IBM Security
Network Intrusion Prevention System



Network Intrusion Prevention System User Guide

Version 1 Release 4.5

Copyright statement

© Copyright IBM Corporation 2003, 2012

U.S. Government Users Restricted Rights — Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Publication Date: August, 2012

Contents

Homologation statement - regulation notice. v

Safety, environmental, and electronic emissions notices vii

Preface xvii

About IBM Security Network IPS appliance documentation xvii

Customer support. xvii

Chapter 1. Introducing the IBM Security Network Intrusion Prevention System . . 1

Intrusion prevention 1

Appliance interface modes. 2

Responses 3

IPv6 5

Chapter 2. Managing the appliance . . . 7

Using IPS Local Management Interface 8

Managing with SiteProtector 8

Health and sensor alerts 10

Capacity planning 11

NTP servers 11

Chapter 3. Configuring firewall settings 15

Configuring firewall rules 15

Firewall rules language 17

Chapter 4. Working with security events 21

Configuring security events 21

Viewing security event information 21

Chapter 5. Configuring other intrusion prevention settings 23

Managing quarantined intrusions 23

Configuring connection events 23

Configuring user-defined events 24

User-defined event contexts 24

Regular expressions in user-defined events . . . 27

Configuring tuning parameters 29

Configuring OpenSignatures. 29

Configuring SNORT 33

Configuring response filters 37

Configuring a LEEF system log. 38

Chapter 6. X-Force protection modules 39

PAM. 39

Using X-Force default blocking 39

Using data loss prevention signatures 40

Using web application protection 40

Chapter 7. Using protection domains 43

Working with protection domains 43

Best practices for protection domains 44

Chapter 8. Configuring appliances for high availability 45

HA configuration options. 45

Deployment for standard high availability 46

Deployment for geographical high availability. . . 49

Chapter 9. General information 51

Compatibility. 51

Appliance partitions 51

Cumulative updates and rollbacks. 51

Notices 53

Trademarks 54

Index 55

© Copyright IBM Corp. 2003, 2012

Homologation statement - regulation notice

This product is not intended to be connected directly or indirectly by any means whatsoever to interfaces of public telecommunications networks.

本製品は、電気通信事業者の通信回線への直接、またはそれに準ずる方法での接続を目的とするものではありません。

© Copyright IBM Corp. 2003, 2012

Safety, environmental, and electronic emissions notices

Safety notices may be printed throughout this guide. **DANGER** notices warn you of conditions or procedures that can result in death or severe personal injury. **CAUTION** notices warn you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous. **Attention** notices warn you of conditions or procedures that can cause damage to machines, equipment, or programs.

DANGER notices

DANGER

To prevent a possible shock from touching two surfaces with different protective ground (earth), use one hand, when possible, to connect or disconnect signal cables. (D001)

DANGER

Overloading a branch circuit is potentially a fire hazard and a shock hazard under certain conditions. To avoid these hazards, ensure that your system electrical requirements do not exceed branch circuit protection requirements. Refer to the information that is provided with your device or the power rating label for electrical specifications. (D002)

DANGER

If the receptacle has a metal shell, do not touch the shell until you have completed the voltage and grounding checks. Improper wiring or grounding could place dangerous voltage on the metal shell. If any of the conditions are not as described, STOP. Ensure the improper voltage or impedance conditions are corrected before proceeding. (D003)

DANGER

An electrical outlet that is not correctly wired could place hazardous voltage on the metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock. (D004)

DANGER

© Copyright IBM Corp. 2003, 2012

When working on or around the system, observe the following precautions:

Electrical voltage and current from power, telephone, and communication cables are hazardous. To avoid a shock hazard:

- **Connect power to this unit only with the IBM® ISS provided power cord. Do not use the IBM ISS provided power cord for any other product.**
- **Do not open or service any power supply assembly.**
- **Do not connect or disconnect any cables or perform installation, maintenance, or reconfiguration of this product during an electrical storm.**
- **The product might be equipped with multiple power cords. To remove all hazardous voltages, disconnect all power cords.**
- **Connect all power cords to a properly wired and grounded electrical outlet. Ensure that the outlet supplies proper voltage and phase rotation according to the system rating plate.**
- **Connect any equipment that will be attached to this product to properly wired outlets.**
- **When possible, use one hand only to connect or disconnect signal cables.**
- **Never turn on any equipment when there is evidence of fire, water, or structural damage.**
- **Disconnect the attached power cords, telecommunications systems, networks, and modems before you open the device covers, unless instructed otherwise in the installation and configuration procedures.**
- **Connect and disconnect cables as described in the following procedures when installing, moving, or opening covers on this product or attached devices.**

To disconnect:

1. **Turn off everything (unless instructed otherwise).**
2. **Remove the power cords from the outlets.**
3. **Remove the signal cables from the connectors.**
4. **Remove all cables from the devices.**

To connect:

1. **Turn off everything (unless instructed otherwise).**
2. **Attach all cables to the devices.**
3. **Attach the signal cables to the connectors.**
4. **Attach the power cords to the outlets.**
5. **Turn on the devices.**

(D005)

CAUTION notices

CAUTION:
Data processing environments can contain equipment transmitting on system links with laser modules that operate at great than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)

CAUTION:
The battery contains lithium. To avoid possible explosion, do not burn or charge the battery.

Do not:

- **Throw or immerse into water**
- **Heat to more than 100°C (212°F)**
- **Repair or disassemble**

Exchange only with the IBM ISS-approved part. Recycle or discard the battery as instructed by local regulations. In the United States, IBM ISS has a process for the collection of this battery. For information, call 1-800-426-4333. Have the IBM ISS part number for the battery unit available when you call. (C003)

CAUTION:
For 19" rack mount products:

- **Do not install a unit in a rack where the internal rack ambient temperatures will exceed the manufacturer's recommended ambient temperature for all your rack-mounted devices.**
- **Do not install a unit in a rack where the air flow is compromised. Ensure that air flow is not blocked or reduced on any side, front, or back of a unit used for air flow through the unit.**
- **Consideration should be given to the connection of the equipment to the supply circuit so that overloading the circuits does not compromise the supply wiring or overcurrent protection. To provide the correct power connection to a rack, refer to the rating labels located on the equipment in the rack to determine the total power requirement of the supply circuit.**
- *(For sliding drawers)* **Do not pull or install any drawer or feature if the rack stabilizer brackets are not attached to the rack. Do not pull out more than one drawer at a time. The rack might become unstable if you pull out more than one drawer at a time.**
- *(For fixed drawers)* **This drawer is a fixed drawer and must not be moved for servicing unless specified by the manufacturer. Attempting to move the drawer partially or completely out of the rack might cause the rack to become unstable or cause the drawer to fall out of the rack.**

(R001 Part 2 of 2)

Product handling information

One of the following two safety notices may apply to this product. Please refer to the specific product specifications to determine the weight of the product to see which applies.

CAUTION:
This part or unit is heavy but has a weight smaller than 18 kg (39.7 lb). Use care when lifting, removing, or installing this part or unit. (C008)

CAUTION:
The weight of this part or unit is between 18 and 32 kg (39.7 and 70.5 lb). It takes two persons to safely lift this part or unit. (C009)







Product safety labels

One or more of the following safety labels may apply to this product.

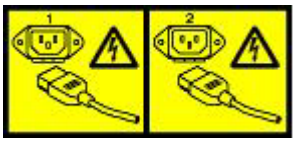
DANGER

Hazardous voltage, current, or energy levels are present inside any component that has this label attached. Do not open any cover or barrier that contains this label. (L001)



DANGER

Multiple power cords. The product might be equipped with multiple power cords. To remove all hazardous voltages, cdisconnect all power cords. (L003)



World trade safety information

Several countries require the safety information contained in product publications to be presented in their national languages. If this requirement applies to your country, a safety information booklet is included in the publications package shipped with the product. The booklet contains the safety information in your national language with references to the US English source. Before using a US English publication to install, operate, or service this IBM ISS product, you must first become familiar with the related safety information in the booklet. You should also refer to the booklet any time you do not clearly understand any safety information in the US English publications.

Laser safety information

The following laser safety notices apply to this product:

CAUTION:
This product may contain one or more of the following devices: CD-ROM drive, DVD-ROM drive, DVD-RAM drive, or laser module, which are Class 1 laser products. Note the following information:

- **Do not remove the covers. Removing the covers of the laser product could result in exposure to hazardous laser radiation. There are no serviceable parts inside the device.**
- **Use of the controls or adjustments or performance of procedures other than those specified herein might result in hazardous radiation exposure. (C026)**

CAUTION:
Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. For this reason, never look into the end of an optical fiber cable or open receptacle. (C027)

Laser compliance

All lasers are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for class 1 laser products. Outside the U.S., they are certified to be in compliance with IEC 60825 as a class 1 laser product. Consult the label on each part for laser certification numbers and approval information.

Product recycling and disposal

This unit must be recycled or discarded according to applicable local and national regulations. IBM encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist equipment owners in recycling their IT products. Information on IBM ISS product recycling offerings can be found on IBM's Internet site at http:// www.ibm.com/ibm/environment/products/prp.shtml.

Esta unidad debe reciclarse o desecharse de acuerdo con lo establecido en la normativa nacional o local aplicable. IBM recomienda a los propietarios de equipos de tecnología de la información (TI) que reciclen responsablemente sus equipos cuando éstos ya no les sean útiles. IBM dispone de una serie de programas y servicios de devolución de productos en varios países, a fin de ayudar a los propietarios de equipos a reciclar sus productos de TI. Se puede encontrar información sobre las ofertas de reciclado de productos de IBM en el sitio web de IBM http:// www.ibm.com/ibm/environment/products/prp.shtml.



Notice: This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable through the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.

In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling, and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

注意: このマークは EU 諸国およびノルウェーにおいてのみ適用されます。

この機器には、EU 諸国に対する廃電気電子機器指令 2002/96/EC(WEEE) のラベルが貼られています。この指令は、EU 諸国に適用する使用済み機器の回収とリサイクルの骨子を定めています。このラベルは、使用済みになった時に指令に従って適正な処理をする必要があることを知らせるために種々の製品に貼られています。

Remarque: Cette marque s'applique uniquement aux pays de l'Union Européenne et à la Norvège.

L'etiquette du système respecte la Directive européenne 2002/96/EC en matière de Déchets des Equipements Electriques et Electroniques (DEEE), qui détermine les dispositions de retour et de recyclage applicables aux systèmes utilisés à travers l'Union européenne. Conformément à la directive, ladite étiquette précise que le produit sur lequel elle est apposée ne doit pas être jeté mais être récupéré en fin de vie.

Battery return program

This product contains a lithium battery. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to http://www.ibm.com/ibm/environment/products/ batteryrecycle.shtm or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426- 4333. Please have the IBM part number listed on the battery available prior to your call.

For Taiwan:



Please recycle batteries 廢電池請回收

For the European Union:

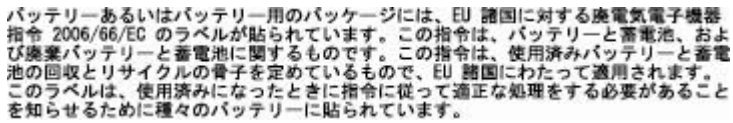


Notice: This mark applies only to countries within the European Union (EU).

Batteries or packing for batteries are labeled in accordance with European Directive 2006/66/EC concerning batteries and accumulators and waste batteries and accumulators. The Directive determines the framework for the return and recycling of used batteries and accumulators as applicable throughout the European Union. This label is applied to various batteries to indicate that the battery is not to be thrown away, but rather reclaimed upon end of life per this Directive.

Les batteries ou emballages pour batteries sont étiquetés conformément aux directives européennes 2006/66/EC, norme relative aux batteries et accumulateurs en usage et aux batteries et accumulateurs usés. Les directives déterminent la marche à suivre en vigueur dans l'Union Européenne pour le retour et

le recyclage des batteries et accumulateurs usés. Cette étiquette est appliquée sur diverses batteries pour indiquer que la batterie ne doit pas être mise au rebut mais plutôt récupérée en fin de cycle de vie selon cette norme.


バッテリーあるいはバッテリー用のパッケージには、EU 諸国に対する廃電気電子機器指令 2006/66/EC のラベルが貼られています。この指令は、バッテリーと蓄電池、および廃棄バッテリーと蓄電池に関するものです。この指令は、使用済みバッテリーと蓄電池の回収とリサイクルの骨子を定めているもので、EU 諸国にわたって適用されます。このラベルは、使用済みになったときに指令に従って適正な処理をする必要があることを知らせるために種々のバッテリーに貼られています。

In accordance with the European Directive 2006/66/EC, batteries and accumulators are labeled to indicate that they are to be collected separately and recycled at end of life. The label on the battery may also include a symbol for the metal concerned in the battery (Pb for lead, Hg for the mercury, and Cd for cadmium). Users of batteries and accumulators must not dispose of batteries and accumulators as unsorted municipal waste, but use the collection framework available to customers for the return, recycling, and treatment of batteries and accumulators. Customer participation is important to minimize any potential effects of batteries and accumulators on the environment and human health due to potential presence of hazardous substances. For proper collection and treatment, contact your local IBM representative.

For California:

Perchlorate Material - special handling may apply. See http://www.dtsc.ca.gov/ hazardouswaste/ perchlorate.

The foregoing notice is provided in accordance with California Code of Regulations Title 22, Division 4.5, Chapter 33. Best Management Practices for Perchlorate Materials. This product, part, or both may include a lithium manganese dioxide battery which contains a perchlorate substance.

Electronic emissions notices

The following statements apply to this IBM product. The statement for other IBM products intended for use with this product will appear in their accompanying manuals.

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. this equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions contained in the installation manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Note: Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, by installation or use of this equipment other than xvi IBM Internet Security Systems as specified in the installation manual, or by any other unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

Note: This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Canadian Department of Communications Compliance Statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité aux normes du ministère des Communications du Canada

Cet appareil numérique de las classe A est conform à la norme NMB-003 du Canada.

European Union (EU) Electromagnetic Compatibility Directive

This product is in conformity with the protection requirements of EU Council Directive 2004/108/ EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM ISS cannot accept responsibility for any failure to satisfy the protection requirements resulting from a non-recommended modification of the product, including the fitting of non-IBM ISS option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to European Standard EN 55022. The limits for Class equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Warning:

This is a Class A product. In a domestic environment, this product may cause radio interference in which case the user may be required to take adequate measures.

European Community contact:

IBM Technical Regulations
Pascalstr. 100, Stuttgart, Germany 70569
Telephone: 0049 (0) 711 785 1176
Fax: 0049 (0) 711 785 1283
e-mail: tjahn@de.ibm.com

EC Declaration of Conformity (In German)

Deutschsprachiger EU Hinweis: Hinweis für Geräte der Klasse A EU-Richtlinie zur Elektromagnetischen Verträglichkeit

Dieses Produkt entspricht den Schutzanforderungen der EU-Richtlinie 89/336/EWG zur Angleichung der Rechtsvorschriften über die elektromagnetische Verträglichkeit in den EUMitgliedsstaaten und hält die Grenzwerte der EN 55022 Klasse A ein.

Um dieses sicherzustellen, sind die Geräte wie in den Handbüchern beschrieben zu installieren und zu betreiben. Des Weiteren dürfen auch nur von der IBM empfohlene Kabel angeschlossen werden. IBM übernimmt keine Verantwortung für die Einhaltung der Schutzanforderungen, wenn das Produkt ohne Zustimmung der IBM verändert bzw. wenn Erweiterungskomponenten von Fremdherstellern ohne Empfehlung der IBM gesteckt/eingebaut werden.

EN 55022 Klasse A Geräte müssen mit folgendem Warnhinweis versehen werden: "Warnung: Dieses ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funk-Störungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Maßnahmen zu ergreifen und dafür aufzukommen."

Deutschland: Einhaltung des Gesetzes über die elektromagnetische Verträglichkeit von Geräten

Dieses Produkt entspricht dem “Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG)“. Dies ist die Umsetzung der EU-Richtlinie 89/336/EWG in der Bundesrepublik Deutschland.

Zulassungsbescheinigung laut dem Deutschen Gesetz über die elektromagnetische Verträglichkeit von Geräten (EMVG) vom 18. September 1998 (bzw. der EMC EG Richtlinie 89/336) für Geräte der Klasse A.

Dieses Gerät ist berechtigt, in Übereinstimmung mit dem Deutschen EMVG das EGKonformitätszeichen - CE - zu führen.

Verantwortlich für die Konformitätserklärung nach Paragraf 5 des EMVG ist die IBM Deutschland GmbH, 70548 Stuttgart.

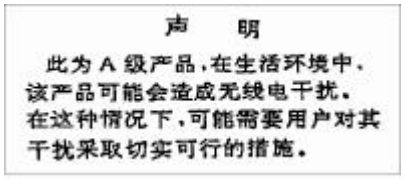
Informationen in Hinsicht EMVG Paragraf 4 Abs. (1) 4:

Das Gerät erfüllt die Schutzanforderungen nach EN 55024 und EN 55022 Klasse A

update: 2004/12/07

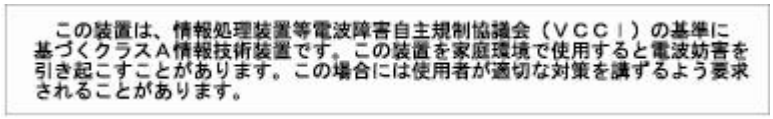
People's Republic of China Class A Compliance Statement:

This is a Class A product. In a domestic environment, this product may cause radio interference in which case the user may need to perform practical actions.

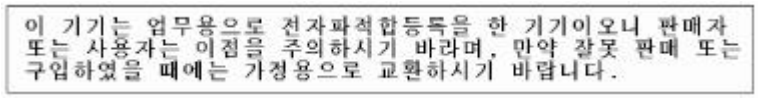


Japan Class A Compliance Statement:

This product is a Class A Information Technology Equipment and conforms to the standards set by the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). In a xviii IBM Internet Security Systems domestic environment, this product may cause radio interference in which case the user may be required to take adequate measures.



Korean Class A Compliance Statement:



Preface

This guide describes the features and capabilities of IBM Security Network Intrusion Prevention System (IPS) for your IBM Security Network IPS GX and GV appliances.

Audience

This guide is intended for network security system administrators responsible for setting up, configuring, and managing the intrusion prevention system in a network environment. A fundamental knowledge of network security policies and IP network configuration is helpful.

Supported appliance models

This firmware release supports the following appliance models:

- GX3002
- GX4000 series
- GX5000 series
- GX6000 series
- GX7000 series
- GV200
- GV1000

About IBM Security Network IPS appliance documentation

This guide describes the concepts and capabilities of IBM Security Network Intrusion Prevention System (IPS). Refer to the online help for procedural and "how to" information for configuring and managing appliances.

Latest publications

For the latest documentation, go to the IBM Security Product Information Center at http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp.

Support knowledge base

The IBM Support knowledge base is a valuable source of information. Visit the knowledge base at IBM Support knowledge base.

License agreement

For licensing information on IBM Security products, download the IBM Licensing Agreement from http://www.ibm.com/services/us/iss/html/contracts_landing.html .

Customer support

IBM Security Solutions provides technical support to customers that are entitled to receive support.

Customer support

The IBM Support Home provides the following information:

- Registration and eligibility requirements for receiving support

© Copyright IBM Corp. 2003, 2012

- Customer support telephone numbers for the country in which you are located
- Information you must gather before contacting customer support

Chapter 1. Introducing the IBM Security Network Intrusion Prevention System

This chapter introduces the IBM Security Network Intrusion Prevention System (IPS) and describes how its features protect the network with minimum configuration. It also describes other IBM Security Network IPS features you can implement to customize your network security.

Intrusion prevention

The IBM Security Network Intrusion Prevention System (IPS) automatically blocks malicious attacks while preserving network bandwidth and availability. The IBM Security Network IPS appliances are purpose-built, Layer 2 network security appliances that you can deploy either at the gateway or the network to block intrusion attempts, denial of service (DoS) attacks, malicious code, backdoors, spyware, peer-to-peer applications, and a growing list of threats without requiring extensive network reconfiguration.

With flexible deployment options and out-of-the-box functionality, these appliances ensure accurate, high-performance protection at both the network perimeter and across internal networks and internal network segments.

Protection features

IBM Security intrusion prevention features include proven detection and prevention technologies, along with the latest security updates. These appliances understand the logical flow and state of traffic, resulting in unsurpassed protection against network threats, including trojans, backdoors and worms.

IBM Security Network IPS offers the following features to protect your network against threats:

- **Dynamic blocking**

  IBM Security Network IPS uses vulnerability-based attack identification to enable an immediate and reliable blocking response to unwanted traffic while allowing legitimate traffic to pass unhindered. It employs a deep traffic inspection process that uses detection-based blocking to stop both known attacks and previously unknown attacks.
- **Firewall rules**

  You can create firewall rules that enable the appliance to block incoming packets from particular IP addresses, port numbers, protocols, or VLANs. These rules block many attacks before they affect your network.
- **Automatic security content updates based on the latest security research**

  You can automatically download and activate updated security content. The security updates you receive are a result of IBM ISS X-Force Research and Development Team's ongoing commitment to provide the most up-to-date protection against known and unknown threats.
- **Quarantine and block responses**

  Inline appliances use the quarantine response to block traffic for a specified amount of time after an initial attack, and they use the block response to block and reset a connection in which an event occurs or to drop the packet that triggered an event.
- **Virtual Patch™ protection**

  IBM Security 's Virtual Patch® capability provides a valuable time buffer, eliminating the need for you to immediately patch all vulnerable systems. You can wait until you are ready to manually update servers or until scheduled updates occur, rather than having to patch and restart systems.
- **SNMP and SNMPv3 support**

© Copyright IBM Corp. 2003, 2012

Using SNMP-based traps, you can monitor key system problem indicators or respond to security or other appliance events using SNMP and SNMPv3 responses.

- **IPv6**

  Network IPS appliances support IPv6 networks for many features, including Firewall Rules, Connection Events, and Quarantine Rules.

- **SNORT**

  Network IPS appliances include an integrated SNORT system that processes packets, sends alerts, logs events, and sends responses according to specific configuration contents and rules.

Appliance interface modes

The inline appliances includes three interface modes as follows:

- Inline protection
- Inline simulation
- Passive monitoring

You selected one of these operation modes when you configured the appliance settings. Using the Configuration menu, if you like, you can use the default operation mode and select a different one later.

Interface modes

Inline protection

Inline protection mode allows you to fully integrate the appliance into the network infrastructure. In addition to the block and quarantine responses, all firewall rules are enabled, and the full security policy you applied is enabled.

Note: This is the default mode of the appliance

Inline simulation

Inline simulation mode allows you to monitor the network using the appliance without affecting traffic patterns. In addition to the traditional block response, the appliance uses the quarantine response. Packets are not dropped when these responses are invoked, and the appliance does not reset TCP connections by default. Events that would have been blocked are reported with the status "Simulated Block." This mode is helpful for baselining and testing your security policy without affecting network traffic.

Passive monitoring

Passive monitoring mode replicates traditional passive intrusion detection system (IDS) functionality, monitoring network traffic without sitting inline. If the appliance encounters suspicious network activity, it sends a reset to block a TCP connection. This mode is helpful for determining what type of inline protection your network requires.

Changing appliance interface modes

If you change between the passive monitoring mode and the inline simulation or inline protection mode, you must change the network connections to your appliance. An appliance operating in passive monitoring mode requires a connection to a tap, hub, or SPAN port.

If you change the appliance interface mode from inline simulation to inline protection, you may need to modify some advanced parameters to set them appropriately for inline protection.

Responses

Responses control how the appliance reacts when it detects intrusions or other important events. The appliance offers many predefined responses. In addition, you can configure your own responses and then apply them to events as necessary.

Block response

The Block and Ignore responses are always available as responses to intrusions. The **Block response** is a default response that blocks attacks by dropping packets and sending resets to TCP connections. The Block response differs depending on the appliance's operation mode, as follows:

In this mode	The appliance
Passive Monitoring	Sends resets to block TCP connections, but performs no other blocking. If not required, resets can be disabled by a tuning parameter or by disabling the block responses in the security events policy. Resets can also be disabled by changing default x-force blocking option to "never."
Inline Simulation	Monitors network traffic and generates alerts but does not block the offending traffic.
Inline Protection	Blocks attacks by dropping packets and sending resets to TCP connections.

Ignore response

The ignore response tells the appliance to disregard packets that match criteria specified within an event. You can set this response to ignore events for specified traffic through a response filter or you can use it to ignore certain events for a protection domain. If you select this response when you create response filters or security events, the appliance does not act when it detects the matching packets.

The Ignore response takes precedence over any other responses you may have configured. If you select Ignore, no other response actions are taken for a particular event.

Important: Use the Ignore response to filter only security events that do not threaten the network.

Configurable responses

You can create additional responses to use in conjunction with the Block and Ignore responses. The following tables lists the types of responses you can configure.

Configurable Response	Description
Email	You can configure the appliance to e-mail notifications to individuals or groups when events occur. You can select the event parameters to include in the message to provide important information about detected events.

Configurable Response	Description
Log Evidence	You can configure the appliance to save a copy of a single packet that triggers an event or to save all packets on a session that triggers an event. Identify the capture log file by its event name and event ID. Evidence logs show you what an intruder tried to do to the network. The appliance logs packets that trigger events to the `/cache/packetlogger/logevidence` folder. Download or delete the packet files from IPS Local Management Interface.
Quarantine	You can create quarantine responses to block intruders when the appliance detects security, connection, or user-defined events. Quarantine responses are effective at blocking worms and trojans. Quarantine responses take effect only when you have configured the appliance to run in Inline Protection mode. The appliance generates its own quarantine rules in response to detected intruder events. These dynamic quarantine rules appear on the Quarantined Intrusions page and are in addition to any quarantine rules that you create manually. **Note:** Some predefined quarantine responses are already in place. You cannot rename, modify, or remove predefined responses.
SNMP	You can configure Simple Network Management Protocol (SNMP) notification responses for connection, security, and user-defined events. SNMP responses pull values relevant to the event and send them to an SNMP manager.
User Specified	You can configure your own responses to events, such as executing an application or script. You can use a Linux binary or shell script, including any command-line options or arguments (such as event name or source address). After you create the response, you must manually copy the executable to the appliance.

Predefined quarantine responses

Response	Description
Quarantine Intruder	Stops inbound network traffic to a target from a specific intruder. This response adds a quarantine rule to block the matching protocol traffic from the intruder IP address to the target IP address. Use this response to prevent a known malicious intruder from establishing communication with a server. This response is not suitable for blocking network sweep security events. If enabled, a sweep of a subnet by an intruder adds so many quarantine rules that the response does not effectively block the sweep.
Quarantine Trojan	Provides a method to stop all network communication for a potentially infected host. This response adds a quarantine rule to block traffic to a certain TCP or UDP port on a single victim for the specified duration of time. Before using this option, consider the false positive risks. Use this option for times when zero-day or high impact Trojans spread across the Internet. **Note:** This response does not apply to ICMP traffic.

Response	Description
Quarantine Worm	Provides a method to minimize the spread of a network worm that is attempting to propagate itself. This response adds a quarantine rule to block traffic to a certain TCP or UDP port from a single intruder for the specified duration of time. It is suitable for blocking a BotNet attempting to establish a conversation with a zombie or a potential vulnerable network service. **Note:** This response does not apply to ICMP traffic.
Quarantine DDOS (Distributed Denial-of-Service)	Blocks traffic from an intruder that is related to a specific attack. This response is suitable for blocking DDOS events while reducing the reporting load. The matching events from the same intruder are silently blocked and are not reported again while the quarantine rule is active. **Note:** The **Quarantine DDOS (Distributed Denial-of-Service)** predefined response functions for security events only and not for any other type of event.

Response objects in SiteProtector™

If you are managing the appliance through SiteProtector and you want to configure responses for events, use Response Objects. Response objects enable you to centralize data. If the data changes, you can modify the response object instead of each instance of the data.

Note: If you are using SiteProtector to manage the appliance, you can use Central Responses to create event responses. See "Configuring Central Responses" in the SiteProtector Help for more information.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Response Tuning** > **Responses**

In SiteProtector:

- **Shared Objects** > **Response Objects**

IPv6

IBM Security Network Intrusion Prevention System (IPS) protects IPv6 networks from attacks. Any special considerations related to IPv6 support are noted in the appliance help.

IPv6 is intended to replace IPv4 as the standard internet protocol. As you prepare your networks for the transition to IPv6 traffic, the Network IPS appliance can be configured to support IPv6 traffic while it continues to support IPv4 traffic.

The appliance supports IPv6 addresses for the following features.

- User defined events
- Protection domains
- Connection events
- Quarantine rules
- Response filters
- Firewall rules
- High availability
- Management interface

- Agent Manager for SiteProtector
- SNMP notifications (informs and traps)
- NTP servers

Chapter 2. Managing the appliance

You can create and deploy security policies, manage alerts, and apply updates for your appliances either locally or through a central appliance management system.

IBM Security Network IPS offers you the following tools for managing appliances:

- IPS Local Management Interface (for managing appliances individually and locally)
- SiteProtector (for managing appliances from a central management console)

IPS Local Management Interface

IPS Local Management Interface is a browser-based graphical user interface (GUI) for local, single appliance management. You can use IPS Local Management Interface to manage the following functions:

- Monitoring appliance's status
- Configuring operation modes
- Configuring firewall settings
- Managing appliance settings and activities
- Reviewing alert details
- Configuring high availability
- Managing security policies with protection domains

SiteProtector

SiteProtector is the IBM ISS central management console. With SiteProtector, you can manage components and appliances, monitor events, and schedule reports. By default, your appliance is set up to be managed through IPS Local Management Interface. If you are managing a group of appliances along with other sensors, you may prefer the centralized management capabilities that SiteProtector provides.

When you register your appliance with SiteProtector, SiteProtector controls the following management functions of the appliance:

- Firewall settings
- Intrusion prevention settings
- Alert events
- Appliance and security content updates

After you register the appliance with SiteProtector, you can *view* these functions in IPS Local Management Interface but you can *change* them only from SiteProtector.

Reference: For instructions on managing the appliance through SiteProtector, see the SiteProtector user documentation at http://publib.boulder.ibm.com/infocenter/sprotect/v2r8m0/index.jsp or the SiteProtector Help.

What you manage in SiteProtector or IPS Local Management Interface

You must manage certain local functions directly on the appliance. However, you can control other functions on SiteProtector after registering with SiteProtector.

Note: After you register the appliance with SiteProtector, some areas of IPS Local Management Interface become read-only. When you unregister the appliance from SiteProtector, IPS Local Management Interface

© Copyright IBM Corp. 2003, 2012

becomes fully functional again.
The following table lists functions you control using SiteProtector or IPS Local Management Interface.

Functions	SiteProtector	IPS Local Management Interface
Alert events	✓	✓
Firewall settings	✓	✓
Installation settings	✓	✓
Intrusion prevention settings	✓	
Manual updates		✓
Quarantine rule management		✓
SiteProtector management		✓
Update settings	✓	✓

Using IPS Local Management Interface

IPS Local Management Interface is the web-based management interface for IBM Security Network IPS appliances. Use IPS Local Management Interface to configure and manage an appliance locally.

Java™ Runtime Environment

The appliance must have the correct version of the Java Runtime Environment (JRE) installed in order to run IPS Local Management Interface. See the readme document to get the version number of the supported JRE.

Accessing IPS Local Management Interface

Access IPS Local Management Interface using a Web browser. Type **https://**<*appliance IP address*> to access the appliance using its IP address. If you are using a DNS server, type **https://**<*host name*>.

Managing with SiteProtector

SiteProtector is the IBM ISS management console. With SiteProtector, you can manage components and appliances, monitor events, and schedule reports. By default, your appliance is set up for you to manage it through IPS Local Management Interface. If you are managing a group of appliances along with other sensors, you may prefer the centralized management capabilities that SiteProtector provides.

SiteProtector management options

When you register the appliance with a SiteProtector group, you can do the following:

- Allow the appliance to inherit sensor group settings
- Manage some or all of settings for a single appliance in the group independently in SiteProtector, so that the appliance maintains those individual settings regardless of group settings

How the SiteProtector Agent Manager works

When you enable SiteProtector management, you assign the appliance to an Agent Manager. Agent Managers manage the command and control activities of various agents and appliances registered with SiteProtector and facilitate data transfer from appliances to the Event Collector, which manages real-time events it receives from appliances.

The Agent Manager sends any policy updates to appliances based on their policy subscription groups. (A subscription group is a group of agents or appliances that share a single policy.) Decide which group the appliance should belong before you register it with SiteProtector. Eventually, the group's policy is shared down to the appliance itself.

For more information about the Agent Manager, see the SiteProtector documentation or online Help.

How the appliance communicates with SiteProtector

When you register the appliance with SiteProtector, the appliance sends its first *heartbeat* to the Agent Manager to let the Agent Manager know that it exists. A heartbeat is an encrypted, periodic HTTP request the appliance uses to indicate it is still running and to allow it to receive updates from the Agent Manager. When you register the appliance with SiteProtector, you set the time interval (in seconds) between heartbeats.

Asset grouping

When the Agent Manager receives the heartbeat, it places the appliance in the group you specified when you set up registration. If you did not specify a group, it places the appliance in the default group "G-Series" or "Network IPS," depending on your version of SiteProtector. If you clear the group box when you register the appliance, it places the appliance in Ungrouped Assets.

Local settings or group settings

If you opted to allow local appliance settings to override group settings, then the appliance maintains its local settings at the first heartbeat. If you did not allow local appliance settings to override group settings, then the Agent Manager immediately "pushes" the group's policy files to the appliance, even if the group's policy settings are undefined. For example, if you set firewall rules on the appliance and then you register the appliance with a group that has no firewall rules defined, the group policy overwrites the local policy, and the appliance no longer has firewall rules enabled.

At the second heartbeat and each heartbeat thereafter, the Agent Manager "pushes" the group policy to the appliance. However, you can change some local appliance settings through SiteProtector. Any local policy settings you change for the appliance take precedence over the group policy settings for that appliance only; the group policy settings remain in effect for all other appliances in the group.

How appliance updates work with SiteProtector

After you register the appliance with SiteProtector, you must continue to update it regularly to maximize performance and to ensure it runs the most up-to-date firmware, security content, and database. Consider scheduling automatic database updates, security content updates, and firmware update downloads and installations.

Note: You can download and install firmware updates in IPS Local Management Interface even if the appliance is registered with SiteProtector.
Use the Update Settings page to schedule the following automatic update options:

- Downloading and installing firmware updates
- Downloading and installing security content updates
- Updating the database

How appliance events are handled in SiteProtector

You can specify the events that generate and deliver an alert to SiteProtector. When an event occurs, the appliance sends an alert to SiteProtector. You can use the event information in the alert to create valuable

reports. Alerts sent to SiteProtector still appear in the Alerts page in IPS Local Management Interface if the alerts are configured for logging.

Health and sensor alerts

Use the Alerts section to configure sensor and health alerts on your Network IPS appliance. You can configure sensor and health alerts to present in SiteProtector.

Sensor alerts

You can configure alert messages that notify you of appliance-related events. Determine what action the appliance takes when an event causes an alert, such as sending an SNMP trap in response to the event.

Table 1. Sensor alerts.

Alert	Description
Sensor Error	Alerts you when a sensor system error has occurred.
Sensor Warning	Alerts you about potential system problems.
Sensor Informative	Alerts you about user actions, such as changing passwords, downloading logs, or editing parameters.

Health alerts

You can configure alert messages that notify you of the health of the appliance. Determine what action the appliance takes when an event causes an alert, such as sending an email to the appliance administrator in response to the event.

Table 2. Health alerts.

Alert	Description
Health Error	Alerts you when the health of the appliance (system, security, network, and SiteProtector) has failed. For example, a health error alerts when an internal process has failed.
Health Warning	Alerts you when the health of the appliance (system, security, network, and SiteProtector) has the potential to fail. For example, a health warning alerts you when your license has expired.
Health Informative	Alerts you when the health of the appliance (system, security, network, and SiteProtector) is normal. For example, a health informative alerts you that the health of the appliance is normal because an expired license was updated.

In the Policy

In IPS Local Management Interface:

- **Manage System Settings** > **Appliance** > **Alerts Settings**

In SiteProtector:

- **Alerts** policy

Capacity planning

Use throughput graphs, driver statistics, and the SNMP GET request to gather information for capacity planning.

Throughput graphs

Throughput graphs show the sum of traffic, in megabits, into or going out of your network. Throughput graphs also show the totals for unanalyzed traffic and secured traffic moving through your network. Unanalyzed traffic has not been inspected by the Protocol Analysis Module (PAM). Secured traffic has been inspected by PAM, but this does not necessarily mean that the traffic is suspect. You can customize these graphs to show statistics for an hour, a day, a week, or a month. Choose the time period that best helps you to view the capacity of traffic and analysis on your appliance.

Find throughput graphs in **Monitor Health and Statistics** > **Network**.

Driver statistics

Driver statistics help with capacity planning because they report the totals for secured traffic, unanalyzed traffic, packets received, and packets transmitted for drivers. The four specific driver statistics that help with capacity planning are:

- **adapter.bytes.secured**: Lists the total number of bytes secured.
- **adapter.bytes.unanalyzed**: Lists the total number of bytes unanalyzed.
- **adapter.0.packets.received**: Lists the number of packets received on adapter 0 (adapter A).
- **adapter.0.packets.transmitted**: Lists the number of packets transmitted (forwarded from inline partner or injected) on adapter 0 (adapter A).

Find driver statistics in **Monitor Health and Statistics** > **Network**.

SNMP GET request and the MIB file

The SNMP GET request helps with capacity planning because you can configure it to retrieve statistics from the management information base (MIB) file. The MIB file includes information about these items:

- **network.driver.stats**: Contains all of the statistics found in **Monitor Health and Statistics** > **Network** > **Network Driver Statistics**.
- **protection.analysis.stats**: Contains all of the statistics found in **Monitor Health and Statistics** > **Security** > **Protection Analysis Statistics**.
- **network.protection.stats**: Contains all of the statistics found in **Monitor Health and Statistics** > **Security** > **Network Protection Statistics**.
- **ipmi.chassis.status** (only for GX7000 series appliances): Contains information about the status of the chassis along with information about power failures and driver failures. You can view this status in the Intelligent Platform Management Interface (IPMI).

Configure the SNMP GET request in **Manage System Settings** > **Appliance** > **SNMP**. Then use an SNMP tool to get the MIB file contents from MIB: NET-SNMP-EXTEND-MIB:nsExtendOutput1Table.

Note: The capacity planning feature is available only when the SNMP GET request is enabled and configured.

NTP servers

You can add Network Time Protocol (NTP) servers to your Network IPS appliance. NTP servers get the correct time of day from a specified source and synchronize the time of day for multiple components on your network.

NTP servers are useful for managing the time of day on networks that span different time zones and different continents. You can configure and manage the NTP policy from SiteProtector and apply it to all of your Network IPS appliances. The NTP policy uses symmetric key and autokey exchanges to authenticate.

Symmetric key

The server and the client use a common secret key for authentication. The advantages of symmetric key exchanges include minimal computing power usage, a relatively quick processing time, and the ability for both the sender and receiver to encrypt or decrypt. To configure symmetric key exchange, you need the key identifiers (key IDs), key types, and key values for your NTP servers. This option is available for only NTP versions 3 and 4.

Autokey

If both the server and the client are located on the outside of the firewall, they can use the autokey authentication. Autokey authentication uses certificate-based key exchanges that are also known as "challenge/response" exchanges. This method of authentication is best used to authenticate servers to clients. For example, this method works well if a central server outside the firewall authenticates to several lower strata servers that are also outside the firewall. These lower strata servers use internal hardware pieces (NICs) to provide NTP access to clients inside the firewall. This option is available for only NTP version 4.

Autokey exchanges use identity schemes to prove the identity of a remote system. Using identity schemes helps to prevent man-in-the-middle attacks. The appliance supports three identity schemes: Schnorr (IFF), Guillou-Quisquater (GQ), Mu-Varadharajan (MV).

FIPS and the NTP policy

The NTP policy is developed to meet the Federal Information Processing Standard (FIPS) 140-2. Before you configure the NTP policy to use FIPS options, ensure that the firmware version and hardware are FIPS certified. There is no advantage to configuring the NTP policy with FIPS options if your network does not need to comply with FIPS 140-2.

Symmetric key: To be FIPS-compliant, use only the cryptographic hash function SHA-1 in your symmetric key content. MD5 is not FIPS-compliant.

Autokey: To be FIPS-compliant, use the following options:

Setting	FIPS-compliant option
Message Digest Algorithm	SHA-1
Encryption Scheme	DSA-SHA-1

For specific information about IBM Security products that are FIPS certified, consult the IBM Security FIPS 140 Security Policy documents. Find these documents on the National Institute of Standards and Technology (NIST) web site, in the Module Validation Lists section: http://csrc.nist.gov/groups/STM/cmvp/index.html.

In the Policy

In IPS Local Management Interface:

- **Manage System Settings** > **Appliance** > **NTP Configuration**

In SiteProtector:

- **NTP Configuration** policy

Chapter 3. Configuring firewall settings

Using rule statements, you can configure firewall rules to block attacks based on various source and destination information in the packet. In addition, you can filter out traffic you do not want to be inspected if you are not interested in seeing it.

Configuring firewall rules

Firewall rules block attacks based on various source and target information in the packet. You can add firewall rules manually, or you can enable the appliance to construct rules using the values you specify. This feature offers you greater flexibility when configuring firewall settings.

Firewall rules behave differently in each mode. The following table describes how the appliance applies firewall rules according to the monitoring mode.

Table 3. Firewall rules and monitoring modes

Mode	Firewall rule behavior
Inline mode	An appliance in inline mode applies firewall rules to passing traffic according to the specified configurations.
Passive mode	An appliance in passive mode works like a traditional sensor and is not in the direct path of the packets. However, an appliance in passive mode can filter out traffic you do not want the appliance to inspect. To use firewall rules as a filter in passive mode, select the ignore response for your firewall rules
Inline simulation	An appliance in inline simulation mode still passes packets, but no actions are taken. Instead, the appliance reports what action would have been taken if it had been in inline mode.

Firewall rule criteria

You can define firewall rules using any combination of the following criteria:

- Interface
- VLAN range
- Protocol (TCP, UDP, or ICMP)
- Source or target IP address and port ranges

Firewall rule order

The appliance processes firewall rules in the order in which they are listed (from top to bottom). Correct ordering is mandatory. When a connection matches a firewall rule, further processing for the connection stops. The appliance ignores any additional firewall rules you have set.

Example:

Use the following statements to block all connections to a network segment except those destined for a specific port on a specific host:

© Copyright IBM Corp. 2003, 2012

```
adapter any ip src addr any dst addr 1.2.3.4 tcp dst port 80
(Action = "ignore")
adapter any ip src addr any dst addr 1.2.3.1-1.2.3.255
(Action = "drop")
```

Explanation:
The first rule allows all traffic to port 80 on host 1.2.3.4 to pass through to a Web server as legitimate traffic. All other traffic on that network segment is dropped.

If you reverse the rule order, all traffic to the segment is dropped, even the traffic to the Web server on 1.2.3.4.

Changing the order of firewall rules

To change the order of firewall rules use the **UP** or **Down** icons to move the rule.

Firewall rules and actions

The firewall supports several different actions that describe how the firewall reacts to the packets matched in the rules, or *statements*. The following table defines these actions:

Rule	Description
Ignore (Permit)	Allows the matching packet to pass through, so that no further actions or responses are taken on the packet. No further inspection is performed on the session.
Protect	Packets that match this rule are processed by PAM. Enables matching packets to be processed by normal responses, such as (but not limited to) logging, the block response, and quarantine response.
Monitor	Functions as an IP whitelist. Allows to packets that match the statements bypass the quarantine response and bypass the block response. However, all other responses still apply to the packet.
Drop (Deny)	Drops the packets as they pass through the firewall. Because the firewall is inline, this action prevents the packets from reaching the target system. The connection most likely makes several retry attempts, and then the connection eventually times out.
Drop and Reset	Functions in the same manner as the drop action, but sends a TCP reset to the source system. The connection terminates more quickly (because it is automatically reset) than with the drop action. **Note:** For all other protocols other than TCP, this option functions like the drop action.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Firewall** > **Firewall Rules**

In SiteProtector:

- **Firewall** policy

Firewall rules language

A firewall rule consists of several statements (or clauses) that define the traffic for which the rule applies. When you manually create firewall rules, use the syntax in this topic.

Firewall clauses

A firewall rule consists of several clauses chained together to match specific criteria for each packet. The clauses represent specific layers in the protocol stack. Each clause can be broken down into conditions and expressions. The expressions are the variable part of the rule in which you plug in the address, port, or numeric parameters.

You can use the following firewall clauses:

- **Adapter clause**

  Specifies a set of adapters (also known as "interfaces") from A through P that attaches the rule to a specific adapter. The adapter clause indicates a specific adapter where the rule is applied. The supported adapter expressions are any and the letters A through P. If you do not specify an adapter clause, the rule matches packets on any adapter.

  ```
  adapter (adapter-id)
  adapter A
  adapter any
  adapter A,C
  adapter A-C
  ```

- **Ethernet clause**

  Specifies either a network protocol type or virtual LAN (VLAN) identifier to match the 802.1 frame. You can use the Ethernet clause to filter 802.1q VLAN traffic or allow/deny specific types of Ethernet protocols. You can find the list of protocol types at http://www.iana.org/assignments/ethernet-numbers. Ethernet protocol constants can be specified in decimal, octal, hexadecimal, or alias notation. To make it easier to block specific types of Ethernet traffic, you can specify an alias instead of the well-known number. In some cases, the alias blocks more than one port (for example, IPX and PPPoE).

  ```
  ether proto (protocol-id)
  ether proto {arp|aarp|atalk|ipx|mpls|netbui|pppoe|rarp|sna|xns}
  ether vid (vlan-number)
  ether vid (vlan-number) proto (protocol-id)
  ```

  ```
  ether proto !arp
  ether vid 1 proto 0x0800
  ether vid 2 proto 0x86dd
  ether vid 3-999 proto 0x0800,0x86dd
  ```

- **IPv4 datagram clause**

  Specifies IPv4 addresses and the transport level filtering fields such as TCP/UDP source or destination ports, ICMP type or code, or a specific IP protocol number. The IP datagram clause identifies the protocol that resides inside the IP datagram and the protocol-specific conditions that must be satisfied in order for the statement to match. Currently, only ICMP, TCP, and UDP conditions are supported, but you can specify filters based on any IP protocol. If you do not specify an IP datagram clause, the statement will match any IP datagram protocol.

  The first and second statements below match IP packets that match the IP address expression. The third statement below matches IP packets that match the IP address expression. The fourth statement below matches IP packets that match the protocol type. The fifth statement is a combination of the first and second statements. The sixth statement is a combination of the first, second, and fourth statements.

  ```
  1. ip src addr <IPv4-addr>
  2. ip dst addr <IPv4-addr>
  3. ip addr <IPv4-addr>
  4. ip proto <protocol-type>
  5. ip src addr <IPv4-addr> dst addr <IPv4-addr>
  6. ip src addr <IPv4-addr> dst addr <IPv4-addr> proto <protocol-type>
  ```

Examples:

```
ip addr 192.168.10.1/24
ip addr 192.168.10.0-192.168.10.255
```

- **IPv6 datagram clause**

  The IPv6 datagram clause identifies the protocol that resides inside the IPv6 datagram and the protocol-specific conditions that must be satisfied in order for the statement to match. Currently, only ICMPv6, TCP, and UDP conditions are supported, but filters can be specified based on any IPv6 protocol. If no IPv6 datagram clause is specified, the statement will match any IPv6 datagram protocol. The first and second statements below block source and destination IPv6 packets that match IPv6 address expression. The third statement below blocks source or destination IPv6 packets that match IPv6 address expression. The fourth statement below blocks IPv6 packets that match the protocol type. The fifth statement is a combination of the first and second statements. The sixth statement is a combination of the first, second, and fourth statements.

  ```
  ipv6 src addr <ipv6-addr>
  ipv6 dst addr <ipv6-addr>
  ipv6 addr <ipv6-adr>
  ipv6 proto <protocol-type>
  ipv6 src addr <ipv6-addr> dst addr <ipv6-addr>
  ipv6 src addr <ipv6-addr> dst addr <ipv6-addr> proto <protocol-type>
  ```

  Examples:

  ```
  ipv6 addr FF01:0:0:0:0:0:0:101
  ipv6 addr 12AB:0:0:CD30::/60
  ipv6 addr FF01::101-FF01:0:0:0:0:0:0:200
  ```

Firewall conditions

TCP and UDP Conditions

You can specify TCP and UDP port numbers in decimal, octal, or hexadecimal notation. The port's value range is 0 through 65535.

```
tcp src port <TCP-UDP-port>
tcp dst port <TCP-UDP-port>
tcp dst port <TCP-UDP-port> src port <TCP-UDP-port>
udp src port <TCP-UDP-port>
udp dst port <TCP-UDP-port>
udp dst port <TCP-UDP-port> src port <TCP-UDP-port>
```

ICMPv4 conditions

You can specify ICMP conditions in decimal, octal, or hexadecimal notation. You can find the valid number for type and code at http://www.iana.org/assignments/icmpparameters.

```
icmp type (protocol-type)
icmp code (message-code)
icmp type (protocol-type) code (message-code)
```

ICMPv6 conditions

You can specify ICMPv6 conditions in decimal, octal, or hexadecimal notation. You can find the valid number for type and code at http://www.iana.org/assignments/icmpparameters.

```
icmpv6 type <protocol-type>
icmpv6 code <message-code>
icmpv6 type <protocol-type> code <message-code>
```

Expressions

An expression describes a list of header values that must match the clause's protocol parser. Each clause is directly responsible for matching a specific layer in the protocol stack. The syntax and accept range of values is controlled by the clause. The expression can be a single value, a comma separated list of values, or a range set. Currently, expressions exist to specify adapter numbers, IPv4 addresses, IPv6 addresses, TCP and UDP port numbers, ICMP message type and codes, and IP datagram protocol numbers.

```
(value)
(value), (value)
(value)-(value)
```

Expressions that begin with an exclamation marks (!) are called *a not-expressions*. Not-expressions will match all values except those you specify. Not-expressions that do not match any values will generate an error.

IPv4 address expression examples

The <n> can be either hex or decimal number in a range from 0 to 255. All hex numbers must have a 0x prefix.

Single address

```
n.n.n.n
```

Address list

```
n.n.n.n, n.n.n.n
```

Specific address using CIDR format; netmask value must range from 1 to 32

```
n.n.n.n/<netmask>
```

Address range, where first value is smaller than last

```
n.n.n.n - n.n.n.n
```

IPv6 address expression examples

The value for <n> must be a hexidecimal digit (0-F). Any four-digit group of zeroes within an IPv6 address may be reduced to a single zero or omitted.

Single address

```
nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn
```

Address list

```
nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn, nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn
```

Specific address using CIDR format; netmask value must range from 1 to 128

```
nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn/<prefix>
```

Address range, where first value is smaller than last

```
nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn - nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn:nnnn
```

TCP/UDP ports, protocol identifiers, or numbers

The values listed for any constant must be within the fields required range; otherwise the parser will refuse the parse clause.

```
0xFFFF
65535
0, 1, 2
0 - 2
! 3 - 65535
```

Complete firewall rule examples

The following statements are examples of complete firewall rules. If you do not specify a protocol, the rule uses the *any* protocol.

- `adapter A ip src addr <ip_address>`
- `adapter A ip src addr <ip_address> dst addr any tcp src port 20 dst port 80`
- `adapter any ip src addr any dst addr <ip_address>`
- `adapter any ip src addr any dst addr any icmp type 8`
- `tcp`
- `adapter B icmp`
- `udp`
- `adapter A ipv6 src addr <ipv6_addr>`
- `adapter A ipv6 src addr <ipv6_addr> dst addr any tcp src port 20 dst port 80`
- `adapter any ipv6 src addr any dst addr <ipv6_addr>`
- `adapter any ipv6 src addr any dst addr any icmpv6 type 128`
- `ipv6 tcp`
- `adapter B icmpv6`

Chapter 4. Working with security events

This chapter describes how to configure security events and response filters. Use these features to control how the appliance responds to and reports security events that occur on the network.

Configuring security events

The Security Events page lists hundreds of events by attack type and audit. A security event is network traffic with content that can indicate an attack or other suspicious activity. These events are triggered when the network traffic matches one of the events in the active security policy. You can edit events in the security policies to meet the network's needs.

Editing multiple security events

You can select multiple security events by:

- Selecting multiple events, pressing CTRL, and then selecting each event
- Selecting a range of events, pressing SHIFT, and then selecting the first and last events in the range

Note: Every item you edit is changed for every selected event.

Visual indication of changes

A blue triangle icon appears next to any item in the selected events that has a different value. If you change the value of a field with this icon, the value changes to the new setting for all selected events and the blue triangle icon no longer appears next to the field.

For example, you have two events that you have not marked to block. You enable the block response for one event. A blue triangle appears next to Block for the edited event. If you enable the block response for the other event, then both events have blocking enabled, and the blue triangle disappears.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Advanced IPS** > **Security Events**

In SiteProtector:

- **Security Events**

Viewing security event information

The Security Events page lists hundreds of security events according to attack type and audit. You can customize how events appear to make viewing and searching easier.

About filters and regular expressions

Security event filters use regular expressions to limit the number of events displayed.

Regular expressions (also known as regex) are sets of symbols and syntax that you can use to search for text that matches the patterns you specify.

At the most basic level, the following wildcard search types are supported:

© Copyright IBM Corp. 2003, 2012

Search value	Returns
.*	All events
http.*	All events that begin with “http”
.*http	All events that end in “http”
.*http.*	All events that contain “http”

Regular expressions search all columns in the Security Events list. If you search for http*, for example, the search returns all events that match the http protocol column and all events that begin with http.

Displaying and grouping security events

Before selecting or grouping security events, click the appropriate icon. This action displays a window in which you decide what columns you want to display or group.

Viewing security events

You can use Filter feature to help you focus on the security events that interest you most. Click the Filter checkbox, and type the regular expression by which you want to filter.

Chapter 5. Configuring other intrusion prevention settings

This chapter describes how to configure and manage other intrusion prevention settings, such as user-defined events, connection events, and OpenSignature events. It discusses how to manage quarantined intrusions, view global tuning parameters for the appliance, and monitor X-Force blocking.

Managing quarantined intrusions

The Quarantined Rules page shows quarantine rules that were dynamically generated in response to detected intruder events. When quarantine response is enabled, the rules specify the packets to block and the length of time to block them. They prevent worms from spreading, and deny access to systems infected with backdoors or trojans. You can manually add and delete your own quarantine rules. However, you cannot edit existing rules.

Single-click blocking

From the Security Alerts Logs page, you can click an event and select to **Block Intruder**. This adds a rule to the Quarantine Rules page for the source IP address reported in the event. The appliance blocks all traffic to and from that IP address for the time specified in the rule. Delete quarantine rules added by the single-click blocking feature when they no longer apply. Otherwise, the appliance removes the rules automatically when the rules expire.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Response Tuning** > **Quarantine Rules**

Important: You can view or remove quarantined intrusions only through IPS Local Management Interface.

Configuring connection events

Connection events are user-defined notifications of open connections to or from particular addresses or ports. They are generated when the appliance detects network activity at a designated port, regardless of the type of activity, or the content of network packets exchanged.

The Connection Events page lists pre-defined connection events for different connection types, such as WWW, FTP, or IRC. Use this page to customize these events or to create your own events to cover the traffic you need to monitor.

For example, you can define a rule that causes a connection event to alert the console whenever someone connects to the network using FTP.

Note: The connections are always registered against the destination port you specify, so to monitor an FTP connection, you must use the FTP port. One entry per connection is sufficient for traffic in each direction.

How connection events work

Connection events occur when network traffic connects to the monitored network through a particular port, from a particular address, with a certain network protocol. The appliance detects these connections

© Copyright IBM Corp. 2003, 2012

using packet header values. Connection events do not necessarily constitute an attack or other suspicious activity, but they are network occurrences that might interest a Security Administrator.

Note: Connection events do not monitor the network for any particular attack signatures. You use security events to monitor for these types of attacks. See "Configuring security events" on page 21 for more information.

About removing connection events

You can remove any connection event from the list. However, if you edit a pre-defined connection event and later decide you want to remove it, be aware that the event is not returned to its pre-defined state. The event is removed from the list entirely. If you want to use this event again, it is no longer available.

Best practice for disabling instead of deleting

Consider disabling the event and keeping it in the list. This way, if you want to use it again at another time, the event is still available to you in some form.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Advanced IPS** > **Connection Events**

In SiteProtector:

- **Connection Events** policy

Configuring user-defined events

The events that are enabled in a policy control what an appliance detects. Create user-defined events around contexts, which specify the type and part of a network packet you want the appliance to scan for events.

New user-defined events

As you add user-defined events, the new events appear at the bottom of the list.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Advanced IPS** > **User Defined Events**

In SiteProtector:

- **User Defined Events** policy

User-defined event contexts

When you create a user-defined event, select a context that tells the appliance the type and particular part of a network packet to monitor for events. After you specify the context, add a string that tells the appliance exactly what to look for when it scans the packet. See "Regular expressions in user-defined events" on page 27 for more information.

The following user-defined event contexts are available:

- DNS_Query context

- Email_Receiver context
- Email_Sender context
- Email_Subject context
- File_Name context
- News_Group context
- Password context
- SNMP_Community context
- URL_Data context
- User_Login_Name context
- User_Probe_Name context

The following table list each user-defined event context, describes what each context monitors, and provides examples.

Context	Examples
DNS_Query. Monitors the DNS name in DNS query and DNS reply packets over UDP and TCP. The appliance compares the information in the String box to the expanded (human-readable) version of the domain name in these packets. If a user accesses a site directly using an IP address, the DNS lookup does not occur, and the appliance cannot detect the event. **Note:** To monitor for a particular URL, remember that the domain name is only the first element. For example, //www.news.com is the first element in http://www.news.com/ stories. Use the URL_Data context (see "URL_Data context") to detect the rest of the URL.	You could use the DNS_Query context along with a string value of `www.microsoft.com` to monitor users accessing the Microsoft Web site. If you are concerned about users on your site accessing hacker-related materials on the Internet, you could monitor access to domains such as the following: • hackernews.com • rootshell.com
Email_Receiver. Monitors incoming or outgoing e-mail to a particular recipient by monitoring the receiver address part of the e-mail header using the SMTP, POP, IMAP protocols. When the appliance detects an event that matches a signature using the Email_Receiver context, you can determine which protocol the e-mail used by examining the details of the event. **Note:** This context does not monitor e-mail sent with the MAPI protocol.	If you suspect that someone is using "social engineering" to manipulate certain employees, you can monitor inbound e-mail to those employees' addresses and log the source IPs. Or if you suspect someone is leaking proprietary information within your company to a particular outside e-mail address, you could track e-mail to that address.
Email_Sender. Use the Email_Sender context to monitor incoming or outgoing e-mail from a particular sender. The Email_Sender context monitors the sender address part of the e-mail header using the SMTP, POP, IMAP protocols. When the appliance detects an event that matches a signature using the Email_Sender context, you can examine the details of the event to control which protocol the e-mail used. **Note:** This context does not monitor e-mail sent with the MAPI protocol.	Use the Email_Sender context to detect instances of social engineering or other employee manipulation (inbound) or to detect information leaks from your company (outbound).
Email_Subject. Monitors the subject line in the e-mail header of messages using the SMTP, POP, and IMAP protocols. **Note:** This context does not monitor e-mail sent with the MAPI protocol.	You can create events to detect information leaks by monitoring for important project names or file names. You can use Email_Subject to detect viruses, such as the I LOVEYOU virus. **Tip:** Because viruses and other attacks have developed programs that systematically change the subject line, use the Email_Content context to track these virus types.

Context	Examples
File_Name. The File_Name context detects when a person or a program attempts to remotely read a file or write to a file with any of the following protocols: • TFTP • FTP • Windows file sharing (CIFS or Samba) • NFS **Note:** NFS can open files without directly referencing the file name. Using this context to monitor NFS access to a file may not be 100% effective.	When the Explorer worm of 1999 propagated over a Windows network, it attempted to write to certain files on remote Windows shares. With a worm like this, you can monitor for attempts to access files and stop the worm from propagating locally.
News_Group. Use the News_Group context to monitor the names of news groups that people at your company access. The News_Group context monitors people accessing news groups using the NNTP protocol.	You can use the context to detect subscriptions to news groups, such as hacker or pornography groups, that are inappropriate according to your company's Internet usage policy.
Password. Use the Password context to identify passwords passed in clear text over the network. When a password is not encrypted, an attacker can easily steal it by monitoring traffic with a sniffer program from another site. The Password context monitors programs or users sending passwords in clear text using the FTP, POP, IMAP, NNTP or HTTP protocols. You can use the Password context to do the following: • Monitor compromised accounts to gain forensic data • Monitor the accounts of terminated employees • Detect the use of default passwords **Note:** This context does not monitor encrypted passwords.	• **Monitoring compromised accounts:** After cancelling a compromised account, you can create an event to monitor outside attempts to use it and find the person that accessed the compromised data. • **Monitoring terminated employee accounts:** Add searches for terminated employees' passwords to detect unauthorized remote access attempts to their closed accounts. • **Detecting the use of default passwords:** Set up events to look for default passwords relevant to your site to detect attackers probing for common vulnerabilities. **Note:** The X-Force database contains many records detailing the names of such accounts. For more information about default passwords, look up passwords in the X-Force database at http://xforce.iss.net.
SNMP_Community. Monitors the use and possible abuse of SMNP community strings. The SNMP_Community context monitors any packet containing an SNMP community string. An SNMP community string is a clear text password in an SNMP message. This password authenticates each message. If the password is not a valid community name, then the message is rejected. If an unauthorized person gains knowledge of your community strings, that person could use that information to retrieve valuable configuration data from your equipment or even to reconfigure your equipment. **Important:** Consider using highly unique community strings that you reconfigure periodically.	• **Detects people trying to use old strings:** If you change the SNMP community strings, create an event using this context to have the appliance search for people trying to use the old strings. • **Detects the use of default strings:** The X-Force database contains information about several vulnerabilities involving default community strings on common equipment. Attackers can attempt to access to your equipment by using these default passwords. To have the appliance detect this activity, create events using this context to monitor for the default passwords relevant to the equipment at your site. These events can detect attackers attempting to probe for these common vulnerabilities. **Note:** If you scan your network using Internet Scanner®, a rule using this context to check for SNMP community strings may detect many instances of this event in response to a SNMP scan. **Reference:** For more information about default passwords, look up SNMP in the X-Force database at http://xforce.iss.net.

Context	Examples
URL_Data. Monitors various security issues or policy issues related to HTTP GET requests. An HTTP GET request occurs when a client, such as a Web browser, requests a file from a Web server. The HTTP GET request is the most common way to retrieve files on a Web server.The URL_Data context monitors the contents of a URL (minus the domain name or address itself) for particular strings, when accessed through an HTTP GET request. **Note:** This context does not monitor the domain name associated with an HTTP GET request.	Use this context to have the appliance monitor for attacks involving vulnerable CGI scripts. IBM ISS Advisory #32, released on August 9, 1999, describes how to use this context to search for an attempt to exploit a vulnerability in a Microsoft Internet Information Server component. **Reference:** For more information, see Vulnerabilities in Microsoft Remote Data Service at http://xforce.iss.net/alerts/advise32.php.You could use this context to generically search whether employees are using computers to access company-banned sites, such as pornography sites.
User_Login_Name. Detects user names exposed in plain text during authentication requests. This context works for many protocols, so you can use it to track attempts to use a particular account no matter what protocol the attacker uses. The User_Login_Name context monitors for plain text user names in authentication requests using the FTP, POP, IMAP, NNTP, HTTP, Windows, or R* protocols.	Use this context to track attempts to use compromised accounts or if you suspect recently dismissed employees have attempted to access their old accounts online. If you know the account named "FredJ" was compromised in an attack, configure an event using this context to search for attempts to access the account.
User_Probe_Name. Identifies attempts to access to computers on your network using default program passwords. The User_Probe_Name context monitors any user name associated with FINGER, SMTP, VRFY, and SMTP EXPN. An attacker can use these default accounts to access to your servers or other computers in the future.	Like the Password and SNMP_Community contexts, you can use the X-Force database to build a list of default accounts and passwords relevant to the systems and software on your network. **Reference:** For more information about default passwords, look up SNMP in the X-Force database at http://xforce.iss.net.

Regular expressions in user-defined events

Regular expressions (strings) are a combination of static text and variables that the appliance uses to detect patterns in the network packets (contexts) that you specify for user-defined events. Use regular expressions if you want the appliance to detect more than a single static text string.

Limitations for regular expressions

Some limitations apply to user defined expressions.

- The limit for regular expressions is 128 bytes.
- The number of regular expressions for a single context is limited to 16.

These values are subject to change. For the latest values, see the IBM Support Portal at http://www.ibm.com/support/entry/portal. Search for technote 1435274.

Regular expression library

The appliance uses a custom IBM Security regular expression library called Deterministic Finite Automata or DFA regular expression.

Changing the order of precedence

Use parentheses in these regular expressions to offset the standard order of precedence.

Example: The natural order of precedence would interpret 4+2*4 as 12, because in the natural order of precedence, multiplication takes precedence over addition. However, you can use parentheses to change

this precedence. For example, if you use (4+2)*4, the answer would be 24 instead of 12. This example describes a mathematical use of the order of precedence, but many other non-numerical uses exist.

Reference: For more information about the order of precedence or other information about using regular expressions, see *Mastering Regular Expressions: Powerful Techniques for Perl and Other Tools (O'Reilly Nutshell)* by Jeffrey E. Friedl (Editor), Andy Oram (Editor).

Regular expression syntax

You can use the following regular expression syntax in a user-defined event:

Meta-Character	Description
(r)	Matches r
x	Matches x
xr	Matches x followed by r
\s	Matches either a space or a tab (not a newline)
\d	Matches a decimal digit
\"	Matches a double quote
\'	Matches a single quote
\\	Matches a backslash
\n	Matches a newline (ASCII NL or LF)
\r	Matches a carriage return (ASCII CR)
\t	Matches a horizontal tab (ASCII HT)
\v	Matches a vertical tab (ASCII VT)
\f	Matches a formfeed (ASCII FF)
\b	Matches a backspace (ASCII BS)
\a	Matches a bell (ASCII BS)
\ooo	Matches the specified octal character code
\xhhh	Matches the specified hexidecimal character code
.	Matches any character except newline
\@	Matches nothing (represents an accepting position)
“”	Matches nothing
[xy-z]	Matches x, or anything between y and z inclusive (character class)
[^xy-z]	Matches anything but x, or between y and z inclusive • The caret must be the first character, otherwise it is part of the set literally • Enter the dash as the first character if you want to include it
“text”	Matches text literally without regard for meta-characters within, and the text is not treated as a unit
r?	Matches r or nothing (optional operator)
r*	Matches zero or more occurrences of r (kleene closure)
r+	Matches one of more occurrences of r (positive kleene closure)
r{m,n}	Matches r at least m times, and at most n times (repeat operator)

Meta-Character	Description
r|l	Matches either r or l (alternation operator)
r/l	Matches r only if followed by l (lookahead operator)
^r	Matches r only at the beginning of a line (bol anchor)
r$	Matches r only at the end of the line (eol anchor)
r, l	Matches any arbitrary regular expression
m, n	Matches an integer
x,y,z	Matches any printable or escaped ascii character
text	Matches a sequence of printable or escaped ascii characters
ooo	Matches a sequence of up to three octal digits
hhh	Matches a sequence of hex digits

Tip for DNS name search

Since a period is a wildcard character that matches any character, escape any periods using a back slash in a DNS name search. **Example**: www\.ibm\.com

Configuring tuning parameters

Tuning parameters affect intrusion prevention settings at the group and site levels.

Edit and configure tuning parameters for groups of appliances managed through SiteProtector. View the parameters that affect a specific appliance at the site level with IPS Local Management Interface.

You can tune the following components on a group of appliances:

- Intrusion prevention responses
- Intrusion prevention security risks
- Firewall logging
- Updates

Default Values

Tuning parameters consist of name/value pairs. Each name/value pair has a default value. For example, the parameter **np.firewall.log** is a parameter that determines whether to log the details of packets that match firewall rules you have enabled. The default value for this parameter is **on**.

Commonly used tuning parameters are listed on the Tuning Parameters page. You can add tuning parameters to the list on this page and to the list of advanced parameters on the Update Settings page. Even if a tuning parameter is not listed on either page or not enabled, its behavior is still controlled by the default values defined for it. To change the behavior of a tuning parameter, you must configure it, enable it, and then apply a default value that includes the wanted behavior.

Configuring OpenSignatures

OpenSignatures use a flexible rules language to allow you to write customized, pattern matching IDS signatures to detect specific threats that are not already preemptively covered in IPS products. This feature is integrated into the IBM Protocol Analysis Module (PAM) as a rule interpreter.

Risks associated with OpenSignatures

The capabilities of custom signature development are very broad. With this flexibility comes added risk. Poorly written rules or signatures could impact sensor performance or have other consequences. Risks of using your own custom signatures include but are not limited to the following:

- Unacceptable appliance performance
- Throwing PAM into an infinite loop
- Blocking all network traffic to a specific segment (inline mode with or without bypass)

CAUTION:
IBM Security does not guarantee appliance performance if you choose to use OpenSignatures. Enable this functionality at your own risk. IBM Customer Support is not available to help you write or troubleshoot custom rules for your environment. If you require assistance to create custom signatures, please contact IBM Professional Services.

OpenSignatures syntax

The syntax options for each custom rule are as follows:

(action): alert

(protocol): tcp, udp, icmp, ip

(IP and netmask): single IP address (a.b.c.d), range of IP addresses (a.b.c.d-w.x.y.z), network address using CIDR notation (a.b.c.0/24)

Important: If you have improperly formatted an OpenSignature rule, you may receive a PAM configuration error response. However, PAM configuration error responses are not enabled by default. Consider enabling this feature in the security events policy because this ensures you receive notifications concerning improper syntax in OpenSignature rules.

The negation operator

The negation operator is indicated with an '!':

alert tcp ! 192.168.1.0/24

An alert prompts you when anything other than what is indicated with the '!' is used.

Enabling the OpenSignatures Parser

Use the settings indicated in the following table to enable the OpenSignatures Parser.

Setting	Description

Name	Type either of the following to enable OpenSignatures:
	`engine.opensignature.enabled`
	`pam.trons.enabled`

Value	Type the following: `true`

The default response for OpenSignatures

The default response for all OpenSignature events is **DISPLAY**. Due to this default response, IPS Local Management Interface and SiteProtector both report the OpenSignature event. If you want to edit the default response, use tuning parameters. With tuning parameters, you can configure features such as notification and protection responses.

Examples:

```
np.opensignature.user.response=DISPLAY:WithoutRaw;EMAIL:admin,Block:Default
np.opensignature.response=block-connection'
```

In the policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **OpenSignatures**

In SiteProtector:

- **OpenSignature Events** policy

Configuring SNORT

SNORT is an open source intrusion prevention and detection system that is integrated into the Network IPS appliance. You can use this integrated system along with the appliances native Protocol Analysis Module (PAM) to protect the network from intrusions.

SNORT on the Network IPS

Note: For specific configuration information about the integrated SNORT system, see the help system in IPS Local Management Interface or in SiteProtector.

Along with offering its own capabilities, the integrated system sends responses for SNORT activity. It lists SNORT event information and generates quarantine rules for these events. The system supports SNORT rules with TCP reset commands. It includes the rule profiling feature to report performance metrics for SNORT rules.

The integrated SNORT system on the Network IPS appliance includes three sections: command-line functions, configuration contents, and rules.

Section	Description
Command-line	Enables the SNORT engine to run and dictates command-line options such as rule order processing, expressions, and packet capture features.
Configuration contents	Includes configuration contents and the configuration file that contains variable definitions, preprocessors, output modules, and other objects to control operations. This piece also contains a rule profiling option.
Rules	Includes the rule files and lists the SNORT rules designed to protect the vulnerabilities on the network.

Risks

If you know how to use SNORT, the system offers customized protection against a vast range of threats. However, if not used properly, the SNORT system can burden the appliance with errors and hinder its

performance. Do not use the integrated SNORT system if you are not familiar with SNORT. IBM Customer Support is not available to help write or troubleshoot custom SNORT rules and configuration contents.

Use the information in this document to manage the integrated SNORT system on the Network IPS appliance. For the latest information on SNORT itself, including rules, documentation, and community forums, go to https://www.snort.org.

Considerations

SNORT rules

- Use an appropriate SNORT rule syntax checker to review the integrity of your rules because the integrated system does not check rule syntax.
- Import no more than 9000 SNORT rules from a rules file. Importing more rules at one time impacts IPS Local Management Interface and SiteProtector console performance.
- Import SNORT rules files no bigger than 5 MB. Importing bigger SNORT rules files impacts IPS Local Management Interface and SiteProtector console performance.
- The Network IPS appliance does not support the use of dynamic rules for SNORT.
- The integrated system supports quarantine rules for actively responding to unwanted traffic. It also supports the use of SNORT TCP reset rules for actively responding to unwanted traffic.
- The integrated system processes rules with duplicate SIDs and revision numbers by inspecting traffic with the rule that was last entered. The system ignores the previous rule.
- Use event filters in the configuration file to manage SNORT rules that cause an excessive number of alerts.

SNORT configuration

- The Network IPS appliance does not support the use of third-party preprocessors.
- Review and adjust the settings and directories in the configuration file (either the default configuration file or an imported configuration file) so that the file works for your environment.
- If you import a SNORT.conf file, delete rule path variables. Examples of rule path variables:
  - `var PREPROC_RULE_PATH ../preproc_rules`
  - `var WHITE_LIST_PATH /etc/snort/rules`

Performance

- **Important:** Use SNORT rule profiling only when needed because it can affect SNORT engine performance.
- High SNORT rule activity can burden the appliance. Use the secured and unanalyzed throughput statistics to determine the capacity of your SNORT rule activity. Find these throughput statistics in the Network Dashboard. Low secured traffic and high unanalyzed traffic might indicate high SNORT rule activity.

General

- The integrated system does not support the block response because the integrated SNORT system is not inline. It is in IDS mode.
- The SNORT system sends TCP resets in response to unwanted TCP connections through the TCP reset port.
- The SNORT system sends ICMP port unreachable messages in response to unwanted UDP connections through the TCP reset port.

SNORT and PAM

SNORT and PAM (Protocol Analysis Module) analyze the same data packets independently. Due to this design, unexpected behavior is possible from each system.

The appliance delivers a single queue of packets to PAM and to the integrated SNORT system. The appliance does not apply a processing order to the queue. The system to get to the packet first, analyzes it first. If the first system alters the packet or responds to it, then the second system analyzes a modified packet or responds to a packet that was already responded to. The outcome of this relationship is that you might see unexpected events or quarantine rules.

Action	Outcome
PAM analyzes first	PAM analyzes a packet before SNORT does, and PAM drops the packet. SNORT analyzes the same packet later, and generates an event. The unexpected outcome is that SNORT generated an unnecessary event from a packet that PAM dropped earlier.
SNORT analyzes first	SNORT analyzes a packet before PAM does, and SNORT generates an event. A quarantine rule is created from the event. It is a packet that PAM drops after it analyzes it but PAM has yet to reach the packet. SNORT sees the same packet because PAM did not yet respond. SNORT generates another event and another quarantine rule is created. PAM analyzes the packet later and drops the traffic. The unexpected outcome is that SNORT generated duplicate events and duplicate quarantine rules were created before PAM responded to the packet.

SNORT and high availability (HA) mode

You have the option of configuring the SNORT system to inspect or not inspect mirrored ports in HA mode. The following table outlines the behavior for each option:

Option	Action
Inspect (Enable)	The SNORT systems running on appliances in an HA pair inspect packets from mirrored ports. This behavior applies to pairs running in inline protection or inline simulation mode. This option increases the possibility of duplicate global responses and SiteProtector alerts. However, this option decreases the chance of the SNORT systems to miss attacks because the systems analyze all packets, including packets from mirrored ports.
Not inspect (Disable)	The SNORT systems running on appliances in an HA pair do not inspect packets from mirrored ports. This behavior applies to pairs running in inline protection or inline simulation mode. This option minimizes the possibility of duplicate global responses and SiteProtector alerts. However, this option limits the ability of the SNORT systems to analyze all traffic. **Important:** When this option is disabled, it is possible for one of the SNORT systems to miss an attack. Also, the quarantine rules generated from SNORT events might be out of sync on the appliances in the HA pair.

Troubleshooting SNORT Errors

The integrated SNORT system identifies errors one error at a time. Due to this process flow, you must troubleshoot and fix each error to successfully apply the SNORT policy.

Errors: SNORT errors occur when the integrated system detects configuration contents or rules that it identifies as invalid. In IPS Local Management Interface and in SiteProtector, the appliance presents a message that the policy failed to apply if you submit settings with errors on the **SNORT Configuration** or **SNORT Rules** tab. The error message includes information from SNORT to help fix the issue. For SNORT rule errors, the message lists the SID and message string. The system reports the policy failure as a significant event.

Tip: Use a syntax checker on SNORT rules to help decrease the number of invalid rules.

Troubleshooting: Troubleshooting the integrated SNORT system is an iterative process because it identifies one error at a time. When the system detects an error, it fails to apply the policy settings and reports the failure. You must troubleshoot the error before you can successfully apply the policy settings. After you fix the error, you must reapply the settings. If the system finds no other errors in the configuration contents or in the rules, then it reapplies the policy settings successfully. However, if the system finds other errors, it repeats this process for each one.

Note: To find the health status of the SNORT engine, go to **Monitor Health and Statistics** > **Security** > **Dashboard**.

The SnEP

The SNORT event processor (SnEP) scrapes errors from the SNORT system and reports SNORT errors in the following ways:

- The SnEP generates a significant event. The SnEP identifies the event as a `[SNORT ERROR]` and SNORT dictates the error message string.
- The SnEP logs the error to the system.
- The SnEP sends an alert to SiteProtector.

SNORT and quarantine functions

Configure quarantine rules and send quarantine responses for events generated from suspicious activity identified by the integrated SNORT system. Quarantine responses block intruders, including worms and Trojan horses, when the system detects events. Quarantine rules are manually added and dynamically generated in response to detected intruder events. These rules prevent worms from spreading and deny access to systems that are infected with back doors or Trojan horses. These rules also help prevent data leakage after an attack.

Importing and deleting SNORT rules

The Network IPS appliance imports and manages SNORT rules from a rules file according to customized settings and programmed behavior.

Customizing attributes to imported rules: When you import SNORT rules from a rules file, the appliance groups those rules by file name. You can customize these attributes of the imported rules:

- Enabled
- Rule String

  Note: You can change the rule string attribute. However, if you import an updated version of the rule file, the appliance does not reapply the changes. Changes to this attribute are lost.
- Comment
- Display
- Severity
- Responses (Email, Quarantine, SNMP, User Specified)

The Network IPS appliance stores these customized attributes so that it can reapply them all (except the rule string) after you import an updated file.

Reimporting updated or changed rules files: The appliance stores customized attributes because, in certain situations, it is necessary to reimport rules files that contain updates and changes. The appliance processes rules in reimported files in the following ways:

- If a rule is new to the updated file, the appliance adds the rule to the group.

- If a rule is deleted from the updated file, the appliance deletes that rule from the group. You must add the rule using the Add icon if you still need the rule.
- If a rule continues to exist in the updated file, the appliance applies the customized attributes to the updated version of the rule.

Note: The integrated system processes rules with duplicate SIDs and revision numbers by inspecting traffic with the rule that was last entered. The system ignores the previous rule.

Deleting SNORT rules: The appliance does not keep a record of past and deleted rules. If you delete a rule and then reimport a rules file that contains the deleted rule, the appliance adds the rule back into the SNORT policy.

SNORT rule profiling

Important: Use the SNORT rule profiling feature only when needed because it can impact SNORT engine performance.

Use SNORT rule profiling to analyze the performance of your SNORT rules and for troubleshooting possible performance issues. When enabled, the appliance produces a SNORT rule profiling file that you can view or download. This file includes performance statistics for the rules with the most offensive numbers. Some considerations concerning SNORT rule profiling are:

- You can access this feature through IPS Local Management Interface only.
- You must enable the SNORT engine and SNORT rule profiling for this feature to work.
- You do not have to enter contents or preprocessors for this feature. The Network IPS appliance already includes this feature.

You can sort your SNORT rule profiling file by the following statistics:

Statistic	Description
Checks	The number of times the SNORT engine checks for rule options after the SNORT engine completes an initial analysis to group and pre-screen traffic.
Matches	The number of times the SNORT engine finds traffic matching all rule options.
No Matches	The number of times the SNORT engine finds no traffic that matches all rule options.
Average Ticks (Avg/Check)	The average time the SNORT engine takes to check each packet against the listed rule.
Average Ticks Per Match (Avg/Match)	The average time the SNORT engine takes to check each packet that matches all rule options.
Average Ticks Per No Match (Avg/Nonmatch)	The average time the SNORT engine takes to check each packet that did not generate an event. **Note:** This statistic represents wasted time spent checking clean traffic.
Total Ticks	The rules responsible for using the most processing time.

For detailed information about SNORT rule profiling statistics, go to https://www.snort.org.

Unsupported SNORT configuration options

The Network IPS appliance does not support these options for SNORT configuration.

```
config alert_with_interface_name
config alertfile
config chroot
config daemon
config daq
config daq_dir
```

```
config daq_list
config daq_mode
config daq_var
config interface
config logdir
config no_promisc
config nolog
config pkt_count
config policy_mode
config profile_rules
config quiet
config response
config snaplen
config umask
config min_ttl
config new_ttl
include
output
preprocessor normalize_ip4
preprocessor normalize_ip6
preprocessor normalize_icmp4
preprocessor normalize_icmp6
preprocessor normalize_tcp
```

SNORT expression examples

Set SNORT expressions in the command-line area, which is on the **SNORT Execution** tab. SNORT expressions are like TCPDump expressions. An expression has one or more primitives. A primitive includes an ID (name or number) preceded by one or more qualifiers. The three main qualifiers in expressions are type, dir, and proto.

Qualifiers	Types
type	Identifies what the ID name or number refers to. Examples: • host: Looks for traffic based on IP address.**host 1.2.3.4** • net: Captures an entire network using CIDR notation. **net 1.2.3.0/24** • port: Inspects traffic to or from a certain port. **port 3389** • portrange: Inspects traffic on any port in a range. **portrange 21-23**
dir	Specifies the direction. Examples: • src: Finds traffic from a source only and eliminates one side of a host conversation. **src 2.3.4.5** • dst: Finds traffic from a destination only and eliminates one side of a host conversation. **dst 3.4.5.6**
proto	Restricts matches to particular protocols. You do not have to type proto. Examples: • tcp: Restricts matches to TCP traffic. **tcp** • icmp: Restricts matches to ICMP traffic.**icmp** • udp: Restricts matches to UDP traffic. **udp**

Examples of combining all three qualifiers:

- **src port 1025 and tcp**

- **udp and src port 53**

In the policy

In IPS Local Management Interface:
- **Secure Protection Settings** > **Advanced IPS** > **SNORT Configuration and Rules**
- **Review Analysis and Diagnostics** > **Diagnostics** > **SNORT Rule Profiling**

In SiteProtector: **SNORT Configuration and Rules** policy

Configuring response filters

Use response filters to control the number of events to which the appliance responds and the number of events reported to the management console.

Use response filters to do the following:
- Configure responses for security events that trigger based off network criteria specified in the filter
- Reduce the number of security events an appliance reports to the console

Example

If you have hosts on the network that are secure and trusted or hosts that you want the appliance to ignore for any other reason, you can use a response filter with the IGNORE response enabled.

Attributes of response filters

Response filters have the following configurable attributes:
- Interface
- Virtual LAN (VLAN)
- Source or target IP address
- Source or target port number (all ports or a port associated with a particular service) or ICMP type/code (one or the other will be used)

Filters and other events

When the appliance detects traffic that matches a response filter, the appliance executes the responses specified in the filter. Otherwise, the appliance executes the responses as specified in the event itself.

Note: If a security event is disabled, its corresponding response filters are disabled.

Response filter order

The response filters follow rule ordering. For example, if you add more than one filter for the same security event, the appliance executes the responses for the first match. The appliance reads the list of filters from top to bottom.

In the Policy

In IPS Local Management Interface:
- **Secure Protection Settings** > **Response Tuning** > **Response Filters**

In SiteProtector:
- **Response Filters** policy

Configuring a LEEF system log

To help integrate different components on your network, the appliance can send event information to a security incident event manager (SIEM) by using the log event extended format (LEEF).

About this task

When this feature is enabled, the appliance converts IPS, SNORT, health alert, and system alert events into LEEF for transmission to a SIEM. You can retrieve the log file from IPS Local Management Interface at **Review Analysis and Diagnostics** > **Downloads** > **Logs and Packet Captures**.

Note: IPS events include events from the security events, connection events, user defined events, and OpenSignatures policies.

This feature was tested with the QRadar SIEM developed by Q1 Labs. For more information go to http://q1labs.com. Customers of Q1 Labs can go to http://partners.q1labs.com and sign in to DocCentral to view the documentation.

Procedure

1. Enable the LEEF system log feature.
   a. In IPS Local Management Interface, go to **Secure Protection Settings** > **Advanced IPS** > **Tuning Parameters**. In SiteProtector Management, select the **Tuning Parameters** policy.
   b. Add the tuning parameter **crm.leef.enabled** and set it to **True**.
2. Optional: Set the size of the LEEF system log file.
   a. In IPS Local Management Interface, go to **Secure Protection Settings** > **Advanced IPS** > **Tuning Parameters**. In SiteProtector Management, select the **Tuning Parameters** policy.
   b. Add the tuning parameter **crm.leef.logsize** and set it to a number between **1-100** MB. The default is **10** MB.
3. Use the Secure Shell (SSH) protocol to communicate with the SIEM. Configure the SIEM to use the secure copy (SCP) command to get the LEEF system log from the appliance. The log file is located at `/var/iss/leef.log`. The QRadar SIEM gets the LEEF system log every 15 minutes.

Chapter 6. X-Force protection modules

The IBM X-Force research and development teams study and monitor the latest threat trends. They deliver security modules and content that work with your appliances to protect your network from threats.

PAM

PAM, the Protocol Analysis Module, provides the information that the appliance uses to protect the network against intrusions. PAM is a database that stores handling specifications for a comprehensive list of intrusions. IBM Security keeps PAM information current with X-Press Updates (XPUs), which you can apply through IPS Local Management Interface or SiteProtector's X-Press Update Server. To control PAM, use tuning parameter configurations.

Using X-Force default blocking

When you use X-Force Default Blocking, the block and quarantine responses are enabled automatically for events that X-Force recommends. The appliance enables or disables recommended settings depending on the options configured on the X-Force Virtual Patch page. Options include:

Table 4. X-Force options and actions

Option	Action when enabled
Always	When you apply X-Press Updates (XPUs), the appliance enables the block and quarantine responses to new events that are defined in the XPU.
Through XPU	When you apply XPUs, the appliance sets the block and quarantine responses to new events that are defined up to and including a specified XPU version. Use this option to control the application of XPU content updates. You can set this option to an XPU version you tested, so the appliance does not apply later XPU versions. Use this option so you can review X-Force recommendations first, so you can decide if you want them applied or not to new events.
Never	When you apply XPUs, the appliance does not set the block and quarantine responses to new events that are defined in the XPU.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Security Modules** > **X-Force Virtual Patch**

In SiteProtector:

- **Shared Objects** > **X-Force Virtual Patch** policy

© Copyright IBM Corp. 2003, 2012

Using data loss prevention signatures

Use the Data Loss Prevention feature to inspect and analyze packets for Personal Identifiable Information (PII) or other confidential information moving through and out of your network. You can use this feature with Predefined Events, User-Combined Events, and User Defined Events on your appliance.

How Data Loss Prevention works

Data Loss Prevention inspects data packets as they move across the network, detecting the transmission of many types of confidential information. The feature can identify patterns such as credit card numbers, names, dates, dollar amounts, e-mail addresses, social security numbers, United States phone numbers, and United States postal addresses in various protocols and content.

In addition to the preset signatures, you can create up to eight custom user-defined signatures. You can also create up to eight user-combined signatures by grouping combinations of preset and user-defined signatures. A user-combined signature functions as a single dataset.

Performance and tuning

With all Data Loss Prevention signatures and protocols turned on, you may notice some impact to network performance. Few enterprises need this level of protection, and your performance numbers are likely to improve as you identify the subset of signatures and protocols you need.

You can use Data Loss Prevention for either auditing or blocking. Most enterprises use audit mode while they are tuning policies. This approach helps security managers understand the kinds of data that they *could* be blocking without disrupting business operations. Other enterprises find that audit mode is sufficient, and they have no plans to deploy in blocking mode.

You may see a large number of events based on certain signatures and content types. You can reduce the number of events by making changes in your Data Loss Prevention policy.

Note: If you need assistance with your policies, our professional security consultants are available to help.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Security Modules** > **Data Loss Prevention**

In SiteProtector:

- **Data Loss Prevention** policy

Using web application protection

Web Application Protection (WAP) uses attacks, audits, and parameter names (keywords) from the IBM Security Protocol Analysis Module (PAM) engine to provide overall protection against web application security attacks.

WAP can help protect your network from the following types of web application security attacks.

Attack	Description
Injection Attack	Allows an attacker to inject code into a program or query, or to inject malware onto a computer in order to execute remote commands that can read or modify a database, or change data on a web site

Attack	Description
Malicious File Execution	Allows an attacker to execute code remotely, install a root kit remotely, compromise the entire system, and compromise the internal system on Windows systems) through the use of SMB file wrappers for the PHP scripting language
Cross-site Request Forgery (CSRF)	Sends unauthorized commands from a user that a web site trusts
Information Disclosure Attack	Attempts to acquire system specific information about a web site, including software distribution, version numbers, and patch levels. The acquired information might also contain the location of backup files or temporary files.
Path Traversal Attack	Forces access to files, directories, and commands that are located outside the web document root directory or CGI root directory
Authentication	Targets and attempts to exploit the authentication process that a web site uses to verify the identity of a user, service, or application
Buffer Overflow	Floods a target with excessive data in order to cause the buffer to overflow. Then an attacker can run remote shell on the computer and gain the same system privileges granted to the application being attacked.
Brute Force	Uses trial and error in order to programmatically guess a person's username, password, credit card number, or cryptographic key
Directory Indexing Attack	Exploits a function of the web server that lists all the files within a requested directory if the normal base file is not present
Miscellaneous Attack	Exploits vulnerable web servers by forcing cache server or web browsers into disclosing user-specific information that might be sensitive and confidential

PAM-Controlled Security Events and Response Filters

The Protocol Analysis Module (PAM) controls X-Force Virtual Patch recommendations, which means that PAM controls many security events. PAM overrides settings configured for some security events in the Web Application Protection (WAP) policy. If you want to override WAP policy settings for security events that PAM controls, use response filters. The response filter overrides the PAM settings so that the WAP policy responds to activity according to the needs of your network.

Important: You cannot change the WAP policy settings that PAM controls from the Web Application Protection page or from the Security Events page. You must use response filters.

Example: change block to ignore

PAM has configured the **HTTP_Unknown_Protocol** event parameter to use the block response, but you want this event to use the ignore response. You go to the Security Events page and look for the **HTTP_Unknown_Protocol** parameter to change it, but it is not there. Go to the Response Filter page and create a response filter for the event name. Then select the **Ignore Events** check box. The response filter setting overrides the PAM setting, and the **HTTP_Unknown_Protocol** event parameter now uses the ignore response.

Example: change enabled to disabled

PAM has enabled the **HTTP_Get_CreateTable** parameter, but this does not meet the needs of your network so you want to disable it. You go to the Security Events page and look for the **HTTP_Get_CreateTable** parameter to reconfigure it, but it is not there. Go to the Response Filter page and create a response filter for the event name. Then clear the **Enabled** check box. The response filter setting overrides the PAM setting, and the **HTTP_Get_CreateTable** parameter is now disabled.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Security Modules** > **Web Application Protection**

In SiteProtector:

- **Web Application Protection** policy

Chapter 7. Using protection domains

With custom protection domains, you can use one appliance to monitor multiple network segments, even if those segments require different security settings. Protection domains function like virtual sensors, as though you had several appliances monitoring the network. You can use custom protection domains to define different security settings for different network segments.

Global protection domain

Each appliance has a global protection domain that cannot be deleted. All events are listed under the global protection domain. Use the global policy to configure events to be applied across all segments of the network. When the appliance uses the global policy, it handles events in the same way for all areas of your network.

If you want to configure polices for specific segments on your network, create protection domains for each segment.

Note: Always enable rules for flood and sweep events in the global protection domain. Flood and sweep attacks generally affect multiple targets which are potentially spread across protection domains. Enable these rules in the global protection domain to help ensure these attacks are detected and reported correctly.

Additional protection domains

Create custom protection domains when you want to use a single appliance to monitor multiple network segments with varying security requirements. Use these protection domains to apply different security policies to different network segments.

You can define protection domains using ports, VLANs, or IP address ranges.

In the Policy

In IPS Local Management Interface:

- **Secure Protection Settings** > **Advanced IPS** > **Protection Domains**

In SiteProtector:

- **Shared Objects** > **Protection Domains**

Working with protection domains

Use protection domains to define security policies and user-defined policies for different network segments monitored by a single appliance.

Policies that use protection domains

You can use the global protection domain or custom protection domains with the following policies:

- Security Events
- User Defined Events
- Data Loss Prevention
- Web Application Protection
- Response Filters

© Copyright IBM Corp. 2003, 2012

Protection domains and events

By default, the appliance uses the global protection domain to manage security. If your network requires different security settings for different segments, define custom protection domains and assign security settings as appropriate for each domain.

Notes:

- Do not use the same name for events that have different contexts and query strings. If you do use the same name, it will be difficult to determine which event occurred.
- If you have two events with the same name, one assigned to the global protection domain and one assigned to a custom protection domain, only the event assigned to the custom domain generates an alert if the alert details occur within the defined network segment. Otherwise, the appliance reports the event assigned to the global protection domain.
- If you have two user-defined events that are the same but have different names, each event generates its own alert.

Protection domains and IPv6 support

You can specify an IPv6 address or range of addresses to define protection domains. Protection domains are fully supported in an IPv6 environment, but may be more challenging to use. For example, a portable asset such as a laptop computer can have multiple IPv6 addresses, depending on where it is connected in the network.

Best practices for protection domains

Protection domains can be a valuable tool extending your network protection if you understand them and use them correctly.

Use the global protection domain when possible

If you want to apply security settings to all network segments protected by a single appliance, use the global protection domain. This approach is faster and easier than setting the same policies in multiple protection domains.

Protecting against flood and sweep attacks

Certain flood and sweep attacks might not be recognized by custom protection domains. These attacks generally affect multiple targets, which are potentially spread across protection domains. Enable these events for the global protection domain to help ensure that these attacks are detected and reported correctly.

Deleting protection domains

If you delete a protection domain, user-defined events, security events, and response filters assigned to that protection domain may remain active, and events associated with the deleted protection domain may still fire. Before you delete a protection domain, delete or reassign all user-defined events, security events, and response filters associated with the protection domain.

Chapter 8. Configuring appliances for high availability

High availability (HA) support is a configuration arrangement between two cooperating appliances. HA mode enables two comparable appliances to work together in an existing high availability environment to provide added protection for your network. Two appliances that are connected and configured to operate in HA mode are called HA partners or an HA pair.

HA and SiteProtector management

You can view HA configurations in IPS Local Management Interface, but you should use SiteProtector to manage appliances in inline HA configurations. Both appliances in an HA pair must be in the same SiteProtector group. SiteProtector can then synchronize appliance updates, including XPUs and policy updates.

You can apply content updates and firmware updates serially so that one appliance is always operational in order to maintain network connectivity, particularly when both appliances are configured to fail closed.

Each appliance reports to SiteProtector using a unique ID.

Licensing

Licensing for an HA configuration is identical to licensing for a non-HA appliance. Each individual appliance must have its own license. If you are using Site Protector to manage HA appliances, each appliance requests a license from Site Protector.

Limitations

In HA mode, you cannot use interface parameters as part of the firewall rules. You cannot define protection domains based on interfaces. Because the same traffic may flow on different interfaces in an HA environment, using interface parameters may cause HA partner appliances to become unsynchronized.

Important: You must select all interfaces when you define protection domains and constructed firewall rules. Do not use the interface keyword when you create firewall rule definitions.

HA considerations

- You cannot mix models in a single HA environment. For example, you cannot use a GX5208 appliance and a GX6116 appliance as an HA pair.
- Ensure the firmware level and X-Press Update (XPU) level on appliances in an HA pair match.
- Manage appliances in an HA pair in the same SiteProtector group.

In the Policy

In IPS Local Management Interface:

- **Manage System Settings** > **Network** > **Security Interfaces**

HA configuration options

IBM Security Network IPS offers the following approaches to high availability (HA) configuration: standard HA and geographical HA.

© Copyright IBM Corp. 2003, 2012

In a **standard HA** configuration, the protection ports for two appliances are cabled so that each appliance mirrors traffic from the other appliance. Half of the available ports on each appliance are used as "inline ports" and half of the ports are "mirror ports" to the other appliance. While this configuration helps maximize network availability and protection, it has some limitations. The appliances that make up the HA pair must be located within cabling distance of each other, and half of the protection ports for each appliance are given up to serve as mirror ports.

In a **geographical HA** configuration, two appliances share their quarantine states, but do not mirror traffic. Quarantine rules that are created on one appliance in the pair are forwarded to the other appliance. Appliances that make up an HA pair communicate through their management ports and use the management network to communicate. Proximity for cabling is not an issue.

High availability modes

In an HA configuration, an appliance can operate in only inline simulation or inline protection mode. Passive monitoring mode is not supported. When you select an HA mode, all inline interfaces are put in the corresponding interface mode automatically.

HA does not address the availability or fault-tolerance of the appliances themselves. No separate high availability solution exists for appliances configured and wired for passive monitoring mode. You can configure appliances using the following high availability modes, as indicated in the following table.

Setting	Description
Normal mode (HA off)	HA is disabled, and each appliance operates on its own. Appliances can be configured to run in inline protection, inline simulation, and passive monitoring modes at the interface level only.
HA Simulation mode (standard HA)	Both HA partner appliances monitor traffic inline, but do not block any traffic. Instead, both appliances monitor traffic and provide passive notification responses. The appliances monitor traffic on each other's segment via mirror links, ready to take over notification in case of network failover.
HA Protection mode (standard HA)	Both HA partner appliances monitor traffic inline, and each report and block the attacks configured with block response, quarantine response, and firewall rules. The appliances monitor traffic on each other's segments using mirror links, ready to take over reporting and protection in case of network failover.
Geographical HA	Each appliance in the HA pair monitors its own traffic, and passes new quarantine rules to its partner.

Deployment for standard high availability

The IBM Security Network IPS High Availability (HA) feature enables appliances to work in an *existing* high availability network environment. The appliances pass all traffic between them over mirroring links, ensuring that both appliances see all of the traffic over the network and thus maintain state. This approach allows the appliances to see asymmetrically routed traffic in order to fully protect the network.

HA support is limited to two cooperating appliances. Both appliances process packets inline, block attack traffic that arrives on their inline protection ports, and report events received on their inline ports to the management console.

Supported appliances

You can use the following appliance models in an existing HA environment:

- GX5000 series appliances
- GX6000 series appliances
- GX7000 series appliances

Important: You cannot *mix* models in a single HA environment. For example, you cannot use a GX5208 appliance and a GX6116 appliance as an HA pair.

Supported network configurations

High availability networks are typically configured in one of two ways:

Existing HA configuration	Description
Primary / Secondary	With this configuration, the traffic flows only on one of the redundant network segments and the primary devices on the network handle all of the traffic until one of the devices fails, at which point the traffic fails over to the secondary redundant network segment and the secondary devices take over.
Clustering	With this configuration, the traffic is load balanced and both sets of devices are active and see traffic all of the time.

The HA feature in IBM Security Network IPS supports both of these network configurations. In order to accomplish this, both IBM Security appliances must maintain identical states. The appliances are connected by mirror links that consist of multiple connections over multiple ports. These mirror links pass all traffic an appliance receives on its inline ports to the other appliance, ensuring the protocol analysis modules on both appliances process all of the network traffic. In addition, the appliances process asymmetrically routed traffic. This approach ensures that there is no gap in protection during failover.

Note: If you run IPS Setup when the HA feature is enabled, you cannot modify network settings.

HA processing, blocking, reporting, and generating responses

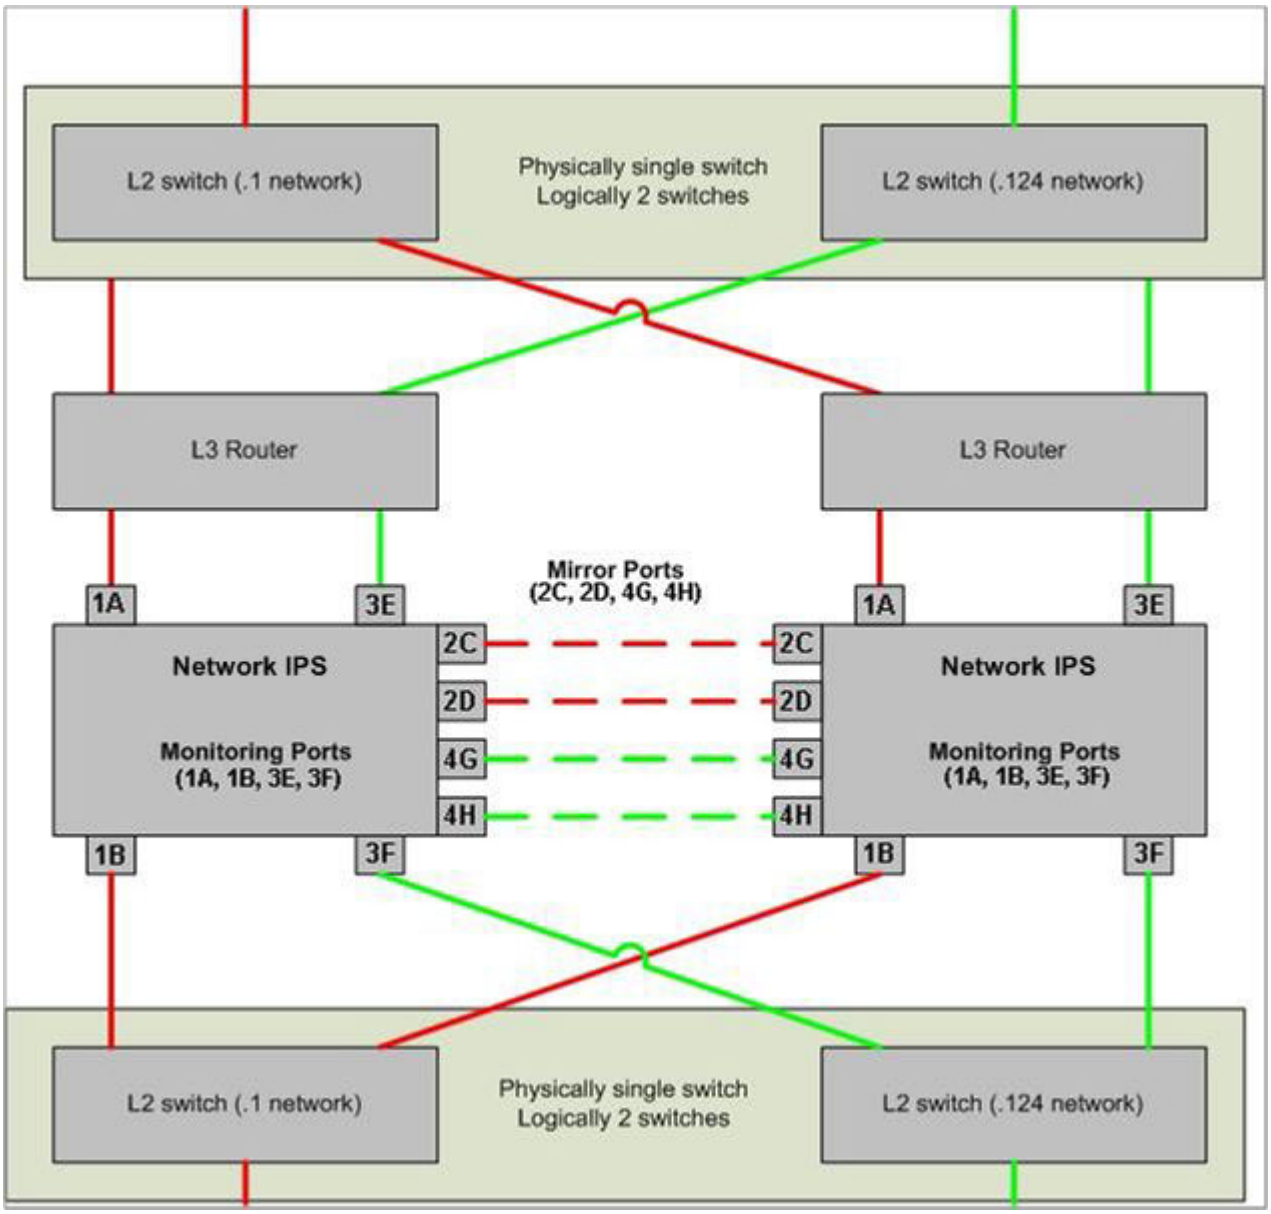
Appliances in an HA pair process all packets received from inline ports and mirror ports. However, the appliances only block attacks, report events, and generate responses for events occurring on their inline ports. They do not block, report, or generate responses for traffic occurring on mirror ports. The appliances only process mirror port traffic.

Both appliances see all traffic at all times. There is no lapse in security if a failover occurs. Both appliances maintain current state, so if one HA network segment fails, the other appliance receives all packets on its inline ports. The network remains protected without interruption.

Note: A small number of attacks, particularly sweep attacks such as Port Scans, can generate duplicate events, one from each appliance in a clustered configuration.

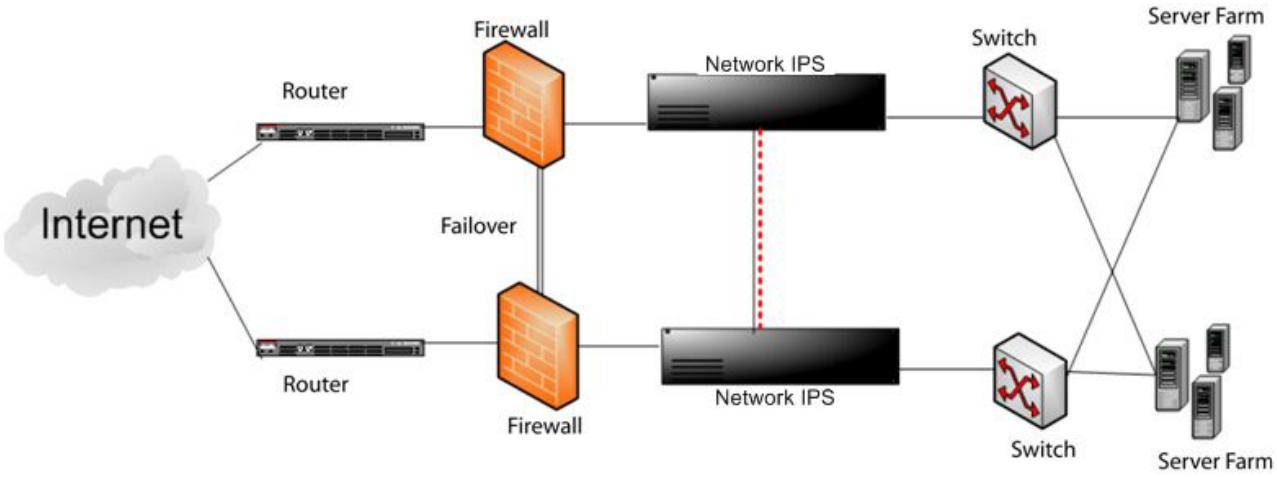
Standard HA Deployment: Logical Diagram

If you use SiteProtector to manage the appliances, you can manage the HA cluster from the SiteProtector Agent Manager. A Logical HA diagram is shown in the following figure.


L2 switch (.1 network)
Physically single switch
Logically 2 switches
L2 switch (.124 network)
L3 Router
L3 Router
Mirror Ports
(2C, 2D, 4G, 4H)
1A
3E
1A
3E
Network IPS
Monitoring Ports
(1A, 1B, 3E, 3F)
2C
2D
4G
4H
2C
2D
4G
4H
Network IPS
Monitoring Ports
(1A, 1B, 3E, 3F)
1B
3F
1B
3F
L2 switch (.1 network)
Physically single switch
Logically 2 switches
L2 switch (.124 network)

Standard HA Deployment: Physical Diagram

A physical network diagram of a typical HA deployment scenario is shown in the following figure.



Deployment for geographical high availability

In a geographical HA configuration, two appliances share their quarantine states, but do not mirror traffic. Appliances that make up an HA pair communicate through their management ports. Quarantine rules that are created on one appliance in the pair are forwarded to the other appliance over the management port. The HA pair uses the management network to communicate, and therefore proximity for cabling is not an issue.

Supported appliances

You can use the following appliance models in an existing HA environment:

- GX3000 series appliances
- GX4000 series appliances
- GX5000 series appliances
- GX6000 series appliances
- GX7000 series appliance
- GV series virtual appliances

Important: You cannot *mix* models in a single HA environment. For example, you cannot use a GX5208 appliance and a GX6116 appliance as an HA pair.

Communication between HA partners

In a geographical HA configuration, the HA partners communicate with each other over the management network. All communication between the partner appliances is encrypted. You need certificates to enable communications.

The following communication options are available:

- **First-time trust:** When each appliance comes online, it requests the necessary encryption certificates from its partner appliance.
- **Explicit trust:** You must manually copy encryption keys to both appliances to enable communications.

Changing host name or time/date settings in a geographical HA pair

First-time-trust: Before you change the host name or time/date settings, disable geographical HA mode on the partner appliance. This step triggers the appliances in a first-time-trust configuration to automatically download new encryption keys when you re-enable geographical HA.

Explicit-trust: If the HA pair is set to use explicit trust, then you must copy the keys from the changed appliance to its HA partner to enable communication.

Re-imaging an appliance in a geographical HA pair

First-time-trust: Before you re-image an appliance or reset it to the factory default configuration, disable geographical HA mode on the partner appliance. This step triggers the appliances in a first-time trust configuration to automatically download new encryption keys when you re-enable geographical HA.

Explicit-trust: If the HA pair is set to use explicit-trust, then you must copy the keys from the re-imaged appliance to its HA partner to enable communication.

System time

Verify that the system times on both appliances is correct before you enable geographical HA. Otherwise, the encryption keys may not be created correctly.

Chapter 9. General information

This section contains general information about IBM Security Network IPS appliances.

Compatibility

The following topic lists the web browsers and Java Runtime Environment (JRE) versions currently supported by the appliance.

Web browser compatibility

The following browser's are supported:

- Internet Explorer 8
- Internet Explore 9
- Firefox 13

Java Runtime Environment compatibility

JRE 1.6 and 1.7 are supported. Do one of the following actions when using JRE.

Important: JRE 1.7 works for only 32-bit Windows systems. It does not work with 64-bit Windows systems

- Clear the Java cache, often.
- Disable the Java console from keeping temporary files on the computer.
- Set the Java cache maximum space to zero.

Appliance partitions

Partitions and file systems are listed in the following table.

Partition	File system
/ *Root partition*	• Operating system • Intrusion prevention system modules • Databases
`/boot`	• Operating system
`/rboot`	• Operating system
`/cache`	• Log files
`/restore`	• Backup images • Factory default images

Cumulative updates and rollbacks

After you install an update, the appliance deletes the update package and the downloaded package is no longer on your appliance. If you roll back the update, then the appliance finds the update available for download and installation the next time you find updates or at the next scheduled automatic update.

© Copyright IBM Corp. 2003, 2012

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corp. 2003, 2012

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Project Management
C55A/74KB
6303 Barfield Rd.,
Atlanta, GA 30328
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Index

A

adapter clause 17
adapter modes
  inline simulation 2
  passive monitoring 2
Agent Manager 8
appliance
  interface modes 2
  protection features 1
  SiteProtector 8
appliance partitions 51
attacks
  Flood 44
  Sweep 44
autokey 12
autokey authentication 12

B

Block response 3

C

capacity planning
  driver statistics 11
  MIB file 11
  SNMP GET request 11
  throughput graphs 11
certificate-based key exchange 12
connection events 23
customer support web site xvii
customer support, IBM Security Solutions xvii

D

data loss prevention 40
  considerations 40
  signatures 40

E

email responses 3
Email_Receiver context 25
Email_Sender context 25
ethernet clause 17
events
  connection 23
  SiteProtector 9
  user-defined 24

F

filters
  response 37
  security events 21
FIPS 140-2 12
firewall clauses 17
firewall clauses *(continued)*
  adapter clause 17
  ethernet clause 17
  IP datagram clause 17
firewall conditions
  ICMP conditions 18
  TCP and UDP conditions 18
firewall expressions 19
firewall rules 15
  actions 16
  criteria 15
  examples 20
  expressions 19
  firewall clauses 17
  firewall conditions 18
  language 17
  rule order 15, 16
Flood attacks 44

H

health alerts
  error 10
  informative 10
  warning 10
high availability
  blocking 47
  considerations 45
  processing 47
  reporting 47
  responses 47
high availability (HA)
  clustering 47
  licensing 45
  limitations 45
  modes 46
  primary/secondary configurations 47
  SiteProtector management 45

I

IBM Security Solutions
  customer support xvii
ICMP conditions 18
ICMP port unreachable 32
Ignore response 3
inline protection mode 2
inline simulation mode 2
interface modes
  inline protection 2
Internet Scanner
  SNMP_Community context 26
intrusion prevention 21
  connection events 23
  OpenSignatures 29
  quarantined intrusions 23
  responses 3
  security events 21
  user-defined events 24
intrusion prevention *(continued)*
  X-Force default blocking 39
IP datagram clause 17
IPS Local Management Interface
  compatibility 51
  IPS Local Management Interface
    compatibility 51
    supported Java 51
  supported browsers 51
IPv6 5

J

Java
  actions 51
  JRE 51
Java compatibility 51

K

key IDs 12

L

LEEF (log event extended format) 38
LEEF syslog 38
LEEF system log 38
licensing
  high availability (HA) 45
log event extended format (LEEF) 38
log evidence responses 4

M

MIB file 11
modes
  high availability (HA) 46
  inline protection 2
  inline simulation 2
  passive monitoring 2

N

negation operator 30
Network Time Protocol (NTP) 12
News_Group context 26
NTP 12
NTP configuration 12
NTP policy 12
NTP servers 12
NTP version 4 12

O

OpenSignatures 29
  default responses 31
  parser 30
  risks 30

© Copyright IBM Corp. 2003, 2012

OpenSignatures *(continued)*
syntax 30

P

PAM,
protocol analysis module 39
partitions
file systems 51
passive monitoring mode 2
Password context 26
policies
security 21
predefined quarantine responses 4
DDOS (distributed denial-of-service 5
intruder 4
trojan 4
worm 5

Q

quarantine responses 4
DDOS (distributed denial-of-service 5
intruder 4
trojan 4
worm 5
quarantine rules
single-click blocking 23
quarantined intrusions 23

R

regular expressions 27
library 27
limitations 27
precedence 27
syntax 28
response filters 37
event attributes 37
order 37
responses 3
Block 3
email 3
executables 4
Ignore 3
log evidence 4
quarantine 4
response objects 5
shell scripts 4
SNMP 4
user specified 4

S

safety notices vii
security events 21
filters 21
security incident event manger (SIEM) 38
security policy documents
where to find 12
sensor alerts
error 10
sensor alerts *(continued)*
informative 10
warning 10
SIEM (security incident event manager 38
single-click blocking 23
SiteProtector
Agent Manager 8
appliance events 9
appliance management 8
high availability (HA) support 45
management options 8
response objects 5
updates 9
SNMP
responses 4
SNMP responses 4
SNMP_Community context 26
Internet Scanner 26
SNORT 31
considerations 32
errors 33
HA mode 33
health status 34
high availability (HA), disable 33
high availability (HA), enable 33
high availability mode 33
ICMP port unreachable 32
PAM 33
Protocol Analysis Module 33
quarantine responses 34
quarantine rules 32, 33, 34
rule profiling 31, 35
SiteProtector alerts 33
TCP reset port 32
TCP resets 31, 32
troubleshooting 34
unsupported configuration options 35
SNORT configuration
unsupported options 35
SNORT configuration file
default 32
imported file 32
SNORT errors 33, 34
SNORT event process (SnEP)
SNORT errors 34
SNORT event processor (SnEP) 34
SNORT rule
capacity 32
SNORT rules
delete 35
import 34
maximum number 32
SNORT rules rile
maximum file size 32
Sweep attacks 44
symmetric key authentication 12
symmetric key IDs 12

T

TCP and UDP conditions 18
tuning parameters
default values 29
PAM 39
protocol analysis module 39

U

updates
SiteProtector 9
URL_Data context 27
user defined events
global compared to custom protection domains 44
user specified responses 4
shell scripts 4
User_Login_Name context 27
User_Probe_Name context 27
user-defined event contexts 24
Email_Receiver context 25
Email_Sender context 25
File_Name context 26
News_Group context 26
Password context 26
SNMP_Community context 26
URL_Data context 27
User_Login_Name context 27
User_Probe_Name context 27
user-defined events 24
event contexts 24
regular expressions 27

W

web application protection
authentication attack 41
brute force attack 41
buffer overflow 41
cross-site request forgery (CSRF) 41
directory indexing attack 41
information disclosure attack 41
injection attack 40
malicious file execution 41
miscellaneous attacks 41
path traversal attack 41
response filters 41
WAP 40
web browser compatability 51

X

X-Force default blocking 39
options 39



Printed in USA