


ProCurve Networking

Automating Network Defense with PCM+ and NIM



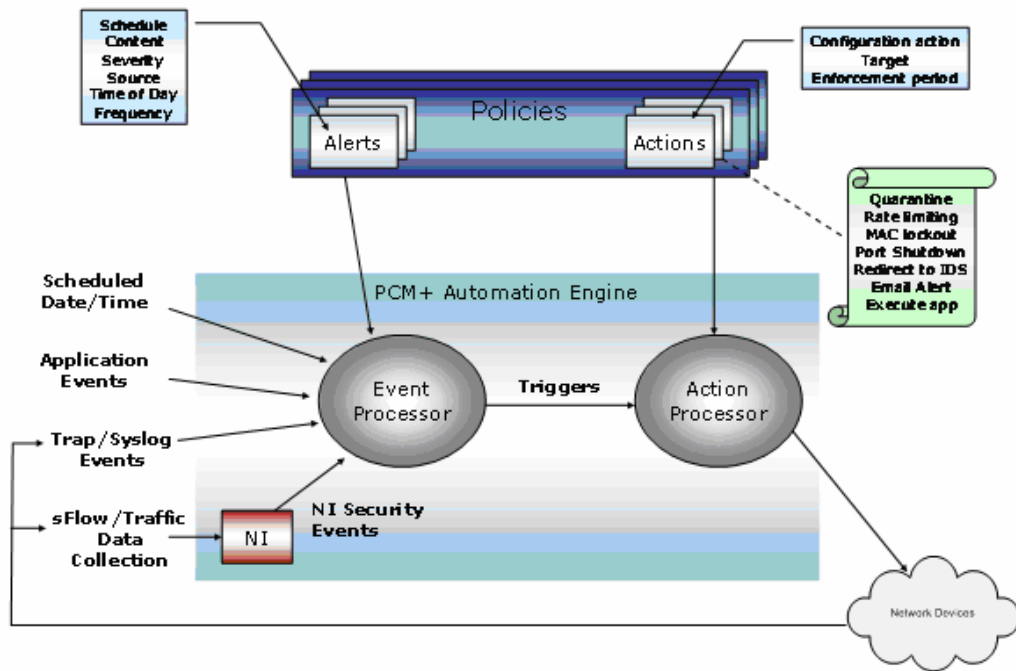
Technical Brief

Introduction	3
Policy Manager	3
Terminology	3
Enhancements	4
Policy Manager Usage	4
Policy Manager Dialog	5
Schedule-Driven Policies	5
Scheduled configuration scans	6
Defining the alert	6
Defining the action	7
Defining the policy	8
Scheduled port provisioning	9
Defining the alert	9
Defining the action	9
Defining the policy	10
Schedule the execution of any report	11
Defining the alert	11
Defining the action	11
Defining the policy	12
General Event-Driven Policies	12
Notification	13
Defining the alert	13
Defining the action	15
Defining the Policy	16
Automated Response	16
Security Policies	17
Security Alerts (ProCurve and External)	17
Security actions	18
Logical source and target	18
Security GUI Usage	18
Alerts GUI	19
Actions GUI	19

Introduction

In today's IT environment, network administrators face the difficult task of performing reactive activities by hand. The purpose of this paper is to enlighten the reader about possibilities of automating network responses in their ProCurve Manager Plus (PCM+) and Network Immunity Manager (NIM) environments. For up-to-date product configuration information and advanced features, please view product manuals at <http://www.hp.com.rnd/support/manuals/index.htm>.

Policy Manager



The diagram above depicts the high-level structure of the automation system. As shown, the system consists of user- and system-defined policies that are comprised of alerts and actions. As incoming events arrive, the automation engine checks those events against the list of enabled active alerts. Should an alert be triggered, all actions called for by policies associated with that alert are taken.

Terminology

The following definitions are provided to help the reader become familiar with terms used throughout this paper:

Policy – A set of one or more actions to be taken (i.e., enforced), either at a scheduled time (with optional recurrence rate), in response to an event (SNMP trap, syslog or PCM/NIM/IDM-generated events), or when explicitly enforced by the network administrator.

Alert – A time or series of events of interest to the network administrator. An alert is used to automatically trigger any policies appropriate to the time or event being monitored.

Action – An operation to be performed by PCM+ when called upon by a policy.

Target – Any device to be acted upon by an action.

Physical Source – In the case of an event-based alert, the physical source is the network device that sent the event that triggered the alert.

Logical Source – In the case of some event-based alerts, the network device of interest is not the physical source of the event, but rather, the device identified by the contents of the event.

For example, when a virus throttling (VT) trap is received, the physical source is the device that sent the trap, but the logical source is the edge port to which the offending host is connected. This is a useful distinction when defining an alert and the actions to be taken in response to that alert (e.g., alter the VT sensitivity on the physical source, as opposed to disabling the edge port identified as the logical source of the alert).

Rollback – Refers to the process of “undoing” any configuration changes made by certain actions that support rollback after a user-defined time period has elapsed since the action first went into effect.

Enhancements

While PCM+ always has included facilities allowing various levels of automated network response, this functionality has been enhanced greatly with the PCM+ 2.2 release. The enhancements provided by this release provide the network administrator with far more flexibility and control, including:

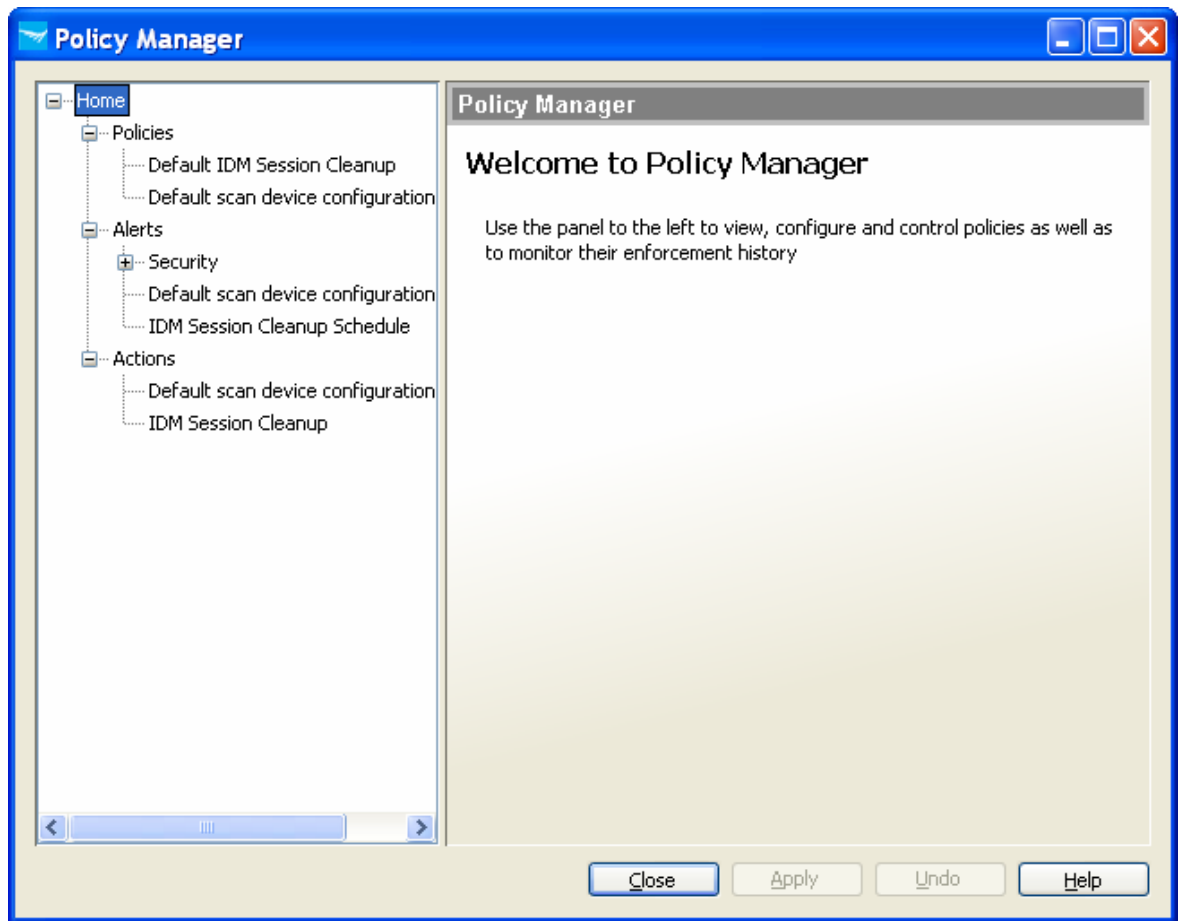
- The ability to reuse previously configured alerts and actions, speeding up the process of configuring new policies
- The ability to restrict the times and days when a policy is in effect, allowing reactive policies to be put in place specific to the needs and expectations of the network during any given time
- The ability to take more than one action when a policy is enforced, reducing the number of policies that need to be defined and put in place to react to any given event
- The ability to restrict the activation of an event-driven policy based on the groups to which the source device(s) belong
- The ability to rollback certain action types, allowing a temporary response to dynamic events
- The ability to view the areas of the network generating alerts, as well as what actions were taken as a result of those alerts

The following section describes the user interfaces used to define, monitor and control alerts, actions and policies.

Policy Manager Usage

While users of previous releases of Policy Manager may be familiar with many of the concepts discussed here, much of the user interface has been enhanced. This section briefly describes the key user interfaces used. For more information, please consult the PCM+ Administrator’s Guide and online help provided with your copy of PCM+.

Policy Manager Dialog



The Policy Manager Dialog is the main interface used to define, control and monitor active policies at a global level. This dialog can be displayed by pressing the “Policy Manager” button on the global toolbar or by selecting “Policy Manager” in the tools menu.

As shown above, this dialog displays a tree structure containing all currently defined policies, alerts and actions. By selecting items on the tree, the user can create new elements of the selected type and also, view and edit the specific properties of any preexisting element.

The “Policy” node of the tree can be used to display a table showing all defined policies (whether enabled or not) and their current states. Also available at this level is a table showing the history of all alerts observed (up to a limit defined in the “Policy Manager” portion of the PCM preferences panel), what actions were taken in response to those events, and the state of those actions. By selecting a row on this table, the user can see a detailed description of the alert and the action selected.

The contents of the “Alerts” and “Actions” portions of the tree can be considered to be the reusable tools in the network administrator’s toolbox. These items can be reused in numerous different policies (e.g., monitoring different locations on the network for the same alert, but taking different actions specific to that location).

When beginning to define a new policy, the administrator can take either of two approaches. In the first approach, the necessary alerts and actions can be created; then, the policy is defined and associated with those alerts and actions. In the second approach, the policy can be defined and any necessary alerts and actions can be created in the course of defining the policy.

The following sections give examples of schedule- and event-driven policies. These are followed by examples of more advanced security-oriented policies.

Schedule-Driven Policies

The following example demonstrates the steps used to create a policy that is used to scan the configuration of one or more devices on a regular schedule, freeing the administrator from

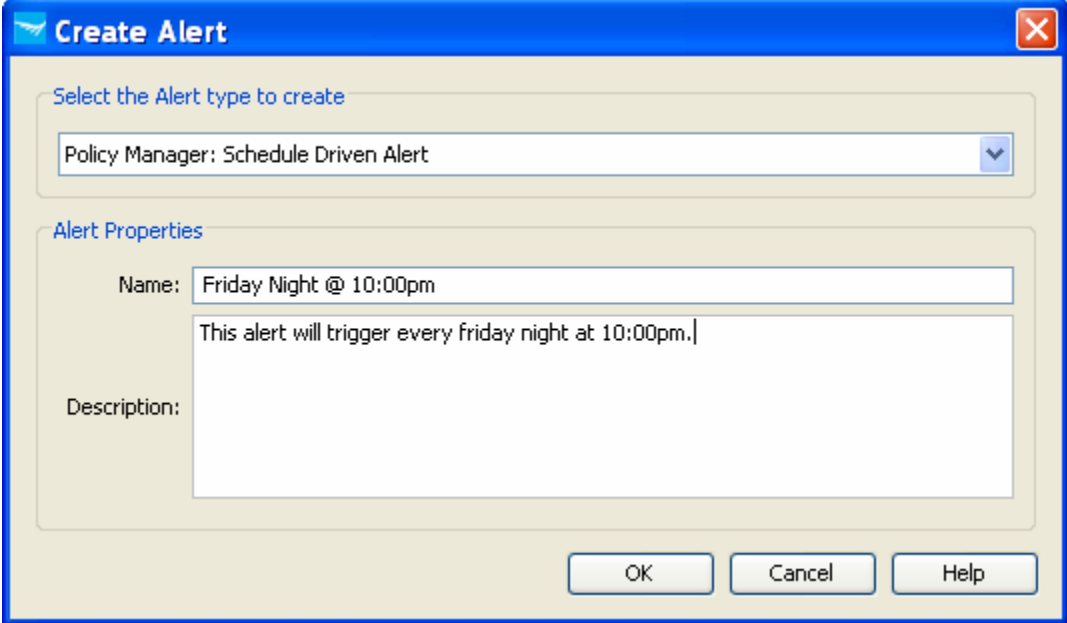
needing to manually invoke the scan. This policy can be useful for archiving device configurations as well as detecting undesired or malicious changes to the network infrastructure.

Scheduled configuration scans

The policy being created to perform a scheduled configuration scan will consist of one alert and one action. In this example, we will begin by defining the alert and action individually. Then, a policy will be created that utilizes these alerts and actions.

Defining the alert

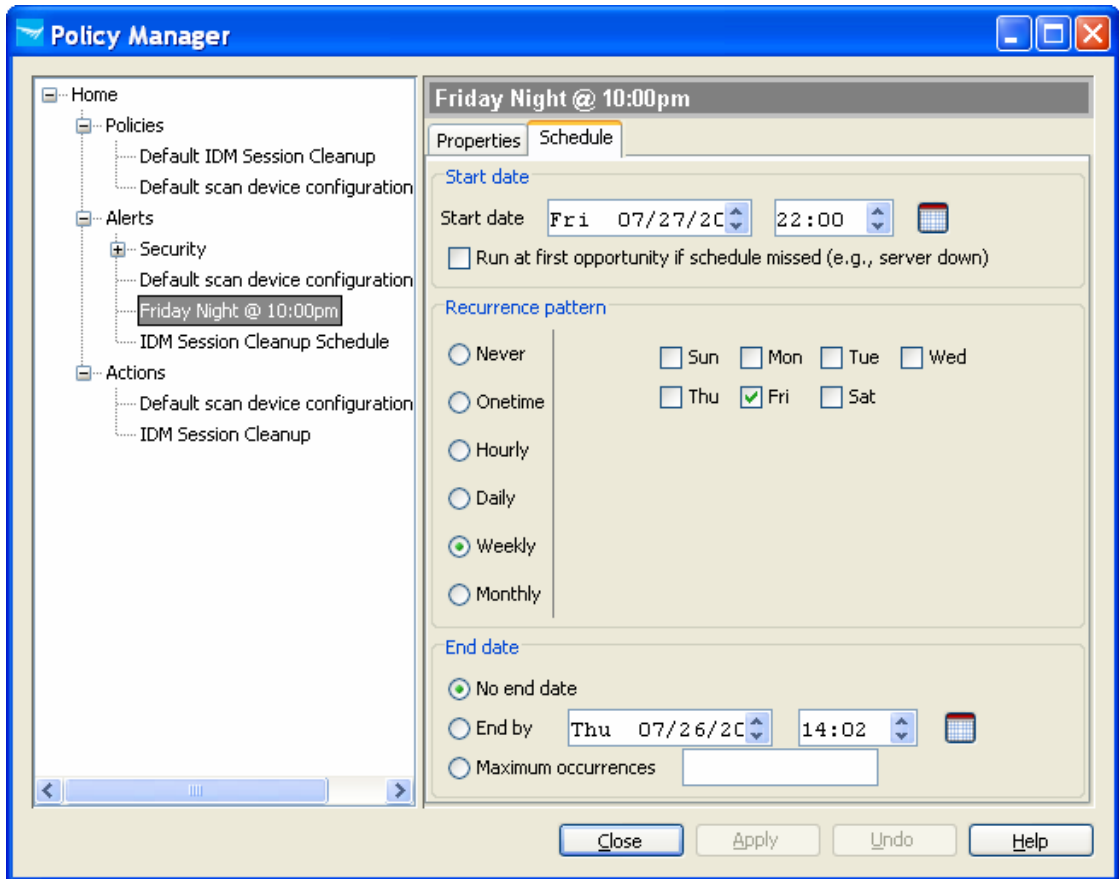
Begin by opening the policy dialog and selecting the "Alerts" level of the tree. The table displayed will show the current defined alerts. Press the "New" button above this table to begin defining a new alert.



The screenshot shows a "Create Alert" dialog box. The title bar is blue and contains the text "Create Alert" and a close button. Below the title bar, there is a section titled "Select the Alert type to create" with a dropdown menu showing "Policy Manager: Schedule Driven Alert". Below that is the "Alert Properties" section, which includes a "Name" field containing "Friday Night @ 10:00pm" and a "Description" text area containing "This alert will trigger every friday night at 10:00pm.". At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

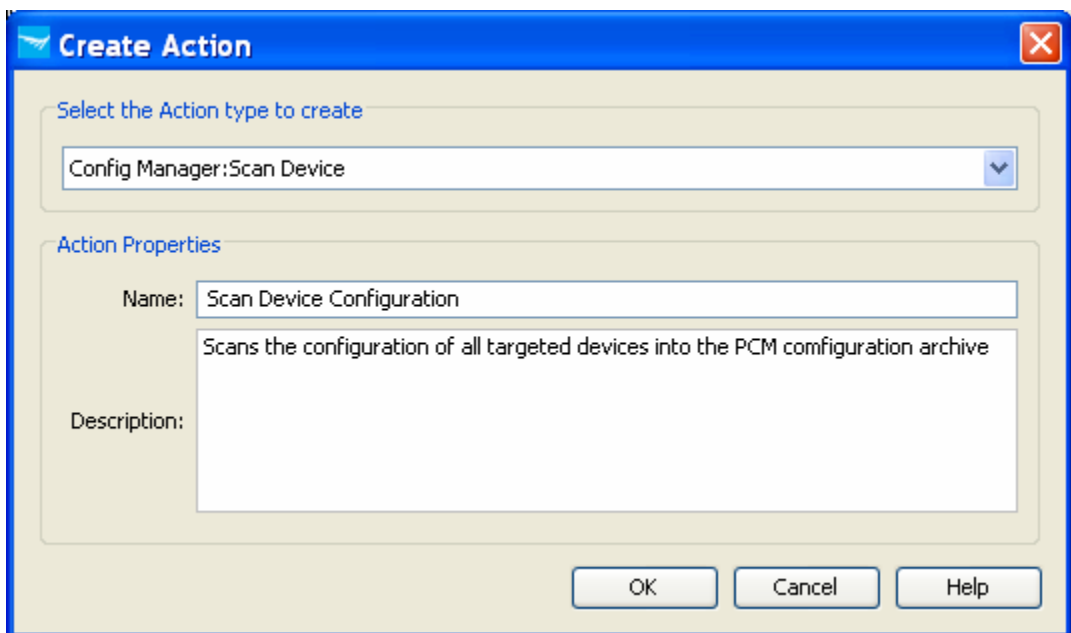
As shown above, select the alert type, "Policy Manager: Schedule Driven Alert". Provide a name (required) for this alert so it can be distinguished easily from others when you are configuring policies. An optional description may be provided to help track the purpose of this alert.

Upon pressing the "OK" button, a new alert with the given name will be entered into the tree and selected for you. The specific schedule for this alert still needs to be provided, so select the "Schedule" tab on the right-hand panel. In the fields of that tab, provide the desired schedule as demonstrated below. Once finished, press the "Apply" button to save these settings to PCM.

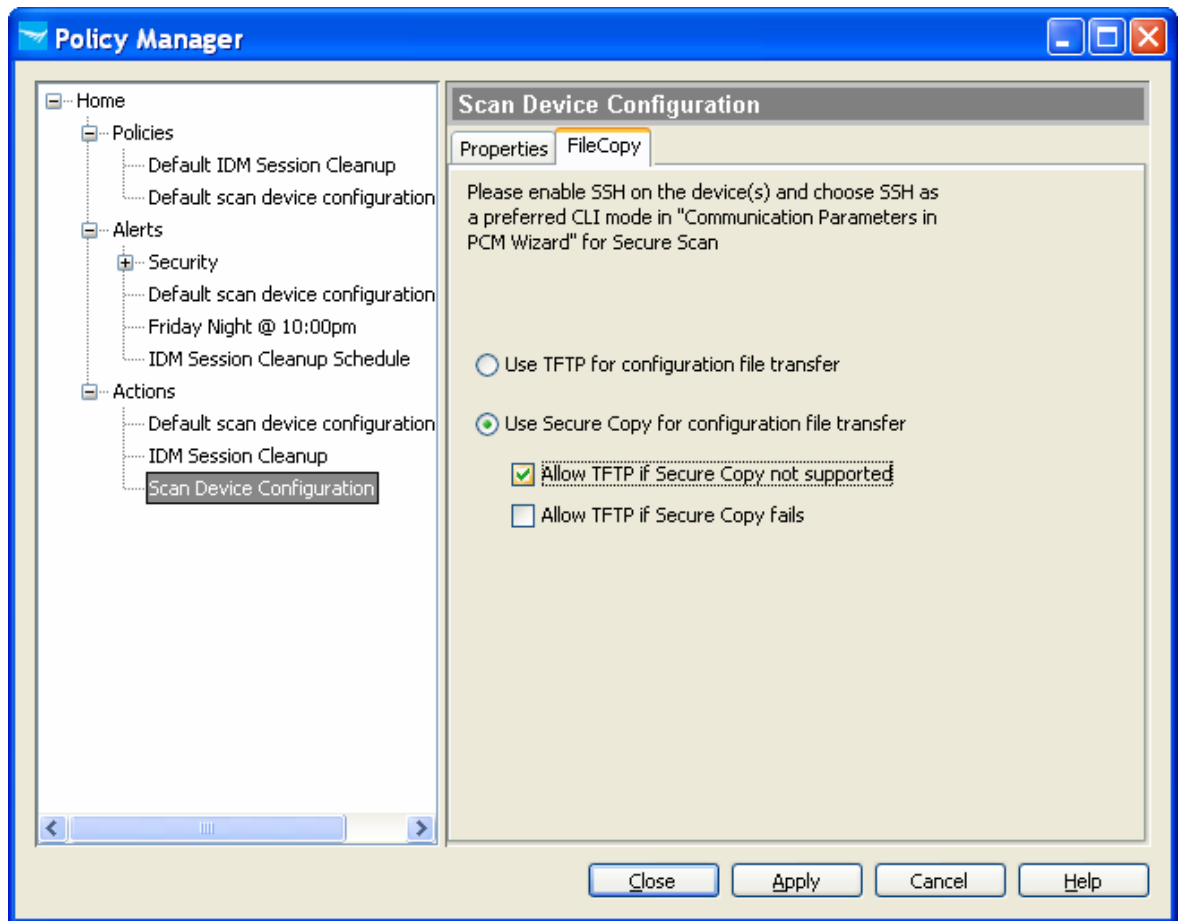


Defining the action

Now that the schedule is defined, we'll move on to defining the action. To do this, select the "Actions" node of the tree and select "New" on the table presented. As before, select the desired item type, in this case, "Config Manager: Scan Device" (as shown below) and provide a name and optional description.




As with the alert, a new action will be created and selected in the tree. Check the "File Copy" tab of the scan action to ensure that the appropriate settings are made for your network infrastructure (see below for an example). Press "Apply" to save these settings to PCM.



Defining the policy

Now that we have the necessary “tools” in our toolbox (that is, the alert and action to be used by the policy), we can define the policy. Begin this process by selecting the “Policies” node of the tree and pressing the “New” button on the displayed table.

In the displayed window, provide a name and optional description for the policy and indicate whether, once defined, it should be enabled automatically.



Once the policy is created and added to the tree, you will need to visit each tab to ensure it is configured as desired. The tabs define the following properties:

- **Times** – This setting allows the network administrator to define what days and times the policy is allowed to be enforced automatically. It is used primarily for event-based policies that the user may wish to restrict. In the case of this schedule-based policy, we will leave the default setting of “any time,” thereby placing no restrictions on the policy, as we’ll rely on the schedule defined in the alert. Note however that you could use this tab even in the case of a schedule policy should you wish to restrict the execution based on company holidays or other scheduled dates.
- **Sources** – This tab allows the administrator to define what event locations on the network an event must originate from to be a valid trigger for alerts associated with the policy. In this case, the only alert used is a schedule alert, so this tab has no impact.
- **Targets** – This tab defines the devices to be acted upon when this policy is enforced. For the purpose of this example, we have a single group defined: “Main campus.” This contains all devices to be scanned. (You may provide as many groups as needed.)
- **Alerts** – This tab defines which predefined alerts are to cause the enforcement of this policy when triggered. In this case, we want to select the “Friday Night @ 10:00pm” alert that was just defined and move it over to the “Selected Alerts” list.
- **Actions** – This tab defines the actions to be taken when the alert is generated and the policy enforced. For this example, select the “Scan Device Configuration” action that was just defined and move it over to the “Selected Action” list.

The policy is now fully defined. Press the “Apply” button to save this policy and add it to the list of active policies.

You now have a policy in place that will scan the configuration of all devices in the selected group once per week at the scheduled time. To see the policy in action, you can select the “Policies” node, select the desired policy in the table, and press the “Enforce” button.

Scheduled port provisioning

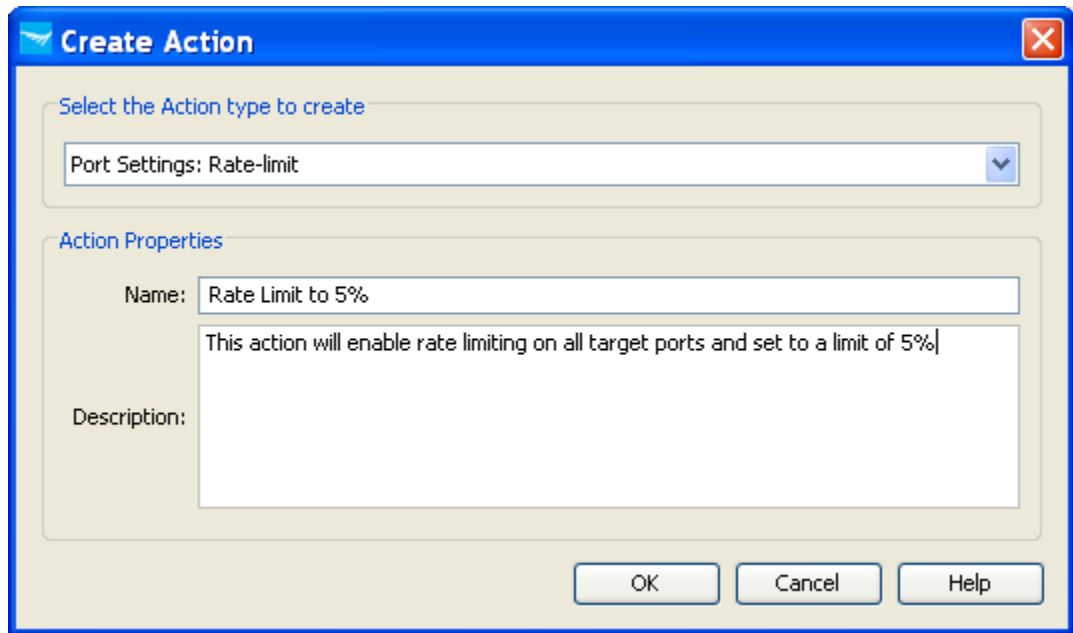
With this example, we will demonstrate how a policy can be defined that will automate the work required to enable and set rate limiting for a set of edge ports at a specific time and revert that setting back later in the day (e.g., to ensure that students in the dorm are not busy on the network when they should be in class).

Defining the alert

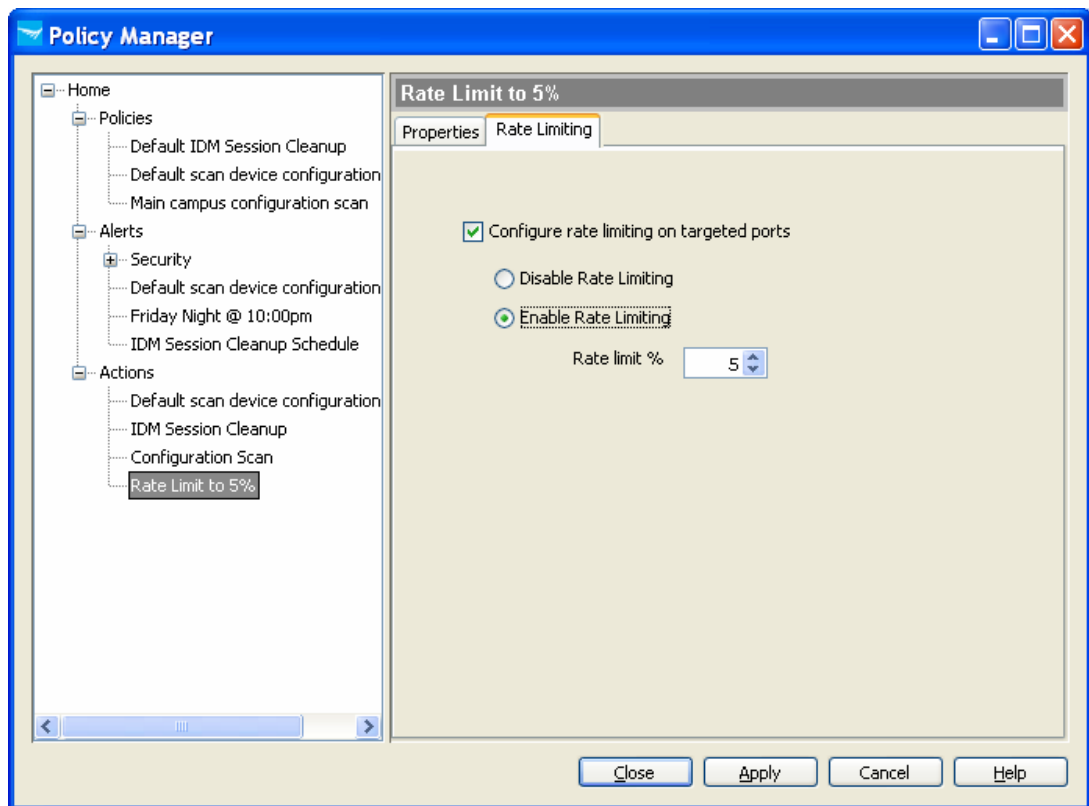
For this example, we will use a schedule-based alert set to trigger every weekday at 10:00 a.m. Use an approach similar to that described in the preceding example to specify the schedule to be used by this alert. Note that this will trigger the enforcement of the policy at that time, and the rollback of the rate limit setting will be accomplished by the settings made during configuration of the policy.

Defining the action

Next, we will define an action that enables rate limiting and sets the rate limit for all targeted ports to 5 percent. Begin by selecting the “Actions” node of the tree and pressing the “New” button on the displayed table. Then, select the action type, “Port Settings: Rate-limit” from the list of available actions. Provide the requisite name and optional description as shown below.



Once the action is created, select the “Rate Limiting” tab of the displayed panel and configure as shown below. Press “Apply” to save these settings to PCM.

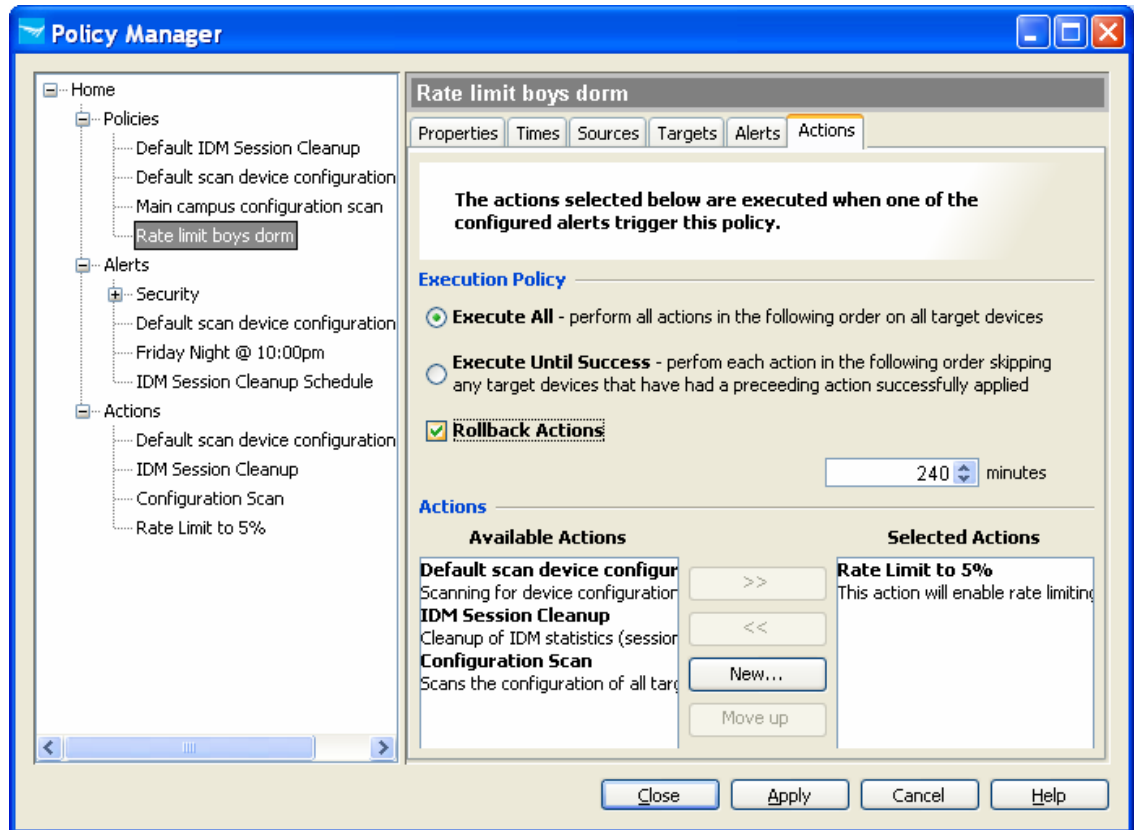


Defining the policy

As with the previous example, we will now associate the alert and action with the policy. Begin by selecting the “Policies” node, pressing the “New” button, and providing the desired name and description.

As before, review the settings of each tab to ensure that it is configured as desired. The only differences between this policy and the previous schedule-based example will be on the “Targets” tab and the “Actions” tab.

On the targets tab, select the group previously defined by the network administrator to contain the ports to be rate-limited during class hours.



As shown above, on the “Actions” tab you now will configure the “Rollback Actions” setting to a value of 240 minutes (4 hours), assuming that the rate-limit policy is only to be in effect from 10:00 a.m. to 2:00 p.m. After that time elapses, the rate-limit setting will be reverted back to whatever value was configured on each target port before the policy was enforced. Note that rollback is supported only by specific action types. Please consult your PCM+ documentation for specifics.

Upon pressing the “Apply” button, this policy will be saved to PCM and enabled for enforcement at the next scheduled time. Now, you have automated the provisioning of your network edge based on the specific needs of your users/customers.

Schedule the execution of any report

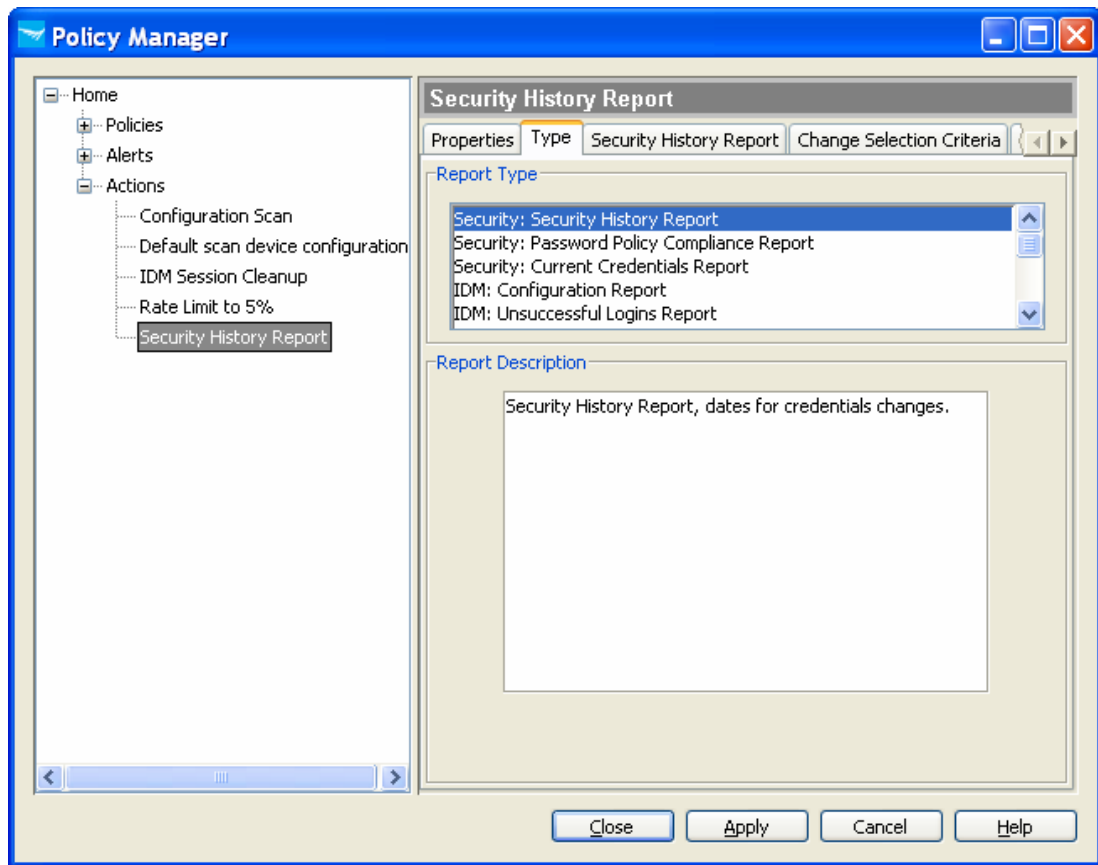
This last example of a schedule-based policy demonstrates how to go about automating the generation of any of the report types supported by PCM and its plug-in products. This can be useful for ensuring that a report is ready for you and your team when needed.

Defining the alert

For the purposes of this example, we will reuse either of the previously defined schedule-based alerts. This is intended to demonstrate the reusability of the alert and action items, in that this same alert can be used to trigger more than one policy. This frees you from needing to define identical alerts should the desired properties exist already.

Defining the action

This example will utilize the “Report Manager: Generate Report” action type. After creating a new action of this type, the panels displayed will allow you to specify the desired report type and any necessary criteria. For the purpose of this example, we will use one of the report types provided with NIM, though any report type provided with your installation can be used.



As shown above, on the action's "Type" tab, select the report type, "Security: Security History Report." This will cause new tabs to be displayed that are specific to that report type.

Security History Report – Please configure this tab with the device group for which you would like the report generated. This can be a device family-based group provided by PCM or any custom group you have configured.

Change Selection Criteria – Please select whether the report should track devices with configurations that have changed in the given time frame, or those for which credentials have not changed.

Format – Please select the file format in which you would like the report generated: PDF, HTML or CSV.

Delivery – Please choose the way you would like the report to be delivered: FTP, file saved to local disk or FTP.

Defining the policy

With the action defined, the policy now can be created. Follow the same procedure described in the previous examples and select either of the previously configured alerts to be the trigger for this policy. Select the report generation action you just created and press "Apply."

You now have a policy in place that will generate a security credential change report at the given schedule to be delivered, as specified.

General Event-Driven Policies

In addition to schedule-based policies, PCM+ provides the powerful ability to have policies enforced in response to new events. These events can be SNMP traps, syslog messages or application events (i.e., notices generated by PCM and its plug-in products). As the following examples will show, the process for creating these event-driven policies is identical to that previously used for schedule-based policies, with the one difference being the alert type used.

Notification

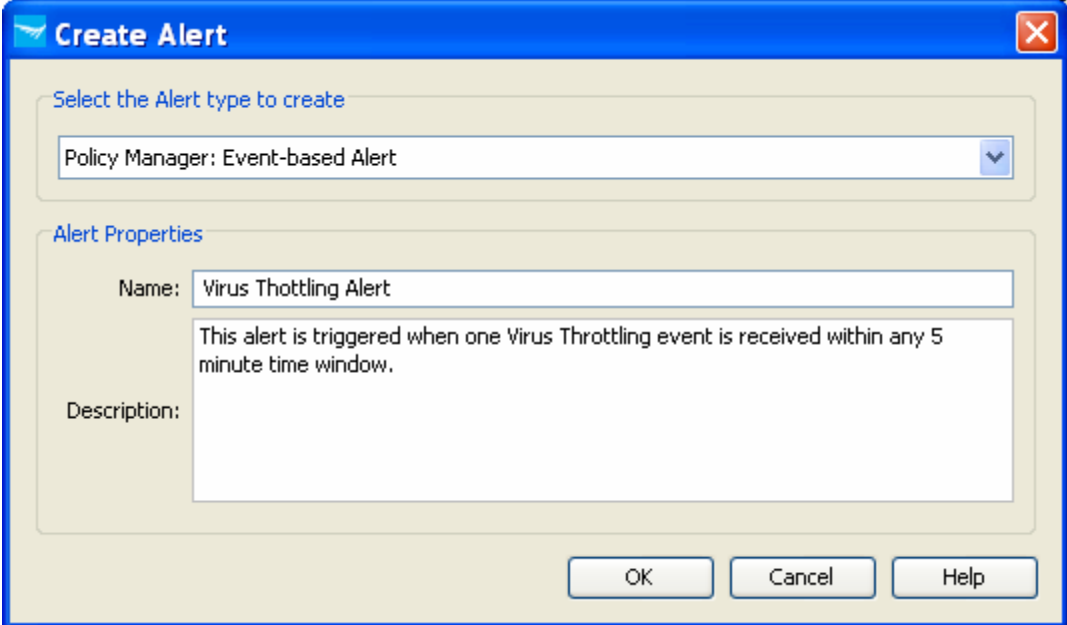
One powerful feature provided by event-based policies is the ability to have a notification generated to inform you or the appropriate contact that an event of interest has occurred. PCM+ provides a variety of mechanisms by which a notification can be generated:

- SNMP trap – generates an SNMP trap to be sent to any SNMP trap receiver on your network
- E-mail – sends an e-mail via SMTP
- Execute Command – executes any user-defined system command/script (e.g., generate IM message, send page, etc.)

As policies can contain more than one action, any or all of these notification types can be configured and attached to any policy of interest.

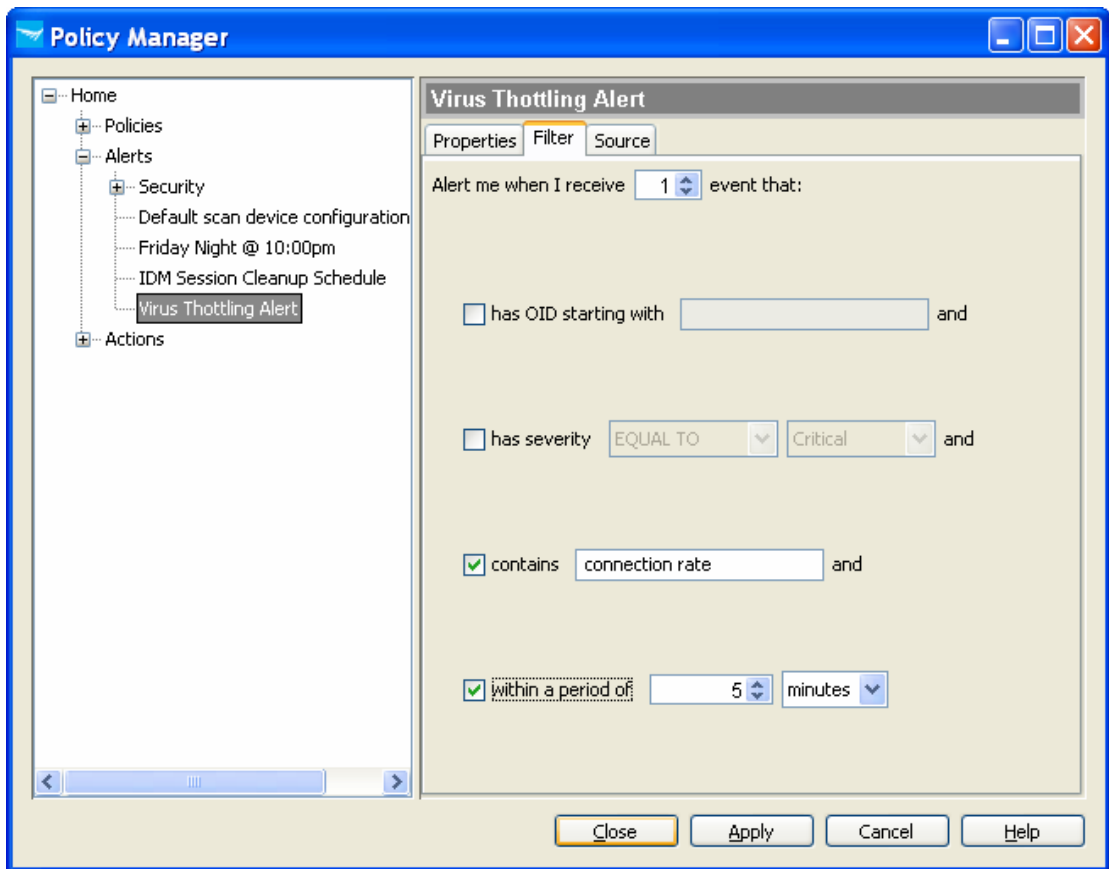
Defining the alert

For the purpose of this example, we will create an event-based alert looking for VT events on the network. To do this, begin by selecting the “New” button on the “Alerts” table and selecting the alert type, “Policy Manager: Event-based Alert.” As shown below, provide a name and optional description for this alert and press the “OK” button.



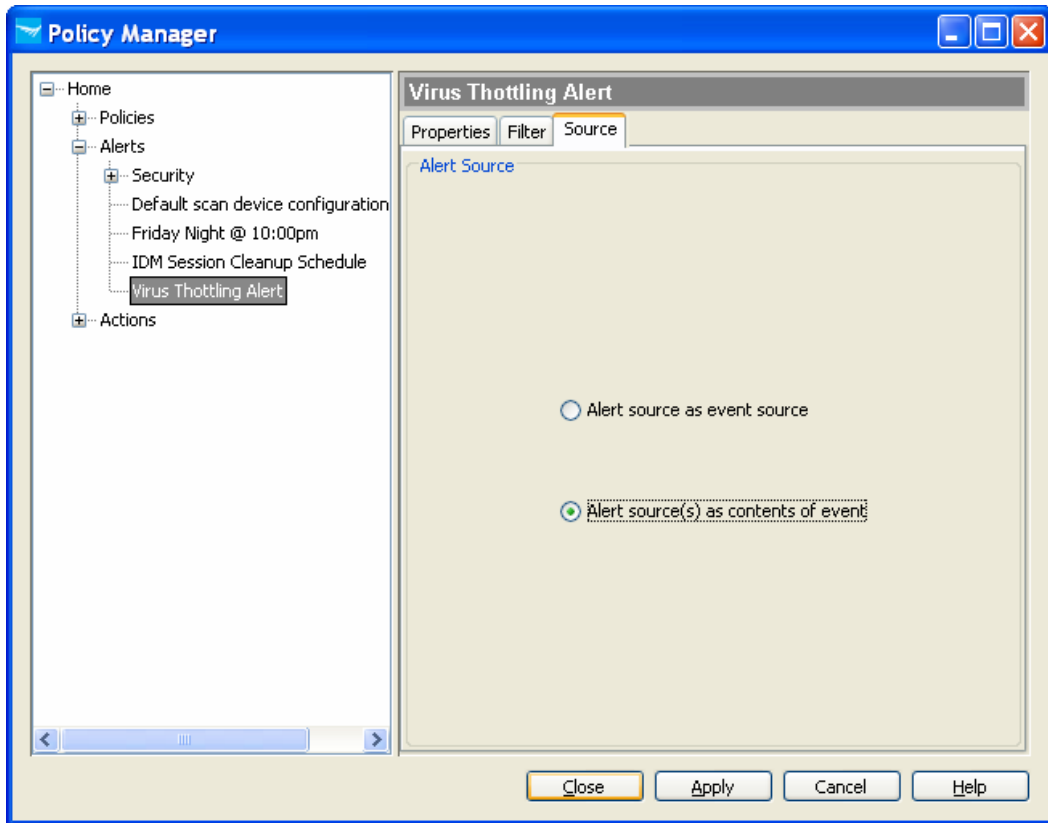
The screenshot shows a "Create Alert" dialog box. The title bar is blue and contains the text "Create Alert" and a close button. The main area is light beige. It has two sections: "Select the Alert type to create" with a dropdown menu showing "Policy Manager: Event-based Alert", and "Alert Properties" with a "Name" field containing "Virus Throttling Alert" and a "Description" field containing "This alert is triggered when one Virus Throttling event is received within any 5 minute time window." At the bottom are "OK", "Cancel", and "Help" buttons.

Once the alert is placed in the tree, the specific criteria for the events this alert will be triggered by can be provided. As shown below, this alert is configured to trigger when one or more events arrive that contain the text, “connection rate.” A further criteria is specified to restrict this alert from firing more than once every five minutes; this can help to prevent too great a number of notifications should a high number of these events arrive at any given time.



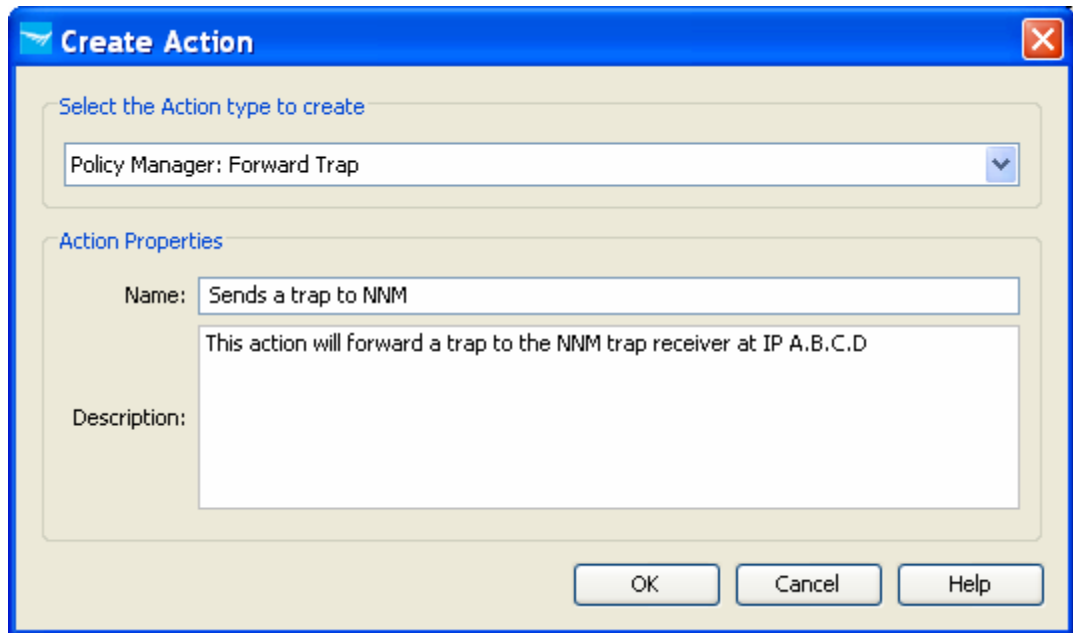
The next step is to define what PCM should treat as the source of the event. In many cases, the device that sent the event is the only reasonable option. However, in other cases, such as the VT event, some scenarios may require that the source be the device that sent the trap (e.g., the VT-capable network device that identified the problem host), while other uses may require that the source be the device(s) identified by the contents of the triggering event (e.g., the edge port to which the host identified by a VT trap is connected).

As shown below, for this example we will consider the source of the alert to be the edge port to which the VT identified host is connected. This has a bearing on the way in which these alerts are counted and reported on the various PCM+ and NIM alert activity screens and reports. Please consult your PCM+ and NIM documentation for more details on viewing and reporting alert activity by source.

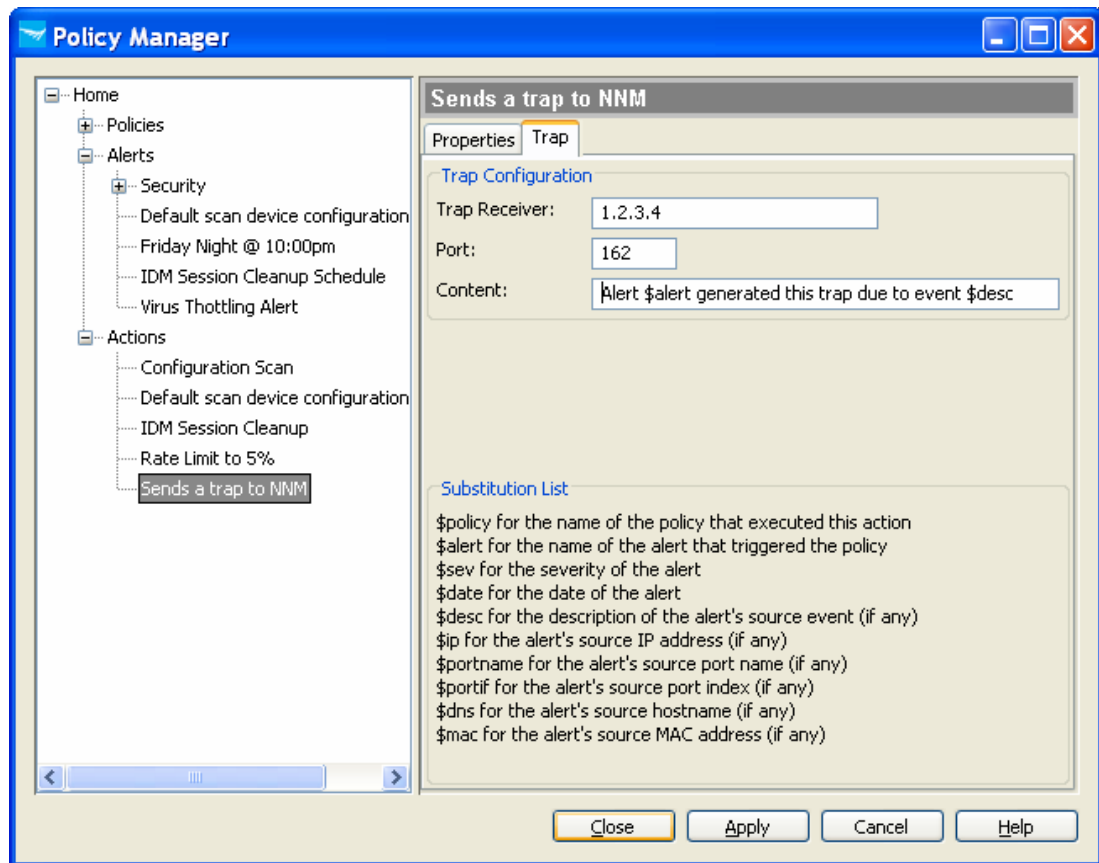


Defining the action

With the event-based alert in place, the action to be used to generate the desired notification now can be configured. Start by selecting the “Actions” node of the policy manager tree and pressing the “New” button. As shown below, select the action type, “Policy Manager: Forward Trap” and provide the name and description.



Then, the action should be configured to direct the trap to the appropriate receiver IP address and port. As shown below, this event can contain the text of your choosing, including the various substitution variables noted on the panel.



Defining the Policy

With the event-based alert and trap-forwarding actions defined, you now can create a new policy using the same steps described in the previous actions. When the alert is triggered by a new incoming VT event, the notification trap will be generated and sent the trap receiver defined in the action. Now, you have automated the integration between PCM+ and any other network management system to ensure that the appropriate teams are notified when a possible virus is detected on the network. Furthermore, the “Policy Activity” and “Security Activity” screens provided by PCM+ and NIM will give you a quick way to identify the areas on the network with the most VT activity.

Automated Response

PCM events can be logged internally from PCM+ or they can come in as SNMP traps from external devices such as ProCurve switches. Within PCM2.2, you have the ability to define thresholds on traffic and set triggers for PCM to create an event. This would enable an administrator to prevent oversubscription of network ports or even a Denial of Service (DoS) attempt against a network switch. This can be accomplished by setting thresholds on individual ports, such as interconnect switch links, and defining a policy that monitors the event log for these events. It is possible to have PCM take an action, such as rate limiting the port where the device generating traffic is connected.

The network response could be automated further by processing SNMP traps from ProCurve switches. There are many features in ProCurve switches that generate a trap when a condition is met. Most traps would need to be enabled through the CLI in the feature’s configuration area. The switch also will need to be configured to send SNMP traps to the management station.

```
ProCurve(config)# snmp-server host <IP_of_PCM> public all
```

The response can range from actions such as shutting down a port, locking out a device’s MAC address, or simply sending an e-mail notification to the network administrator. Below is a list of some supported SNMP traps from ProCurve switches:

SNMP Security Access Violations	Virus Throttling (VT)
Authentication Failures	Clear CLI Password
DHCP Snooping	Instrumentation Monitoring
Dynamic ARP Protection	Cold Start
Dynamic IP Lockdown	LLDP Topology Changes
Power over Ethernet (PoE) Thresholds	VRRP Transitions or Authentication Failures
MAC Lockdown Violation	OSPF Authentication Failures
MAC Lockout Violation	Link Status
RMON Alarms	STP bpdu Protection

Security Policies

Installing NIM on top of PCM+ provides an administrator the ability to create policies on security-based alerts that enable PCM to take action on an offender.

The security policies are built on the same automation framework discussed throughout this paper. Below is the typical flow for creating a security policy for NIM:

1. Create a "Security" alert to watch
2. Configure an action to take when policy is triggered
3. Define from where these logical "Sources" will be monitored
4. Define what logical targets will be executed
5. Apply the intended "Alerts" and "Actions" to a policy

Security Alerts (ProCurve and External)

NIM receives security alerts from two sources. The first is called "ProCurve" and consists of Network Based Anomaly Detection Protocol (NBAD) and VT events. NBAD is built into NIM, and VT is a HP proprietary feature that runs on select ProCurve switches. Currently, there are seven NBAD alerts from which to choose.

NBAD Engine	Description
IP Fanout	One IP sending traffic to many other IPs within a specified window of time.
TCP/UDP Fanout	One source IP communicating with many ports on a destination IP within a specified window of time.
Duplicate IP	One IP appearing from more than one MAC address appearing in the specified time window.
IP Spoof	One MAC address with more than one IP address appearing within the specified time window.
Packet Size Deviation	Occurs when the engine detects a statistically unusual change in the average size of sent and/or received packets.
Protocol Anomaly	Occurs when a host sends traffic containing unusual properties that would not normally be expected to occur on the network.
DNS Tunnel	Excessive traffic on port 53.

The second alert source is called "External" and this includes any SNMP trap received from an external device. More specifically, this could be an Intrusion Detection/Prevention System (IDPS) or Universal Threat Management (UTM) device.

Policies triggered by security alerts will post their results to the “Security Activity” screens to provide an administrator or manager a view of security activity over time. NIM gives a customer the ability to customize their own policies based on any SNMP trap PCM is capable of decoding.

Security actions

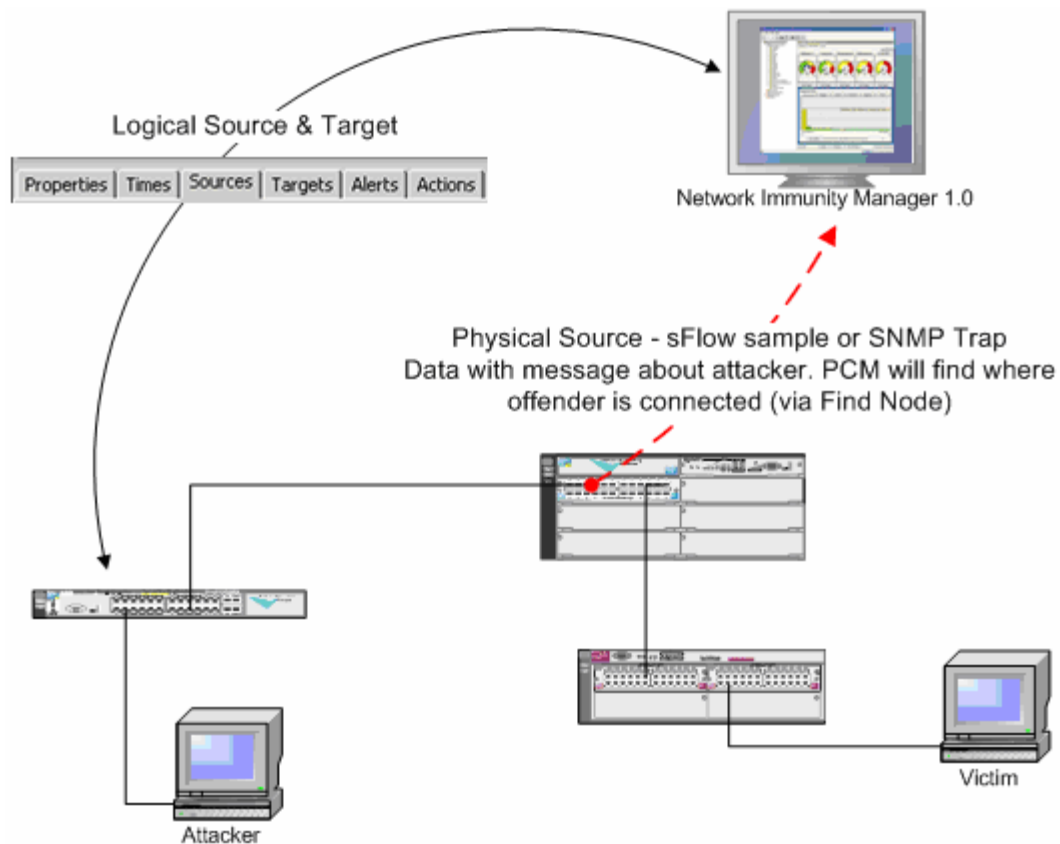
Security actions can range from locking out the MAC address of the offender, to quarantining them to a VLAN, or even shutting down the switch port where they are connected. It also can take less invasive actions, such as sending an e-mail to the administrator or mirroring malicious traffic to an IDPS system for further analysis.

Logical source and target

The physical source is the device that sends the sFlow packet or SNMP trap. A “logical source” is the device where the offender is connected. Therefore, it would make sense to use “Interconnect Devices” as your source, because that is the logical source where users and potential offenders may be connected. The sFlow packet can come from any managed port where that packet traverses.

When choosing your logical source for security policy, consider using any device to which the offender might be connected. When the security event is received, PCM will use “Find Node” to track down the port to which the offender is connected.

When choosing a target, you should choose select: “Target all alert sources (devices & ports) that trigger this policy.” This is the primary selection for mitigating an attack. For sources, choose any group to which the offender might be connected. Note that interconnect devices encompass all discovered and managed ProCurve switches.

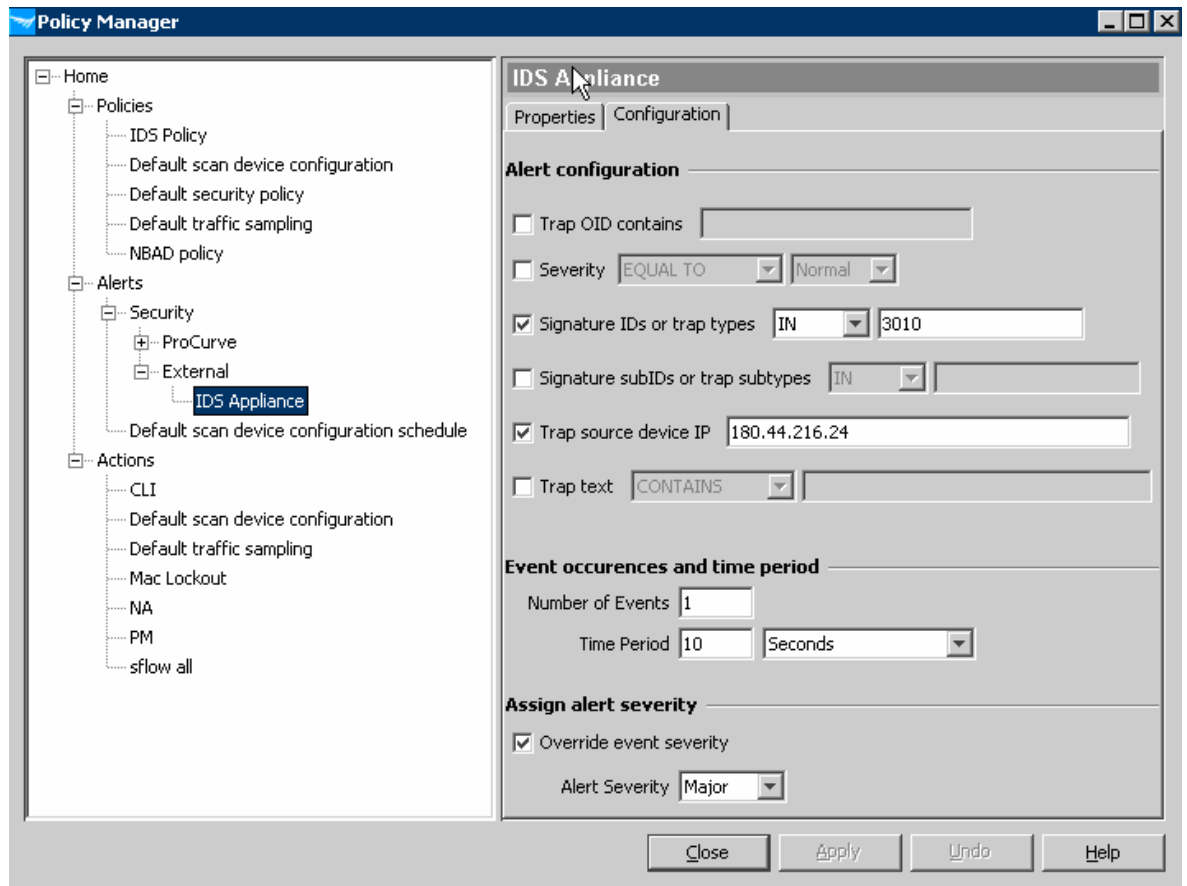


Security GUI Usage

The examples below are intended to walk an administrator through configuring a security policy to take action on an external event.

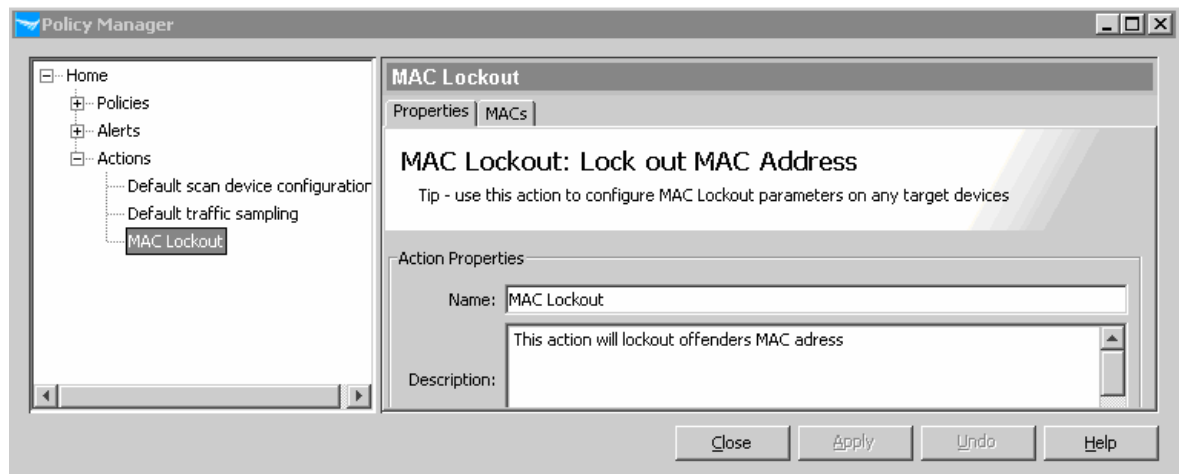
Alerts GUI

When NIM receives one trap on signature #3010 from IDS at physical IP address 180.44.216.24, it will trigger an alert. This alert, shown below, is designed to meet those criteria and also, is configured to override severity of the event to be “Major.”



Actions GUI

Once policy is triggered on the alert above, it will look for the offender IP in the trap. Next, it will call “Find Node” within PCM to obtain the MAC address of the offender and lock it out at the edge port of any logical source (interconnect devices). If the attacker information is not an event, the policy will not trigger. It is important to note that if the alert policy is set to trigger on one event in 10 seconds, this policy will fire only once per 10 seconds, at most.



To find out more about
ProCurve Networking
products and solutions,
visit our web site at

www.procurve.com



© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA1-5627ENW, 10/2007