corero
NETWORK SECURITY

IPS CONTROLLER

Administrator’s Guide

Corero Network Security, Inc.
990-0244-00

Legal Information

Contents

Figures

Tables

Preface

This guide contains information about using the Corero IPS Controller (IPSC) software and management application. The IPS Controller is used to manage Corero Network Devices, including IPS 5500 and the DDS 5500 Units.

NOTE

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 managing network security equipment. It assumes the reader has a high level understanding of network operations, and familiarity with Corero Network Devices.

Revision Information

This is a new book.

Related Books

This guide is part of the Corero IPS Controller documentation set, which includes the following:

Documentation	Description
Corero IPS Controller Release Notes	Information on known problems, bug fixes, and technical tips.
Corero IPS Controller Administrator’s Guide (This Guide)	Conceptual and Procedural information for configuring and managing the integration of the Corero Unit into your network. The guide describes network and port role settings, bridging, system setup, and configuration, management, and monitoring of traffic security features.
IPS Controller Online Help	Available through the IPS or DDS management application, the online help system provides detailed descriptions of configuration parameters, procedures, and notes regarding use of the IPS Controller product features.

Accessing the IPS Controller Documentation

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

You access the online help system from within the IPS Controller management application 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 management application.

- To view context-sensitive help which pertains to the management application dialog box you are currently viewing, click the Help button on the dialog box.

Accessing Information on the CD-ROM

Each IPS Controller product includes a documentation CD-ROM. This CD-ROM includes both documentation and product software. 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 component 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.

Corero Product License and Warranty Information

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

Corero Services and Support

Corero Network Security offers two options for contacting Customer Services and 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. Glossary - list of terms used w/in the portal, customer feedback, customer-specific history, profile, and service requests track customer contact and support history.

When you contact Customer Services and Support 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 contacting Customer Services and 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 to our Customer Service Department, so that we can continue to improve our products.

Please send your comments and suggestions to the following address:

Customer Service Department
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

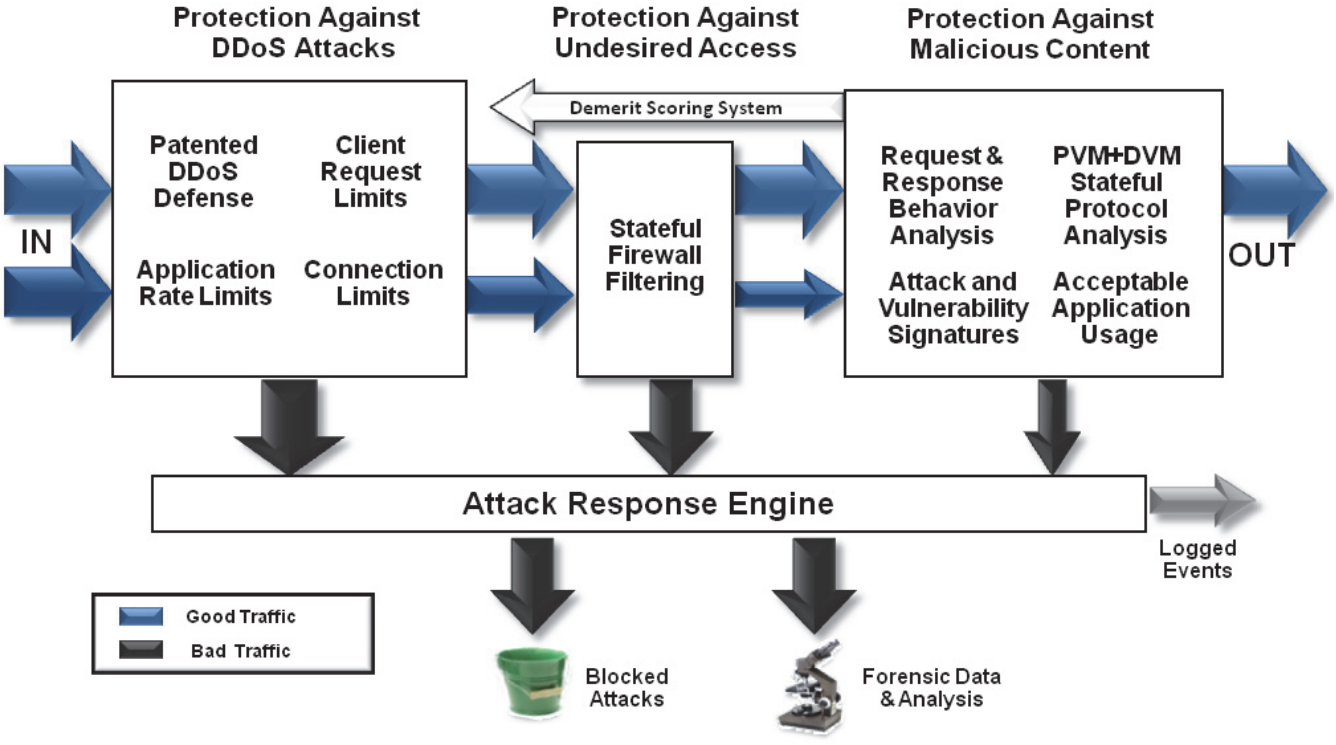
Protection Against DDoS Attacks
Protection Against Undesired Access
Protection Against Malicious Content
Demerit Scoring System
IN
Patented DDoS Defense
Client Request Limits
Application Rate Limits
Connection Limits
Stateful Firewall Filtering
Request & Response Behavior Analysis
PVM+DVM Stateful Protocol Analysis
Attack and Vulnerability Signatures
Acceptable Application Usage
OUT
Attack Response Engine
Logged Events
Good Traffic
Bad Traffic
Blocked Attacks
Forensic Data & Analysis

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.

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

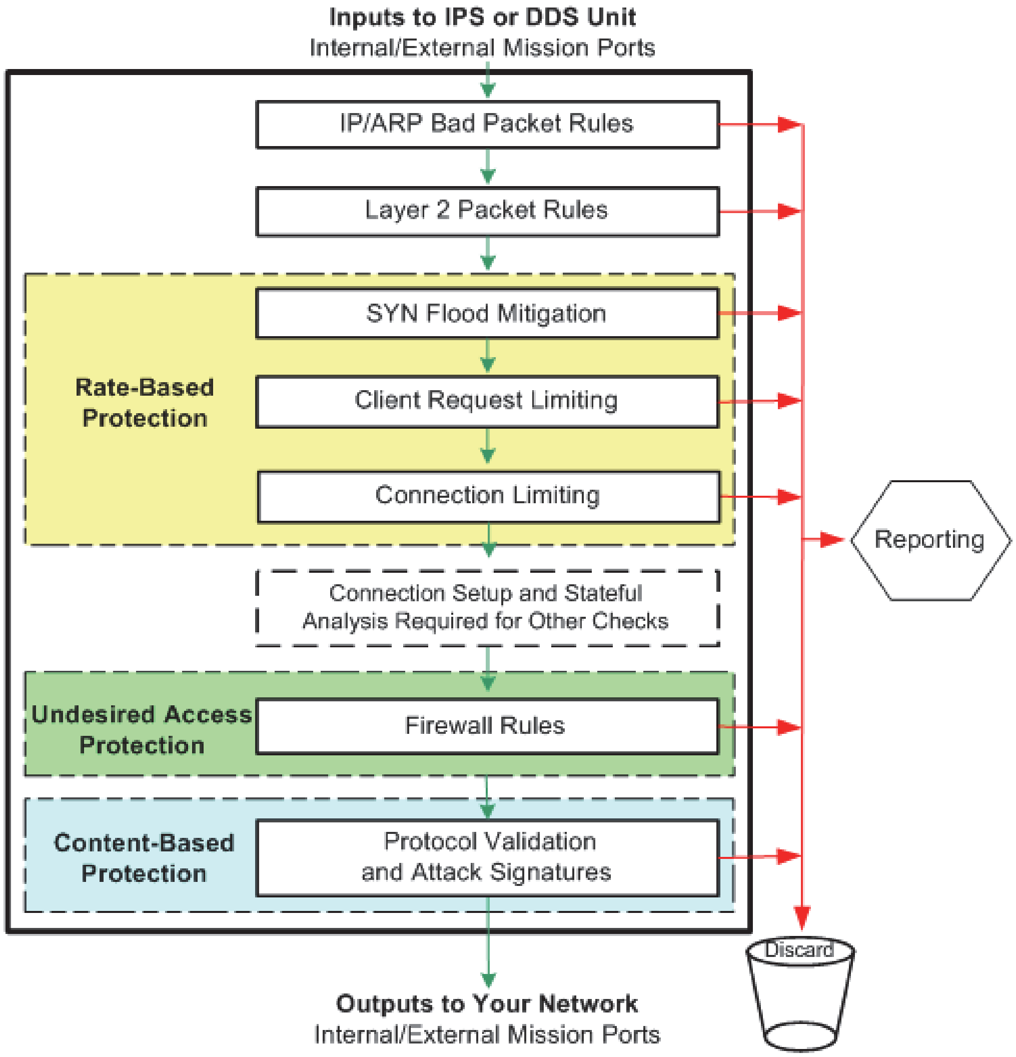
Inputs to IPS or DDS Unit
Internal/External Mission Ports
IP/ARP Bad Packet Rules
Layer 2 Packet Rules
Rate-Based Protection
SYN Flood Mitigation
Client Request Limiting
Connection Limiting
Connection Setup and Stateful Analysis Required for Other Checks
Undesired Access Protection
Firewall Rules
Content-Based Protection
Protocol Validation and Attack Signatures
Reporting
Discard
Outputs to Your Network
Internal/External Mission Ports

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

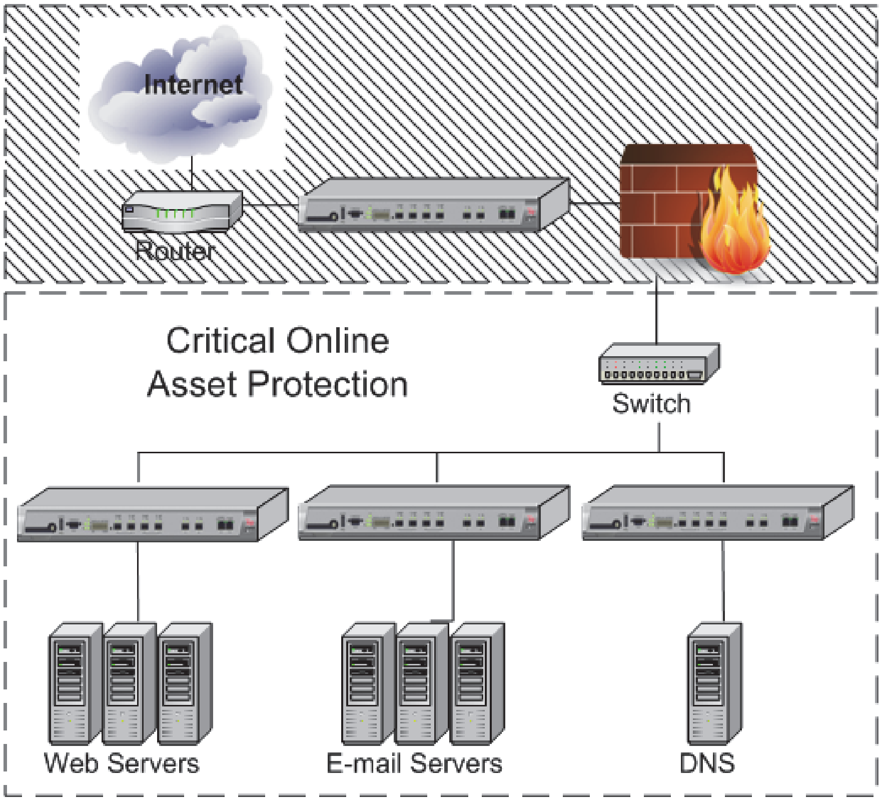

High Volume Configuration (Single Inline with Peer)

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

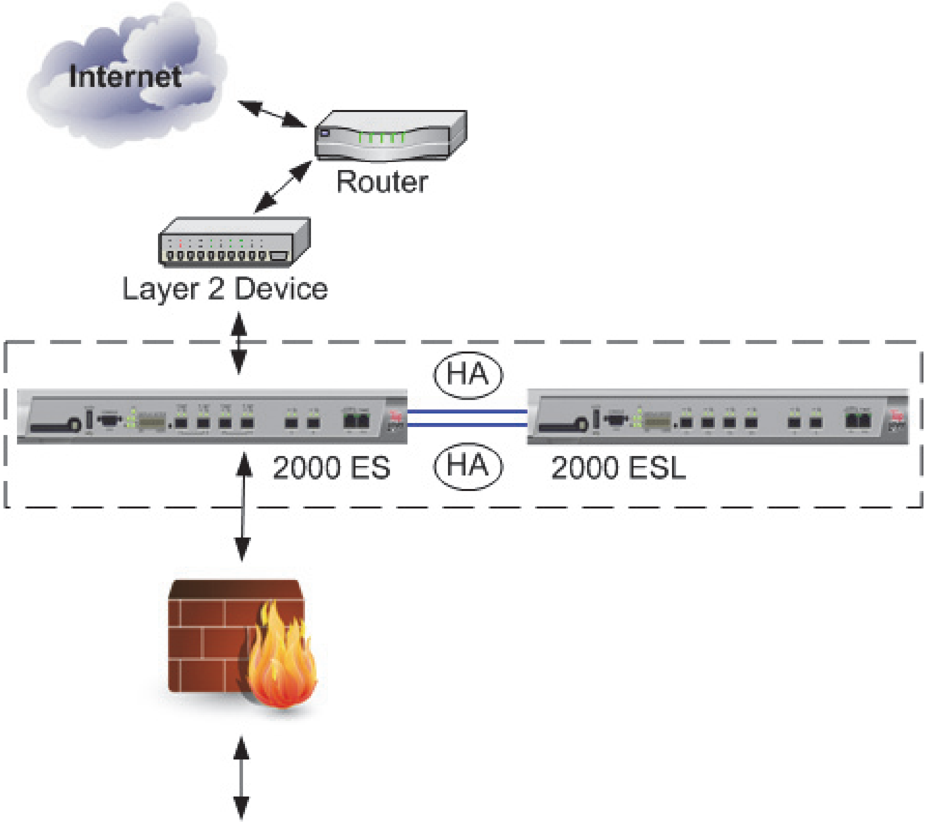

ProtectionCluster Configuration

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

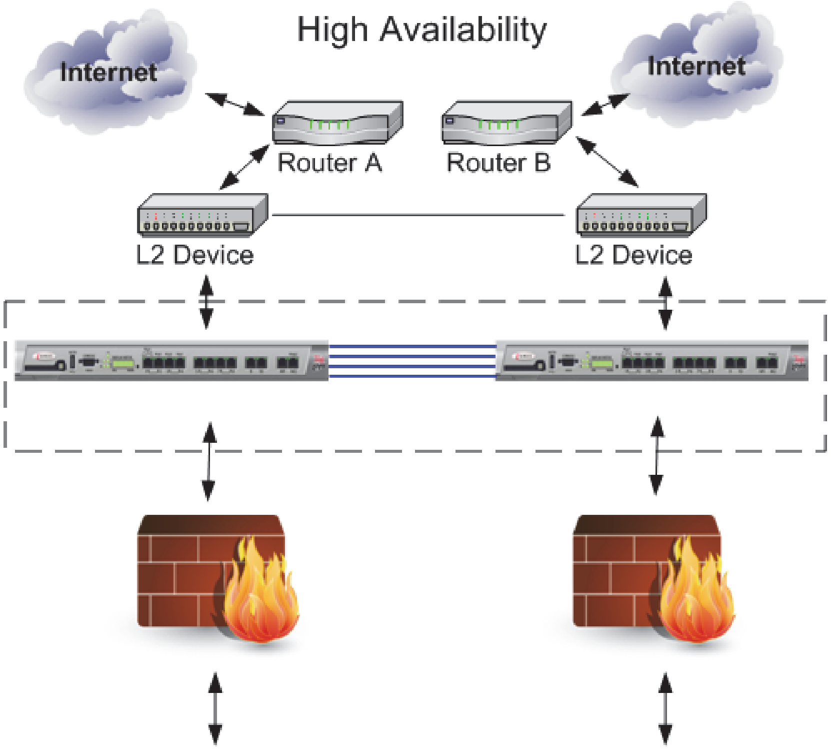

A ProtectionCluster refers to a network configuration option that provides higher bandwidth and redundancy. This configuration connects multiple Corero Network Devices together. On the 5100-Series Corero Network Devices, this is done by using up to four of the 10/100/1000 ports that can be configured 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

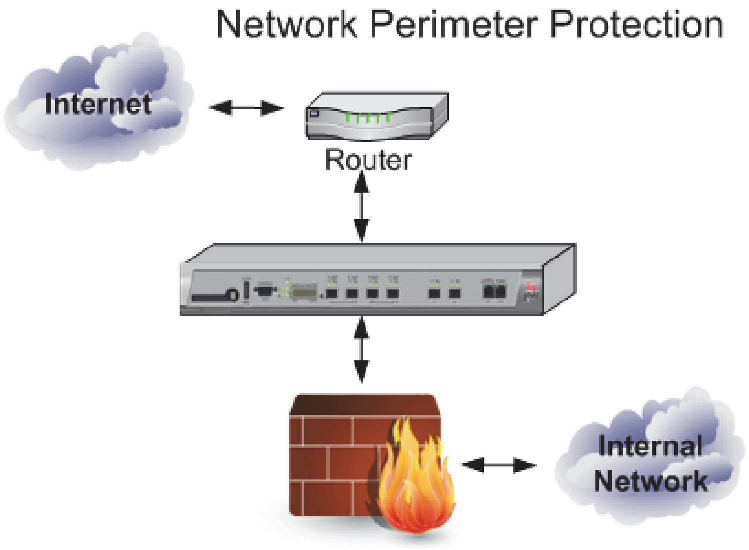

Protect Your Hosting Center

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

Figure 1-7: Protecting a Hosting Center

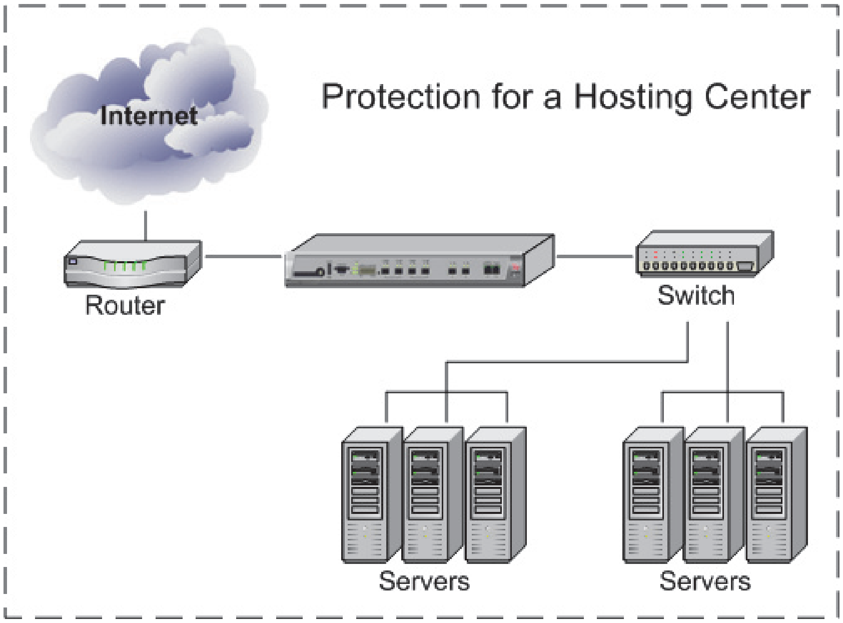

Protect Servers in Your Enterprise

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

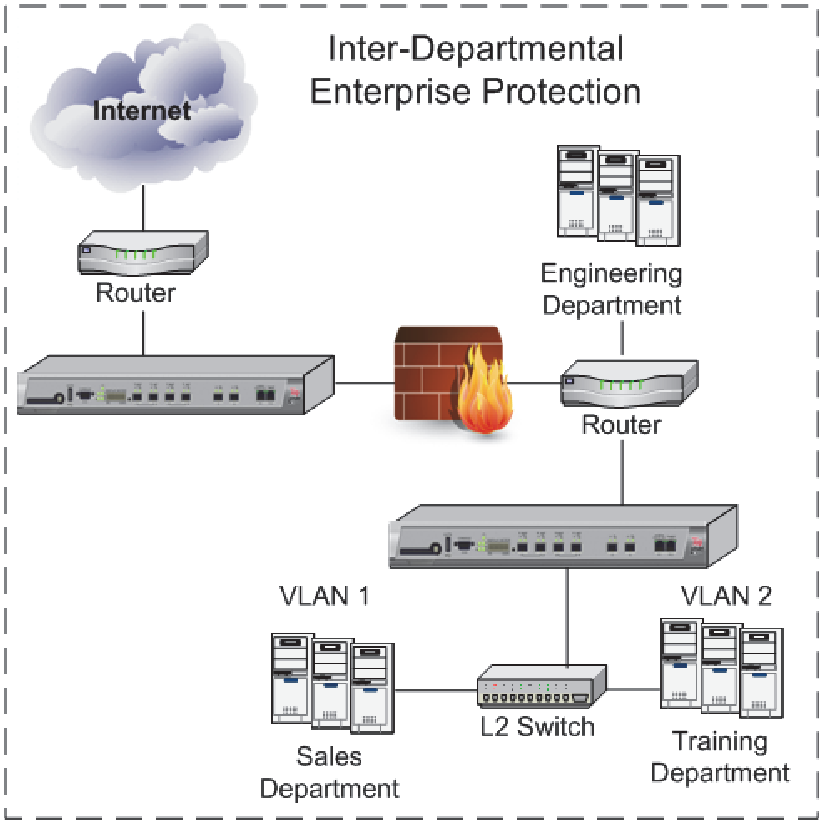

Chapter 2
IPS Controller Overview

Corero's IPS 5500 family of network intrusion prevention systems is targeted for medium to large enterprises, service providers, education providers, financial institutions, and governments. Corero's products act quickly to mitigate the effects of an attack, allowing customers and employees to access services even while an attack is happening. A critical requirement of these multiple-device deployments is the ability to configure and manage devices from a central location. The IPS Controller provides a centralized management solution for Corero Network Devices. This chapter describes the features of the IP Controller.

This chapter contains the following sections:

IPS Controller Product Overview

The IPS Controller is a software product that allows centralized Corero Network Device management, multi-unit real-time incident response, and automated TopResponse Protection Pack updates. An IPS Controller system is a computer system which runs the Linux operating system and has the IPS Controller software installed on it. The IPS Controller software is a Corero software product that is installed on a Linux computer system.

NOTE

For detailed information on system requirements, see Chapter 3, "IPS Controller Pre-Installation Requirements".

The IPS Controller maintains secure connections to the Corero Network Devices (IPS and DDS Units) that it manages. You can manage both the Corero Network Devices and the IPS Controller itself using a Java Web Start management application that works in a similar fashion to the IPS 5500 and DDS 5500 local device management applications.

The IPS Controller enables you to manually or automatically download Corero TopResponse protection pack updates that contain protection against new security threats, technical support, and security advisories. The TopResponse server is maintained by Corero. The IPS Controller must be able to connect to all IPS and DDS Units within your network in order to ensure the IPS and DDS Units are configured with the latest security updates.

NOTE

In order to use the TopResponse features of the IPS Controller, your IPS Controller system must have a valid TopResponse key and internet access.

The IPS Controller enables you to do the following for managed devices.

- View a high-level dashboard summary that provides real-time, at-a-glance status.
- Create User-defined groupings (called policy groups) and centrally manage them.
- Acquire and distribute operating software and threat update and protection packs.
- View, manage, and modify device configurations
- View, manage, modify, validate, and distribute security policies.
- View aggregated statistics and generate aggregated data reports.

Corero Network Device Features

The IPS Controller's purpose is to manage Corero Network Devices, including both IPS Units and DDS Units. These units provide an integrated approach to intrusion prevention. Leveraging Corero's patented and award-winning DDoS defense technology, Corero Network Devices are specifically focused and highly optimized to identify and deflect attacks while allowing genuine traffic to pass with minimum disruption. These devices provide maximum protection for critical IT assets while allowing full access to legitimate users and applications.

The rest of this section describes:

- Protocol Anomaly Detection (page 2-3)
- Data File Inspection (page 2-3)
- Acceptable Application Usage (page 2-3)
- Signature Matching (page 2-3)
- Real-time Shunning (page 2-4)
- Robust Protection (page 2-4)
- Attack Mitigation (page 2-4)
- Deep Packet Inspection (page 2-5)

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 each Corero network device.

Real-time Shunning

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 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. 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 Corero Network Device family consists of IPS and DDS Units with the performance and capacity to handle throughputs from 100Mbit/sec to 8Gbit/sec, with transaction rates up to 40,000 sessions per second.

Attack Mitigation

Corero Network Devices offer 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.

Devices provide 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.

Corero Network Devices provide 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

Corero Network Devices provide 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

Corero Network Devices ensure 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

Corero Network Devices are 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, these devices protect 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

Chapter 3
IPS Controller Pre-Installation Requirements

This chapter describes the system requirements and software prerequisites for installing the IPS Controller software on a dedicated system.

This chapter contains the following sections:

IPS Controller System Requirements

You can choose to install the IPS Controller software either on a dedicated computer system, or on a VMware system.

System requirements differ depending on how many IPS or DDS Units the IPS Controller is managing. Management of a maximum of 64 IPS or DDS Units is supported. Note that each IPS or DDS Unit counts as one unit toward the 64-unit limit, regardless of whether or not it belongs to a ProtectionCluster.

The requirements for these separate options are described in the following sections:

- Hardware Requirements (page 3-2)
- Virtual Machine Requirements (page 3-2)

Hardware Requirements

The hardware requirements for the IPS Controller vary depending upon the number of IPS or DDS Units under management.

NOTE

Corero recommends that, if you intend to manage eight or more Corero Network Devices, you install the IPS Controller software on a dedicated computer system. Managing a large number of devices may require additional system resources.

Table 3-1 describes the requirements for a standard computer system on which IPS Controller software is run.

Table 3-1: IPS Controller Computer System Requirements

Requirement	Managing up to 8 IPS or DDS Units	Managing More Than 8 IPS or DDS Units
CPU	Single 2.8GHz Pentium 4 or greater	Dual 3.2 GHz Xeon or greater
RAM	2 GB (dedicated)	4 GB (dedicated)
Free Disk Space	200 GB of available hard disk space	200 GB of available hard disk space
Network Connection	1	1
Optional		Redundant Power and RAID as specified by customer application re.quirements for network management solutions

Virtual Machine Requirements

The IPS Controller is supported on a virtual machine. Note that there is minimal performance impact to the hosting server when running the IPS Controller software on a virtual machine.

Table 3-2 describes the requirements when running an IPS Controller on a virtual machine.

Table 3-2: IPS Controller Virtual Machine Requirements

Requirement	Managing up to 8 IPS or DDS Units	Managing More Than 8 IPS or DDS Units
System	VMware Workstation	VMware ESX Server
CPU	Dual 3.2 GHz Xeon or greater	Dual 3.2 GHz Xeon or greater
RAM	10 GB	10 GB

Table 3-2: IPS Controller Virtual Machine Requirements *(Continued)*

Requirement	Managing up to 8 IPS or DDS Units	Managing More Than 8 IPS or DDS Units
Disk Space	230 GB	230 GB
Drive	CD or DVD	CD or DVD
Network Connections	2	2

IPS Controller Software Requirements

When you purchase the IPS Controller software, you must install the software on a Linux System that meets the following requirements.

Operating System Requirements

The following Linux operating systems are supported by the IPS Controller software, for 32 bit machines only:

- Red Hat Enterprise
  - Release 6
  - Release 5.3
- CentOS
  - Release 6
  - Release 5.3
- Fedora
  - Release 15
  - Release 10
- Fedora Core 6

For instructions on how to install required Linux software, see Preparing the Linux System for IPS Controller Installation (page 3-6).

Management Application Java Requirements

The IPS Controller management application runs as a stand-alone Java Web Start application which is used to configure the IPS Controller software. Refer to the *IPS Controller Release Notes* for detailed information.

IPS Controller Network Requirements

The IPS Controller software uses several dedicated ports for communication. These ports are described in Table 3-3.

NOTE

For information on installation and cabling of your Corero Network Device, see the device-specific Hardware Installation Guide.

Table 3-3: IPS Controller Port Requirements

Port Number	Use	Considerations
TCP/2616 (Fixed)	Used to securely communicate with IPS or DDS Units.	Make sure that there is no network device such as a firewall blocking this port.
TCP/2080 (Default)	Used for HTTP communication with the IPS Controller web-based management application.	You can configure a different port, but once entered during initial configuration, the port is then fixed.
TCP/2443 (Default)	Used for HTTPS communication with the IPS Controller web-based management application.	You can configure a different port, but once entered during initial configuration, the port is then fixed. Corero recommends that you use HTTPS to communicate with the IPS Controller.
TCP/443 (Fixed)	Used for SSL communication with the TopResponse server on the Internet.	Make sure that there is no network device such as a firewall blocking this port. Optionally, the IPS Controller can be configured to use a local proxy server with a different port if necessary.
TCP/1812 (Fixed)	Used for radius authentication.	This is the industry-standard Radius communication port.

Preparing the Linux System for IPS Controller Installation

Before installing the IPS Controller software on a Linux machine, you must first ensure the Linux system has been prepared by verifying that software requirements have been met or, if they have not, by meeting them.

Corero suggests that you install yum (Yellowdog Updater, Modified) , a Linux software installation and management utility, on the IPS Controller system. Whether this utility is available on your system depends on the type of Linux you use.

- For Fedora, yum comes preinstalled.
- For Red Hat, yum is generally not preinstalled and may be difficult to find. If this is the case, either use rpm to install the binary file, or compile the source file(s) using gcc.

> NOTE
>
> The following procedure assumes that yum is being used for software installation and management. If you are using a different utility, you will need to modify the commands to conform to your installation utility.

To prepare your system for IPS Controller installation:

1. Verify that sharutils is installed. To see whether it is installed, run the following command as root:

   ```
   [root@your system]# which uuencode
   ```

   If it is installed, you will see the following:

   ```
   /usr/local/bin/uuencode
   ```

   If sharutils is not installed, **do one of the following**. Options are listed in order of increasing difficulty.

   - Run the following command: `yum install sharutils`
   - Locate and download the sharutils rpm from the installation disk or the Internet. Move or copy the binary rpm file into a local directory, then run the following command: `rpm -ivh sharutils`
   - Locate and download the sharutils source code. The steps to compile the code are documented in the INSTALL script included within the source code for sharutils.

     Ensure you read the list of dependencies beforehand and ensure that the required components are installed in the correct order. Then compile with gcc, issuing the following commands while in the source code directory:

     ```
     configure
     make
     makefile
     ```

2. Verify that cvs is installed. To see whether it is installed, run the following command as root:

   ```
   [root@your ips system ]# which cvs
   ```

   If cvs is installed, you will see the following:

   ```
   /usr/bin/cvs
   ```

   If cvs is not installed, **do one of the following**. Options are listed in order of increasing difficulty.

   - Run the following command: `yum install cvs`
   - Locate and download the cvs rpm from the installation disk or the Internet. Move or copy the binary rpm file into a local directory, then run the following command: `rpm -ivh cvs`

- Locate and download the cvs source code. The steps to compile the code are documented in the INSTALL script included within the source code for cvs.

  Ensure you read the list of dependencies beforehand and ensure that the required components are installed in the correct order. Then compile with gcc issuing the following commands while in the source code directory:

  ```
  configure
  make
  make install
  ```

3. Verify that crontabs is installed. To see whether it is installed, run the following command as root:

   ```
   [root@your ips system]# which crond
   ```

   If crontabs is installed, you will see the following:

   ```
   /usr/sbin/crond
   ```

   If crontabs is not installed, **do one of the following**. Options are listed in order of increasing difficulty.

- Run the following command: `yum install crontabs`
- Locate and download the crontabs rpm from the installation disk or the Internet. Move or copy the binary rpm file into a local directory, then run the appropriate command.

  The command is **either** `rpm -ivh {crontabs}` **or** specific crontab rpm name.

- Locate and download the crontabs source code. The steps to compile are documented in the INSTALL script included within the source code for crontabs.

  Ensure you read the list of dependencies beforehand and ensure that the required components are installed in the correct order. Then compile with gcc issuing the following commands while in the source code directory:

  ```
  configure
  make
  make install
  ```

4. To avoid any conflict over TCP port 80, ensure that the Apache web server is either not installed or the service is not running on the system.

Chapter 4
Installing and Upgrading the IPS Controller Software

This chapter describes how to install and upgrade the IPS Controller software.

CAUTION

Before you install, the IPS Controller software, ensure the system on which you want to install the software meets all of the requirements in Chapter 3, “IPS Controller Pre-Installation Requirements”. Failure to do so may result in an unsuccessful installation.

This chapter contains the following sections:

- IPS Controller Software Upgrade Considerations (page 4-2)
- Installing or Upgrading the IPS Controller Software from the Command Line (page 4-4)
- Migrating IPS Controller Operations From One Server to Another (page 4-6)
- Uninstalling the IPS Controller Software (page 4-7)

IPS Controller Software Upgrade Considerations

When new versions of the IPS Controller software are released, you will need to upgrade the IPS Controller software.

NOTES

1. When a new version of the IPS Controller software is released at the same time as a new version of the IPS or DDS software, you should always upgrade the IPS Controller before you upgrade any IPS or DDS Units to their new software version.
2. To install the new IPS Controller software onto a different server than the one currently running the IPS Controller software, you will need to migrate the software. For detailed instructions, see IPS Controller Software Upgrade Considerations (page 4-2).

There are different procedures for upgrading the IPS controller software. Which one you use will depend on what version of the IPS Controller software you are currently running.

- Upgrading Process From a Previous 4.x Release (page 4-2)
- Upgrading From a 3.20 or 3.10 Release (page 4-2)

Upgrading Process From a Previous 4.x Release

To upgrade an existing IPS Controller from a previous 4.x release to the current release:

1. For safety, backup the IPS Controller, then copy the backup file to a safe location, preferably on a different system. This will enable you to restore the current system should you encounter unexpected difficulties while upgrading. For detailed instructions, see Backing Up and Restoring the IPS Controller (page A-5).
2. Install the new IPS Controller software by following the standard installation instructions, which will upgrade the IPS Controller software while preserving your existing configuration information. For detailed upgrade instructions, see Installing or Upgrading the IPS Controller Software from the Command Line (page 4-4).
3. Once you have successfully completed the upgrade process, back up the IPS Controller, then copy the backup file to a safe location.

Upgrading From a 3.20 or 3.10 Release

Because of differences between this release and earlier 3.20 and 3.10 releases, configuration information from these earlier releases cannot be moved forward to the current release. The following steps are recommended to upgrade from a release prior to 4.10 to this current release.

1. Review the requirements listed in Chapter 3, “IPS Controller Pre-Installation Requirements” to ensure that the system on which you intend to install this version of the IPS Controller meets these requirements.

   If your system does not meet the necessary requirements, modify the chosen system or obtain a different system that does.

2. Back up all the IPS Controller information using the procedure described in Backing Up and Restoring the IPS Controller (page A-5). Copy the backup information to a separate location. Should you need to downgrade from the new version of the IPS Controller to the previous version, you will now be able to restore your configuration.
3. Log in as root.
4. Stop the IPS Controller process with the following command:

   ```
   /sbin/service tlnipscd stop
   ```

5. Uninstall the IPS Controller software from your IPS Controller system as described in Installing an SSL Certificate on the IPS Controller (page 6-2).
6. Install the new version of the IPS Controller software as described in Installing or Upgrading the IPS Controller Software from the Command Line (page 4-4).
7. Once you have successfully installed the new software version, back up the IPS Controller, then copy the backup file to a safe location.
8. If your 3.20 or 3.10 configuration included one or more ProtectionClusters, you must import these ProtectionClusters into the new IPS Controller configuration. To do so, see Importing ProtectionClusters (page 14-19).

Installing or Upgrading the IPS Controller Software from the Command Line

The IPS Controller software installation kit can be used for initial installation, as well as to upgrade and downgrade the software on the IPS Controller system.

> NOTE
>
> Before initial installation, verify that the system you wish to use as an IPS Controller meets Corero's system requirements. These requirements are described in Chapter 3, "IPS Controller Pre-Installation Requirements".

To install the IPS Controller Software:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Verify that the Linux machine meets all software requirements for the version of the IPS Controller software you wish to install. Software requirements are described in Chapter 3, "IPS Controller Pre-Installation Requirements".
3. If you are upgrading the system, Corero recommends that you back up your IPS Controller configuration before you upgrade. For detailed instructions, refer to Backing Up and Restoring the IPS Controller (page A-5).
4. Copy the software to the Linux machine and modify the permissions on the installation script, then install the software.

   The name of the installation script varies, and is based on the release number of the software you are installing. In the example below, the version number is 4.70.049. If you are installing a different release, replace this information with the version-appropriate values.

   ```
   [root@your-system linux]# ftp myhost.mycompany.com
   ftp> get ipscd-V470049.scr
   ftp> quit
   [root@your-system linux]# chmod +x ipscd-V470049.scr
   [root@your-system linux]#./ipscd-V470049.scr
   ```

   The software installation proceeds.

5. Install the TopResponse license key onto the server, replacing the representative key below with your actual license key. If you do not have a TopResponse license key, contact Corero Customer Services for assistance.

   ```
   /sbin/service tlnipscd topkey 1111-2222-3333-4444-5555-6666-7777-8888
   ```

6. Restart the IPS Controller service on the new server to activate the TopResponse license key using the following commands:

   ```
   /sbin/service tlnipscd stop
   /sbin/service tlnipscd start
   ```

   > NOTE
   >
   > For more information on managing the IPS Controller service, see Appendix A, "IPS Controller System Management".

7. From another machine in your network, open a supported Web browser.
8. Specify the management interface URL. This includes the IP address of your IPS Controller system in either IPv4 or IPv6 format. You can specify the URL using either http or https. Corero recommends you use https for security purposes.

   ```
   http://<ip-address>:2080
   ```

or

```
https://<ip-address>:2443
```

9. Log into the IPS Controller software. The default username is admin and the default password is blank.

   NOTE

   When connecting to your IPS Controller system over SSL, you may see one or more warning messages. These warning message are normal. You may click the button(s) that accept this SSL certificate, and the messages will not display again.

10. If you are the first user to launch a version of the software that contains a new End User License Agreement (EULA), the EULA text will display. You must accept the EULA before you can continue.
11. The IPS Controller management application opens.

    Context-sensitive help is available from every window. You can access the entire online help system by choosing Help > Help Topics from the menu bar.

    NOTE

    Corero strongly recommends that the default SSL certificate supplied with the IPS Controller be updated with a customer supplied certificate as soon as possible. For more information on installing SSL certificates, see Installing an SSL Certificate on the IPS Controller (page 6-2).

Migrating IPS Controller Operations From One Server to Another

If you increase the number of Corero Network Devices your IPS Controller is managing, at some point, you may need to migrate your IPS Controller software and configuration information to a higher capacity system.

To perform this migration, complete the following steps:

1. Review the requirements listed in Chapter 3, “IPS Controller Pre-Installation Requirements” to ensure that the system on which you intend to install this version of the IPS Controller meets these requirements.

   If your system does not meet the necessary requirements, modify the chosen system or obtain a different system that does.

1. On the server being moved, log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Back up the IPS Controller on this system by doing one of the following:
   - If you want to back up all of the IPS Controller files for transfer to the new IPS Controller system, use the following command:

     ```
     /sbin/service tlnipscd backup <optional path>
     ```

   - If you only want to backup (then transfer) the configuration files to the new IPS Controller system, use the following command:

     ```
     /sbin/service tlnipscd configbackup <optional path>
     ```

3. Copy the backup files to the new server, and also to a secondary backup location for safekeeping.
1. On the new system, log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. If the new server is currently being used as an IPS Controller, you should backup the configuration directory in case there is a need to abort the migration and restore the old configuration. To back up the configuration directory, use the following command:

   ```
   /sbin/service tlnipscd configbackup <optional path>
   ```

3. Restore the backup files onto the new server using the following commands:

   ```
   /bin/gunzip <fname>.tgz
   /bin/tar -xvf <fname>.tar
   ```

   Whether you originally backed up the whole system, or simply backed up the configuration files, these commands create the /usr/local/tlnipscd/config directory and fill it with the contents from the original IPS Controller system.

   If the full backup was performed, all other backed up directories are restored as well.

4. Install the Top Response license key onto the new server, replacing 1111-2222-3333-4444-5555-6666-7777-8888 with your actual license key

   ```
   /sbin/service tlnipscd topkey 1111-2222-3333-4444-5555-6666-7777-8888
   ```

5. Stop and restart the IPS Controller service on the new server to activate the Top Response license key using the following commands:

   ```
   /sbin/service tlnipscd stop
   /sbin/service tlnipscd start
   ```

6. View displays and perform procedures on the new IPS Controller to verify that the configuration was successfully transferred.

Uninstalling the IPS Controller Software

If you want to repurpose a system running the IPS Controller software, you will want to remove all IPS Controller software components first.

CAUTION

Executing the uninstall script from any version's bin directory removes all copies of all versions of the IPS Controller in all bin directories. This includes all saved configurations, groups, and IPS Controller users.

To remove the IPS Controller software:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Navigate to the directory where the uninstall script is located. This script is typically found in the bin/<sw-release-name> directory. For example:

   ```
   /usr/local/tlnipscd/bin/V470049
   ```

   where V470049 is the version number of the IPS Controller release that you want to remove. If you want to remove a different version then use the number for that specific version.
3. Enter the following command:

   ```
   [root@your-system linux]# ./cms_uninstall.sh
   ```

   The software, service, saved configurations, user groups, and users are all removed.

Chapter 5
Getting Started With the IPS Controller

The IPS Controller's Graphical User Interface (GUI) is a Java Web Start™ application that runs as a stand alone application. The GUI is the primary management application for IPS Controller 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:

NOTE

A maximum of eight GUI sessions can be opened simultaneously. The recommended limit is 6 because if a GUI session is temporarily disconnected from the IPS Controller and then reconnects, it may use up 2 of the allotted 8 sessions for a short period of time.

Accessing the Management Application

To launch the IPS Controller management application:

1. The IPS Controller'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 Controller Release Notes* for JRE version and availability.
2. To display the IPS Controller's main window, point your browser at the IP address you assigned to the IPS Controller 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 Controller downloads the Java Web Start Application to your computer.
4. The login window displays. The default management username is admin. The default password is blank (there is no default password).

   NOTE

   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 recently-installed 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.

   NOTE

   After you have accepted the EULA, you can access the EULA text from the About dialog box. For more information, see Viewing "About..." Information (page 22-5).

6. The main window displays (Figure 5-1).

Multiple Management Application Sessions

If there are several users currently accessing the IPS Controller management application, consider the following:

- Configuration changes made in one session are reflected in and affect all other open sessions.
- Graphs, statistics, and monitoring activities are specific to each individual open session. The data you view reflects only the policy groups, ProtectionClusters, or devices you selected from the Policy Group tree.

Figure 5-1: IPS Controller Management Application Main Window

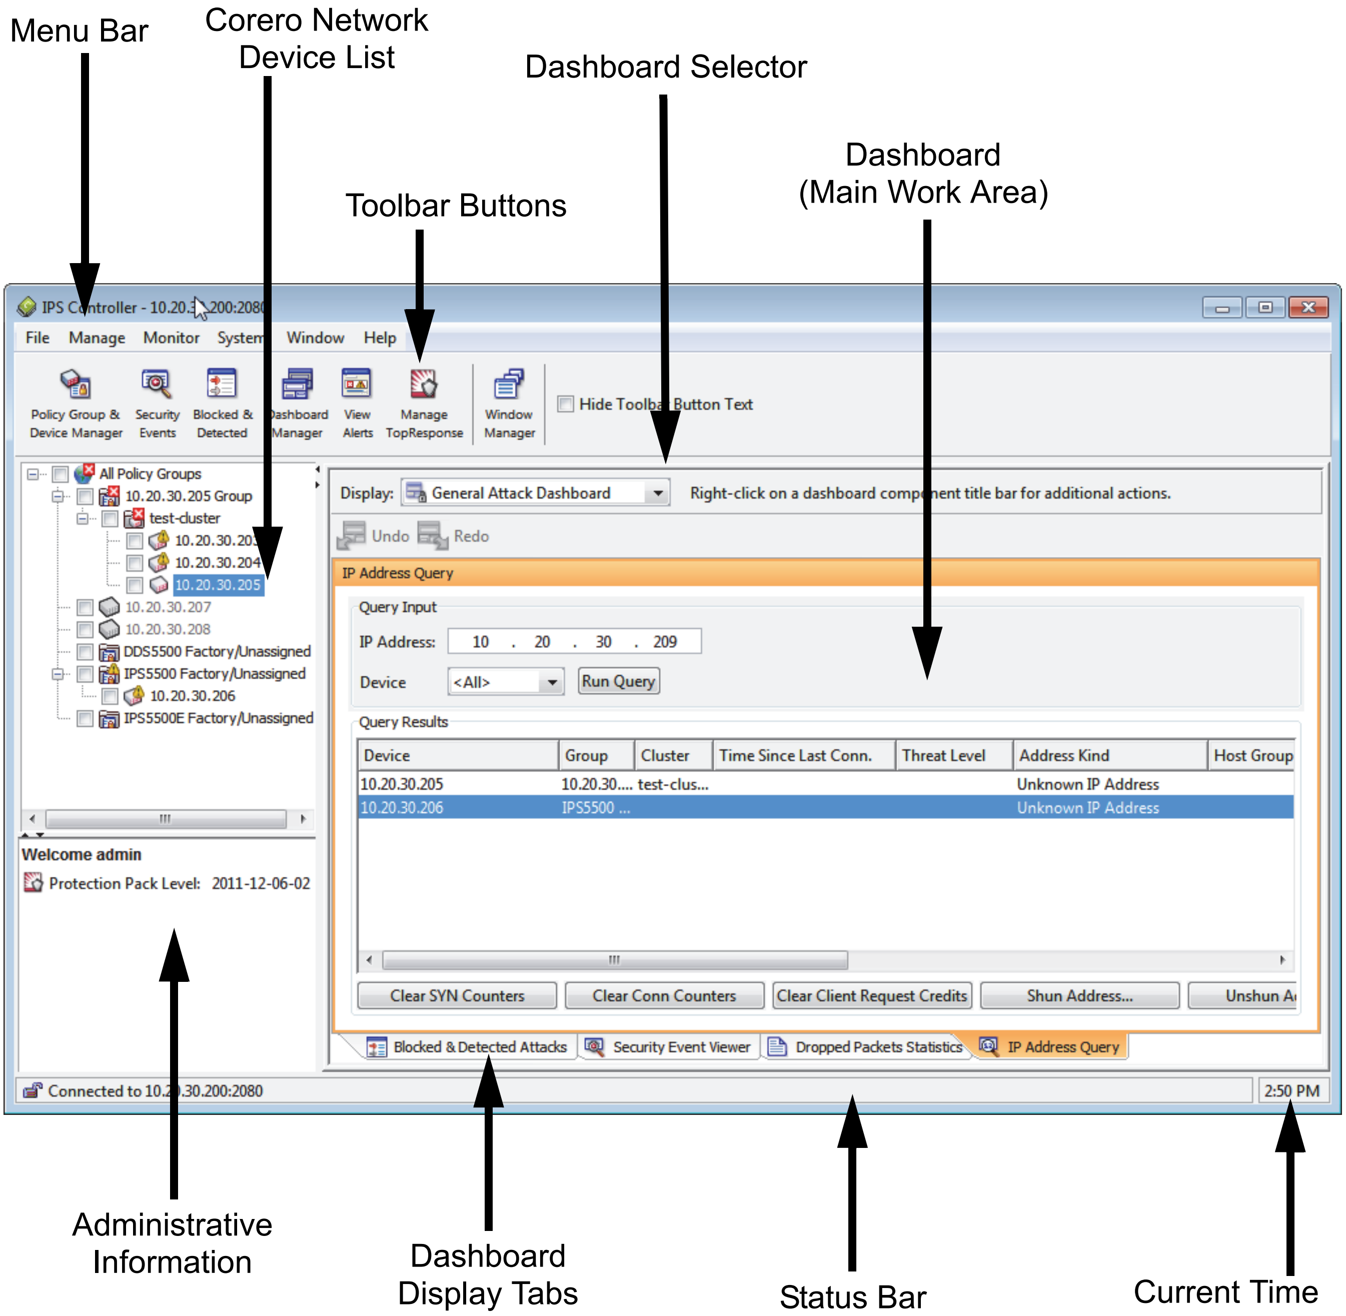

Display: General Attack Dashboard

Right-click on a dashboard component title bar for additional actions.

IP Address: 10 . 20 . 30 . 209

Device <All>

Device	Group	Cluster	Time Since Last Conn.	Threat Level	Address Kind	Host Group
10.20.30.205	10.20.30....	test-clus...			Unknown IP Address	
10.20.30.206	IPS5500 ...				Unknown IP Address	

Welcome admin

Protection Pack Level: 2011-12-06-02

Connected to 10.20.30.200:2080

2:50 PM

Using the Menu Bar

The Menu Bar, located at the top of the main window, provides access to system configuration and management windows as well as security configuration and management.

The Navigation Tree groups IPS Controller features into a series of top-level choices, which are described in Table 5-1.

Table 5-1: Menu Bar Options

Menu	Option	Description
File	Logout	Logout from the current session (but leave the session running), and display the login window.
	Exit	Exit (close) this session.
Manage	Policy Groups	There are two tabs in the Policy Groups area of the Policy Group and Device Manager dialog box: • Membership - Enables you to create and manage policy groups, initiate IPS Controller management of Corero Network Devices, and assign Corero Network Devices to policy groups. • Settings - When you select a device or policy group you can modify its associated firewall, rate-based, and environmental security settings; push settings out to one or more Corero Network Devices; pull device settings that you want to use as the basis for a policy group's settings; and apply Corero Protection Packs. For more information on policy groups, see Chapter 8, “Managing Policy Groups”.
	Devices	There are five tabs in the Devices area of the Policy Group and Device Manager dialog box: • Ports - Select a Corero Network Device and manage port-related settings, access a real-time view of the device’s front panel display, or perform device management tasks such as rebooting or changing the operating (bypass) mode. • NTP (Network Time Protocol) Servers - Identify up to four NTP servers from which a Corero Network Device can receive the correct time in order to automatically update its system clock. • Syslog Servers - Identify up to four Syslog servers to which you want the IPS Controller to send log data. Syslog Information can then be processed by the Network Security Analyzer program. • Software Upgrades - Install software updates to the Corero Network Devices you select. • Diagnostics & Reports - Upload and display security reports or diagnostics from a Corero Network Device. For more information, see Chapter 6, “Initial IPS Controller Configuration Tasks” and Chapter 11, “Viewing and Configuring Ports”.

Table 5-1: Menu Bar Options *(Continued)*

Menu	Option	Description
Monitor	Security Event Viewer	The Security Event Viewer is 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. For more information, see About the Security Event Viewer (page 23-18)
	IP Address Query	Provides details about the behavior of the selected host as it is requesting and completing connections. You can filter the data so only the host behavior associated with the selected device(s) displays. You can also reset the SYN flood or connection counters for the selected IP address. For more information about querying IP addresses, see Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).
	ProtectionCluster Status	Opens a window containing detailed ProtectionCluster-related information about each cluster member and its availability (connection) to other members of the cluster. ProtectionCluster members communicate with one another over High Availability (HA) links. Every cluster member is connected to every other cluster member by 2 high availability links. For more information about ProtectionClusters, see Chapter 14, “ProtectionCluster Configuration”.
	Alerts	Opens a searchable view of the current Alerts list. This window displays alerts associated with the IPS Controller and the Corero Network Devices it manages. For more information about viewing alerts, see Viewing the IPS Controller Management Alert Table (page 21-5).
	Statistics	Provides a statistical view of key traffic data such as blocked and detected attacks, dropped packets, and application connections handled by a Corero Network Device. For more information about viewing statistics see Viewing Blocked and Detected Attacks (page 23-16), Viewing Dropped Packet Statistics (page 23-31), and Viewing Port Statistics (page 23-33).
	Reports	Obtain various reports that were generated on Corero Network Devices. These reports display detailed security information. You can upload either scheduled reports (periodic reports), or reports generated on demand (immediate reports). For more information on reports, see Chapter 20, “Generating and Viewing Security Reports.

Table 5-1: Menu Bar Options *(Continued)*

Menu	Option	Description
System	Manage TopResponse Updates	Corero's software subscription service delivers frequent protection pack updates and security advisories. Protection packs typically include data about ill-behaved IP addresses and updated vulnerability and attack signatures. These protection packs can be automatically or manually downloaded to the IPS Controller from Corero's secure TopResponse server. From the IPS Controller, protection packs can be automatically or manually delivered to your Corero Network Devices. For more information on TopResponse, see Chapter 9, "Using and Managing TopResponse".
	Manage IPS Device Software Packages	You can download Corero Network Device software upgrade packages to the IPS Controller and deliver those packages for installation on IPS and DDS Units. For more information on managing Corero Network Device software, see Managing Corero Network Device Software (page 7-9).
	View/Edit Login Banner	You can specify text for a customer banner that is displayed on the IPS Controller login screen. This banner text can be used to communicate current information about the IPS Controller and its managed Corero Network Devices. For more information on modifying the login banner, see Modifying the Login Banner (page 6-11)
	Configure Audit Log Settings	The IPS Controller has an audit function that logs every change to the system's configuration. For more information, see
	View Audit Logs	
	Users	Enables you to create and manage user groups and users, set privilege levels, control user status, manage passwords and apssword settings, and manage authentication and radius settings.
Window	Window Manager	Access the Window manager to select and display currently open windows in the IPS Controller interface, or access the dashboard manager.
	Dashboard Manager	Access the Window manager to select and display currently open windows in the IPS Controller interface, or access the dashboard manager.
Help	IPS Controller Help	Access the complete help system;
	About IPS Controller	obtain information about the current version of the IPS Controller software.

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 5-2.

NOTE

You can hide the text associated with the toolbar buttons, leaving only the associated icon, by clicking the Hide Toolbar Button Text check box.

Table 5-2: Toolbar Buttons

Icon	Label	Corresponding Menu Option	Description	For More Information, See...
	Policy Group & Device Manager	Any option under Manage > Policy Groups or Manage > Devices	Display the Policy Group and Device Manager window and manage policy groups and devices.	Chapter 6, "Initial IPS Controller Configuration Tasks" Chapter 8, "Managing Policy Groups" Chapter 11, "Viewing and Configuring Ports" Chapter 21, "Managing Security Logs"
	Security Events	Monitor > Security Event Viewer	Display the Security Event Viewer to examine and manage network security events.	About the Security Event Viewer (page 23-18)
	Blocked & Detected	Monitor > Statistics > Blocked & Detected	Display the blocked and detected window to examine network security events.	Viewing Blocked and Detected Attacks (page 23-16)
	Dashboard Manager	Window > Dashboard Manager	Enables you to select and manage dashboards.	Using the Dashboard Display (page 5-9)
	View Alerts	Monitor > Alerts	Displays detailed alert information, and enables you to search for specific alerts, and acknowledge alerts.	Viewing the IPS Controller Management Alert Table (page 21-5)
	Manage TopResponse	System > Manage TopResponse Updates	Enables you to view, acquire, and apply TopResponse Protection Packs.	Chapter 9, "Using and Managing TopResponse"
	Window Manager	Window > Window Manager	Display the Window Manager tool. For more information.	Managing Application Windows (page 5-13)

Using the Online Help

The Graphical User Interface (GUI) provides access to detailed conceptual and procedural information in a richly cross-referenced online help system. 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 Display

The dashboard consists of one or more individual graphs and charts that are displayed simultaneously, providing snapshots of Corero Network Device operation and attack mitigation results. The dashboard is located in the work area of the main window.

You can select a dashboard from the Display drop-down 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.

You can also create your own dashboard using the Dashboard Manager, which you can access by clicking the Dashboard Manager toolbar button. When configuring a custom dashboard, consider the following:

- The individual components in a dashboard can be moved around within that dashboard to suit the user's requirements.
- If you find you have made a modification and you do not want to keep it, you can click Undo to remove your change. You can also click Redo to reinstate a change you undid.
- Components can be overlapped with one another 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 on the dashboard.

Available default dashboards are described in Table 5-3.

Table 5-3: IPS Controller Default Dashboards

Dashboard	Contents
Activity	Displays information on connection setup rates, dropped packets, and current application connections.
Aggregate Charts Dashboard	Displays aggregate information on IPS unit health, CPU statistics, dropped packet statistics, and flow statistics for the selected devices.
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 5-4.

Table 5-4: 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.

Table 5-4: Available Dashboard Components *(Continued)*

Chart	Description
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.
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.
Chart: IPS Unit Health	Displays the status of the selected IPS unit(s), showing whether they are operational and synchronized, operational but needing a push or pull, or not connected.
Chart: IPS Dropped Packet Statistics	Provides a breakdown showing dropped packet counts and the reasons why a Corero Network Device dropped packets.
Chart: IPS CPU Statistics	Displays statistics associated with CPU resource usage.
Chart: IPS Flow Statistics	Displays statistics associated with traffic flows.

Using the Policy Group Tree

The IPS Controller management application enables you to organize Corero Network Devices into policy groups. You can then specify, modify, and apply settings to the group as a whole. In addition, for high availability or high throughput purposes, you can combine between two and eight Corero Network Devices to form a ProtectionCluster.

When you manage multiple Corero Network Devices, the management application enables you to group them in order to simplify the management process. You can view the hierarchy of organized devices in the Policy Group Tree. This tree structure displays to the left of most main windows in the IPS Controller management application (Figure 5-2).

Figure 5-2: Policy Group Tree

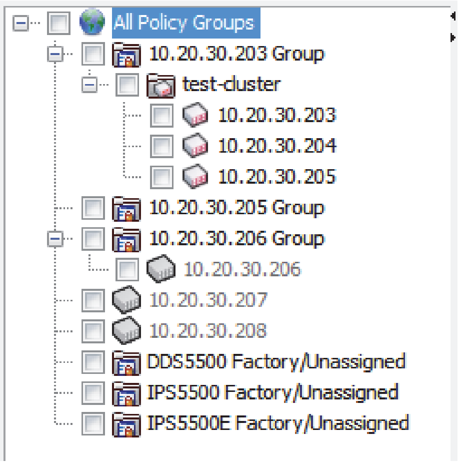

```
⊟ ☐ All Policy Groups
   ⊟ ☐ 10.20.30.203 Group
      ⊟ ☐ test-cluster
            ☐ 10.20.30.203
            ☐ 10.20.30.204
            ☐ 10.20.30.205
     ☐ 10.20.30.205 Group
   ⊟ ☐ 10.20.30.206 Group
         ☐ 10.20.30.206
     ☐ 10.20.30.207
     ☐ 10.20.30.208
     ☐ DDS5500 Factory/Unassigned
     ☐ IPS5500 Factory/Unassigned
     ☐ IPS5500E Factory/Unassigned
```

Icons to the left of each item in the tree help you identify whether it represents a policy group, ProtectionCluster, IPS Unit, or DDS Unit. You can click the check boxes associated with items in the Policy Group Tree to display real-time data for the selected items. The data displays to the right in the form of security information charts and graphs and other windows such as the Security Event Viewer. If your selections include more than one item, the IPS Controller management application displays aggregate data from them.

You can also right-click elements (and their icons) in the Policy Group Tree to quickly access related features. These operations apply to the selected items or group. For more information on policy groups, see Chapter 8, “Managing Policy Groups”.

Using the Policy Group and Device Manager Window

When you need to view or modify settings for the IPS Controller and the Corero Network Devices it manages, you do so through the Policy Group and Device Manager dialog box. You can access this window by clicking the Policy Group & Device Manager toolbar button.

When you launch the Policy Group and Device Manager dialog box, the Membership tab displays (Figure 5-3).

Figure 5-3: Membership Tab

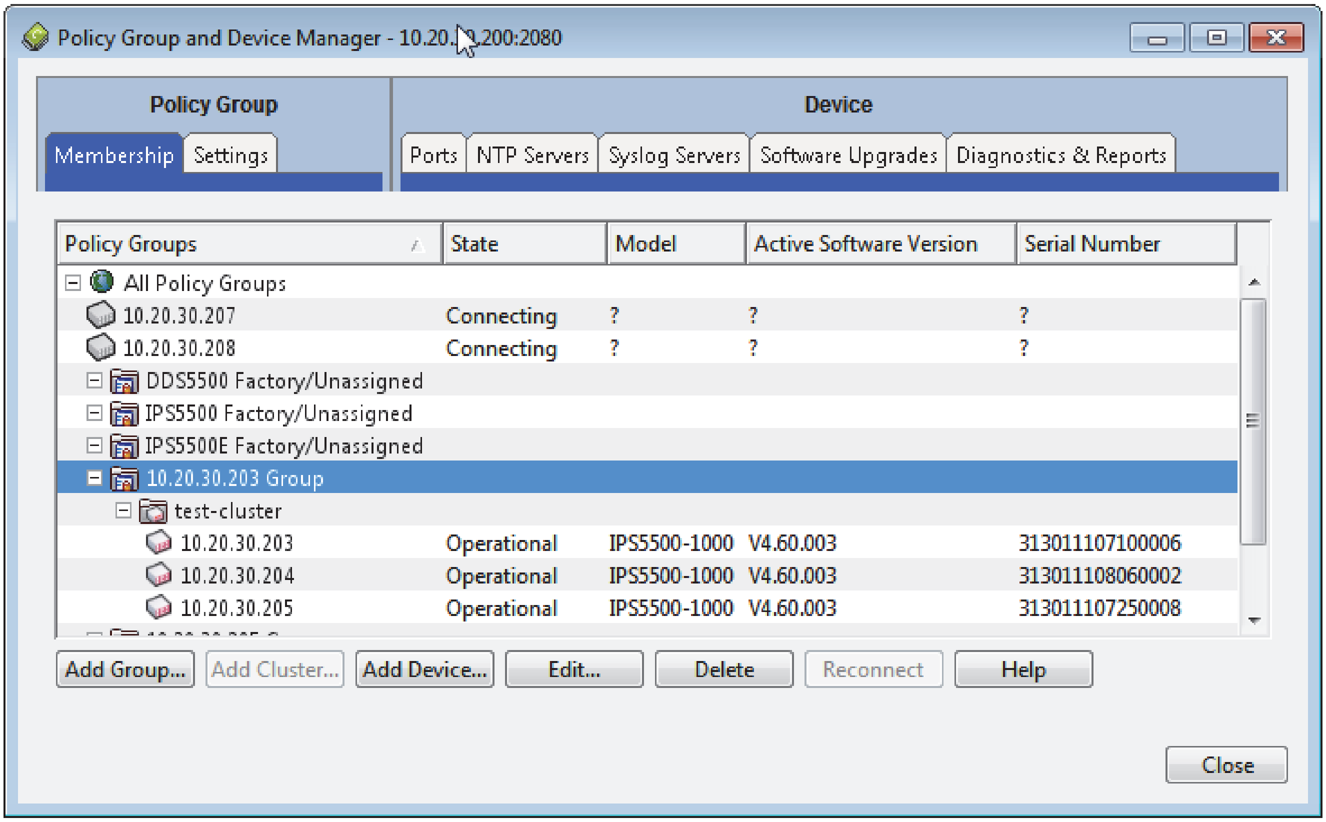

Policy Groups	State	Model	Active Software Version	Serial Number
⊟ All Policy Groups				
10.20.30.207	Connecting	?	?	?
10.20.30.208	Connecting	?	?	?
⊟ DDS5500 Factory/Unassigned				
⊟ IPS5500 Factory/Unassigned				
⊟ IPS5500E Factory/Unassigned				
⊟ 10.20.30.203 Group				
⊟ test-cluster				
10.20.30.203	Operational	IPS5500-1000	V4.60.003	313011107100006
10.20.30.204	Operational	IPS5500-1000	V4.60.003	313011108060002
10.20.30.205	Operational	IPS5500-1000	V4.60.003	313011107250008

The tabs at the top of the screen provide access to IPS Controller features. On the left, in the Policy Group area, you can use the Membership tab to view and modify policy group and ProtectionCluster membership. The Settings tab enables you to modify settings associated with security policies and protection packs, and view IPS Controller alerts. To the right is a group of tabs that enable you to view and manage both Corero Network Devices and the IPS Controller itself.

Each tab displays and provides access to managed items using a tree view of the available policy groups. You can expand or collapse policy groups or ProtectionClusters to see which Corero Network Devices belong to them. To the right of policy group entries, each tab displays feature-specific status and setting information. At the bottom of each tab, you can use the buttons perform actions on selected devices, ProtectionClusters, or policy groups. For more information about policy groups, see Chapter 8, “Managing Policy Groups”.

Helpful Hints

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

Maximum Management Sessions

A maximum of eight GUI sessions can be opened simultaneously. The recommended limit is 6 because if a GUI session is temporarily disconnected from the IPS Controller and then reconnects, it may use up 2 of the allotted 8 sessions for s short period of time.

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 which enables you to view a list of the currently open management windows (including graphs), and switch to a given window.

1. To manage IPS Controller windows, click the Window Manager tool bar button in the top right section of the main window.
2. To jump to a particular window, select the window from the list, then click Show.

Saving Changes

When you modify IPS Controller settings, your changes are automatically applied and saved. However, they are not applied to Corero Network Devices until you push the settings to them.

Reserved Words

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

- Any

- IP
- TCP
- UCP
- ICMP
- Inbound
- Outbound
- Internal
- External

Chapter 6
Initial IPS Controller Configuration Tasks

After you have completed IPS Controller installation, you must perform several initial configuration tasks prior to initial use.

This chapter describes the following initial configuration tasks:

- Installing an SSL Certificate on the IPS Controller (page 6-2)
- Installing the TopResponse License Key (page 6-3)
- Configuring User Authentication (page 6-4)
- Managing Syslog Servers (page 6-7)
- Managing Audit Logs (page 6-8)
- Managing Network Time Protocol (NTP) Servers (page 6-9)
- Modifying the Login Banner (page 6-11)

NOTE

For information on features related to management access, refer to Chapter 8, “Management Access”.

Installing an SSL Certificate on the IPS Controller

SSL stands for Secure Sockets Layer, a protocol used for transmitting information back and forth securely over the Internet. SSL uses a cryptographic system that uses two keys to encrypt data: a public key known to everyone and a private or secret key known only to the recipient of the message. The private key is only revealed when the public key is given. Once a device has a private key of its own, the device can establish a secure connection using the SSL protocol.

When working with SSL Certificates on the IPS Controller, consider the following:

- The SSL certificate on the IPS Controller, used to provide secure (https) access to the IPS Controller management application, is located in the following directory:

  ```
  /usr/local/tlnipscd/ssl/
  ```

  The IPS Controller first looks in this directory for a user-defined certificate in a file called user.pem. If this file does not exist, it uses the default certificate that is installed by Corero in the file called default.pem.

- Corero strongly recommends that the default SSL certificate supplied with the IPS Controller be updated with a customer supplied certificate as soon as possible.
- IPS Controller SSL certificate files are PEM formatted and must contain both an unencrypted RSA private key and the SSL certificate. If a certificate chain is to used, this must be put into the user.pem file after the RSA key and the host SSL certificate. The Certificate Authority is the last entry in the file.
- If a new IPS Controller version is installed, the user.pem file is retained, and the default.pem file is changed to that provided by the new IPS Controller version.
- If the IPS Controller software is removed, the whole SSL directory will be deleted, removing any user.pem and default.pem files.

To install an SSL certificate:

1. Create a user.pem file that contains the new key and certificate (and optional certificate chain).
2. Place the file in the ssl directory (/usr/local/tlnipscd/ssl/).

   It is important that the new user.pem file have the same access rights as the default.pem file. You can use the following command to modify the access rights:

   ```
   chmod 774 user.pem
   ```

You must restart the IPS Controller service to read the modified certificate files. For instructions on how to restart the service, see Managing the IPS Controller Service (page A-2).

Installing the TopResponse License Key

TopResponse is an Automated Protection Update Service that provides Corero Network Device customers with advanced security services to maximize security, availability, and performance of their network. TopResponse provides automated updates, technical support, and security advisory and software subscription services, along with access to Corero's Knowledge Base.

> NOTE
>
> Any existing TopResponse license key is automatically removed when a new license key is installed

In order to access TopResponse features, you must install a TopResponse license key on the IPS Controller system.

To install the TopResponse license key:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. On the IPS Controller Linux system, issue the following command:

   ```
   > /sbin/service tlnipscd topkey 1111-2222-3333-4444-5555-6666-7777-8888
   ```

   replacing 1111-2222-3333-4444-5555-6666-7777-8888 with your actual license key.
3. Restart the tlnipscd service as described in Managing the IPS Controller Service (page A-2).

Configuring User Authentication

The IPS Controller enables 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 6-1.

Table 6-1: User Authentication Methods

Authentication Method	Description
Local Only	Users are authenticated against a locally maintained database on the IPS Controller 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 IPS Controller searches for the Radius server. For information on how to define Radius servers, see Managing Radius Servers (page 6-5).

You can identify up to eight Radius servers for each device. 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 6-2.

Table 6-2: 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 6-2: 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 6-5).
2. Choose System > Users> 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 6-1.
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 6-2.
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 for use by your IPS Controller.

To configure Radius servers:

1. Select System > Users > Radius Settings from the menu bar. 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 6-3.
4. When finished specifying Radius server parameters, click OK.

Table 6-3: 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 6-3: 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 IPS Controller 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. 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 Corero Network Devices configured to use the Round Robin algorithm for an alternate radius server search method. It only applies to devices configured to use the Priority-based algorithm for an alternate radius server search method. For more information on Radius server search methods, see Configuring User Authentication (page 6-4)
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.

Managing Syslog Servers

The IPS Controller management application enables you to add, view, and modify Syslog server (host) information for managed Corero Network Devices.

> NOTE
>
> You can define up to four Syslog servers on each Corero Network Device, but you cannot define remote Syslog servers for the IPS Controller itself. The IPS Controller can use the local Linux Syslog server for storing its own audit log messages.

To view, add, or modify Syslog servers for IPS and DDS Units:

1. From the menu bar, choose System > Syslog Servers. The Policy Group and Device Manager dialog box displays with the Syslog Servers tab highlighted.

   The Syslog Servers tab displays the current state of each managed Corero Network Device. This information is followed by the IP address for each of the Syslog servers defined for that device. Each IP address is followed by the UDP Port, which is, by default, the well known UDP port 514.

2. To add or modify a Syslog server to which a Corero Network Device should send event messages:
   a. Select the desired device, then click Edit Servers. The Syslog Hosts dialog box displays.
   b. To modify an existing Syslog host, you can edit the existing IP Address, Port, and Admin Status (enabled or disabled) for that host.
   c. To add a new Syslog host, in the uppermost row displaying an IP address of all zeros (0), enter the desired IP Address and Port information, and select the desired Admin Status (enabled or disabled).
   d. To remove a Syslog host, modify the IP Address for the device to be 0.0.0.0.
   e. To change the Admin Status (enabled or disabled) for a specified Syslog host, use the Admin Status drop-down to select a different status for the desired Syslog host.

3. When finished, click OK.

Managing Audit Logs

The IPS Controller has an audit function which logs every change to the system's configuration. These log items are kept in an Audit log file, and can also optionally be sent to the local Linux system's Syslog server.

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

Audit messages are kept in an audit log file stored on the IPS Controller itself. Up to 10 audit files are maintained, the oldest being deleted to make available space for a new one when required.

To enable audit logging on the IPS Controller:

1. From the menu bar, select System > Configure Audit Log Settings. The Audit Logging dialog box displays.
2. To enable audit logging, select the Enable Audit Logging check box.
3. To additionally send audit log messages to the local Linux Syslog server, once you have enabled audit logging, select the Send Audit Data to Syslog Server check box.

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

4. When finished, click OK.

For information on how to view audit logs, see Viewing Audit Logs (page 21-8).

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 managed Corero Network Devices. 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 a Corero Network Device loses NTP sync, this information will be logged to the system log file and Syslog server (if configured) for that device.

NOTE

The IPS Controller obtains its time settings and synchronization from the local Linux server. The NTP Servers tab only enables you to modify settings for managed Corero Network Devices.

With NTP enabled, Corero Network Devices determine the system time by receiving NTP broadcast or multicast messages or by querying an NTP server at the time interval you configure. The IPS or DDS Unit 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 via NTP from a stratum 1 time server. Corero Network Devices support 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

Corero Network Devices do not accept a response from an NTP server that is more than 45 minutes different 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 Corero Network Device’s time as close to the real time as possible.

NOTE

Because each 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 the Configuration and Management Guide for your Corero network Device.

1. From the menu bar, choose Manage > Devices > NTP. The NTP Servers tab of the Policy Group and Device Manager dialog box displays.

   The NETP Servers tab displays the current state of each managed Corero Network Device, followed by NTP settings information including:

   - Whether or not the Corero Network Device should Receive NTP broadcast messages.
   - Whether or not the Corero Network Device should query an NTP Server for an updated time.
   - If you have configured the device to query NTP servers, the Query Interval indicates how frequently the device requests the current time from an NTP server.
   - The NTP Server(s) currently specified for the device.

2. To modify the NTP settings for a Corero Network Device, select the desired device, then click Edit NTP Settings. The Network Time Protocol dialog box displays.

3. 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.
4. Specify whether or not the Corero Network Device should query an NTP Server for an updated time.

   You specify the Query Interval, which indicates how frequently the device requests the current time from the NTP server.

5. To add a new NTP server, enter the IP Address for the NTP server, then click Add.
6. To delete an NTP server, select the server in the list and click Delete.
7. When finished, click OK.

Modifying the Login Banner

You may find you wish to communicate technical or time-sensitive information to IPS Controller users. The IPS Controller enables you to specify and display a login banner. Banner text displays on the IPS Controller login screen for all users.

To modify the Login Banner:

1. From the menu bar, choose System > View/Edit Login Banner. The View/Edit Login Banner dialog box displays.
2. Type the message you want to display in the message area.

   You can enter plain text, or you can enter text with simple HTML tagging.

3. You can turn display of the banner on and off using the Display Banner at Login check box.
4. When finished, click OK.

Chapter 7
Managing Corero Network Devices

Once you have performed the initial configuration tasks for the IPS Controller, you will need to add the Corero Network Devices you wish to manage. Once the IPS Controller is managing your IPS and DDS devices, you can use the IPS Controller management application as a central location to monitor and manage multiple devices.

This chapter contains the following sections:

NOTE

For information on how to organize the IPS Controller-managed Corero Network Devices into policy groups, see Planning Policy Groups (page 8-12)

Management Prerequisites

Before you perform initial configuration on the IPS Controller:

- Ensure all of the Corero Network Devices (IPS and DDS Units) you will be managing with the IPS Controller are installed and cabled as described in the device-specific Hardware Installation Guides.
- Ensure you have assigned (and noted) the following information for the device(s) you wish to manage:
  - IP address
  - Shared management key
- Ensure that the device has been configured to allow the IPS Controller to manage it. By default, Corero Network Devices are set to block the IPS controller from managing the device.
- Ensure you have performed all initial configuration tasks for the Corero Network Devices you will be managing with the IPS Controller as described in the product-specific Configuration and Management Guides for those devices.

  For detailed information on preparing Corero Network Devices for management, refer to the *Configuration and Management Guide* for your device.

NOTE

By default Corero Network Devices have port 2616 open for communication with the IPS Controller. If you have blocked this port, you must open it on the management port or the IPS Controller will not be able to manage the Corero Network Device.

About Shared Management Keys

The IPS Controller uses a shared key to authenticate itself with a Corero Network Device (IPS or DDS Unit) you wish to manage or monitor from the controller. Authentication is part of the process by which the Corero Network Device becomes operational, that is, able to be managed from the controller. Once authentication is complete. the IPS Controller communicates with each Corero Network Device through a secure tunnel.

By default, a Corero Network Device does not have a shared management key, so you will need to set one before you can initiate IPS Controller management. For instructions on how to set the shared management key for a Corero Network Device, refer to the device-specific Configuration and Management Guide.

Adding Corero Network Devices to the IPS Controller

Adding a Corero Network Device to the IPS Controller means that you identify the device to the controller so that you can monitor and manage the device from the controller. When you first add the device to the IPS Controller, the controller places the device into a generic policy group based on the device's type and model. Once you create your own policy group, you must move the device to the policy group you have defined for it.

Once you have added a Corero Network Device to the IPS Controller, the IPS Controller begins its connection process to the IPS or DDS Unit. It uses the shared key for authentication. As soon as it determines the device's type, it places the device in the proper default "Factory/Unassigned" policy group.

CAUTION

Once you are managing a Corero Network Device with an IPS Controller, all device-specific configuration changes must be made through the IPS Controller. Under central (IPS Controller) management, attempts to modify configuration settings using the Corero Network Device management application will be denied.

Adding an IPS or DDS Unit to the IPS Controller

Before you can create a policy group and use it to organize your Corero Network Devices, you must add the devices to the IPS Controller using the management interface. Adding devices to the IPS Controller puts them under controller management.

The IPS Controller communicates with IPS and DDS Units using a shared management key. This key is a

To add a Corero Network Device to the IPS Controller:

1. Ensure you have assigned an IP address to the Corero Network Device management port.
2. Ensure you have assigned a shared management key to the Corero Network Device that you wish to add to the IPS Controller. When performing this task, consider the following:
   - By default, Corero Network Devices do not have a shared management key. Use the Corero Network Device's management interface to add the shared key. For more information on setting shared keys, see the Configuration and Management Guide for your Corero Network Device.
   - Be sure to remember or record the shared management key for your Corero Network Device. You will be required to enter it before you can manage that device from the IPS Controller. If you do not remember the shared management key for a device, contact Corero Support for assistance.
   - If you change the shared key for a Corero Network Device after the device has been added to the IPS Controller, the Controller will be unable to connect to the device until you change the shared key to match using the IPS Controller management application.

   For more information on using shared management keys, see Changing a Corero Network Device Shared Management Key (page 7-7).
3. By default, the Corero Network Device is set to block the IPS Controller from managing the device. You must modify this setting on the Corero Network Device using the device's management application. For instructions on how to do so, see the Configuration and Management Guide for your Corero Network Device.

   NOTE

   Corero Network Devices use port 2616 as the management port for communication with the IPS Controller.
4. To access the Membership tab of the Policy Group and Device Manager, do one of the following:

- Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Membership tab.
- From the menu bar, choose Manage > Policy Groups > Membership.

5. From the Membership tab of the Policy Group and Device Manager window, click Add Device. The Add Device dialog box displays.
6. In the IP Address box, enter the IP address of the IPS or DDS Unit you want to add.
7. Do not attempt to modify the default Port value (2616). This is currently the only port that the IPS Controller uses to communicate with Corero Network Devices.
8. In the Shared Key text box, enter the shared key you previously defined for the Corero Network Device you are adding.
9. Do one of the following:
   - Click Add to add this device and leave the dialog box open so you can add another.
   - Click Done to add the device and close the dialog box.
10. After you add the device to the controller, the IPS controller immediately attempts to establish a management link to the device. Watch the connection process in the State column of the Policy Group and Device Manager Membership tab to verify the device has been successfully added. The appropriate connection process messages are listed in Table 7-1.

Table 7-1: Corero Network Device Connection Messages

State	Description
Connecting	The IPS Controller attempts to connect to the IP address you specified for the Corero Network Device.
Authenticating	Once connected, the IPS Controller uses the shared management key to authenticate itself with the device.
Synchronizing	Once authenticated, the IPS Controller attempts to synchronize with the device by capturing a copy of all settings or parameters on the device.
Operational	Once synchronized, the device becomes operational. The IPS Controller is now able to manage changes to device settings and monitor the device’s security events.

NOTE

If you find the device does not proceed through all phases culminating an an Operational state, see Troubleshooting Corero Network Device Connection Issues (page B-4) for troubleshooting information.

11. When the connection is successfully completed, the Corero Network Device is added to the default policy group for that device type (IPS or DDS).

Viewing and Modifying Corero Network Device Settings

You can modify a majority of device configuration parameters for managed Corero Network Devices directly from the IPS Controller management application. For example you can modify port settings, create or modify security policies, or configure access information for Syslog servers.

NOTE

Corero Network Device setting modifications can only be performed if a direct connection from the managed device to the IPS Controller is up and available.

When you use the IPS Controller to manage settings on Corero Network Devices, you will find there are some settings that can only be performed on a single device, and some settings that can only be performed on a policy group. Table 7-2 lists the types of settings and how you can manage them. It also provides a cross-reference to additional information on that topic.

Table 7-2: Modifying Corero Network Device Settings

Settings	Can Only Modify for a Single Device	Can Only Modify for a Policy Group	For More Information, See...
Ports, including Bypass Settings	X		• Chapter 11, “Viewing and Configuring Ports • Chapter 10, “Understanding Ports • Chapter 13, “Advanced Port Configuration
NTP Servers	X		• Managing Network Time Protocol (NTP) Servers (page 6-9)
Syslog Servers	X		• Managing Syslog Servers (page 6-7)
Reports	X		TBD
Security Policies	X		• Chapter 15, “About Security Policies • Chapter 16, “Managing FW+IPS Security Policies
Host Groups		X	• Chapter 17, “Managing Host Groups
Services		X	• Chapter 18, “Managing Services
IPS Rule Sets		X	• Chapter 19, “Managing Rules and Rule Sets
Connection Limiting		X	• Chapter 24, “SYN Flood and Connection Limiting Security
SYN Flood Limiting		X	• Chapter 24, “SYN Flood and Connection Limiting Security
Client Rate Limiting		X	• Chapter 25, “Client Rate Limiting

Reconnecting to a Corero Network Device

If you find that a managed Corero Network Device is unable to reach Operational status, you can instruct the IPS Controller to attempt to reconnect to the device. Note that you should avoid disrupting a device connection if system operations are underway, such as a push or pull operation.

To reconnect to a Corero Network Device:

1. To access the Membership tab on the Policy Group and Device Manager dialog box, do one of the following:
   - Click the Policy Group and Device Manager tool bar button. When the dialog box displays, click the Membership tab.
   - From the menu bar, choose Manage > Policy Groups > Membership.
2. From the Policy Groups tree, select the Corero Network Device to which you want the IPS Controller to reconnect, then click Reconnect.
3. The IPS controller immediately attempts to renew its management link to the device. Watch the connection process in the State column of the Policy Group and Device Manager Membership tab to verify the device has been successfully added. The appropriate connection process messages are listed in Table 7-1.

Changing a Corero Network Device Shared Management Key

Before the IPS Controller can manage a Corero Network Device, it must authenticate itself with that device. Authentication is part of the process by which the device becomes manageable by the controller.

Authentication only occurs if the shared management key specified on the Corero Network Device (using the device's management application) matches the key you specify in the IPS Controller management application.

If you modify the shared management key on a managed Corero Network Device, you will also need to change the management key information for that device in the IPS Controller management application. This procedure is described below.

> NOTE
>
> Data created by a Corero Network Device is sent directly to the device's configured Syslog servers. Changing the shared key, which temporarily breaks the connection between the device and the IPS Controller, does not affect data flow from the device to its associated Syslog servers.

To change the key for a managed Corero Network Device:

1. To access the Membership tab on the Policy Group and Device Manager dialog box, do one of the following:
   - Click the Policy Group and Device Manager tool bar button. When the dialog box displays, click the Membership tab.
   - From the menu bar, choose Manage > Policy Groups > Membership.
2. From the Policy Groups tree, select the Corero Network Device whose shared management key you wish to change, then click Edit. The Edit Device window displays.
3. To edit the shared key used by the IPS Controller to communicate with the device, select the Edit Shared Key check box.
4. In the Old Shared Key box, enter the existing shared key currently assigned to the Corero Network Device on the IPS Controller.
5. In the Shared Key box, enter the new shared key for this device.
6. In the Confirm Shared Key box, reenter the shared key.
7. Click OK. Since the connection between a Corero Network Device and the IPS Controller depends on their both having the same shared key, one of the following happens:
   - If you changed the shared key on the device before changing it on the IPS Controller, the IPS Controller will have lots its connection to the device at that time. Once you have modified the shared key on the IPS Controller, the controller will now attempt to reconnect with the device using the new shared key.
   - If you changed the shared key on the IPS Controller before changing it on the device, the controller will lose connectivity with the device until you enter the new shared key on the device's management interface.

Deleting a Corero Network Device

At some point, you may find you need to delete a Corero Network Device from the IPS Controller, discontinuing IPS Controller management of the device.

When preparing to delete a Corero Network Device from the IPS Controller, consider the following:

- The Delete action completely removes the device from policy group tree, and ends IPS Management of the device. If you want to move an IPS or DDS Unit to a different policy group, see Modifying the Membership of a Policy Group (page 8-15).
- You cannot delete an individual Corero Network Device from a ProtectionCluster. You must delete (dissolve) the cluster before you can delete individual cluster members from the IPS Controller.

To delete an IPS or DDS Unit from a policy group, complete the following steps:

1. To access the Membership tab on the Policy Group and Device Manager dialog box, do one of the following:
    - Click the Policy Group and Device Manager tool bar button. When the dialog box displays, click the Membership tab.
    - From the menu bar, choose Manage > Policy Groups > Membership.
2. From the policy groups tree, select the Corero Network Device you want to remove, then click Delete.
3. A confirmation dialog box displays. Click Yes to proceed with the deletion or click No if you want to cancel it.

NOTE

At this point, if you want to continue using the device you just removed from the IPS Controller as an autonomous (unmanaged) device, you should change the device's IPS Controller management setting from "Allow the IPS Controller to Manage This Device" to "Block the IPS Controller from Managing This Device". For information on how to do this, refer the Configuration and Management Guide associated with your device.

Managing Corero Network Device Software

When uploading an IPS device software package file to the IPS Controller, Corero recommends that you first copy the software package file to the local computer running the IPS Controller management application. Then use the IPS Controller GUI to add the package file to the IPS Controller. This avoids any network related issues that may cause problems when moving the file onto the local computer.

CAUTION

If you are upgrading the software on a managed Corero Network Device, ensure you perform the software upgrade from the IPS Controller, rather than performing the upgrade using the management application on the Corero Network Device. Upgrading the software using the device's management application can yield unpredictable results.

Software upgrade on Corero Network Devices is a three step process.

1. Download the software package to the IPS Controller.
2. Install the software on one or more Corero Network Devices.
3. Activate the software on the Corero Network Devices.

NOTE

The process you use to upgrade software on the IPS Controller itself is different from the process used to upgrade software on managed Corero Network Devices. For information on upgrading software on the IPS Controller, see Chapter 4, "Installing and Upgrading the IPS Controller Software.

It can take up to ten minutes to upload and apply new software and reboot the device. How traffic is handled depends on your bypass-related device settings. You can use bypass settings to specify whether traffic is passed unhindered (Always Bypass or Bypass on System Reset) during system restarting, or whether traffic is blocked (Never Bypass). For more information on bypass settings, see Bypass Settings (page 10-6).

The rest of this section describes the following:

- Viewing Software Version Information for a Corero Network Device (page 7-9)
- Making a Software Package Available on the IPS Controller (page 7-10)
- Upgrading a Software Package on a Corero Network Device (page 7-10)
- Activating a Software Package on a Corero Network Device (page 7-11)
- Removing Unneeded Software Packages from Your Corero Network Devices (page 7-11)
- Removing Unneeded Software Packages from the IPS Controller (page 7-12)

Viewing Software Version Information for a Corero Network Device

Before you install software on a Corero Network Device that is under IPS Controller management, you will need to view the installed software version that is currently running on that device.

To view the active software version for a Corero Network Device:

1. Click the Policy Group & Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
2. Click the Software Upgrades tab. This tab displays the information listed in Table 7-3

Table 7-3: Software Upgrades Tab Information

Column	Description
Policy Groups	The policy group tree can be expanded to list managed Corero Network Devices.
State	The current state of the management connection between the IPS Controller and the Corero Network Device. These states are described in Table 7-1.
Model	The hardware model of the Corero network Device.
Active Software Version	The software version that is currently running on the Corero Network Device.
Pending Software Version	The pending software version is the version of the software that will be activated after the next reboot.
Status	Indicates the software package installation progress. Software is downloaded and unpacked prior to installation.
Installed Software	Lists all resident versions of software available on the Corero Network Device. Corero Network Devices allow you to have multiple versions of software resident, with only one selected for activation (use) at any time.

Making a Software Package Available on the IPS Controller

If this is the first time you will be installing a software package, you will need to download the software package to your IPS Controller before you can install the software on any of your Corero Network Devices

To download a new software package to the IPS Controller:

1. Download the software (in the form of a .pkg file) from Corero and store it in a location that you can access from the IPS Controller.
2. From the menu bar, select System > Manage IPS Device Software Packages. The IPS Device Software Packages dialog box displays.

   Note that you can also access this dialog box by clicking the Manage Packages button on the Software Upgrades tab of the Policy Group and Device Manager dialog box.
3. Click Upload. The Upload IPS Device Software Package dialog box displays.
4. Navigate to the location where you placed the software package and upload it to the IPS Controller.

Upgrading a Software Package on a Corero Network Device

Once you have uploaded a Corero Network Device software package to the IPS Controller, you must then install the software upgrade.

When upgrading software on Corero Network Devices, consider the following:

- Since ProtectionCluster members must always run the same software version, you cannot upload new software to individual cluster members. You can only upload software on all members of the cluster at once.
- Although you must upload software to all ProtectionCluster members simultaneously, you can stagger the activation and rebooting of cluster members over a period of time.

Once the software package has been added to the IPS Controller, to upgrade software on Corero Network Devices:

1. Click the Policy Group & Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.

2. Click the Software Upgrades tab. The Software Upgrades tab displays information about managed Corero Network Devices and their Corero software.
3. Select the Corero Network Device(s) where you want the software installed, then click Install Package.
4. Review the list of available software packages and select the one you want to install on the device(s), then click OK.

   The IPS Controller sends the software package to the selected devices one by one. The status column indicates the progress of each package transfer.
5. When the installation is complete, the software package will be installed on the device(s), but will not be running (active) on that device. You will see that the information in the Pending Software Version column on the Software Upgrades tab displays a different version from the Active Software Version column.
6. When you are ready, you can activate the new software.

Activating a Software Package on a Corero Network Device

Corero Network Devices can hold multiple software versions in their non-volatile memory. Once you have uploaded software to a Corero Network Device, the software is now resident on the device, but has not yet been activated.

CAUTION

You must reboot the Corero Network Device in order to activate (apply) the pending software.

To activate (use) the pending software:

1. Click the Policy Group & Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
2. Click the Software Upgrades tab. The Software Upgrades tab displays information about managed Corero Network Devices and their Corero software.
3. Select the device whose software you wish to activate, then click Reboot.
4. The Corero Network Device reboots, and as it does it loads the latest software.

NOTE

After you activate the new software version on a Corero Network Device, if you are running the management application for that individual device (which is separate from the IPS Controller interface), you should restart the management application to ensure you have loaded the device management interface from the software package you activated.

Removing Unneeded Software Packages from Your Corero Network Devices

When you have finished installing software to one or more Corero Network Devices, if desired, you can remove older software packages (ones you know that will no longer need to install on devices) from the Corero Network Device on which it was previously installed.

Because the volume of compact flash memory on a Corero Network Device is limited, Corero recommends that you regularly store only 1 or 2 software versions on your Corero Network Device. You may store a temporary maximum of 3.

NOTE

A different process is used to remove unneeded software packages from the IPS Controller itself. For detailed instructions, see Removing Unneeded Software Packages from the IPS Controller (page 7-12).

To remove unneeded software packages from the IPS Controller:

1. Click the Policy Group & Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
2. Click the Software Upgrades tab.
3. Select the software package in the list, then click Delete Package. The Delete Software dialog box displays.
4. Visually verify that the device (listed by its IP address) is the one you intended.
5. Once you have confirmed the version, select the desired software package from the drop-down list, then click OK.

   The software is deleted from the Corero Network Device.

Removing Unneeded Software Packages from the IPS Controller

When you have finished installing software to one or more Corero Network Devices, if desired, you can remove older software packages (ones you know that will no longer need to install on devices) from the IPS Controller.

NOTE

If you want instructions on how to remove unneeded software packages from a Corero Network Device, see Removing Unneeded Software Packages from Your Corero Network Devices (page 7-11).

Note that removing the package from the IPS Controller does not delete the package from the directory location where you placed it on your network prior to uploading it to the IPS Controller. You will need to delete that copy of the software package manually.

To remove unneeded software packages from the IPS Controller:

1. From the menu bar, select System > Manage IPS Device Software Packages. The IPS Device Software Packages dialog box displays.
2. Select the software package in the list, then click Remove. The software is deleted from the IPS Controller.

Chapter 8
Managing Policy Groups

A policy group is a collection of Corero Network Devices that you would like to have the same security settings. A policy group is typically comprised of devices that share a common role or location at your site.

The purpose of a policy group is to allow you to create and modify policies and environmental configuration items and apply them as a group of settings to one or more IPS or DDS Units. You can also modify the settings within a policy group by automatically updating certain environmental configuration items using TopResponse Protection Packs.

The policy group settings of a Corero Network Device being managed by the IPS Controller are changed using the IPS Controller. Once you are satisfied with your changes, you can download them to the Corero Network Devices in that policy group. This enables you to modify the policy group settings at any time, and then choose specifically when you want to push, or deliver, the settings to the policy group members.

You can only create policy groups containing similar Corero Network Device models. For example, you cannot create a host group that contains both DDS Units and IPS Units. In addition, you may add a ProtectionCluster to a policy group, but all ProtectionCluster members must belong to the same policy group.

This chapter contains the following sections:

NOTE

Before you can add a Corero Network Device to a policy group, you must add that device to the IPS Controller, initiating device management. For instructions on how to initially add Corero Network Devices to the IPS Controller, see Management Prerequisites (page 7-2)

Policy Group Overview

A policy group is a collection of Corero Network Devices that should have the same security settings. A policy group can contain one or more individual devices and/or one or more ProtectionClusters(s). Each policy group contains all of the settings required to make security decisions as the devices examine traffic flows.

By collecting all of these settings into a policy group, you can easily manage the common settings you want to apply to specific sets of Corero Network Devices.

Each policy group contains all of the settings required by Corero Network Devices to make traffic security decisions as they examine traffic flows. By collecting devices into a policy group, you can easily manage the settings that you apply to specific sets of IPS or DDS Units.

A policy group enables you to specify the following types of settings across all policy group members:

- Firewall and intrusion protection policies, including shunning
- Rate-based policies
- Traffic rate limits
- Security environment configuration items, including:
  - Host groups
  - Services (a service is a combination of a network protocol, a port number, and a server group)
  - IPS rule sets (rules used by Corero Network Devices to classify traffic as good or bad)
  - IPS rule details (specific settings that apply to certain rules)
  - IPS signatures (specific data patterns that identify problem traffic)

The overall process for creating a policy group includes the following steps:

1. Plan the policy group, as described in Planning Policy Groups (page 8-12).
2. Add the policy group members to the IPS Controller, as described in Management Prerequisites (page 7-2).
3. Create the policy group, as described in Creating Policy Groups (page 8-13).
4. Add the desired members to the policy group, as described in Modifying the Membership of a Policy Group (page 8-15)

Viewing Policy Group Membership Information

Each Corero Network Device under management by the IPS Controller belongs to one, and only one, policy group. Each policy group contains a full set of security policy parameters that the IPS Controller applies to all members of that group.

When you first add a Corero Network Device to the IPS Controller, the controller places the device into a default policy group based on the device's type (IPS or DDS). The IPS Controller automatically defines a default policy group for each product family. Default product-specific policy groups include the following:

- IPS5500 Factory/Unassigned policy group
- IPS5500E Factory/Unassigned policy group
- DDS5500 Factory/Unassigned policy group

NOTE

These are the three primary types of Corero Network Devices. You cannot combine devices of different types in a single policy group.

After you add a Corero Network Device to the controller, at your convenience you can then move it out of its default policy group and into one of the policy groups that you define.

To view policy group membership information, go to the Membership tab of the Policy Group and Device Manager dialog box by doing one of the following:

- Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Membership tab.
- From the menu bar, choose Manage > Policy Groups > Membership.

The Membership tab displays (Figure 8-1).

Figure 8-1: Membership Tab

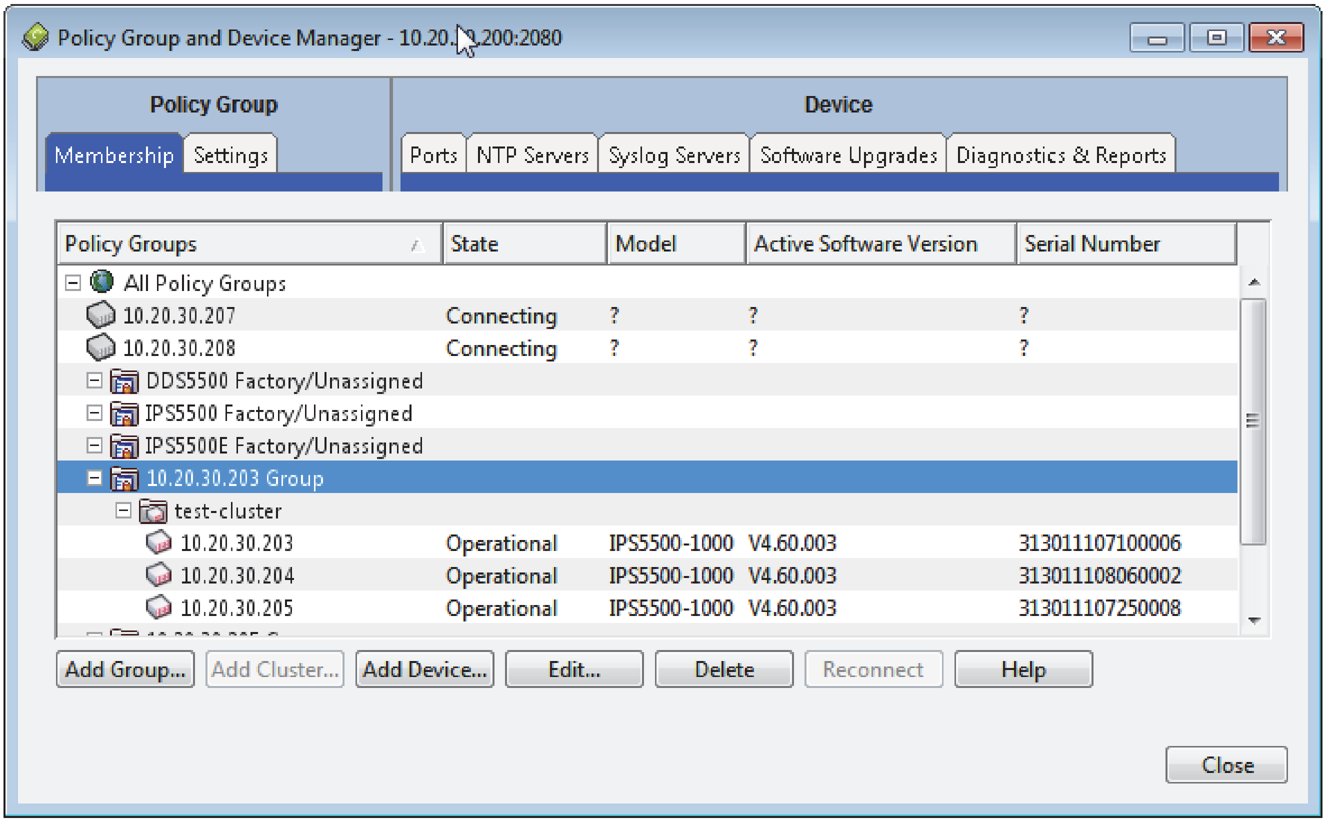

Policy Groups	State	Model	Active Software Version	Serial Number
⊟ All Policy Groups				
10.20.30.207	Connecting	?	?	?
10.20.30.208	Connecting	?	?	?
⊟ DDS5500 Factory/Unassigned				
⊟ IPS5500 Factory/Unassigned				
⊟ IPS5500E Factory/Unassigned				
⊟ 10.20.30.203 Group				
⊟ test-cluster				
10.20.30.203	Operational	IPS5500-1000	V4.60.003	313011107100006
10.20.30.204	Operational	IPS5500-1000	V4.60.003	313011108060002
10.20.30.205	Operational	IPS5500-1000	V4.60.003	313011107250008

The Membership tab displays the information listed in Table 8-1

Table 8-1: Membership Tab Columns

Column	Description
Policy Groups	The Policy Groups tree displays all devices, and lists them under the ProtectionClusters and Policy Groups to which they belong. Click on a group to expand or collapse the group's member list of Corero Network Devices and ProtectionClusters.
State	Indicates the current connection state between the listed Corero Network Device and the IPS Controller. Refer to Table 22-4 for a description of these states.
Model	The model name and number of this Corero Network Device.
Active Software	The current software revision installed on the Corero Network Device.
Serial Number	The device's unique serial number.

The buttons at the bottom of the Membership tab are described in Table 8-2.

Table 8-2: Membership Tab Buttons

Button	Use	For More Information, See...
Add Group	Enables you to add a policy group to the IPS Controller.	Creating Policy Groups (page 8-13)

Table 8-2: Membership Tab Buttons *(Continued)*

Button	**Use**	**For More Information, See...**
Add Cluster	Enables you to create a ProtectionCluster.	Creating a ProtectionCluster Using the IPS Controller (page 14-14)
Add Device	Enables you to add an IPS or DDS Unit to the IPS Controller	Management Prerequisites (page 7-2)
Edit	Depending on what item you select in the policy group tree, this button enables you to edit a policy group, ProtectionCluster, or Corero Network Device.	Modifying the Membership of a Policy Group (page 8-15) Managing an Existing ProtectionCluster (page 14-17) Viewing and Modifying Corero Network Device Settings (page 7-5)
Delete	Depending on what item you select in the policy group tree, this button enables you to delete a policy group, ProtectionCluster, or Corero Network Device.	Deleting a Policy Group (page 8-17) Deleting a ProtectionCluster (page 14-21) Deleting a Corero Network Device (page 7-8)
Reconnect	Enables you to re-establish a connection between the IPS Controller and a Corero Network Device.	Reconnecting to a Corero Network Device (page 7-6)
Help	Launches the IPS Controller online help system.	Using the Online Help (page 5-8)

Viewing Policy Group Settings

The IPS Controller informs you if the settings in the policy group have been changed and are out of synch with the settings on policy group members by placing an icon in the Policy Group tree and an icon and a warning message beside the listing for the device that is out of sync with the policy group.

To view policy group settings, go to the Settings tab of the Policy Group and Device Manager dialog box by doing one of the following:

- Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Settings tab.
- From the menu bar, choose Manage > Policy Groups > Settings.

The Settings tab displays (Figure 8-2).

Figure 8-2: Settings Tab

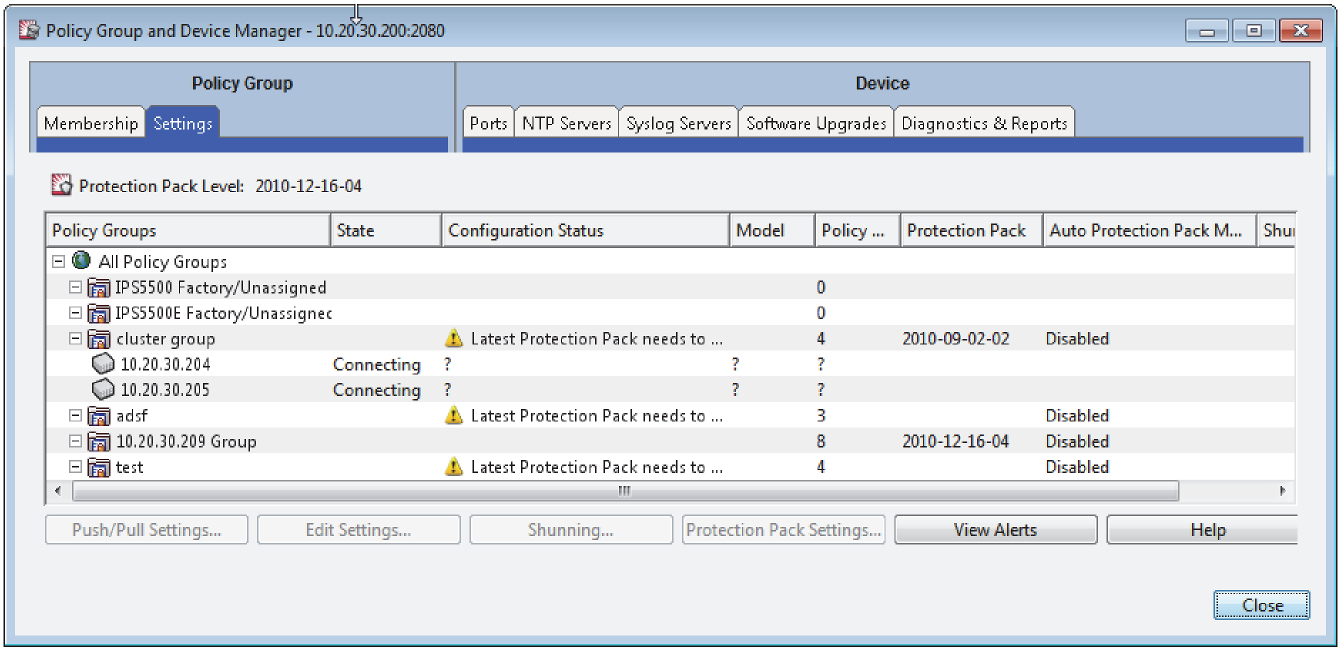

Protection Pack Level: 2010-12-16-04

Policy Groups	State	Configuration Status	Model	Policy ...	Protection Pack	Auto Protection Pack M...	Shu
All Policy Groups							
IPS5500 Factory/Unassigned				0			
IPS5500E Factory/Unassignec				0			
cluster group		Latest Protection Pack needs to ...		4	2010-09-02-02	Disabled	
10.20.30.204	Connecting	?	?	?			
10.20.30.205	Connecting	?	?	?			
adsf		Latest Protection Pack needs to ...		3		Disabled	
10.20.30.209 Group				8	2010-12-16-04	Disabled	
test		Latest Protection Pack needs to ...		4		Disabled	

Protection Pack information displays at the top of the tab. This information lists the most recent protection pack version that was downloaded from Corero's secure TopResponse server.

The table on the Settings tab displays the information listed in Table 8-3

Table 8-3: Settings Tab Columns

Column	Description
Policy Groups	The Policy Groups tree displays all devices, and lists them under the ProtectionClusters and Policy Groups to which they belong. Click on a group to expand or collapse the group's member list of Corero Network Devices and ProtectionClusters.
State	Indicates the current connection state existing between the listed Corero Network Device and the IPS Controller. Refer to Table 7-1 for a description of these states.
Configuration Status	Displays status information for the group, ProtectionCluster, or device listed in the table.

Table 8-3: Settings Tab Columns *(Continued)*

Column	Description
Model	The type and model of the IPS or DDS Unit.
Policy Rev	This IPS Controller-generated value indicates how many times the settings for the policy group have been updated. This number increases each time you either apply a protection pack or use the Edit Settings button to access the policy settings windows, make changes, and save them. When you notice a member of a policy group has a lower policy revision than the policy group to which it belongs, this indicates the member will need to have the current security policy settings pushed (delivered) to it from the IPS Controller.
Protection Pack	Indicates whether a Protection Pack has been applied to a policy group, and if one has been installed, displays the release level of the Protection Pack.
Auto Protection Pack Mode	Indicates whether protection packs are automatically downloaded and applied to the policy group, automatically downloaded to the policy group but not applied, or automation is disabled and protection packs must be acquired manually.
Shun Sync Status	Shun sync column displays status information describing whether the shunning settings on the Corero Network Device are synchronized with the current settings on the IPS Controller. This synchronization is not triggered by pushing settings, rather, it is automatically performed whenever the shun configuration changes. This column displays one of three status messages: • In Sync The shunning settings on the Corero Network Device have been synchronized with the current settings on the IPS Controller. • Syncing The shunning settings on the Corero Network Device are in the process of being synchronized with the current settings on the IPS Controller • Out of Sync The shunning settings on the Corero Network Device are not currently in synch with the current settings on the IPS Controller. This could be because the current policy has not yet been pushed to the device, or the device is not currently connected.

The buttons at the bottom of the Membership tab are described in Table 8-2.

Table 8-4: Settings Tab Buttons

Button	Use	For More Information, See...
Push/Pull Settings	Manually push settings and protection packs to a policy group, or manually pull settings from a Corero Network Device to the IPS Controller.	About Pushing and Pulling Policy Group Settings (page 8-9) How to Push Policy Group Settings (page 8-10) How to Pull Policy Group Settings (page 8-11)
Edit Settings	Change the security settings for the selected policy group. These settings enable you to modify firewall, intrusion protection, rate-based, and environmental settings for the policy group.	Chapter 15, "About Security Policies" Chapter 16, "Managing FW+IPS Security Policies"

Table 8-4: Settings Tab Buttons *(Continued)*

Button	Use	For More Information, See...
Shunning	Enables you to view and manage shun labels in order to shun (block) malicious IP addresses.	About Using IP Address Shunning to Stop an Attack (page 23-4) Shunning IP Addresses (page 23-6)
Protection Pack Settings	Enables you to apply a protection pack to the Corero Network Device and view protection pack application settings.	Manually Downloading and Delivering a Protection Pack (page 9-5) Configuring Protection Pack Settings (page 9-6)
View Alerts	Displays a searchable list of IPS Controller alerts.	Viewing the IPS Controller Management Alert Table (page 21-5)
Help	Launches the IPS Controller online help system.	Using the Online Help (page 5-8)

About Pushing and Pulling Policy Group Settings

One of the primary benefits of the IPS Controller is providing centralized management of Corero Network Devices. The IPS Controller tracks all configuration and settings changes on the managed devices and alerts you if the settings on the device and the settings currently stored on the IPS Controller are out of sync.

When these settings are out of sync, the IPS Controller enables you to upload (pull) security settings from a device, or download (push) security settings from the IPS Controller to policy groups and ProtectionClusters.

The IPS Controller enables you to manage policy group settings in two ways:

- Pull - Settings on a Corero Network Devices are acquired (pulled) from an IPS or DDS Unit to the IPS Controller. The primary reason for pulling settings from a device is to duplicate those settings in other members of a policy group.
- Push - Settings modified on the IPS Controller are taken from the current IPS Controller configuration and delivered (pushed) to the specified device(s) or policy group(s).

NOTE

The IPS Controller software is unable to push settings to a Corero Network Device until you have applied an initial protection pack to that device.

You can also use these features in sequence. For example, you can pull the settings from one device, and push them out to an entire policy group.

The IPS Controller requires you to push settings to policy groups and ProtectionClusters whenever the settings on the IPS Controller differ from those on the group or cluster. For example, this can happen when you have acquired the latest Protection Pack; the system automatically alerts you to the fact that you need to push updated settings all devices in policy group or ProtectionCluster.

You perform both the push and pull operations using the left-most button (labeled Push Settings or Pull Settings, depending on the selected device, ProtectionCluster, or policy group) at the bottom of the Settings tab on the Policy Group and Device Manager dialog box. The label of the button will change depending on the context, making the desired action available.

The button label changes as follows:

- When no items are selected in the Policy Groups tree, the button is unavailable (grayed out) and labeled Push/Pull Settings.
- When you have selected a single Policy Group or a ProtectionCluster in the Policy Groups tree, the button is labeled Push Settings.
- When you have selected a single device in the Policy Groups tree, the button is labeled Pull Settings.

How to Push Policy Group Settings

When you download or deliver policy settings from the IPS Controller to a group of Corero Network Devices, this process is called a push. Pushing policy settings takes the changes you have made to devices in ProtectionClusters or policy groups on the IPS Controller, and sends them out to the selected Corero Network Devices, where they take effect immediately. In addition to policy settings, the IPS Controller also enables you to push protection packs from the IPS Controller to groups of devices.

> NOTE
>
> Device-specific settings, such as port, bypass, and time settings, apply to an individual Corero Network Device rather than a policy group. These settings can be made on the IPS Controller, then pushed to the device.

When the policy settings for a device need to be pushed from the IPS Controller to the device, you will see a Configuration Status message on the Settings tab indicating that a push is required. Once the push operation completes and the security configuration is updated on the device, the IPS Controller removes the warning status message.

When pushing a policy group's settings, consider the following:

- You can only push settings to a single policy group or ProtectionCluster. You cannot push policy settings to a single device, to multiple devices that do not comprise a single group or cluster, or to multiple groups or clusters.
- You cannot push policy settings down to a policy group or ProtectionCluster until *after* one TopResponse protection pack has been applied to it. For information about applying protection packs, see Chapter 9, "Using and Managing TopResponse".
- The IPS Controller can comfortably push configuration settings to up to eight Corero Network Devices at a time.
- You can only push settings to devices that are available (running and connected to the IPS Controller).

> NOTE
>
> If a member of a policy group is unavailable when settings are pushed to that group, the IPS Controller management application will indicate that a push is required once that device is available.

To push a current configuration settings out to a selected policy group or ProtectionCluster:

1. Go to the Settings tab of the Policy Group and Device Manager dialog box by doing one of the following:
   - Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Settings tab.
   - From the menu bar, choose Manage > Policy Groups > Settings.
2. From the Settings tab of the Policy Group and Device Manager window, select the policy group or ProtectionCluster to which you want to push settings.
3. Click Push Settings to initiate the download.
4. When prompted, confirm the operation. The IPS Controller pushes the current configuration information out to each device individually. Once the security configuration is updated on a device, the IPS Controller removes the warning status message from the Settings tab.

How to Pull Policy Group Settings

When you upload policy settings from a device to the IPS Controller, this process is called a pull. Pulling policy settings brings the current settings on a Corero Network Device up to the IPS Controller, so the controller has the latest information about the device settings. Pulling settings is often intentionally done in preparation for pushing those same settings to another policy group.

If you need to pull settings from a device, the device will display a status message indicating that the device configuration was locally modified, and that a push or pull operation is required in order to synchronize the configuration settings across all members of that policy group.

When pulling a device’s settings to the IPS Controller, consider the following:

- You can only pull settings from a single device.
- You con only pull settings from a device that is available (running and connected to the IPS Controller).

To pull a device's current configuration settings into the IPS Controller:

1. Go to the Settings tab of the Policy Group and Device Manager dialog box by doing one of the following:
    - Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Settings tab.
    - From the menu bar, choose Manage > Policy Groups > Settings.
2. From the Policy Group Settings tab of the Policy Group and Device Manager window, select the device from which you want to pull settings.
3. Click Pull Settings to initiate the upload.
4. When prompted, confirm the operation. The IPS Controller pulls configuration information from the device. Once the pull operation completes and the device’s security configuration is updated on the controller, the controller removes the warning status message.

Planning Policy Groups

One of the main functions of the IPS Controller is to enable you to group Corero Network Devices in a logical manner so you can easily manage a common set of security settings. These device groups are called policy groups. You will determine the organizational structure for policy group members for your site. Note that a device can only belong to one policy group at a time. There are two typical ways for organizing Corero Network Devices into groups that share common security settings: by location and by function. Table 8-5 describes different ways you might choose to organize your devices.

Table 8-5: Policy Group Device Organization

Category	Organization	Description
Location	Physical Location	Organizing Corero Network Devices by physical location means creating groups of devices that are co-located. For example, several Corero Network Devices in a particular main office could be grouped together based on the need to allow considerable flexibility to the local administrators due to a large variety of internal network uses. You could group these devices together, push out the initial policy group settings, then turn off the IPS Controller's access to these units (by setting the Corero Network Devices to Block the IPS Controller from Managing this Device) and then allow the local administrators to make needed security policy changes. In this example, the local administrator would have to apply Protection Packs to individual Corero Network Devices, since the IPS Controller would not be able to push out the updates to devices not under its control.
	Geographic Location	A company has nine branches. Six of those branches have Corero Network Devices with similar traffic requirements and security conditions, but also have less-than-expert administrators. You create a policy group and strictly control the devices in these branches, pushing out identical policies and updating them from your central office. You also push out Protection Packs and software updates from the central office.
Function	Edge Devices	You have several Corero Network Devices that serve as edge devices in various offices throughout your enterprise. You want these devices maximized to deal with rate-based issues such as SYN flood controls and connection limits for certain services. You create an Edge Device policy group and customize the policy for the rate-based functions as they apply to these edge devices, and you turn off deep packet inspections for these Corero Network Devices. Then you continue to manage, monitor, and update these devices using the IPS Controller.
	Internal Network Support	You have several Corero Network Devices that are protecting servers within your enterprise network, but some units protect a combination of HTTP, mail, and FTP servers while other devices focus on SQL servers. You set up two policy groups and customize them to focus on these two different sets of requirements.
	Varying Needs for Access Control	You have certain internal networks that should be allowed virtually unlimited outbound traffic, and other internal networks that should only be allowed Web, DNS, FTP, and Mail services. You create two policy groups: High Security Outbound and Medium Security Outbound, then customize each group’s security policies to fit these security needs.
	Special ProtectionCluster Traffic Requirements	Throughout your enterprise you have several instances where multiple Corero Network Devices are installed as ProtectionClusters for high availability. You create a High Availability policy group and then create ProtectionClusters for each HA instance. You would do the same for Corero Network Devices installed in a cluster to handle high traffic volumes.

Creating Policy Groups

Once you have planned your policy groups as described in Planning Policy Groups (page 8-12), you can create them. In order to create a policy group you first need to acquire the initial security settings for the policy group. Afterwards, you must specify the devices that will become the group's members.

There are two ways to add a policy group to the policy group tree:

- Creating a Policy Group Based on a Copy of an Existing Group (page 8-13)
- Creating a Policy Group by Pulling the Settings from a Corero Network Device (page 8-14)

Creating a Policy Group Based on a Copy of an Existing Group

To create a policy group based on a copy of an existing group:

1. Go to the Membership tab of the Policy Group and Device Manager dialog box. To do so, do one of the following:
   - Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Membership tab.
   - From the menu bar, choose Manage > Policy Groups > Membership.
2. To add a group, from the Policy Group Membership tab, click the Add Group button. The Add Group dialog box displays.
3. From the Policy Group to Copy drop-down window, select the group you wish to use as the basis for your new policy group.
4. In the Name field, enter a name for the new policy group.
5. Select the desired mode from the Auto Protection Pack Mode drop-down. Available settings are listed in Table 8-6.

Table 8-6: Automatic Protection Pack Settings

Setting	Description
Disabled	Protection packs will not be automatically applied to the selected policy groups on the IPS Controller. They will need to be manually applied using the Apply Protection Pack button on this screen. You may choose to manually apply protection packs if your company or industry requires that you only apply protection packs at specific times, or on specific dates. This is the default setting.
Apply	Protection packs are automatically applied to the policy group on the IPS Controller, but you will need to manually push these settings to the Corero Network Devices that comprise the policy group.
Apply and Push	Once downloaded, protection packs are automatically applied to the policy group. Also, provided all IPS or DDS Units in the policy group have the current policy group configuration, the protection pack is automatically pushed to all policy group members.

6. When finished, click OK.

When you have finished, the new policy group displays in the Policy Groups tree view of the management interface, and also displays the new group in any other IPS Controller management sessions currently running.

Creating a Policy Group by Pulling the Settings from a Corero Network Device

At times you may want to create a new policy group based on the security settings of a specific Corero Network Device.

> NOTE
>
> During this process, the template device is moved to the newly created group. If you do not wish the template device to belong to the new group, you can move it to a different group at the end of the process.

To create a policy group by pulling the settings from a connected Corero Network Device:

1. Go to the Settings tab of the Policy Group and Device Manager dialog box. To do so, do one of the following:
   - Click the Policy Group & Device Manager toolbar button. When the dialog box displays, click the Settings tab.
   - From the menu bar, choose Manage > Policy Groups > Settings.
2. Select (highlight) an operational IPS or DDS Unit that you want to use as a security settings template for the new policy group.
3. Click Pull Settings. Examine the confirmation window and click Yes to confirm the operation.

   The IPS Controller creates a new policy group using the ID of the selected Corero Network Device, and the new policy group has security settings identical to those of the device.
4. At this time, if desired, you can:
   - Rename the policy group, as described in Changing the Name of a Policy Group (page 8-16).
   - Move the Corero Network Device back to its original policy group, or to a different policy group containing similar model devices. For more information, see Modifying the Membership of a Policy Group (page 8-15).

Modifying the Membership of a Policy Group

At times, you may need to change the Corero Network Device membership of a policy group, either by adding new members, or by moving current members to either a user-defined policy group or a default policy group.

When modifying the membership of a policy group, remember the following:

- You perform this task on a per-device basis.
- All members of a single policy group must be the same Corero Network Device model type (IPS 5500, IPS5500-E, or DDS 5500).
- You cannot add individual ProtectionCluster members to a policy group; you must add the ProtectionCluster as a whole.
- You can only move devices to an existing policy group. If you want to move a device to a new policy group, you must create the new policy group before you can move the device.

NOTE

If you attempt to move a Corero Network Device to a policy group whose members are a different model (product family), or, if the IPS Controller has not yet determined the device's model (has not yet connected to the device), you will receive an Incompatible Product Family error.

To modify policy group membership:

1. To access the Membership tab on the Policy Group and Device Manager dialog box, do one of the following:
   - Click the Policy Group and Device Manager tool bar button. When the dialog box displays, click the Membership tab.
   - From the menu bar, choose Manage > Policy Groups > Membership.
2. From the Policy Groups tree, select the Corero Network Device whose policy group membership you wish to change, then click Edit. The Edit Device dialog box displays.
3. Select the desired policy group from the Policy Group drop-down list, then click OK.

NOTE

Once you have moved a Corero Network Device to a new policy group, you will typically want to push the policy group settings for the new policy group to the moved device. For more information on how to do this, see How to Push Policy Group Settings (page 8-10).

Changing the Name of a Policy Group

If you are modifying the constituency or purpose of the Corero Network Devices that comprise a policy group, you may want to change the name of the policy group to reflect this.

To change the name of a policy group:

1. To access the Membership tab on the Policy Group and Device Manager dialog box, do one of the following:
   - Click the Policy Group and Device Manager tool bar button. When the dialog box displays, click the Membership tab.
   - From the menu bar, choose Manage > Policy Groups > Membership.
2. From the Policy Groups tree, select the policy group whose name you wish to change, then click Edit. The Edit Group window displays.
3. In the Name field, enter the new name for this policy group, then click OK.

NOTE

You can also use this screen to modify the protection pack mode for this policy group. For additional instructions on how to do this, see Configuring Protection Pack Settings (page 9-6).

Deleting a Policy Group

At times, you may want to delete a policy group. Note that you can only delete an empty policy group, so if you want to delete a policy group, you must remove all members before doing so, as described in the procedure below.

To delete a policy group:

1. To access the Membership tab on the Policy Group and Device Manager dialog box, do one of the following:
   - Click the Policy Group and Device Manager tool bar button. When the dialog box displays, click the Membership tab.
   - From the menu bar, choose Manage > Policy Groups > Membership.
2. From the policy group tree, select the policy group you plan to delete and expand it to see if the policy group has any members. You can only delete an empty policy group.
3. If the policy group contains any ProtectionClusters, you must delete (dissolve) them before you can proceed. For information on deleting a ProtectionCluster, see Deleting a ProtectionCluster (page 14-21).
4. If the policy group contains any individual Corero Network Devices, you must either move them to a different policy group or delete them from the IPS Controller.
   - For information about moving devices to a different policy group, see Modifying the Membership of a Policy Group (page 8-15).
   - For information about deleting devices from the IPS Controller (so they are unmanaged), see Deleting a Corero Network Device (page 7-8)
5. Once the policy group is empty, click Delete.
6. You are asked to confirm the deletion. Click Yes to delete the policy group, or click No to retain it.

Chapter 9
Using and Managing TopResponse

A major component of Corero Network Security's protection is Corero's subscription-based TopResponse Threat Update Service. This service provides automated updates, technical support, security advisory and software subscription services, along with access to Corero's Knowledge Base and special delivery programs.

TopResponse updates are made available to customers through their IPS Controller. The software subscription service delivers frequent protection pack updates and security advisories. TopResponse Protection Packs offer proactive protection from zero-day threats and resolution to security issues.

These protection packs can be automatically or manually downloaded to the IPS Controller from Corero's secure TopResponse server. From the IPS Controller, protection packs can be automatically or manually delivered to your Corero Network Devices. The TopResponse service will not remove or alter user-specified security features, such as defined address filters, address ranges, client groups, signatures, and so forth.

CAUTION

> If you use the IPS Controller to manage and update your IPS or DDS Units, do not use Corero's standalone TopResponse Update Manager software. This can cause unsynchronized and inconsistent protection pack delivery, which may result in suboptimal protection.

This chapter contains the following sections:

TopResponse Protection Pack Overview

The TopResponse Automatic Update Service is a comprehensive technical support service that provides IPS customers with advanced security support services to maximize the security, availability, and performance of their network. Specifically, TopResponse delivers protection pack updates that contain items such as new signatures and new security rules. Protection packs provide protection against new security threats, and contain technical support information and security advisories. Protection packs are downloaded from Corero's secure TopResponse server.

In order to keep your Corero Network Devices configured with the latest protection packages, the IPS Controller must have network access to Corero's secure TopResponse Server. The IPS Controller must also be able to connect to Corero units within your network.

Protection packs are released at regular intervals, and during particularly dynamic attack periods, they can be released several times a day. Note that when protection packs are applied, your system retains all of the settings you have configured for your system; protection packs never overwrite custom settings.

You can download a protection pack to the IPS Controller manually or set up automatic downloads. Once you download a protection pack, you have complete control over the policy groups and individual IPS or DDS Units that receive the changes. You can configure the members of policy group to be automatically updated with protection packs when they are available on the IPS Controller, or you can choose to apply the protection packs manually.

There is a three-step process to obtain and apply a protection pack to one or more Corero Network Devices.

NOTE

You can choose to perform each step in the process either manually or automatically.

1. Download the protection pack from Corero's TopResponse Server to the IPS Controller.
2. Select and apply the latest protection pack to a policy group on the IPS Controller.
3. Push the applied protection pack settings out to the IPS or DDS Units within a policy group.

Viewing Current Protection Pack Information

The IPS Controller management application displays protection pack level and status information for policy groups, ProtectionClusters, and Corero Network Devices, including suggested user actions like “Latest Protection Pack Needs to be Applied”.

Figure 9-1 shows protection pack status information on the Settings tab of the Policy Group and Device Manager dialog box.

Figure 9-1: Device Protection Pack Status

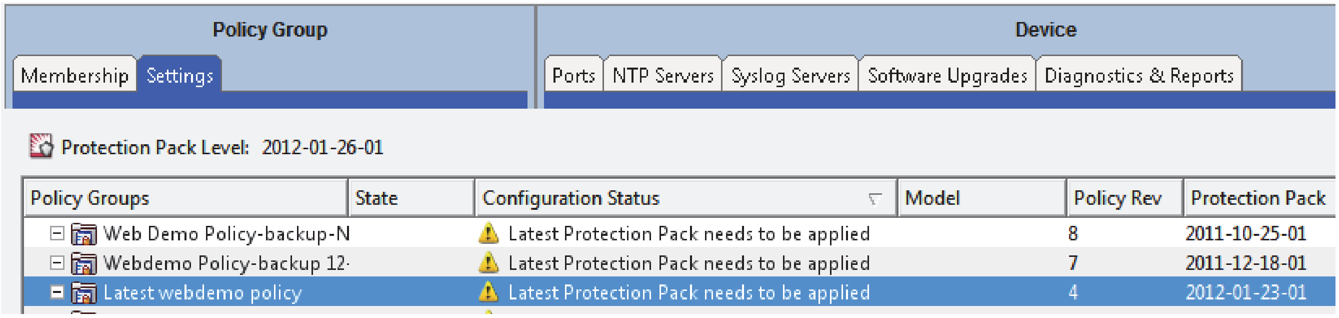

Protection Pack Level: 2012-01-26-01

Policy Groups	State	Configuration Status	Model	Policy Rev	Protection Pack
Web Demo Policy-backup-N		⚠ Latest Protection Pack needs to be applied		8	2011-10-25-01
Webdemo Policy-backup 12·		⚠ Latest Protection Pack needs to be applied		7	2011-12-18-01
Latest webdemo policy		⚠ Latest Protection Pack needs to be applied		4	2012-01-23-01

To view protection pack settings:

1. You can view the current protection pack level of the IPS Controller by doing any of the following:
   - Open the main window of the IPS Controller. You can view the current protection pack level in the display area at the lower left corner. This level reflects the most recent protection pack that was acquired from Corero’s secure TopResponse server.
   - Open the Settings tab on the Policy Group and Device Manager dialog box by choosing Manage > Policy Groups > Settings from the menu bar. You can view the current protection pack level in a display area above the table.
   - Open the TopResponse Update Manager dialog box by clicking the Manage TopResponse toolbar button, or by choosing System > Manage TopResponse Updates from the menu bar.
2. Before you can decide whether you need to perform any protection-pack-related tasks, you must first determine a policy group’s current protection pack status.

   To do this, open the Settings tab on the Policy Group and Device Manager dialog box by doing one of the following:

   - Click the Policy Group and Device Management toolbar button, then click the Settings tab on the Policy Group and Device Manager dialog box.
   - Choose Manage > Policy Groups > Settings from the menu bar.

   You can view the current protection pack level for any device or policy group in the Protection Pack column.

   For those policy groups or Corero Network Devices that do not list the same protection pack as the IPS Controller, the Configuration Status column indicates that the latest protection pack needs to be applied.
3. You can view information about whether protection packs are downloaded, applied, or pushed manually or automatically by looking in the Auto Protection Pack Mode column. Available settings are listed in Table 9-1.

Table 9-1: Automatic Protection Pack Settings

Setting	Description
Disabled	Protection packs will not be automatically applied to the selected policy groups on the IPS Controller. They will need to be manually applied using the Apply Protection Pack button on this screen. You may choose to manually apply protection packs if your company or industry requires that you only apply protection packs at specific times, or on specific dates. This is the default setting.
Apply	Protection packs are automatically applied to the policy group on the IPS Controller, but you will need to manually push these settings to the Corero Network Devices that comprise the policy group.
Apply and Push	Once downloaded, protection packs are automatically applied to the policy group. Also, provided all IPS or DDS Units in the policy group have the current policy group configuration, the protection pack is automatically pushed to all policy group members.

Manually Downloading and Delivering a Protection Pack

Overall, there is a three-step process to obtain and apply a protection pack to one or more Corero Network Devices.

1. Download the protection pack from Corero's TopResponse Server to the IPS Controller.
2. Select and apply the latest protection pack to one or more policy groups.
3. Push the new settings out to the IPS or DDS Units within a policy group.

By default, both download and delivery of protection packs is done manually. You can configure the IPS Controller to automatically download protection packs from the TopResponse server. You can also configure the IPS Controller to automatically apply and push these protection packs to Policy Groups. For detailed instructions how to do this, see Configuring Protection Pack Settings (page 9-6).

NOTE

Corero Network Devices must be online and operational in order to push settings to them.

To manually download and deliver a protection pack to a policy group:

1. Do one of the following:
   - Click the Manage TopResponse toolbar button.
   - Select System > Manage TopResponse Updates from the menu bar.

   The TopResponse Update Manager dialog box displays.
2. To acquire the latest protection pack from Corero's TopResponse server, click Update to Latest.
3. Before applying the protection pack, review information about the new protection pack by clicking View Prot. Pack Notes.
4. Do one of the following:
   - Click the Manage TopResponse toolbar button. Then, in the Update Settings area of the TopResponse Update Manager dialog box, click Modify. The Protection Pack Settings dialog box displays.
   - Click the Policy Group and Device Manager toolbar button. Then click the Settings tab on the Policy Group and Device Manager dialog box, select a policy group, then click Protection Pack Settings. The Protection Pack Settings dialog box displays with the selected policy group highlighted.
5. Select one or more policy groups, then click Apply Protection Pack Now.
6. Return to the Settings tab of the Policy Group and Device Manager dialog box. You can see which policy groups need updated settings because the Configuration Status will indicate "Latest Protection Pack Needs to be Applied".
7. Select the policy groups to which you just applied the protection pack, then click Push Settings.

   The new protection pack is pushed to the selected policy groups.

NOTE

You cannot push settings to individual devices in a ProtectionCluster. You must push settings to the entire cluster.

Configuring Protection Pack Settings

There are two steps to applying protection pack settings to a policy group: applying the settings to the policy group on the IPS Controller, and pushing the settings to the Corero Network Devices that comprise the policy group.

To configure protection pack settings for the IPS Controller:

1. If you like, you can manage update settings. To do so:
   a. Do one of the following
   - Click the Manage TopResponse toolbar button.
   - Choose System > Manage TopResponse Updates from the menu bar.

     The TopResponse Update Manager dialog box displays.

   b. Specify the desired update settings. Update settings are described in Table 9-2.

Table 9-2: TopResponse Update Settings

Setting	Description
Update Method	Select from the following options: • Manual To manually download the latest Protection Pack (set of security protection updates) from Corero's secure TopResponse server, select this option from the TopResponse Update Manager window. • Automatically Update Protection Packs Every Hour The IPS Controller automatically checks for Protection Pack updates every hour. If there is a new Protection Pack, it is downloaded automatically. • Automatically Update Protection Packs Daily At... The IPS Controller automatically checks for Protection Pack updates once a day at the specified time. Select this option and then select the time from the drop-down menu. If there is a new Protection Pack, it is automatically downloaded to the IPS Controller at the specified time of day. **Note:** The automatic application time is specified in Coordinated Universal Time (UTC), also known as Greenwich Mean Time (GMT) time, not the local time of your IPS Controller.
TopResponse Update Server Proxy	If the IPS Controller must use a proxy server to connect with the TopResponse server, you can (optionally) enable the use of a proxy server, then specify the IP address and port of the proxy server. If you are using a secure proxy server, select Enable Proxy Authentication, and, specify the user name and password.

2. To manage protection pack settings, do one of the following:
   - Click the Manage TopResponse toolbar button. Then, in the Update Settings area of the TopResponse Update Manager dialog box, click Modify.
   - Click the Policy Group and Device Manager toolbar button. Then click the Settings tab on the Policy Group and Device Manager dialog box, select a policy group, then click Protection Pack Settings.

   The Protection Pack Settings dialog box displays. Now, you should:

   a. Select the desired policy groups from the Policy Group list.

   b. From the Auto Protection Pack Mode drop-down on the Protection Pack Settings dialog box, select the desired setting. Settings are described in Table 9-1.
3. When you are finished, click OK.

Chapter 10
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 11, “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 10-8).
- 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 10-1.
  - 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 10-1 lists all of the port role types available for Corero Network Devices.

Table 10-1: 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 11-11).
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. 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 10-2 displays a summary of port role features.

Table 10-2: 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 10-2).

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 10-3. Choose the mode that reflects the combination of normal and failure operation that you want to occur.

Table 10-3: 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 10-3: 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 10-2).

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 10-1.

Table 10-4 explains port roles for 5100 Series Corero Network Devices.

Table 10-4: 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)	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 10-5 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 10-1.

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

Table 10-5: 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 10-5: 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 10-1.

Table 10-6 explains port roles for 5200 Series Model 2000 ES devices.

Table 10-6: 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 10-7 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 10-1.

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

Table 10-7: 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 10-7: 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 10-1.

Table 10-8 explains port roles for 5200 Series Model 2000 ESL Corero Network Devices.

Table 10-8: 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 10-9 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 10-1.

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

Table 10-9: 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 10-11).
- 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 10-13).

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 11
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 10, "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 of a Corero Network Device on the IPS Controller:

1. Do one of the following:
   - Click the Policy Group & Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab.
   - From the menu bar, choose Manage > Devices > Ports.

   The Ports tab displays.

2. Select a Corero Network Device, then click View Front Panel.

   The Front Panel View for that Corero Network Device displays.

For a complete description of the Front Panel view and its features, see Using the Front Panel View (page 22-2).

Figure 11-1 shows the Front Panel View for an IPS Unit.

Figure 11-1: IPS Front Panel View

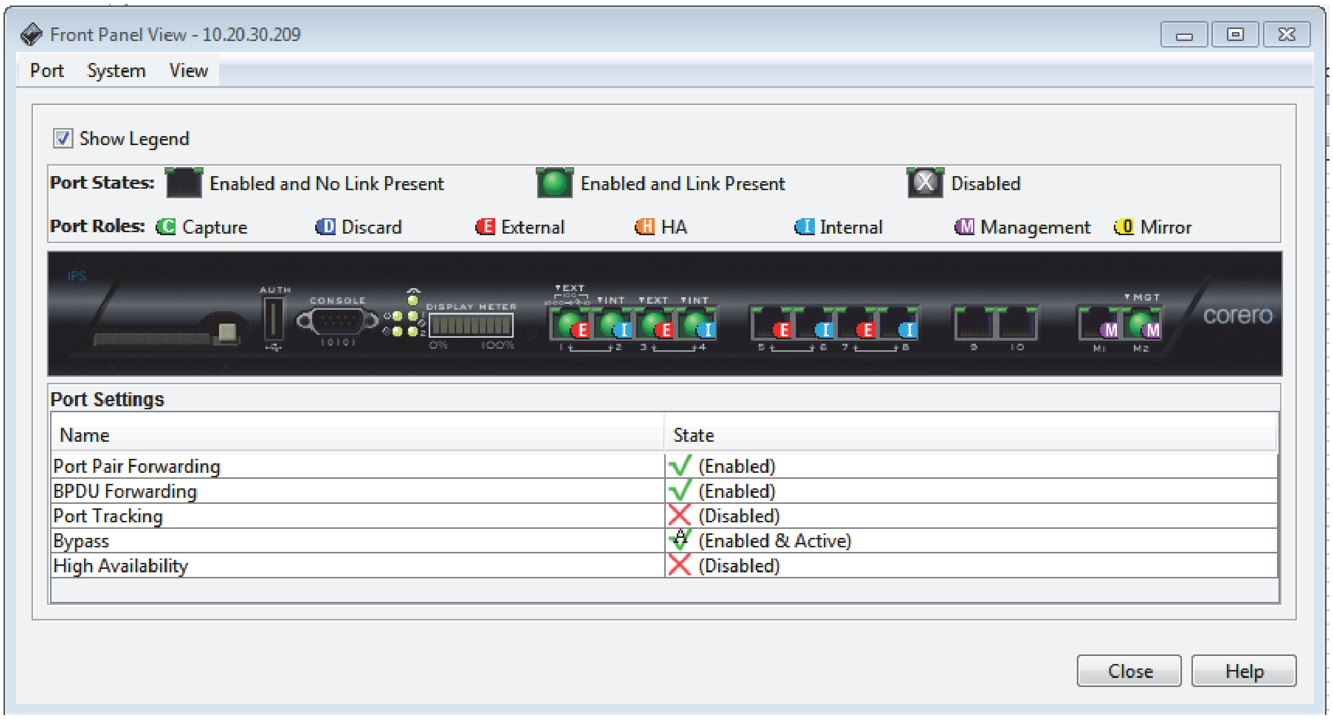

☑ Show Legend

Port States: Enabled and No Link Present   Enabled and Link Present   Disabled

Port Roles: C Capture   D Discard   E External   H HA   I Internal   M Management   O Mirror

Port Settings

Name	State
Port Pair Forwarding	✓ (Enabled)
BPDU Forwarding	✓ (Enabled)
Port Tracking	✗ (Disabled)
Bypass	✓ (Enabled & Active)
High Availability	✗ (Disabled)

Figure 11-2 shows the Front Panel View for a DDS Unit.

Figure 11-2: DDS Front Panel View

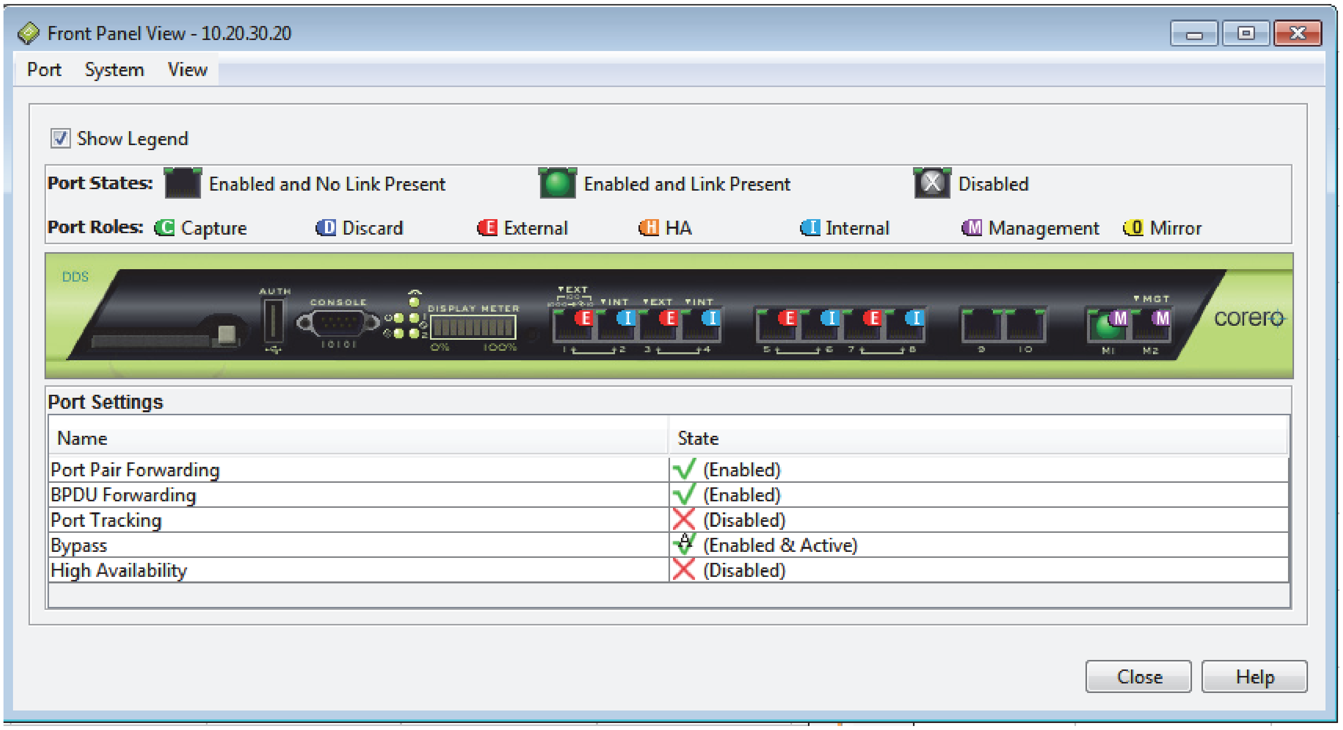

☑ Show Legend

Port States: ■ Enabled and No Link Present ● Enabled and Link Present ☒ Disabled

Port Roles: C Capture D Discard E External H HA I Internal M Management 0 Mirror

Port Settings

Name	State
Port Pair Forwarding	✓ (Enabled)
BPDU Forwarding	✓ (Enabled)
Port Tracking	✗ (Disabled)
Bypass	✓ (Enabled & Active)
High Availability	✗ (Disabled)

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 10, “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.

NOTE

If you are configuring ports on a DDS Unit, the Getting Started Wizard accessible through the DDS management application contains several additional configuration screens beyond those available through the IPS Controller being used to manage the DDS Unit. In order to configure Host Groups, Default Firewall Policies for Web Servers, Default Firewall Policies for DNS Servers, and Rate-Based Protection through the Getting Started Wizard, use the wizard available from the DDS management application.

To modify port settings for a managed Corero Network Device using the Getting Started Wizard in the IPS Controller management interface:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab in the Device area.
   - Choose Manage > Devices > Ports from the menu bar.
2. Select the device whose settings you want to modify, then click Role Wizard. The Getting Started Wizard displays.
3. 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.

4. 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 10-2).

   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 Corero Network Device, 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 11-9)

5. 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 10-8).

   Bypass Settings specify the hardware-based bypass feature between Mission port-pairs. The term bypass indicates that traffic is not being mitigated by the Corero Network Device You can choose whether traffic never bypasses the Corero Network Device, whether traffic always bypasses the device, or whether traffic only bypasses the device when the system is down or being reset. For more information on bypass settings, see Bypass Settings (page 10-6).

   Specify your Mission Port Pair settings, then click Next.

6. 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 Corero Network Device model. For more information on port roles, . For information on specifying a port for capture, see Modifying Traffic Capture Settings (page 11-11)

   CAUTION

   Discard ports use significant system resources. For this reason, discard port use may unintentionally affect mission traffic.

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

7. 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.

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 10-6).

To view port information using the IPS Controller management application:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab in the Device area.
   - Choose Manage > Devices > Ports from the menu bar.
2. Select one or more Corero Network Devices in the Policy Groups tree.
3. Click Port Settings. The Ports dialog box displays (Figure 11-3).

Figure 11-3: IPS Controller Ports Dialog Box

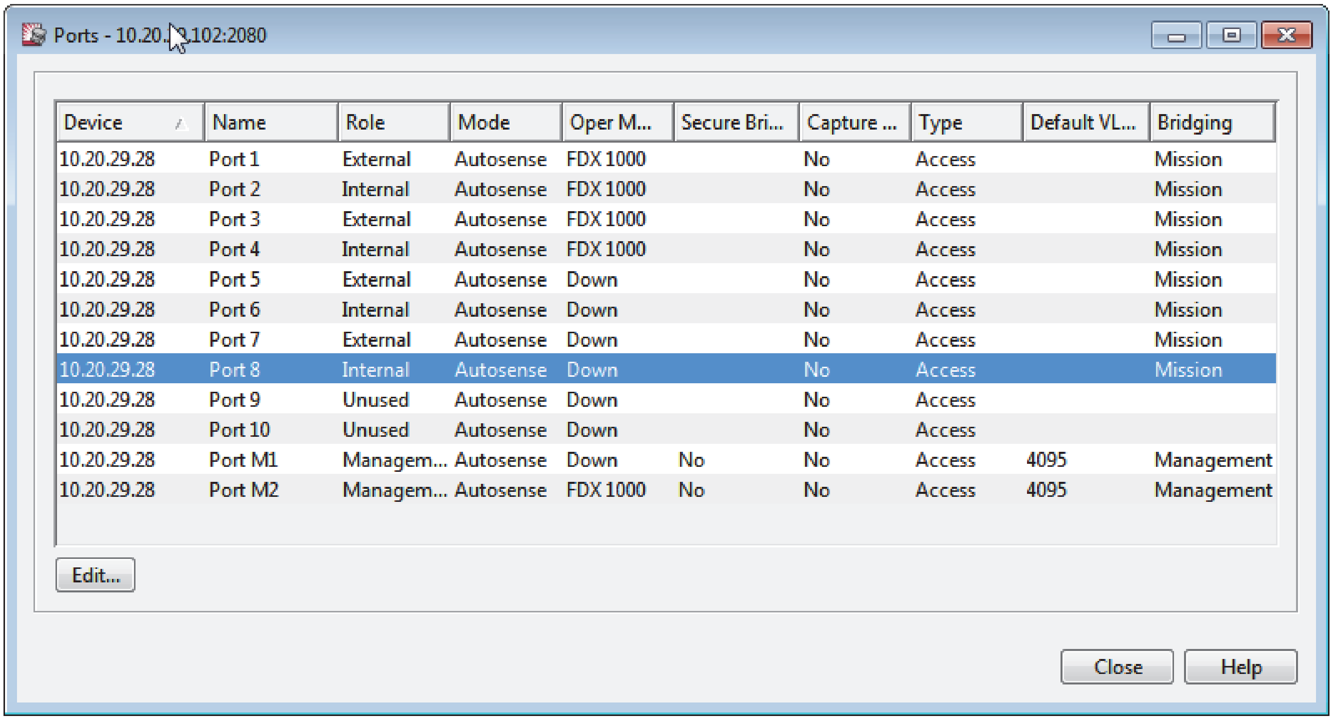

Ports - 10.20.29.102:2080

Device	Name	Role	Mode	Oper M...	Secure Bri...	Capture ...	Type	Default VL...	Bridging
10.20.29.28	Port 1	External	Autosense	FDX 1000		No	Access		Mission
10.20.29.28	Port 2	Internal	Autosense	FDX 1000		No	Access		Mission
10.20.29.28	Port 3	External	Autosense	FDX 1000		No	Access		Mission
10.20.29.28	Port 4	Internal	Autosense	FDX 1000		No	Access		Mission
10.20.29.28	Port 5	External	Autosense	Down		No	Access		Mission
10.20.29.28	Port 6	Internal	Autosense	Down		No	Access		Mission
10.20.29.28	Port 7	External	Autosense	Down		No	Access		Mission
10.20.29.28	Port 8	Internal	Autosense	Down		No	Access		Mission
10.20.29.28	Port 9	Unused	Autosense	Down		No	Access		
10.20.29.28	Port 10	Unused	Autosense	Down		No	Access		
10.20.29.28	Port M1	Managem...	Autosense	Down	No	No	Access	4095	Management
10.20.29.28	Port M2	Managem...	Autosense	FDX 1000	No	No	Access	4095	Management

Table 11-1 summarizes the information in the Ports dialog box.

Table 11-1: Ports Dialog Box

Column	Description
Device	The IP address of the selected Corero Network Device.
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.
Role	Indicates the assigned function for this port. For more information about port roles, see Port Role Overview (page 10-2).
Mode	Indicates the port's transmit/receive speed/mode setting. For more information, see Viewing and Modifying Port Settings (page 11-6).
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 13, "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.

4. In the Ports dialog box, select the port you want to edit, then click Edit. The Edit Port Settings dialog box displays.
5. 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.

6. 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 13-3).

   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 13, "Advanced Port Configuration".

7. 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 IPS Controller management application:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab in the Device area.
   - Choose Manage > Devices > Ports from the menu bar.
2. Select one or more Corero Network Devices in the Policy Groups tree.
3. Click Port Pairs. 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).
4. 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.

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 10-6).

To modify bypass settings using the IPS Controller management application:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab in the Device area.
   - Choose Manage > Devices > Ports from the menu bar.
2. Select one or more Corero Network Devices in the Policy Groups tree.
3. Click Bypass Settings. The Bypass Settings dialog box displays.
4. 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.

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.

NOTE

Although you can use the IPS Controller management application to configure (allocate) a capture port using the Getting Started Wizard, you cannot enable port capture (activate) or specify a port to capture (designate) from the IPS Controller management application. You can only accomplish this from the management application for a Corero Network Device, as described below.

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 12
Managing Users

The IPS Controller maintains 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 12-1.

Table 12-1: 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 12-3), and some are global, as described in Configuring Global User Security Settings (page 12-7).

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 12-7).
- 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 12-2.

Table 12-2: 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 12-3.

Table 12-3: 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 12-7).

To manage users:

1. Choose System > Users > Manage from the menu bar.

   The Management Users dialog box displays (Figure 12-1).

Figure 12-1: Management Users Dialog Box

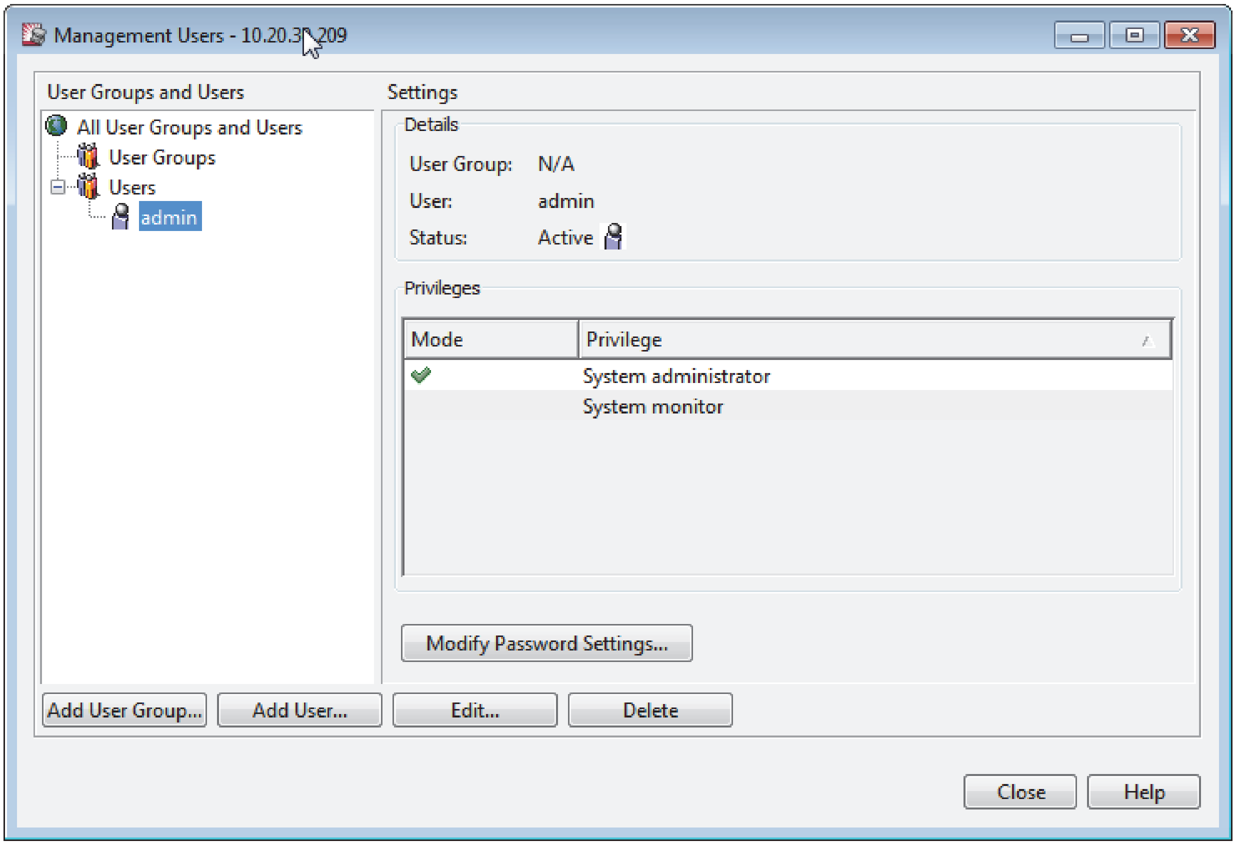

Management Users - 10.20.3 209

User Groups and Users

- All User Groups and Users
  - User Groups
  - Users
    - admin

Settings

Details

User Group: N/A

User: admin

Status: Active

Privileges

Mode	Privilege
✓	System administrator
	System monitor

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 12-2. For more information on this feature, see Configuring Global User Security Settings (page 12-7).

   d. Select the user's privilege level. Privilege levels are described in Table 12-3.

   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 12-2).

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.

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 12-3).

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. Choose System > Users > Manage from the menu bar.

   The Management Users dialog box displays (Figure 12-1).

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 12-3).
   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.

> 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. Choose System > User > Settings from the menu bar.

   The User Security Settings dialog box displays.

2. Modify the settings as desired, then click OK. The settings are described in Table 12-4.

Table 12-4: 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 13
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 13-1.

Figure 13-1: VLAN Forwarding Algorithm

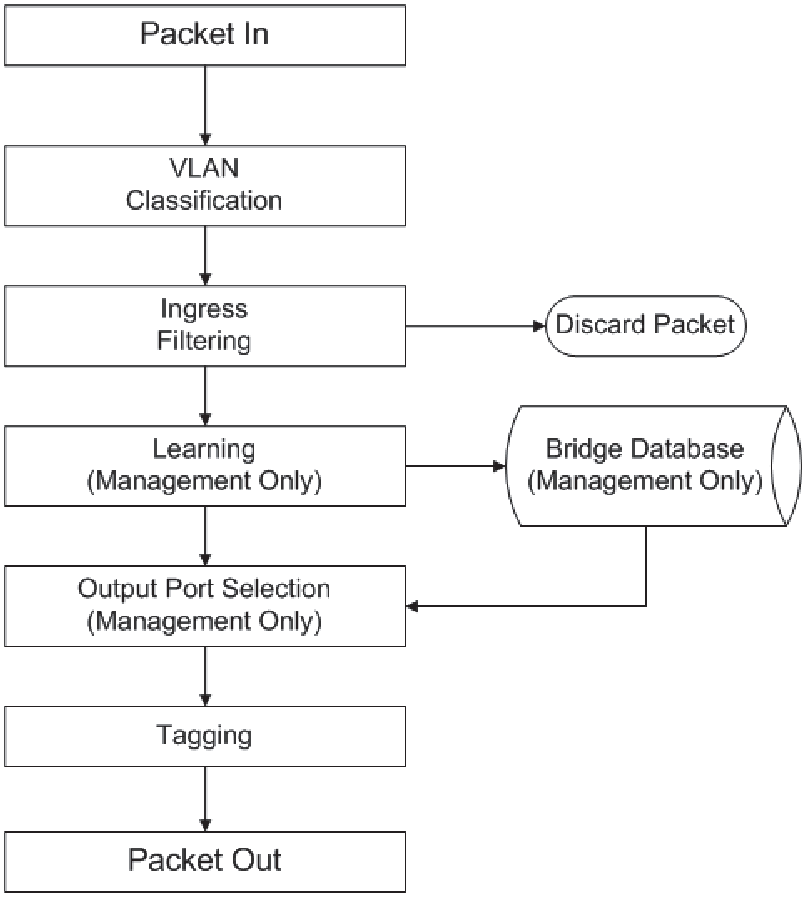
Packet In
VLAN Classification
Ingress Filtering
Discard Packet
Learning (Management Only)
Bridge Database (Management Only)
Output Port Selection (Management Only)
Tagging
Packet Out

The forwarding algorithm steps shown in Figure 13-1 are described in Table 13-1.

Table 13-1: 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 13-1: 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 10, “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 13-4). 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 13-4).

- 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 13-4).

Changing Management Entity VLAN ID

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

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**

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 13-7).

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 B-2).

To view and modify the VLAN ID for a device-specific management port using the IPS Controller management application:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab in the Device area.
   - Choose Manage > Devices > Ports from the menu bar.
2. Select one or more Corero Network Devices in the Policy Groups tree.
3. Click Port Settings. The Edit Port Settings dialog box displays.
4. Select a Management Port and click Edit.
5. 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 11, ‘‘Viewing and Configuring Ports”.

6. When finished, click OK.

Chapter 14
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:

CAUTION

You cannot modify the device membership of an existing ProtectionCluster.
In order to change the membership of an existing ProtectionCluster, you must dissolve (delete) the existing ProtectionCluster, then create a new ProtectionCluster with the desired list of cluster members. This process has significant system and network consequences. For information and considerations on deleting and recreating ProtectionClusters, see Deleting a ProtectionCluster (page 14-21) and Creating a ProtectionCluster Using the IPS Controller (page 14-14).

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 14-1 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 14-1: 5100-Series High Availability Ports: 5 Through 8

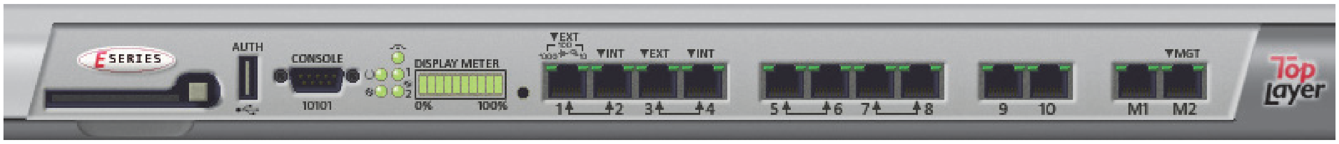

5200-Series hardware requires that you interconnect two HA ports between ProtectionCluster members. Ports 5 and 6 are dedicated to high availability, as shown in Figure 14-2.

Figure 14-2: 5200-Series High Availability Ports: 5 and 6

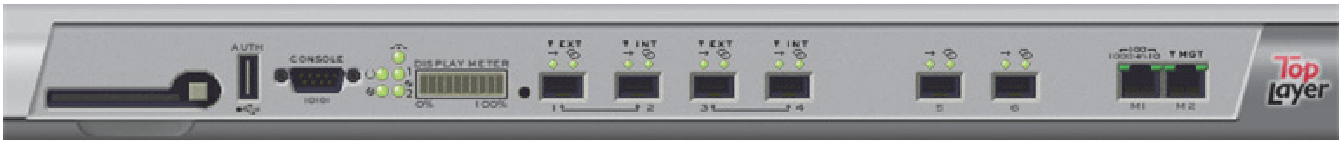

High Availability ProtectionCluster Configurations

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 14-3 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 14-3: Existing Customer High Availability Configuration

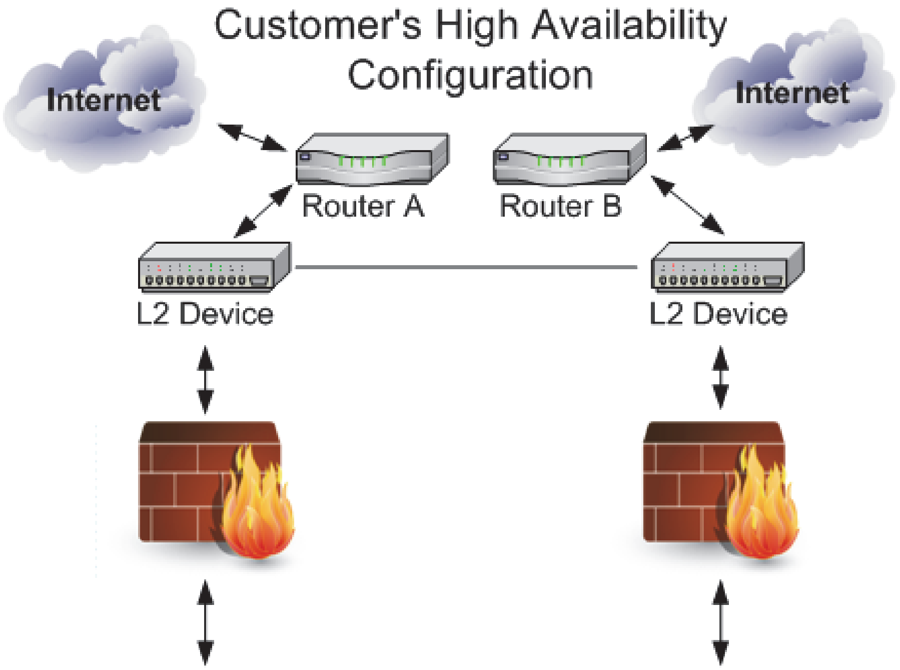

Figure 14-4 depicts two Corero Network Devices inserted into an existing redundant network configuration. The devices are deployed inline in the network.

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 14-4 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 14-4: Dual Inline High Availability Configuration

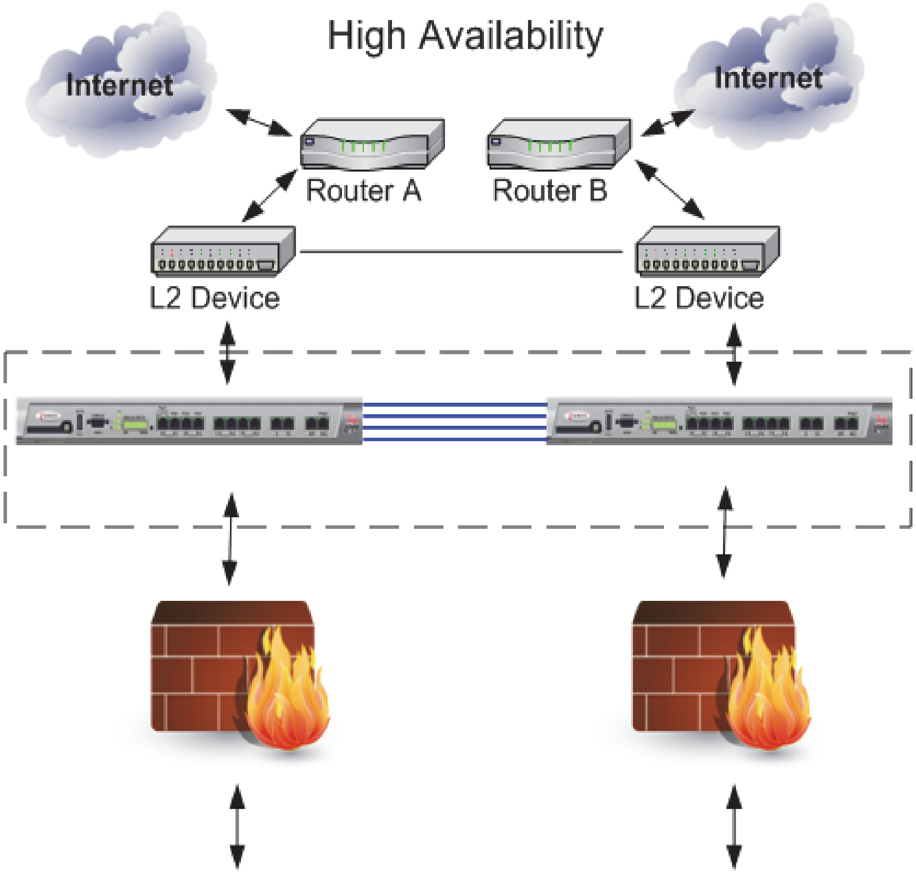

High Capacity ProtectionCluster Configurations

You can also use a ProtectionCluster to increase processing power and inspection throughput on your network. In the example shown in Figure 14-5, 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 14-5: High Capacity 2000 ES and ESL Configuration

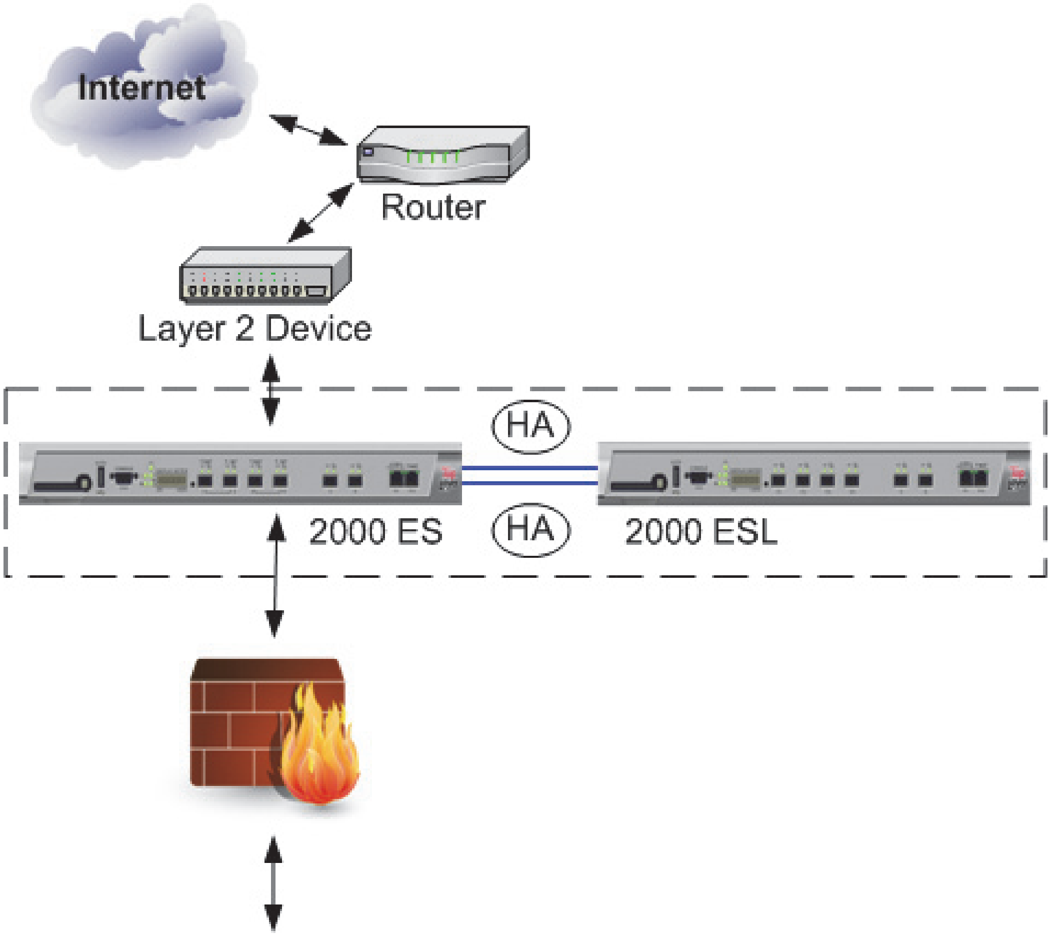

If you want to implement a dual-inline configuration for high capacity, and you have an existing Corero Network Device 5200 Model 2000 ES, you can either add a Model 2000 ES or a Model 2000 ESL to your existing unit. However, if you are installing 5200 units for the first time, you can purchase the 5200 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 14-6 shows the connections required between a Model 2000 ES DDS Unit and a Model 2000 ESL DDS Unit.

Figure 14-6: DDS 5500 Model 2000 ES and Model 2000 ESL Interconnections

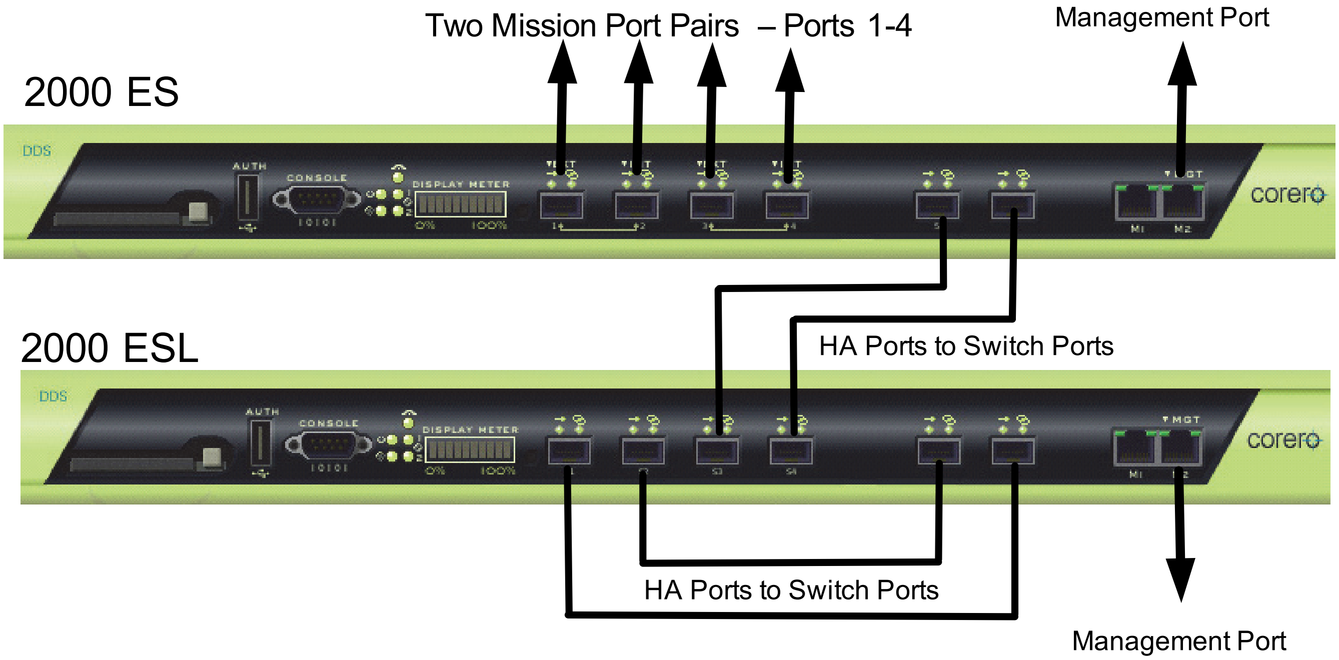

Figure shows how two IPS 5500 Model 2400 ES units could be interconnected in a ProtectionCluster configuration. You could also connect a single Model 2400 ES unit an additional 2000 ES unit.

Figure 14-7 shows how two DDS 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 14-7: DDS 5500 2400 ES to 2400 ES Interconnections

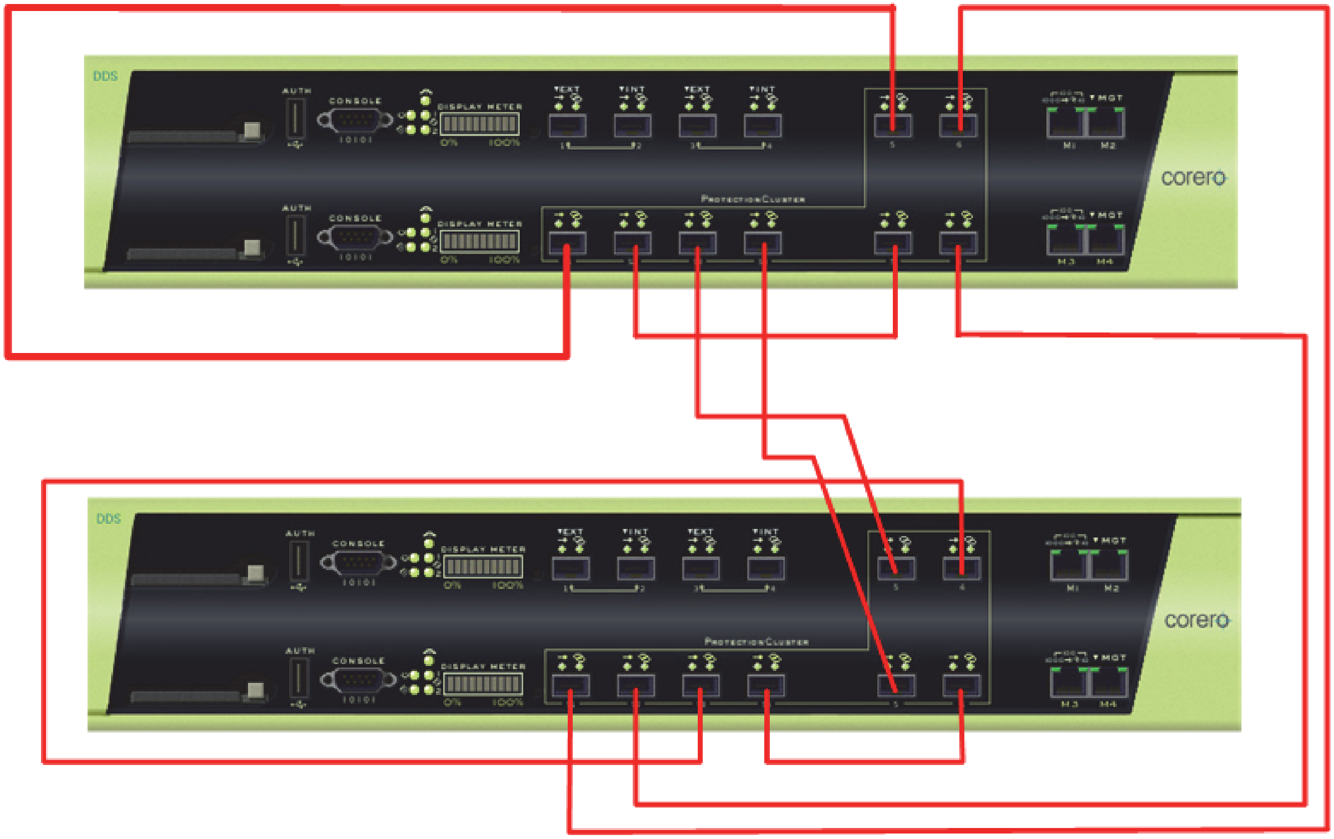

Figure 14-8 shows how a high-throughput ProtectionCluster comprised of two older IPS 5500 Model 2400 ES units could be used for perimeter defense in a customer network. You could use the same configuration for two DDS 5500 Model 2400 ES units.

Figure 14-8: High-Throughput Perimeter Defense ProtectionCluster Configuration

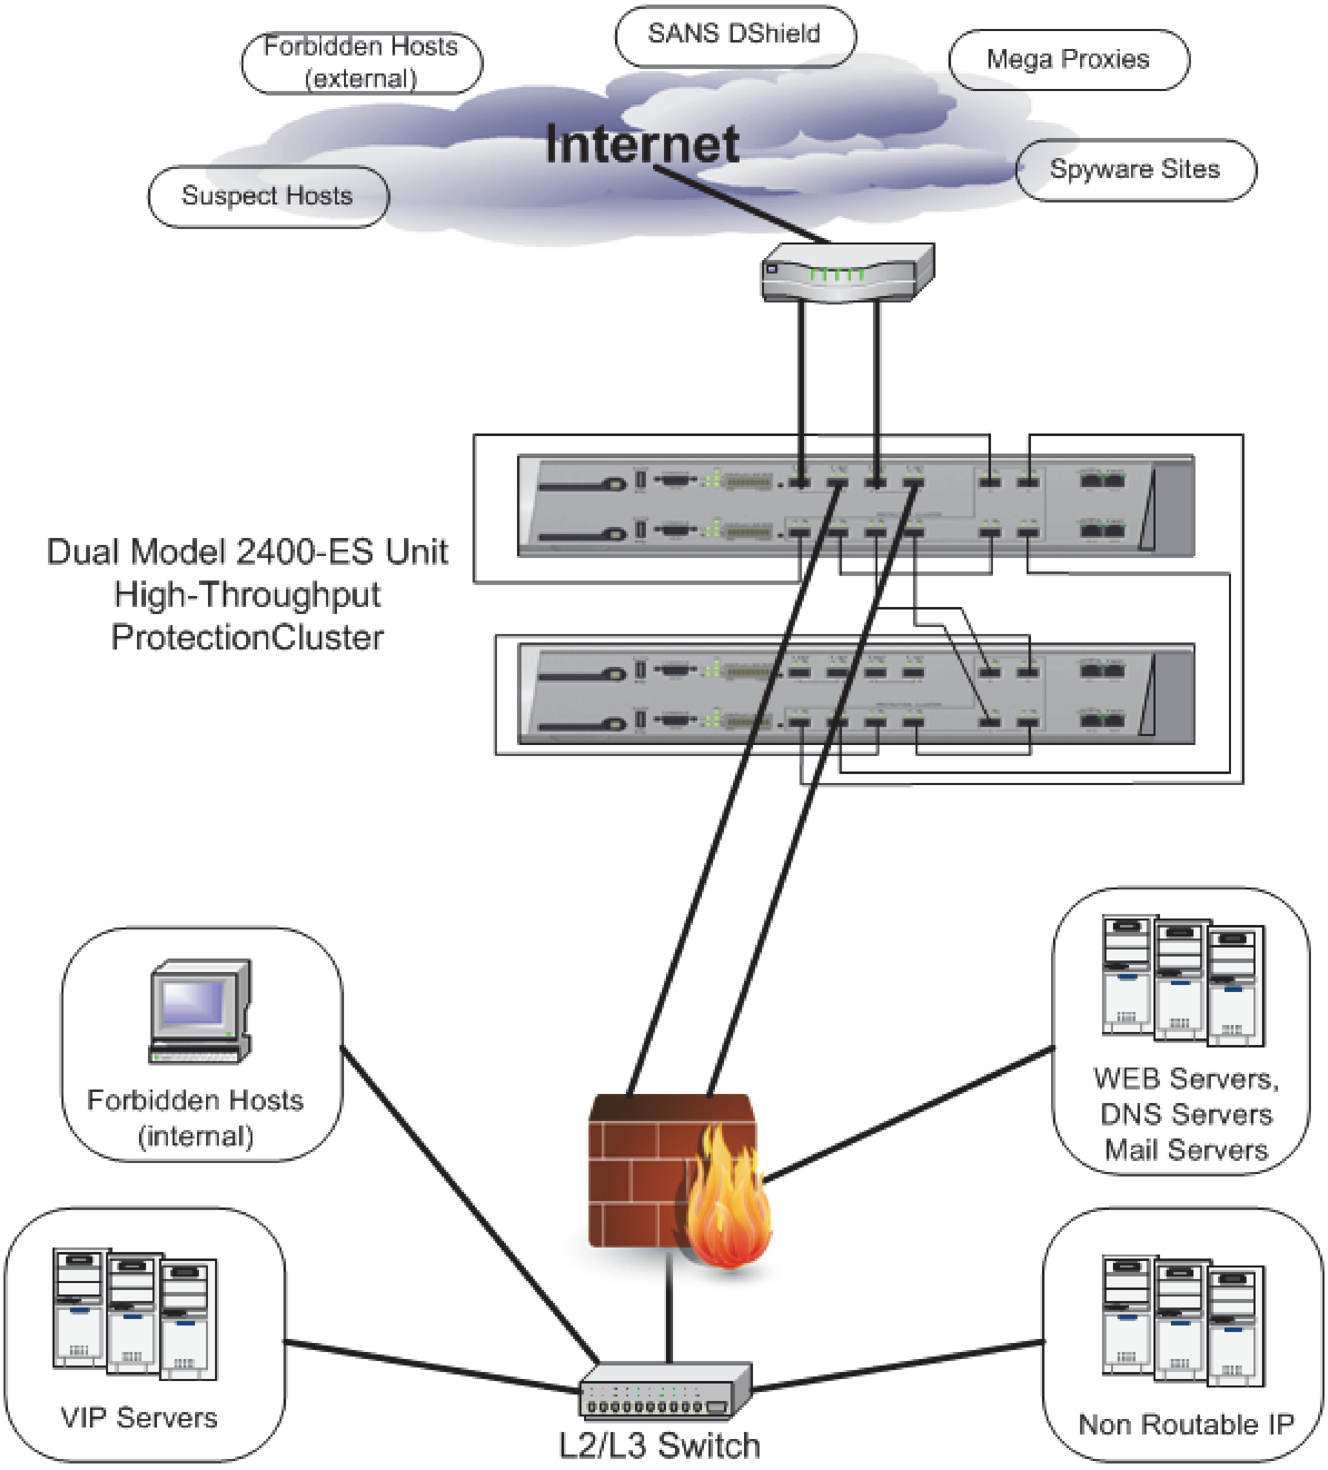

Viewing ProtectionCluster Status

You can view the current status of a ProtectionCluster through the IPS Controller management application. If there are issues with a ProtectionCluster, the management application generates alerts.

> NOTE
>
> For assistance with troubleshooting ProtectionCluster issues, see Appendix B, “Troubleshooting IPS Controller Issues”.

ProtectionCluster members communicate with one another over High Availability (HA) links. Every cluster member is connected to every other cluster member by two HA links in newer models, and four HA links in older models.

To view ProtectionCluster Status:

1. The management interface also visually flags problem clusters and devices with an alert icon. ⚠ ☒
   A yellow alert icon indicates a warning, and a red alert icon indicates a critical issue.
2. If you see an alert icon on a ProtectionCluster device, there are two ways you can get additional information:
   - Right-click the device, and choose View Alerts from the pop-up menu. The Management Alert Table displays, showing only alerts associated with the selected device.
   - Click the Policy Group & Device Manager toolbar button. Click the Settings tab on the Policy Group and Device Manager page. Detailed configuration status information will display for ProtectionCluster and cluster members.
3. Once you are aware there is an issue, you can view additional information in the IPS Controller management application.

   From the menu bar, choose Monitor > ProtectionCluster Status. The ProtectionCluster Status dialog box displays information in the form of a table. This information is described in Table 14-1.

Table 14-1: ProtectionCluster Status Columns

Column	**Description**
Policy Groups	The expanded policy group navigational tree.
State	The current state of each ProtectionCluster member. Available states include: • Connecting - The IPS Controller attempts to connect to the IP address you specified for the Corero Network Device. • Authenticating - Once connected, the IPS Controller uses the shared management jet to authenticate itself with the device. • Synchronizing - Once authenticated, the IPS Controller attempts to synchronize with the device by capturing a copy of all settings or parameters on the device. • Operational - Once synchronized, the device becomes operational. The IPS Controller is now able to manage changes to device settings and monitor the device’s security events.
HA Status	The High Availability status of the ProtectionCluster member. Status messages include: • ? (question mark) - The ProtectionCluster member could not be reached, so the current status is not known. • Active with peers - The ProtectionCluster member is operational, and is able to communicate with one other cluster member. • Active with no peers - The ProtectionCluster member is operational, but is unable to communicate with any other cluster member.

Table 14-1: ProtectionCluster Status Columns *(Continued)*

Column	Description
Size	The total number of members in the ProtectionCluster.
Peers	The number of other members currently available in the ProtectionCluster. When all members are available, this number is one less than the cluster size. The number of other members should be identical on all members of the cluster. A 0 (zero) in this column indicates that there this ProtectionCluster member cannot communicate with other peers in the cluster.

4. To view additional information about a single ProtectionCluster, select the cluster label in the Policy Groups column, then click Details. The ProtectionCluster Device Status dialog box displays status information about the connections between the ProtectionCluster members:
   - Green (with an up arrow) visually indicates the link is active (available).
   - Red (with a down arrow) visually indicates the link is not active (unavailable).
5. You can view operational status information for ProtectionCluster members by clicking the Policy Group & Device Manager toolbar button and viewing the contents of the Membership tab.

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 14-2

Table 14-2: 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 ProtectionCluster Using the IPS Controller

The simplest way to create a ProtectionCluster is to use the IPS Controller management application.

> NOTE
>
> On rare occasions, you may find you need to add a pre-existing ProtectionCluster (one that was previously unmanaged by an IPS Controller) to an IPS Controller. If you need to do this, please refer to Managing an Existing ProtectionCluster (page 14-17) for complete instructions.

To create a ProtectionCluster using the IPS Controller management application:

1. Verify that the Corero Network Devices you want to include in the ProtectionCluster meet the following criteria:
   - Ensure all of the devices you want to include in the ProtectionCluster are the same model, are running the same version of software, and are members of the same Policy Group.
   - Ensure all devices you want to include in the ProtectionCluster have been installed, fully configured, are available for proper cluster configuration, and are be connected to the management network.
   - Ensure all devices you want to include in the ProtectionCluster are being managed by the IPS Controller, and display a status of Operational.
   - In addition to the usual network and management connections, members of a ProtectionCluster must be connected to one another through their High Availability (HA) ports. When creating a ProtectionCluster, you must physically connect the high-availability (HA) ports between all devices. If you are clustering two devices, connect the high availability ports point-to-point.

   > NOTE
   >
   > For detailed information on the ProtectionCluster member cabling process, refer to the Hardware Installation Guide for your Corero Network Device.

   If you are connecting three or more Corero Network Devices, you must do so using a high-speed switch. This switch must provide a jumbo-frame-tolerant L2 path between all of the devices in the ProtectionCluster, and must provide non-blocking interconnect performance, adequately buffered for large bursts.

   > NOTE
   >
   > The four HA Interconnect ports built in to the Model 2000 ESL are designed specifically for this purpose, meeting the product-specified switching requirements. For a list of approved switches, contact Corero.

2. On the IPS Controller, click the Policy Group & Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
3. Click the Membership tab.
4. From the Policy Groups tree, expand the group containing the devices you want to add to the new ProtectionCluster.
5. Hold down the Ctrl key and select each of the devices you want added to the ProtectionCluster. The Add Cluster button becomes available for selection.
6. Click Add Cluster. The Add Cluster dialog box displays.
7. Verify that you have selected the desired cluster members for the ProtectionCluster, and provide a cluster name.

8. From the Health-Check Frequency drop-down, select the health interval to use for this ProtectionCluster. If you are unsure what frequency to select, use the default value.
9. Click OK. The IPS Controller reboots the selected units, configures the ProtectionCluster, and displays it in the Policy Groups tree.

CAUTION

After configuring a cluster, all of the devices in the cluster will automatically reboot. This is required in order for clustering to properly function.

10. Go to the Settings tab on the Policy Group and Device Manager dialog box. You must ensure that the Policy Group revision number is different than the revision number for the devices that are members of the new ProtectionCluster before proceeding.

    If necessary, perform a minor edit, such as modifying the comment associated with a policy row.

CAUTION

If you attempt to push the policy from a policy group to devices that have the same policy revision number, the policy information will not be pushed and you will need to restart the IPS Controller from an SSH session into the machine hosting the IPS Controller software.

11. Once you have verified that the revision number for the policy group containing the cluster is later than the revision number for the new cluster members, push the policy out to the new ProtectionCluster. Note that pushing initial policy information to a ProtectionCluster can take more than five minutes.
12. To view the status of the ProtectionCluster, at any time, choose Monitor > ProtectionCluster Status on the IPS Controller menu bar.

NOTE

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

Moving a ProtectionCluster to a New Policy Group

The only modification you can make to a ProtectionCluster is to reassign the ProtectionCluster to a different Policy Group. To do this:

1. Click the Policy Group and Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
2. On the Policy Group and Device Manager dialog box, select the Policy Group Membership tab.
3. From the Policy Groups tree, expand the policy group containing the ProtectionCluster you want to move.
4. Select the name of the ProtectionCluster, then click Edit. The Edit Cluster dialog box displays.
5. In the Edit Cluster dialog box, select the policy group to which you would like the ProtectionCluster to belong, then click OK.
6. If needed, push the settings from the new policy group out to the ProtectionCluster.

Managing an Existing ProtectionCluster

On rare occasions, you may find you need to add a pre-existing ProtectionCluster to an IPS Controller. You may need to do this if:

- You are installing an IPS Controller for the first time at a site that contains one or more ProtectionClusters that were not previously managed by an IPS Controller.
- If you are re-installing an IPS Controller, and are replacing a pre-existing IPS Controller that managed one or more ProtectionClusters.

> NOTE
>
> If you want to create a ProtectionCluster out of IPS or DDS Units that are not already in a functioning ProtectionCluster, refer to Creating a ProtectionCluster Using the IPS Controller (page 14-14).

The process used to identify a pre-existing, already functioning, ProtectionCluster to the IPS Controller, and then manage the newly identified cluster, is different from creating a cluster from unassociated devices in two ways:

- When you add the devices to the IPS Controller, the controller realizes that they are already in a ProtectionCluster and places them in an unassigned policy group (as it usually does) but automatically defines them as members of a ProtectionCluster.
- Instead of creating a policy group for these devices by copying an existing policy group, you will mostly like want to pull the configuration from one of the devices in the ProtectionCluster to create a policy group that has a starting point that is identical to the devices in the cluster.

To identify a ProtectionCluster composed of previously associated Corero Network Devices, create a policy group from one of the devices, and move the devices into the new policy group (or a different group), complete the following steps:

1. Click the Policy Group and Device Management toolbar button. The Policy Group and Device Manager dialog box displays.
2. Select the Membership tab.
3. Click Add Device. The Add Device dialog box displays.
4. Enter the IP address, port, and shared management key for the first device you are adding. Repeat this process with all other members of the ProtectionCluster.

   The IPS Controller does three things:

   - Automatically identifies each Corero Network Device you add as belonging to an existing ProtectionCluster
   - Places the devices into one of the unassigned policy groups (based on the device's type)
   - Places the devices into a ProtectionCluster.
5. Create the policy group you want to use to hold the ProtectionCluster as described in Planning Policy Groups (page 8-12). You can do this two ways:
   - Create a policy group by pulling the settings from one of the devices in the pre-existing cluster.
     If you create the policy group by pulling the settings from one of the devices, and you wish to rename the policy group, see Modifying the Membership of a Policy Group (page 8-15) for detailed instructions.
   - Create a policy group based on the settings from another policy group.
6. To move the entire ProtectionCluster into the new policy group, click on one of the devices in the cluster and then click Edit Device. The Edit Device dialog box displays.

7. From the Policy Group drop-down menu, select the policy group to which you would like to add the ProtectionCluster, then click OK.
8. To ensure that all Corero Network Devices in the ProtectionCluster have the same security settings, push the policy group settings to all the devices in the cluster. For information on pushing policy group settings, see About Pushing and Pulling Policy Group Settings (page 8-9).

Importing ProtectionClusters

If you have a ProtectionCluster that was created by a version of the IPS Controller that was below 4.10, and the cluster members are running a Corero software release below V4.40, after you upgrade the IPS Controller software, you must import the ProtectionClusters using a special procedure.

NOTE

Please contact the Customer Services Center if you require assistance with this process.

1. Log in to each Corero Network Device and ensure you are running code to V4.40.026 or greater. Each device must be running this minimum code version in order to proceed.
2. Place each Corero Network Device (IPS or DDS unit) in the ProtectionCluster into “Always Bypass” mode either through its local management application or by using the IPS Controller (Version 3.20 or lower).

   Save this setting.
3. Do one of the following:
   - For a two-unit ProtectionCluster, dissolve the cluster configuration by disabling High Availability on each Corero network Device. To do this, choose Configure System > High Availability from the management interface navigation tree, then clearing (deselecting) the Enable High Availability check box.
   - For ProtectionClusters that have more than two units assigned, log in to the IPS Controller (Version 3.20 or lower) if that is currently managing the Corero Network Devices and delete the ProtectionCluster. Once the cluster is dissolved, the devices will automatically reboot.
4. Verify that each Corero Network Device and check to make sure that it is no longer in a ProtectionCluster using the State Browser utility to verify that the units are no longer part of an active cluster. Contact the Customer Services Center for more information on this procedure.
5. Save the configuration and reboot all Corero Network Devices that were members of the ProtectionCluster you dissolved.

   CAUTION

   If you have **not** dissolved the existing ProtectionCluster and rebooted the Corero Network Devices and you have already added the ProtectionCluster members to the current version of the IPS Controller, stop what you are doing immediately and contact the Corero Customer Services Center at +1 978 212 1500 before proceeding.
6. Log in to each Corero Network Device after it has rebooted and configure it with the shared key that will enable it to be managed by the IPS Controller. For more information on shared keys, refer to the Configuration and Management Guide for your Corero Network Device.
7. Verify that the new IPS Controller has been installed correctly on the designated system.
8. Log in to the new IPS Controller and add all of the Corero Network Devices that you will be using to create the cluster. The IPS Controller connects to each Corero Network Device and places them in the unassigned Policy Group.
9. Using the IPS Controller management interface, pull the configuration from each newly-added Corero Network Device. On import, the IPS Controller creates a new Policy Group for each Corero Network Device, using the IP Address of the device as the name of the Policy Group.
10. Pick one of the Policy Groups and give it a name of your choosing that complies with your site's naming convention.
11. Pull the configuration from the Corero Network Device that you added to the new Policy Group.
12. Move all other cluster members into the Policy Group you created.

13. Push the policy configuration to all Corero Network Device in the new Policy Group.

    NOTE

    Note that only policy information is pushed to each Corero Network Device. None of the device information (LAN port settings, etc) will be changed. Corero recommends that, if needed, you make any required changes to the device configuration of each Corero Network Device before proceeding to the next step.

14. Reboot all of the ProtectionCluster members.
15. After the Corero Network Device have finished rebooting, in the IPS Controller management application, select all of the Corero Network Device that you want to configure into a cluster, then click Add Cluster.

    All of the highlighted Corero Network Device will be placed in to a cluster under the Policy Group name.

    When this process completes, all of the Corero Network Device will automatically reboot.

The ProtectionCluster is now operational and can be managed from this version of the IPS Controller.

Deleting a ProtectionCluster

You cannot delete a device from a ProtectionCluster, you can only delete (dissolve) the ProtectionCluster as a whole and recreate it with the desired membership.

Deleting a ProtectionCluster is a complex task. If you choose to delete a ProtectionCluster, consider the following:

- Deleting a ProtectionCluster should only be done using the IPS Controller, and it should always be done with great care.
- Deleting a cluster that is servicing redundant links in an asymmetric routing environment may cause networking problems.
- Cluster deletion is typically done prior to redeploying devices. If this is the case, the devices must first be taken out of the in-line environment prior to deleting the cluster.
- Deleting a ProtectionCluster produces the following effects on the Corero Network Devices that were in the ProtectionCluster:
  a. Causes the Corero Network Devices to stop functioning as a coordinated group and to become stand-alone Corero Network Devices.
  b. Causes the Corero Network Devices that comprised the cluster to immediately reboot.
  c. After the reboot, each device restarts as a stand-alone Corero Network Device. The devices will remain in the policy group to which the cluster previously belonged.

Note that after you delete a ProtectionCluster, if you wish, you can use the former cluster members in a new ProtectionCluster.

To delete a ProtectionCluster:

1. Click the Policy Group and Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
2. On the Policy Group and Device Manager dialog box, select the Policy Group Membership tab.
3. From the Policy Groups tree, expand the policy group containing the ProtectionCluster you want to delete.
4. Select the name of the cluster you want to delete.
5. Click the Delete button. You are prompted to confirm the action.

   At this time, all members of the ProtectionCluster will reboot.

Chapter 15
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 (page 15-2)
- Elements of a Security Policy (page 15-4)
- Elements of a Firewall + IPS Security Policy (page 15-9)
- Default FW+IPS Policy Operation (page 15-12)
- Elements of a Rate Based Security Policy (page 15-14)

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 23, "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 15-1.

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

Table 15-1: 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 • 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 15-1 shows how policies are used to mitigate traffic.

Figure 15-1: Three Dimensional Protection

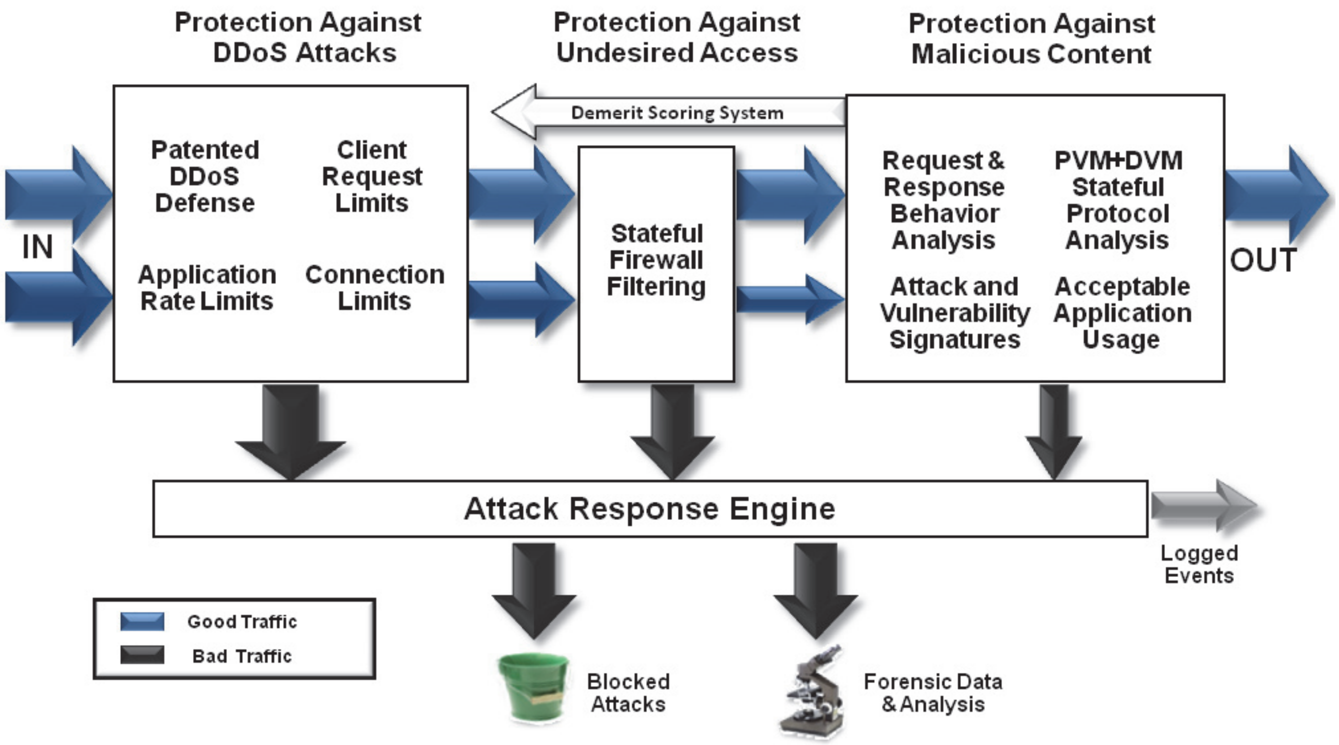
Protection Against DDoS Attacks
Protection Against Undesired Access
Protection Against Malicious Content
Demerit Scoring System
IN
Patented DDoS Defense
Client Request Limits
Application Rate Limits
Connection Limits
Stateful Firewall Filtering
Request & Response Behavior Analysis
PVM+DVM Stateful Protocol Analysis
Attack and Vulnerability Signatures
Acceptable Application Usage
OUT
Attack Response Engine
Logged Events
Good Traffic
Bad Traffic
Blocked Attacks
Forensic Data & Analysis

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 15-2 describes the parts that comprise a security policy.

Table 15-2: 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.
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 15-2), 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 15-2: FW+IPS Tab Columns Indicate Policy Components

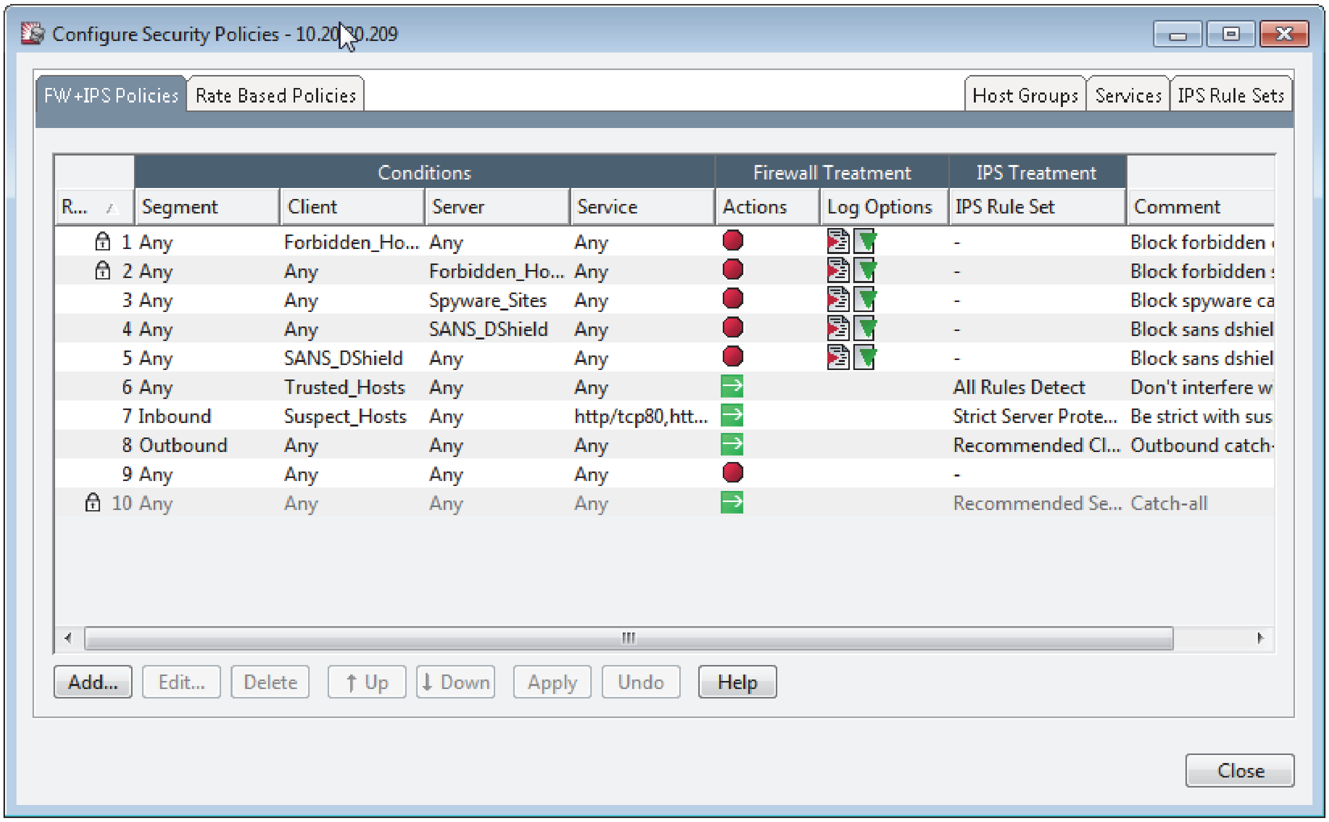

	Conditions				Firewall Treatment		IPS Treatment	
R...	Segment	Client	Server	Service	Actions	Log Options	IPS Rule Set	Comment
🔒 1	Any	Forbidden_Ho...	Any	Any	●	▤ ▼	-	Block forbidden
🔒 2	Any	Any	Forbidden_Ho...	Any	●	▤ ▼	-	Block forbidden
3	Any	Any	Spyware_Sites	Any	●	▤ ▼	-	Block spyware ca
4	Any	Any	SANS_DShield	Any	●	▤ ▼	-	Block sans dshiel
5	Any	SANS_DShield	Any	Any	●	▤ ▼	-	Block sans dshiel
6	Any	Trusted_Hosts	Any	Any	→		All Rules Detect	Don't interfere w
7	Inbound	Suspect_Hosts	Any	http/tcp80,htt...	→		Strict Server Prote...	Be strict with sus
8	Outbound	Any	Any	Any	→		Recommended Cl...	Outbound catch-
9	Any	Any	Any	Any	●		-	
🔒 10	Any	Any	Any	Any	→		Recommended Se...	Catch-all

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

- Segments (page 15-5)
- Host Groups (page 15-6)
- Services (page 15-7)
- Rules (page 15-7)
- Rule Sets (page 15-8)

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 15-3.

Figure 15-3: Select Segments Dialog Box

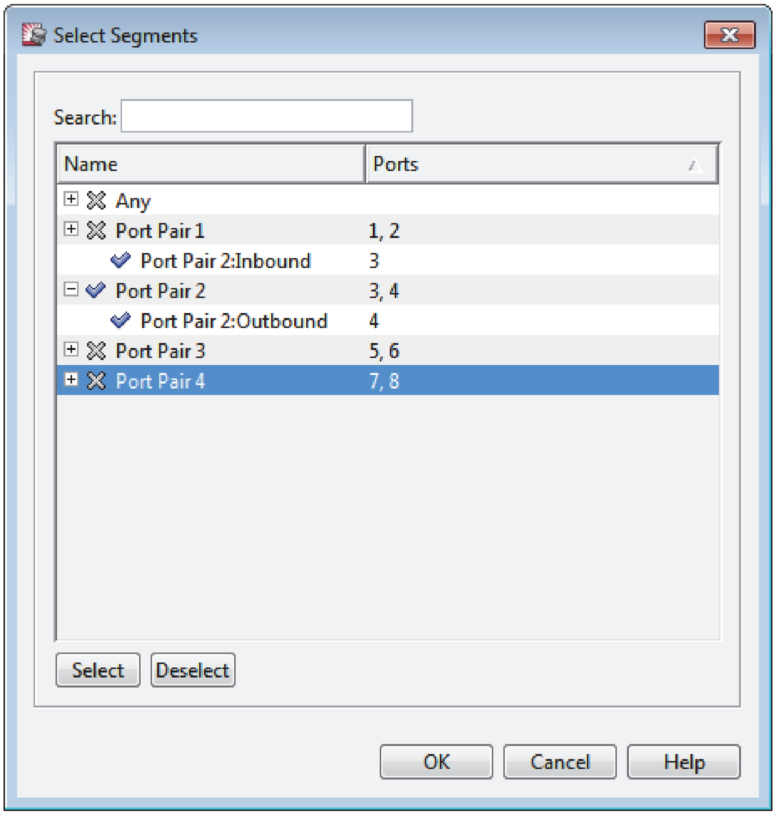

Name	Ports
Any	
Port Pair 1	1, 2
Port Pair 2:Inbound	3
Port Pair 2	3, 4
Port Pair 2:Outbound	4
Port Pair 3	5, 6
Port Pair 4	7, 8

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 15-4).

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

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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).

- 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 15-3.
- 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 15-3: 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 19, ‘‘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 19, ‘‘Managing Rules and Rule Sets

Elements of a Firewall + IPS Security Policy

There are two primary types of security policy:

- Firewall Security Policy (Figure 15-4)
  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 15-5)
  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 15-4: Elements of a Firewall Policy

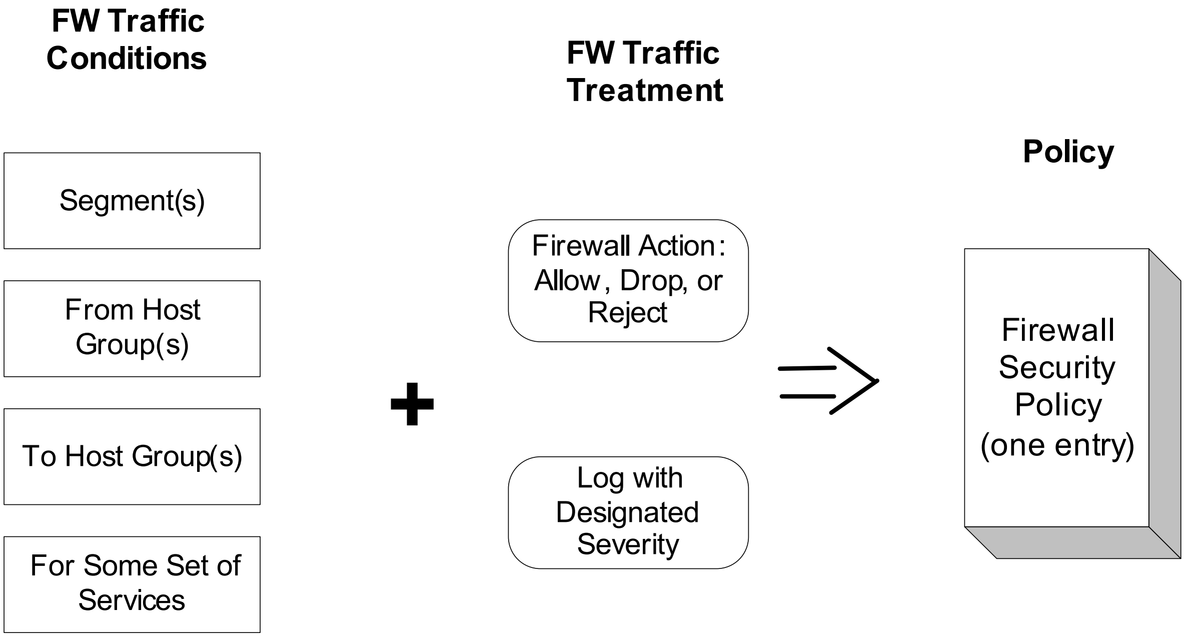

Figure 15-5 shows the elements used in a firewall policy.

Figure 15-5: Elements of an IPS Policy

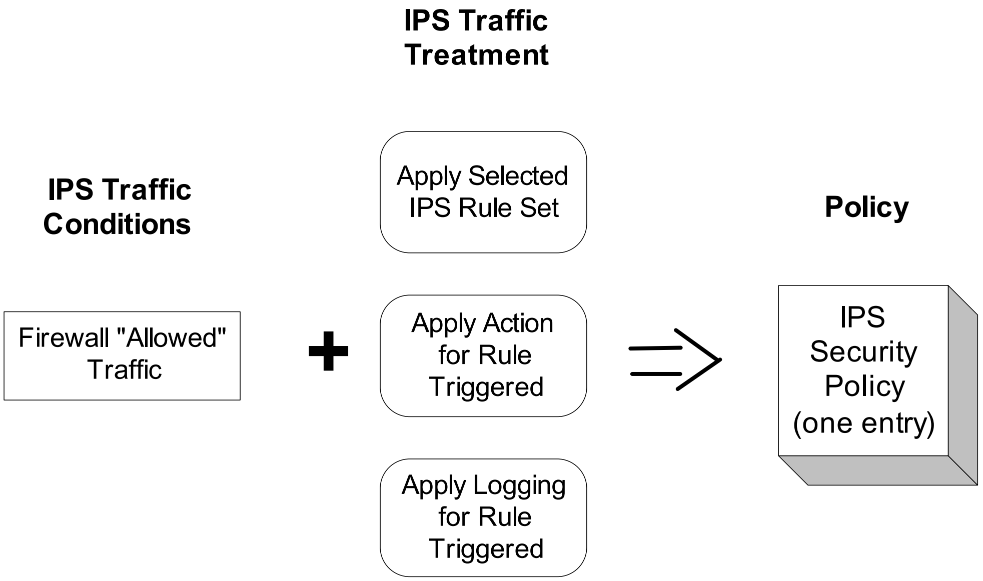
IPS Traffic Treatment
IPS Traffic Conditions
Apply Selected IPS Rule Set
Policy
Firewall "Allowed" Traffic
+
Apply Action for Rule Triggered
IPS Security Policy (one entry)
Apply Logging for Rule Triggered

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 15-6.

Figure 15-6: Elements in a Combined FW+IPS Policy

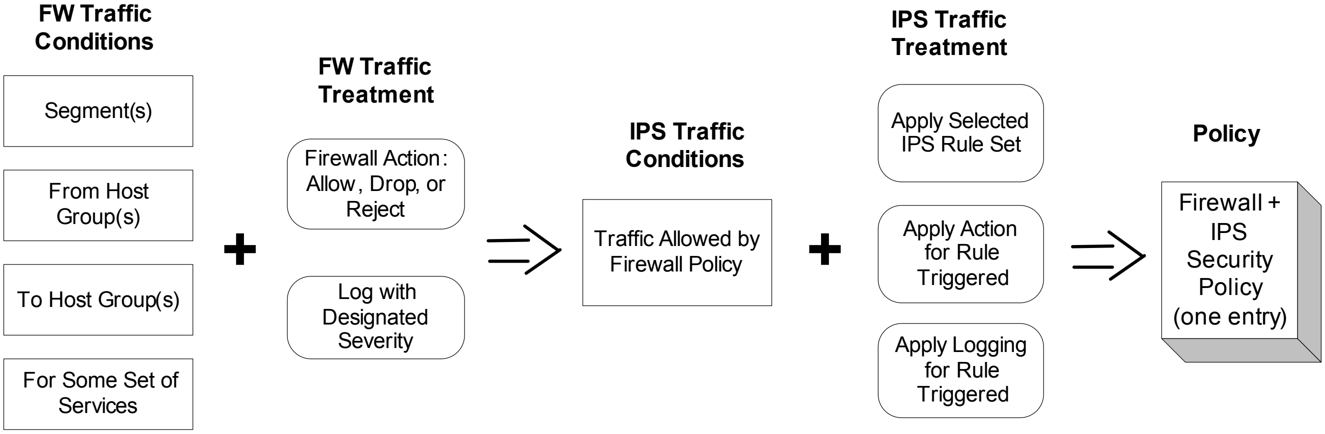
FW Traffic Conditions
Segment(s)
From Host Group(s)
To Host Group(s)
For Some Set of Services
+
FW Traffic Treatment
Firewall Action: Allow, Drop, or Reject
Log with Designated Severity
IPS Traffic Conditions
Traffic Allowed by Firewall Policy
+
IPS Traffic Treatment
Apply Selected IPS Rule Set
Apply Action for Rule Triggered
Apply Logging for Rule Triggered
Policy
Firewall + IPS Security Policy (one entry)

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 15-4.

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 15-7.

Figure 15-7: Default Strict Server Protection Policy

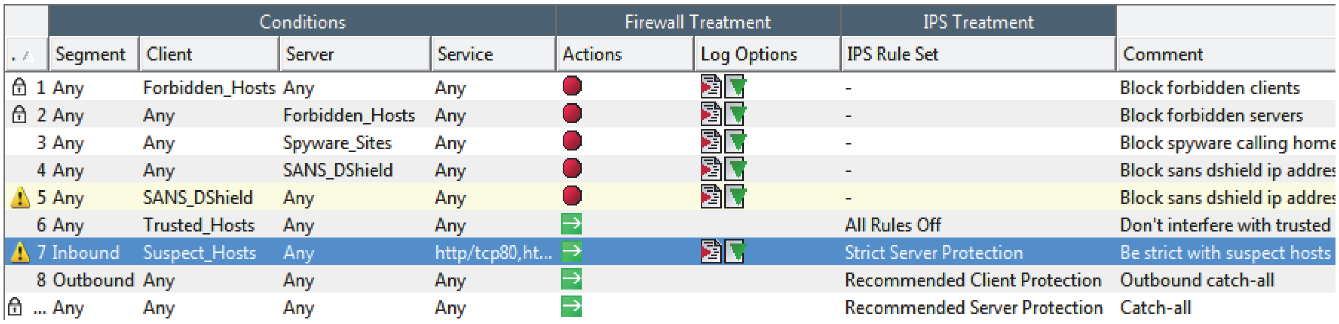

	Conditions				Firewall Treatment		IPS Treatment	
	Segment	Client	Server	Service	Actions	Log Options	IPS Rule Set	Comment
🔒 1	Any	Forbidden_Hosts	Any	Any	●	📄▼	-	Block forbidden clients
🔒 2	Any	Any	Forbidden_Hosts	Any	●	📄▼	-	Block forbidden servers
3	Any	Any	Spyware_Sites	Any	●	📄▼	-	Block spyware calling 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 15-4 describes the settings for the policy shown in Figure 15-7, and how they apply to traffic.

Table 15-4: 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.

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 15-5, and block or restrict much of the obviously questionable traffic. The table below describes each entry.

Table 15-5: 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.
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.

Table 15-5: Default FW+IPS Policies *(Continued)*

Row	Description	Summary
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 24, ‘‘SYN Flood and Connection Limiting Security”, Chapter 25, ‘‘Client Rate Limiting” and Chapter 22, ‘‘Advanced Client Rate Limiting”.

Chapter 16
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 15, “About Security Policies”.

This chapter includes the following topics:

NOTE

For information on managing rate-based policies, see Chapter 25, “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 16-1 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 16-1: FW+IPS Policy Example

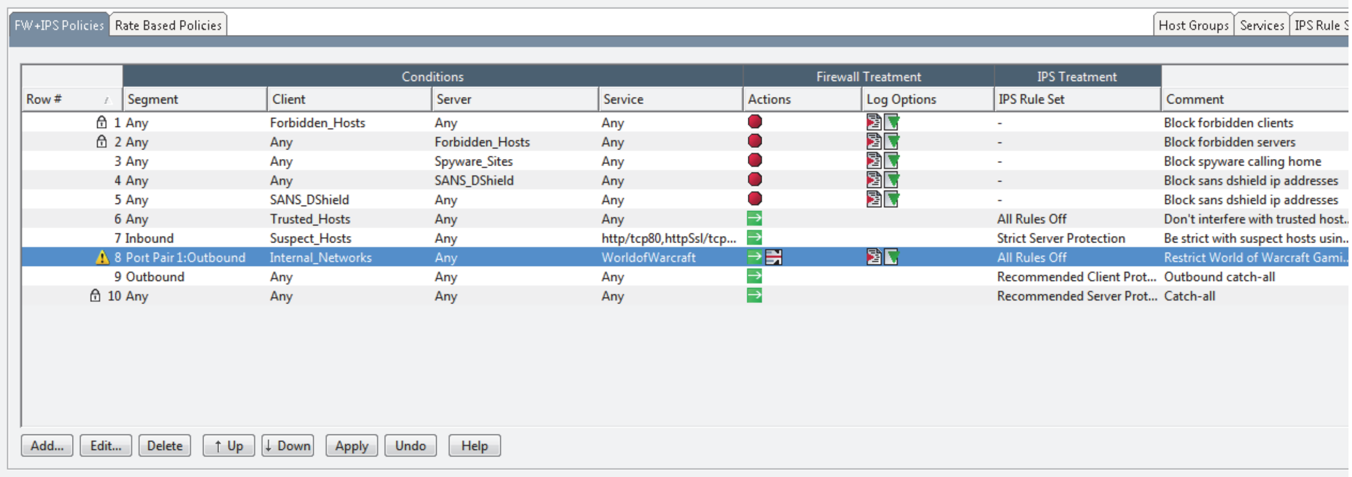

	Conditions				Firewall Treatment		IPS Treatment	
Row #	Segment	Client	Server	Service	Actions	Log Options	IPS Rule Set	Comment
1	Any	Forbidden_Hosts	Any	Any	●	▤ ▼	-	Block forbidden clients
2	Any	Any	Forbidden_Hosts	Any	●	▤ ▼	-	Block forbidden servers
3	Any	Any	Spyware_Sites	Any	●	▤ ▼	-	Block spyware calling home
4	Any	Any	SANS_DShield	Any	●	▤ ▼	-	Block sans dshield ip addresses
5	Any	SANS_DShield	Any	Any	●	▤ ▼	-	Block sans dshield ip addresses
6	Any	Trusted_Hosts	Any	Any	→		All Rules Off	Don't interfere with trusted host..
7	Inbound	Suspect_Hosts	Any	http/tcp80,httpSsl/tcp...	→		Strict Server Protection	Be strict with suspect hosts usin..
⚠ 8	Port Pair 1:Outbound	Internal_Networks	Any	WorldofWarcraft	→ ⇥	▤ ▼	All Rules Off	Restrict World of Warcraft Gami..
9	Outbound	Any	Any	Any	→		Recommended Client Prot...	Outbound catch-all
10	Any	Any	Any	Any	→		Recommended Server Prot...	Catch-all

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 IPS Controller management application:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the FW+IPS Policies tab (Figure 16-1).

The columns displayed for each policy are listed in Table 16-1.

Table 16-1: 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 16-1: 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 IPS Controller management application enables you to modify a policy before pushing (applying) it to one or more members of a Policy Group.

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 policy group members until you push the policy changes to one or more Corero Network Devices in the selected policy group.

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 16-2).
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. When you are finished making your changes, click Done.
6. Any time after you are finished with your modifications, you can click Push Changes on the Policy Group settings tab to push your changes to the Policy Group.

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 17, “Managing Host Groups”.
   - Services - for detailed instructions, see Chapter 18, “Managing Services”.
   - IPS Rule Sets - for detailed instructions, see Chapter 19, “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 16-2).
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 16-2 shows the Add FW+IPS Policy dialog box. The Edit FW+IPS Policy dialog box contains the same tabs and options.

Figure 16-2: Configuring a Policy: Conditions Tab

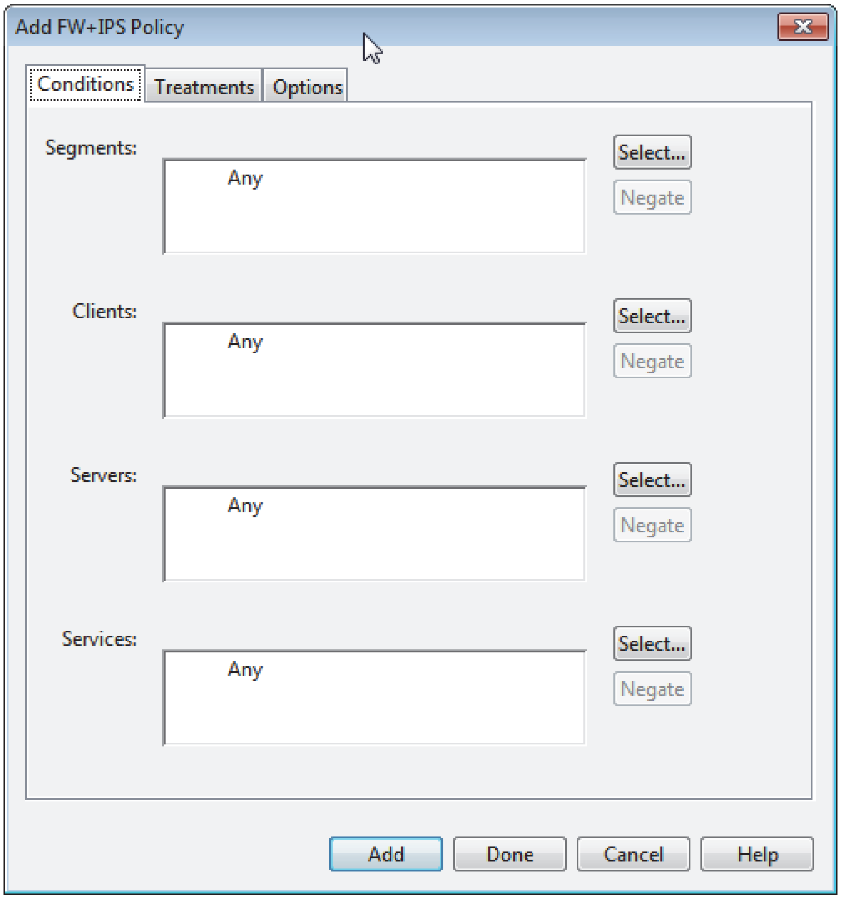

Add FW+IPS Policy

Segments: Any

Clients: Any

Servers: Any

Services: Any

4. On the Conditions tab (Figure 16-2), 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 16-2), 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 16-2), 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 16-2), 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 16-3). 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 16-3: Configuring a Policy: Treatments Tab

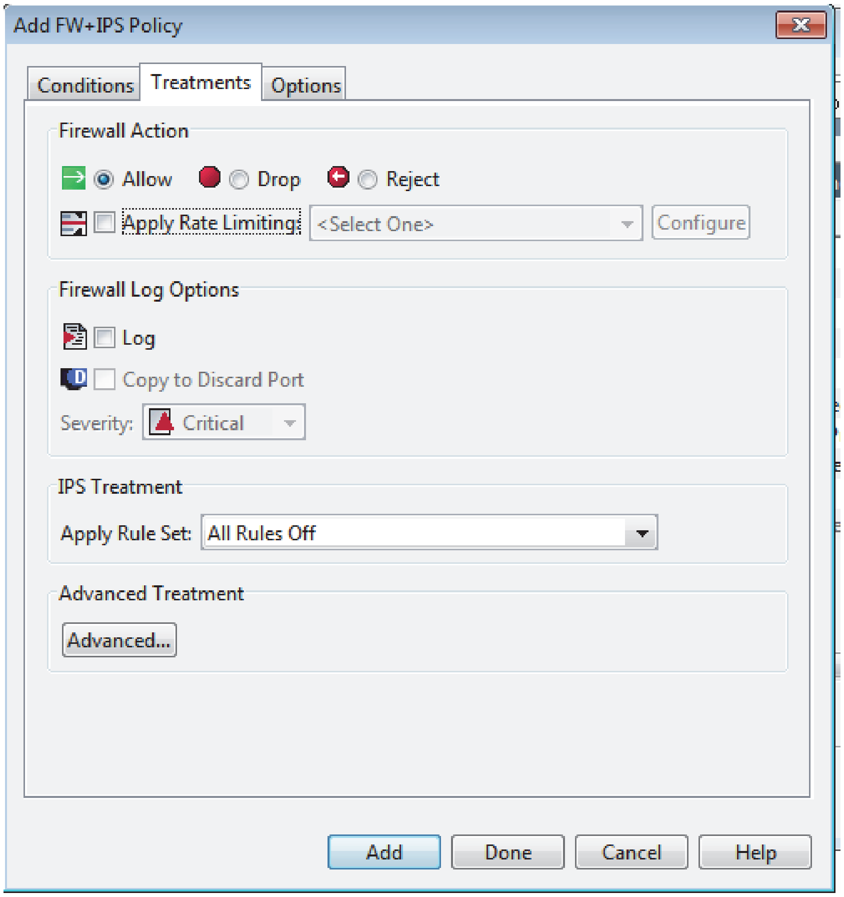

Firewall Action: ● Allow ○ Drop ○ Reject

□ Apply Rate Limiting: <Select One> Configure

□ Log

□ Copy to Discard Port

Severity: Critical

Apply Rule Set: All Rules Off

9. On the Treatments tab (Figure 16-3), 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 16-3), 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 16-3), 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 19, “Managing Rules and Rule Sets”.
12. On the Treatments tab (Figure 16-3), 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 16-2

Table 16-2: 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 IPS Controller 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. Any time after you finish modifying the policy settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

Chapter 17
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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).

- 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 17-1.
- 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 17-1: 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 17-2. 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 17-2: Default Host Groups

Host Group Name	Pre-Populated or Empty	Default FW Policy	Default IPS Policy	Other Policy Notes
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.
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.

Table 17-2: Default Host Groups *(Continued)*

Host Group Name	**Pre-Populated or Empty**	**Default FW Policy**	**Default IPS Policy**	**Other Policy Notes**
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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Host Groups tab for the selected Policy Group (Figure 17-1).

Figure 17-1: IPS Controller Host Groups Tab

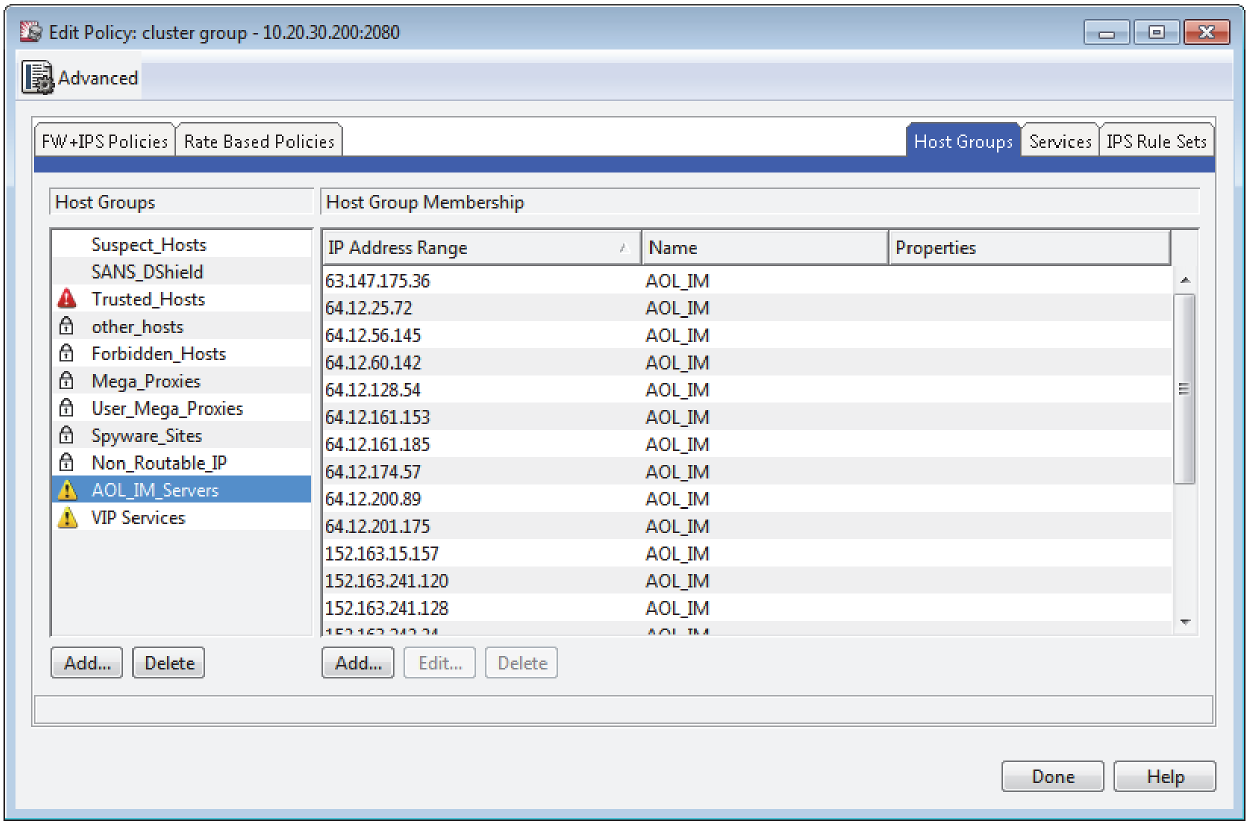

Host Groups

- Suspect_Hosts
- SANS_DShield
- ⚠ Trusted_Hosts
- 🔒 other_hosts
- 🔒 Forbidden_Hosts
- 🔒 Mega_Proxies
- 🔒 User_Mega_Proxies
- 🔒 Spyware_Sites
- 🔒 Non_Routable_IP
- ⚠ AOL_IM_Servers
- ⚠ VIP Services

Host Group Membership

IP Address Range	Name	Properties
63.147.175.36	AOL_IM	
64.12.25.72	AOL_IM	
64.12.56.145	AOL_IM	
64.12.60.142	AOL_IM	
64.12.128.54	AOL_IM	
64.12.161.153	AOL_IM	
64.12.161.185	AOL_IM	
64.12.174.57	AOL_IM	
64.12.200.89	AOL_IM	
64.12.201.175	AOL_IM	
152.163.15.157	AOL_IM	
152.163.241.120	AOL_IM	
152.163.241.128	AOL_IM	

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 17-10)

To add or modify a host group:

1. Go to the Host Groups tab as described in Viewing Host Groups (page 17-6).
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 IPS Controller management application, click Done.

7. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

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 17-6).
2. Select the host group you want to remove, then click Delete.
3. When you have finished deleting the host group in the IPS Controller management application, click Done.
4. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

Chapter 18
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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Services tab (Figure 18-1).

Figure 18-1: Services Tab

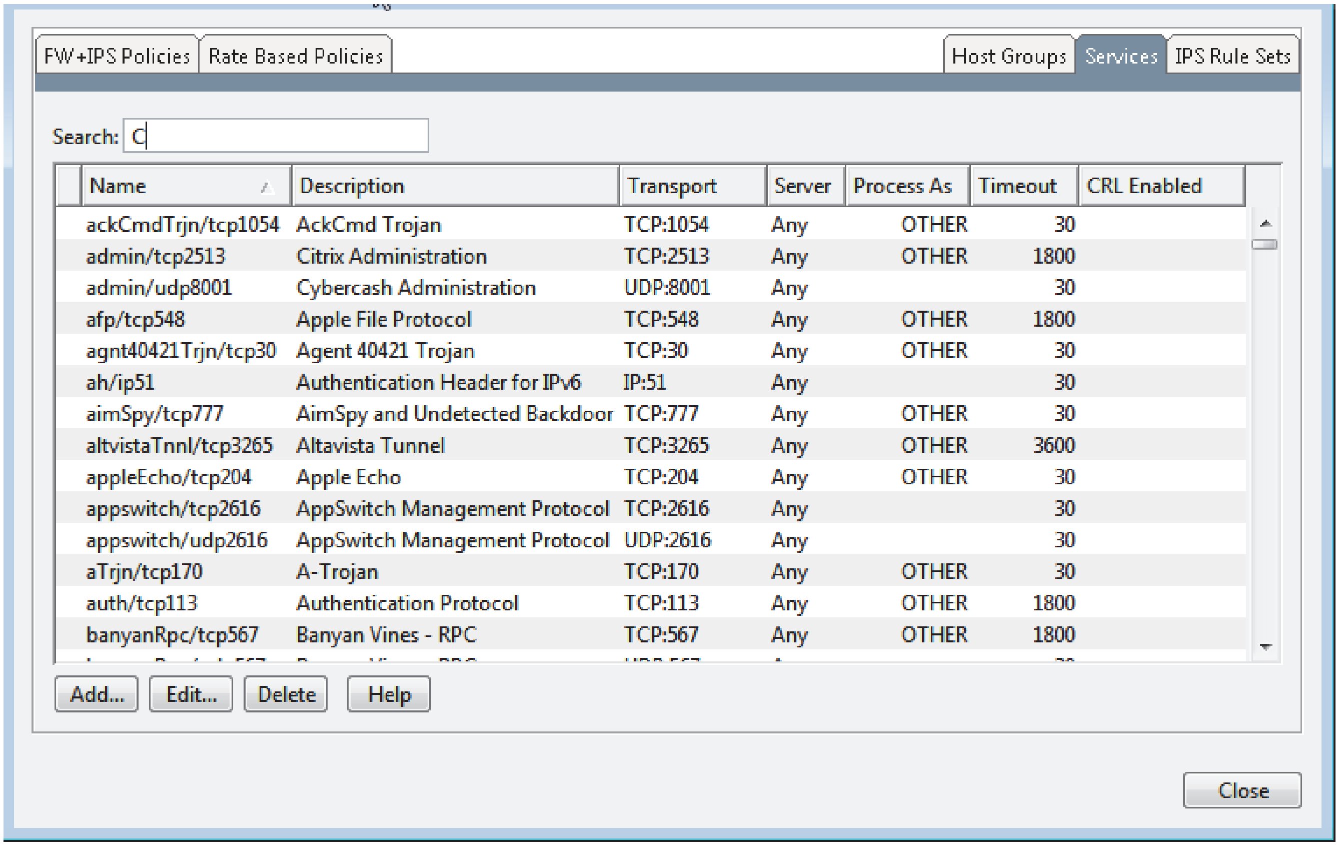

Search: C

Name	Description	Transport	Server	Process As	Timeout	CRL Enabled
ackCmdTrjn/tcp1054	AckCmd Trojan	TCP:1054	Any	OTHER	30	
admin/tcp2513	Citrix Administration	TCP:2513	Any	OTHER	1800	
admin/udp8001	Cybercash Administration	UDP:8001	Any		30	
afp/tcp548	Apple File Protocol	TCP:548	Any	OTHER	1800	
agnt40421Trjn/tcp30	Agent 40421 Trojan	TCP:30	Any	OTHER	30	
ah/ip51	Authentication Header for IPv6	IP:51	Any		30	
aimSpy/tcp777	AimSpy and Undetected Backdoor	TCP:777	Any	OTHER	30	
altvistaTnnl/tcp3265	Altavista Tunnel	TCP:3265	Any	OTHER	3600	
appleEcho/tcp204	Apple Echo	TCP:204	Any	OTHER	30	
appswitch/tcp2616	AppSwitch Management Protocol	TCP:2616	Any		30	
appswitch/udp2616	AppSwitch Management Protocol	UDP:2616	Any		30	
aTrjn/tcp170	A-Trojan	TCP:170	Any	OTHER	30	
auth/tcp113	Authentication Protocol	TCP:113	Any	OTHER	1800	
banyanRpc/tcp567	Banyan Vines - RPC	TCP:567	Any	OTHER	1800	

NOTE

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

4. 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.
5. You can view the settings available for each service. Settings are described in Table 18-1

Table 18-1: 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 18-2).
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 18-5).
8. When you have finished specifying service settings in the IPS Controller 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. Any time after you finish modifying the service settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

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 18-4).
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 18-2).

Figure 18-2: Advanced Service Settings Dialog Box

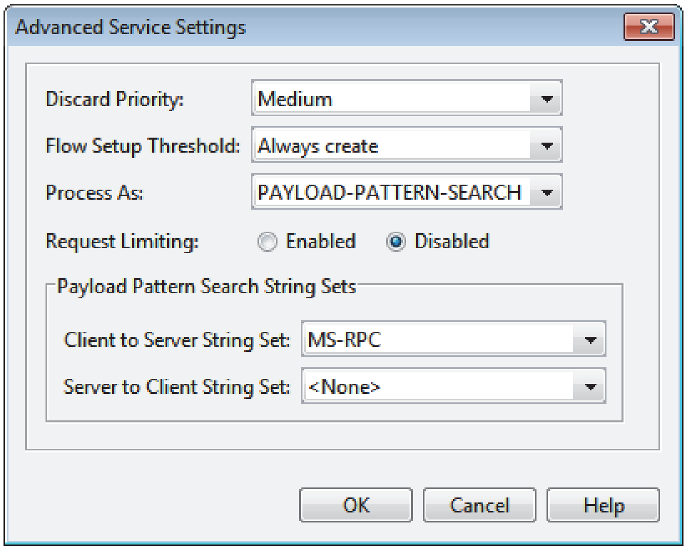

Advanced Service Settings

Discard Priority: Medium

Flow Setup Threshold: Always create

Process As: PAYLOAD-PATTERN-SEARCH

Request Limiting: ○ Enabled ● Disabled

Payload Pattern Search String Sets

Client to Server String Set: MS-RPC

Server to Client String Set: <None>

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

   - For a DDS Unit, you can select from the following protocol types:

DNS
ECHO
FTP
HTTP
LDAP
OTHER
PPS
SIP
SMTP
SSH
SSL
Telnet

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 19-20).

8. When finished, click OK.
9. When you have finished specifying service settings in the IPS Controller 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. Any time after you finish modifying the service settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

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 18-2).
2. Select the service you want to remove, then click Delete.
3. Any time after you delete the service, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

Chapter 19
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 19-1.

Table 19-1: 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 19-2.

Table 19-2: 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 19-15).

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 15-9).

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.

DDS Units provide Server Protection rule sets, but do not provide Client 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 19-3 describes the default rule sets.

Table 19-3: 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.

Table 19-3: Default Rule Sets *(Continued)*

Rule Set	Description
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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. 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 19-1 shows the IPS Rule Sets tab.

Figure 19-1: IPS Rule Sets Tab

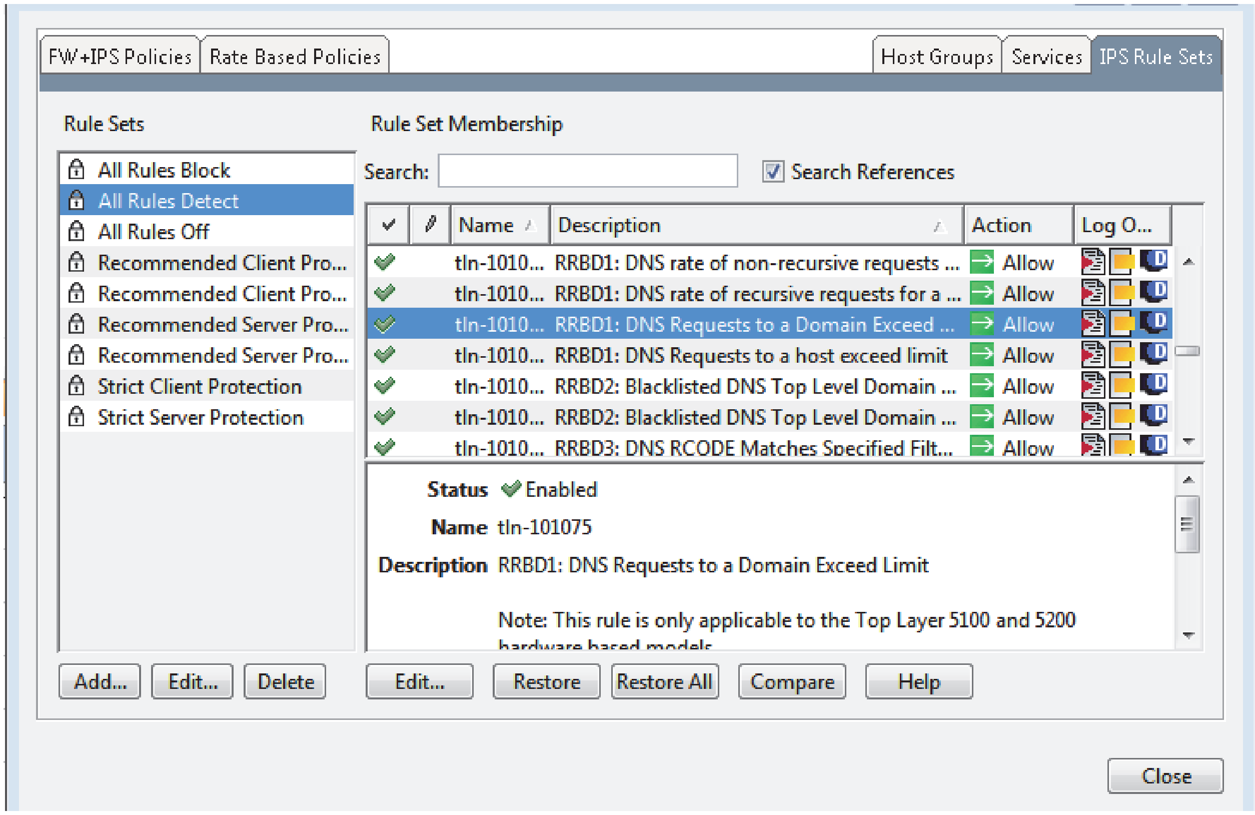

Rule Sets

- All Rules Block
- All Rules Detect
- All Rules Off
- Recommended Client Pro...
- Recommended Client Pro...
- Recommended Server Pro...
- Recommended Server Pro...
- Strict Client Protection
- Strict Server Protection

Rule Set Membership

Search: ☑ Search References

✓	✎	Name	Description	Action	Log O...
✓		tln-1010...	RRBD1: DNS rate of non-recursive requests ...	Allow	
✓		tln-1010...	RRBD1: DNS rate of recursive requests for a ...	Allow	
✓		tln-1010...	RRBD1: DNS Requests to a Domain Exceed ...	Allow	
✓		tln-1010...	RRBD1: DNS Requests to a host exceed limit	Allow	
✓		tln-1010...	RRBD2: Blacklisted DNS Top Level Domain ...	Allow	
✓		tln-1010...	RRBD2: Blacklisted DNS Top Level Domain ...	Allow	
✓		tln-1010...	RRBD3: DNS RCODE Matches Specified Filt...	Allow	

Status ✓ Enabled

Name tln-101075

Description RRBD1: DNS Requests to a Domain Exceed Limit

Note: This rule is only applicable to the Top Layer 5100 and 5200 hardware based models

4. 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 19-4

Table 19-4: 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 19-1 for a listing of the prefixes.

Table 19-4: 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.

5. 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.

6. 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 19-4). For more information on managing rules, see Rules Customization (page 19-24).

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 19-7).

   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 19-2).

Figure 19-2: Edit Rule Set Dialog Box

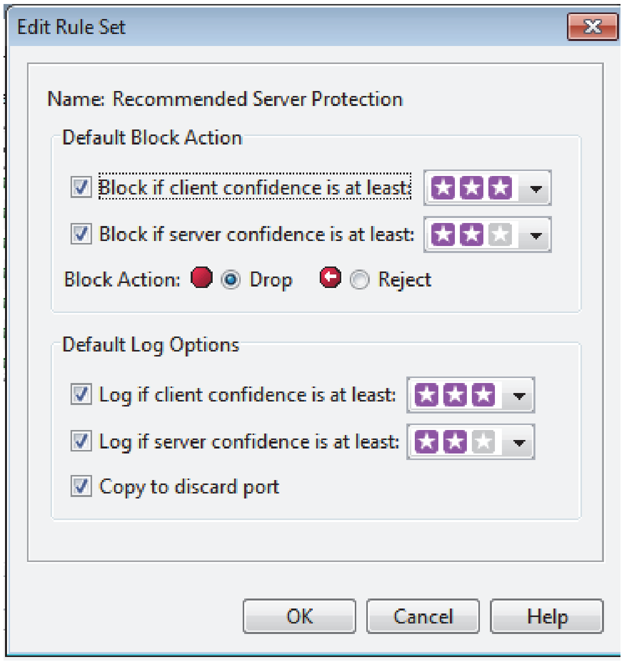

Name: Recommended Server Protection

Default Block Action

☑ Block if client confidence is at least: ★★★

☑ Block if server confidence is at least: ★★☆

Block Action: ◉ Drop ○ Reject

Default Log Options

☑ Log if client confidence is at least: ★★★

☑ Log if server confidence is at least: ★★☆

☑ Copy to discard port

b. Specify the desired settings for the rule set. These settings are described in Table 19-5.

Table 19-5: 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 IPS Controller management application, click Done.

6. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

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 24, "SYN Flood and Connection Limiting Security".

Packet-based rules are executed first, followed by Rate-Based Protection rules.

Figure 19-3 shows where these checks fit into the device's security subsystems.

Figure 19-3: Packet-Based Checks

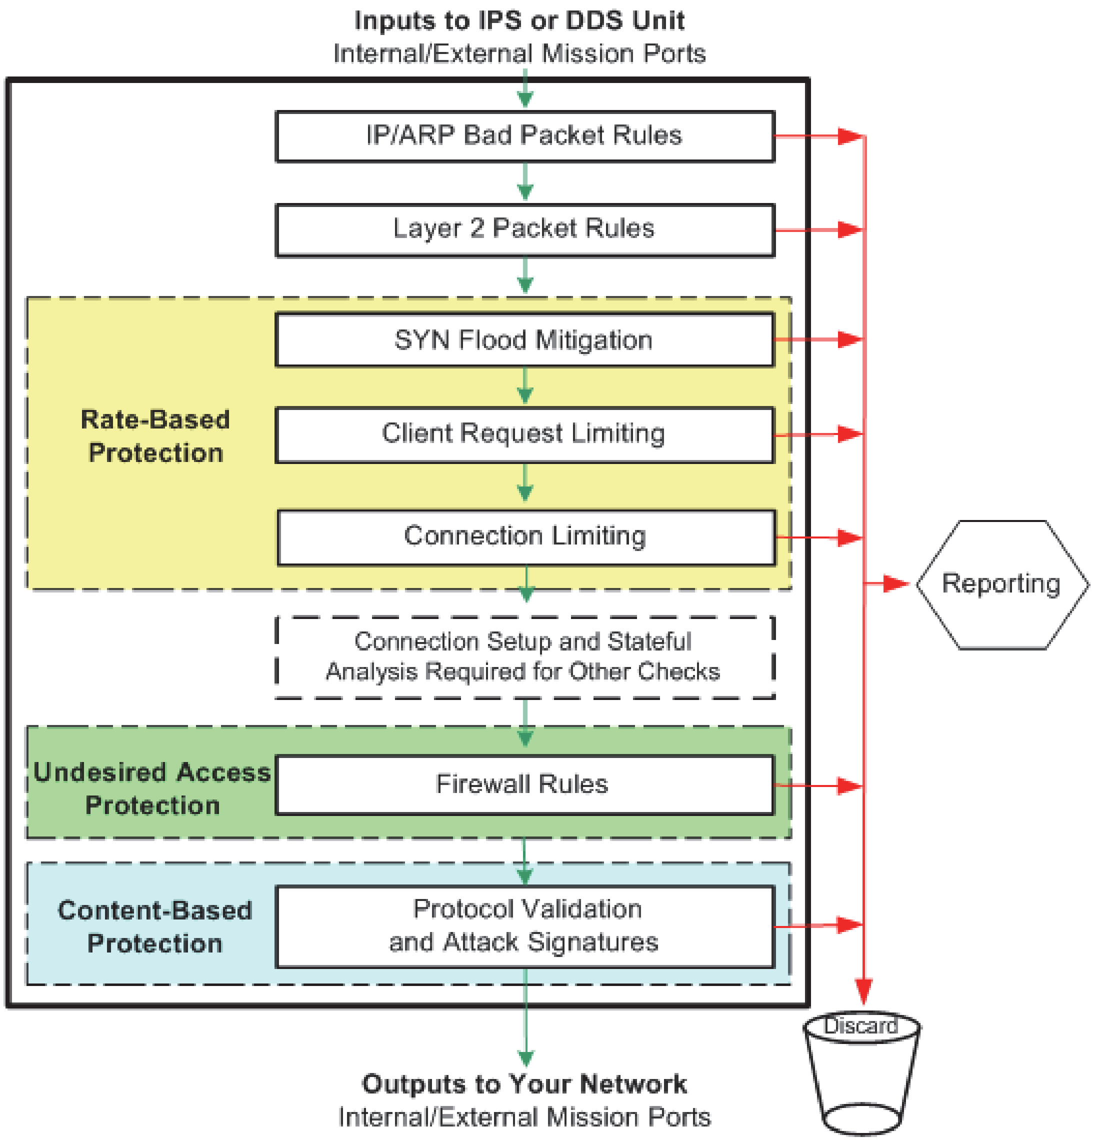
Inputs to IPS or DDS Unit
Internal/External Mission Ports
IP/ARP Bad Packet Rules
Layer 2 Packet Rules
Rate-Based Protection
SYN Flood Mitigation
Client Request Limiting
Connection Limiting
Reporting
Connection Setup and Stateful Analysis Required for Other Checks
Undesired Access Protection
Firewall Rules
Content-Based Protection
Protocol Validation and Attack Signatures
Discard
Outputs to Your Network
Internal/External Mission Ports

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:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. At the top of the Edit Policy dialog box, click the Advanced button, and choose either Packet Based Rules or Rate Based Rules from the menu.

   The appropriate rule dialog box displays. Figure 19-4 shows the Packet Based Rules dialog box, but the Rate Based Rules dialog box displays similar information.

Figure 19-4: Packet Based Rules Dialog Box

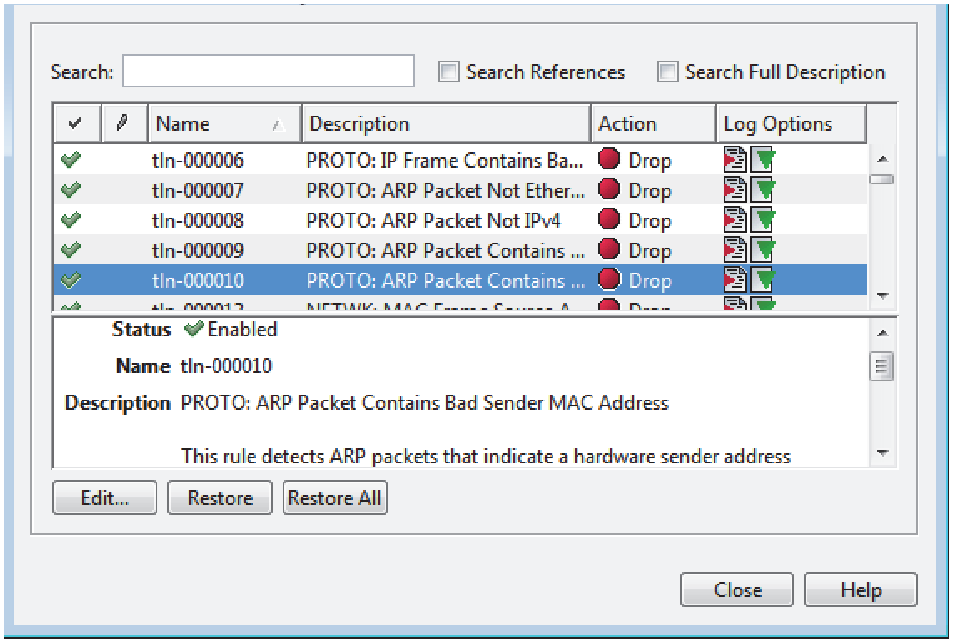

Search: □ Search References □ Search Full Description

✓		Name	Description	Action	Log Options
✓		tln-000006	PROTO: IP Frame Contains Ba...	● Drop	
✓		tln-000007	PROTO: ARP Packet Not Ether...	● Drop	
✓		tln-000008	PROTO: ARP Packet Not IPv4	● Drop	
✓		tln-000009	PROTO: ARP Packet Contains ...	● Drop	
✓		tln-000010	PROTO: ARP Packet Contains ...	● Drop	

Status ✓ Enabled

Name tln-000010

Description PROTO: ARP Packet Contains Bad Sender MAC Address

This rule detects ARP packets that indicate a hardware sender address

4. 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 IPS Controller management application, do one of the following:

In order to...	You must...
Modify packet-based rule settings	1. Click the Policy Group and Device Manager toolbar button. The Policy Group and Device Manager dialog box displays. 2. Click the Settings tab. 3. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays. 4. Click the Advanced button and choose Packet Based Rules from the drop-down list. The Packet Based Rules dialog box displays.
Modify rate-based rule settings	1. Click the Policy Group and Device Manager toolbar button. The Policy Group and Device Manager dialog box displays. 2. Click the Settings tab. 3. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays. 4. Click the Advanced button and choose Rate Based Rules from the drop-down list. The Rate Based Rules dialog box displays.
Modify IPS rule settings	1. Click the Policy Group and Device Manager toolbar button. The Policy Group and Device Manager dialog box displays. 2. Click the Settings tab. 3. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays. 4. Click the IPS Rule Sets tab. 5. 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 19-5).

Figure 19-5: Edit Rule Settings Dialog Box

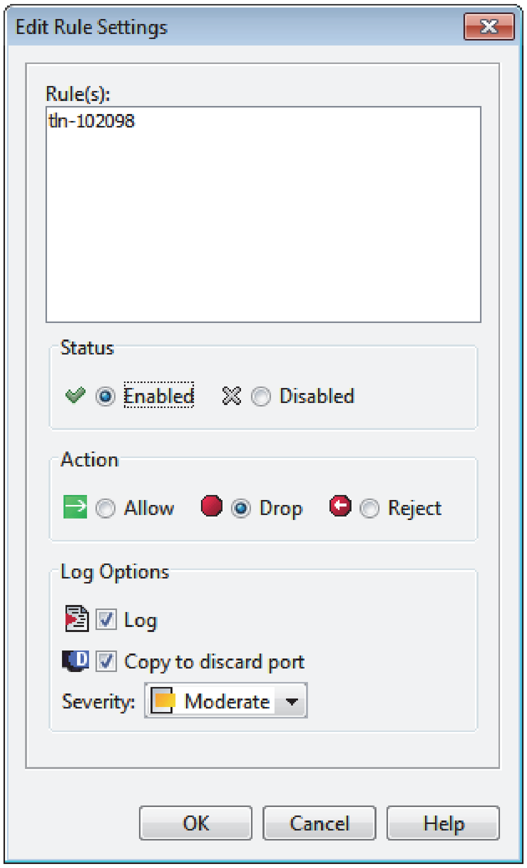

Edit Rule Settings

Rule(s):
tln-102098

Status

● Enabled ○ Disabled

Action

○ Allow ● Drop ○ Reject

Log Options

☑ Log

☑ Copy to discard port

Severity: Moderate

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 19-4). 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 IPS Controller management application, click OK.
9. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

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 23-16).

Comparing Two Rule Sets

If you want to view the differences between settings on two rule sets:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the IPS Rule Sets tab.
4. Select a rule set in the Rule Sets list.
5. Click Compare. The Compare IPS Rule Sets dialog box displays.
6. 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:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the IPS Rule Sets tab.
4. Select the desired rule set in the Rule Sets list.

   NOTE

   You can only restore rule settings for one rule list at a time.

5. 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.
6. You are prompted to confirm your selection.
7. When you have finished specifying rule settings in the IPS Controller management application, click Done.
8. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

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 DDS Unit provides 18 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 19-23).

To manage attack signature payload patterns:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Advanced button, and choose IPS Rules Customization from the drop-down list.

   The IPS Rules Customization dialog box displays (Figure 19-6).
4. View available patterns by choosing Attack Signatures > Patterns in the Filters area.
5. To modify a pattern, select the Signature name, then click Edit.
6. In the Signature field, specify the signature to be used as a filter.
7. 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 19-23).
8. Specify an ID label, which is a text string that describes the signature pattern.
9. 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.
10. Click Apply to save your pattern changes.
11. When you have finished specifying patterns in the IPS Controller 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.
12. Any time after you finish modifying the pattern settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

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. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Advanced button, and choose IPS Rules Customization from the drop-down list.

   The IPS Rules Customization dialog box displays (Figure 19-6).
4. View available string sets by choosing Attack Signatures > Sets in the Filters area.
5. To modify a signature set, select the Signature Set name, then click Edit.
6. If desired, enter a new name for the signature set.
7. Use the check box to indicate whether the signatures in this set are case sensitive.
8. Click OK.
9. When you have finished specifying rule settings in the IPS Controller management application, click Close.
10. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

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, 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. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a policy group, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Advanced button, and choose IPS Rules Customization from the drop-down list.

   The IPS Rules Customization dialog box displays (Figure 19-6).

Figure 19-6: IPS Rules Customization Dialog Box

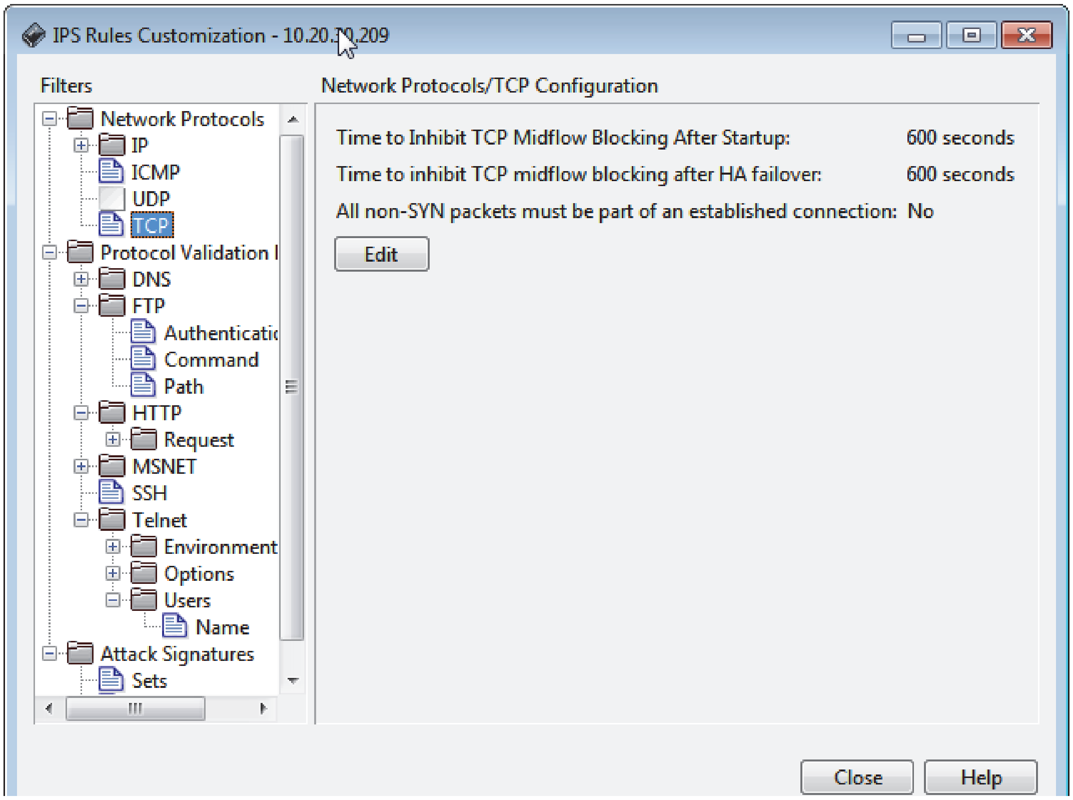

IPS Rules Customization - 10.20.10.209

Network Protocols/TCP Configuration

Time to Inhibit TCP Midflow Blocking After Startup: 600 seconds

Time to inhibit TCP midflow blocking after HA failover: 600 seconds

All non-SYN packets must be part of an established connection: No

4. Select the parameter in the left pane and view its settings in the right pane.
5. If desired, click Edit, and modify the parameters.
6. When finished, click OK.
7. When you have finished specifying rule settings in the IPS Controller management application, click Close.
8. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

Chapter 20
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 20-1 lists the available report templates:

Table 20-1: 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.
Gigashield Report	Includes all of the information contained in the Security Overview Report, and also includes CPU utilization, session statistics, IP address summary, LAN port utilization, and packet analysis details.
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 20-2) 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 20-2: 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 20-2: 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 20-1 shows the first page of a sample report.

Figure 20-1: Sample Periodic Security Report

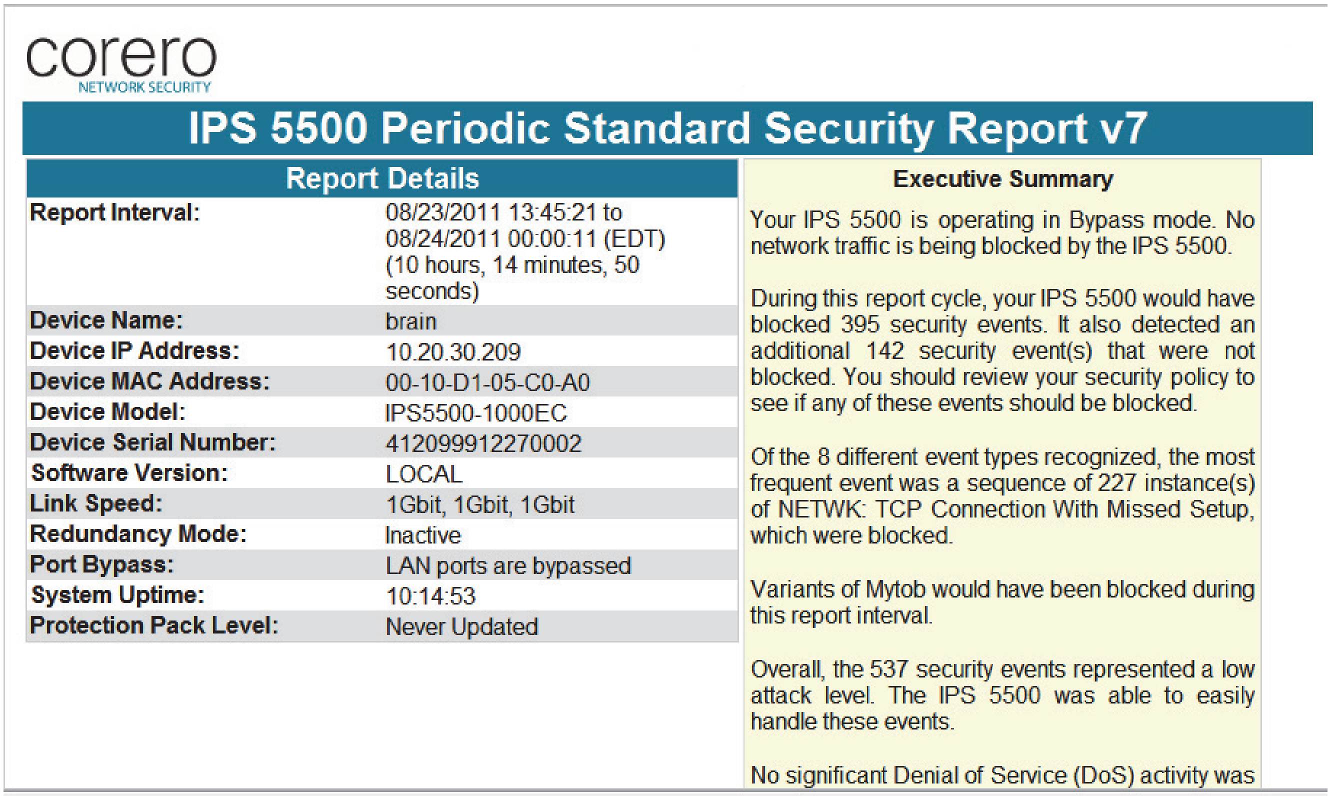

corero
NETWORK SECURITY

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 20-3 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.

NOTE

The Gigashield report, available only on the IPS Controller, displays all of the information available in the Standard Report.

Table 20-3: 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 20-3: 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 20-3: 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 20-3: 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. Do one of the following:
   - From the menu bar, choose Monitor > Reports > Devices.
   - Click the Policy Group and Device Manager tool bar button. Then click the Diagnostics & Reports tab on the Policy Group and Device Manager dialog box.
2. Select a device, then click View Immediate Report.

   The View Immediate Security Report dialog box displays.
3. Select the desired Report Template. Report templates are listed in Table 20-1.
4. 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 20-3).
- An IPS Controller can store an unlimited number of reports.
- 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 from the IPS Controller management application:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Diagnostics and Reports tab.
   - Choose Monitor > Reports > Devices from the menu bar.
2. Select a device, then click Report Settings. The Security Report Settings dialog box displays.
3. In the Periodic Report Generation Interval area, 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.
4. The Periodic Report Settings area lists the following information:
   - The name of the report and a description.
   - The report filename and timestamp.
   - Listing (statistical) information and whether or not the report will be periodically generated.
   - Whether or not the report template was factory defined. Report templates that are not factory defined are, by definition, user configured.
5. If you want to modify the list of data that will be included in the periodic report, or enable a report template for periodic generation, click Settings. 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.

7. Once you have specified your settings, click OK.

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. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Diagnostics and Reports tab.
   - Choose Monitor > Reports > Devices from the menu bar.

   The Diagnostics and Reports tab displays.

2. Click Reports. The Reports dialog box displays, listing all saved reports.

   On an IPS Controller, previously generated reports are listed in chronological order. Each row in the table displays the device, the report file name and template, whether the report was an Immediate report, the date and time it was generated, and the policy group and cluster information associated with the report.

3. Select the desired report.
4. 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:

- An IPS Controller can store an unlimited number of reports.
- 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. Do one of the following:
    - Click the Policy Group and Device Manager toolbar button, then click the Diagnostics and Reports tab.
    - Choose Monitor > Reports > Devices from the menu bar.

    The Diagnostics and Reports tab displays.
2. Click Reports. The Reports dialog box displays, listing all saved reports.
3. If desired, from the Display list, select the device or policy group whose reports you want to include in the list.
4. If desired, enter text in the search box to filter the list of reports.
5. Select the desired report. You can select multiple reports using Ctrl-Click and Shift-Click.
6. Click Delete. You are prompted to confirm your selection.

Chapter 21
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? (page 21-2)
- Event Logging System Outputs (page 21-3)
- Message Control Hierarchy (page 21-3)

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. Before you can view the events log for a Corero Network Device from the IPS Controller management application, you need to download diagnostic information from the device to the controller. This process is described in Downloading Diagnostic Information (page B-3).
2. Once you have generated and opened the zip file containing diagnostic information for a specific device, open the events.log text file.
3. 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 IPS Controller Management Alert Table

The Management Alert Table displays a summary of operational events, conditions, and issues related to Corero Network Devices, ProtectionClusters, and policy groups.

To view Management Alerts:

1. Do one of the following:
    - To view all alerts, from the menu bar, choose Management > Alerts.
    - To view only alerts associated with a specified device, group, or ProtectionCluster, right-click in the Policy Groups area and choose View Alerts from the pop-up menu.
      When you see an alert icon on a device, cluster, or group in the Policy Group tree, if you right-click and choose View Alerts, the Management Alert Table will only display those alerts associated with the item you selected.

   The Management Alert Table displays (Figure 21-1).

Figure 21-1: IPS Controller Management Alert Table

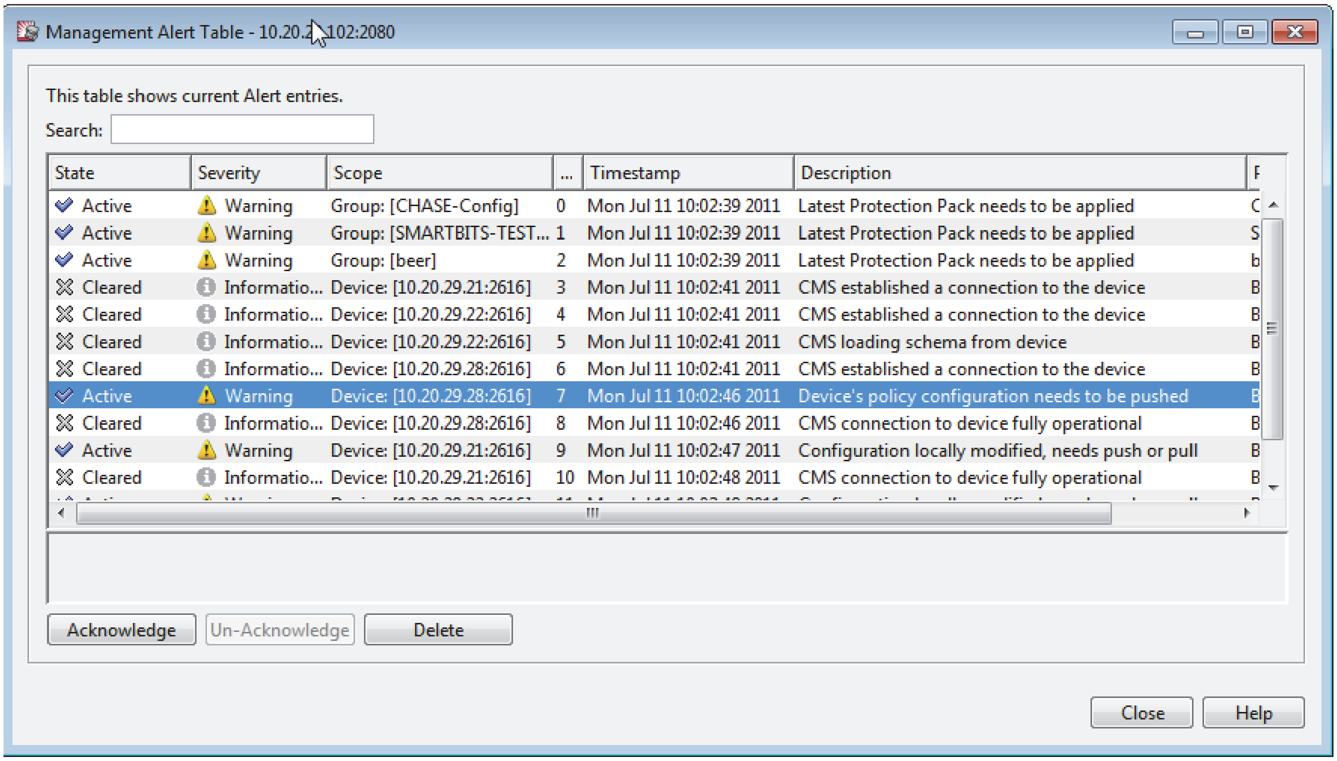

Management Alert Table - 10.20.2[illegible]102:2080

This table shows current Alert entries.

Search:

State	Severity	Scope	...	Timestamp	Description	F
✓ Active	⚠ Warning	Group: [CHASE-Config]	0	Mon Jul 11 10:02:39 2011	Latest Protection Pack needs to be applied	C
✓ Active	⚠ Warning	Group: [SMARTBITS-TEST...	1	Mon Jul 11 10:02:39 2011	Latest Protection Pack needs to be applied	S
✓ Active	⚠ Warning	Group: [beer]	2	Mon Jul 11 10:02:39 2011	Latest Protection Pack needs to be applied	b
✕ Cleared	ⓘ Informatio...	Device: [10.20.29.21:2616]	3	Mon Jul 11 10:02:41 2011	CMS established a connection to the device	B
✕ Cleared	ⓘ Informatio...	Device: [10.20.29.22:2616]	4	Mon Jul 11 10:02:41 2011	CMS established a connection to the device	B
✕ Cleared	ⓘ Informatio...	Device: [10.20.29.22:2616]	5	Mon Jul 11 10:02:41 2011	CMS loading schema from device	B
✕ Cleared	ⓘ Informatio...	Device: [10.20.29.28:2616]	6	Mon Jul 11 10:02:41 2011	CMS established a connection to the device	B
✓ Active	⚠ Warning	Device: [10.20.29.28:2616]	7	Mon Jul 11 10:02:46 2011	Device's policy configuration needs to be pushed	B
✕ Cleared	ⓘ Informatio...	Device: [10.20.29.28:2616]	8	Mon Jul 11 10:02:46 2011	CMS connection to device fully operational	B
✓ Active	⚠ Warning	Device: [10.20.29.21:2616]	9	Mon Jul 11 10:02:47 2011	Configuration locally modified, needs push or pull	B
✕ Cleared	ⓘ Informatio...	Device: [10.20.29.21:2616]	10	Mon Jul 11 10:02:48 2011	CMS connection to device fully operational	B

2. The table displays detailed information about each alert. This information is described in Table 21-1:

QUESTION: WHAT DOES CLEARED MEAN? CAN YOU MANUALLY GET AN ALERT INTO THE CLEARED STATE, OR DOES IT ONLY HAPPEN AUTOMATICALLY? WHAT TRIGGERS CLEARING?

Table 21-1: Management Alerts Table Information

Item	**Description**
State	Indicates the current state of the alert. Available states include: • Active The alert is an issue that has not yet been addressed or acknowledged; one that needs further action. Alerts in the Active state display alert icons for their associated devices, clusters, or policy groups in the Policy Group tree. When an alert first appears, it is in the Active state. • Acknowledged The alert is an issue you have seen and noted. When you manually acknowledge an alert, the alert icon no longer appears in the Policy Group tree. When you select an Active alert and click Acknowledge, the alert is placed in the Acknowledged state. When you select an acknowledged alert and click Un-Acknowledge, the alert is returned to the Active state. • Cleared The IPS Controller automatically handled this alert.
Severity	Indicates the severity of the alert. Severity levels include: • ☒ Critical - A serious problem has occurred. • ⚠ Warning - A typical maintenance action is needed. • ℹ Informational - The system has taken an action automatically. **QUESTION: WHAT COLOR IS THE INFORMATION ICON SUPPOSED TO BE? I CAN'T SEEM TO FIND IT ANYWHERE IT HAS NOT BEEEN CLEARED.** Active alerts with a severity of Critical or Warning also display on the Policy Group tree, adjacent to the device, cluster, or policy group to which the alert applies.
Scope	Indicates whether the alert is associated with a policy group, a ProtectionCluster, or a device.
ID	A system-assigned ID that informs you of the order in which the alert was generated with respect to other alerts in the table.
Timestamp	The date and time with the IPS Controller generated the alert.
Description	Text describing the cause of the alert.
Policy Group	The policy group associated with the alert.

3. To search the table, type the search text in the Search box and press Enter. The table is filtered to display only those alerts that meet the search criteria.
4. After performing a search, to view all alerts in the table, delete the text in the Search box and press Enter.
5. To change the state of an alert, do one of the following:
   - To move an alert from the Active state to the Acknowledge state, select the alert and click Acknowledge.
   - To move an alert from the Acknowledge state back to the Active state, select the alert and click Unacknowledge.

NOTE

If you want to remove the alert icon associated with an alert from the Policy Group view, Acknowledge the alert.

6. To remove an alert from the table, select the alert and click Delete. You are prompted to confirm your selection.

NOTE

The IPS Controller allows you to delete alerts in any state, including the Active state. Ensure you visually verify that you have selected the desired alert before you confirm the deletion.

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 Controller, see Managing Audit Logs (page 6-8).

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

Table 21-2: 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 System > View Audit Logs 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.

Chapter 22
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 for a Corero Network Device on the IPS Controller management application:

1. Do one of the following:
    - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
    - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Select a device, then click View Front Panel.
4. The Front Panel View displays.

Figure 22-1 shows the Front Panel View for an IPS Unit.

Figure 22-1: IPS Front Panel View

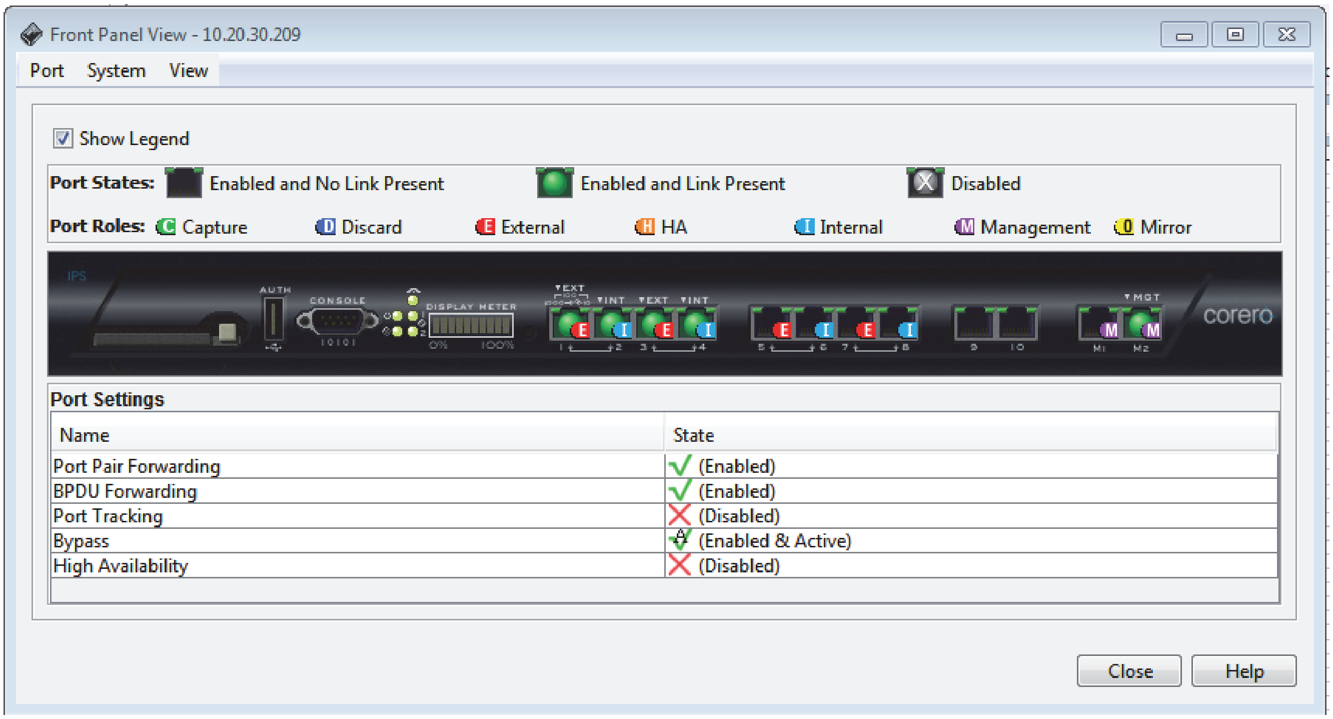

Front Panel View - 10.20.30.209

Port System View

☑ Show Legend

Port States: Enabled and No Link Present · Enabled and Link Present · ☒ Disabled

Port Roles: C Capture · D Discard · E External · H HA · I Internal · M Management · 0 Mirror

Port Settings

Name	State
Port Pair Forwarding	✓ (Enabled)
BPDU Forwarding	✓ (Enabled)
Port Tracking	✗ (Disabled)
Bypass	✓ (Enabled & Active)
High Availability	✗ (Disabled)

Figure 22-2 shows the Front Panel View for a DDS Unit.

Figure 22-2: DDS Front Panel View

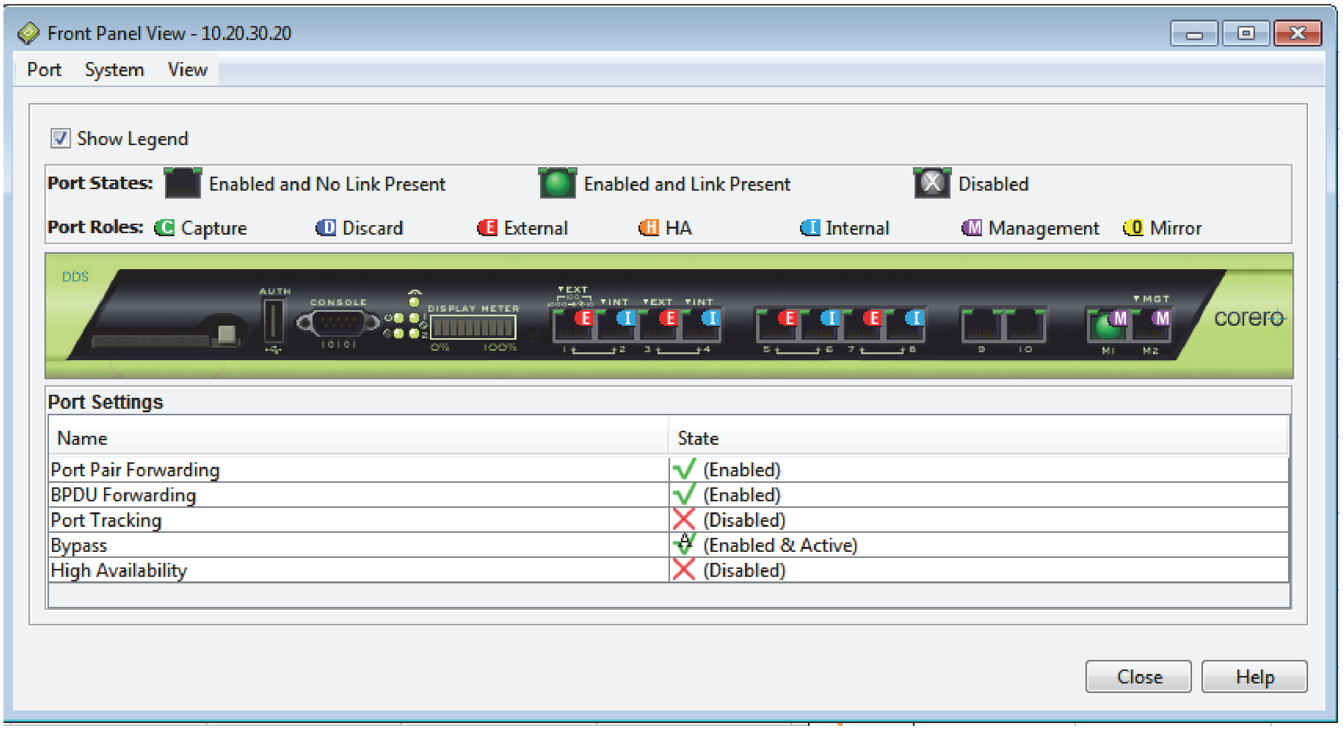

☑ Show Legend

Port States: ■ Enabled and No Link Present ● Enabled and Link Present ☒ Disabled

Port Roles: C Capture D Discard E External H HA I Internal M Management O Mirror

Port Settings

Name	State
Port Pair Forwarding	✓ (Enabled)
BPDU Forwarding	✓ (Enabled)
Port Tracking	✗ (Disabled)
Bypass	✓ (Enabled & Active)
High Availability	✗ (Disabled)

The Front Panel displays the information listed in Table 22-1.

Table 22-1: 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.

Table 22-1: Port Icons on the Front Panel View *(Continued)*

Port Icon	Description
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 11-10). 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 Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Select a device, then click View Front Panel.

   The Front Panel View displays.

   Figure 22-1 shows the IPS Front Panel View.

   Figure 22-2 shows the DDS Front Panel View.
4. 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.
5. 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 22-6).
6. 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 IPS Controller management application, choose Help > About IPS Controller from the menu bar.

   The "About..." dialog box displays the information fields listed in Table 22-2.

Table 22-2: "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.
Software Version	The currently-running version of the software.
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 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 using the IPS Controller management application, do one of the following:
   - Click the Policy Group and Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab in the Device area.
   - Choose Manage > Devices > Ports from the menu bar.
2. Select one or more devices, then click Port Statistics.

   The Port Statistics dialog box displays. For details on the information displayed in the Port Statistics dialog box, see Table 22-3.
3. To clear (zero) all statistics for a specific port, select the port and click Clear Statistics.

Table 22-3: 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.

Table 22-3: LAN Port Information *(Continued)*

Column	Description
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.
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 menu bar, choose Monitor > Statistics > 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 Corero Network Device Connectivity Status from the IPS Controller

At any time, you can view the state of connectivity between the IPS Controller and Corero Network Devices.

To view the connectivity status between a Corero Network Device and the IPS Controller:

1. In the IPS Controller management application, click the Policy Group & Device Manager toolbar button. The Policy Group and Device Manager dialog box displays.
2. Select either the Membership tab or the Settings tab.
3. The State column indicates the current connection state for each managed device. When there is connectivity, the status is Operational. Connection states are described in Table 22-4.

Table 22-4: Corero Network Device Connection Status Messages

Step	State	Description
1	Connecting	The IPS Controller attempts to connect to the IP address you specified for the Corero Network Device.
2	Authenticating	Once connected, the IPS Controller uses the shared management key to authenticate itself with the device.
3	Synchronizing	Once authenticated, the IPS Controller attempts to synchronize with the device by capturing a copy of all settings or parameters on the device.
4	Operational	Once synchronized, the device becomes operational. The IPS Controller is now able to manage changes to device settings and monitor the device’s security events.

NOTE

If you find the device persists in any state other than Operational, there is a connectivity issue. Refer to Troubleshooting Corero Network Device Connection Issues (page B-4) for troubleshooting information.

Viewing ProtectionCluster Status

To view the status of ProtectionCluster links between cluster members:

1. From the menu bar, choose Monitor > ProtectionCluster Status. The ProtectionCluster Status dialog box displays.

   This dialog box displays the information listed in Table 22-5

Table 22-5: ProtectionCluster Status Information

Column	Description
Policy Groups	Displays the Policy Group tree, fully expanded.
State	The current state of the Corero Network Device. Available states include: • Connecting - The cluster member is attempting to connect with other members of the cluster. • Operational - The cluster member is successfully connected to other cluster members.
HA Status	The status of the ProtectionCluster member in relation to the cluster as a whole. Available status messages include: • ? - The status of the member is unknown. • Active with peers - The member is active (operational), and other members of the cluster (peers) are available as well. • Active with no peers - The member is active (operational), but has no connectivity with other cluster members.
Size	The total number of members assigned to this ProtectionCluster. This displays as a fixed number regardless of whether all members are available or not.
Peers	The number of peers (other cluster members) with which this device is communicating. You can often identify problems by noting how many other cluster members a device can “see”.

2. For additional ProtectionCluster information, click Details. The ProtectionCluster Device Status dialog box displays. This dialog box lists each individual link between one cluster member and another. For each link, the dialog box displays the current link status. A green circle with an arrow pointing up indicates that this link is up, whereas a red circle with an arrow pointing down indicates that this link is down.

   NOTE

   For information on troubleshooting ProtectionClusters, see Troubleshooting ProtectionCluster Issues (page B-5)

Chapter 23
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 23-1 lists the primary views used to identify security issues.

Table 23-1: 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 23-16)
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 23-18) • Viewing Security Events and Security Event Details (page 23-21) • Security Event Viewer Filter Tool (page 23-23)

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 23-4).

Table 23-2 summarizes the tools you can use to research issues you discover using the Blocked and Detected Attacks page or the Security Event Viewer.

Table 23-2: 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 20, “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 21, “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.	• Using IP Address Query to Learn About a Host and Clear Counters (page 23-26)

Table 23-2: 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 23-31)
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 23-35)

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 23-3.

Table 23-3: 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 23-21) .
- Using the IP addresses from all the events matching the current filter in the Security Event Viewer. See Security Event Viewer Filter Tool (page 23-23).
- 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 23-4 lists some common reasons to use the shunning function.

Table 23-4: 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 23-1).
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 23-2).
   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 23-1).

Figure 23-1: Shun Attackers Dialog Box

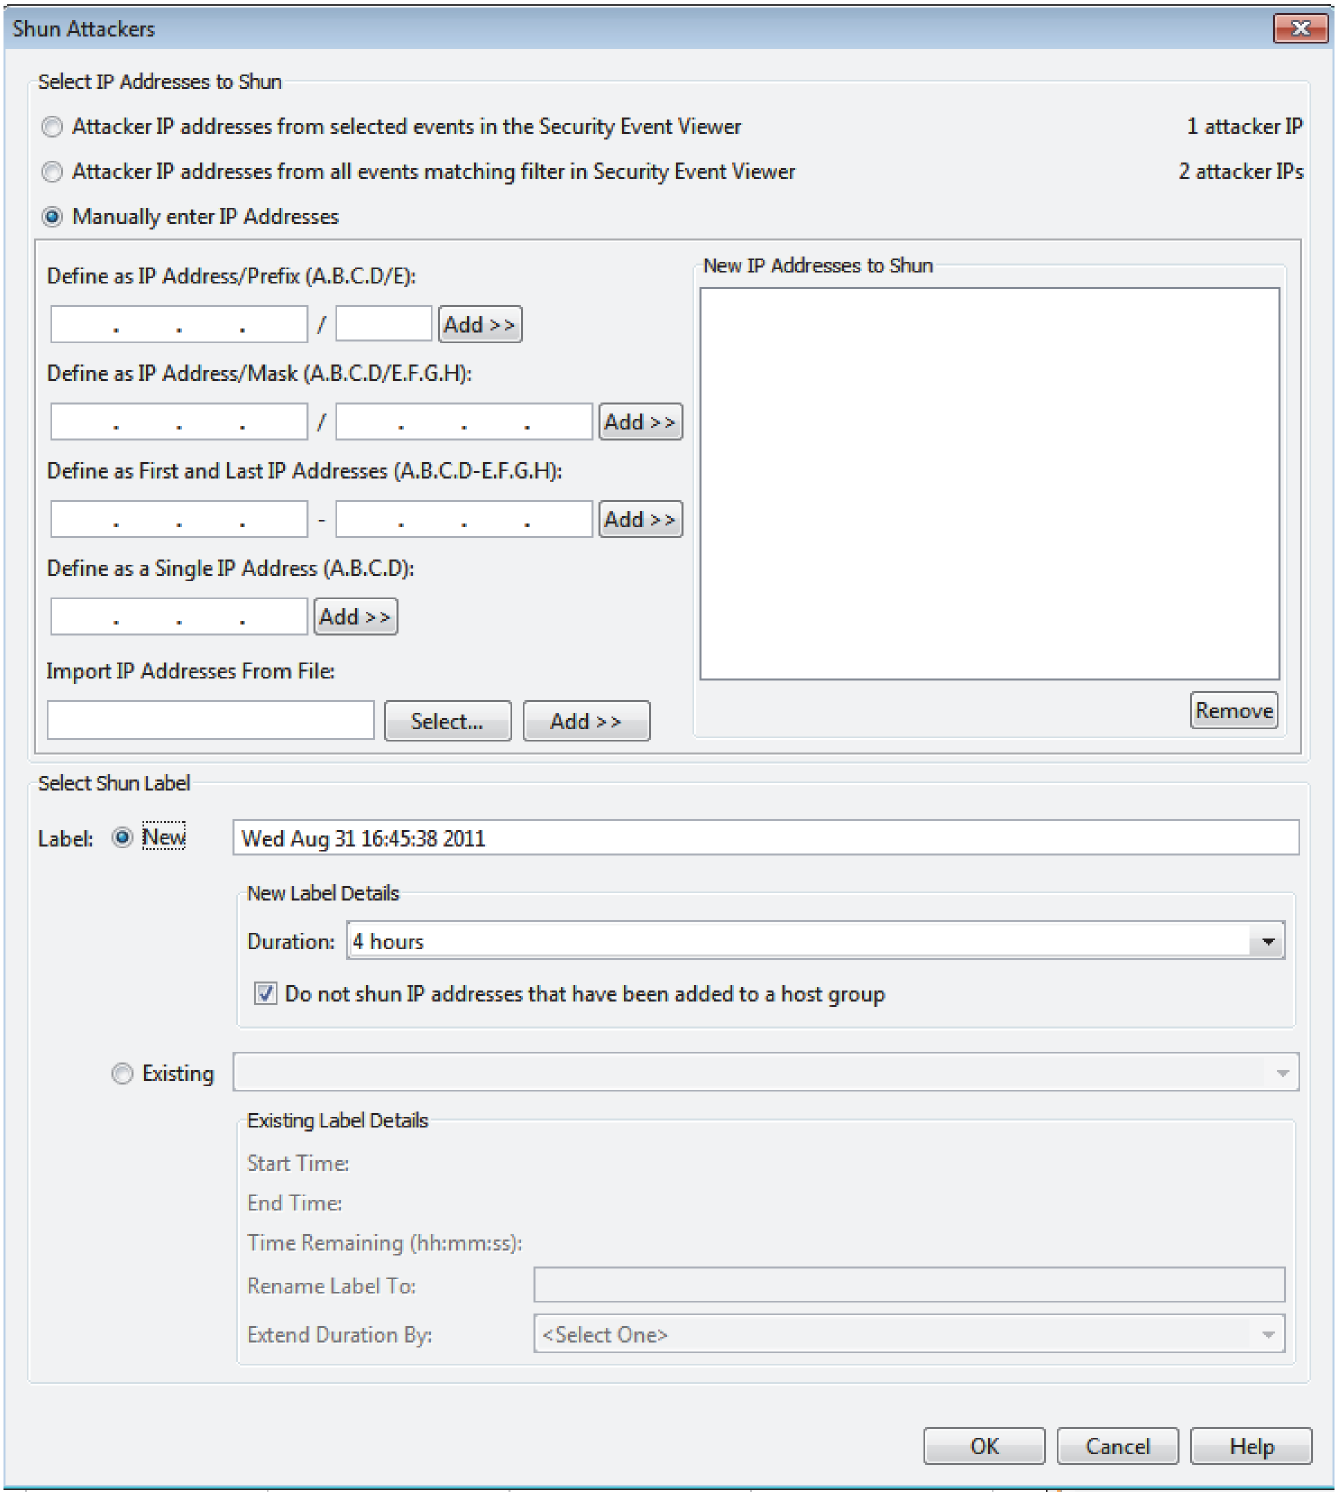

Shun Attackers

Select IP Addresses to Shun

○ Attacker IP addresses from selected events in the Security Event Viewer — 1 attacker IP

○ Attacker IP addresses from all events matching filter in Security Event Viewer — 2 attacker IPs

● Manually enter IP Addresses

Define as IP Address/Prefix (A.B.C.D/E):

Define as IP Address/Mask (A.B.C.D/E.F.G.H):

Define as First and Last IP Addresses (A.B.C.D-E.F.G.H):

Define as a Single IP Address (A.B.C.D):

Import IP Addresses From File:

New IP Addresses to Shun

Select Shun Label

Label: ● New — Wed Aug 31 16:45:38 2011

New Label Details

Duration: 4 hours

☑ Do not shun IP addresses that have been added to a host group

○ Existing

Existing Label Details

Start Time:

End Time:

Time Remaining (hh:mm:ss):

Rename Label To:

Extend Duration By: <Select One>

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.

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. On the IPS Controller, choose Manage > Policy Groups > Settings, and click the Shunning button.

   The Shunned Address Viewer displays (Figure 23-2).

Figure 23-2: Shunned Address Viewer

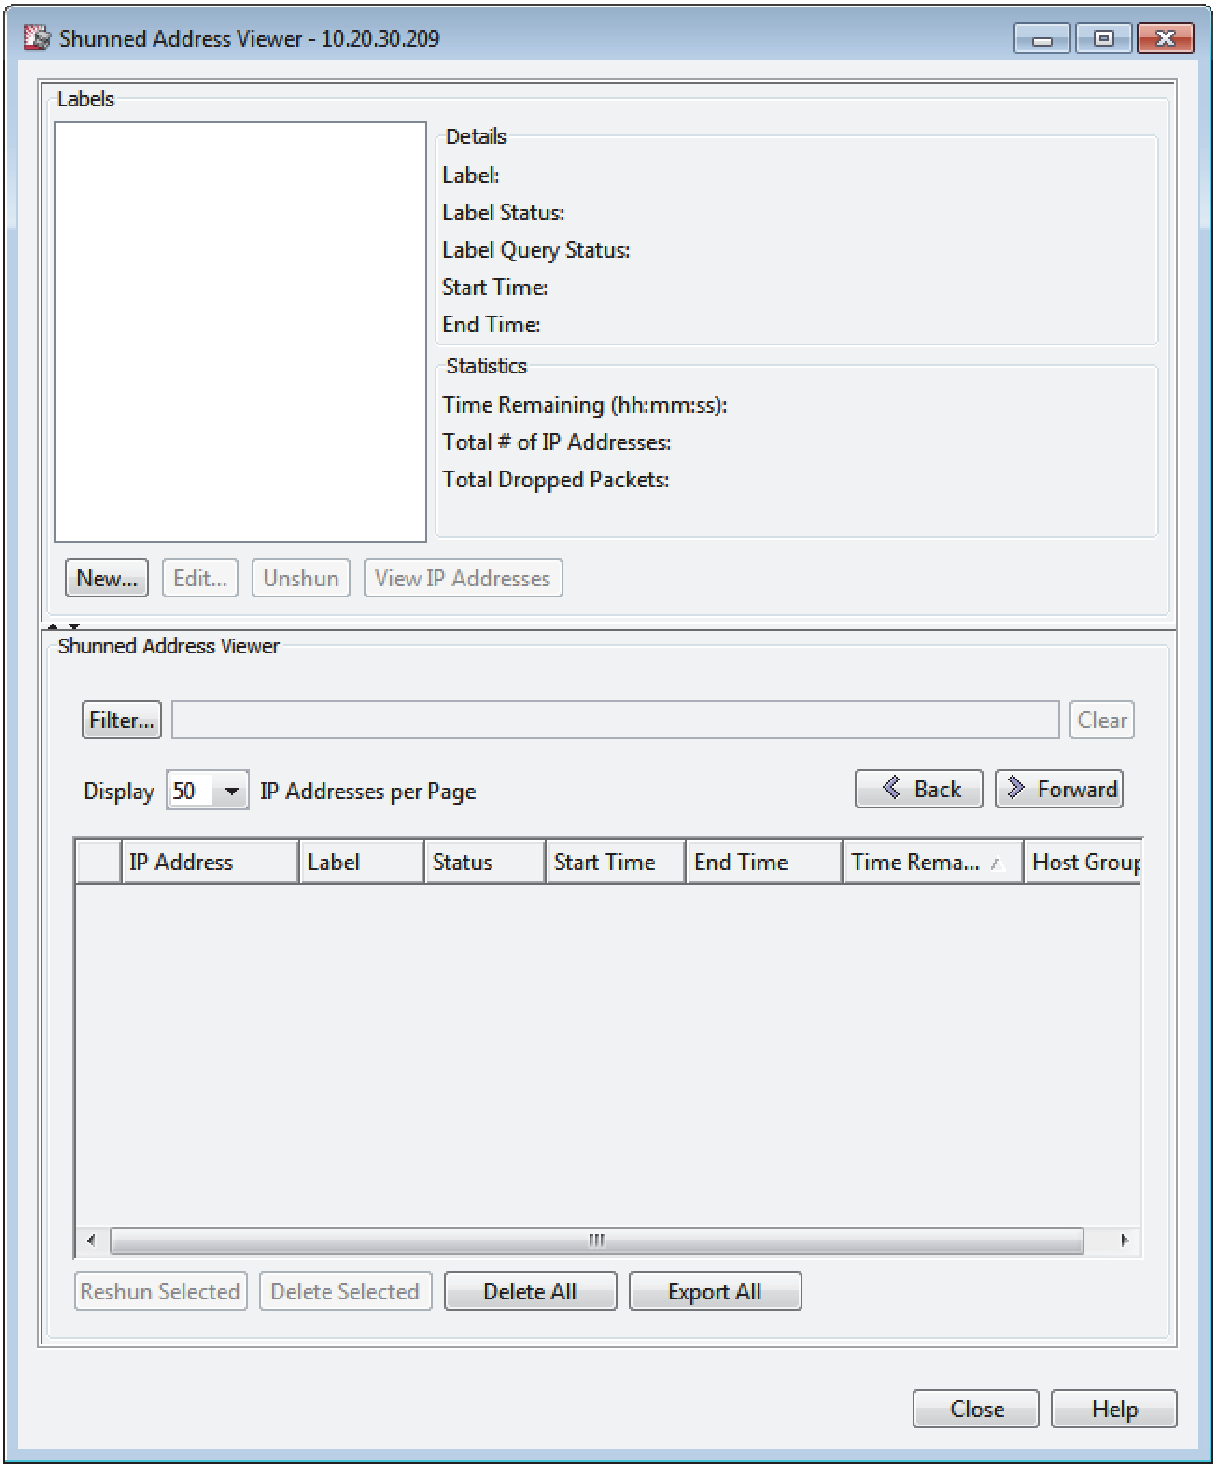

Shunned Address Viewer - 10.20.30.209

Labels

Details

Label:

Label Status:

Label Query Status:

Start Time:

End Time:

Statistics

Time Remaining (hh:mm:ss):

Total # of IP Addresses:

Total Dropped Packets:

Shunned Address Viewer

Display 50 IP Addresses per Page

	IP Address	Label	Status	Start Time	End Time	Time Rema...	Host Group

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 23-5 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 23-5: 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 23-6).
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 23-6).
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 23-6.

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 23-6: 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 23-14)

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 23-14).

   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 23-14).

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 23-14).

 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.

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. On the IPS Controller, choose Manage > Policy Groups > Settings, and click the Shunning button.

   The Shunned Address Viewer dialog box displays (Figure 23-2).

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 23-3).

Figure 23-3: Shunned Address Filter Dialog Box

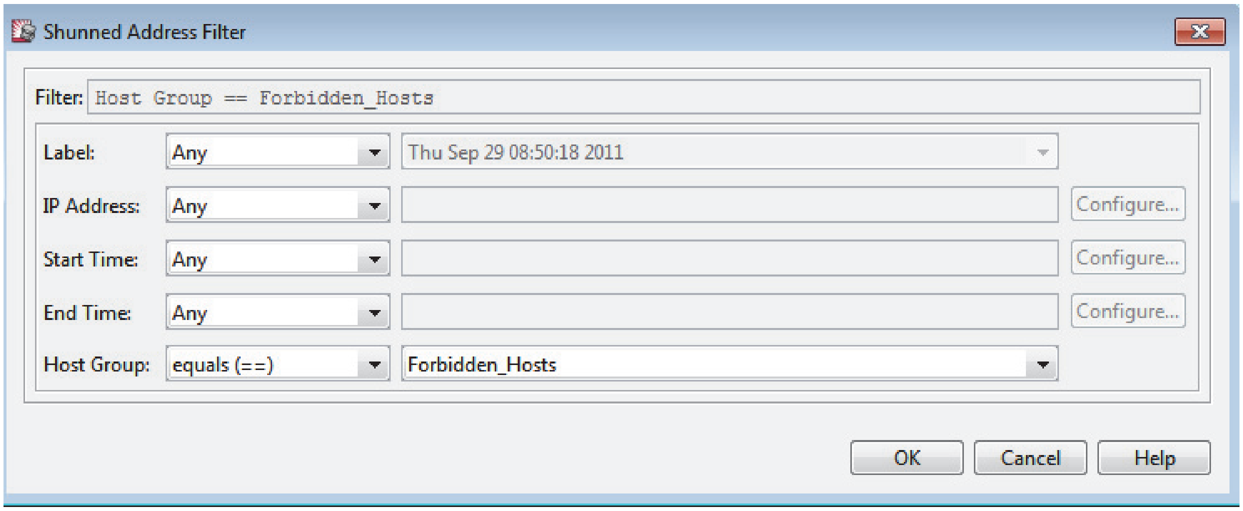

Shunned Address Filter

Filter: `Host Group == Forbidden_Hosts`

Label:	Any	Thu Sep 29 08:50:18 2011	
IP Address:	Any		Configure...
Start Time:	Any		Configure...
End Time:	Any		Configure...
Host Group:	equals (==)	Forbidden_Hosts	

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 23-7).

Table 23-7: 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 > Statistics > Blocked and Detected from the menu bar.
   - Click the Blocked & Detected toolbar button.

   The Blocked and Detected Attacks dialog box displays (Figure 23-4).

Figure 23-4: Blocked and Detected Attacks Dialog Box

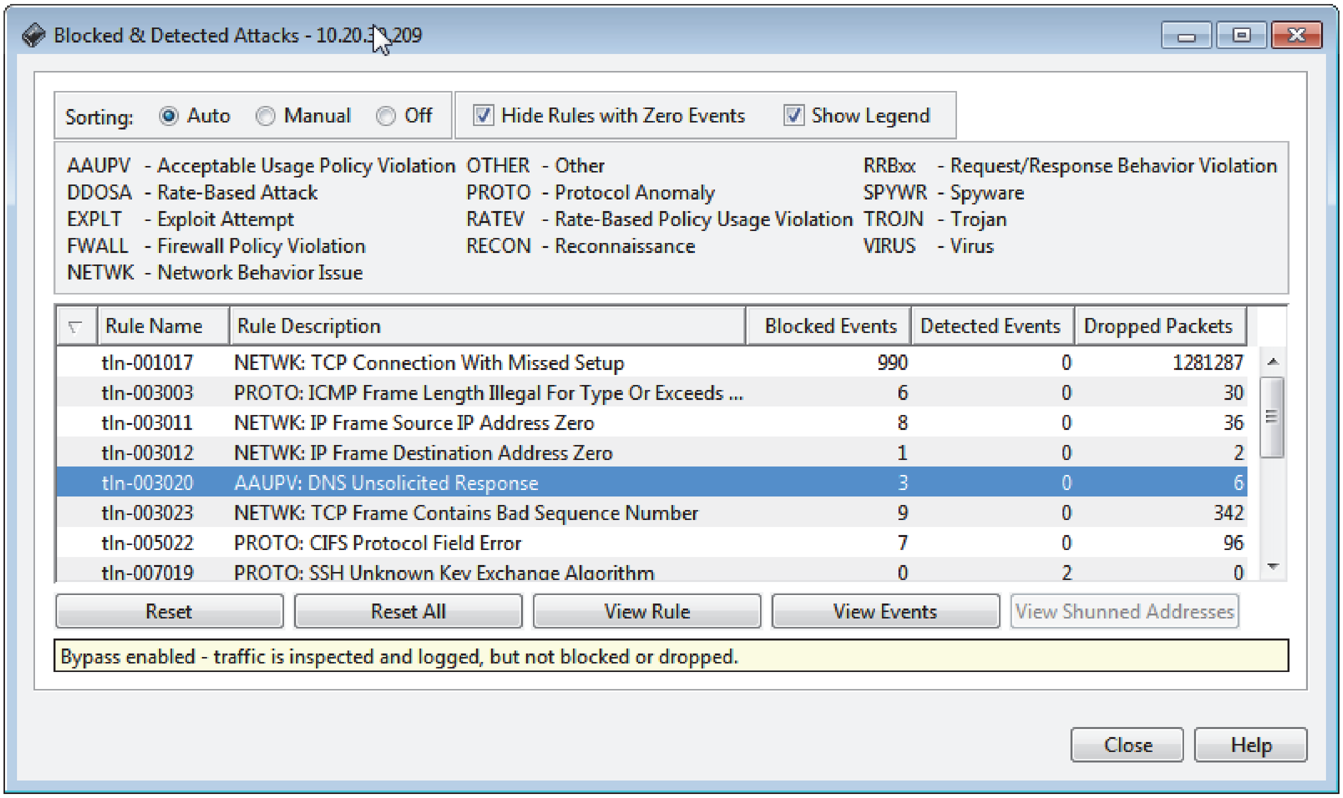

Blocked & Detected Attacks - 10.20.30.209

Sorting: ● Auto ○ Manual ○ Off ☑ Hide Rules with Zero Events ☑ Show Legend

AAUPV	- Acceptable Usage Policy Violation	OTHER	- Other	RRBxx	- Request/Response Behavior Violation
DDOSA	- Rate-Based Attack	PROTO	- Protocol Anomaly	SPYWR	- Spyware
EXPLT	- Exploit Attempt	RATEV	- Rate-Based Policy Usage Violation	TROJN	- Trojan
FWALL	- Firewall Policy Violation	RECON	- Reconnaissance	VIRUS	- Virus
NETWK	- Network Behavior Issue				

▿	Rule Name	Rule Description	Blocked Events	Detected Events	Dropped Packets
	tln-001017	NETWK: TCP Connection With Missed Setup	990	0	1281287
	tln-003003	PROTO: ICMP Frame Length Illegal For Type Or Exceeds ...	6	0	30
	tln-003011	NETWK: IP Frame Source IP Address Zero	8	0	36
	tln-003012	NETWK: IP Frame Destination Address Zero	1	0	2
	tln-003020	AAUPV: DNS Unsolicited Response	3	0	6
	tln-003023	NETWK: TCP Frame Contains Bad Sequence Number	9	0	342
	tln-005022	PROTO: CIFS Protocol Field Error	7	0	96
	tln-007019	PROTO: SSH Unknown Kev Exchanae Alaorithm	0	2	0

Bypass enabled - traffic is inspected and logged, but not blocked or dropped.

The Blocked and Detected Attacks page displays the information listed in Table 23-8.

Table 23-8: 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 23-8: 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.
   - 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 23-21).
6. To reset the Blocked, Detected, and Dropped Packet counters for an individual detection rule, select the attack and click the Reset button.
7. 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 23-5 shows the Security Event Viewer in the IPS Controller management application.

Figure 23-5: IPS Controller Security Event Viewer

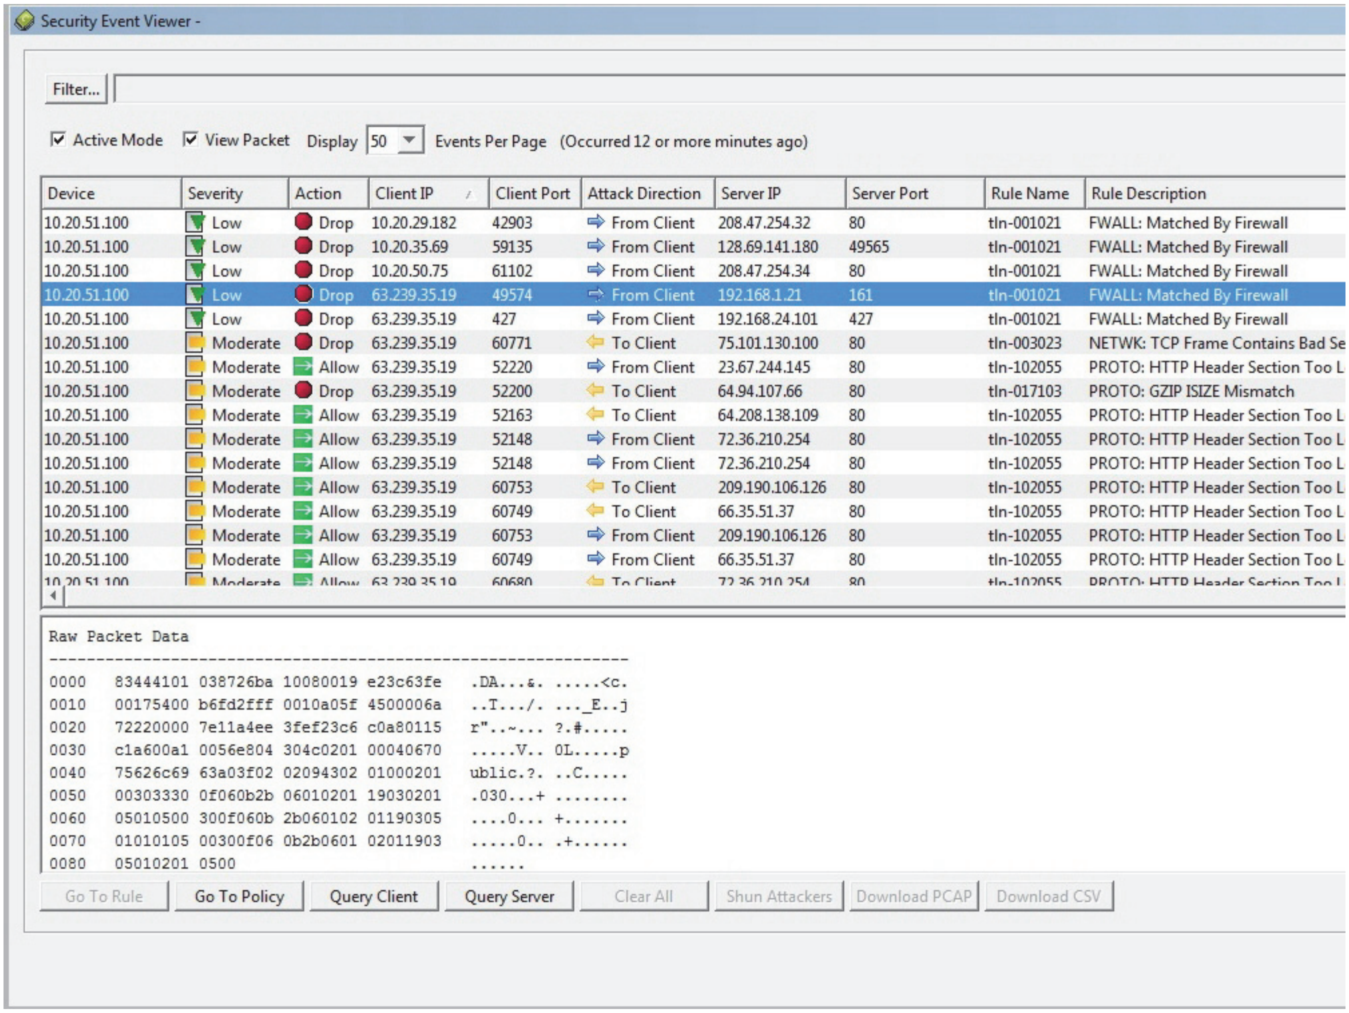

☑ Active Mode ☑ View Packet Display 50 Events Per Page (Occurred 12 or more minutes ago)

Device	Severity	Action	Client IP	Client Port	Attack Direction	Server IP	Server Port	Rule Name	Rule Description
10.20.51.100	Low	Drop	10.20.29.182	42903	From Client	208.47.254.32	80	tln-001021	FWALL: Matched By Firewall
10.20.51.100	Low	Drop	10.20.35.69	59135	From Client	128.69.141.180	49565	tln-001021	FWALL: Matched By Firewall
10.20.51.100	Low	Drop	10.20.50.75	61102	From Client	208.47.254.34	80	tln-001021	FWALL: Matched By Firewall
10.20.51.100	Low	Drop	63.239.35.19	49574	From Client	192.168.1.21	161	tln-001021	FWALL: Matched By Firewall
10.20.51.100	Low	Drop	63.239.35.19	427	From Client	192.168.24.101	427	tln-001021	FWALL: Matched By Firewall
10.20.51.100	Moderate	Drop	63.239.35.19	60771	To Client	75.101.130.100	80	tln-003023	NETWK: TCP Frame Contains Bad Se
10.20.51.100	Moderate	Allow	63.239.35.19	52220	From Client	23.67.244.145	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Drop	63.239.35.19	52200	To Client	64.94.107.66	80	tln-017103	PROTO: GZIP ISIZE Mismatch
10.20.51.100	Moderate	Allow	63.239.35.19	52163	To Client	64.208.138.109	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	52148	From Client	72.36.210.254	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	52148	From Client	72.36.210.254	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	60753	To Client	209.190.106.126	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	60749	To Client	66.35.51.37	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	60753	From Client	209.190.106.126	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	60749	From Client	66.35.51.37	80	tln-102055	PROTO: HTTP Header Section Too L
10.20.51.100	Moderate	Allow	63.239.35.19	60680	To Client	72.36.210.254	80	tln-102055	PROTO: HTTP Header Section Too L

```
Raw Packet Data
-------------------------------------------------------------
0000   83444101 038726ba 10080019 e23c63fe   .DA...&. .....<c.
0010   00175400 b6fd2fff 0010a05f 4500006a   ..T.../. ..._E..j
0020   72220000 7e11a4ee 3fef23c6 c0a80115   r"..~... ?.#.....
0030   c1a600a1 0056e804 304c0201 00040670   .....V.. 0L.....p
0040   75626c69 63a03f02 02094302 01000201   ublic.?. ..C.....
0050   00303330 0f060b2b 06010201 19030201   .030...+ ........
0060   05010500 300f060b 2b060102 01190305   ....0... +.......
0070   01010105 00300f06 0b2b0601 02011903   .....0.. .+......
0080   05010201 0500                         ......
```

The information shown in the Security Event Viewer is described in Table 23-9.

Table 23-9: Security Event Viewer Information

Column	Description
Device	The Corero Network Device associated with the security event.
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 19-15). NOTE: To view a description of the rule that triggered an event, select the event and click the View Rule button.

Table 23-9: 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 19-15).
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 19-1 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 menu bar, choose Monitor > Security Event Viewer.
   - From the Toolbar, click the Security Events button.

   The IPS Controller Security Event Viewer dialog box displays (Figure 23-5).

   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 23-23).

   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 23-10 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 23-10: 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 23-10: 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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).
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 23-4). **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 23-11, 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 23-6: Security Event Filter Tool

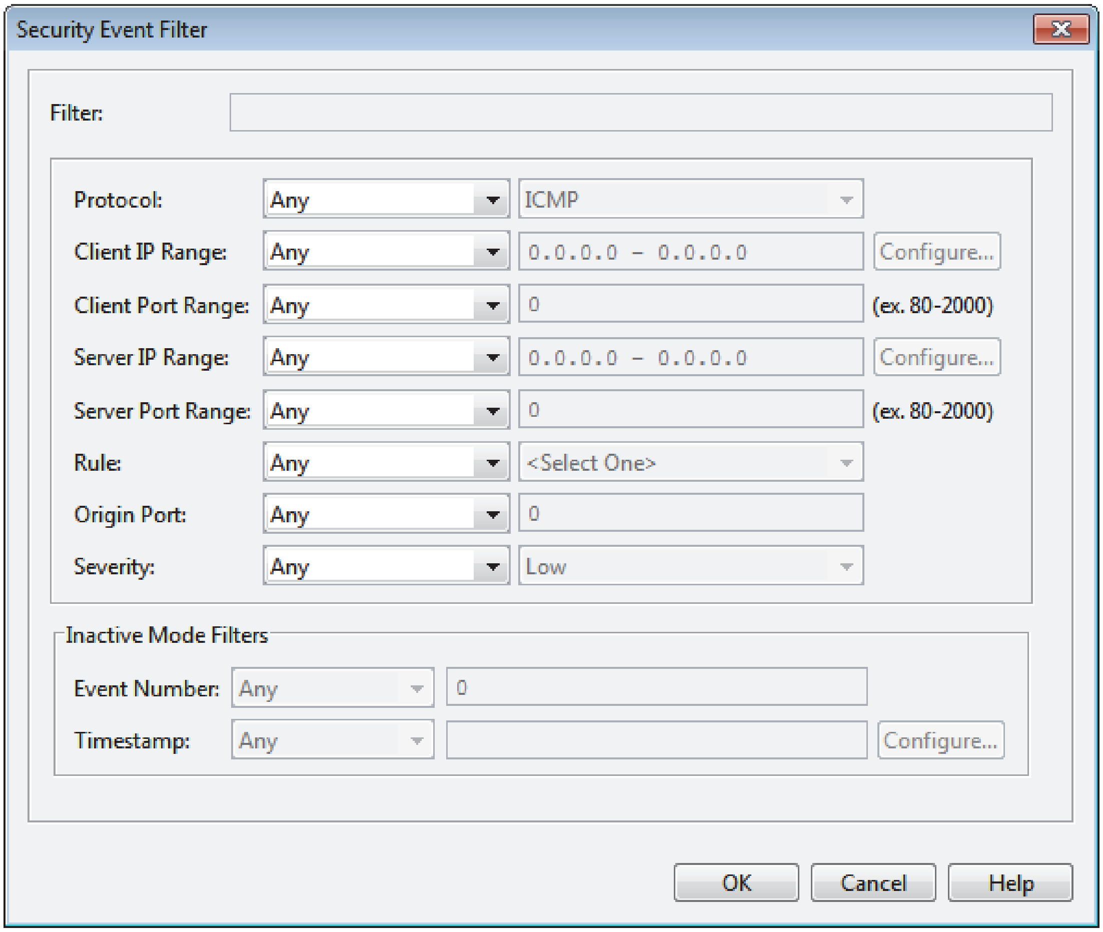

Security Event Filter

Filter:

Protocol: Any ICMP

Client IP Range: Any 0.0.0.0 - 0.0.0.0 Configure...

Client Port Range: Any 0 (ex. 80-2000)

Server IP Range: Any 0.0.0.0 - 0.0.0.0 Configure...

Server Port Range: Any 0 (ex. 80-2000)

Rule: Any <Select One>

Origin Port: Any 0

Severity: Any Low

Inactive Mode Filters

Event Number: Any 0

Timestamp: Any Configure...

To access the Security Event Viewer Filter tool,

1. Do one of the following:
   - From the menu bar, choose Monitor > Security Event Viewer.
   - From the Toolbar, click the Security Events button.

   The IPS Controller Security Event Viewer dialog box displays (Figure 23-5).

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 23-6).

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 23-11).

Table 23-11: 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.

Table 23-11: 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 IPS Controller Management Application:

1. 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 23-7).

Figure 23-7: IPS Controller IP Address Query Dialog Box

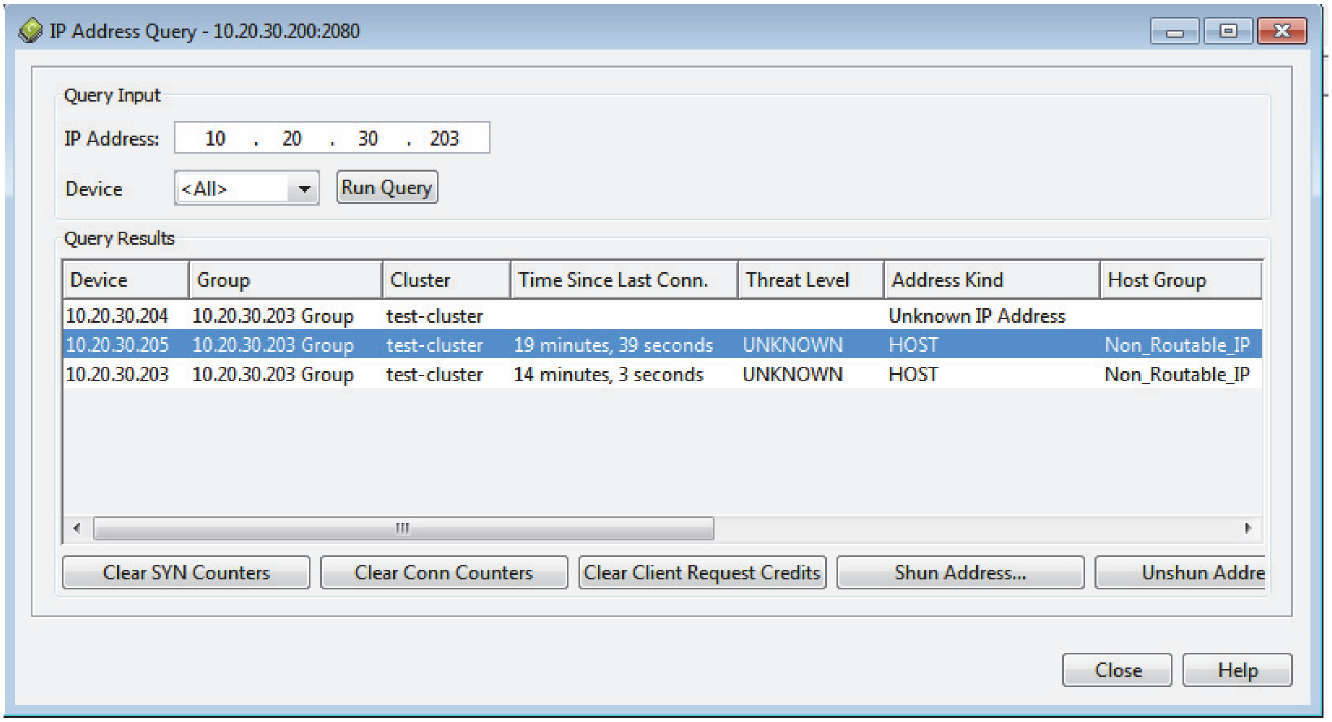

IP Address Query - 10.20.30.200:2080

Query Input

IP Address: 10 . 20 . 30 . 203

Device <All>

Query Results

Device	Group	Cluster	Time Since Last Conn.	Threat Level	Address Kind	Host Group
10.20.30.204	10.20.30.203 Group	test-cluster			Unknown IP Address	
10.20.30.205	10.20.30.203 Group	test-cluster	19 minutes, 39 seconds	UNKNOWN	HOST	Non_Routable_IP
10.20.30.203	10.20.30.203 Group	test-cluster	14 minutes, 3 seconds	UNKNOWN	HOST	Non_Routable_IP

2. Enter the IP address of the host you want to learn about, select the desired device(s) from the drop-down list, then click Run Query. The IP Address Query dialog box displays the query results.

   The IP Address Query dialog box displays the information listed in Table 23-12.

Table 23-12: IP Address Query Information

Item	Description
Device	The Corero Network Device that received traffic from the specified IP address.

Table 23-12: IP Address Query Information *(Continued)*

Item	Description
Group	The policy group to which the Corero Network Device belongs.
Cluster	The cluster (if any) to which the Corero Network Device belongs.
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.

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 counters, 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 also resets the threat level of the IP address to Unknown, and sets the completed connections to zero.

   NOTE

   For more information on clearing counters, see Reset (Clear) SYN Flood and Connection Counters (page 23-29).

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 23-4).

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 23-8).

> NOTE
>
> Alternatively, you can use the IP Address Query window to reset counters for a specific IP address. For more information, see Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).

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 23-8).

Figure 23-8: Clear Counters Dialog Box

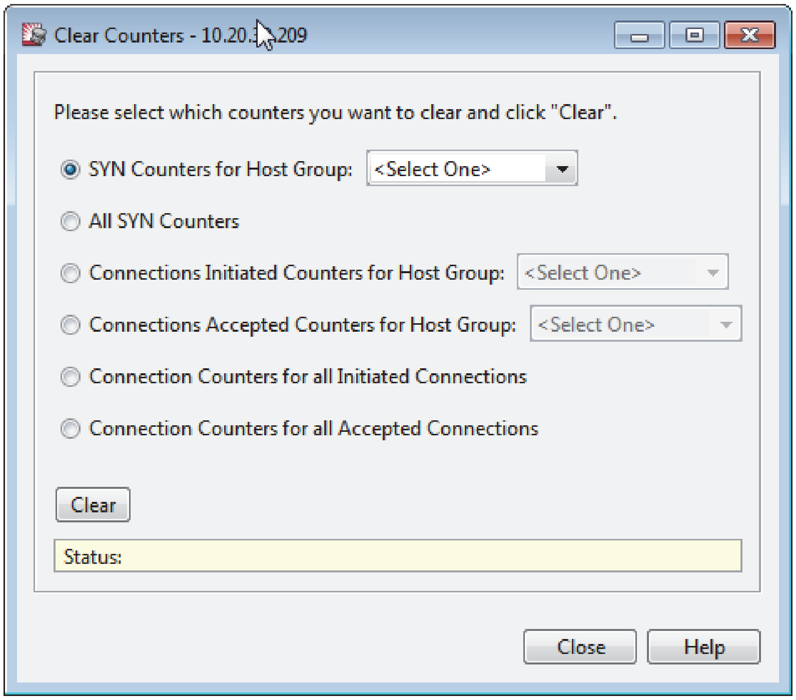

Clear Counters - 10.20. 209

Please select which counters you want to clear and click "Clear".

● SYN Counters for Host Group: <Select One>

○ All SYN Counters

○ Connections Initiated Counters for Host Group: <Select One>

○ Connections Accepted Counters for Host Group: <Select One>

○ Connection Counters for all Initiated Connections

○ Connection Counters for all Accepted Connections

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 menu bar, choose Monitor > Statistics > Dropped Packets.

   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 23-13 lists the reasons why packets might be dropped.

Table 23-13: 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, Corero Network Devices require 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 device dropped because it entered DDoS Rejection mode and was applying this test to new clients.

Table 23-13: 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. Do one of the following:
   - From the menu bar, choose Manage > Devices > Ports.
   - Click the Policy Group & Device Manager toolbar button. On the Policy Group and Device Manager dialog box, click the Ports tab.

   The Port Statistics dialog box displays the information listed in Table 23-14.

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 23-14 lists the information displayed in the Port Statistics dialog box table.

Table 23-14: Port Statistics Information

Column	Description
Device	The Corero Network Device associated with the port.
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.

Table 23-14: Port Statistics Information *(Continued)*

Column	Description
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.
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. You view graphs as part of the available dashboards. You can select different dashboards from the Display drop-down to view different charts. For more information about dashboards, see Using the Dashboard Display (page 5-9).

   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 23-15.

Table 23-15: 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 23-15: 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 Setup	Displays the device’s current rate of setup for various types of connections, including TCP, UDP, and Other IP.
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.

Chapter 24
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 25, "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 24-1.
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 24-1.

Table 24-1: 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 24-2.
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 24-2.

Table 24-2: SYN Flood Protection Profile Settings

Setting	Description
Name	The profile name.

Table 24-2: 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 24-2: 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 24-3 describes the differences between these processes.

Table 24-3: Enabling SYN Flood, Connection, and Client Request Limiting

To enable...	Do this...	And consider the following...
SYN Flood Limiting	On the multi-tabbed Edit Policy 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 24-8). 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 24-8). 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 24-8).
Connection Limiting	On the multi-tabbed Edit Policy 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 24-8). 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 24-8). 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 24-8).
Client Rate Limiting for Client host groups.	On the multi-tabbed Edit Policy 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 25, “Client Rate Limiting”.

Table 24-3: 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 Edit Policy 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 24-8)
- Step 2: Enabling Request Limiting for a Service (page 24-11)
- Step 4: Creating a Rate Based Policy for a Specific Host Group (page 24-13)
- Step 4: Creating a Rate Based Policy for a Specific Host Group (page 24-13)
- Step 5: Enabling the Relevant Rules for the Desired Rule Set (page 24-16)

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 17, “Managing Host Groups”.

To prepare host groups:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Host Groups tab (Figure 24-1). 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.
4. To add a new host group, in the Host Groups area, click Add, then specify a name for the new host group.
5. To add or edit the IP address membership of a host group (Figure 24-1):
   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.

6. To Delete an IP Address Range from the Host Group, select the IP Address Range, then click Delete.
7. When you have finished specifying host group settings in the IPS Controller management application, click Done.
8. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

Figure 24-1: Modifying Host Group Membership

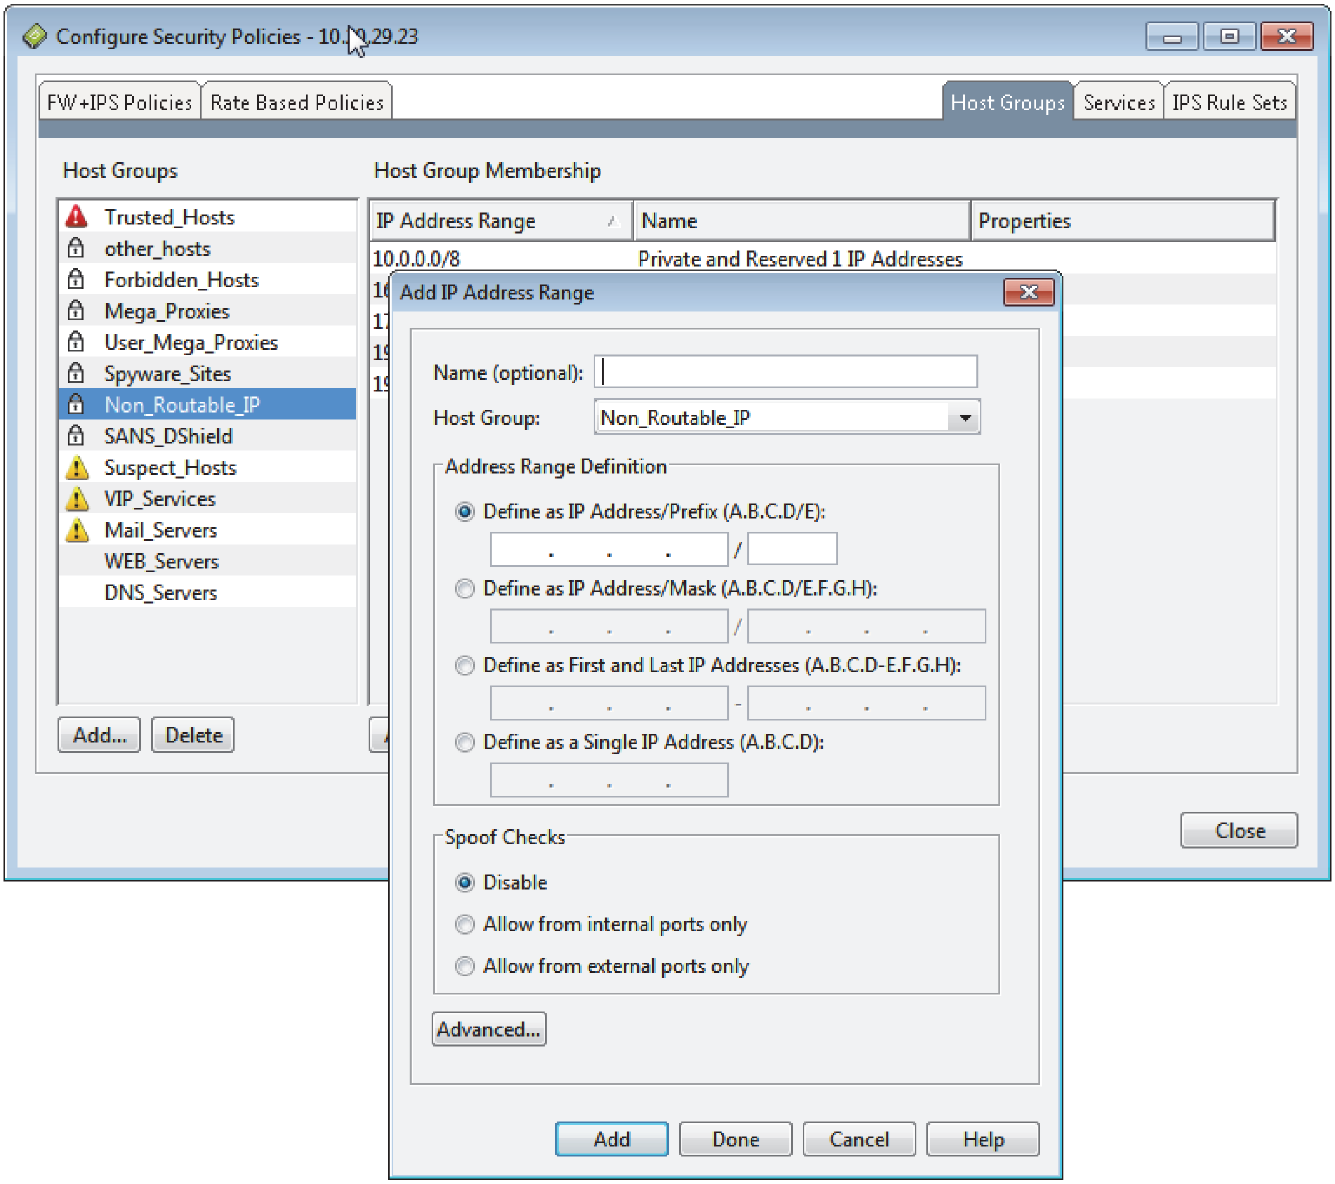

Host Groups

- ⚠ Trusted_Hosts
- 🔒 other_hosts
- 🔒 Forbidden_Hosts
- 🔒 Mega_Proxies
- 🔒 User_Mega_Proxies
- 🔒 Spyware_Sites
- 🔒 Non_Routable_IP
- 🔒 SANS_DShield
- ⚠ Suspect_Hosts
- ⚠ VIP_Services
- ⚠ Mail_Servers
- WEB_Servers
- DNS_Servers

Host Group Membership

IP Address Range	Name	Properties
10.0.0.0/8	Private and Reserved 1 IP Addresses	

Name (optional):

Host Group: Non_Routable_IP

Address Range Definition

- ● Define as IP Address/Prefix (A.B.C.D/E):
- ○ Define as IP Address/Mask (A.B.C.D/E.F.G.H):
- ○ Define as First and Last IP Addresses (A.B.C.D-E.F.G.H):
- ○ Define as a Single IP Address (A.B.C.D):

Spoof Checks

- ● Disable
- ○ Allow from internal ports only
- ○ Allow from external ports only

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 18, "Managing Services".

To enable request limiting for a service:

enable request limiting for a service:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Services tab (Figure 24-2). 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 18, "Managing Services".

4. Select the service to which you want to apply client rate limiting, then click Edit. The Edit Service dialog box displays (Figure 24-2).
5. Click Advanced. The Advanced Service Settings dialog box displays.
6. For request limiting, click the Enabled radio button for Request Limiting, then click OK.
7. When you have finished specifying service settings in the IPS Controller 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.
8. Any time after you finish modifying the service settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

Figure 24-2: Modifying Services

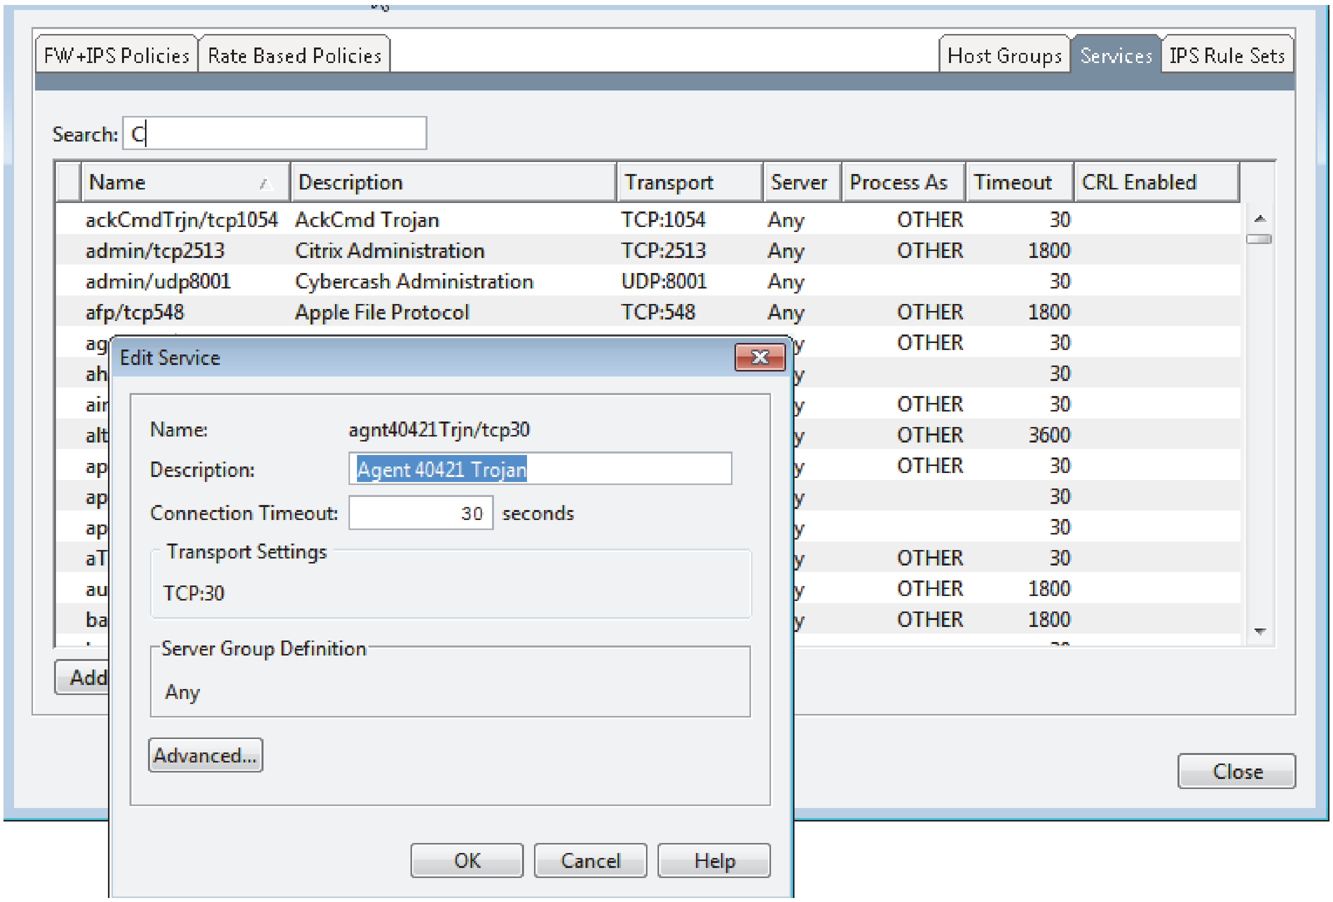

Search: C

Name	Description	Transport	Server	Process As	Timeout	CRL Enabled
ackCmdTrjn/tcp1054	AckCmd Trojan	TCP:1054	Any	OTHER	30	
admin/tcp2513	Citrix Administration	TCP:2513	Any	OTHER	1800	
admin/udp8001	Cybercash Administration	UDP:8001	Any		30	
afp/tcp548	Apple File Protocol	TCP:548	Any	OTHER	1800	
ag			y	OTHER	30	
ah			y		30	
air			y	OTHER	30	
alt			y	OTHER	3600	
ap			y	OTHER	30	
ap			y		30	
ap			y		30	
aT			y	OTHER	30	
au			y	OTHER	1800	
ba			y	OTHER	1800	

Edit Service

Name: agnt40421Trjn/tcp30

Description: Agent 40421 Trojan

Connection Timeout: 30 seconds

Transport Settings

TCP:30

Server Group Definition

Any

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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Rate Based Policies tab.
4. Click either the Clients or Servers category, depending on the host group to which you want the limits applied.
5. 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 24-3).
   - If you selected Servers, the Edit Server Limits dialog box displays. The SYN Flood and Connection Limit options are on the General tab.

Figure 24-3: Rate Based Policies Tab

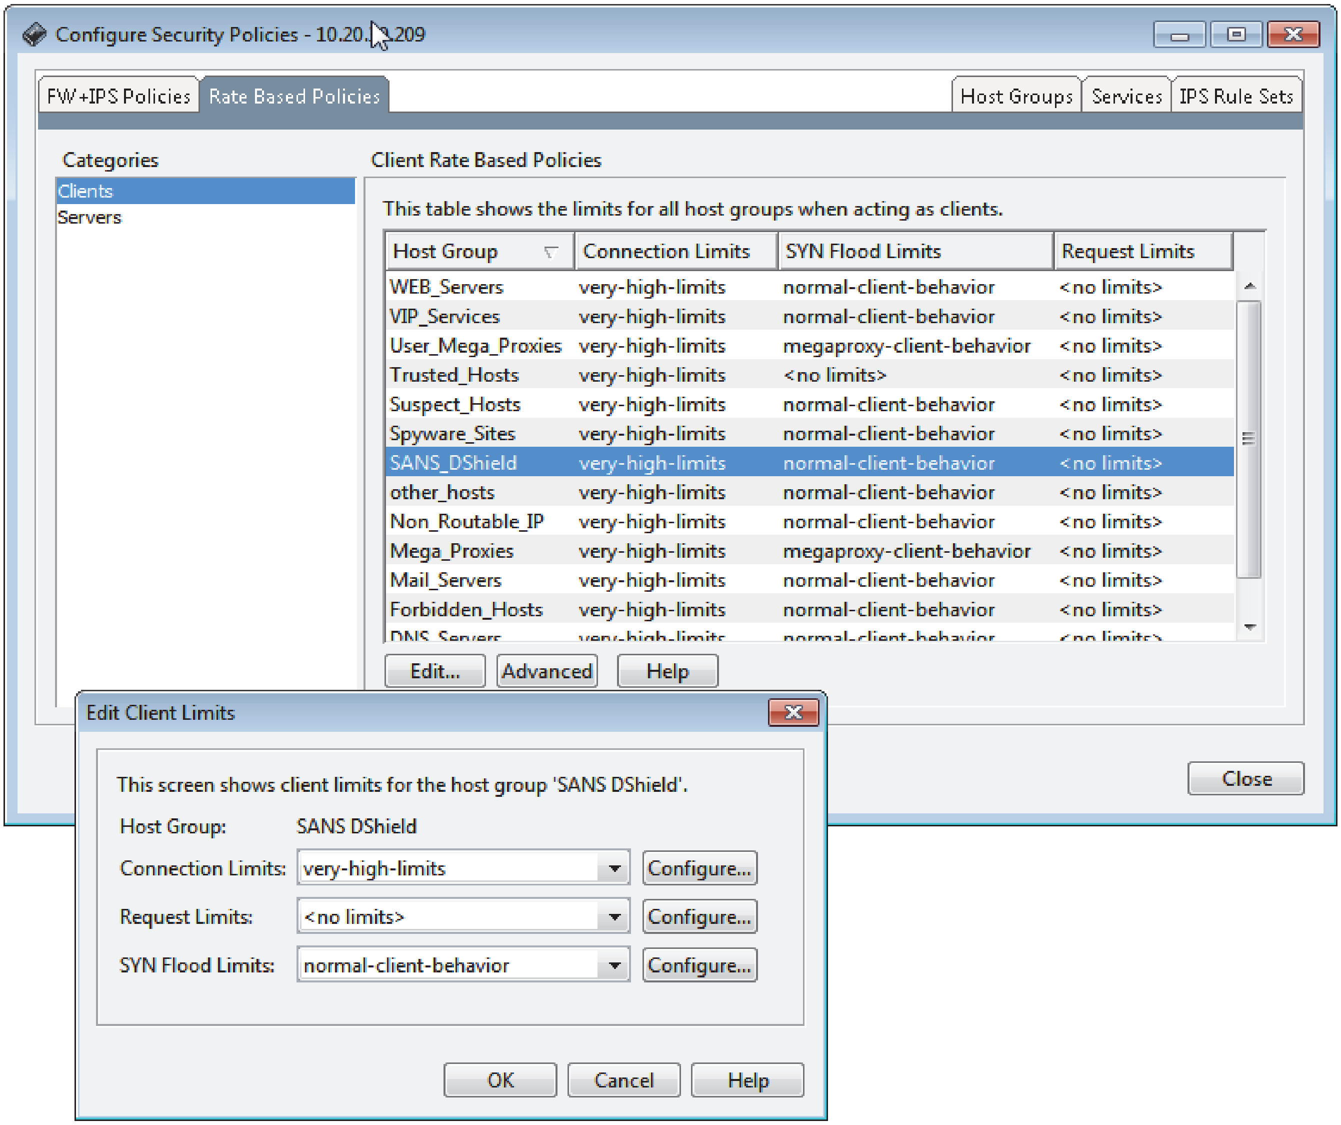

Client Rate Based Policies

This table shows the limits for all host groups when acting as clients.

Host Group	Connection Limits	SYN Flood Limits	Request Limits
WEB_Servers	very-high-limits	normal-client-behavior	<no limits>
VIP_Services	very-high-limits	normal-client-behavior	<no limits>
User_Mega_Proxies	very-high-limits	megaproxy-client-behavior	<no limits>
Trusted_Hosts	very-high-limits	<no limits>	<no limits>
Suspect_Hosts	very-high-limits	normal-client-behavior	<no limits>
Spyware_Sites	very-high-limits	normal-client-behavior	<no limits>
SANS_DShield	very-high-limits	normal-client-behavior	<no limits>
other_hosts	very-high-limits	normal-client-behavior	<no limits>
Non_Routable_IP	very-high-limits	normal-client-behavior	<no limits>
Mega_Proxies	very-high-limits	megaproxy-client-behavior	<no limits>
Mail_Servers	very-high-limits	normal-client-behavior	<no limits>
Forbidden_Hosts	very-high-limits	normal-client-behavior	<no limits>
DNS_Servers	very-high-limits	normal-client-behavior	<no limits>

Edit Client Limits

This screen shows client limits for the host group 'SANS DShield'.

Host Group: SANS DShield

Connection Limits: very-high-limits

Request Limits: <no limits>

SYN Flood Limits: normal-client-behavior

6. 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.

7. 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 24-1.

   d. When finished, click Close. Now you can select it in the drop-down list.

8. 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 24-2.

   d. When finished, click Close. Now you can select it in the drop-down list.

9. 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.

10. When finished, click OK.

11. When you have finished specifying host group settings in the IPS Controller management application, click Done.

12. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the IPS Rule Sets tab. Select the desired rule set.

   NOTE

   For more information on managing rule sets, see Chapter 19, "Managing Rules and Rule Sets".

4. 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 the IPS Controller, click the Advanced button, and choose Rate Based Rules from the drop-down list.
5. To modify the settings for a rule, select the rule, then click Edit. The Edit Rule Settings dialog box displays (Figure 24-4).

Figure 24-4: Modifying Rule Settings

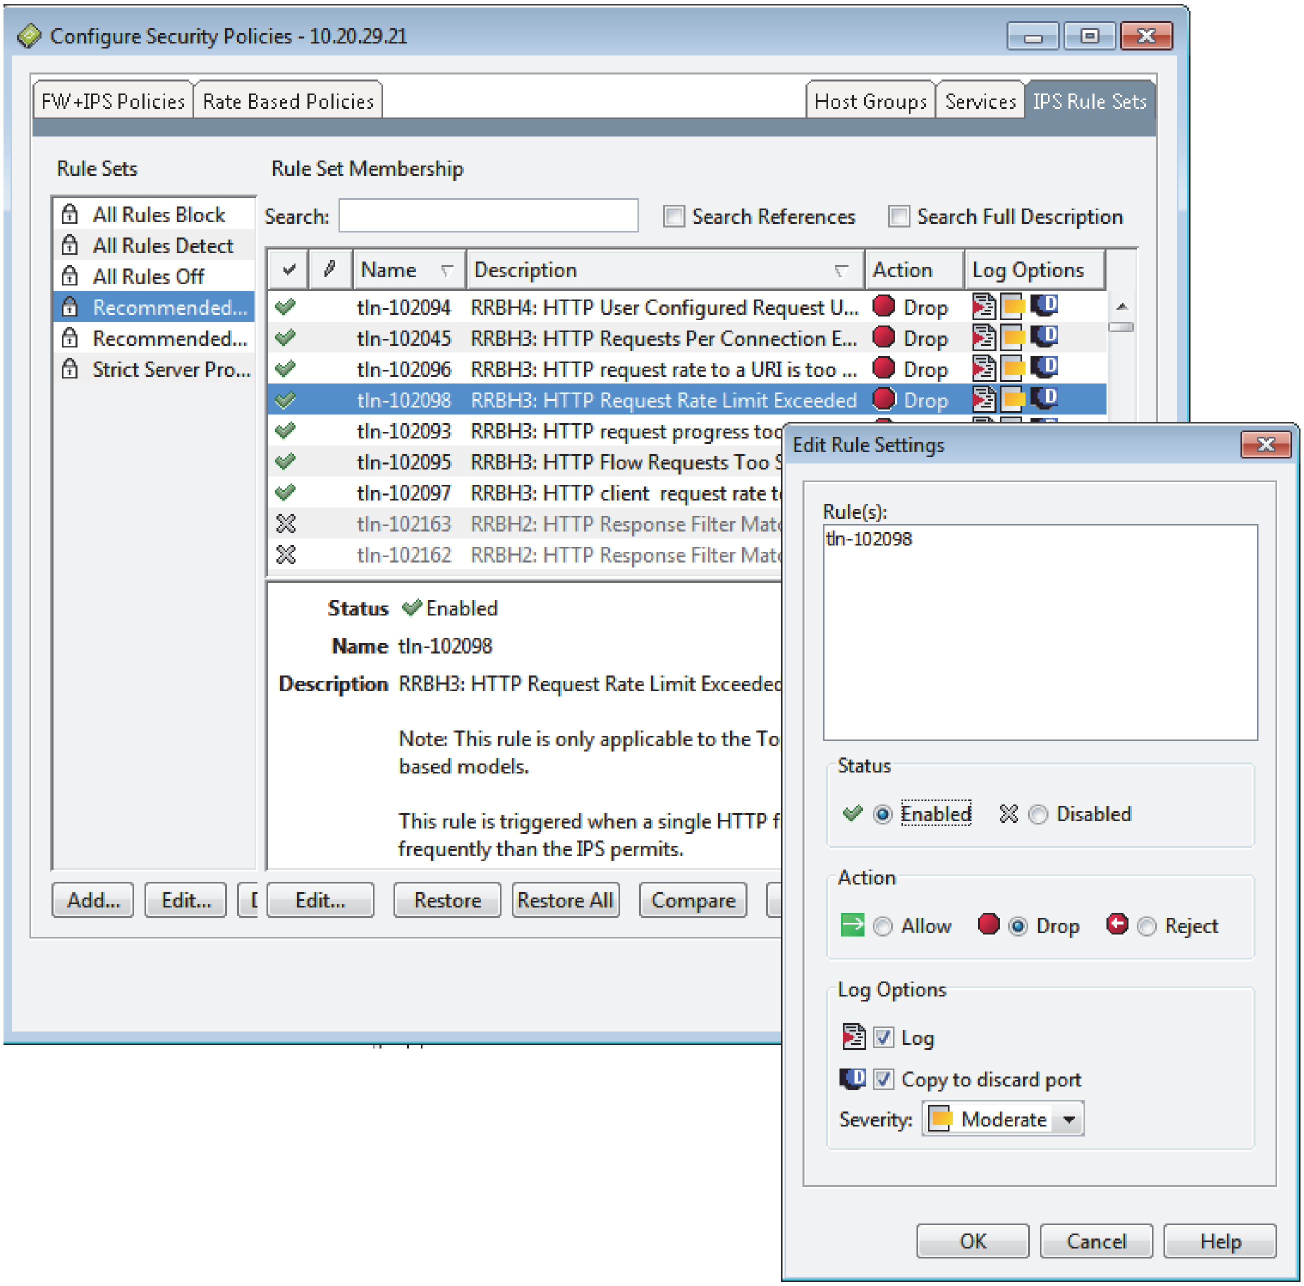

✓		Name	Description	Action	Log Options
✓		tln-102094	RRBH4: HTTP User Configured Request U...	● Drop	
✓		tln-102045	RRBH3: HTTP Requests Per Connection E...	● Drop	
✓		tln-102096	RRBH3: HTTP request rate to a URI is too ...	● Drop	
✓		tln-102098	RRBH3: HTTP Request Rate Limit Exceeded	● Drop	
✓		tln-102093	RRBH3: HTTP request progress too		
✓		tln-102095	RRBH3: HTTP Flow Requests Too S		
✓		tln-102097	RRBH3: HTTP client request rate t		
☒		tln-102163	RRBH2: HTTP Response Filter Matc		
☒		tln-102162	RRBH2: HTTP Response Filter Matc		

Status ✓ Enabled

Name tln-102098

Description RRBH3: HTTP Request Rate Limit Exceeded

Note: This rule is only applicable to the To based models.

This rule is triggered when a single HTTP f frequently than the IPS permits.

Rule(s): tln-102098

Status: ● Enabled ○ Disabled

Action: ○ Allow ● Drop ○ Reject

Log Options: ☑ Log ☑ Copy to discard port

Severity: Moderate

6. 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 19-4). Severity levels include:
  - Low (green)
  - Moderate (yellow)
  - Critical (red)

7. When you have finished specifying rule settings in the IPS Controller management application, click OK.
8. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

   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 IPS Controller management application, select Monitor > IP Address Query from the menu bar.

   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 IPS Controller management application, click Run 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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).

Chapter 25
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:

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
>
> 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 25-1 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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26)

Table 25-1: 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 25-1: 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 (page 25-6)
- Step 2: Enabling Request Limiting for a Service (page 25-8)
- Step 3: Creating a Rate Based Client Limit Policy for a Client Host Group (page 25-10)
- Step 4: Enabling the Relevant Rules for the Desired Rule Set (page 25-12)

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 17, "Managing Host Groups".

To prepare host groups:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. 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.
4. To add a new host group, in the Host Groups area, click Add, then specify a name for the new host group.
5. To add or edit the IP address membership of a host group (Figure 25-1).
   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.

6. To Delete an IP Address Range from the Host Group, select the IP Address Range, then click Delete.

7. When you have finished specifying host group settings in the IPS Controller 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.

8. Any time after you finish modifying the service settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

Figure 25-1: Modifying Host Group Membership

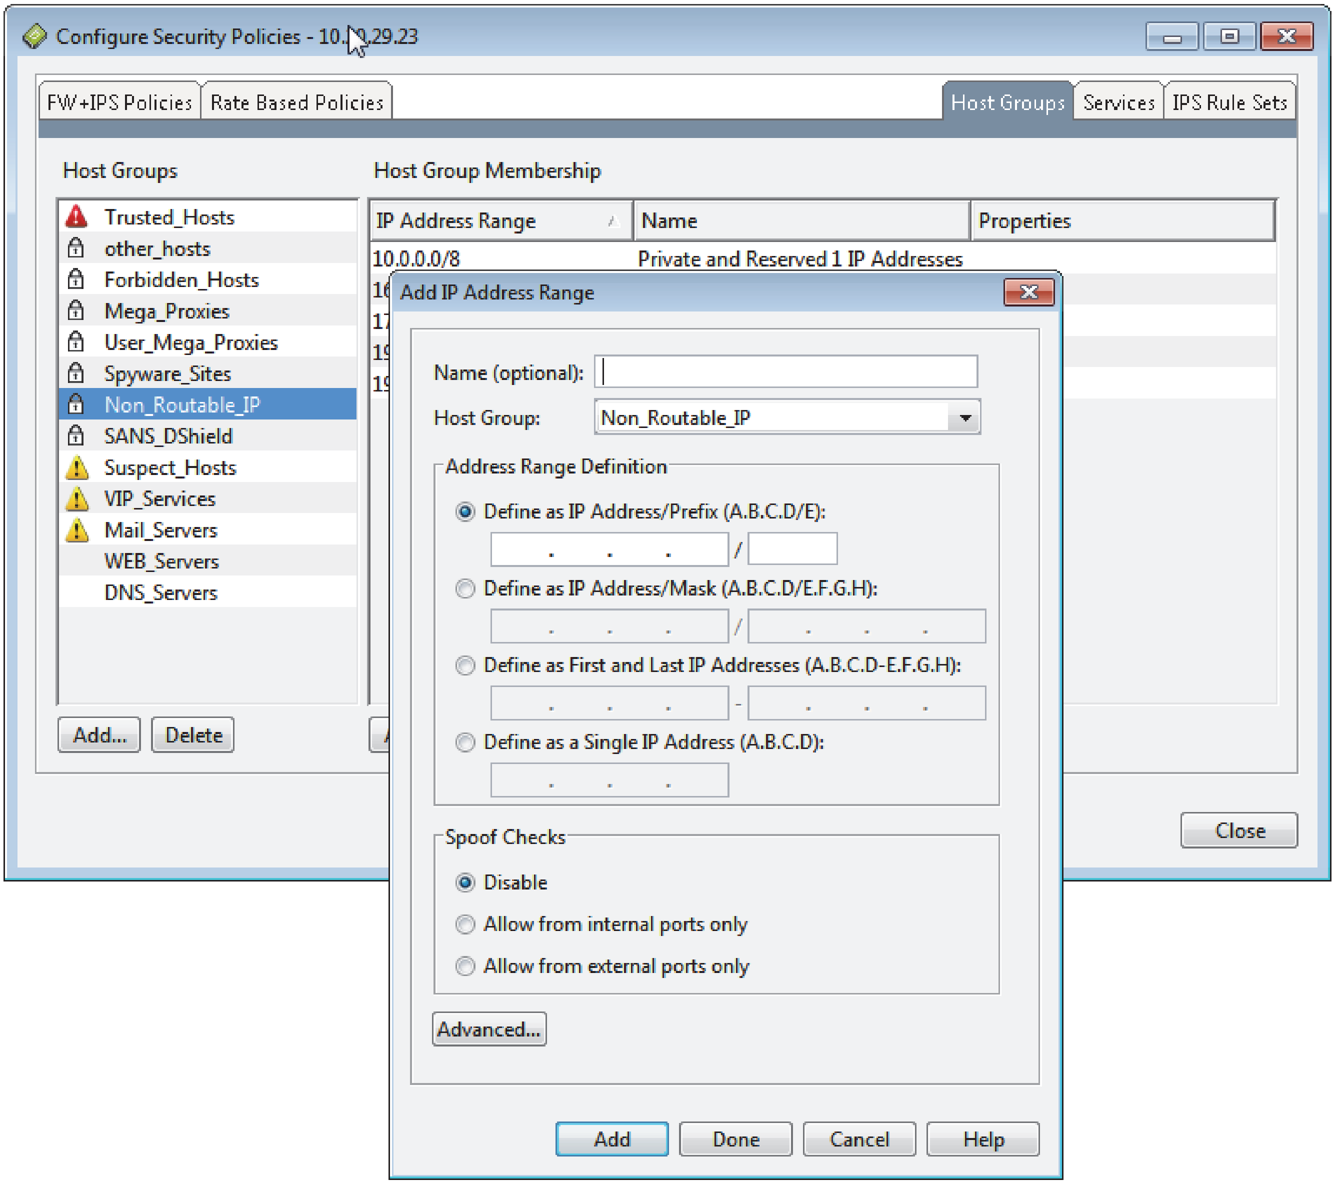

Host Groups

- Trusted_Hosts
- other_hosts
- Forbidden_Hosts
- Mega_Proxies
- User_Mega_Proxies
- Spyware_Sites
- Non_Routable_IP
- SANS_DShield
- Suspect_Hosts
- VIP_Services
- Mail_Servers
- WEB_Servers
- DNS_Servers

Host Group Membership

IP Address Range	Name	Properties
10.0.0.0/8	Private and Reserved 1 IP Addresses	

Name (optional):

Host Group: Non_Routable_IP

Address Range Definition

● Define as IP Address/Prefix (A.B.C.D/E):

○ Define as IP Address/Mask (A.B.C.D/E.F.G.H):

○ Define as First and Last IP Addresses (A.B.C.D-E.F.G.H):

○ Define as a Single IP Address (A.B.C.D):

Spoof Checks

● Disable

○ Allow from internal ports only

○ Allow from external ports only

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 18, "Managing Services".

To enable request limiting for a service:

1. Do one of the following:
   - Click the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. 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 18, "Managing Services".

4. Select the service to which you want to apply client rate limiting, then click Edit. The Edit Service dialog box displays (Figure 25-2).
5. Click Advanced. The Advanced Service Settings dialog box displays.
6. For request limiting, click the Enabled radio button for Request Limiting, then click OK.
7. When you have finished specifying service settings in the IPS Controller 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.
8. Any time after you finish modifying the service settings, you can push the updated settings out to the Policy Group by clicking Push Settings.

Figure 25-2: Modifying Services

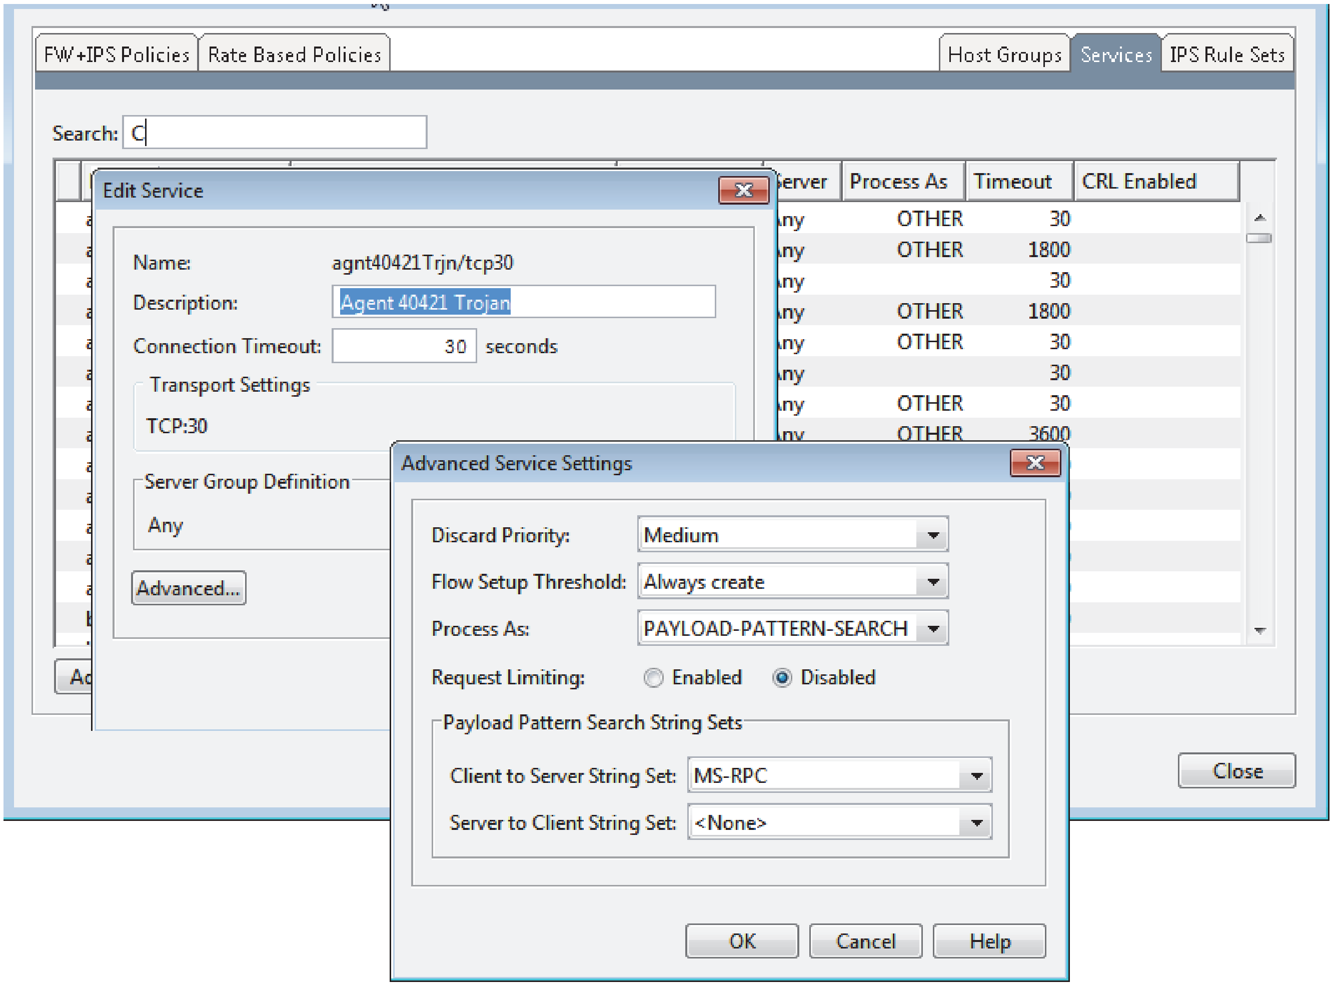

Search: C

erver	Process As	Timeout	CRL Enabled
ny	OTHER	30	
ny	OTHER	1800	
ny		30	
ny	OTHER	1800	
ny	OTHER	30	
ny		30	
ny	OTHER	30	
ny	OTHER	3600	

Name: agnt40421Trjn/tcp30

Description: Agent 40421 Trojan

Connection Timeout: 30 seconds

Transport Settings

TCP:30

Server Group Definition

Any

Discard Priority: Medium

Flow Setup Threshold: Always create

Process As: PAYLOAD-PATTERN-SEARCH

Request Limiting: ○ Enabled ● Disabled

Payload Pattern Search String Sets

Client to Server String Set: MS-RPC

Server to Client String Set: <None>

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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the Rate Based Policies tab.
4. Click Clients, then select the host group whose client requests you want to rate limit
5. Click Edit. The Edit Client Limits dialog box displays (Figure 25-3).
6. 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.
7. 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 24, “SYN Flood and Connection Limiting Security.
8. When finished, click OK.
9. When you have finished specifying host group settings in the IPS Controller management application, click Done.
10. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

Figure 25-3: Rate Based Policies Tab

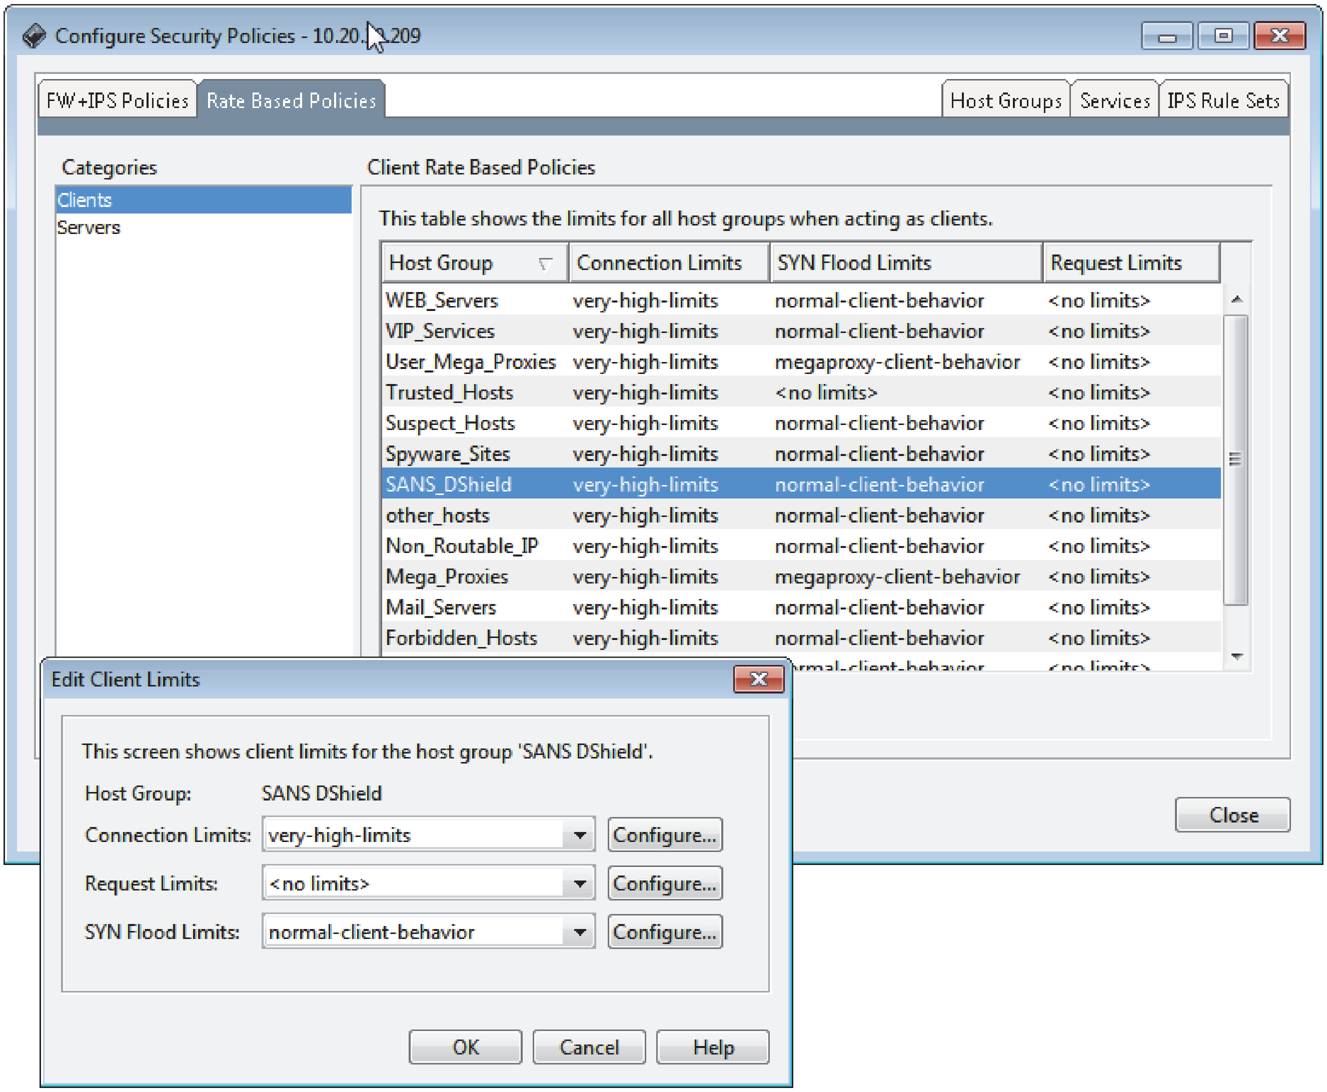

This table shows the limits for all host groups when acting as clients.

Host Group	Connection Limits	SYN Flood Limits	Request Limits
WEB_Servers	very-high-limits	normal-client-behavior	<no limits>
VIP_Services	very-high-limits	normal-client-behavior	<no limits>
User_Mega_Proxies	very-high-limits	megaproxy-client-behavior	<no limits>
Trusted_Hosts	very-high-limits	<no limits>	<no limits>
Suspect_Hosts	very-high-limits	normal-client-behavior	<no limits>
Spyware_Sites	very-high-limits	normal-client-behavior	<no limits>
SANS_DShield	very-high-limits	normal-client-behavior	<no limits>
other_hosts	very-high-limits	normal-client-behavior	<no limits>
Non_Routable_IP	very-high-limits	normal-client-behavior	<no limits>
Mega_Proxies	very-high-limits	megaproxy-client-behavior	<no limits>
Mail_Servers	very-high-limits	normal-client-behavior	<no limits>
Forbidden_Hosts	very-high-limits	normal-client-behavior	<no limits>

This screen shows client limits for the host group 'SANS DShield'.

Host Group: SANS DShield

Connection Limits: very-high-limits

Request Limits: <no limits>

SYN Flood Limits: normal-client-behavior

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 the Policy Group and Device Manager toolbar button, then click the Settings tab.
   - Choose Manage > Policy Group > Settings from the menu bar.
2. Select a Policy Group on the Settings tab, then click Edit Settings. The Edit Policy dialog box displays.
3. Click the IPS Rule Sets tab. Select the desired rule set.

   NOTE

   For more information on managing rule sets, see Chapter 19, "Managing Rules and Rule Sets".
4. 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 the IPS Controller, click the Advanced button, and choose Rate Based Rules from the drop-down list.
5. To modify the settings for a rule, select the rule, then click Edit. The Edit Rule Settings dialog box displays (Figure 25-4).

Figure 25-4: Modifying Rule Settings

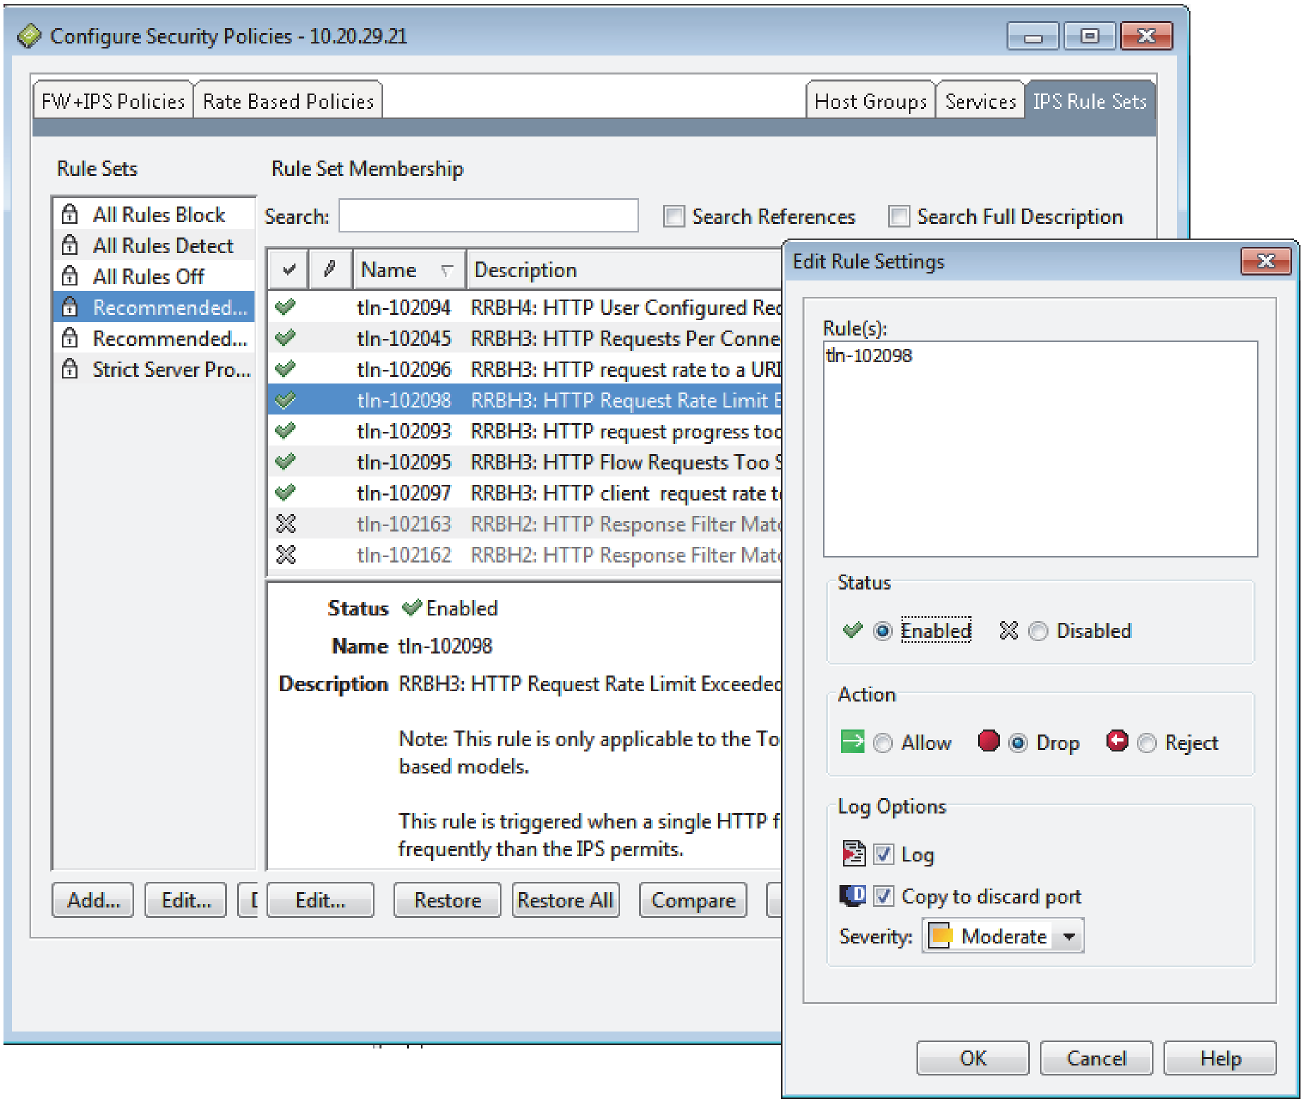

✓		Name	Description
✓		tln-102094	RRBH4: HTTP User Configured Rec
✓		tln-102045	RRBH3: HTTP Requests Per Conne
✓		tln-102096	RRBH3: HTTP request rate to a URI
✓		tln-102098	RRBH3: HTTP Request Rate Limit E
✓		tln-102093	RRBH3: HTTP request progress too
✓		tln-102095	RRBH3: HTTP Flow Requests Too S
✓		tln-102097	RRBH3: HTTP client request rate t
☒		tln-102163	RRBH2: HTTP Response Filter Matc
☒		tln-102162	RRBH2: HTTP Response Filter Matc

Status ✓ Enabled

Name tln-102098

Description RRBH3: HTTP Request Rate Limit Exceeded

Note: This rule is only applicable to the To
based models.

This rule is triggered when a single HTTP f
frequently than the IPS permits.

Rule(s):
tln-102098

Status
✓ ● Enabled ☒ ○ Disabled

Action
○ Allow ● Drop ○ Reject

Log Options
☑ Log
☑ Copy to discard port
Severity: Moderate

6. 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 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.

CAUTION

If you disable a rule, packets that would normally trigger the rule are passed, and no logging is performed.

7. 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 19-4). Severity levels include:
     - Low (green)
     - Moderate (yellow)
     - Critical (red)
8. When you have finished specifying rule settings in the IPS Controller management application, click OK.
9. Any time after you finish, you can push the updated settings out to the selected Policy Group by clicking Push Settings on the Settings tab of the Policy Group and Device Manager dialog box.

   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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26). To query an IP address:

1. On the IPS Controller management application, select Monitor > IP Address Query from the menu bar.

   The IP Address Query dialog box displays (Figure 23-7).

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 IPS Controller management application, click Run 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 Using IP Address Query to Learn About a Host and Clear Counters (page 23-26).

Appendix A
IPS Controller System Management

The IPS Controller 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:

- Managing the IPS Controller Service (page A-2)
- Backing Up and Restoring the IPS Controller (page A-5)

For information on managing the IPS Controller’s software, see Chapter 4, “Installing and Upgrading the IPS Controller Software”.

Managing the IPS Controller Service

The IPS Controller runs as a service on your Linux computer system. By default, the tlnipscd (IPSC) service working directory is `/usr/local/tlnipscd`. You must be logged in as root in order to manage IPS Controller service operation and parameters.

The tlnipscd service is controlled using standard Linux commands, such as chkconfig and service.

The rest of this section describes the following:

- Starting the IPS Controller Service (page A-2)
- Stopping the IPS Controller Service (page A-2)
- Restarting the IPS Controller Service (page A-2)
- Preventing the IPS Controller Service from Starting (page A-3)
- Obtaining the Status of the IPS Controller Service (page A-3)
- IPS Controller Service Port Settings (page A-4)

Starting the IPS Controller Service

To start the IPS Controller service:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Enter the following command:

```
> /sbin/service tlnipscd start
```

As an example, the system might respond with this:

```
Starting tlnipscd daemon:/usr/local/tlnipscd 2080 2443 0
```

The system's response shows that the service was started and also displays important service elements:

- /usr/local/tlnipscd - the working root directory
- 2080 - the HTTP service port
- 2443 - the HTTPS service port
- 0 - the TELNET service port. Zero means this service is disabled.

Stopping the IPS Controller Service

To stop the IPS Controller service:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Enter the following command:

```
> /sbin/service tlnipscd stop
```

The system responds:

```
Stopping tlnipscd daemon:
```

Restarting the IPS Controller Service

At times, it may be necessary to restart the IPS Controller service. For example, you might need to do this because the HTTP port was changed

To restart the IPS Controller service:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Enter the following command:

   ```
   > /sbin/service tlnipscd restart
   ```

The system responds:

```
Stopping tlnipscd daemon:
Starting tlnipscd daemon:/usr/local/tlnipscd 2080 2443 0
```

A force reload has the same effect as restart:

```
> /sbin/service tlnipscd force-reload
```

There is another command called `condrestart` that issues a restart only if the service is currently running. If the service is not running then this command has no effect.

Preventing the IPS Controller Service from Starting

At times, you may wish to prevent the IPS Controller service from starting. To do so:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Issue the following command:

   ```
   > chkconfig --del tlnipscd
   ```

Obtaining the Status of the IPS Controller Service

To obtain the current status of the IPS Controller status:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Issue the following command.

   ```
   > /sbin/service tlnipscd status
   ```

The system responds by displaying the status of the IPS Controller service, followed by detailed information about the IPS Controller:

```
tlnipscd is started.
tlnipscd version = V470049
Management session count = 0
Installation directory = /usr/local/tlnipscd
HTTP port = 2080
HTTPS port = 2443
TELNET port = 2023
TopResponse license is installed
TopResponse update version = 2011083101
```

The status information that displays includes:

- tlnipscd service status
- tlnipscd version
- Management session count: This number indicates how many management users are logged in to the service. One user is counted for each user logged into the management interface, and each user logged in to the State Browser.
- Installation directory (file system root for the service)

- HTTP port
- HTTPS port
- TELNET port: Corero recommends that this port is so seldom used that you should disable it by setting it to zero. It is possible that a user may be asked to use TELNET by customer support for diagnostic purposes, but otherwise it is typically unused.
- TopResponse license (indicates the presence of a TopResponse license)
- TopResponse update version (currently loaded TopResponse Protection Pack)

IPS Controller Service Port Settings

When you modify the IPS Controller service port settings, the changes do not take affect until the service is restarted.

To set the HTTP port used by the IPS Controller service:

1. Log into the IPS Controller Linux machine as root. You must be logged in as root to perform this procedure.
2. Issue the following command:

```
> /sbin/service tlnipscd httpport port#
```

Where *port#* is the number of the port you want to assign to the service.

To set the HTTPS port used by the IPS Controller service issue the following command:

```
> /sbin/service tlnipscd httpsport port#
```

To set the TELNET port used by the IPS Controller service issue the following command. Setting the TELNET port to zero disables the TELNET service.

```
> /sbin/service tlnipscd telnetport port#
```

Backing Up and Restoring the IPS Controller

Corero recommends that you regularly back up IPS Controller working directory. The frequency with which you back up this data depends on your site's IT policies and the frequency with which you modify Corero product configuration information.

CAUTION

The IPS Controller will be stopped during the backup process. When you initiate a backup, the IPS Controller automatically stops the IPS Controller service, performs the backup, then restarts the service.

From an administrative standpoint, you should only backup the IPS Controller when there are no management sessions in progress. The service status command can be used to determine if there are users logged in to the IPS Controller. For more information on managing the tlnipscd service, see Managing the IPS Controller Service (page A-2).

Backing Up the IPS Controller

To back up the tlnipscd working directory, while logged in as root, issue the following command:

```
/sbin/service tlnipscd backup <optional path>
```

If you only need to back up the configuration files on the IPS Controller, while logged in as root, use the following command:

```
/sbin/service tlnipscd configbackup <optional path>
```

By default, the backup is stored in `/usr/local/tlnipscd/backup`.

If you wish, you can specify an optional path where you would like the backup file stored.

The system responds:

```
Stopping tlnipscd daemon:
(... lists files as they are backed up)
Backup written to /usr/local/tlnipscd/backup/04-Aug-09-10_11_28.tgz
Starting tlnipscd daemon: /usr/local/tlnipscd 3080 3443 3023
```

The backup process creates a g-zipped tar file named `dd-mm-yy-hh-mm-ss>.tgz`.

The backup includes the following content:

- The service control script and configuration file.

  The service configuration file is named `/etc/tlnipscd.conf`. It contains the working directory and the service ports used by the tlnipscd daemon.

- All tlnipscd service executable code.
- Factory and user configuration data.
- TopResponse license key and downloaded protection packs.
- Installed IPS 5500 upgrade packages.

A variation of the backup command is available which performs the same actions as the full backup command, but only backs up the configuration files. This command is:

```
tlnipscd configbackup <optional path>
```

Restoring the IPS Controller

An IPS Controller backup can be restored on any machine that has met the prerequisites for installation. These prerequisites are described in Chapter 3, “IPS Controller Pre-Installation Requirements”.

1. Log into your Linux machine as root.
2. Use the cd command to change to the root directory.
3. To restore the files, use the following command:

   ```
   tar -xvf filename
   ```

   Where *filename* is the name of the backup tar file.

4. Start the IPS Controller by issuing the following command:

   ```
   service tlnipscd start
   ```

Appendix B
Troubleshooting IPS Controller Issues

The IPS Controller management application is useful in troubleshooting network issues. Primary views into network behavior include checking port statistics, CPU usage, connection utilization, connection graphs, and top attackers to look for anomalous traffic or behavior.

This chapter contains the following sections:

Recovering Connectivity When You Accidentally Lose Management Access

If you attempted to change the Management VLAN ID, and you have unintentionally isolated yourself from managing the Corero Network Device through its management application, you can access the device through its CONSOLE port and use its Command Line Interface to temporarily disable VLAN enforcement. This will enable you to reconnect with the device.

To recover connectivity:

1. Connect a serial cable with DB-9 connectors to the CONSOLE port on the Corero Network Device and your management PC.
2. On your management PC, using a terminal emulation program, connect to the device using the following settings:
   - Bits per Second: 115200
   - Data Bits: 8
   - Parity: None
   - Stop Bits: 1
   - Flow Control: None
3. Press <Enter> if necessary until the prompt displays.\
4. At the prompt enter:

   ```
   >me set vlan noenforce <Enter>
   ```

   This command turns VLAN enforcement off.
5. You should now be able to connect to the Graphical User Interface and modify VLAN settings as needed.

   NOTE

   The above change is not persistent. If the device reboots, the device reverts to the default behavior of enforcing VLAN checks.
6. To restore VLAN enforcement, do one of the following:
   - From the Navigation Tree, choose Configure System > Advanced System Config > Management Bridging > VLAN Settings, and enable VLAN enforcement.
   - From the CLI, you can use the following command:

     ```
     >me set vlan enforce
     ```

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 for a Corero Network Device using the IPS Controller management application:

1. Do one of the following:
    - From the Menu bar, choose Monitor > Reports > Devices.
    - Click the Policy Group & Device Manager tool bar button. On the Policy Group and Device Manager dialog box, click the Diagnostics & Reports tab.

    The Diagnostics and Reports tab displays.

2. Select the desired device, then click Download Diagnostics.

    The IPS Controller gathers information from the selected device. You can view the status of diagnostic generation and download in the Download Status column. You can view progress in the % Done column.

3. To view the diagnostic information files, once the diagnostic download has been completed, ensure the device is selected and click Open Diags Package. The zip package containing all of the diagnostic information files opens.

Troubleshooting Corero Network Device Connection Issues

When you add the device to the controller, the IPS controller immediately attempts to establish a management link to the device. You can watch the connection process in the State column of the Policy Group and Device Manager Membership tab to verify the device has been successfully added. The appropriate connection process messages are listed in Table B-1.

Table B-1. Corero Network Device Connection Messages

Step	State	Description
1	Connecting	The IPS Controller attempts to connect to the IP address you specified for the Corero Network Device.
2	Authenticating	Once connected, the IPS Controller uses the shared management key to authenticate itself with the device.
3	Synchronizing	Once authenticated, the IPS Controller attempts to synchronize with the device by capturing a copy of all settings or parameters on the device.
4	Operational	Once synchronized, the device becomes operational. The IPS Controller is now able to manage changes to device settings and monitor the device’s security events.

NOTE

Data flows from a Corero Network Device to its configured Syslog servers even if a device’s connection to the IPS Controller is broken. Data created by the IPS or DDS Unit is sent directly to its configured Syslog servers. Changing the shared key (which means temporarily breaking the connection between the IPS or DDS Unit and the controller), does not affect data flow from a Corero Network Device to its associated Syslog servers.

If the connection state for a Corero Network Device remains in any state other than Operational for an extended period of time, this indicates the IPS Controller is unable to connect to the device. If this occurs, assess the following:

1. Verify that management by the IPS Controller is enabled for the Corero Network Device. You can check and modify this setting using the device’s management application.
2. Ensure that the IP address you entered for the device in the ISPC management application is correct.
3. Ensure that there is network connectivity between the IPS Controller and the device. You can do this in two ways:
    - If you find that a managed Corero Network Device is unable to reach Operational status, you can instruct the IPS Controller to attempt to reconnect to the device. For instructions on how to reconnect to a device, see Reconnecting to a Corero Network Device (page 7-6).
    - Verify that there is physical network connectivity between the Corero Network Device and the IPS Controller. Ensure all cables are firmly connected, and check to ensure the network path is operational between the controller and the managed device.
4. Ensure the Corero Network Device is on and fully operational.
5. If the IPS Controller is not able to authenticate its connection with the device, ensure you entered the correct shared management key for the device. These keys must match exactly and it, for any reason, you need to modify the shared key on a device, you will need to perform a corresponding key update in the IPS Controller management application.

Troubleshooting ProtectionCluster Issues

If you look in the IPS Controller management interface, the display will indicate whether or not a ProtectionCluster, or a ProtectionCluster member, has a problem. The space before the cluster or device will display an icon indicating the severity of the problem.

There are two types of status severity icons:

- A yellow alert icon indicates a warning. ⚠
- A red alert icon indicates a critical issue. ❌

NOTE

For information on monitoring ProtectionCluster status, see Viewing ProtectionCluster Status (page 14-11)

When a problem occurs with a ProtectionCluster, one or more alerts display. For this reason, troubleshooting begins with viewing alerts. You can view alerts in two ways:

- To view all alerts on the IPS Controller, choose Monitor > Alerts from the menu bar.
- To view alerts on a particular device or ProtectionCluster, right-click the item in the policy group tree, then choose View Alerts from the pop-up menu.

NOTE

For detailed information on viewing alerts on the IPS Controller, see Viewing the IPS Controller Management Alert Table (page 21-5).

describes the cluster-related alerts, what causes them, and what to do to correct them.

Table B-2. Troubleshooting ProtectionCluster Alerts

Alert Message	Probable Causes	Suggested Actions
Configuration locally modified, need push or pull.	You may have moved the a ProtectionCluster into a policy group that has a different configuration than one or more of the devices used previously. Or, the configurations did not match up on the devices in the pre-existing cluster and so the policy group you created by pulling the configuration from one of the devices does not match the other devices. Or, you have changed the configuration of the policy group and need to push the changes out to the devices.	For the ProtectionCluster to work properly, all member devices must have the same security configuration. Push the policy group's configuration out to all the devices.
Cluster contains one failed device.	one device in a cluster is unreachable.	Be sure the unreachable device is turned on and that its cables are properly connected. Troubleshoot network connectivity to the device in question.

Table B-2. Troubleshooting ProtectionCluster Alerts *(Continued)*

Alert Message	Probable Causes	Suggested Actions
Cluster members claim to have no operational peers.	Clustered devices are functioning properly, but are unaware of each other.	Connectivity between devices on the high-availability Gigabit Ethernet links is broken and must be restored. Check the connecting cables.
Cluster does not have enough devices.	Not all devices in a cluster are seen by the IPS Controller.	Typically, this occurs when you are identifying a cluster that already existing and is operating within the network. Before you click the Add Cluster button, be sure you have added all the devices to the IPS Controller and selected all the devices in the cluster.
High availability link error.	One of the devices has a missing, loose, or damaged cable linking its high-availability Gigabit Ethernet link to its peers.	Check the HA link cables between the devices.
Cluster nonfunctional on device.	A device within a cluster indicates that its cluster-state is Disabled.	The device needs to be rebooted.
Cluster members disagree on cluster size.	one of the devices thinks there is a different number of devices in the cluster than there really are.	Delete the cluster and add it again. Refer to Create and Monitor a ProtectionCluster for prerequisites and detailed add instructions.
Too many devices claim to be in this cluster.	The total number of devices in the cluster is not being reported properly.	Delete the cluster and add it again. Refer to Create and Monitor a ProtectionCluster for prerequisites and detailed add instructions.
Multiple cluster members claim the same cluster index.	The cluster could not be created properly.	You may have added a pre-existing cluster incorrectly, or perhaps the devices were not clearly in a cluster state when you added them. Delete the cluster and add it again using Create and Monitor a ProtectionCluster for prerequisites and detailed add instructions.
Device configured for too many cluster members.	Somehow the same device is acting as multiple members of the cluster.	Delete the cluster and add it again. Refer to Create and Monitor a ProtectionCluster for prerequisites and detailed add instructions.
Cluster members disagree on definition of cluster peers.	The cluster definition did not get created properly	Delete the cluster and add it again. Refer to Create and Monitor a ProtectionCluster for prerequisites and detailed add instructions.

Index