

Network Immunity Manager Event Interpretation



Technical Brief

Introduction	3
Network Immunity Manager.....	3
Network-Based Anomaly Detection	3
General Troubleshooting Steps.....	3
False Positive Avoidance.....	5
Duplicate IP	5
IP Spoofing	5
IP Fanout.....	5
TCP/UDP Fanout	6
DNS Tunneling.....	6
Protocol Anomalies	7
IP Anomalies	7
TCP Anomalies.....	7
UDP Anomalies	8
ICMP Anomalies	8
Average Packet Size Deviation	8
Deployment Scenario.....	9
Switch Events	9
Third-Party Events.....	9

Introduction

This document is intended to inform administrators about ProCurve Network Immunity Manager (NIM) events and how to interpret them. It will also provide a guide to follow while optimizing NIM in your unique environment. This product contains several features that rely on a carefully planned strategy for network deployment. While some customers may turn on NIM and see it trigger events immediately on suspect traffic, others may not.

NIM receives security alerts from two sources. The first is called "ProCurve," which consists of Network Based Anomaly Detection (NBAD) and Virus Throttling (VT) events. The second is called "External," and includes any trap received from a supported external device. More specifically, this could be security devices or applications (e.g., IDS, UTM, etc.). It is important to note that the supported third-party device must be configured to send SNMP Traps to NIM. For up-to-date product configuration information and advanced features, please view product manuals at <http://www.hp.com.rnd/support/manuals/index.htm>.

Network Immunity Manager

Network Immunity Manager (NIM) is a security plug in that must be installed on top of ProCurve Manager Plus (PCM+) network management software. NIM uses statistics from sampling technologies such as sFlow and XRMON to detect and trigger NBAD alerts. The switch on which you want to detect malicious traffic must have an sFlow agent or XRMON support. Networks with switches that do not support these technologies can still work with NIM; however, the malicious traffic must traverse an sFlow/XRMON switch to detect an offender. Actions can be taken on any switch that the PCM server is authorized to manage and the switch supports the action taken (e.g., rate limit).

NIM compiles and analyzes these samples over time. It is designed to trigger an event when network traffic meets the criteria of one of the seven NBAD engines. NIM out of the box is simply alerting the network administrator of the occurrence that network traffic has met the criteria set. Network administrators are ultimately responsible for determining if this is normal traffic in their environment or malicious.

Each network is unique, and NIM provides a customizable interface for NBAD sensitivity through the preferences panel. The event browser and security views enable you to analyze your network alerts and actions over time. With this data, you can make informed decisions on the level of sensitivity that you may need for each NBAD engine and decide what sensitivity best fits your traffic model. We recommend that you exclude certain devices that inappropriately trigger NBAD alerts such as proxies, routers and mail servers.

Network-Based Anomaly Detection

This section explains the thought process that can be used to increase the value of events.

General Troubleshooting Steps

These troubleshooting steps are relevant to all of the NBAD engines. Later in this document, we will present additional steps that are required specifically for each particular engine. The following is an example of the thought process a network administrator might use to investigate an event:

- What is the offender IP and MAC address?
 - This information is present in the NBAD events.
- What switch port is the offender using?
 - Can be found with the PCM+ Find Node utility using the offender data in the NBAD events.
- Is the protocol mix unusual for the offender's port?
 - Can be answered using Traffic top talkers / port details.
- Is the traffic volume unusual for the offender's port?
 - Can be answered using Traffic top talkers / port details.
- What kind of device / operating system is the offender?
 - Can be determined using a network access control (NAC) device such as ProCurve NAC 800 as shown below in Figure 1.

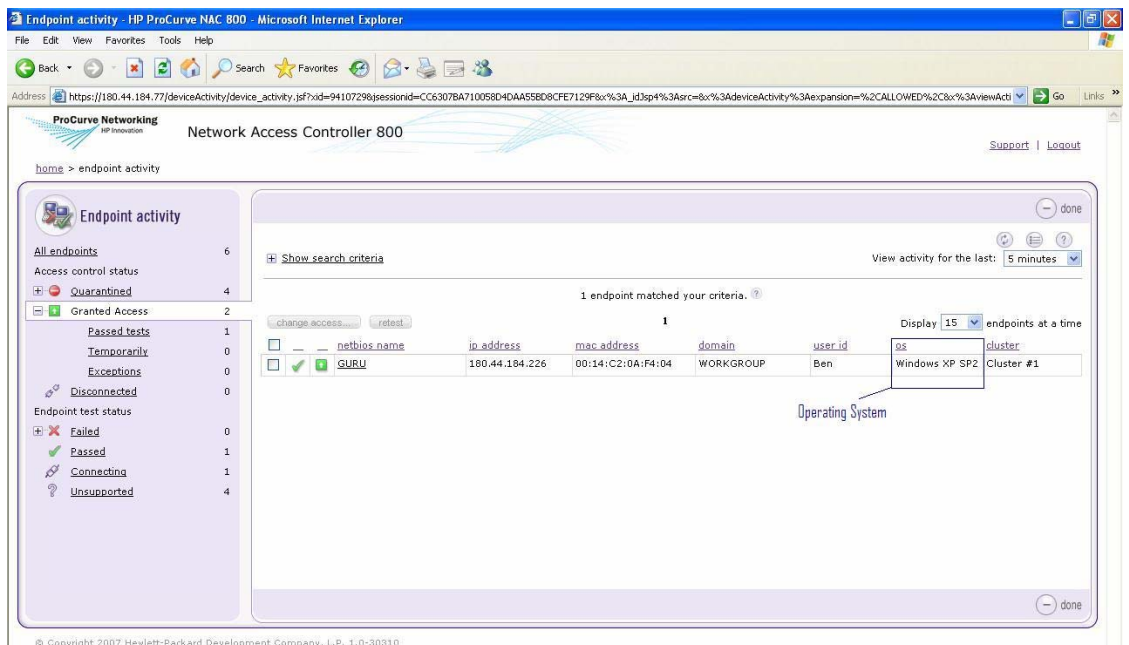


Figure 1

- Is the machine properly patched?
 - Can be determined using a network access control (NAC) device such as ProCurve NAC 800 as shown below with ProCurve IDM.

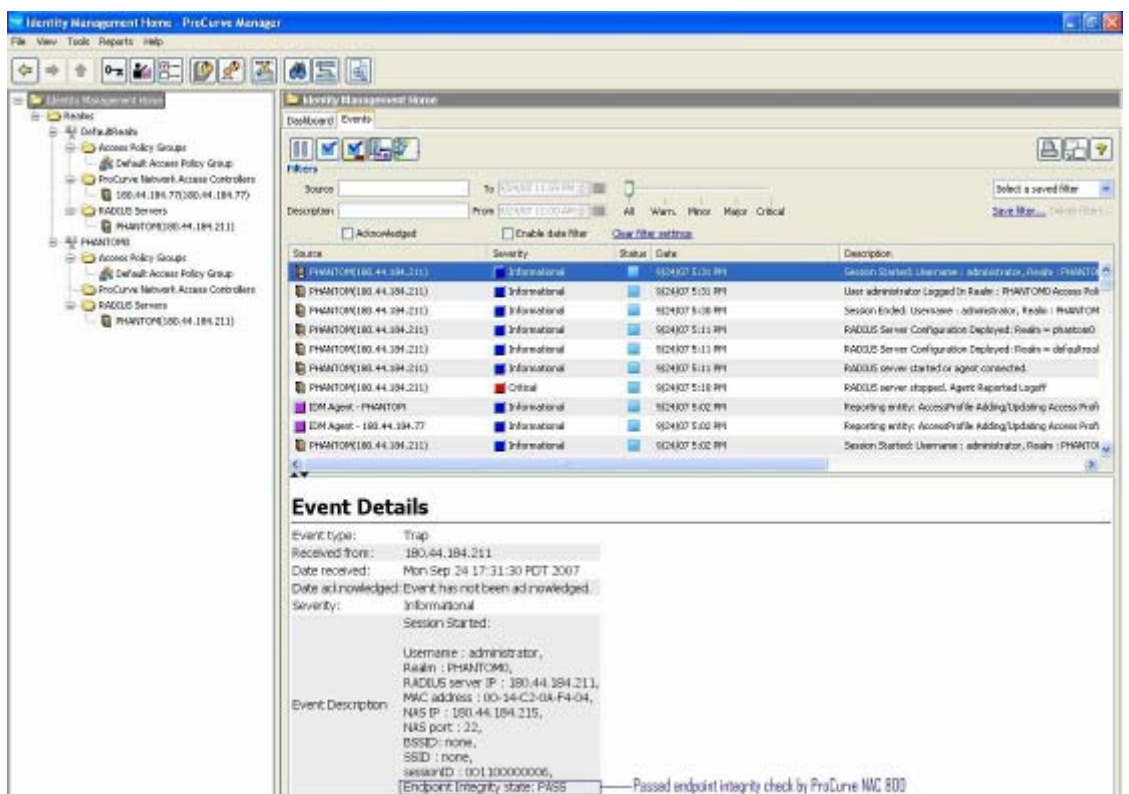


Figure 2

- What network programs are running on the machine?
 - Can be determined using a vulnerability scanner or physically connecting to the machine.
- What appears if you mirror and log the offender's traffic with a packet capture application or examine it with an IDS or UTM?
 - Using ProCurve Remote Port Mirroring capability, one IDPS or UTM anywhere in the network can analyze suspect traffic looking for security problems.

- Third-party SNMP events can be used to trigger appropriate mitigation actions by NIM.

False Positive Avoidance

The algorithms for NIM use aging tables to track offenders, building trust and confidence over time to ensure that NBAD has enough samples to make an informed decision. In addition to those mechanisms to provide False Positive Avoidance, NIM will use thresholds for triggering each event. When you move the sensitivity bar for an NBAD event, the thresholds affect all of these mechanisms internally to make the right decision based on individual network needs. The security monitor sensitivity bar can be found via the preferences panel. Choosing “1” sets the NBAD engines to be the least sensitive while setting the bar to “5” causes NBAD engines to be the most sensitive.

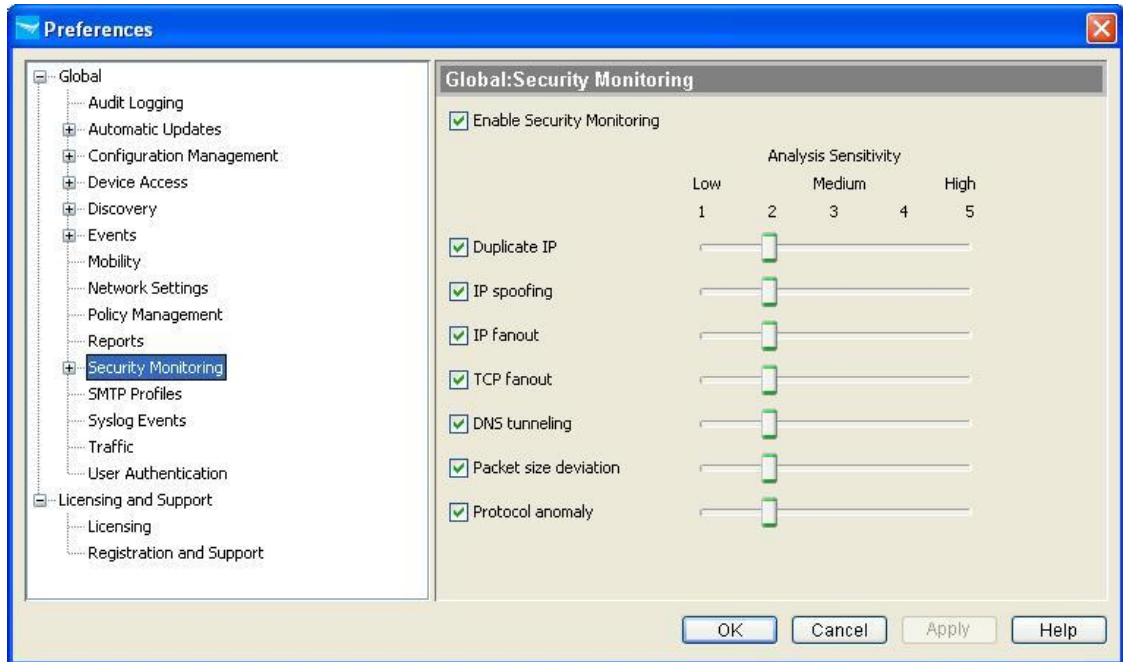


Figure 3

Duplicate IP

This occurs when two network packets are sampled that have different source MAC addresses, but the same source IP address. This can be reported for innocuous reasons when traffic is sampled on both sides of a router. By default, NIM will automatically exclude ProCurve devices that are discovered with routing enabled. As a guideline, all non-ProCurve routers and gateways should be excluded manually.

IP Spoofing

This occurs when two network packets are sampled that have the same source MAC address, but different source IP addresses. This can be reported for innocuous reasons when traffic is sampled on either side of a router; packets forwarded by the router will have the source MAC address of the router interface and the source IP address of the devices that sent the packet. By default, NIM will automatically exclude ProCurve devices that are discovered with routing enabled. As a guideline, all non-ProCurve routers, gateways, NAT devices and proxies should be excluded manually.

IP Fanout

This occurs when one IP is sending traffic to an excessive number of other IPs within a specified time window. Normal client/server applications where many clients innocuously contact a server simultaneously can cause this engine to trigger. These servers should be added to the exclude list. This could include DNS, Antivirus, e-mail and management servers such as PCM. The PCM server is excluded by default.

Because an IP fanout represents an attempt to scan a network or infect a range of machines with worms, a user who is trying to debug an IP fanout event will be wondering the following:

- How many offenders are generating IP fanouts at the same time?
 - If one machine, this might indicate a worm infection or an attempt to scan the network.
 - If many machines, this might indicate a worm that has infected multiple machines and has started to spread.
- What is the transport protocol?
 - TCP indicates a TCP worm or TCP port scan across a wide range of systems – for example, port scanning an entire subnet at once.
 - UDP indicates a UDP worm or UDP port scan across a wide range of systems. These are especially dangerous because the offender's IP can often be spoofed to mask the real offender because there is no protocol handshake.
 - ICMP indicates some kind of host detection scan across a wide range of systems. This could be a legitimate discovery tool such as PCMP+, or it could be a malicious attempt to find live IPs that might be checked later for vulnerability to various security exploits.
- If the transport layer has source and destination ports (TCP or UDP), how many destination ports and which destination ports are appearing in the IP fanout events?
 - If the destination port is always identical, or one of a small group of ports, this might indicate a worm that attacks a specific security hole in a network service, or it could indicate a scanning and attacking tool that is specific to certain network services, such as an attempt to run a Metasploit plugin against a wide range of IPs looking for any that are vulnerable to the specific exploit.
 - If the destination port seems to be a lot of random values, this would indicate some kind of port scan across a wide range of systems, such as port scanning an entire subnet at once.

TCP/UDP Fanout

This occurs when one source IP is communicating with many ports on a destination IP. Anything that does a vulnerability assessment or a port scan can cause this to trigger. This could include endpoint integrity checks.

Because a TCP fanout represents an attempt to scan a single machine for network services or overload the machine with excessive requests, a user who is trying to debug a TCP fanout event will be wondering the following:

- How many offenders are generating TCP fanouts at the same time?
 - If one machine is generating TCP fanouts, this indicates an attempt to scan a specific machine for network services or overload a specific machine with excessive requests.
 - If many machines are generating TCP fanouts, then:
 - if the attack is directed at one machine, this indicates a distributed denial-of-service attack;
 - if the attack is directed at many machines, this indicates an attempt to cover up a scan by generating fake scans from fake machines. One might try to look for IP spoofing at this point, but it could be difficult if the offender changes MACs and IPs simultaneously and defeats the spoofing detection.
- How many TCP destination ports and which destination ports are appearing in the TCP fanout events?
 - If the destination port is always the same port, or one of a small group of ports, this would indicate a small port scan from a normally quiet system, or an attempt to create a denial of service for particular network services.
 - If the destination port seems to be a lot of random values, this would indicate some kind of traditional port scan, or an attempt at overall denial of service for that host, such as a TCP SYN flood.

DNS Tunneling

Significant DNS traffic such as a zone transfer can cause this engine to trigger. A network administrator has the option of lowering the sensitivity, but it is recommended that DNS servers be excluded for this engine and the IP fanout engine.

Because a DNS tunneling event represents an attempt to create a covert channel, a user who is trying to debug a DNS tunneling event will be wondering the following:

- Is the offender a mail server?
 - Mail servers use techniques that appear to be tunneling to send and receive SPF spam prevention DNS records.
 - If this happens, exclude the mail servers from DNS tunneling analysis.
- Is the problem originating on a guest network or wireless network?
 - Publicly accessible networks with restrictions are a common target because it is easy to get in, but the restrictions must be bypassed before the network can be used to its full potential.
- Are DNS requests happening unusually often and/or consuming a large amount of bandwidth on internal clients that should not perform much DNS?
 - This can be investigated using Traffic top talkers / port details under the traffic tab of the device.
- Tunneling almost always uses weird lookups in Base 32/64 encoding that are ugly and contain more entropy than traditional DNS lookups.
 - Bogus DNS tunneling requests quite simply look different.

Protocol Anomalies

This occurs when a host sends traffic containing unusual properties that would not normally be expected to occur on the network.

IP Anomalies

The "Both IP Addresses Identical" anomaly, in which the source and destination IP addresses in the packet are identical, indicates a "Land Attack," which is a known attempt to attack obsolete TCP/IP stacks that would go into infinite processing loops upon receiving such packets.

The various "IP Option" anomalies generally indicate misconfigured hosts that are sending obsolete IP packets containing strange option settings and should be handled by investigating the host using a packet capture tool.

The "Unknown IP Protocol" anomaly indicates a host that is sending unusual data encapsulated in IP packets. This could indicate an attempt to run an uncommon routing protocol or an attempt to create a covert channel.

Some networking protocols trigger IP anomalies. For example, Protocol Independent Multicast (PIM) can trigger IpUnknownProtocol, and Internet Group Multicast Protocol (IGMP) can trigger IpOptionSecurity anomalies. Once NIM informs you what is running on your network and you have validated their presence, it is good practice to add these known multicast IP addresses to your exclude list.

TCP Anomalies

TCP anomalies generally indicate severe and/or covert TCP port scan methods. The advantage of finding them with NIM instead of a port scan detector is that NIM works on every switch port that has sFlow enabled, rather than just the key choke points where an IDPS or UTM is available.

Anomaly Type	Meaning
TcpFlagsAllBitsSet	XMAS Tree Scan
TcpFlagsFinSetButNoAck	TCP FIN Port Scan
TcpFlagsNoBitsSet	TCP NULL Port Scan
TcpFlagsSynFinBitsSet	TCP Port scan with SYN and FIN flags on

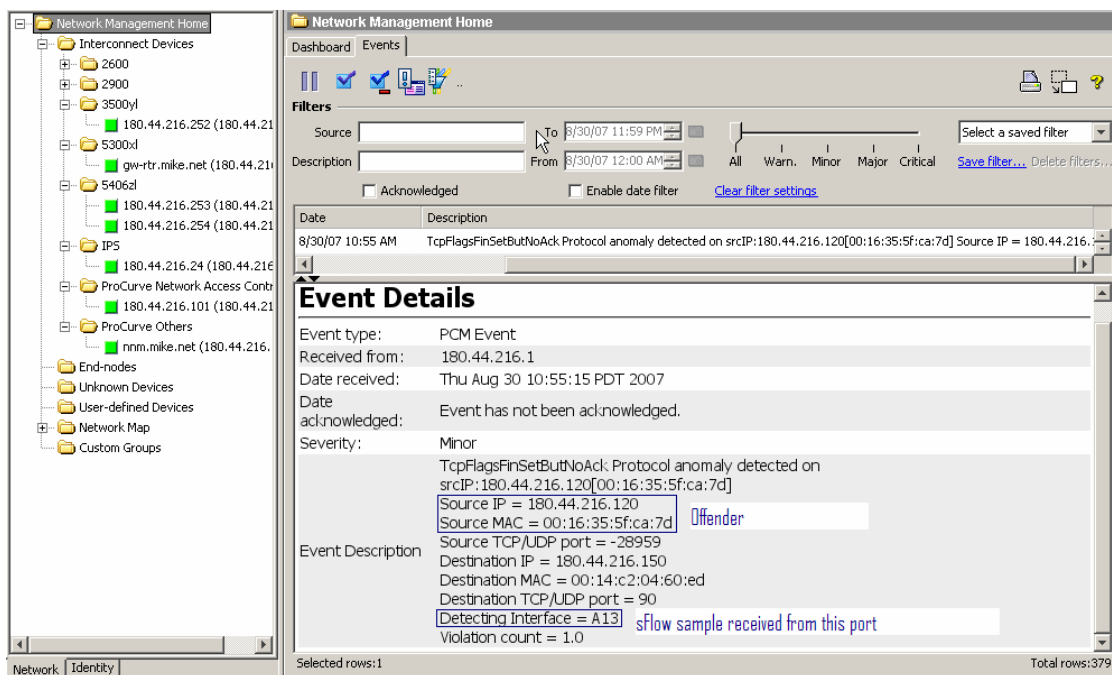


Figure 4

UDP Anomalies

The "UDP Bomb" anomaly indicates an attempt to attack obsolete UDP/IP stacks that would crash and cause denial of service upon receiving such packets. Other UDP anomalies indicate an attempt to take advantage of UDP's stateless nature and some obsolete services: chargen (RFC 864), daytime (RFC 867), discard (RFC 863), echo (RFC 862) and time (RFC 868).

These services make it possible to send spoofed packets that create infinite traffic loops between these services. These infinite loops consume bandwidth and system resources until the machines are located and the services are disabled. Any machines that are seen as offenders or victims of these attacks should be carefully investigated to ensure these services are disabled; however, this should be taken with a grain of salt because it is possible for a third party to spoof the source IP address of these packets in an attempt to evade detection.

ICMP Anomalies

"ICMP Echo" anomalies generally indicate attempts to scan a network looking for open machines that could later be probed for open network services and potential vulnerabilities.

The "Fragmented ICMP Packet," "Oversized ICMP Packet" and "ICMP Ping Of Death" anomalies indicate attempts to attack obsolete IP stacks that would crash and cause denial of service upon receiving such packets.

Other ICMP anomalies generally indicate attempts to disrupt the network or other hosts by generating unexpected error states or confusion for users and for protocol stacks or network services receiving the unexpected ICMP messages.

Average Packet Size Deviation

Average packet size deviation events are different from the other events because they are not meant for direct human interpretation. These events are intended to indicate a difference in the level of activity on a switch port so that NIM can look for freshly active ports on the network.

This engine examines traffic data collected from every monitored device port, irrespective of whether the data is acquired via sampling or polling. On ports that are being polled, traffic data is obtained each minute using the IF-MIB interface counters (or an equivalent set of counters). On ports that are being sampled, traffic data is continuously obtained from arriving sFlow datagrams and XRMON traps.

Once the counter data is retrieved, the collector uses the packet and octet counters to calculate the average size of each packet. Over time, the collector looks for new packet sizes that are outside of the statistical confidence interval established around the packet size.

Ports whose average packet size has exceeded the maximum value in the confidence interval are considered to be freshly active ports. By default, NIM triggers automatic sFlow

sampling on such ports, at a priority above that employed by the PCM+ auto-sampling, but still below that applied when the user configures manual sampling.

Deployment Scenario

One methodology an administrator could follow for protecting the network is a three-phase process.

Phase 1 – Establish a baseline

Install NIM to see what type of anomalous traffic is reported on the network; it should take about three days for NIM to fully learn traffic patterns of your network and build up statistical confidence over time. This also allows a network administrator to track down the sources of NBAD events and determine if the traffic that caused the events is expected. Once this data has been gathered, sources of and destinations of allowable traffic should be added to the exclude list so that they're not repeatedly identified as offenders.

Phase 2 – Take action on alerts

The next phase involves configuring NIM to take action on an alert. A network administrator has the option to create an action to mirror suspect traffic to an IDPS. In this case, a second policy should be setup to act on an alert from this IDPS device. There are many actions to choose from, but for the sake of this scenario, MAC Lockout with a rollback could be configured, which will block malicious traffic from a specific MAC address but not make any permanent changes to network switches. The actions that network administrators choose should be in line with their company's security policy. It is recommended that a network administrator consider the severity, location and impact on the network when choosing an action.

Phase 3 – Fine-tuning Network Immunity Manager

After one month of operation, a network administrator should have a pretty good idea of his/her network traffic patterns. In this final phase, NIM should be fine-tuned for the network that it is monitoring using the NBAD engine sensitivity settings, exclude list and the definition of new security triggers and policies. It is recommended that custom groups be used to define non-mission-critical areas that NIM should take automated action against. This typically would be edge switches with client computers connected that are susceptible to malicious programs or hackers. Any deployment scenario of NIM should always comply with an individual company's security policies.

Switch Events

The switch trap most commonly used by NIM is the Virus Throttling trap. This trap, and the event that results from it in PCM+, can be used to trigger NIM policies that mitigate a threat in more sophisticated ways than are possible without NIM, such as sending email notifications or configuring richer responses such as port shutdowns, MAC lockouts, VLAN quarantines or rate limits. Without help from NIM, switches can throttle or block only a subset of the traffic generated by the offender.

In addition, depending on the network topology, the switch that generates the Virus Throttling trap might not be the offender's edge switch. However, once receiving a Virus Throttling trap from anywhere in the network, NIM is capable of mitigating the offender at the offender's edge port, which prevents the virus from spreading as far as it could have spread otherwise.

Third-Party Events

Select third-party IDPS and UTM security devices are already supported by default in NIM. Additional documents such as "Customizing ProCurve Manager Plus" explain how new sources of third-party SNMP events can be integrated with NIM to provide additional sources of security information. Useful third-party events generally include the following information:

- Offender MAC
- Offender IP
- Victim MAC
- Victim IP
- Severity
- Signature ID & Sub-ID

- Description

Network Immunity Manager uses this data to track security alerts over time across the whole network, giving the administrator a convenient console for considering all sources of security events. Once NIM is configured to accept a given type of third-party event, all of the mitigation actions available (such as emailing the administrator, or disabling suspect ports or MAC addresses) can be configured to automatically respond to third-party notifications. This enhances the capability of NIM and the third-party security device by allowing the mitigation to be performed at the edge of the network, thus offering greater protection than what each could achieve alone.

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

www.procurve.com



© 2007 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.

4AA1-5630ENN, 10/2007