credit deductions differently depending on whether they are HTTP rules or DNS rules:

NOTE

By design, RRBxx Rules only trigger only once per flow, but, depending on the rule, a cost may be incurred for each additional incident after the rule is triggered.

HTTP RRB rules behave as follows:

- When HTTP rules are configured to Drop, the rules will trigger only once per flow, and incur a single Client Request Credit deduction. The connection is then dropped, and, as a result of the loss of this traffic, the device will not incur any additional costs associated with this flow.
- When HTTP rules are configured to Allow, the client will see a Client Request Credit deduction of the appropriate value for every packet over the threshold that causes one of the rules to trigger.

DNS RRB rules behave as follows:

- DNS rules treat every packet independently, regardless of whether they are part of the same flow or not. The end result of this is that DNS will decrement the client credits for every offending packet over a rule's configured limit, regardless of whether every additional packet is part of the same flow or not.
- Unlike HTTP rules which may act to block the flow, DNS makes per-packet decisions whether to block or not, and, again, will decrement client credits for every packet the device decides to block.
- When DNS rules are configured to Allow, the Client Request Credit behavior is the same as when the rules are configured to drop. The only behavioral difference is that the offending packets are allowed through.

Client Credit Deductions Based on Per Packet Costs

By default, the per-packet cost for all packet types is set to 1. This should be sufficient for those services for which you enable client rate limiting.

You can modify the per-packet cost for each packet type across all profiles. You can also use the Cost Scale Factor which acts as a multiplier for *all* costs associated with a specific profile.

NOTE

Per packet costs apply to the entire device, not to any particular host group or policy, and they apply for any service for which client rate limiting is enabled.

CAUTION

Use caution when increasing the per-packet cost of any packet type. Increasing the cost of a specific packet type can significantly increase the packet cost of a session, and when you multiply that by many sessions, the total cost can quickly accumulate. This may result in unintentionally blocked legitimate traffic.

Table 22-2 lists the packet costs associated with a common client/server interaction. The third column lists the relevant packet costs associated with the default per-packet cost settings. The fourth column lists the relevant packet costs associated with the modified per-packet cost settings shown in Figure 22-1.

NOTE

No cost (0) is associated with any traffic sent from the server.

Table 22-2: Client Credit Costs During Traffic Flow

Device	Network Traffic Type	Default Per-Packet Cost	Example of Modified Per-Packet Cost
Client	Sends SYN to server	1	5
Server	Sends SYN-ACK to client. Per-packet costs are not associated with server-initiated packets.	0	0
Client	Sends ACK to server	1	3
Client	Sends GET to server	1	1
Server	Sends ACK-PUSH to client	0	0
Client	Sends another GET to server	1	1
Server	Sends another ACK-PUSH to client	0	0
Client	Sends FIN to server	1	3
Server	Sends FIN ACK to client	0	0
Client	Sends RST to server	1	3
Server	Sends RST to client	0	0
	Note the Different in the Total Cost for the Session	**6 (Total)**	**16 (Total)**

Figure 22-1 shows how the default per-packet costs would be modified in the example shown in Table 22-2.

Figure 22-1: Example Settings for Per-Packet Costs

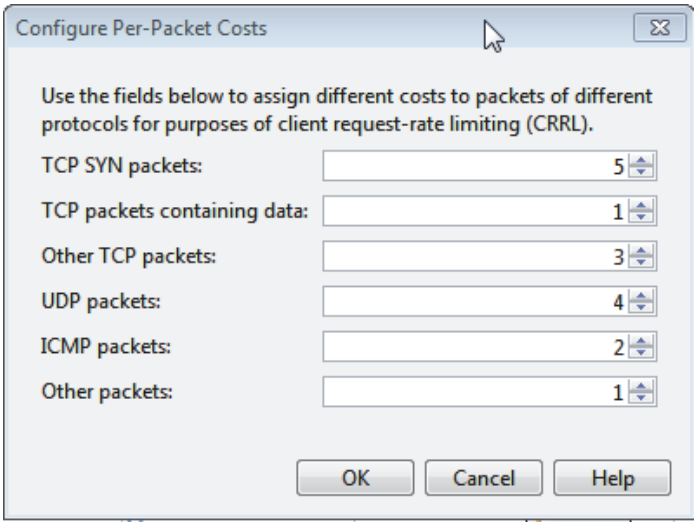

How Profile Settings Affect Rate Limiting Behavior

For most new profiles, there are two components:

- The limit - the maximum number of incidents the system will pass without triggering a rule and taking the specified action.
- The sample period - the maximum amount of time during which incidents will be counted. If a setting has a sample period, once the sample period is over, the incident counter resets to zero, and counting begins again.

You specify the limit beyond which you want the rule triggered. If you want the rule to trigger on the 7th incident of a particular behavior, you set the limit to 6. If you want the rule to trigger on the first incident, you set the limit to 0. Note that if you have set the limit to zero, the time limit is immaterial because traffic will be blocked on the first instance.

You may also specify information you want the Corero Network Device to search for in various information associated with a packet. For example, you can specify a header string, a URI, or a top level domain. Note that searching for a longer specified string is faster than searching for a shorter string, so ensure you specify the longest string possible that will suit your needs.

Note that limits are only assessed, and costs are only incurred, on services that have Client Rate Limiting enabled. For information on how to do this, see Specifying Advanced Service Settings (page 14-5).

NOTE

You can neither delete nor modify a factory default profile. You cannot delete a user-specified profile that is currently in use.

Table 22-4 shows an example of the contents of an HTTP Header String Limit profile.

How Rule Settings Affect Rate Limiting Behavior

HTTP rules and DNS rules affect credit deductions differently, depending on whether their disposition is set to Drop or Allow. These differences are by design.

- HTTP rules tend to block the flow as a whole when rules-based limits are exceeded.
- DNS rules tend to make decisions whether or not to block on a packet-by-packet basis, and will decrement client credits for every packet that is blocked.

These differences are described in greater detail in Table 22-3.

Table 22-3: How Rule Type and Disposition Affect Rate Limiting Behavior

When the Rule Type Is...	And the Rule Disposition Is...	Then the Following Occurs...
HTTP (RRBHx)	Allow	As long is client credits are available, traffic from a client will be allowed regardless of whether or not the rule has been triggered. Once a rule has triggered, client credits will be decremented for every packet over the threshold that caused a rule to trigger. If the number of client credits falls to or below zero, additional penalties may be incurred (flow dropped, packets blocked, and so forth).
HTTP (RRBHx)	Drop or Reject	HTTP rules will trigger “once per flow”. Each rule will only trigger one Drop or Reject event per flow, resulting in a single reduction in client credits. You will not see additional decrements for each dropped packet in the flow beyond the initial one that triggered the rule. If an HTTP packet triggers a rule that is set to block, the entire flow will be blocked.
DNS (RRBDx)	Allow	DNS will decrement client credits for every offending packet over a rule’s configured limit, regardless of their flow. DNS treats every packet independently, regardless of whether or not they are part of the same flow. Traffic will be allowed by the rule, but when the number of credits reaches zero or below, additional penalties may be incurred (flow dropped, packets blocked, and so forth).
DNS (RRBDx)	Drop or Reject	DNS will decrement client credits for every offending packet over a rule’s configured limit, regardless of their flow. DNS treats every packet independently, regardless of whether or not they are part of the same flow. So If a DNS packet triggers a rule that is set to block, only the offending packet is discarded, not the entire flow. Traffic will be dropped or rejected when the rule is triggered.

Maximum Limits Per Profile

When configuring profiles, keep in mind the maximum number of values for particular parameters. These limits include:

- Maximum number of header strings configured in an HTTP Header String Limits profile: 16
- Maximum length of a specified header string in an HTTP Header String Limits profile: 64 characters
- Maximum number of URI rate limits in an HTTP URI profile: 16
- Maximum number of HTTP response codes in an HTTP Response profile: 64

Configuring HTTP and DNS Profiles

The overall process for configuring HTTP and DNS profiles is the same as the process for configuring other rate-based profiles, except the manner in which you specify client limits is more powerful and more specific.

This process includes:

- Step 1: Preparing Host Groups (page 21-6)
- Step 2: Enabling Request Limiting for a Service (page 21-8)
- Configuring Client Request Limit Profiles
  For overview information, see Chapter 21, "Client Rate Limiting"
  For specific procedures, see the detailed instructions below.
- Step 4: Enabling the Relevant Rules for the Desired Rule Set (page 21-12)

However, instead of following the procedure for specifying Client Request settings in Chapter 21, "Client Rate Limiting", you must configure rate and cost information for the available HTTP and DNS-specific rules, as described in the procedures that follow.

CAUTION

Profile names must be unique, not only across profiles of the same type (for example, HTTP Header String Limit profiles) on the Corero Network Device, but across all profile types on that unit.

To configure HTTP and DNS profiles:

1. Go to the Configure Security Policies window either by clicking on the Security Policies toolbar button, or by choosing Manage Security > Security Policies from the navigation tree.
2. Click the Rate Based Policies tab (Figure 21-3).
3. Click Servers, then select the server host group whose clients you want to rate limit.
4. Click Edit. The Edit Server Limits dialog box displays. There are three tabs on this dialog box: General, HTTP Limits and DNS Limits.

   NOTE

   For information on specifying SYN Flood limits on the General tab, see Chapter 20, "SYN Flood and Connection Limiting Security". For information on specifying Client Request limits on the General tab, see Chapter 21, "Client Rate Limiting".

5. You can specify HTTP Header String Limits for client request limiting. Select the desired profile.

   Or, if you need to modify HTTP Header String Limit profiles or their contents:

   a. On the Edit Server Limits dialog box, click the HTTP Limits tab.
   b. In the HTTP Header String Limits area, click Configure. The Configure HTTP Header String Profiles dialog box displays.

      From here, you can add or delete HTTP Header String Profiles, or modify the profiles' contents.

   c. To add a profile, click Add, and provide the profile name.
   d. To delete a profile, select the profile in the Profiles list, then click Delete.
   e. To view the contents of a specific profile, select the profile name. The contents display to the right.

f. To add new contents to the selected profile, click Add in the right pane. Or, if you want to modify existing profile contents, select the profile contents from the list in the right pane, then click Edit. Profile Contents are described in Table 22-4.

g. To delete contents from a profile, select the profile contents in the right pane, then click Delete.

h. When finished modifying profiles, click Close.

i. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

NOTE

You must click Apply in order for your changes to take effect.

j. If desired, in the HTTP Header String Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

Table 22-4: HTTP Header String Limit Profile Contents

Parameter	Description	Associated Rule
Header String	HTTP Headers have the following format: <header-name>:[<space><header-values>] The Header String represents the text before the colon. You can choose whether you want to specify a header string, or not. If you specify a header string with no matched value, the profile will identify all packets with that type of header string.	None
Matched Value	HTTP Headers have the following format: `<header-name>:[<space><header-values>]`, for example: `User-Agent: HTTP Retriever` The Matched Value represents any text after the colon. You can choose whether or not to specify a matched value. If you specify a matched value with no header string, the profile will identify all packets with that value for any type of header string. **NOTE:** Searching for a matched value with no header string consumes more system resources than searching with a header string.	None
Cost Per Instance, Rate Limit Exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	None
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None

Table 22-4: HTTP Header String Limit Profile Contents *(Continued)*

Parameter	Description	Associated Rule
Trigger if Header is Absent	Use this check box to specify when a rule will trigger: • Unchecked (cleared) - Trigger when the specified header string or value is present. • Checked (selected) - Trigger when the specified header string or value is absent **NOTE:** When this check box is selected, the search will consume additional system resources.	• Header is Present: Rules tln-105010 through tln-105025. For additional information see RRBH1: HTTP Header or String Found in Request (tln-105010 Through tln-105025) (page 22-30). • Header is Absent: Rules tln-105030 through tln-105045. For additional information, see RRBH1: HTTP Header or String Missing From Request (tln-105030 through tln-105045) (page 22-31).

6. You can specify HTTP Response Limits for client request limiting. Select the desired profile.

   Or, if you need to modify HTTP Response Limit profiles or their contents:

   a. On the Edit Server Limits dialog box, click the HTTP Limits tab.

   b. In the HTTP Response Limits area, click Configure. The Configure HTTP Response Profiles dialog box displays.

      From here, you can add or delete HTTP Response Profiles, or modify the profiles' contents.

   c. To add a profile, click Add, and provide the profile name.

   d. To delete a profile, select the profile in the Profiles list, then click Delete.

   e. To view the contents of a specific profile, select the profile name. The contents display to the right.

   f. To add new contents to the selected profile, click Add in the right pane. Or, if you want to modify existing profile contents, select the profile contents from the list in the right pane, then click Edit. Profile Contents are described in Table 22-5.

   g. To delete contents from a profile, select the profile contents in the right pane, then click Delete.

   h. When finished modifying profiles, click Close.

   i. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

      NOTE

      You must click Apply in order for your changes to take effect.

   j. If desired, in the HTTP Response Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

Table 22-5: HTTP Response Profile Contents

Parameter	Description	Associated Rule
HTTP Response Code	Specify the HTTP Response Code (sent from the server to the client in response to a client request) whose instances you want to limit. RFC2616 provides an authoritative list of HTTP response codes. You can specify a wildcard (asterisk *) as the last character of the response code. For example, 40* would specify codes 400 through 409, including overloaded IIS response codes such as 404;1. Entering 4* would include codes 400 through 499. **Note:** You can only specify a wildcard (asterisk *) as the last character, and not anywhere before the last character. If you specify both an explicit match (404, for example), and a wildcard match (40*, for example) in the same profile that might match the same response code, the explicit (exact) match will be used, even if it is not the first entry in the profile definition.	Rules tln-102100 through tln-102163. You can specify up to 64 response codes (with associated costs) per profile. For additional information see RRBH2: HTTP Response Filter Match (tln-102100 Through tln-102163) (page 22-29).
Cost Per Instance, Rate Limit Exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	None
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None

7. You can specify HTTP Request Parameter Limits for client request limiting. Select the desired profile.

   Or, if you need to modify HTTP Request Parameter Limit profiles or their contents:

   a. On the Edit Server Limits dialog box, click the HTTP Limits tab.
   b. In the HTTP Request Param Limits area, click Configure. The Configure HTTP Response Profiles dialog box displays.
   c. From here, you can add or delete HTTP Request Parameter Profiles, or modify the profiles' contents.
   d. To add a new profile, click Add, and specify the profile name. Or, if you want to modify an existing profile, select the profile, then click Edit. Profile Contents are described in Table 22-6.
   e. To delete contents from a profile, select the profile contents in the right pane, then click Delete.
   f. When finished modifying profiles, click Close.
   g. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

   NOTE

   You must click Apply in order for your changes to take effect.

h. If desired, in the HTTP Request Parameter Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

Table 22-6: HTTP Request Parameter Profile Contents

Parameter	Description	Associated Rule
Exceeded Limits Tab - Rule tln-102098		
Cost Per Instance, Rate Limit Exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-102098 triggers when the HTTP request rate limit is exceeded by a client IP address during the specified time period. For additional information see RRBH3: HTTP Request Rate Limit Exceeded (tln-102098) (page 22-28).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Exceeded Limits Tab - Rule tln-102045		
Cost, too many requests per flow	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-102045 triggers when there are too many HTTP requests per flow from an IP source address during the specified time period. For additional information see RRBH3: HTTP Requests Per Connection Exceeded Specified Maximum (tln-102045) (page 22-22).
Exceeded Limits Tab - Rule tln-102096		
Cost per Repeat URI Request, Rate Limit Exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-102096 triggers when there are too many repeat URI requests from an IP source address during the specified time period. For additional information see RRBH3: HTTP Request Rate to a URI is Too High (tln-102096) (page 22-26).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Exceeded Limits Tab - Rule tln-102097		
Cost per URI Request from the Same Client, Rate Limit Exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-102097 triggers when there are too many URI requests from the same IP source address during the specified time period. For additional information see RRBH3: HTTP Client Request Rate to a URI is Too High (tln-102097) (page 22-27).

Table 22-6: HTTP Request Parameter Profile Contents *(Continued)*

Parameter	Description	Associated Rule
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Unmet Limits Tab - Rule tln-102095		
Cost Per Request, Rate Limit not Met	The cost (credit reduction) incurred when the rate limit is not met.	Rule tln-102095 triggers when there are not enough HTTP requests within the specified time period. For additional information see RRBH3: HTTP Flow Requests Too Low (tln-102095) (page 22-25).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Unmet rate trigger lifetime (seconds)	The number of seconds you want the device to wait for a follow-up request. You can specify this value so, if the follow-up to a request takes longer than the specified rate limit (in seconds), the rule will still trigger. For example, if you specified a rate limit of 5 seconds and an unmet rate trigger of 10 seconds, if a follow-up to a request took 8 seconds to arrive, even though it exceeded the rate limit in seconds, the rule will still trigger.	None
Unmet Limits Tab - Rule tln-102093		
Cost Per HTTP Request Packet, Request Proceeding Too Slowly	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-102093 triggers when the HTTP flow is proceeding too slowly, and not enough bytes of HTTP request (get) data were received within the specified time period (in milliseconds). For additional information, see RRBH3: HTTP Request Progress Too Slow (tln-102093) (page 22-23).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None

8. You can specify HTTP URI Limits for client request limiting. Select the desired profile.

   Or, if you need to modify HTTP URI Limit profiles or their contents:

   a. On the Edit Server Limits dialog box, click the HTTP Limits tab.

b. In the HTTP URI Limits area, click Configure. The Configure HTTP URI Profiles dialog box displays.

   From here, you can add or delete HTTP URI Profiles, or modify the profiles' contents.

c. To add a profile, click Add, and provide the profile name.

d. To delete a profile, select the profile in the Profiles list, then click Delete.

e. To view the contents of a specific profile, select the profile name. The contents display to the right.

f. To add new contents to the selected profile, click Add in the right pane. Or, if you want to modify existing profile contents, select the profile contents from the list in the right pane, then click Edit. Profile Contents are described in Table 22-7.

g. To delete contents from a profile, select the profile contents in the right pane, then click Delete.

h. When finished modifying profiles, click Close.

i. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

   NOTE

   You must click Apply in order for your changes to take effect.

j. If desired, in the HTTP URI Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

Table 22-7: HTTP URI Profile Contents

Parameter	Description	Associated Rule
URI	The Uniform Resource Identifier whose requests you want to count. The URI is the final part of the URL. The specified value must start with a forward slash (/). The value must not exceed 256 characters, including the forward slash (/). For example, in the URL http://www.corero.com /content/products/index.jsp, the method field would be *http*, the host field would be *www.corero.com*, and the URI would be everything that follows.	Rule tln-102094 triggers when the number of requests for the same URI that occur during the same flow exceeds the limit during the specified time period. For additional information see RRBH4: HTTP User Configured Request URI Rate Limit Exceeded (tln-102094) (page 22-24).
Cost per URI request, rate limit exceeded.	The cost (credit reduction) incurred when the rate limit is exceeded.	None
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None

9. You can specify DNS Parameter Limits for client request limiting. Select the desired profile.

   Or, if you need to modify DNS Parameter Limit profiles or their contents:

   a. On the Edit Server Limits dialog box, click the DNS Limits tab.

   b. In the DNS Param Limits area, click Configure. The Configure DNS Param Profiles dialog box displays.

      From here, you can add or delete DNS Parameter Profiles, or modify the profiles' contents.

c. To add a new profile, click Add, and specify the profile name. Or, if you want to modify an existing profile, select the profile, then click Edit. Profile Contents include are described in Table 22-8.

d. To delete contents from a profile, select the profile contents in the right pane, then click Delete.

e. When finished modifying profiles, click Close.

f. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

NOTE

You must click Apply in order for your changes to take effect.

g. If desired, in the DNS Parameter Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

Table 22-8: DNS Parameter Profile Contents

Parameter	Description	Associated Rule
Profile Name	The name of this profile	None
Rule tln-101078		
Maximum DNS request length, in bytes	The maximum permitted DNS request length.	If a DNS request exceeds the length specified, rule tln-101078 triggers. For more information, see AAUPV: DNS Request Exceeds Maximum Allowed Length in Bytes (tln-101078) (page 22-37).
Rule tln-101079		
Recursive request cost, rate limit exceeded	The cost (credit reduction) incurred when the rate limit is exceeded. The cost is subtracted from the credits of the client IP of the offending request.	Rule tln-101079 triggers when the total number of recursive DNS requests for a given domain exceeds the limit during the specified time period. For additional information see RRBD1: Rate of DNS Recursive Requests for a Domain Has Been Exceeded (tln-101079) (page 22-38).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Rule tln-101077		
Non-recursive request cost, rate limit exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-101077 triggers when the number of non-recursive DNS requests from an IP source address exceeds the limit during the specified time period. For additional information see RRBD1: Rate of DNS Non-Recursive Requests for a Domain Has Been Exceeded (tln-101077) (page 22-36).

Table 22-8: DNS Parameter Profile Contents *(Continued)*

Parameter	Description	Associated Rule
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Rule tln-101075		
Domain request cost, rate limit exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-101075 monitors DNS requests for all domains, and triggers when the number of DNS requests for the same domain coming from the same IP source address exceeds the limit during the specified time period. For additional information see RRBD1: DNS Requests to a Domain Exceed Limit (tln-101075) (page 22-34).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None
Rule tln-101076		
Host request cost, rate limit exceeded	The cost (credit reduction) incurred when the rate limit is not met. The client IP address of the offending request incurs the credit cost.	Rule tln-101076 triggers when the number of DNS requests for all hosts exceeds the limit during the specified time period. For more information, see RRBD1: DNS Requests to a Host Exceed Limit (tln-101076) (page 22-35).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None

10. You can specify DNS Top Level Domain Limits for client request limiting. Select the desired profile.

    Or, if you need to modify DNS Top Level Domain Limit profiles or their contents:

    a. On the Edit Server Limits dialog box, click the DNS Limits tab.

    b. In the DNS TLD Limits area, click Configure. The Configure DNS TLD Profiles dialog box displays.

    From here, you can add or delete DNS TLD Profiles, or modify the profiles' contents.

    c. To add a profile, click Add, and provide the profile name.

    d. To delete a profile, select the profile in the Profiles list, then click Delete.

    e. To view the contents of a specific profile, select the profile name. The contents display to the right.

f. To add new contents to the selected profile, click Add in the right pane. Or, if you want to modify existing profile contents, select the profile contents from the list in the right pane, then click Edit. Profile Contents are described in Table 22-9.

g. To delete contents from a profile, select the profile contents in the right pane, then click Delete.

h. When finished modifying profiles, click Close.

i. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

NOTE

You must click Apply in order for your changes to take effect.

j. If desired, in the DNS TLD Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

Table 22-9: DNS TLD Profile Contents

Parameter	Description	Associated Rule
TLD	Top Level Domain, such as .edu, .com, or .net.	None.
Cost Per TLD, Rate Limit Exceeded	The cost (credit reduction) incurred when the rate limit is exceeded.	Rule tln-101073 triggers when the number of DNS requests for hosts in the same Top Level Domain is exceeded by an IP source address during the specified time period. For additional information see RRBD2: User-Specified Blacklisted DNS Top Level Domain (tln-101073) (page 22-33).
Rate Limit (instances)	the number of occurrences the system will count before triggering the rule.	None
Rate Limit (seconds)	The sample period, in seconds, during which the rate limit may be reached.	None

11. You can specify DNS RCODE Limits for client request limiting. Select the desired profile.

Or, if you need to modify DNS RCODE Limit profiles or their contents:

a. On the Edit Server Limits dialog box, click the DNS Limits tab.

b. In the DNS RCODE Limits area, click Configure. The Configure DNS RCODE Profiles dialog box displays.

From here, you can add or delete DNS RCODE Profiles, or modify the profiles' contents.

c. To add a profile, click Add, and provide the profile name.

d. To delete a profile, select the profile in the Profiles list, then click Delete.

e. To view the contents of a specific profile, select the profile name. The contents display to the right.

f. To add new contents to the selected profile, click Add in the right pane. Or, if you want to modify existing profile contents, select the profile contents from the list in the right pane, then click Edit. Profile Contents are described in Table 22-10.

g. To delete contents from a profile, select the profile contents in the right pane, then click Delete.

h. When finished modifying profiles, click Close.

i. To apply your changes, on the Edit Server Limits dialog box, click Apply. Alternatively, if you do not want the changes applied, click Undo.

NOTE

You must click Apply in order for your changes to take effect.

j. If desired, in the DNS RCODE Limits area, modify the Cost Scale Factor. This value will be used to automatically multiply all costs associated with the selected profile. The default multiplier is 1.

12. When finished modifying your settings, click Apply.
13. Finally, save your changes by clicking the "Save Configuration" Toolbar button.

Table 22-10: DNS RCODE Profile Contents

Parameter	Description	Associated Rule
RCODE	The response code from a DNS server to the client that made the DNS request.	Rule tln-101080 through 101095 trigger when the number of DNS RCODE responses of a given type for an IP source address exceeds the specified limit during the specified time period. For additional information including the list of response codes associated with each rule, see RRBD3: DNS RCODE Matches Specified Filter (tln-101080 through tln-101095) (page 22-39).
Cost	The cost (credit reduction) is incurred when a client IP address has exceeded the limit for the number of RCODES of the specified type generated as a result of traffic from the client IP associated with a rule within this range is received.	

HTTP Client Rate Limiting Rules

The following rules are available for HTTP rate-based policies:

AAUPV: HTTP Requests Outstanding Exceeds Specified Maximum (tln-102036)

Table 22-11 provides detailed information about rule tln-102036.

Table 22-11: AAUPV: HTTP Requests Outstanding Exceeds Specified Maximum (tln-102036)

Rule Description	This rule detects and counts outstanding requests (those that have not yet received a response). When this request count exceeds the maximum configured number allowed for a single connection, this rule triggers.
Attack Countered	An attacker could overwhelm the system by sending so many HTTP requests that they consume available system resources. This could limit the server's ability to process legitimate requests, or could cause a server to crash, creating a denial-of-service condition.
Action on Rule Trigger	When this rule is triggered, the system takes the action specified for this rule in the selected rule set: Allow, Drop, or Reject. **NOTE:** This rule is an AAUPV (Acceptable Application Usage Policy Violation) rule, not an RRBxx rule, and *does not affect* Client Request Credits.
Rule Configuration (Quick Overview)	To configure the maximum allowed number of requests without a response: 1. Choose Configure Security > Advanced Security Config > IPS Rules Customization from the Navigation Tree. 2. Then, within the IPS Rules Customization dialog box, choose Protocol Validation Module > HTTP > Request.
Applicable Profile Limits	None

RRBH3: HTTP Requests Per Connection Exceeded Specified Maximum (tln-102045)

Table 22-12 provides detailed information about rule tln-102045.

Table 22-12: RRBH3: HTTP Requests Per Connection Exceeds Specified Maximum (tln-102045)

Rule Description	This rule detects and counts the number of HTTP requests per connection. When the request count exceeds the maximum configured number allowed for a single connection, this rule triggers. **NOTE:** This rule *does not track* the rate at which requests arrive, it simply counts the number of requests per connection.
Attack Countered	An attacker could overwhelm the system by sending so many HTTP requests that they consume available system resources. This could limit the server's ability to process legitimate requests, or could cause a server to crash, creating a denial-of-service condition.
Action on Rule Trigger	When this rule is triggered, the system takes the action specified for this rule, in the selected rule set: Allow, Drop, or Reject. The credit deduction varies depending on the disposition of the rule: • If the rule is set to Allow, the client will be charged the configured number of credits for every request over that limit, though that traffic will be allowed. • If the rule is set to Drop or Reject, the client is charged the configured number of credits a single time (for exceeding the limit), and traffic for that connection will be dropped or rejected.
Rule Configuration (Quick Overview)	You should configure the maximum allowed number of requests without a response in two locations in the management application. To configure the maximum number of requests: 1. Choose Configure Security > Advanced Security Config > IPS Rules Customization from the Navigation Tree. 2. Then, within the IPS Rules Customization dialog box, choose Protocol Validation Module > HTTP > Request. To configure the credit cost associated with exceeding the limit: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Request Param Limits area, click Configure. 6. Choose to add or edit a profile. 7. Click the Exceeded Limits tab. 8. Specify a value for Cost, Too Many Requests Per Flow.
Applicable Limits	None

RRBH3: HTTP Request Progress Too Slow (tln-102093)

Table 22-13 provides detailed information about rule tln-102093.

This rule differs from rule tln-102095 in that it measures the progress of each individual request based on data sent. In addition, its limit is measured in milliseconds, where other rules' limits are measured in seconds.

Table 22-13: RRBH3: HTTP Request Progress Too Slow (tln-102093)

Rule Description	This rule counts instances where client HTTP (get) requests are being sent too slowly from the client address. This rule measures the progress of each individual request.
Attack Countered	In this type of attack, such as a SlowLoris or OWASP "Post attack", a malicious client can initiate an HTTP connection, and then send data to continue that connection at a very slow rate, leaving the connection open for an unexpectedly long period of time. Multiple open connections of this type can consume system resources, which could limit the server's ability to process legitimate requests, creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period (in milliseconds). If progress does not meet the specified minimum time limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the minimum acceptable number of HTTP requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Request Param Limits area, click Configure. 6. Choose to add or edit a profile. 7. Click the Unmet Limits tab. 8. Specify values for the Request Proceeding Too Slowly section.
Applicable Limits	None

RRBH4: HTTP User Configured Request URI Rate Limit Exceeded (tln-102094)

Table 22-14 provides detailed information about rule tln-102094.

NOTE

This rule differs from rule tln-102096 in that it applies only to user-specified URIs.

Table 22-14: RRBH4: HTTP User Configured Request URI Rate Limit Exceeded (tln-102094)

Rule Description	This rule counts instances where a client made a request for a **user-specified** frequently used URI. The instances are counted across all flows from that client.
Attack Countered	In this type of attack, a client makes repeated requests for a frequently used URI. If enough requests are made, this could limit the server's ability to process legitimate requests, creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum number of requests for a frequently used URI: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP URI Limits area, click Configure. 6. Choose to add or edit the contents of a profile.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBH3: HTTP Flow Requests Too Low (tln-102095)

Table 22-15 provides detailed information about rule tln-102095.

This rule differs from rule tln-102093 in that it measures the rate of completed requests.

Table 22-15: RRBH3: HTTP Flow Requests Too Low (tln-102095)

Rule Description	This rule counts instances where client HTTP (get) requests are being sent too infrequently from the client address for a given flow. This rule measures the rate of completed requests.
Attack Countered	In this type of “slow” attack, a malicious client can initiate an HTTP flow, and then send packets to continue that connection at a very slow rate, thus leaving the connection open for an unexpectedly long period of time. Multiple open connections of this type can consume system resources, which could limit the server’s ability to process legitimate requests, creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the minimum acceptable number of HTTP requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Request Param Limits area, click Configure. 6. Choose to add or edit a profile. 7. Click the Unmet Limits tab.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBH3: HTTP Request Rate to a URI is Too High (tln-102096)

Table 22-16 provides detailed information about rule tln-102096.

This rule differs from rule tln-102094 in that it applies to all URIs.

Table 22-16: RRBH3: HTTP Request Rate to a URI is Too High (tln-102096)

Rule Description	This rule counts the rates of requests to **any** URI from all client IP addresses. When a URI is accessed too frequently, the rule triggers.
Attack Countered	In this type of attack, a large volume of HTTP requests is made to a specified URI in a brief period of time. If this volume of requests is permitted, they could consume web server resources, which could limit the server's ability to process legitimate requests, creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum number of HTTP requests for a given URI: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Request Param Limits area, click Configure. 6. Choose to add or edit a profile. 7. Click the Exceeded Limits tab.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.
Configuration Considerations	You should specify a setting that is consistent with the capacity of the server hosting the URI.

RRBH3: HTTP Client Request Rate to a URI is Too High (tln-102097)

Table 22-17 provides detailed information about rule tln-102097.

Table 22-17: RRBH3: HTTP Client Request Rate to a URI is Too High (tln-102097)

Rule Description	This rule counts the number of HTTP requests to URIs that are made by a single client. If a client requests a URI too frequently, the rule is triggered.
Attack Countered	In this type of attack, a client can repeatedly request the same URI, overwhelming the system's resources. This could limit the server's ability to process legitimate requests, creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum acceptable number of HTTP requests to a URI: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Request Param Limits area, click Configure. 6. Choose to add or edit a profile. 7. Click the Exceeded Limits tab.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.
Configuration Considerations	You should specify a setting that is consistent with the capacity of the server hosting the URI.

RRBH3: HTTP Request Rate Limit Exceeded (tln-102098)

Table 22-18 provides detailed information about rule tln-102098.

Table 22-18: RRBH3: HTTP Request Rate Limit Exceeded (tln-102098)

Rule Description	This rule is triggered when a single HTTP flow has exceeded the maximum number of requests permitted in the user-specified time.
Attack Countered	Attackers can use a single HTTP flow to quickly make a large volume of requests. Once HTTP resources are consumed by this type of attack, this could limit the server's ability to process legitimate requests, creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum acceptable number of HTTP requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Request Param Limits area, click Configure. 6. Choose to add or edit a profile. 7. Click the Exceeded Limits tab.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBH2: HTTP Response Filter Match (tln-102100 Through tln-102163)

Table 22-19 provides detailed information about rule tln-102100 through tln-102163.

You can find a list of valid HTTP response codes in RFC 2616.

Table 22-19: RRBH2: HTTP Response Filter Match (tln-102100 Through tln-102163)

Rule Description	The device inspects HTTP server responses for configured HTTP response codes (404, 501, etc). When an instance of a user-configured response code is found, the rule triggers.
Attack Countered	Attackers can make requests to servers that cause errors or timeout conditions to occur, placing increased or excessive load on the server. In addition, an attacker may request a sequence of URIs, probing to find a file or resource they believe exists. Monitoring the responses from the server, and assigning Client Request Credit deductions for some specific server responses, provides a mechanism for penalizing clients who persistently exhibit this behavior.
Action on Rule Trigger	If a specific HTTP response code is generated, the specified cost is deducted from the credits available for the client who initiated the traffic that triggered the response code.
Rule Configuration (Quick Overview)	To configure the maximum acceptable number of HTTP response codes: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Response Limits area, click Configure. 6. Choose to add or edit the contents of a profile.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBH1: HTTP Header or String Found in Request (tln-105010 Through tln-105025)

Table 22-20 provides detailed information about rules tln-105010 through tln-105025.

Table 22-20: RRBH1: Header or String Found in Request (tln-105010 through tln-105025)

Rule Description	The device observes HTTP request headers in client requests. The device counts instances where a user-specified HTTP header or string ***was*** found in a request sent from a specific client. When the count exceeds the specified limit, the rule triggers. You can specify either a header or a string, or you can specify both. HTTP headers have the following format: <header-name>:[<space><header-values>]. The Header String represents the text before the colon. The matched value represents the text that follows the colon. Note that this is a range of rules, with one rule associated with each header or string search you specify. The first item you specify becomes rule tln-105010, then second becomes tln-105011, and so on.
Attack Countered	Attackers can make requests with specific content in a header that identifies the requests as attacks or non-typical end user application traffic. This rule may be triggered by a robot scanner probing the server, rather than an outright attack. This rule limits the system load generated in either case. You can also use these rules to restrict client access to only certain browsers, or other clients that can be distinguished by a unique header or header string.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the minimum acceptable number of matching HTTP requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Header String Limits area, click Configure. 6. Choose to add or edit the contents of a profile. 7. To ensure you are searching for the presence of the string in the header, ensure the “Trigger of header is absent” check box is cleared (unchecked).
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models. You may specify up to 16 header strings to search for, one for each rule. Note that this rule only searches the header string contents up to the initial colon (:).

RRBH1: HTTP Header or String Missing From Request (tln-105030 through tln-105045)

Table 22-21 provides detailed information about rule tln-105030 through 105045.

Table 22-21: RRBH3: HTTP Header or String Missing from Request (tln-105030 through tln-105045)

Rule Description	The device observes HTTP request headers in client requests. The device counts instances where a user-specified HTTP header or string ***was not*** found in a request sent from a specific client. When the count exceeds the specified limit, the rule triggers. You can specify either a header or a string, or you can specify both. HTTP headers have the following format: <header-name>:[<space><header-values>]. The Header String represents the text before the colon. The matched value represents the text that follows the colon.
Attack Countered	Attackers can make requests with one or more HTTP headers absent from the request, for example, a missing user-agent header. Recognizing these incidents enables the device to identify the requests as attacks or non-typical end user application traffic.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the minimum acceptable number of HTTP requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the HTTP Limits tab. 5. In the HTTP Header String Limits area, click Configure. 6. Choose to add or edit the contents of a profile. 7. To ensure you are searching for the absence of the string in the header, ensure the “Trigger of header is absent” check box is selected (checked).
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models. You may specify up to 16 header strings to search for, one for each rule. Note that this rule only searches the header string contents up to the initial colon (:).

DNS Client Rate Limiting Rules

The following rules are available for DNS rate-based policies:

RRBD2: User-Specified Blacklisted DNS Top Level Domain (tln-101073)

Table 22-22 provides detailed information about rule tln-101073.

Table 22-22: RRBD2: User-Specified Blacklisted DNS Top Level Domain (tln-101073)

Rule Description	This rule counts instances where the requested DNS Top Level Domain (TLD), matches the user-specified black-listed (undesirable) domain. Top Level Domains are the final portion of a web address, such as .com, .gov, or .edu.
Attack Countered	Attacks that make repeated requests to a TLD, exceeding a user specified number of request in a user specified time frame may be identified using this rule. You may use this rule to discourage resolution requests for undesirable TLDs. Also, if the client lives in one TLD, and is repeatedly querying a different TLD, they may be trying to exhaust the resources of the local TLD servers.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure blacklisted DNS top level domains: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS TLD Limits area, click Configure. 6. Select a profile and choose to add or edit the profile's contents.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBD1: DNS Requests to a Domain Exceed Limit (tln-101075)

Table 22-23 provides detailed information about rule tln-101075.

Table 22-23: RRBD1: DNS Requests to a Domain Exceed Limit (tln-101075)

Rule Description	This rule counts the frequency of DNS requests to each domain. If the count for a domain exceeds the limit, the rule is triggered. This rule counts only the information immediately to the left of the Top Level Domain. The Top Level Domain is the last portion of the URL, such as .com, .edu, and .net.
Attack Countered	In this type of attack, a client can send high volumes of DNS requests to a domain. In attempting to service all of those requests, system resources can be consumed, possibly creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum number of DNS requests to a domain: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS Param Limits area, click Configure.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBD1: DNS Requests to a Host Exceed Limit (tln-101076)

Table 22-24 provides detailed information about rule tln-1021076.

Table 22-24: RRBD1: DNS Requests to a Host Exceed Limit (tln-101076)

Rule Description	This rule counts the frequency of DNS requests to each host. If the count for a host exceeds the limit, the rule is triggered. The host consists of the text string immediately to the left of the first dot (.) in the host address.
Attack Countered	In this type of attack, a client can send high volumes of DNS requests to a single host. In attempting to service all of those requests, system resources can be consumed, possibly creating a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum number of DNS requests to a host: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS Param Limits area, click Configure.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBD1: Rate of DNS Non-Recursive Requests for a Domain Has Been Exceeded (tln-101077)

Table 22-25 provides detailed information about rule tln-101077.

Table 22-25: RRBD1: Rate of DNS Non-Recursive Requests for a Domain Has Been Exceeded (tln-101077)

Rule Description	This rule counts the number of times that non-recursive requests are sent to a domain. When the count exceeds the specified limit over the specified period of time, the rule is triggered.
Attack Countered	If the number of non-recursive requests is in the thousands, system resources are consumed, and a denial-of-service condition can occur.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the minimum acceptable number of HTTP requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS Param Limits area, click Configure.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

AAUPV: DNS Request Exceeds Maximum Allowed Length in Bytes (tln-101078)

Table 22-26 provides detailed information about rule tln-101078.

NOTE

This rule is unusual for two reasons: it is not rate-based, as there is no time component, and it has no cost, so there is no credit deduction associated with it. Rather than deducing credits, the system responds to the triggering packet by acting on the rule's disposition.

Table 22-26: AAUPV: DNS Request Exceeds Maximum Allowed Length in Bytes (tln-101078)

Rule Description	This rule totals the arriving DNS request payload in a single flow and compares this value with a maximum length specified by the user.
Attack Countered	In this type of attack, packets of extended length, or repeated fast transmissions of the same packet, could be attack vectors. This rule can sometimes be used to block attacks in which the client is asking the DNS server to resolve numerous names, rather than just one.
Action on Rule Trigger	If the total DNS request length for a flow exceeds the specified limit, the rule is triggered, and the DNS request that exceeded the length is dropped. Note that there is no cost associated with this rule, so there is no credit deduction associated with it. If the rule's disposition is set to Drop or Reject, the request will be dropped.
Rule Configuration (Quick Overview)	To configure the maximum allowed DNS request length: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS Param Limits area, click Configure.
Applicable Limits	you can specify This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBD1: Rate of DNS Recursive Requests for a Domain Has Been Exceeded (tln-101079)

Table 22-27 provides detailed information about rule tln-101079.

Table 22-27: RRBD1: Rate of DNS Recursive Requests for a Domain Had Been Exceeded (tln-101079)

Rule Description	This rule counts the number of times that recursive requests are sent to a domain. When the count exceeds the specified limit over the specified period of time, the rule is triggered.
Attack Countered	Attacks can use spoofed recursive DNS requests. If recursive queries are permitted on the DNS server, the system will attempt to resolve the DNS queries by sending out responses to a target address. When the number of requests is in the thousands, the flood of DNS replies multiplies that volume, which can result in a denial-of-service condition.
Action on Rule Trigger	The system will count the number of times this specific event occurs over a user-specified time period. If this number exceeds the specified limit, the rule is triggered, and the specified cost is deducted from the credits available for that client.
Rule Configuration (Quick Overview)	To configure the maximum number of recursive DNS requests for a domain: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS Param Limits area, click Configure.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

RRBD3: DNS RCODE Matches Specified Filter (tln-101080 through tln-101095)

Table 22-28 provides detailed information about rule tln-101080 through tln-101095.

Table 22-28: RRBD3: DNS RCODE Matches Specified Filter (tln-101080 through tln-101095)

Rule Description	This rule counts receipt of a packet with a specific RCODE, and tracks which client IP address sent the traffic that triggered the response code. When any client IP address generates this RCODE from a server, the rule is triggered. Note that there is no rate information (time or limit specification) associated with this rule.
Attack Countered	Attackers may make DNS requests that cause a DNS server increased or excessive load, normally resulting in an error response from the server. To prevent this, the device monitors DNS RCODEs for requests. You can configure larger client request credit deductions for requests that result in a non successful (anything other than RCODE 0) server response to address this class of DNS attacks. If seen frequently, RCODEs that indicate “Host Not Found” or RCODEs that contain a format error may indicate an attack. DNS requests that cannot be resolved cost additional system resources.
Action on Rule Trigger	The system will watch for specified DNS RCODEs and, when they are found, note which client IP address triggered the RCODE. If that RCODE is seen (even once), the rule is triggered, and the specified cost is deducted from the credits available for the initiating client.
Rule Configuration (Quick Overview)	To configure the minimum acceptable number of DNS RCODE requests: 1. Click the Security Policies toolbar button. 2. On the Configure Security Policies dialog box, click Rate-Based Policies. 3. Click Servers, click the desired Host Group, then click Edit. 4. Click the DNS Limits tab. 5. In the DNS RCODE Limits area, click Configure.
Applicable Limits	This rule is only applicable to IPS 5100 and 5200 E-Series models.

Table 22-29 lists the Corero Network Device RCODE rules.

NOTE

Rate limiting only permits the specification of DNS RCODE values from 0 through 15. These values are used in rules TLN-101080 through 101095. RCODEs above 15 are not currently supported.

Table 22-29: DNS Response Code (RCODE) Rules

Rule Number	Rule Name	Description
tln-101080	RRBD3: DNS RCODE matches specified filter (RCODE 0)	RCODE 0 indicates “no error”. The server was able to interpret the query successfully.
tln-101081	RRBD3: DNS RCODE matches specified filter (RCODE 1)	RCODE 1 indicates a format error. The server was unable to interpret to the query due to a problem with how it was constructed.

Table 22-29: DNS Response Code (RCODE) Rules *(Continued)*

Rule Number	Rule Name	Description
tln-101082	RRBD3: DNS RCODE matches specified filter (RCODE 2)	RCODE 2 indicates a server failure. The name server was unable to process this query due to a problem with the name server.
tln-101083	RRBD3: DNS RCODE matches specified filter (RCODE 3)	RCODE 3 indicates a name error. The name specified in the query does not exist in the domain. This is meaningful only from an authoritative name server.
tln-101084	RRBD3: DNS RCODE matches specified filter (RCODE 4)	RCODE 4 indicates this query is not implemented. The name server does not support this type of query.
tln-101085	RRBD3: DNS RCODE matches specified filter (RCODE 5)	RCODE 5 indicates the query was refused. The server refused to process the query, most likely due to policy reasons.
tln-101086	RRBD3: DNS RCODE matches specified filter (RCODE 6)	RCODE 6 indicates a domain name exists although it should not.
tln-101087	RRBD3: DNS RCODE matches specified filter (RCODE 7)	RCODE 7 indicates that a resource record set exists although it should not.
tln-101088	RRBD3: DNS RCODE matches specified filter (RCODE 8)	RCODE 8 indicates that a resource record set does not exist, although it should.
tln-101089	RRBD3: DNS RCODE matches specified filter (RCODE 9)	RCODE 9 indicates that the server receiving the query is not authoritative for the specified zone.
tln-101090	RRBD3: DNS RCODE matches specified filter (RCODE 10)	RCODE 10 indicates that the name specified in the message is not within the specified zone.
tln-101091	RRBD3: DNS RCODE matches specified filter (RCODE 11)	RCODE 11 is not currently used.
tln-101092	RRBD3: DNS RCODE matches specified filter (RCODE 12)	RCODE 12 is not currently used.
tln-101093	RRBD3: DNS RCODE matches specified filter (RCODE 13)	RCODE 13 is not currently used.
tln-101094	RRBD3: DNS RCODE matches specified filter (RCODE 14)	RCODE 14 is not currently used.
tln-101095	RRBD3: DNS RCODE matches specified filter (RCODE 15)	RCODE 15 is not currently used.

Appendix A
IPS Unit System Management

The IPS Unit is designed to require little or no system management tasks for normal operation. But there may be times when you need to upgrade the software, reboot the system, or manage the unit's configuration files.

This chapter contains the following sections:

System License Management Key

Corero Network Devices now support the use of a system license key. Single-use system license keys can be applied to a device to manage trial evaluation operation and perform in-field model upgrades. The use of each individually generated system license key is restricted to a single Corero Network Device, and will not function on any other Corero unit.

A license key can be installed at any time. When you install a new license key, the restrictions on that key (or the lack thereof) will supersede any previous license key.

NOTE

Pre-existing Corero Network Devices do not require a key, but will accept the entry of one. They will continue to operate as they currently do unless a key with restrictions is applied.

A System License Key is used for one or more of the following purposes:

- To restrict the trial usage of a device by enforcing an expiration date.
- To restrict the trial usage of a device by restricting license application to a specific software version.
- To replace a trial license with a permanent one
- In limited cases, to change the model number of a device.

Trial License Expiration

A warning will display seven days before a system license expires, and will display in the status area until expiration occurs. Expiration occurs at midnight on the expiration date.

When a system license expires, the Corero Network Device will automatically enter Bypass mode. You will not be able to exit Bypass mode until a new license either with a later expiration date, or with no expiration date, has been entered.

CAUTION

If you are entering a duration-limited (trial) system license , ensure your system clock is accurate. If you need to modify the system clock, you must do so before license entry. Duration-limited system licenses are designed to be tamper-proof, and attempts to modify system time after the key is entered may cause the license to expire immediately, rendering the Corero Network Device unusable until a new license is acquired.

If a Corero Network Device is using a trial (time-limited) System License, when the license expires, the system will automatically enter Bypass mode. When locked in Bypass mode, you will be able to modify the bypass settings, but your changes will not take affect until you acquire a new system license key from Corero and apply it to the device.

Viewing System License Key Status

To view system license key status on a Corero Network Device:

1. If the Corero Network Device is using a trial system license, the gray status bar at the bottom of the main window will display trial license information including:
   - When the device is running in Trial mode (before the license expiration date).
   - When the license is approaching its expiration date.

- When the trial license has expired (which locks the device in bypass mode until another license is entered).

2. For additional system license status information, do one of the following:
   - Click the System Information toolbar button.
   - Choose Monitor System > System Info from the Navigation Tree.

   The System Information dialog box displays.

3. On the System Information dialog box, view the information in the License Key Status area. Status messages include:
   - Permanently Unlocked - The Corero Network Device is operating with a permanent system license, with no expiration date or restrictions.
   - Valid Until (expiration date) - The Corero Network Device is operating with a trial system license, and the expiration date is listed. Note that the expiration date is specified in GMT (UMT) time.
   - Locked In Bypass Mode - The Corero Network Device was operating with a trial system license, but the license has expired. The device will remain locked in bypass mode until a new license is applied.
   - Unknown - The Corero Network Device is currently unable to obtain information on the license status. If this problem persists, contact Corero.

Entering a System License Key

You can enter a system license key at any time, so long as you have saved all configuration changes for the Corero Network Device.

CAUTION

Before you enter a duration-limited (trial) system license key, ensure your system clock is accurate. If you need to modify the system clock, you must do so before license key entry. Duration-limited system license keys are designed to be tamper-proof, and attempts to modify system time after the key is entered may cause the license to expire immediately, rendering the Corero Network Device unusable until a new license is acquired.

To enter a system license key on a Corero Network Device:

1. Before you enter a system license key, ensure that the current system time is accurate.
2. Do one of the following:
   - Click the System Information toolbar button.
   - Choose Monitor System > System Info from the Navigation Tree.

   The System Information dialog box displays.

3. Click Enter System License Key. One of the following happens:
   - If you have unsaved configuration changes, a message box displays indicating you need to review and save your configuration before you can enter the key. Once you have saved your changes, click Enter System License Key again.
   - If you do not have any unsaved configuration changes, the Enter System License Key dialog box displays.

4. Enter the key in the License Key field, then click Apply Key.

   One of the following happens:

- After entering a key that specifies a different product model, if the key is accepted a message displays indicating that key application was successful, and informing you that this will cause a reboot of the device. Click Yes to proceed with the reboot. Once the system has rebooted, choose Monitor System > System Info to view the current product model in the System Information dialog box.
- If you have entered a key that does not specify a different product model and the key is accepted, the system displays a confirmation message.

Rebooting (Restarting) the IPS Unit

Restarting the IPS Unit is primarily done after a software upgrade. You can perform this restart by either:

- Using the management application.

or

- Pressing a button on the IPS Unit chassis.

The management application enables you to reboot the IPS Unit, but the hardware reset button allows you to do two things, depending on how long you depress it:

- If you briefly depress the reset button, the IPS Unit will reboot.
- If you depress the button continuously for five seconds or longer, you will reboot the IPS Unit and reset the IPS Unit to its factory configuration settings. For more information, see Resetting the IPS Unit to Factory Defaults (page A-6).

NOTE

A reboot operation will disconnect all user management application sessions to that device.

If you wish to reboot (restart) the IPS Unit, perform the following steps:

1. Verify that the system is using the desired bypass policy. The bypass policy controls whether the IPS Unit passes, or does not pass, traffic during system reset.

   To view or modify the bypass policy, choose Configure System > Management Access > IPS Controller Management > Settings from the Navigation Tree.

   For detailed information on bypass settings, see Bypass Settings (page 5-6).

2. Do one of the following:
   - If you want to reboot the system using the management application, from the Navigation Tree, choose Manage System > Reboot. If you have any unsaved changes, the system asks if you want to save your configuration before rebooting. Do one of the following:
     - To save your configuration changes and reboot, click Yes.
     - To reboot the IPS Unit without saving your configuration changes, click No.
     - To cancel both saving your changes and rebooting the IPS Unit, click Cancel.
   - If you want to reboot the system using the reset button on the IPS Unit chassis, locate and press the button. The Reset button on the IPS Unit is a small, recessed button that is typically located just to the left of the IPS Unit ports. For the exact location of the Reset button on a specific IPS Unit model, refer to the model-specific hardware guide. Because the Reset button is recessed, you cannot press it with your finger. You can depress the Reset button with the point of a pen.
3. The IPS Unit displays progress messages during and after the reboot process. One of the following happens:
   - If you have updated your IPS Unit's software and then rebooted, you must close the management application and reopen it.
   - If you have only rebooted the IPS Unit, when you confirm the reboot complete message, the management application Login window displays.

Resetting the IPS Unit to Factory Defaults

You can reset the IPS Unit's configuration files to their factory-default settings. This will reset most of the system configuration settings to those implemented during manufacturing.

Several of the parameters you established when you used the set up the system are not changed. They include:

- IP address, subnet mask, default route settings and contact information.
  If you need to change this information, refer to the relevant Corero Hardware Installation guide's installation chapter.
- Time zone, date, and time settings.
  For information on how to change date and time settings, see Chapter 4, ''Initial IPS Unit Configuration Tasks".

To reset the IPS Unit to its factory default configuration:

1. From the Navigation Tree, select Manage System > Reset to Factory Defaults.
2. You are prompted to confirm your selection. Click Yes to return to the factory settings
3. Once the default configuration is loaded, the IPS Unit prompts you to apply it. Click the Apply button in the Policy Update section of the management application.

About Configuration Files

The IPS Unit stores its configuration in a set of configuration files. Each file contains specific configuration values for a single subsystem within the IPS Unit. Taken as a set, the files provide a snapshot of the entire IPS Unit configuration.

Configuration for the IPS Unit is partitioned into subsystems. All the settings necessary for a subsystem are stored in that subsystem's configuration file and can be loaded independently from the setting in another subsystem's configuration file.

The IPS Unit saves the last three versions of each configuration file, including the one that is currently active. Older versions are automatically deleted. Each configuration file is displayed with its timestamp, so you can identify configuration files created on a particular date.

The naming convention for a configuration files consists of three parts:

- Two characters that identify the type of configuration file, for example EV which identifies the file containing configuration data for the Event Logging System.
- A six-digit hexadecimal serial number that uniquely identifies the instance of that configuration file; the serial number is incremented by one each time the file is saved. The last three saved versions of the file are actually saved in the IPS Unit's nonvolatile memory.
- The standard text file name extension (txt). A configuration file is a standard ASCII text file and can be read by most ASCII file editors.

NOTE

In addition to these configuration files, there are system configuration files for advanced protocol and data file configuration settings, which are predefined and updated via TopResponse protection packs. These advanced configuration files have the pvc file extension, and are not accessible to the user via the GUI.

Table A-1lists the configuration file naming conventions.

Table A-1. Configuring File Names

File Name	Subsystem	Description
DE######.txt	Device	Device configuration details such as bridge settings, port settings, NTP, SSL, and so forth.
EC######.txt	Environment	Environment configuration information, such as services, host groups, ranges, etc
EV######.txt	Event	Event configuration information including event message settings, message groups, and other Event Logging System information.
HA######.txt	High Availability	Configuration information for the High Availability feature.
LP######.txt	LAN Port	LAN port configuration settings including static MAC addresses, bridging settings, and so forth.
MC######.txt	Management Configuration	Management configuration parameters for management users, Syslog server profiles, NTC server identification, and so forth.
PO######.txt	Policy Rules and Rule Sets	Security policy settings such as policy rules and rule-sets.

Table A-1. Configuring File Names *(Continued)*

File Name	Subsystem	Description
RB######.txt	Rules	User-added traffic security rules and information on content signatures. Note that factory-defined rules are defined in a separate file, which cannot be modified by users.
SE######.txt	Security	Security management information such as user and group access configurations, and remote management service access controls. Security information at both the individual and the group level is protected. You cannot view sensitive information such as passwords, nor can you change this information by editing the configuration file. If you manually edit the file, the hashed value for the edited file will not match the value the IPS Unit expects.
SI######.txt	Signatures	Contains security signatures and signature configuration parameters.

Managing Configuration Files

The management application enables you to view, save, activate, download, and upload configuration files.

The IPS Unit loads and activates a saved configuration under the following conditions:

- All Active Configuration Files — When the IPS Unit reboots.
- A Single Subsystem File — When you select and activate the file.
- A Single Subsystem File — When you upload a single file to the IPS Unit from your management station, such as a protection pack, the file is automatically loaded and activated.

NOTE

The IPS Unit does not need to be rebooted when you change its configuration or activate a different configuration file unless you are modifying its high availability settings.

The Management application enables you to store copies of configuration files on your management station. Reasons why you might want to save configuration files to your PC include:

- To Preserve Your IPS Unit's Configuration when Upgrading the IPS Unit's Software
  The upgrade process should preserve your configuration, but to ensure your configuration files are preserved, it is safer to store a copy of each configuration file on your management station before upgrading the IPS Unit's software.
- To Preserve a Golden or Snapshot Configuration
  Since the GUI only preserves the last three versions of any given configuration file, if you have a specific configuration that you want to preserve for future reference or use, you must download it to your management PC.

Corero recommends that you back up your system's configuration every time you make major changes to it. In order to back up configuration information for all of your Corero products, you can run a backup script on the IPS Controller, which backs up both IPS Controller configuration information, and the corresponding information for all Corero Network Devices (IPS and DDS Units) managed by the IPS Controller. This process is described in the *IPS Controller Release Notes*.

Alternatively, if you do not have an IPS Controller, from the IPS Unit, you can download diagnostic information, as described in Downloading Diagnostic Information (page A-11), and export the current shun information, as described in Viewing and Managing Shunned Addresses (page 19-9).

To manage configuration files:

1. Choose Manage System > Configuration Files from the Navigation Tree to view the list of currently stored configuration files. The Configuration Files dialog box displays (Figure A-1). This page shows you:
   - Which configuration files are currently in use (active).
   - The timestamp when a configuration file was last saved.
   - The user who created or last saved the file.

Figure A-1. Configuration Files

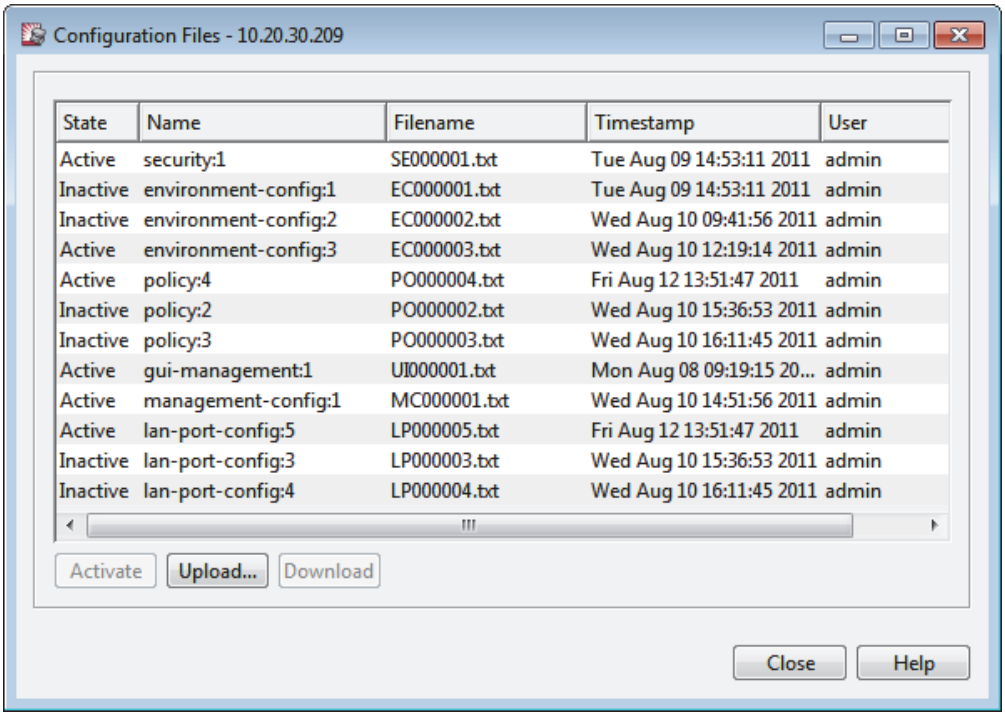

Figure A-1 shows three saved versions of the policy subsystem's configuration file (policy:2, policy:3, and policy:4). Policy:4 is currently the active policy configuration file, and its parameters are the ones the IPS Unit is currently using.

2. To view a configuration file:
   a. Select the file from the list of files and click Download. The file displays in your browser's window.
   b. To save the file to your management station, choose File > Save As from the browser window's menu bar.
3. To upload a previously saved file to your IPS Unit.
   a. On the Configuration Files window, click Upload.
   b. Browse to the file location and select the file, then click Upload.
4. You can instantly revert to a previous configuration file by activating a configuration file. To activate a configuration file (to select and activate a file for current use):
   a. Select the inactive file you want to use from the list of configuration files.
   b. Click Activate.

      The configuration file's contents are activated, and the previous configuration file changes state to Inactive.
5. The IPS Unit implements the configuration changes you make as soon as you make them (except for policy updates, which must be applied). However, it does not automatically save your configuration changes to its nonvolatile memory. You must do this manually.
   a. Whenever you make changes to the IPS Unit's configuration, the Graphical User Interface (GUI) indicates that you have unsaved changes.
   b. To save your configuration changes, click the "Save" Toolbar button ( ). When you click the Save icon, only the configuration files associated with your current changes are saved.

Downloading Diagnostic Information

If you are having problems with your system, Corero Network Security support personnel may ask you to download diagnostic information for forensic use.

To download diagnostic information:

1. From the Navigation Tree, choose Monitor System > Download Diagnostic Information. The Download Diagnostic Information dialog box displays.
2. Choose one of the following:
   a. All Diagnostic Information: Select this option to compile all of your IPS Unit's diagnostic information into a single file.

      Note that this file may take several minutes to generate.

   b. Specify Diagnostic Filename: Select this option if you want to specify the name of a file stored in the IPS Unit's non-volatile memory to download. Options include:
      - alerts.log
      - core.log
      - events.log
      - system.log
3. When finished, click OK.
4. Depending on your selection, do one of the following:
   - If you chose to download All Diagnostic Information, the system prompts you to specify a custom file name and location if you do not want to use the defaults.
   - If you chose to download a specific file, the file displays in your browser window. Use the browser's Save option from the File menu if desired.

Managing the IPS Unit's Software

When new software becomes available, you will want to upgrade the software on the IPS Unit.

CAUTION

If your IPS Unit is under IPS Controller management, you must upgrade the IPS software through the IPS Controller. For detailed instructions, see the *IPS Controller Administrator's Guide*.

Since the IPS Unit can store multiple versions of software in its non-volatile memory, the Management application enables you to choose and activate whichever resident software version you want the IPS Unit to use.

NOTE

Corero recommends that you regularly store 1 or 2 software versions on your IPS Unit, with a temporary maximum of 3.

To manage the IPS Unit software:

1. From the Navigation Tree, choose Manage System > Manage Software. The Manage Software dialog box displays, showing a table listing the current software stored in the IPS Unit's non-volatile memory. For each software version, the following information is displayed:
   - Name — The software release version name.
   - Status — Indicates whether a given version of the software is currently being used by the IPS Unit (active) or not (inactive).
2. To upload and activate a new version of the software:
   a. Download or copy the software upgrade package to your local file system (accessible by the management station).
   b. On the Manage Software dialog box, click Upload.
   c. Browse to and select the software upgrade package.

      NOTE

      If you receive a Java-based security message during this process, click Yes to continue.

   d. Click Upload. The IPS Unit uploads the software package to its non-volatile memory, and makes it available in the Manage Software dialog box.
   e. Once the new software displays in the Manage Software dialog box, select the software and click Activate.
   f. The IPS management application informs you that it will reboot the IPS Unit to complete the activation process.

      CAUTION

      For proper operation, after the reboot, you must close and reopen the IPS management application.

3. Once you have installed a new version of the software, check to see how many software releases are being stored on the IPS Unit. You can view the number of versions on the Manage Software dialog box.

NOTE

It is important to note that the IPS Unit stores software versions differently from the way configuration files are stored. The IPS Unit saves the three most recent configuration files and automatically deletes any older ones, but the IPS Unit never automatically deletes a software version. You should keep a maximum of two releases on the IPS Unit. Storing more than two releases may reduce operational drive space.

To delete older releases from the IPS Unit, select the older release you want to remove, then click Delete.

Index

D

E

F

G

H

T

U

V

W