ProCurve Network Immunity Manager

Security Administrator's Guide

Publication Number

5991-8566

April, 2007
Edition 1.0

Trademark Credits

Microsoft, Windows, Windows 95, and Microsoft Windows NT are registered trademarks of Microsoft Corporation. Internet Explorer is a trademark of Microsoft Corporation. Ethernet is a registered trademark of Xerox Corporation. Netscape is a registered trademark of Netscape Corporation.

Disclaimer

The information contained in this document is subject to change without notice.

HEWLETT-PACKARD COMPANY MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

Warranty

See the Customer Support/Warranty booklet included with the product.

A copy of the specific warranty terms applicable to your Hewlett-Packard products and replacement parts can be obtained from your HP Sales and Service Office or authorized dealer.

Hewlett-Packard Company
8000 Foothills Boulevard, m/s 5551
Roseville, California 95747-5551
http://www.procurve.com

Contents

4 Using Security Policies

1

About ProCurve Network Immunity Manager

Chapter Contents

Introduction

Network usage has skyrocketed with the expansion of the Internet, wireless, and convergence technologies. This increases the burden on network managers working to control network usage. Also, the complexity of large networks makes it difficult to control network access and usage by individual users.

Together with ProCurve Manager Plus (PCM+), the Network Immunity Manager (NIM) provides an affordable, feature-rich, unified, centralized approach to network security management.

Network Immunity Manager provides an affordable, scalable, and easily managed method for providing per port intrusion detection and response to stop malicious network traffic at the edge of the network.

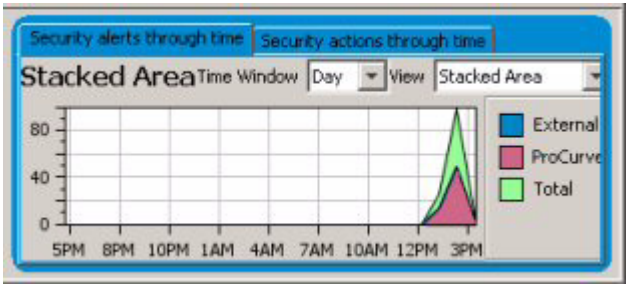

Figure 1-1. ProCurve Network Immunity Manager, dashboard panel

Why Network Immunity Manager?

Network Immunity Manager is an add-on module to ProCurve Manager that supports security policy setting and monitors network activity for anomalous behaviors that may indicate the presence of security threats such as worms, or viruses. NIM detects viruses by behavior anomaly detection on sampled traffic, and it accepts virus throttling alerts from switches using ProCurve's Virus Throttle feature, as well as accepting alerts from other IDS/IPS/UTM security appliances, such as Cisco IPS, etc.

ProCurve Network Immunity Manager can be used to automatically detect and respond to virus attacks on per port basis. It lets you create automated threat response policies, and you can create on-demand reports for Security activity such as alerts generated and actions taken by policies. NI Manager

works in conjunction with ProCurve Virus Throttle technology to allow each port on a switch to act as a security sensor to detect and respond to virus attacks inside the network.

Network Immunity Manager provides an automated response to virus attacks by taking action on the source port of the attack. It can detect zero-day attacks (first attacks by a new virus or worm) and protects against threats from inside the network such as an employee bringing an infected laptop to work.

Used with the Policy Manager features of ProCurve Manager Plus (PCM+), NI Manager can provide automatic virus attack response and mitigation actions against the offenders, including:

- MAC lockout
- Port Rate limiting
- Disable Port
- E-mail notifications of alerts to IT and Network administrators
- Enable sFlow monitoring to more closely track suspect activity.
- Enable a Remote Port Mirror to forward suspect traffic to external IPS/UTM devices for further analysis.

In addition, NI Manager provides Offender tracking details, including IP address, MAC or DNS name, and username (via the Identity Driven Manager module)

NIM Architecture

The Network Immunity Manager software module is included in the PCM+ product CD. When installed, it utilizes the PCM server and Client interface for monitoring and configuration. If the IDM module is installed, NIM also interacts with the IDM server and client to get information on the user connected to ports where an attack is detected.

The following figure illustrates the NIM architecture and how it fits in with PCM.

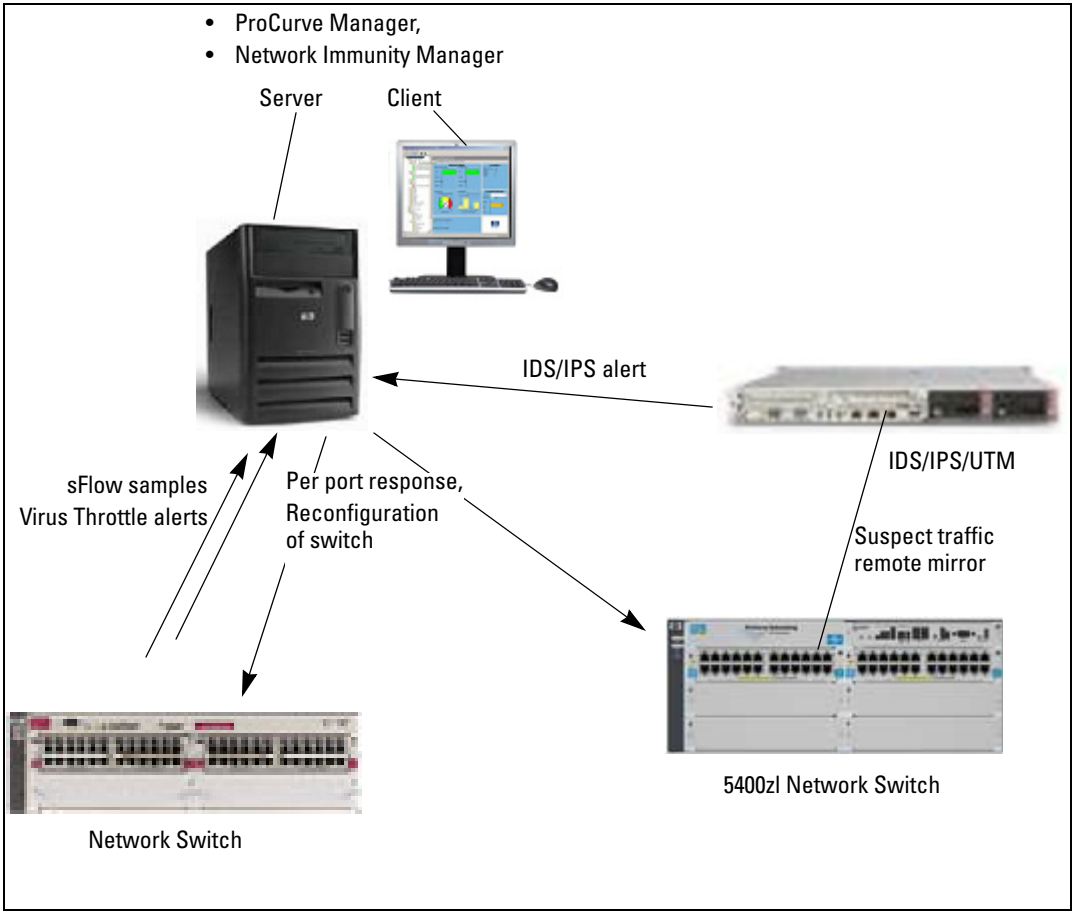

Figure 1-2. NIM Architecture

With Network Immunity Manager, you decide the level of protection to deploy. You can use the system to monitor selected network devices for specific NBAD (network behavior anomaly detection) parameters, then send alerts when suspicious activity occurs. Or you can configure the system to automatically respond to malicious behavior by limiting the bandwidth of an offender's port, moving the offender's port to a quarantine VLAN, executing a MAC lockout of the offender, or disabling the offender's port completely.

NI Manager relies on the traffic sampling engine in PCM to collect data on normal network usage and create profiles for traffic on various network segments or ports. When the traffic flow deviates from the established profile, NI Manager applies specially designed algorithms for false positive avoidance (FPA) to determine if a security alert should be generated. Alternately, NI Manager can send the suspect data information (using remote port mirror feature embedded in select ProCurve switches) to an external UTM/IDS/IPS device for additional analysis. Any traps that result from the external UTM/IDS/IPS device can be consumed by NI Manager as positive identification of an offender, and can be used by security policies as a source to trigger mitigation actions.

Note: If you have disabled the Traffic Manager functionality in PCM, no data will appear in the Security Activity displays. NI Manager uses the Traffic monitor feature to collect sFlow, XRMON, and other traffic statistics for analysis.

You can let NI Manager run in 'monitor' mode, in which it collects data and generate alerts without taking any action or other response. This is the default operation for NI Manager when installed. You can then adjust sensitivity levels and create policies to match your network's security requirements.

Terminology

Action Used with Policies, actions are defined rules or commands applied to the Policy target in response to an event-driven or scheduled alert.

Alert An alert notifies you when certain types of network device or application events occur that meet the alert's filter criteria.

AP Wireless Access Point. Examples include the ProCurve 420wl and AP530 devices.

Bandwidth Amount of network resources available. Generally used to define the amount of network resources a specific user can consume at any given time. Also referred to as rate-limiting. The amount of bandwidth available on a port may be constrained by NI Manager via port rate limiting, and the amount of network resources a specific user can consume should be the maximum throughput available to a specific port or user.

CIP Configurable Integration Platform. This is an API for use with PCM+ to add support for non-ProCurve network devices, and end-node device types such as RADIUS servers or shared printers, as well as launch external applications and decode traps from external devices, such as IDS/IPS/UTM devices.

Client An end-node device such as a management station, workstation, or mobile PC attempting to access the network. Clients are linked to the switch through a point-to-point LAN link, either wired or wireless.

CRF Connection-Rate-Filter, see Virus Throttling

Edge Device A network device (switch or wireless access point) that connects the user to the rest of the network. The edge devices can be engaged in the process of granting user access and assigning a user's access rights and restrictions.

FPA False Positive Avoidance. A proprietary algorithm developed by HP ProCurve to reduce incidence of false positives reported by the NBAD engine.

Firewall A firewall is network security device which is configured to permit, deny or proxy data connections set and configured by the organization's security policy. Firewalls can either be hardware and/or software based.

IDM Identity Driven Manager. This is a network management application that plugs into PCM to control and track which users are permitted to access network resources, primarily for use with RADIUS authentication.

IDS Intrusion Detection System. Devices to detect malicious attacks (traffic) on the network and alert the network security staff to take action.

IPS Intrusion Prevention System. Devices that can both detect and respond in some automated fashion to suspected network attacks. Also known as IDRS or Intrusion Detection and Response System.

MAC Media Access Control (MAC) address is a data link-layer address that is unique for each node on a LAN. MAC addresses consist of a 12-digit hexadecimal number and are designed to be unique and contain a code identifying the manufacturer of the network adapter or interface within the beginning of the address.

PMM ProCurve Mobility Manager, The plug-in module to PCM for monitoring and managing ProCurve wireless network devices.

NBAD Network Behavior Anomaly Detection, method used to detect unauthorized, malicious network traffic, such as worms, virus attacks, or other security threats.

PCM, PCM+ ProCurve Manager and ProCurve Manager Plus. PCM is a network management application for ProCurve network devices that provides utilities for monitoring ProCurve equipment. PCM+ provides advanced device management and configuration features, as well as real-time network traffic monitoring capability, and the ability to create Policies for applying and enforcing configuration and security of network devices. It provides the infrastructure that supports other task-specific modules such as Network Immunity Manager, Mobility Manager, and Identity Driven Manager.

Port Mirroring Port mirroring is used on a network switch to send a copy of all network packets seen on one switch port to a network monitoring connection on the same switch. Remote Port mirroring works in the same manner, the difference being that packets are sent to a port on another (remote) switch or device. This is commonly used with network appliances for monitoring and analysis of network traffic, such as an IDS/IPS/UTM devices.

Quarantine VLAN A port-based Virtual LAN configured on the switch that restricts access, by user's accessing that port, to the rest of the network. In this context, it is a type of Action that can be applied by a security policy in response to a suspected or known security threat.

Rate Limiting See Bandwidth

sFlow The sFlow standard (RFC 3176) describes a mechanism to capture traffic data in switched or routed networks. It uses a sampling technology to collect statistics from the device and is for this reason applicable to high speed networks (at gigabit speeds or higher).

UTM Unified Threat Management. Devices that combine the functionality of an IDS, IPS, and firewall.

VLAN A port-based Virtual LAN configured on the switch. When the client connection terminates, the port drops its membership in the VLAN.

Virus Throttling Also referred to as connection rate filtering, Virus Throttling (VT) can be used to block traffic from a host exhibiting a relatively high incidence of attempts to connect with other devices. You can set VT to block traffic permanently (until administrator re-configures to allow traffic), or to throttle (block) traffic temporarilyfrom the host for a calculated period of time, and then allow traffic to resume. If the undesired behavior persists, the cycle is repeated. You can also use VT in Notify Only mode so that the administrator is advised when VT events occur, without blocking traffic.

NI Manager Device Support Matrix

The matrix below lists ProCurve devices and the Network Immunity detection capabilities each device supports. The matrix also lists which mitigation actions Network Immunity can take on each specific device.

Table 1. NI Device Support and Mitigation Actions

	Switch/AP Detection Capabilities				**Mitigation actions NI can take on switch/AP**				
Device Types	sFlow/ XRMON	VT	Basic Local Mirror	Intel. Remote Mirror	Port shut-down	MAC Lockout	Rate Limit	VLAN	Reconfigure Basic Local Mirror
1600, 2400, 4000, 8000	x		x		x			x	
2524, 2512	x		x		x			x	x
2510			x		x	x		x	x
2626,2650,2608			x		x	x		x	x
4100, 6100			x		x			x	x
3400, 5300	x	x*	x		x	x	x	x	x
2800,2810	x		x		x	x		x	x
6400	x		x		x	x	x	x	x
9300,9400	x		x		x			x	
3500, 5400, 6200	x	x	x	x	x	x	x	x	x
8100			x		x			x	
4200	x		x		x	x		x	x
2900	x		x		x	x		x	x
420, and 520 AP									
530 AP						x			
5300 WESM	x					x			
5400 WESM	x					x			
7000 WAN Router					x				

* The 3400 does not currently support Virus Throttle (VT)

Registering Your NI Manager Software

The ProCurve Manager installation CD includes a fully operable version of the PCM application, and a 30 day trial version of the PCM+, Mobility Manager (MM), Network Immunity Manager (NI), and the Identity Driven Manager (IDM) applications. Until you have registered your PCM+ and NI applications, an Expiring License warning will be displayed each time you log in, similar to the following.

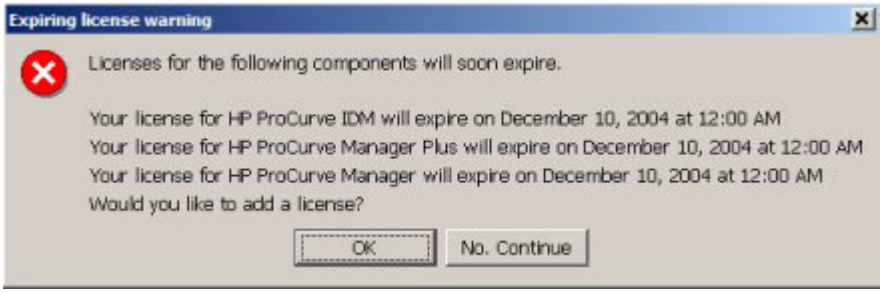

Figure 1-3. ProCurve Expiring License warning dialogue

Click **No, Continue** to close the dialogue and just start the program.

Click **OK** to launch the Licensing administration screen.

NOTE: You must first purchase a copy of ProCurve Network Immunity Manager from your networking reseller to get the Registration ID. *You do not need to re-install the software from the purchased CD, but you need the Registration ID from that CD to complete the registration process.*

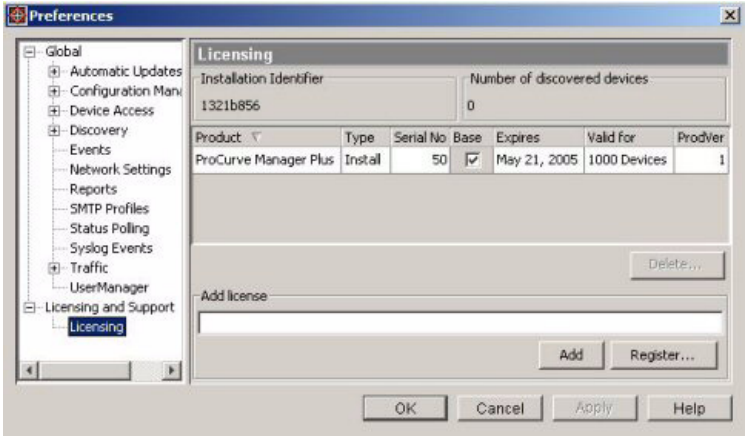

Figure 2. ProCurve License Administration dialogue

You can also get to this screen from the Preferences window which can be accessed from the PCM Tools menu or by clicking on the Preferences icon in the tool bar.

Note:

When registering add-on modules for PCM+, including NIM, PMM, and IDM: You must use the same user account to register the add-on modules as you used when registering PCM+, and you will add the modules the existing PCM+ registration.

To register the PCM+, and NI manager software:

1. Contact your HP Sales Representative or HP Reseller to purchase the PCM+ and NI Manager software. You will receive a Registration ID for the purchased software—either on the Software CD case, or a separate registration card sent with the purchase information.
2. Go to the Licensing window in PCM [Preferences→Licensing and Support→ Licensing]. Write down the Installation Identifier for the software as it appears in the upper left corner of the window. You can also leave this window open and use the "copy and paste" functions to enter the Install ID in the My ProCurve software registration window.
3. Click the Register button to go to the My ProCurve registration web site.
4. If this is an upgrade, log in with your My ProCurve ID and password.
   If you are a new user, click the "Register Here" button, and then enter the required information to create a user account, including user name, password, company name, and E-mail address.

5. In the Registration window:
   a. Select the Software tab, then choose the ProCurve Network Management Software link.
   b. Select the product to register from the Product Type pull-down menu.
   c. Enter the Registration ID, found on the back of the software CD case, or on the registration card you received when you purchased the software.
   d. Enter the Installation Identifier (from the Licensing window in PCM). Be careful not to include any extra spaces as this causes a failure in the registration.
   e. Click the Generate License button.

   The window is refreshed and the registration information, including your License key is displayed. The license key is also sent to you via e-mail.

   f. To get the license key for the next software package, click Generate Another License and repeat the process in step 5, above.
6. When you receive the License key, go back to the Licensing window in PCM.

   Enter the License key number in the Add license field, then click Add.

   To avoid data entry errors, you can copy and paste the number from the e-mail or My ProCurve (My Software) Web page. Be careful not to include any extra spaces as an extra space will cause a failure in the licensing process.

NOTE:

You must first purchase a copy of ProCurve Manager Plus and Network Immunity Manager to get the Registration ID. *You do not need to re-install the software from the purchased CD, but you need the Registration ID to complete the registration process.*

Learning to Use NI Manager

The following information is available for learning to use ProCurve Network Immunity Manager (NI):

- This User's Guide—helps you become familiar with using the application tools for access control management.
- Online help information—provides information through Help buttons in the application GUI that provide context-sensitive help, and a table of contents with hypertext links to additional procedures and reference information.
- *ProCurve Manager*, *Getting Started Guide*—provides details on installing the PCM, PCM+, NI, MM, and IDM applications and licensing, and an start-up tips for PCM and PCM+.
- *ProCurve Manager Network Administrator's Guide* —provides details on using the PCM and PCM+ features for monitoring and managing ProCurve network devices, monitoring network traffic, configure VLANs, using Policies to enforce device configurations and security, and using the Configurable Integration Platform to add support for other device types and applications within PCM. Also includes instructions for using Mobility Manager to mange and configure ProCurve wireless network devices, including configuring Radio options and WLANs.
- For additional information on managing and configuring your ProCurve network devices, refer to the documentation that came with your switch.

Getting ProCurve Documentation From the Web

1. Go to the Procurve Web at http://www.procurve.com.
2. Click on **Technical Support**.
3. Click on **Product manuals**.
4. Click on the product for which you want to view or download a manual.

ProCurve Support

Product support is available on the Web at: www.procurve.com
Click on **Technical Support**. The information available at this site includes:

- Product Manuals
- Software updates
- Frequently asked questions (FAQs)
- Links to Additional Support information.

You can also call your HP Authorized Dealer or the nearest HP Sales and Support Office, or contact the ProCurve Elite Partner nearest you for information on ProCurve Access Control Security solutions. You can find ProCurve Elite partners on the web at:
http://hp.via.infonow.net/locator/us_partner/index.jsp

2

Configuring Network Immunity Manager

Chapter Contents

Security Capabilities Overview

Network behavior anomaly detection (NBAD) is the continuous monitoring of a network for unusual events or trends. The NBAD engine employed by NI Manager tracks critical network characteristics in real time and generates an alert if a strange event or trend is detected.

Analysis is performed on traffic metrics from ProCurve switches to detect internal threats. NBAD builds a picture of network activity over time and over the topology via traffic sampling, and performs analyses to isolate suspicious traffic for unknown or zero-day attacks.

NIM also accepts attack alerts from Virus Throttle™ technology embedded in select ProCurve switches, and accepts alerts from select external (non-ProCurve) IDS/IPS/UTM security devices.

The Network Immunity Manager has been tested to detect the following protocol anomalies and attack types:

- Port scanning techniques:
  - Xmas Tree Scan – Sends a TCP frame to a remote device with the URG, PUSH, and FIN flags set
  - NULL Scan – Turns off all flags, creating a lack of TCP flags
  - FIN Scan - The FIN scan's "stealth" frames are unusual because they are sent to a device without first going through the normal TCP handshaking
- Denial of Service:
  - UDP Bomb - An illegal sent User Datagram Protocol (UDP) packet
  - Land Attack – An attack involving IP packets where the source and destination address are set to address the same device
  - Ping of Death – Sends a malformed or otherwise malicious ping to a computer
- Reconnaissance before an attack, including the following tools:
  - Nessus
  - NMAP
  - Port Scanners and Ping tools
- Network-based attacks
  - DNS Tunneling
  - Unauthorized Network Mapping

  - IP Spoofing
  - Various Worm Propagation techniques
- Anomalous Packet Size, designed to inform NI to:
  - Sample suspicious traffic
  - Detect some covert channels
- Mis-Configured devices, tested to detect:
  - Duplicate IP's
  - Rogue Routers
  - Rogue Proxies

Setting Security Preferences

The Preferences feature in PCM provides the tools to control and adjust the NBAD detection sensitivity and manage the NBAD options to fit your particular environment.

The Security Sensitivity levels will be different for any given network environment. There is no one setting applicable to all networks or sites. You will need to tune the security sensitivity parameters depending on the traffic in your network. You should perform traffic assessment and security planning before deploying security products like NI Manager.

Setting the wrong sensitivity level could generate false positives or false negatives. ProCurve recommends starting with the default sensitivity levels with log only actions, then analyze the alerts to build confidence on the alert settings. Tune the sensitivity levels to the optimum level for their environment and when enough confidence is reached, then enable analysis (sFlow or port mirror) actions and mitigation actions via the Security Policies described in Chapter 4.

Note that security is a process that requires regular tuning. In general, static security settings may fail against new breeds of viruses and worms. Sensitivity settings in NI Manager require regular tuning to match the ever changing nature of attacks by viruses or worms that can occur.

To configure Security Monitoring Preferences for NIM:

1. Go the Preferences menu and select the Security Monitoring option {Preferences -> Security Monitoring] to display the Security preferences configuration window.

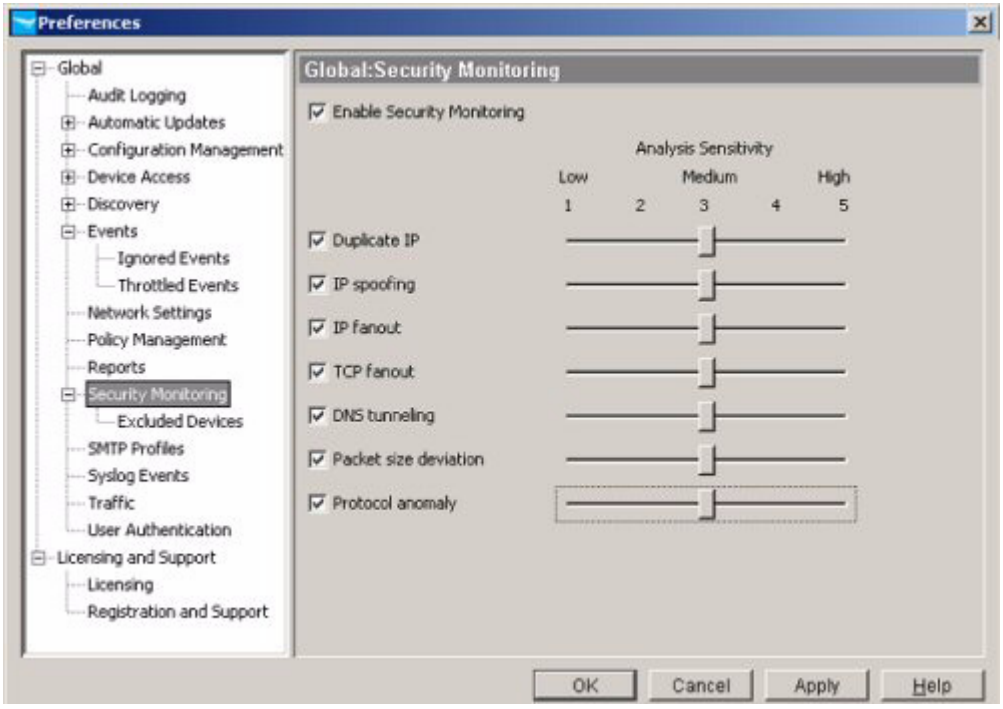

Figure 2-1. Global Preferences, Security Monitoring configuration window.

1. The Enable Security Monitoring option is selected by default. Click the check box to disable Security Monitoring. When Security Monitoring is disabled, the remaining security monitoring options are also disabled.

2. Select the security monitoring options and set the Analysis Sensitivity for each of the selected monitoring conditions.

   Click the slider on the bar of an enabled security monitoring option, and drag it to the right to increase sensitivity or drag it to the left to decrease sensitivity.

   - The default sensitivity for all security monitoring is Medium (value 2 or 3).
   - A High sensitivity (5) reports more security events by requiring a lower confidence level in the analysis before reporting.

- A Low sensitivity(1) reports only security events that result from analysis with a high confidence level.

See “NI Manager Sensitivity Setting Definitions” on page 2-8 for detailed descriptions of sensitivity settings.

You can also enable or disable each individual security monitoring condition. A check indicates the option is enabled and vice versa. If a condition is disabled (not selected), its security related condition may still be analyzed and monitored by the security engine but not reported.

A high sensitivity setting analyzes and reports on more general security events, while low sensitivity analyzes and reports only security events with a "higher confidence" level. You can start by setting higher sensitivity levels to build confidence in the "anomaly" detection, and then later adjust the sensitivity lower to avoid "false positives."

Excluding Devices from Security Monitoring

You can use the Global Preferences for Security Monitoring, Excluded Devices window to configure specific network devices that you do not want monitored by NIM at all. You can also configure a device to monitor for only one or two specific anomaly types.

To review or modify the list of devices excluded from security monitoring

1. Open the Preferences Menu and expand the Security Monitoring node, then select the Excluded Devices node.
   [Preferences->Security Monitoring->Excluded Devices]

   This displays the list of excluded devices.

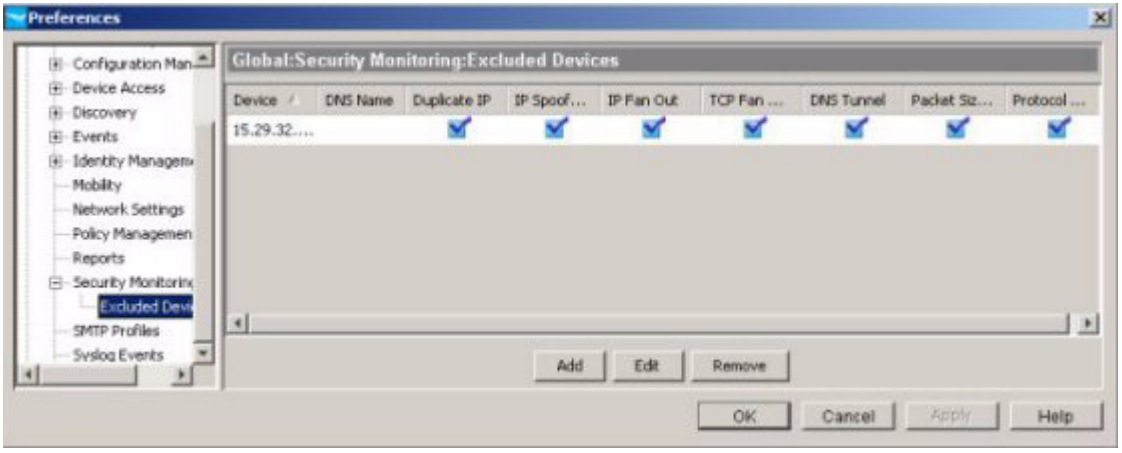

Figure 2-2. Global Preferences: Security Monitoring, Excluded Devices window.

The list displays the Device IP and DNS name. Check marks indicate the security monitoring conditions the device is excluded from.

Note:

The PCM Discovery engine causes NBAD events to be logged. For this reason the PCM Server is autom atically excluded. Routers that PCM can communicate with via SNMP are also auto-excluded. To override the auto-exclude, you must uncheck all the anomalies. If you delete the device, it will be re-excluded at the next discovery cycle.

By default NIM will not take action against inter-switch device ports that are identified as offenders

To add a device to be Excluded from Monitoring:

1. Open the Preferences Menu and expand the Security Monitoring node, then select the Excluded Devices node.
2. Click the **Add** button to launch the Add Excluded Device dialog.

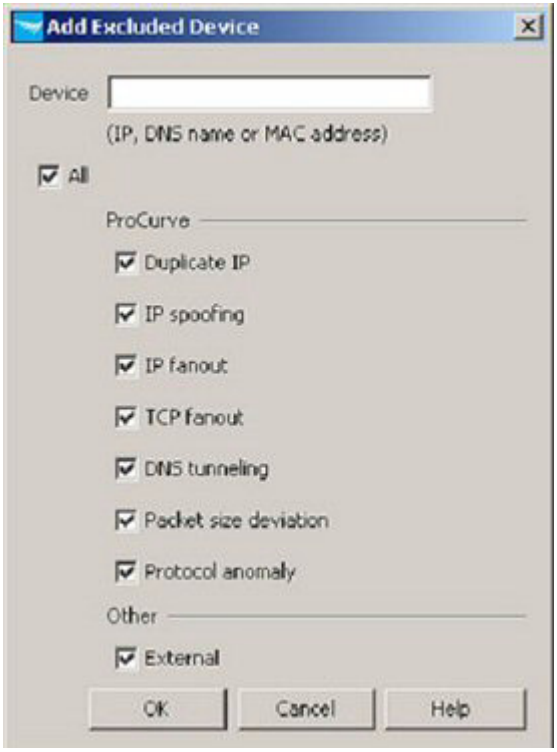

Figure 2-3. Preferences: Security Monitoring, Add Excluded devices window

3. In the **Device** field, type in the identifier for the device to be excluded, any one of the following is supported: IP address, DNS name, or MAC address of the device
4. Click to select or deselect the Security Monitoring conditions to be excluded. **All** conditions are selected by default.
   A check mark in the box indicates the option is selected, and the device will be excluded from monitoring of that condition.
   Other; External will exclude the device from any External Alert types you define in the Policy Manager for Security Policies (see Chapter 4).
5. Click **OK** to save the Excluded Device information and close the window.

The device you just added now appears in the list for Security Monitoring: Excluded Devices.

NOTE: Any Windows system that is multi-homed can generate false IP Spoofing and/or Duplicate IP alerts because of the way that Windows initializes the default routes for each NIC in a multi-homed system. You should be aware of this potential for false positives and should exclude such multi-homed systems from NBAD analysis for IP Spoofing and Duplicate IP addresses as needed.

To remove a device from the Excluded Devices list:

1. Open the Preferences Menu and expand the Security Monitoring node, then select the Excluded Devices node.
   [Preferences->Security Monitoring->Excluded Devices]
2. Click to select the device you want to remove in the list of devices.
3. Click the **Remove** button.
   The list of excluded devices is refreshed and the selected device is removed.

To change security monitoring options on a device in the Excluded Devices list:

1. Open the Preferences Menu and expand the Security Monitoring node, then select the Excluded Devices node.
   [Preferences->Security Monitoring->Excluded Devices]
2. Click to select (highlight) the device you want to edit in the list of devices.
3. Click **Edit** to launch the Edit Excluded Device dialog.
4. Click to enable or disable the security monitoring conditions in the same manner as described for “Excluding Devices from Security Monitoring” on page 2-5.
5. Click the **OK** button to save the changes and close the dialog.

NI Manager Sensitivity Setting Definitions

NBAD type	Data Points	Violation	Sensitivity	Time Window	Event Size
Duplicate IP	-MAC Address -IP Address -Time Window	One IP appearing from more than one MAC within the specified time window	1 2 3 4 5	1 min. 15 min. 1 hour 3 hours 24 hours	Triggers when two or more such samples appear within the time window.
Spoofed IP	-MAC Address -IP Address -Time Window	One MAC appearing with more than one IP within the specified time window.	1 2 3 4 5	1 min. 15 min. 1 hour 3 hours 24 hours	Triggers when two or more such samples appear within the time window.
IP Fan-Out	-Source IP Address -Destination IP Address	One IP communication with *N* other IPs and/or one IP communicating with a statistically unusual number of other IPs in the specified time window.	1 2 3 4 5	1 min. 15 min. 10 min 15 min 30 min	256 IPs 128 96 32 3
TCP/UDP Fan-Out	-Source IP Address -Destination TCP/UDP Ports (per destination IP address)	One source IP communicating with *N* other ports on a given destination IP and/or one source IP communicating with a statistically unusual number of destination ports on a given destination IP in the specified time window.	1 2 3 4 5	1 min. 15 min. 10 min 15 min 30 min	256 Ports per IP 128 10 5 2
Ave. Packet Size Deviation	-Host IP Address -Average Packet Payload Size in Bytes	Occurs when NI Manager detects a statistically unusual change in the average size of sent and/or received packets.	1 2 3 4 5	N/A	6 PSD* from mean 5 *(packet size 4 deviation) 3 2
Protocol Anomaly	-Source IP Address -Packet Header Content	Occurs when the host sends traffic containing unusual properties that would not normally be expected to occur on the network.	1 2 3 4 5	N/A	Any packet matching the approximately 30 anomalous behaviors defined for NBAD immediately creates an event.

Operating Notes for Security Configuration

General Notes: When reporting suspect activity, if the same anomaly is detected by 10 switches, is it counted as one event.

In general, security events are counted by type and identified by offender (attacker).

sFlow Sampling: PCM Traffic sampling (sFlow or XRMON) must be enabled so that NIM can receive traffic samples across the network for anomaly detection. The NBAD engine in NIM builds a picture of network activity over time and over the topology via traffic sampling, and performs analyses to isolate suspicious traffic for unknown or zero-day attacks.

PCM can support up to 500 sFlow ports simultaneously. In the default configuration, PCM can automatically detect inter-switch connection ports and prioritizes those ports for sFlow sampling. You may change the sFlow port selection if there is a need to change PCM's default selection, although every port manually configured for full-time sFlow consumes one of the 500 ports that are available to PCM and NIM for automatic sFlow management.

NI Manager has a default policy in place that invokes sFlow sampling on a port in response to average packet size deviation anomalies (but not other anomaly types) that are observed on that port using statistics polling. NIM-specified ports are "rolled back" after a period so that they do not permanently consume the slots that they occupy among the 500 sampling ports. A port on which NIM-specified sampling is rolled back may still sample when determined to need it by the traffic-selected algorithm.

ProCurve recommends that you employ the PCM and NIM automatic selection except for locking down sFlow on critical edge ports. If you decide to select sFlow ports manually then you should treat core switch ports as a primary priority, distribution switch ports as a secondary priority, and non-critical edge ports as the final priority; you should strive to leave a significant portion of the 500 sampling port budget available for PCM and NIM to manage automatically.

For additional information on configuring PCM Traffic monitoring, see the *ProCurve Manager Network Administrator's Guide*, Chapter 9 "Monitoring Network Traffic".

Virus Throttling: Using the Virus Throttle (VT) feature in select ProCurve switches* provides a first layer defense against anomalies or unknown attacks. You can set VT to notify only so that NIM gets informed of VT alerts,

without impacting overall network performance. When VT alerts are forwarded to NI Manager, it gives the administrator an overview of all VT alerts across the network.

If you want to use the Virus Throttling feature with NI Manager, ProCurve recommends you enable the Virus Throttle (connection rate filter) option on all switches that support it. Based on your network environment you should decide if this is feasible, or if you want to use VT only on critical ports such as uplink or ports connecting to business critical resources.

NI manager handles VT alerts and NBAD alerts separately, and NI can take action on either or both alert types. You can decide to have VT or NIM both take action in the event of a VT alert. For example, if you have VT enabled on a switch, that switch can block an attack using VT, and if the alert is also sent to NIM, it can take a notify action. If VT is only available in switches at the network core, you would need to use a NIM Security Policy to block an offender at the network edge using other available Actions, such as rate-limiting, MAC lockout, Port disable, or VLAN configuration.

Because IP fanout does not work across subnets (by default routers are excluded), you may want to enable Virus Throttle to take action to block the user, and configure a SecurityPolicy to display a message or send an e-mail of the VT activity to the network administror.

Using external IDS devices: If an IDS is installed in the network, have at least one remote mirroring session set up from the edge switches (source switch) that can mirror to the remote destination switch(es) where the IDS connects. This allows NI Manager to reconfigure the mirror ports (via the Port Mirror action in Policy Manager) to copy suspicious traffic to the IDS for deeper analysis. The IDS can then generate alerts of high confidence back to NI Manager. This configuration applies to all switches where you want NI Manager to be able to apply remote port mirroring for deeper analysis of suspect traffic.

Make sure jumbo frames are enabled along the entire path from the source switch to the destination switch, so the IDS gets the actual packets without loss of data.

---

* The virus throttle feature is currently available on ProCurve Series 5300xl, 3500yl, 5400zl, and 6200yl switches.

3

Monitoring Security Activity

Chapter Contents

Using the Security Activity Tab

The Security Activity tab provides a summary of security alerts and resulting actions in a variety of views. These views are typically used to:

- Identify specific groups where alerts are concentrated, groups triggering alerts of a given severity level, and how alerts are distributed across the severity levels.
- Determine the types of actions that have been attempted in response to an alert and the group, device, or port where the action was attempted.
- Pinpoint the top offenders (user, IP address/DNS name, or MAC address suspected of launching an attack over the network).
- Identify alerts that did not result in an action (by filtering alerts). This helps you confirm that policies are configured correctly and that the appropriate action is taken when a specific alert type occurs.

WARNING: NI Manager uses the Traffic monitor feature to collect sFlow, XRMON, and other traffic statistics for security analysis. If you have disabled the Traffic Manager functionality in PCM, no data will appear in the Security Activity displays except Virus Throttle events and alerts received from external UTM/IPS/IDS devices, if used.

The Security Activity tab contains three sub-tabs:

- Alerts: Displays the security alerts that pertain to the selected navigation tree context (group or device) as determined by the device group of the offenders that triggered the alert.
- Actions: Displays the actions that have been executed within the selected navigation tree context (group or device) as a result of a security alert.
- Offenders: Displays the security alerts by offender (based on port or device) within the selected group or device.

Each sub-tab contains security information for the selected device or device group. The displayed data is updated every time PCM receives a new alert and/or action combination.

The Filters panel at the top of the display lets you configure filters that govern the actions and alerts included in the displayed data set. You can also generate a custom report for Security Alerts, with filter criteria similar to the tab display.

Note: The Security Activity tab is not shown for groups with no devices. Tables and charts in the Security Activity tab only show alert types supported by a device.

To display the Security Activity tab:

1. In the navigation tree, select the Interconnect Devices node, a group node, or a device node.
2. Click the Security Activity tab in the (Interconnect) Devices window.

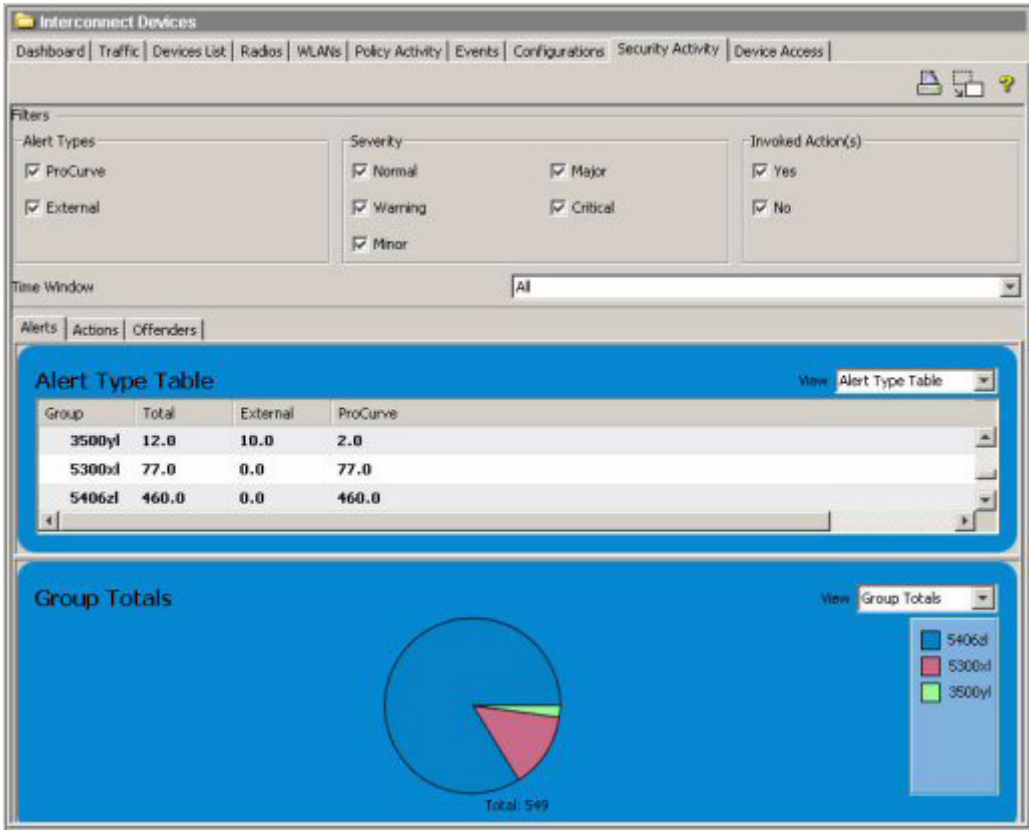

Figure 3-1. Security Activity Tab, default display

To display additional security activity information:

- To display the security information for devices in the selected group or ports in the selected device, double-click the group field or device in the table or chart.

- You can single-click on a table cell (except for the group column) or click a bar in the chart to display details for the selected field. You can click on a slice in the pie chart to display details for that segment.
- Double-click a cell in the table (except in the Group column) to display the associated Policy History for the selected field.

Filtering The Security Activity Display

The default display includes all possible data points. You can filter the display to show only the data you need. Changing the filters affects the selected group or device only. To change the filters for all groups and devices, open the Security Activity tab in the Interconnect Devices window.

The selected filters will remain in effect until you change them. That is, whenever the Security Activity tab for the selected device or group is accessed, the last filters configured will be automatically selected.

To configure a filter:

In the Filters panel of the Security Alert tab, check to select or deselect the data you want to include in the display. A check mark indicates the filter option is selected.

The filter options are context-sensitive. Only the filter options that are applicable to the data in the tab are displayed. Filter options are shared between the Alerts and Offenders subtabs; that is, settings made in the Alerts tab appear on and influence data shown in the Offenders tab, and vice-versa.

Time Window

The Time Window filter is available for all security activity tabs. Use the pull-down menu to select the time period for which you want data, one of:

All - Display entire activity history

Last Hour- Display activity for the previous hour

Last Day - Display activity for the previous 24 hours.

Last Week - Display activity for the previous 7 days.

Last Month - Display activity history for the past calendar month.

Possible filters for the Alerts and Offenders sub-tabs are:

Alert Types

- ProCurve: Alerts based on definitions in Policy Manager-Security: ProCurve Alerts. The security alerts for a group, device, or port that are generated by NI Manager (via NBAD), and ProCurve devices (via VT).
- External: Alerts based on definitions in Policy Manager-Security: External Alerts, which apply to events from external IDS, IPS, UTM devices, or other network devices.

Severity

- Normal: Alerts classified as Normal (Informational) severity
- Warning: Alerts classified as Warning severity
- Minor: Alerts classified as Minor severity
- Major: Alerts classified as Major severity
- Critical: Alerts classified as Critical severity

Invoked Action

- Yes: Action defined in Security Policy was invoked.
- No: Action was not invoked, but event message is logged.

Possible filters for the Actions sub-tab are:

Action Types

- sFlow: Alerts that resulted in the invocation of sFlow sampling on a device port
- Port Rate Limit: Alerts that resulted in changes to the port rate limit settings of an offending port.
- Port Mirror: Alerts that resulted in suspect network traffic being mirrored to a monitor port.
- Quarantine VLAN: Alerts that resulted in assigning an offending port to a specified VLAN.
- Disable Port: Alerts that resulted in disabling the offending port.
- MAC Lockout: Alerts that resulted in locking out an offender's MAC address (blocking all traffic to and from the MAC address) on a specific device
- Other: Alerts that resulted in other actions, typically generated by user scripts.

Completion Status

- Actions(s) Complete: Alerts with an associated action that has been completed, including aborted and interrupted actions
- Actions(s) Incomplete: Alerts with an associated action that has not been completed, usually because of the time set in the associated policy

Creating Security Activity Reports

You can generate a Security Alert report from the Security Activity tab.

1. Click the Print button on the Security Activity tab to launch the Report Wizard.

Figure 3-2. Security Alert Report Wizard, report filters

2. Select the filter criteria for data to be included in the report. These options are the same as for filtering the Security Activity tab views (see page 3-4). In the first window of the wizard you can select the following filters:
   - Time Window (hour, day, week, month, all)
   - Action Types
   - Alert Types
   - Severity

Click the check box to select or deselect the filter option. A check mark indicates the filter option is selected.

3. Click **Next** to continue the Report Filter selection.

Figure 3-3. Security Alert Report Wizard, report filters continued

In the second window you can select the following filters:

- Completion Status
- Invoked Action(s)

4. Click **Finish** to generate a custom report containing the details corresponding to the selected filters.

5. The report will display on the Windows desktop.

- Click the print button in the report display to print the report using the Windows "Print" function.

- Click the save button in the report display to save the report to a file using the Windows "Save" function.

Reviewing the Alerts Sub-tab

The Alerts sub-tab on the Security Activity tab (figure 3-1 on page 3-3) shows the following information panels for security alerts received from the selected group or device during the selected time period:

- Table or Bar Chart Summary for Security Alerts
  You can select the display formats using the View pull-down menu to categorize data by Alert Type or Severity, presented in table format or bar chart.
- Totals
  You can set the Totals display using the View pull-down menu to categorize data by Group, Alert Type, or Severity. This lets you easily find the information you desire.

  For example, on the Alerts sub-tab for Interconnect Devices, select an Alert Type view to see specific groups where the alerts are concentrated, or select a Severity view to find how alerts are distributed across the event severity levels and which groups of devices or ports are triggering the alerts of a given severity level.

To display additional information:

1. To display the Policy History summarized for a value in the table, double-click the value.
2. To display the security activity for subgroups or devices in the selected group or ports in the selected device, click on the group or device in the table or chart.
3. To display additional information for a section of a pie chart, mouse over the section. You can single-click a section of the Totals pie chart to display the Policy History summarized for the selected segment.

Alerts Table or Bar Chart Summary

Click the View drop-down arrow in the top panel of the Alerts sub-tab and select the view you want to display.

The bar chart scrolls horizontally and the table scrolls vertically to accommodate multiple groups, devices, or ports.

You can double-click on any Group/Device in the Alert Type table or bar chart to display the Alerts summary for the selected group.

Double-click on any number in the Alert type table, or bar chart, to view the Policy History summarized by that number.

The Alert Type Display:

Selecting Table (Alert Type) or Bar Chart (Alert Type) displays the following information, regardless of whether you select a table or bar chart:

Information	Description
Group/Device/Port	Depending on the node selected, the name of the groups, devices, or ports.
Total	Total number of alerts for a group, device, or port that meet the filter and time criteria
External	Total number of alerts for a group, device, or port that originated from an external source (such as external IDS/IPS/UTM appliances) and meet the filter and time criteria
ProCurve	Total number of security alerts for a group, device, or port that that are generated internally by NI Manager (via NBAD), and from ProCurve devices (such as VT alerts) and that meet the filter and time criteria.

The Severity Display:

Selecting Table (Severity) or Bar Chart (Severity) displays the following information, regardless of whether you select a table or bar chart:

Information	Description
Group/Device/Port	Depending on the node selected, name of the group, device, or port.
Total	Total number of alerts for a group, device, or port that meet the filter and time criteria
Normal	Total number of Normal (Informational) alerts for the selected group, device, or port that meet the filter and time criteria
Warning	Total number of Warning alerts for a group, device, or port that meet the filter and time criteria
Minor	Total number of Minor alerts for a group, device, or port that meet the filter and time criteria
Major	Total number of Major alerts for a group, device, or port that meet the filter and time criteria
Critical	Total number of Critical alerts for a group, device, or port that meet the filter and time criteria

Alert Totals Display

Depending on the view you select from the View drop-down list, the bottom panel provides a pie chart showing the total alerts for each alert type, alert severity, or selected group or device.

Clicking a wedge in the pie chart displays detailed information about the alerts that pertain to your selection.

- **Alert Type:**
  Selecting Total by Alert Type shows the total number of alerts for each alert type that meet the filter and time criteria.
- **Group**:
  The Totals by Group option shows the total number of alerts for each group, device, or port (depending on the node selected) that meet the filter and time criteria.

  This chart shows totals for the top 10 groups only, which is especially useful when many groups are displayed in the table or bar chart.
- Severity:
  The Totals by Severity option shows the total number of alerts for each alert severity that meet the filter and time criteria.

You can single-click or hover on a value in the Security Activity: Alerts sub-tab Totals to display a Details pop-up, which provides additional information about the selected segment, similar to the following figure.

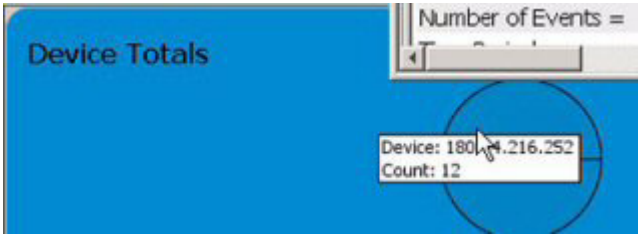

Figure 3-4. Single click (mouse-over) of Bar Chart

Alert Details Display

You can single-click a value in the Alerts table or bar chart to display additional details for the Alert, as shown in the following figure.

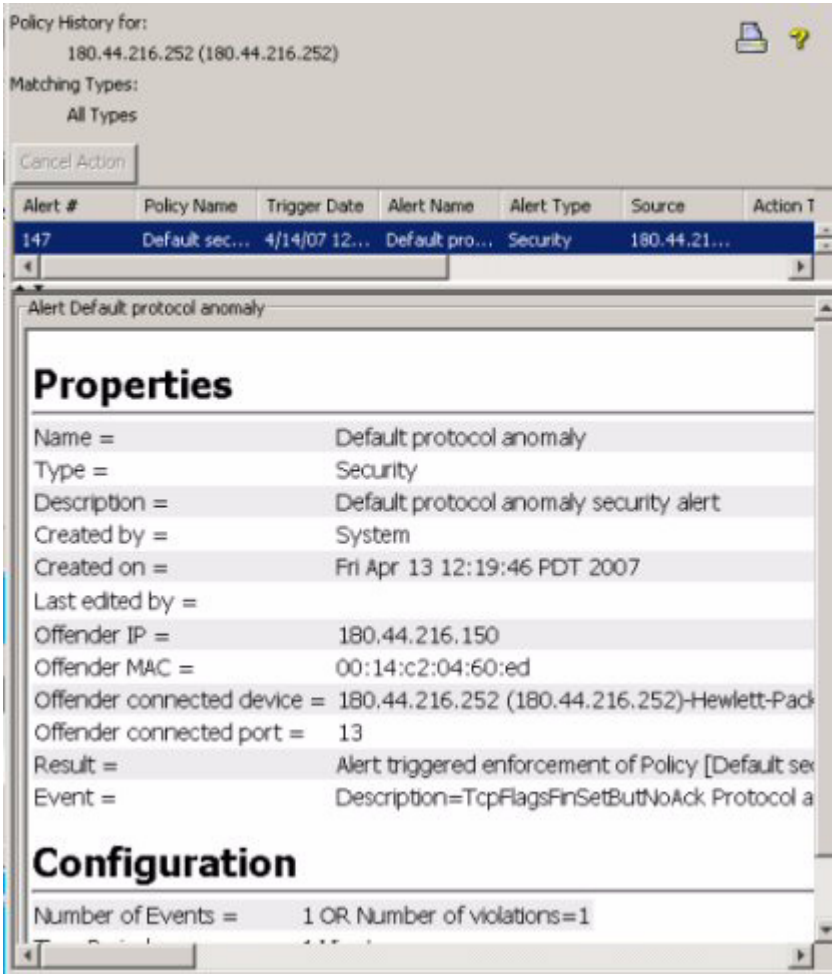

Figure 3-5. Security Alerts: Alert Details display

The top panel displays the alert history for the selected device or group, including the related Policy information and actions taken as a result of the alerts.

The bottom portion of the display lists the Alert Properties, including information on the offender that caused the alert, and the alert Configuration details.

Reviewing the Actions Sub-tab

The Actions sub-tab on the Security Activity tab shows information about actions that were taken in response to security alerts during the selected time. Information is displayed in separate panels for:

- Table or Bar Chart Summary
- Totals

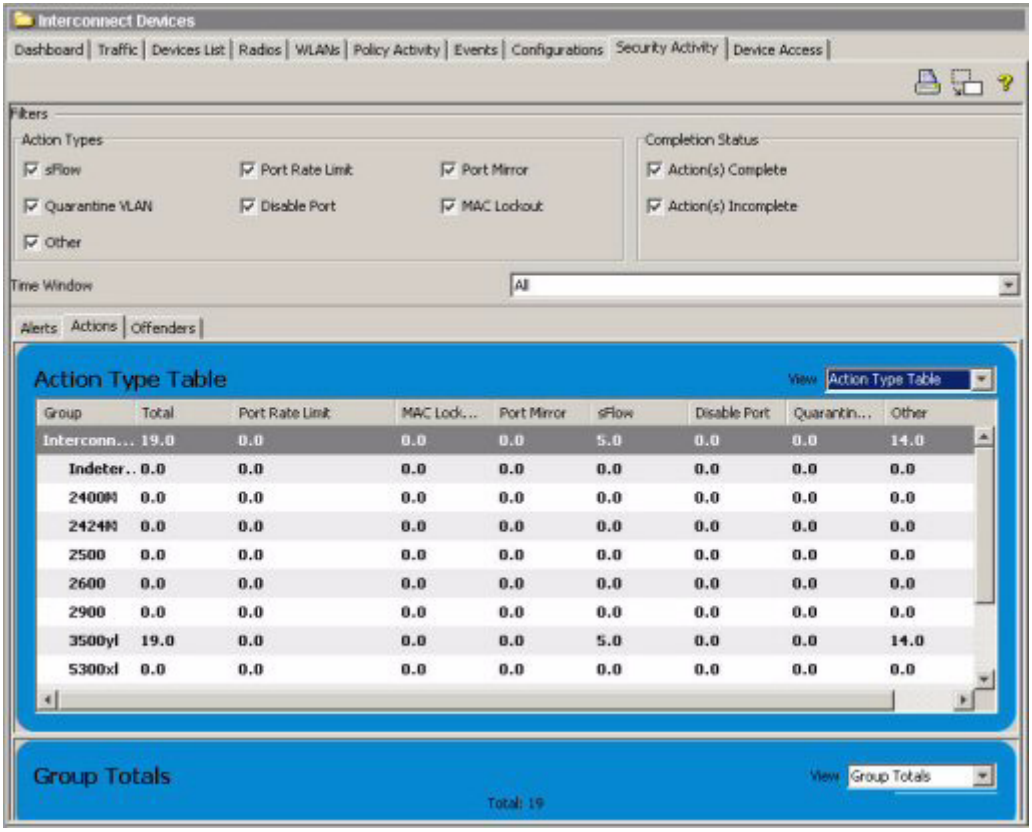

Figure 3-6. Security Activity: Actions sub-tab display

The Actions Table or Bar Chart Summary

The Table or Bar Chart Summary panel offers views that categorize information by group, or device, and action type, which lets you easily determine which actions have been applied and where. In the Actions sub-tab for the Interconnect Devices node, either of the available views will indicate the specific groups of devices where actions were concentrated.

Use the View pull-down menu in the top panel of the Actions sub-tab to select the view you want to display. The bar chart scrolls horizontally and the table scrolls vertically to accommodate multiple groups, devices, or ports.

Selecting Table (Action Type) or Bar Chart (Action Type) displays the following information:

Information	Description
Group/Device/Port	Depending on the node selected, name of the groups, devices, or ports
Total	Total number of actions for a group, device, or port that meets the filter and time criteria
Disable Port	Total number of actions invoked to disable an offending port for a group, device, or port that meets the filter and time criteria
Port Rate Limit	Total number of actions invoked to apply port rate limiting to an offending port for a group, device, or port that meets the filter and time criteria
Quarantine VLAN	Total number of actions invoked to place an offending port on a specified VLAN for a group, device, or port that meets the filter and time criteria
Port Mirror	Total number of actions invoked to mirror suspect network traffic from an offending port to a monitor port, where a third-party appliance can be used for more definitive analysis of the traffic for a group, device, or port that meets the filter and time criteria
sFlow	Total number of actions invoked to enable sFlow sampling on a device port for better visibility to suspect traffic for a group, device, or port that meets the filter and time criteria
MAC Lockout	Total number of actions invoked to lock out an offender's MAC address (blocking all traffic to and from the MAC address) for a group, device, or port that meets the filter and time criteria
Other	Total number of actions invoked to carry out other activities, typically via user scripts for a group, device, or port that meets the filter and time criteria

Reviewing Action Totals

Depending on the view you select from the View drop-down list, the bottom panel provides a pie chart showing the total actions attempted for each action type or selected group, or device.

Clicking a wedge in the pie chart displays detailed information about the actions that pertain to your selection.

View by Action Type:
Selecting Total by Action Type shows the total number of actions attempted for each action type that meets the filter and time criteria.

Group:
The Totals by Group option shows the total number of actions attempted for each group (depending on the node selected) that meets the filter and time criteria.

This chart shows totals for the top 10 groups only, which is especially useful when many groups are displayed in the table or bar chart.

Reviewing Policy History

Double-clicking a value in the Actions sub-tab of the Security Activity tab displays the Policy History window, which lists detailed information about actions for the selected group, device, or port.

Information	Description
Groups	Description of the custom groups affected by the action
Action Types	Action taken in response to the alert
Time	Time and date the action was taken
Group	Name of the custom group affected by the action
Details	Detailed description of the event and action taken

Reviewing the Offenders Sub-tab

The Offenders sub-tab shows information on security alerts by offenders (by user, end node, or MAC address suspected of launching an attack over the network) in the selected group. If used in conjunction with IDM, information available from IDM can also be displayed, including user name.

The Offenders sub-tab shows the following panels of information about offenders during the selected time:

- Table or Bar Chart Summary
- Totals

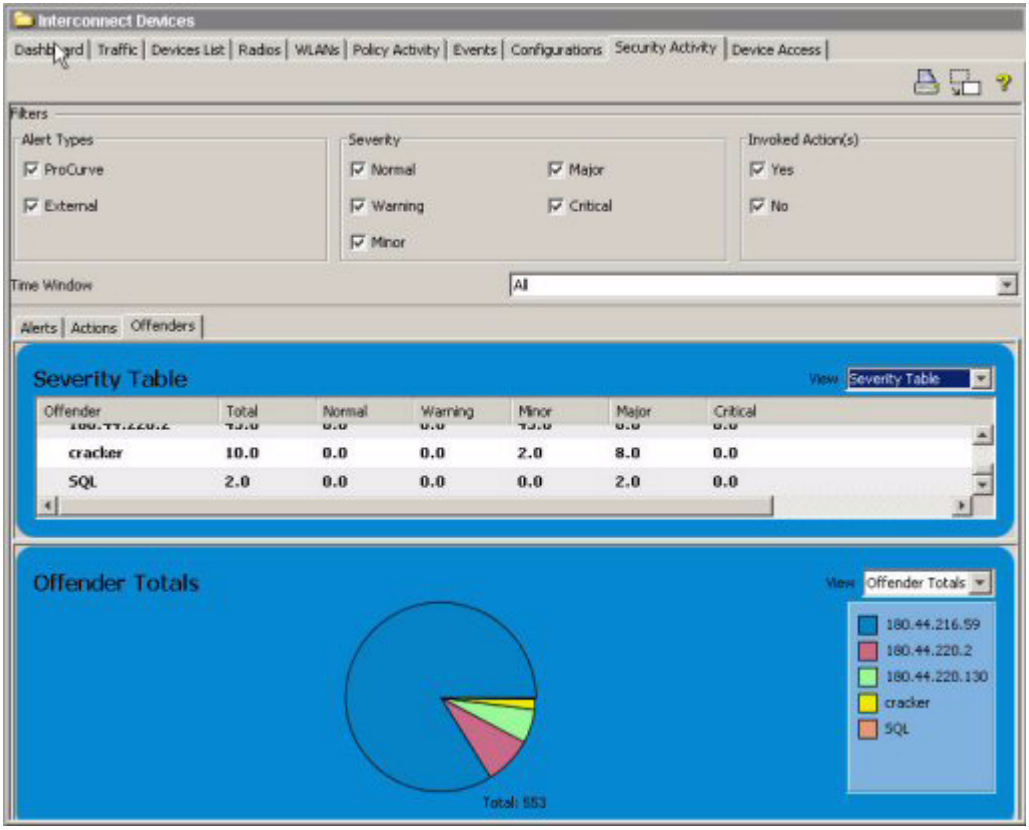

Figure 3-7. Security Activity: Offenders sub-tab display

To display additional information:

- Double-click the value in the Offender column to display the offender details window for a row in the Severity Table.
- Double-click any other value in the Severity table to displaythe Policy History records associated with that value.
- To display the security alerts and policy actions for a section of a pie chart, mouse over the section. Double-clicking a section of a pie chart displays a Policy History window with details for the selected section.

Offenders Table or Bar Chart Summary

The top panel of the Offenders sub-tab displays a summary of alerts by severity for each offender that meet the filter and time criteria. This summary can be viewed in a table or bar chart that provides the following information:

Information	Description
Group/ Offender or Device/ Offender	Depending on the node selected, the top row of the table contains the name of the selected group or device. All other rows contain the offender names associated with the selected group or device.
Total	Total number of alerts for the group, device, or offender that meets the filter and time criteria
Normal	Total number of Normal alerts for the group, device, or offender that meets the filter and time criteria
Warning	Total number of Warning alerts for the group, device, or offender that meets the filter and time criteria
Minor	Total number of Minor alerts for the group that meet the filter and time criteria
Major	Total number of Major alerts for the group, device, or offender that meets the filter and time criteria
Critical	Total number of Critical alerts for the group, device, or offender that meets the filter and time criteria

To display Offender Details for a specific group, device, or offender, click on the entry in the Offender column. Double-click on numeric values in the table to display a Policy History related to the selected value.

Offenders Totals Display

The bottom panel of the Offenders sub-tab displays the total number of alerts that meet the filter and time criteria by offender or severity.

Security Offender Details

You can review the Offender Details by double-clicking on an entry in the Offender column of the Security Activity: Offenders sub-tab display.

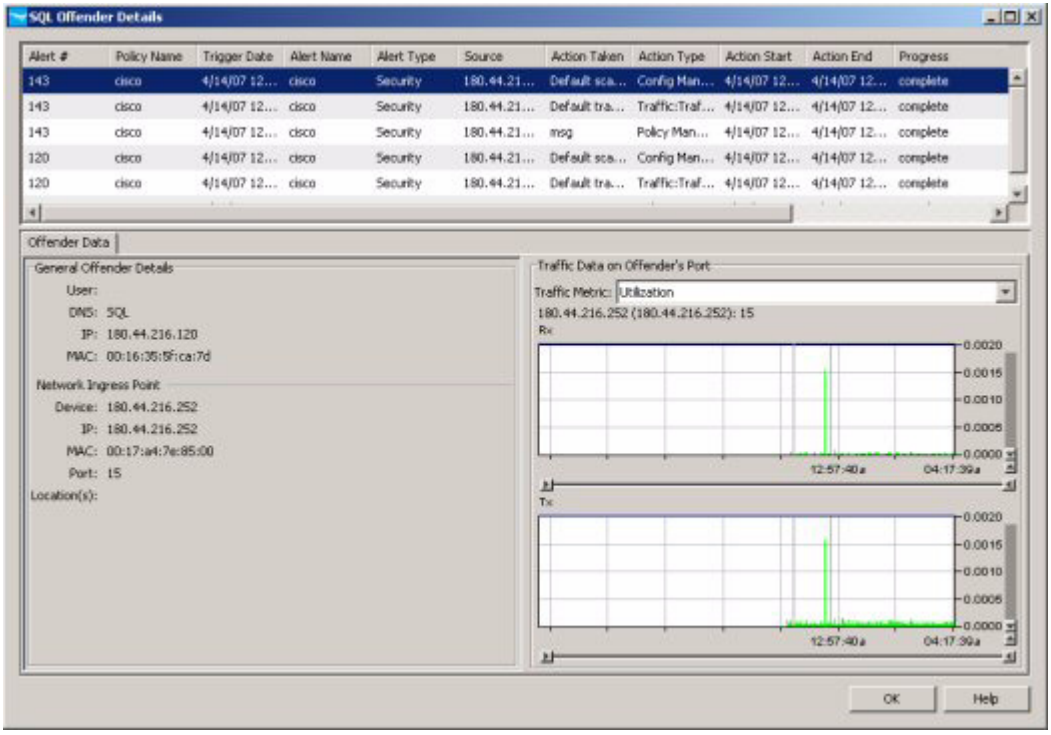

Figure 3-8. Security Activity: Offender Details display

The Offender Details window lets you quickly review pertinent information about a specific offender (user, IP address/DNS name, or MAC address suspected of propagating an attack over the network).

The top panel displays the alert history table for the selected offender, including the related Policy information and actions taken as a result of the alerts generated by the offender. Note that the total number of alerts shown here may be greater than in the Offenders sub-tab display as this window shows activity logged for the selected offender across all devices, not just the Device Group selection from which the Offender Details window was launched.

The lower panel contains an Offender Data tab and an optional Session History tab. These sub-tabs display details about the selected row in the Alert History table at the top of the display.

Offender Data tab provides general details about the offender. The actual level of detail displayed depends on the data provided in the offense and the PCM components installed (e.g., IDM, MM).

- The General Offender Details portion of this panel enumerates data that is most recently known about the offender including IDM user identifier (if applicable), DNS name, IP address, and MAC address. This data comes from the source application that generated the security alert. For example PCM, IDM, IDS, UTM, etc. Note that the User field data may not be available if IDM is not installed and/or depending on how the switch port is configured.

  In some cases and for some types of attack, the General Offender Details will not vary from record to record in the Policy History pane, while in others (for example IP Spoofing attacks) it is common for the MAC address of the offender to be consistent, but the IP and DNS name to vary between Policy History records.

- The Network Ingress Point portion of this panel lists device and port information about the offender's most recent connection, if Find Node was able to discover the location of the end node (MAC or IP address).

If traffic monitoring is enabled on the offender's ingress port, the bottom of the Offender Data tab displays the worst measure of traffic for the discovered switch and port.

If you are also using the IDM module, a second tab is available in the lower panel: a Session History tab, which displays the session history of the offender. It indicates where, when, and for how long the user connected to the network in the past. This data is only available when IDM is in use with PCM+ and NIM, and when there is session data available for the displayed network ingress point.

Using the Security Heatmap

The Network Maps window in PCM also provides an overview of the security state of the managed network based on data from NI Manager. It can display the security totals by category and severity. The example below shows the device view security state by severity.

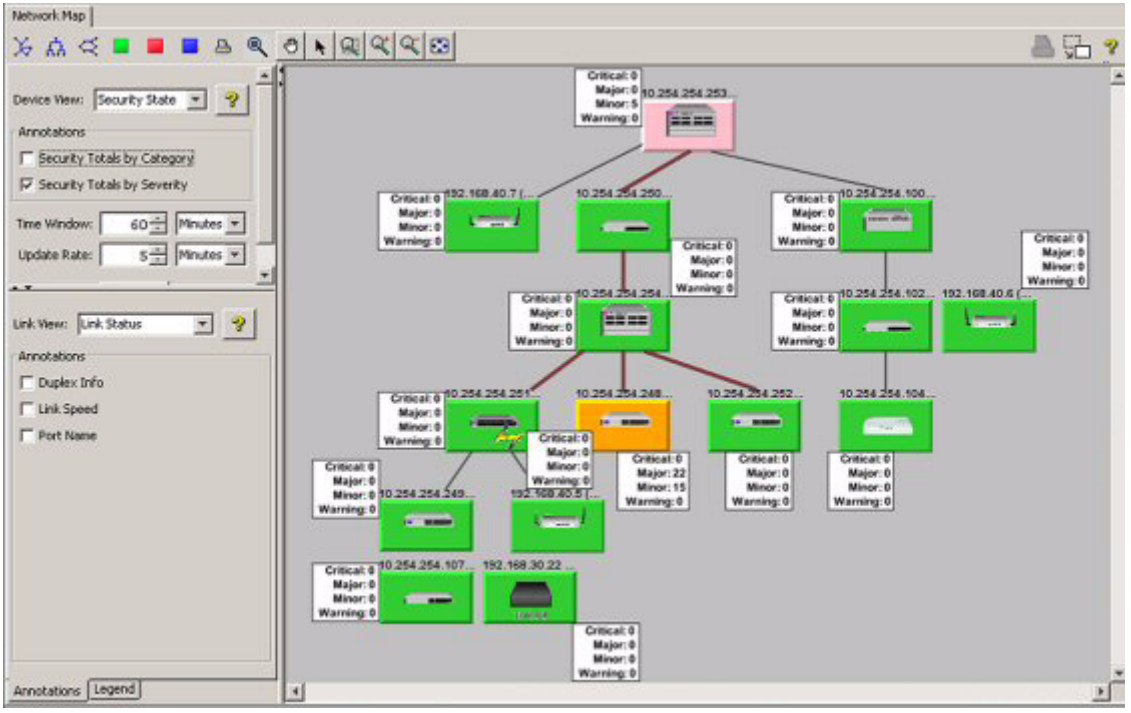

Figure 3-9. Security Map display

You can also set the “Time Window” and “Update Rate” for the map.

- Time Window controls the period of time during which security events in the map occured.
- Update Rate controls how often security events in the map are refreshed. This will affect how often security events are added or removed, relative to the Time Window set.

For additional information on using the Network Maps feature, refer to Chapter 4, "Using Network Maps" in the *ProCurve Manager: Network Administrator's Guide*.

Using Security Policies

Chapter Contents

Security Policies

You can configure NI Manager to just listen and report suspicious activity on the network, or using the Policy Manager feature available from PCM+, you can configure NI Manager to automatically respond to detected threats.

Policy Configuration Overview

Similar to the Policies you configure for PCM, the Security Policies are defined with a combined set of parameters:

- **Times** - Time periods when the policy can be executed. If no time is specified, the policy can execute at any time.
- **Sources** - Devices or ports from which events are received. If no source is selected, the policy will match events from any source. In NI Manager, sources are correlated to the offender connected devices. The events (or traps) do not always come from the switch where the offender is connected. They may come from edge connected switches, sampling devices, the NIM NBAD engine, or external IPS, IDS, UTM devices.
- **Targets** - Devices or ports on which a defined action will be performed in response to an alert, if applicable.
- **Alerts** - A defined trigger used to launch an Action. Alerts can be event-driven, or scheduled to occur at a specified time. Security alerts are event-driven.
- **Action** - The action taken on Targets in response to the Alert. If no action is specified, the alert will generate a Policy Manager event in the Event browser.

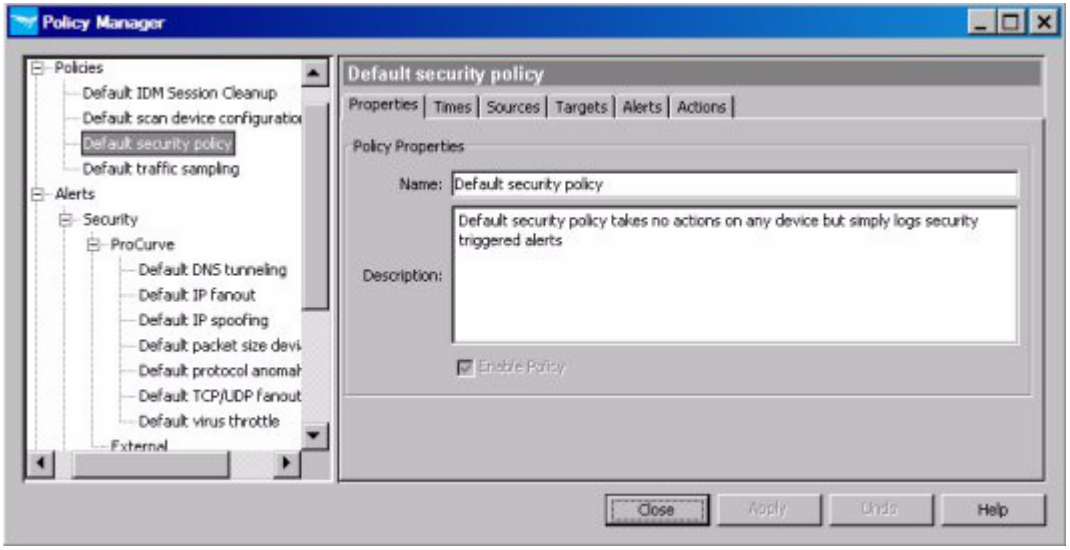

Figure 4-1. Default Security Policy

Multiple parameters of each type can be applied to a Policy. When the Policy is activated, it reads through each set of parameters until a match is found. For the policy to execute, it must find a match for each defined parameter. If there is no match the policy does not execute. For example, if you configure a policy with Times limited to "weekdays", defined as 9:00 am to 5:00 pm and an alert trigger is received at 10:00 pm, the policy will not execute.

You can define custom Times, Alerts, and Actions separate from the Policy configuration process. The new 'definitions' will be available in the selection lists in the Policy Configuration Manager when you create your Policy. NI Manager provides a set of pre-defined alerts for each monitored security threat type, and a pre-defined Security Policy to log alerts generated by ProCurve security events in the PCM Events browser, and security policy history.

Security Alerts Overview

The Alerts for Security policies differ from PCM in that the Alert types are one of the following:

- **ProCurve Alerts** - Defines alerts relative to security events generated by ProCurve devices.
- **External Alerts** - Defines alerts relative to events received from external devices, such as UTM, IDS, and IPS devices.

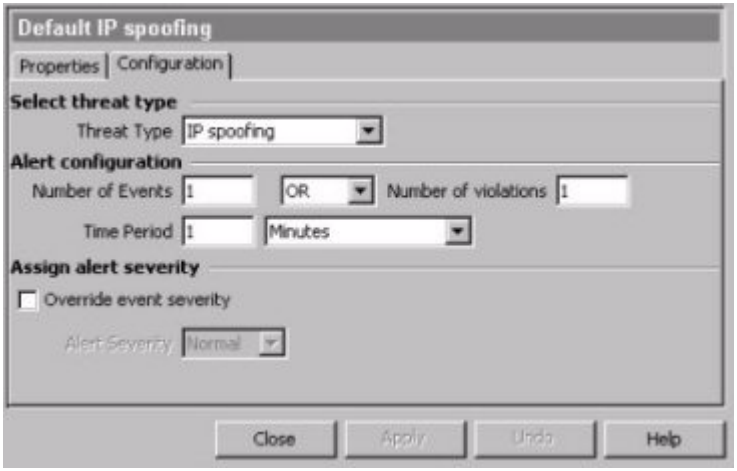

Figure 4-2. Default Alerts Configuration tab

The Security Alerts configuration works together with the Security Analysis Sensitivity Settings to further define when an action will be taken in response to suspect activity.

Security Actions Overview

You can use the Default Security Policy to determine the sensitivity level you want to apply for anomaly detection, and to develop a higher level of confidence in NBAD and NI Manager functionality. Once you have the Security Analysis sensitivity settings configured, and a high confidence level in the NI Manager and NBAD function, you can create policies for automated response and mitigation of security threats.

The following Action Types can be used in Security Policies to provide an automated mitigation response:

- MAC Lockout: Use to block traffic on the target device from the specified MAC address.
- Port Settings: Enable/Disable - Use to enable or disable a port.
- Port Settings: Rate-limit - Set rate limits on a port.
- Remote Mirror: - used to forward suspect traffic to an external IDS, IPS, or UTM device for additional analysis.
- Security: VT Configuration - Sets virus throttle parameters on target device. (For use on 5400zl, 3500yl, 6200yl, and 5300xl devices.)
- VLAN - Lets you create a VLAN on target device or port to restrict access to the rest of the network.

NOTE: The Rate-limit and VLAN action types are also used by IDM. If these actions have already been applied on a device or port by IDM, the NI generated action will not be taken, and the conflict will be logged to the Events browser.

ProCurve recommends that you use Security Manager with only the default 'notification' policy to start. You can alter the Security Monitoring Preferences to tune the Analysis Sensitivity settings. You may also want to try altering the Alert configuration for the Security:ProCurve Alert types to see the interaction between Analysis Sensitivity settings and Alert configuration.

When you first create your own security policies, ProCurve recommends that you use the Policy Manager Preferences for "Configuration Changes" to enable a test mode for security policies. See “Setting Policy Management Preferences” on page 4-25 for details.

When you are confident you have the security monitoring system tuned, and are confident the policy will operate as intended you can create and enable Policies for automated mitigation response to detected threats.

Configuring Security Policies

To configure a Security Policy:

1. Click the Policy Manager icon in the toolbar to launch the Policy Configuration Manager window.

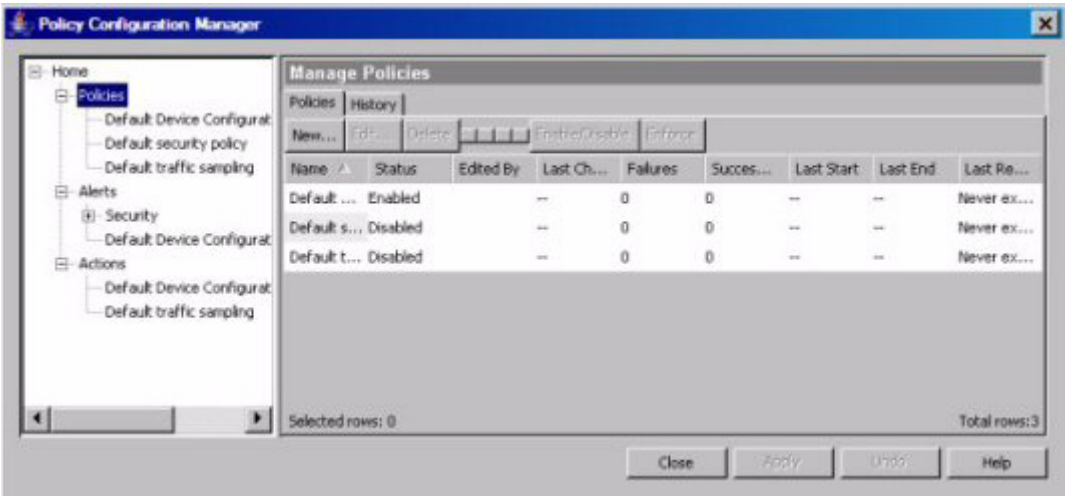

Figure 4-3. Policy Configuration Manager: Manage Policies panel

2. Select the Policies node in the navigation tree to display the Manage Policies panel, then click New... to launch the Create Policy dialog.

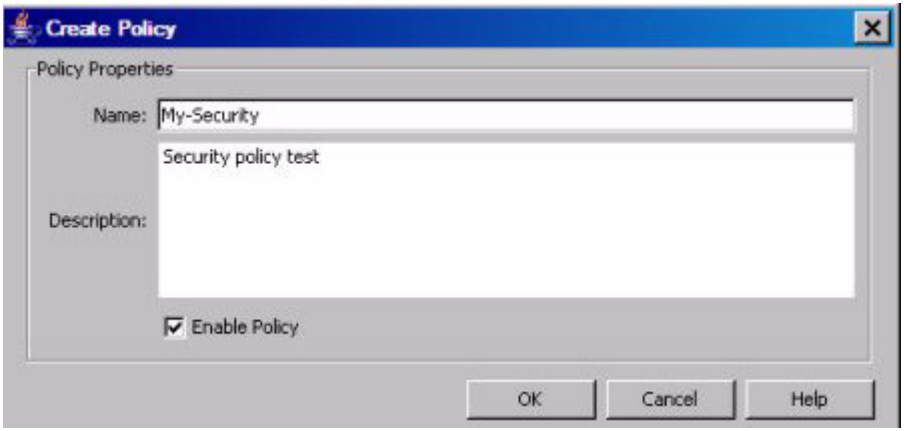

Figure 4-4. Policy Manager: Policy Properties display

3. Fill in the Policy information:
   a. In the Name field, type a name to identify the policy.
   This name will appear as a node in the Policies navigation tree, and in the list in the Manage Policies panel.
   b. In the Description field, type in a brief description to help you identify the policy and what it will do.
   c. Click the Enable Policy check box to enable the policy.
   A check in the box indicates the policy will take effect immediately when its configuration is completed.
   If the check box is empty, the Policy is disabled. It will not take effect until you Enable it.
   d. Click OK to save the Policy Properties and display the Policy Configuration panel for your new policy.

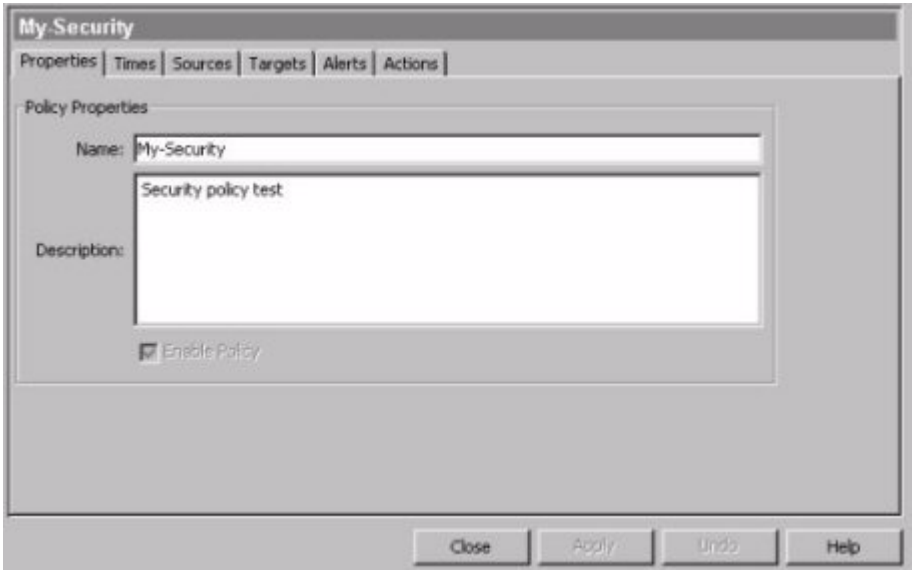

Figure 4-5. Security Policy: Properties tab

4. Click the Times tab to configure the time periods that will be applied for your policy.

   Applying "Times" to a policy restricts the application of the policy to the defined time. If no times are selected, the policy will always be active and can be executed at any time.

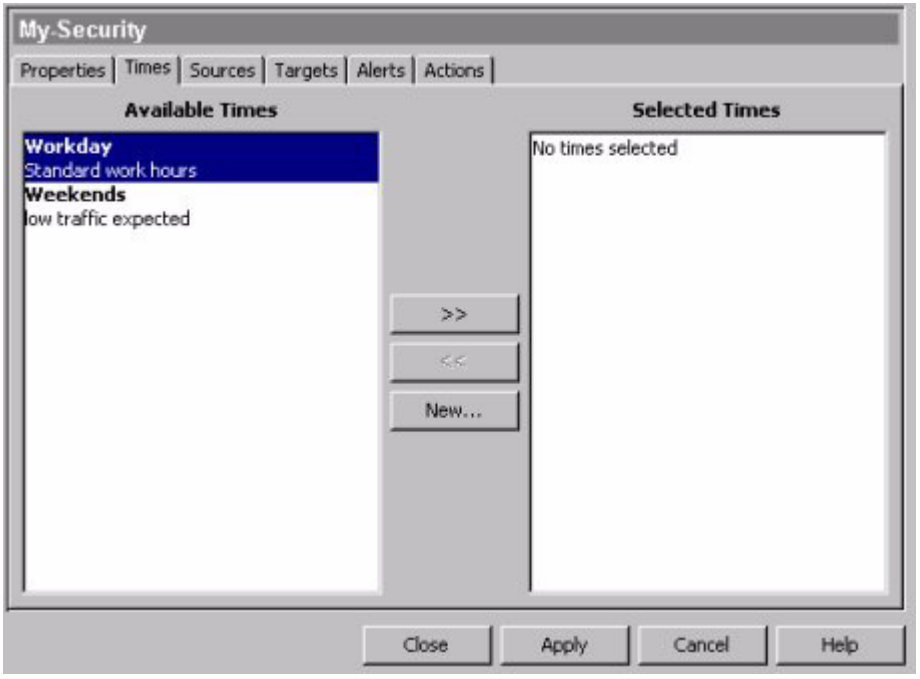

Figure 4-6. Policy Configuration: Times tab

5. To apply a time, select it in the Available Times list on the left, then click **>>** to move it to the list of Selected Times.

   You can apply more than one Time. When the policy is activated, it will read each time until a match is found.

   Click New... to launch the Configure Times dialog. Refer the Policy Manager chapter of the PCM Network Administrator's Guide for details on configuring Times.

6. Click Apply to save the changes.

7. Click the Sources Tab to configure the device groups from which an event trigger will be applied.

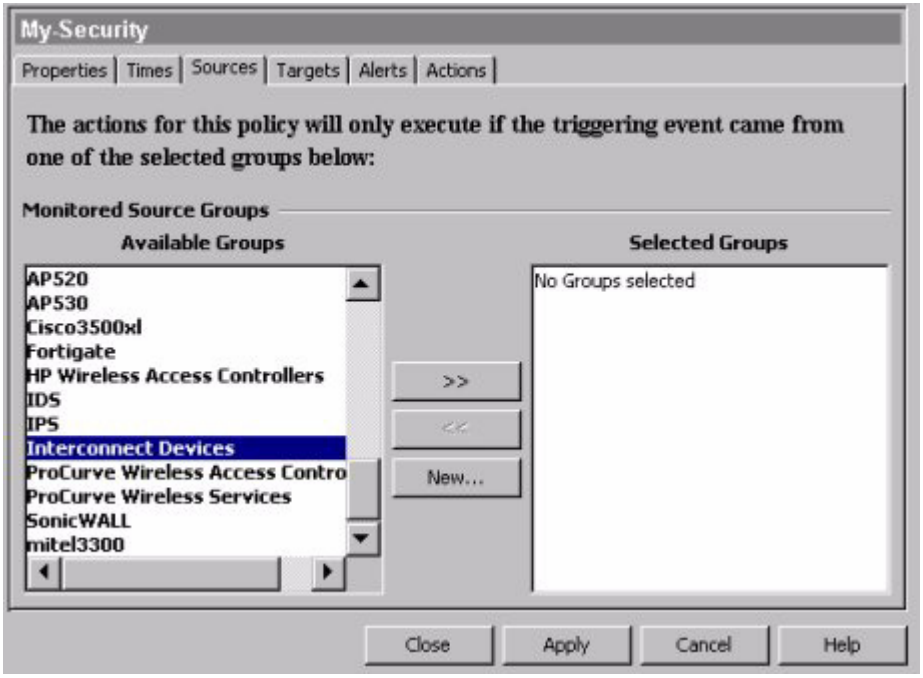

Figure 4-7. Policy Configuration: Sources tab

8. To apply a Group, select it in the Available Groups list on the left, then click **>>** to move it to the list of Selected Groups on the right.

   If no group is selected, the Policy will accept events from any source.

   If you select Interconnect Devices, the Policy will accept events from any of the pre-defined ProCurve Device groups.

   If you select more than one group, the policy will only execute if an event is received from a device in the Selected Groups list.

   If you configured Custom Groups they will appear in the Available Groups list for application as monitored source. You can use a Custom Group to define a group of ports on various devices, rather than all ports on a single device type. See "Custom Groups" on page x-x for details.

9. Click Apply to save the changes.

10. Click the Targets Tab to configure the device groups to which the policy action will be applied.

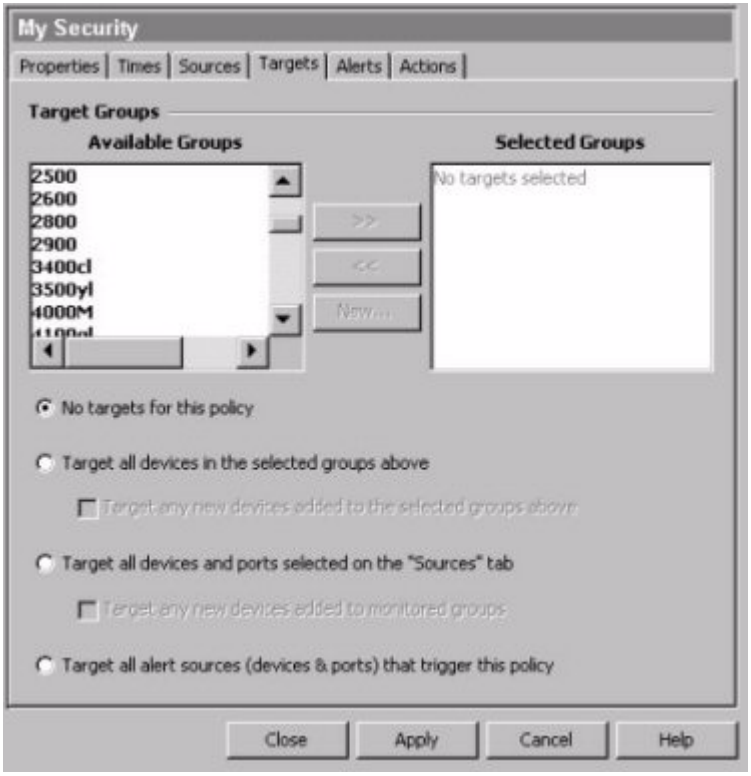

Figure 4-8. Policy Configuration, Targets tab

11. To apply a Group, select it in the Available Groups list on the left, then click **>>** to move it to the list of Selected Groups on the right. The policy will be applied to all discovered devices of that type, unless you select one of the Target qualifiers in the bottom portion of the window.

    If no group is selected, the Policy action will not be applied to any device, and the No targets for this policy option is selected.

    If you select the Interconnect Devices group, the Policy will accept events from any of the pre-defined ProCurve Device groups.

    If you configured Custom Groups they will appear in the Available Groups list for application as monitored source. You can use a Custom Group to define a group of ports on various devices, rather than all ports on a single device type. See "Working with Custom Groups" in the *ProCurve Manager v2.2 Network Administrator's Guide* for details.

    Click New... to launch the Create Group dialog to define a Custom Group and add it to the list of available groups.

12. Apply any target qualifiers by clicking the radio button or check box to select it.
    - Target all devices in the selected groups above will apply the policy to all devices included in the Selected Groups on the Targets tab. Selecting this option enables the check box so you can:
      Target any new devices added to the selected groups above. Use this option to apply the policy to newly discovered devices. This is useful for applying standardized configurations.
    - Target all devices and ports selected on the "Sources" tab will apply the policy to all discovered devices included in the Selected Groups on the Sources tab. Selecting this option enables the check box so you can:
      Target any new devices added to monitored groups will apply the policy to any newly discovered devices in the Selected Groups on the Sources tab.
    - Target all alert sources (devices & ports) that trigger this policy will apply the policy action to any device(s) or port(s) identified by the trigger alert. For example, if a virus throttling event triggers the policy and the alert was configured to use the trap contents as the alert source then the policy will target the edge port to which the host identified by VT is connected.

      For Security policies, the source is the offender connected device. When executing security policies, NI Manager dynamically applies the PCM+ "find-node" feature to determine on which device the offender is connected. If this target option is selected then actions will be taken only on the devices where the offender is connected, against the offender's MAC and, or port as applicable.
13. Click Apply to save the changes.
14. Click the Alerts tab to configure the alerts that will trigger the policy execution.

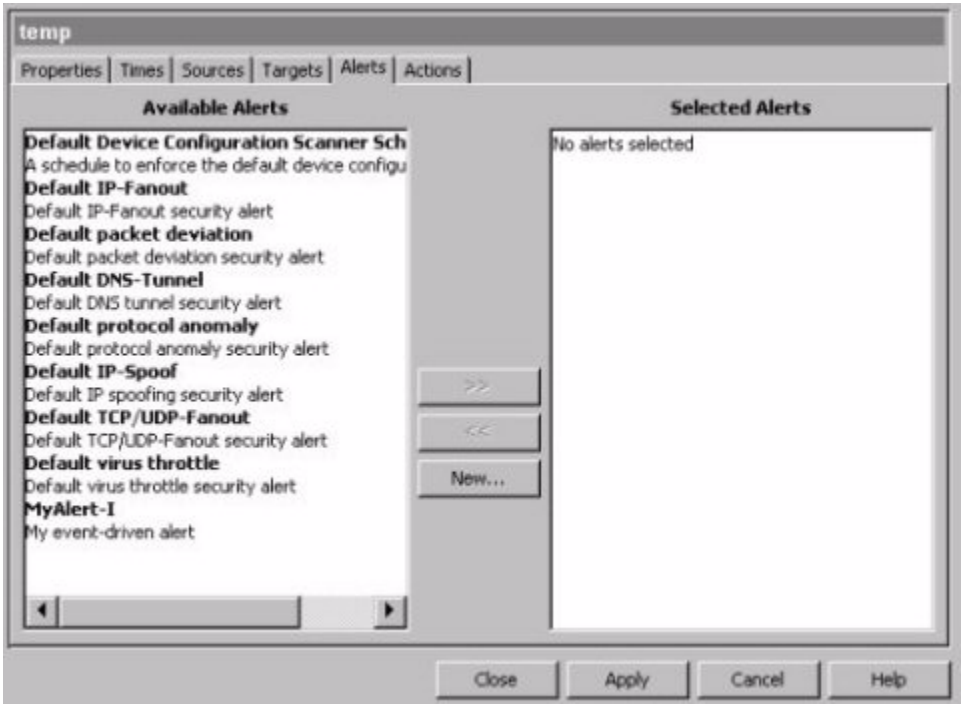

Figure 4-9. Policy Configuration, Alerts tab

15. The Alerts tab lists the pre-configured alerts in the Available Alerts list. To apply an Alert, select it in the Available Alerts list on the left, then click **>>** to move it to the list of Selected Alerts on the right.

    You can select multiple alerts, and when an event is received each of the alerts will be evaluated until a match is found. The policy will execute on the first matching Alert.

    If you configured any custom Alerts they will appear in the Available Alerts list for application as a source

    Click New... to launch the Create Alert dialog to define an Alert and add it to the list of available Alerts. Refer to “Configuring Security Alerts” on page 4-15 for details.

16. Click Apply to save the changes.

17. Click the Actions tab to configure the actions the policy will take when it is executed.

    If you do not specify an Action for the policy, when the policy executes it will log a PolicyManager event in the Event browser.

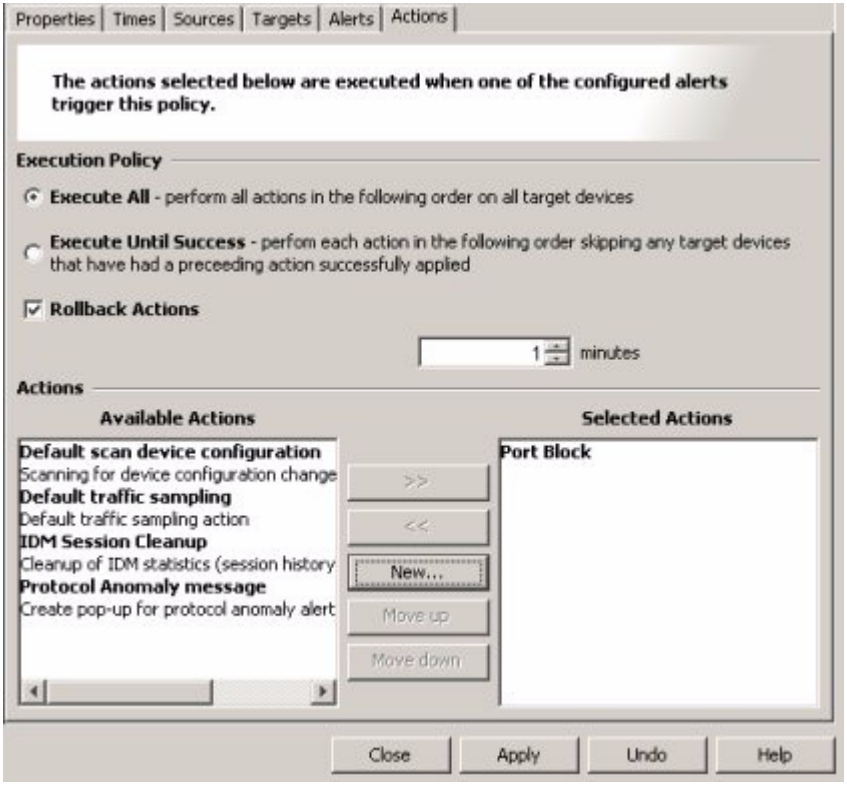

Figure 4-10. Policy Configuration, Actions tab

18. Select the Execution Policy options you want to apply by clicking the radio buttons or check box.
    - Execute All - this is the default setting. Indicates all selected actions will be attempted when the Policy runs.
    - Execute Until Success - this will attempt to execute each selected action on each target device or port in the order listed. As soon as an action completes successfully, the policy moves to the next target device/ port and attempts to execute the selected actions.

      This can be used to create a single policy to mitigate for a confirmed attack across multiple device types on the network. For example:

      i. Create one action for MAC Lockout that uses the MAC supplied in the event source, and one action to Disable the port (Port Settings: Port Status action option).
      ii. Create a Policy that targets all source devices/ports when a Protocol Anomaly alert is generated.

iii. In the Actions tab, select the MAC Lockout action and the Disable Port action, in that order.

When the Policy executes, it will first attempt to use MAC lockout on the target device or port. If the target device does not support the MAC lockout feature, the Policy will attempt the Disable Port action.

- Rollback Actions - For Action types that support a rollback operation, it will stop the action, returning the target of the action to its original state after the time (minutes) specified. This option is not enabled until an action that supports rollback is selected. The rollback feature is supported by the following actions:
  - Port Mirroring
  - MAC Lockout
  - Port Settings: Enable/Disable)
  - Port Settings: Rate Limit
  - Traffic Sampling

Note that you should configure the Alert time window for the policy to be larger than the Rollback Actions time, or you may create a loop state between events and mitigating actions.

19. The Actions tab lists the pre-configured actions in the Available Actions list. To apply an Action, select it in the Available Actions list on the left, then click **>>** to move it to the list of Selected Actions on the right.

    You can select multiple actions to apply when the Policy executes. The actions will be applied according to the Execution Policy options you select.

20. Click Apply to save the changes, then click Close to exit the Policy Configuration Manager window.

    If you click Close before Apply, you will be prompted to save or cancel the changes.

Configuring Security Alerts

NI Manager comes with a set of pre-defined ProCurve Security Alerts for the monitored Threat types. You can edit the pre-defined ProCurve Security Alerts, or create your own. You can also create Security Alerts that work with External event sources, such as IDS, IPS, or UTM appliances.

To Configure ProCurve Security Alerts

1. Click the Policy Manager icon in the toolbar to launch the Policy Configuration Manager window.
2. Click to expand the Alerts ->Security node in the navigation tree
3. Click to expand the ProCurve node and display the Manage ProCurve Alerts panel.

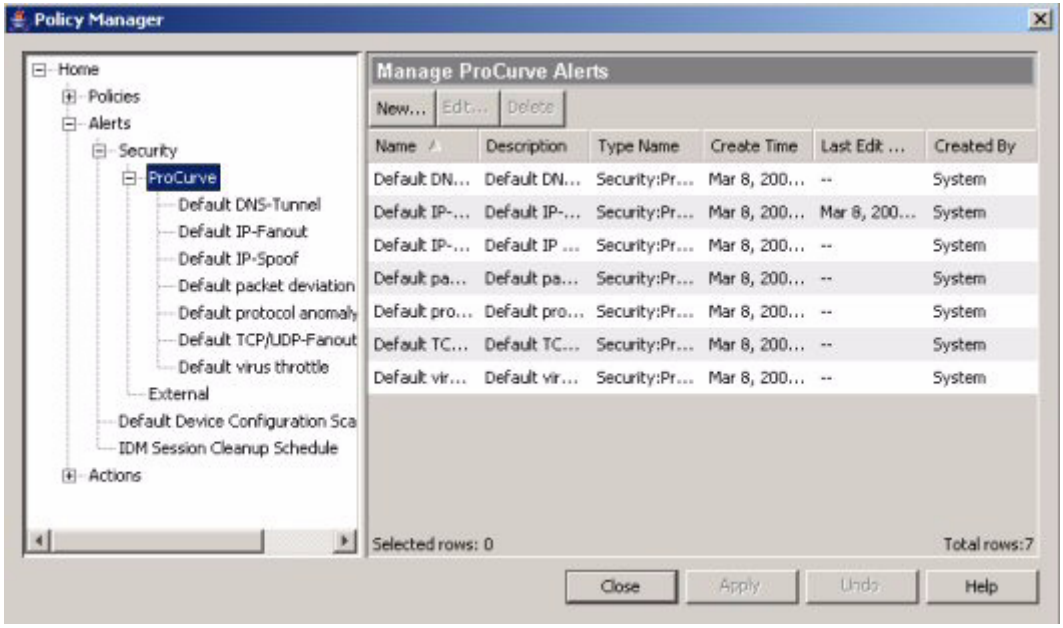

Figure 4-11. Security: Manage ProCurve Alerts display

4. Click New... to launch the Create Alerts dialog.

Figure 4-12. ProCurve Security Alert type selection

5. Enter the Alert properties:
   a. Select the Security:ProCurve Alert option from the pull-down menu.
   b. Type in a Name for the Alert. This name will appear in the list of available alerts for ProCurve Security Policies.
   c. Enter a Description for the Alert if desired.
   d. Click OK to save the changes and display the Alert configuration panel with the Alert properties.
6. Click the Configuration tab to complete the configuration.

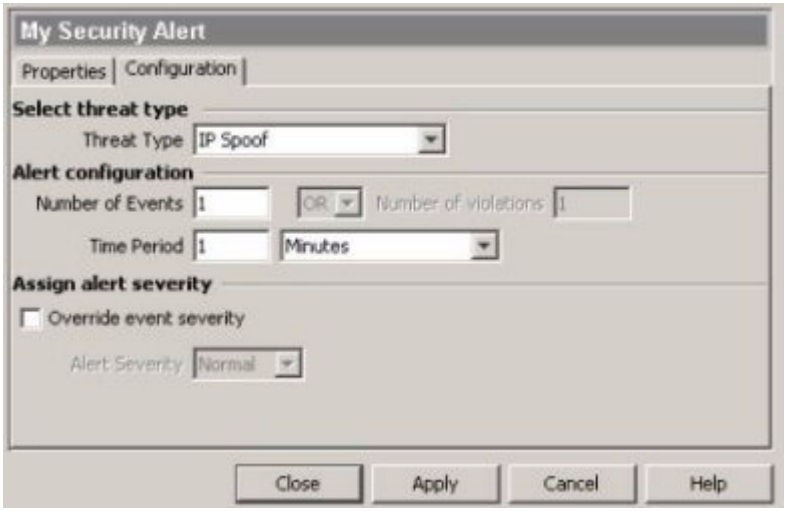

Figure 4-13. ProCurve Security Alert, Configuration tab

7. Select the parameters for the Alert:
   a. Use the pull-down menu to select the Threat Type, one of: IP Spoof, IP Fanout, TCP/UDP Fanout, DNS Tunneling, Average Packet Deviation, Virus Throttle, or Protocol Anomaly.

b. Use the Alert Configuration fields to set event parameters. That is the number of events received within the specified time period that will cause an alert:
   – Time period: type in the number, then select the value, one of: Seconds, Minutes, Hours, or Days.
   – Number of events: type in any number

   This violation count is given by NBAD engine in every event (if NBAD doesn't give a number then I set this to 1 by default). User can configure alerts based upon number events and/or number of violations. This violation count indicates what's the intensity (ex: attacker is trying to attack 100 servers) of the attack

c. Assign alert severity. The default associated with ProCurve generated alerts is normal. To associate a different severity with the alert:
   – Click the check box to select Override event severity and enable the Alert Severity pull-down menu
   – Select the Alert Severity from the menu, one of: Normal, Warning, Minor, Major, or Critical.

8. Click Apply to save your configuration, then click Close to exit the Alert Configuration panel.

   If you click Close before Apply, you will be prompted to save or cancel the changes.

The Alert Configuration parameters work together with the Analysis Sensitivity settings described under "Setting Security Preferences" on page 2-3. The sensitivity settings configure how many instances of suspect activity must be detected before an event is generated. The Alert Configuration parameters specify the number of events that must be logged within the given time period in order for the associated Policy Action to be executed. These two settings work together to help control the rate of alert notification, and to control instances of false positives.

Note that if an alert fires it will not fire again for the configured time period (if any). However, keep in mind that a policy with numerous alerts associated with it can execute as a result of any of the other alerts being triggered. Thus, even though alert-1 is going to stay 'silent' for a two minute time period, it is possible for alert-2 to trigger the policy execution during that time period.

Editing Security Alerts

You can edit the default Security Alert configurations, or any other alerts you create at any time. Simple select the Alert in the navigation tree, then click the Configuration tab in the right panel to edit the alert parameters.

An alternate method:

1. Display the list of alerts in the Manage ProCurve Alerts panel
2. Select the alert in the list, then click the Edit... button.
3. Click the Configuration tab to display the alert parameters for editing.

Refer to step 7 and 8, above, for details on setting the alert parameters.

Configuring Actions for Security Policies

In order to automate a response to detected threats or attacks on your network, you need to define the actions that NI Manager will take when a Security Alert is issued.

To define an Action for Security Policies:

1. Click the Policy Manager icon in the toolbar to launch the Policy Configuration Manager window.
2. Click the Actions node in the navigation tree to display the Manage Actions panel.

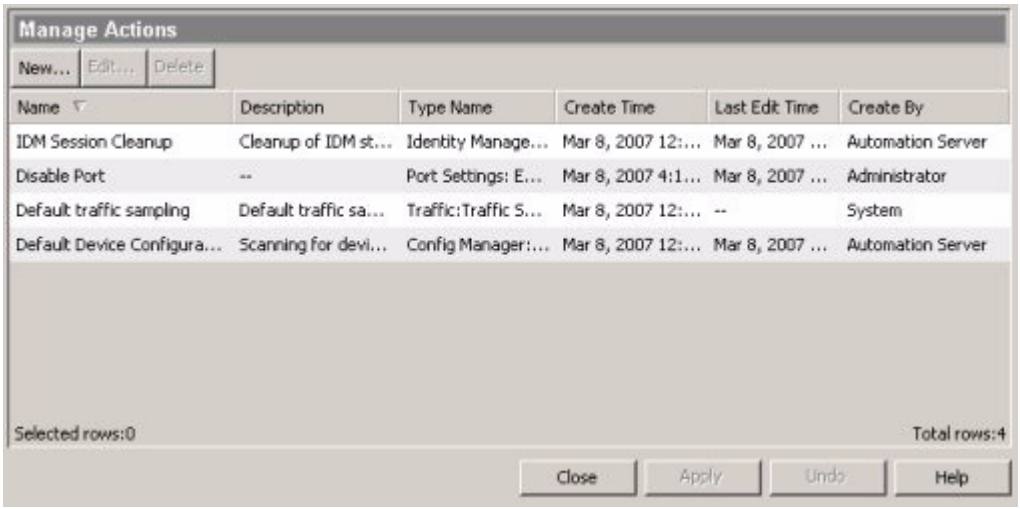

Figure 4-14. Policy Manager: Manage Actions display

The Manage Actions panel displays a table with the following information for each of the Actions already defined:

Name: The name of the action

Description: A brief description of what the action does.

Type Name: The action type. ProCurve Manager comes with a set of pre-defined Action Types that are used when configuring policy actions

Create Time: Date and time the action was originally created

Last Edit Time: Date and time the action was last edited

Create By: Login Name or ID of the creator of the action. "Automation Server" and "System" indicates a pre-configured action supplied with the PCM+ software.

3. Click New... to launch the Create Action dialog.

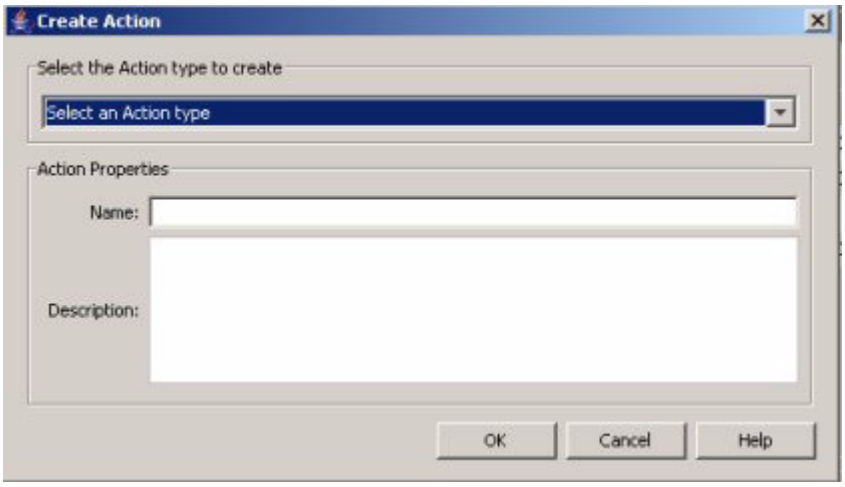

Figure 4-15. Policy Manager: Create Action properties display

4. Use the pull-down menu to select the Action Type. You can select any of the available action types, but there are several that are designed for use in automated security mitigation or threat response:
   - MAC Lockout: Blocks traffic from a specified MAC address on the target device.
   - Port Settings: Enable/Disable - use to enable or disable a port
   - Port Settings: Rate-limit - Set rate limits on target ports.
   - Security: VT Configuration - Sets virus throttle parameters on target device. (For use on 5400zl, 3500yl, 6200yl, 5300xl, and 3400cl devices.)
   - VLAN - Lets you create a VLAN on target device or ports to restrict access to the rest of the network.

NOTE: The Rate-limit and VLAN action types are also used by IDM. If these actions have already been applied on a device or port by IDM, the NI generated action will not be taken, and the conflict will be logged to the Events browser.

5. Enter the Action Properties
   a. Type in a Name for the Action (required).
   b. Type in a Description for the Action (optional).

6. Click OK to save the Action properties and display the Action configuration panel.

   The Properties tab is the default display, with the information you just defined for the Action.

   Click the action type configuration tab to set the parameters applicable to the selected action. The tab name will reflect the Action Type you selected.

   For example, if you selected the Mac Lockout Action Type, then you will click the MACs tab in the configuration panel, as shown in the following figure.

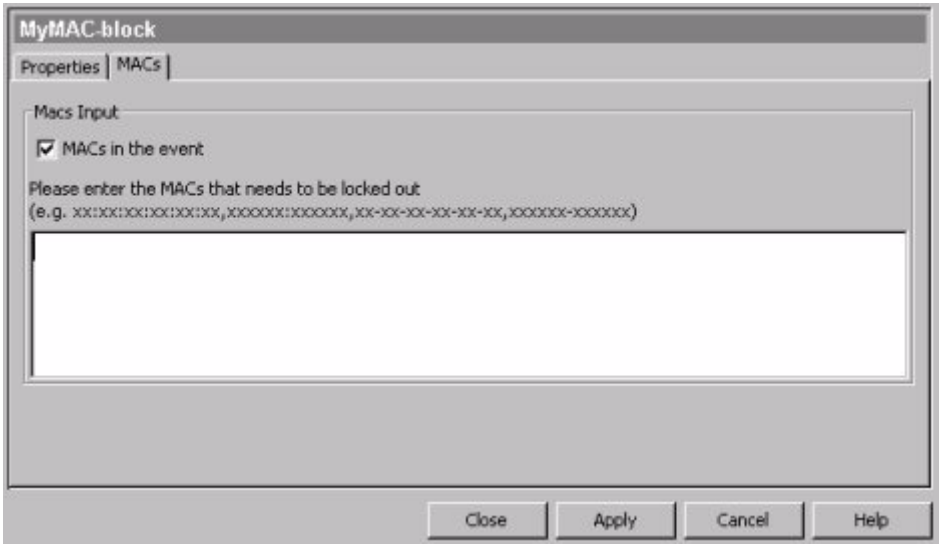

Figure 4-16. Policy Manager: Action Configuration example

7. Fill in the parameters needed for the Action type. (See “Security Action Type Definitions” on page 4-22 for details on the Action type, Tab label, and parameters)

   In the case of MAC lockout, you can click to select the MACs in the event option. This will block the MACs included in listed in the event text when the action when the Security Policy is executed on the target device.

8. Click Apply to save the configuration, then click Close to exit the action configuration panel.

   If you click Close before Apply, you will be prompted to save or cancel the changes.

Security Action Type Definitions

The following tables provide a description of the Action types that are designed for use in automated security mitigation or threat response, along with the tabs and configurable parameters for that action. For a complete listing of Action types, refer to the chapter on "Using Policy Manager Features" in the *ProCurve Manager v2.2 Network Administrator's Guide*.

Note that the "Properties Tab" is not listed as it is the same for all Action types; that is, you use it to select the action type, and enter a name and description for the configured action.

Table 4-1. Policy Manager Actions

Action	Description	Tabs	Parameters
Display Message Dialog	Use to display text pop-up message for the alert	**Message**	• Message text Can use substitution list for variables.
Execute Command on Server	Execute system command on management server	**Command**	• Command text can use substitution list for variables.
Forward Trap		**Trap**	• Trap Receiver (IP address) • Port (default is 162) • Content - enter contents to be included in trap message, can use substitution list for variables.
Send Email	Fwd e-mail with alert details	**Email**	• SMTP Profile* • To: email address • From: email address • Subject: text input, can use variable substitutions. Message Body: text input, can use variable substitutions shown. * Prerequisite: Must create SMTP profile in PCM.

Content Variables for use in Policy Manager Actions: The Substitution List in the tabs for configuring Policy Manager actions describes the variables you can use in the Content and text fields. The variables will be replaced (before the trap or message is forwarded) by data from fields in the event that invokes the alert.

Table 4-2. Port Settings Actions

Action	Description	Tabs	Parameters
Port Setting: Enable/Disable Port	Use to temporarily shut down a port	**Port Status**	• Enabled • Disabled
Port Setting: Guaranteed Minimum Bandwidth (GMB)	Use to set the percentage of bandwidth allocated to the various priority levels of each outbound traffic priority queue of the targeted ports on devices that support GMB.	**Guaranteed Minimum Bandwidth**	• Configure GMB on target port - Disable GMB - Enable GMB • If enable GMB, set Low Priority Queue % Normal Priority Queue % Medium Priority Queue % High Priority Queue %
Port Setting: Quality of Service	Used to set the priority of packets handled by the targeted ports on devices that support Quality of Service (QoS).	**Quality of Service**	• Configure source port QoS settings on targeted port • No override • 802.1p Priority, priority (0 - 7) • DSCP Priority, priority (0 -7) and codepoint (0 - 63). See Operating Notes for QoS below for additional detail.
Port Setting: Rate Limit	Limits the inbound bandwidth on a switch port that a user or device can utilize. Effectively enforces maximum service level commitments granted to network users.	**Rate Limit**	Configure Rate Limiting on target ports • Disable Rate Limiting • Enable Rate Limiting • Rate Limit % : set the maximum percentage of bandwidth to be allocated to the targeted ports.

Operating Notes for QoS:

With No override, QoS does not affect the packet queuing priority or VLAN tagging, and packets are handled as follows:

- If received and forwarded on a tagged VLAN, the 802.1 priority is not changed.
- If received on an untagged VLAN and forwarded on a tagged VLAN, the 802.1 priority is 0 (normal).
- If forwarded on an untagged VLAN, no 802.1 priority is used.

For 802.1p Priority:

Assigns an 802.1p traffic priority setting (0-7) carried by packets moving from one device to another in an 802.1Q tagged VLAN environment. The switch uses the 802.1p priority to determine the queue in the outbound port to use for the packet. If the packet leaves the switch in a tagged VLAN, it carries the 802.1p priority to the next downstream device. If the packet leaves the switch through an untagged VLAN, this priority is dropped, and the packet arrives at the next downstream device without an 802.1p priority assignment. 802.1p priorities range from 0-7 with 7 being the highest priority.

For DSCP Priority:

Associate a handling priority with a codepoint in an incoming IPv4 packet. DSCP priority is not dependent on tagged VLANs to carry priority policy to downstream devices. DSCP priorities range from 0-7 with 7 being the highest priority. Codepoints range from 0-63. The priority selected will be assigned to this codepoint regardless of its current setting.

Table 4-3. Other Actions for Security Management

Action	Description	Tabs	Parameters
Mac Lockout	Use to block access to the target device for the specified MAC address.	**MACs**	• Select option to use MACS in event, or • Type in MAC addresses to be blocked.
Security: VT Configuration	Configure Virus Throttle on target device.	**VT Configuration**	• Disable/Enable, • Set Global sensitivity (low, medium, high, aggressive) • Set VT Action to take. (notify only, throttle, block, no)
Traffic: Traffic Sampling (SFLOW, XRMON	Use to automatically enable or disable traffic sampling (sFlow, XRMON) in response to an event.	**Traffic Sampling State**	• Click to select the sampling option – Enable traffic sampling – Disable traffic sampling
Vlan Manager: Create VLAN		**VLAN Settings**	• Click check box to select the Ignore and reboot options: – Ignore if VLAN not enabled on device – Ignore if max. VLANs reached on device – Ignore VLAN IDs that already exist on device. – Allow device reboot if needed
		VLAN Information	• VLAN name, • IP Config (DHCP or disabled), • Subnet Mask (for dhcp), • VLAN IDs for – Tagged – Untagged – Forbidden VLAN IDs

Setting Policy Management Preferences

Use the Preferences for Global Policy Management to set the parameters that define the number of entries to include in the Policy History, the global setting for execution of device configuration changes by policies, and logging options for policies in the Events browser.

To set Policy Management Preferences:

1. Navigate to the Policy Management Preferences window.

   a. Click the Preferences icon in the toolbar (or use the Tools Menu).
   b. In the Preferences navigation Pane, select Policy Management.

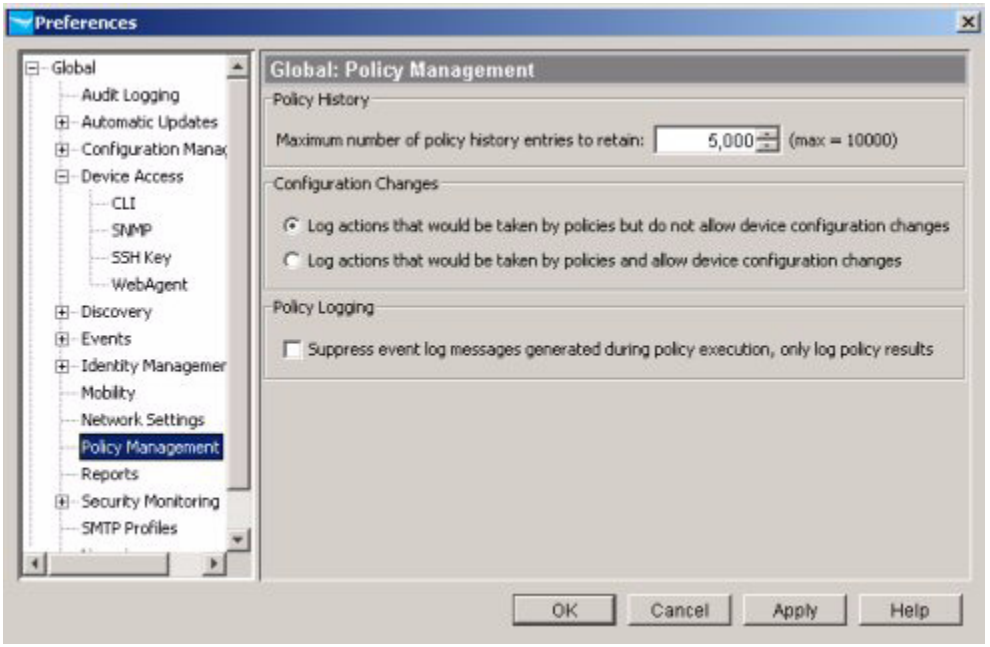

Figure 4-17. Preferences, Global: Policy Management display

1. Select the Maximum number of policy history entries to retain in the Policy History log. The default is 5,000. You can type in a number, or use the buttons to increase or decrease in steps of 100.

2. Click the radio button to select the Configuration Changes option you want to apply to all policies:
   - Log actions that would be taken by policies but do not allow device configuration changes

     This option is useful for monitoring or testing of policies prior to full implementation. It will log the policy activity as if all actions were executed, but it will not actually allow any policy action to change a device configuration.
   - Log actions that would be taken by policies and allow device configuration changes.

     This allows full implementation of the policy, including device configuration changes. Use this option when you have tested the policy and are confident the result of a device configuration is what you intended.
3. Click the check box to enable the Policy Logging option.

   The "Suppress event log messages during policy execution, only log policy results" will trim the reporting of intermediate steps taken during the execution of a policy, and log only the result of the final policy action. ProCurve recommends that you do not suppress Policy Logging until you have tested the policy and fully understand how your policy is operating. Once you are confident the policy is operating as intended you can suppress policy logging to reduce the number of policy activity events in the Events browser.
4. Click OK to save your changes and exit the window

Notes:

- The number of Policy History entries retained is global and effects all policy history tables (Policy Activity tab, Security Activity tab and Policy Manager dialog). The history size chosen will impact the length of history available, as older records will be deleted to make room for new records.
- Policy History entries are not archived, except in the sense that the policy activity events shown in the event browser will be archived.
- When you enable the Policy Logging suppression, you will not be able to recover the suppressed policy events, they are lost forever.

Using External IDS/IPS/UTM Devices

ProCurve Network Immunity Manager can also be used with external IDS/IPS/UTM devices to provide an additional level of analysis for suspected security threats. This helps to improve confidence in attack identification and reduce occurrence of false positives.

Operating Rules:

- You must configure the external IDS, IPS, or UTM device to have the PCM server as a trap receiver.
- Use of external IDS, IPS, or UTM devices requires the ability to enable remote mirroring from the ProCurve managed network (VLAN) to the external device.
- The following ProCurve network switches support Remote Port Mirror (intelligent mirroring): 3500yl, 5400zl, and 6200yl, with switch software version K.12.03 or newer.
- You will need a minimum of two devices that support remote mirroring. One will be used as the remote source (monitored) port and one as the destination mirror port.
- Devices in between the source port and mirror port must have gigabit (1 GB) links, and the VLAN should have jumbo frames enabled, otherwise packet loss may occur between the source and the external device.
- Using Remote Mirror to external devices generally requires two separate policy configurations. One to capture the suspect traffic and forward it to the external device via the Remote Mirror, and a second policy to process the trap.
- You must have the Policy Management Preferences set to allow configuration changes.

Configuring Remote Mirror to External Devices

The first step in configuring a Remote Mirror to an external IDS, IPS, or UTM device is to set the destination port for the mirror. The external IDS, IPS, or UTM device must be connected to a port on a ProCurve device that is remote mirror capable. Make a note of the port where the external device is connected as you will reference this later on.

To assign the destination port in NI Manager (PCM+):

1. Select the device node in the navigation tree, or select the device in the Interconnect Devices list. (this is the switch where the IDS, IPS, or UTM is connected) Click the Port List tab to get to the Port Status sub-tab display.
2. In the Port Status table, select the Port you will use as the Remote Monitoring (destination) port. This is the port on the switch where the external IDS, IPS, or UTM device is connected
3. Select the Configure Mirror Port option from the toolbar pull-down menu.

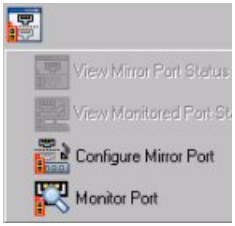

4. The Configure Mirror Port dialog displays, with the selected port ID.

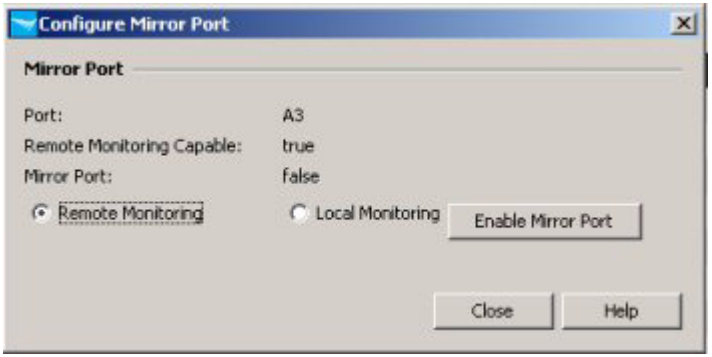

Figure 4-18. Configure Mirror Port dialog

5. Click the radio button to select Remote Monitoring.
6. Click the Enable Mirror Port button.

   This launches the Select Mirror Source dialog.

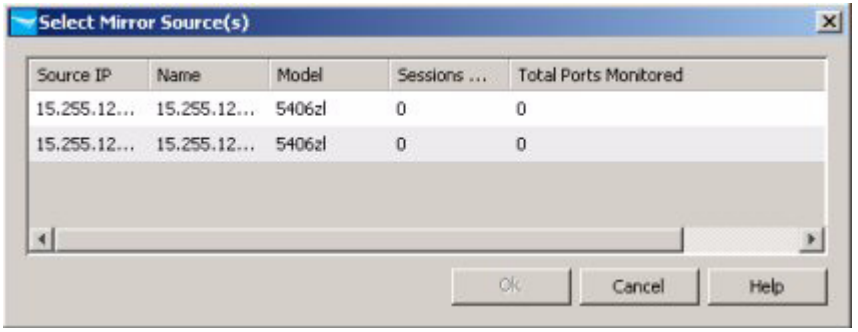

Figure 4-19. Configure Remote Mirror: Select Mirror Source dialog

7. The Select Mirror Source list is auto-populated with all devices that are remote mirror capable that have been discovered by PCM.

   Click to select a device, then click Ok to save the source and close the dialog. This returns you to the Configure Mirror Port dialog.

   The Mirror Port: option changes to true, and the button changes to Disable Mirror Port.

8. Click Close to save the mirror port setting.

Creating the Remote Mirror Action:

The next step is to create the action to mirror suspect traffic from the source port to the destination port.

1. Navigate to the Policy Manager, Create Actions window.

   a. Click the Policy Manager icon in the toolbar to launch the Policy Configuration Manager window.
   b. Click the Actions node in the Policy Manager window to display the Manage Actions panel.
   c. Click New... to launch the Create Action dialog:

2. In the Create Action dialog box, select the Port Mirroring action from the pull-down menu.
3. Enter a name for the Action (for example: UTM Mirror) and then click OK.

   The Policy Manager displays the Port Mirror action configuration tabs.

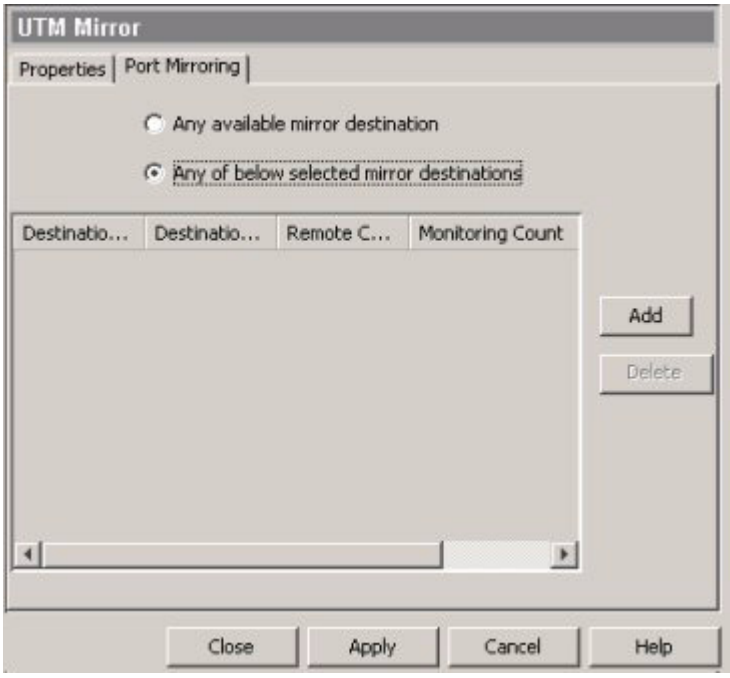

Figure 4-20. Policy Action: Port Mirroring configuration tab

4. Click the Port Mirroring tab and select the "Any of the below selected mirror destinations".
5. Click Add to display the Select Mirror Port dialog.

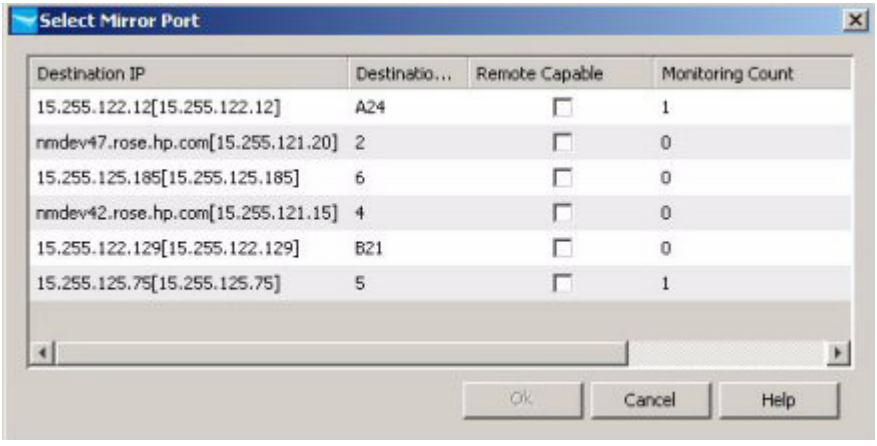

Figure 4-21. Port Mirror Action: Select Mirror Port dialog.

6. Select the (remote) mirror port you configured earlier, then click OK. (This is the port connected to the external IDS, IPS, or UTM device)

   The Select Mirror Port dialog is auto-populated with information on all configured mirror ports that have been discovered. Ports capable of remote mirroring are indicated with a check mark in the Remote Capable column.

   Note if you have multiple IDS, IPS, or UTM devices configured, you can select more than on and NI M will automatically load balance

7. The selected port appears in the mirror destinations list on the Port Mirror configuration tab in the Policy Manager window.
   (the window may have been minimized, just reselect it from the Windows taskbar to continue)

8. Click Apply, and then Close to save the Action configuration.

Configure a Policy to send traffic to the external device:

Now that the remote mirror is configured, you can create a policy to forward suspicious traffic on to the remote IDS, IPS, or UTM device for additional analysis.

1. Select the Policies node in the Policy Manager navigation tree to display the Manage Policies panel, then click New... to launch the Create Policy dialog.

2. In the Sources tab, select the groups for devices capable of remote mirror (3500yl, 5400zl, 6200yl) and move them to the Selected Groups list.

3. In the Targets tab, select the "Target all alert sources (devices & ports) that trigger this policy" option

4. In the Alerts tab, select any or all of the available "security" alerts and move them to the Selected Alerts list.
   (Default DNS Tunneling, Default Protocol Anomaly, etc., or ProCurve security alerts that you have configured).

5. In the Actions tab, select the Remote Mirror action you configured and move it to the Selected Actions list. Set the rollback time to rescind the remote mirror to the desired time (for example 3 minutes)

   You can also create an action to send a message or email and add that to the Selected Actions list. This will serve to immediately notify you when the remote mirror is started.

6. Click Apply to save the Policy.
   Click Close to exit the Policy Manager

Create a Policy to respond to External traps:

The final step in the process is to configure a policy to receive traps from the external IDS, IPS, or UTM device, and if needed provide mitigation of an attack.

First you need to create an External security alert to use in the policy.

1. Navigate to the External alerts configuration window
   - In the Policy Manager navigation tree, click to expand the Alerts ->Security node in the navigation tree,
   - Select External to display the Manage External Alerts pane.
   - Click New to launch the Create Alerts dialog
2. Select the Security: External alert type from the pull-down menu. Type in a name for the alert (for example, UTM Alert), and then click OK to display the Alert configuration tabs.
3. Click the Configuration tab to display the external alert configuration parameters.

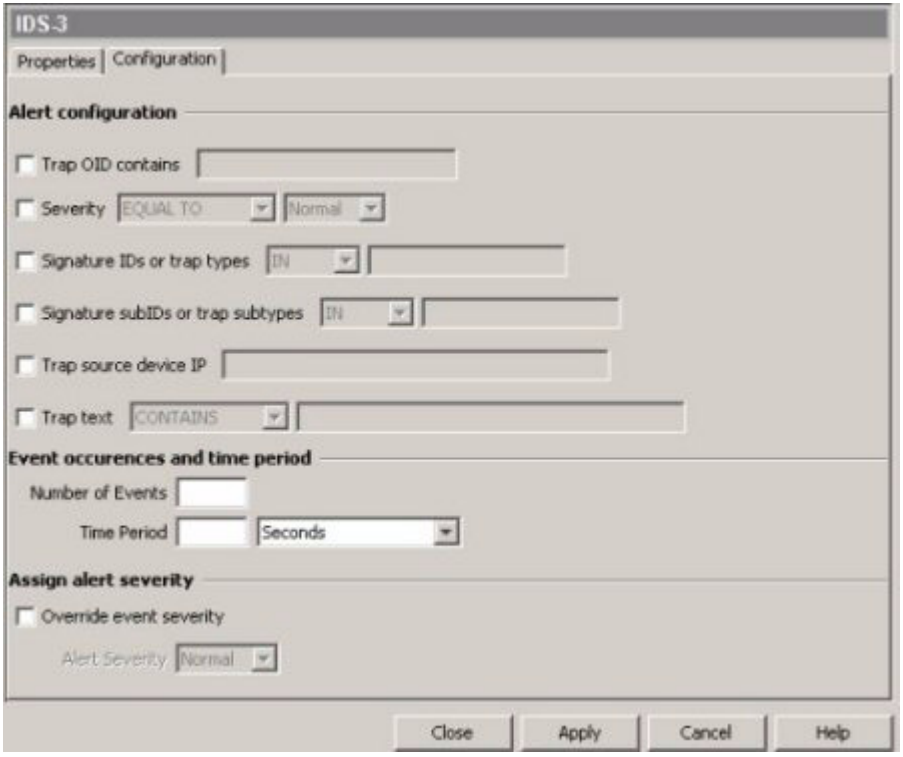

Figure 4-22. Alerts, Security: external alert configuration.

4. Select or enter the Alert Configuration parameters that will be applied. These parameters will be matched to text in the trap being sent to PCM from the external IDS, IPS, or UTM device.
   - To create an alert based on the OID contained in the trap, check the Trap OID contains check box and type any portion of the OID.

   You can find the Trap OIDs in the .trp file for the device, which is stored on the PCM server (<*installdir*>server\config\TrapEventConfig). These are text files that can be read using Notepad. The OID is the number at the top of the file, for example 1_3_6_1_4_1_9_9_383_0_1. You will need to change the underscores to periods (1.3.6.1.4.1 etc.)

   For additional information on creating .trp files for non-ProCurve devices, refer to the *ProCurve Manager (v2.2) Network Administrator's Guide*, chapter titled ""Using the PCM+ Configurable Integration Platform.
   - To create an alert based on the severity included in the trap, check the Severity check box, and then select the qualifier and severity.
   - To create an alert based on the signature or trap type contained in the trap text, check the Signature IDs or trap types check box, select IN or NOT IN, and type the ID or type.
   - To create an alert based on the signature subID or trap subtype contained in the trap text, check the Signature subIDs or trap sub-types check box, select IN or NOT IN, and type the ID or type.
   - To create an alert based on the IP address of the device generating the SNMP trap, check the Trap source device IP check box and type the IP address of the external IDS, IPS, or UTM device.
   - To create an alert based on any other text included in the trap message, check the Trap text check box, select the qualifier, and type the text to match.
   - In the Number of Events field, type the number of events that must occur before an alert is issued.

5. In the Event occurrences and time period field, identify the window of time used to count the minimum number of traps that must be received before an alert is issued. For example, you can issue an alert when more than two traps are received within 5 minutes.

   Note that you should configure the Alert time window for the policy to be larger than the "Rollback Actions" time, or you may create a loop state between events and mitigating actions.

6. To change the default severity assigned to the alert, check the Override event severity check box and use the Alert Severity drop-down arrow to select the desired severity to apply to the alert.

7. Click Apply, then Close to save the Alert and exit the window.

Now create the response or mitigation Actions that you want to apply when the external alert trigger is received. Refer to "Configuring Actions for Security Policies" on page 4-19. To start, ProCurve recommends you create an action that will send an email or pop-up message. Once you are confident you understand how the Policy works, you can add Actions that will disable the port, or block the attackers MAC address on a device.

1. Click the Policy Manager icon in the toolbar to launch the Policy Configuration Manager window.

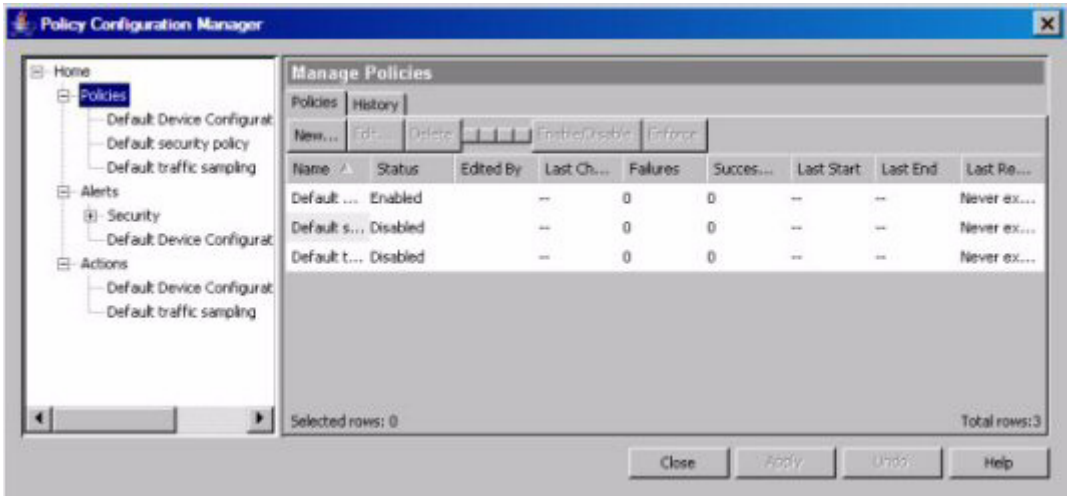

Figure 4-23. Policy Configuration Manager: Manage Policies panel

2. Select the Policies node in the navigation tree to display the Manage Policies panel, then click New... to launch the Create Policy dialog.
3. Enter the Name for the Policy (for example, External Alert Response) in the Create Policy dialog and click OK.
4. Complete the External Alert Response configuration by entering the necessary information in the Targets, Alerts, and Actions tabs.
   a. In the Targets tab, select the "Target all alert sources (devices & ports) that trigger this policy" option
   b. In the Alerts tab, select the "External Alert" you created from the Available Alerts list, and then click >> to move it to Selected Alerts list.
   c. In the Actions tab, move the response or mitigation actions you want to use to the Selected Actions list. These are the policy actions that you configure.
5. Click Apply and Close to save the policy and exit the Policy Manager window.

Index