



Network and Security Manager

Configuring Intrusion Detection and Prevention Devices Guide

Release
2012.2



Published: 2013-01-03
Revision 01

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Network and Security Manager Configuring Intrusion Detection and Prevention Devices Guide

Release 2012.2

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

Revision History

January 2013 —01

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About This Guide	vii
	Objectives	vii
	Audience	vii
	Conventions	vii
	List of Technical Publications	ix
	Requesting Technical Support	x
	Self-Help Online Tools and Resources	x
	Opening a Case with JTAC	xi
Part 1	Getting Started	
Chapter 1	Intrusion Detection and Prevention Device and NSM Installation Overview	3
	Intrusion Detection and Prevention Device Installation Overview	3
	NSM Installation Overview	3
Chapter 2	Understanding Intrusion Detection and Prevention Device Configuration and Integration Overview	5
	NSM and Intrusion Detection and Prevention Device Management Overview	5
	Intrusion Detection and Prevention Services and Device Configurations Supported in NSM	6
	Adding Intrusion Detection and Prevention Devices in NSM Overview	8
	Adding Intrusion Detection and Prevention Clusters in NSM Overview	8
	Using Templates and Configuration Groups in NSM Overview	8
Part 2	Configuring Intrusion Detection and Prevention Devices	
Chapter 3	Configuring Profiler Settings	13
	Configuring Profiler Options (NSM Procedure)	13
	Specifying General Options	14
	Specifying Tracked Hosts	16
	Specifying Context Targets	17
	Specifying Alert Options	18
	Viewing Profiler Logs (NSM Procedure)	20
	Application Profiler	20
	Protocol Profiler	22
	Network Profiler	23

	Violation Viewer	25
	Modifying Profiler Settings (NSM Procedure)	25
	Configuring Profiler Database Preferences (NSM Procedure)	26
	Displaying Profiler Database Information (NSM Procedure)	28
	Querying the Profiler Database (NSM Procedure)	28
	Purging the Profiler Database (NSM Procedure)	28
Chapter 4	Configuring Security Policies	31
	Intrusion Detection and Prevention Devices and Security Policies Overview	31
	Configuring Predefined Security Policies (NSM Procedure)	33
	Creating a New Security Policy (NSM Procedure)	34
	Modifying IDP Rulebase Rules (NSM Procedure)	36
	Specifying Rule Match Conditions	37
	Specifying IDP Rulebase Attack Objects	38
	Specifying Rule Session Action	39
	Specifying Rule IP Action	41
	Specifying Rule Notification Options	42
	Specifying Rule VLAN Matches	42
	Specifying Rule Targets	43
	Specifying Rule Severity	43
	Specifying Rule Optional Fields	44
	Specifying Rule Comments	44
	Configuring Exempt Rulebase Rules (NSM Procedure)	45
	Configuring Backdoor Rulebase Rules (NSM Procedure)	47
	Configuring SYN Protector Rulebase Rules (NSM Procedure)	49
	Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)	51
	Configuring Network Honeypot Rulebase Rules (NSM Procedure)	54
	Configuring Application Rulebase Rules (NSM Procedure)	57
Chapter 5	Working with Attack Objects	63
	Attack Objects in Intrusion Detection and Prevention Security Policies	
	Overview	63
	Loading J-Security-Center Updates (NSM Procedure)	64
	Viewing Predefined Attack Objects (NSM Procedure)	65
	Working with Attack Groups (NSM Procedure)	66
	Creating Dynamic Groups	67
	Creating Static Groups	68
	Creating a Signature Attack Object (NSM Procedure)	69
	Creating a Compound Attack Object (NSM Procedure)	81
	Verifying the Attack Object Database Version (NSM Procedure)	84
	Updating the IDP Detector Engine (NSM Procedure)	85
Chapter 6	Working with Application Objects	87
	Application Objects in Intrusion Detection and Prevention Security Policies	
	Overview	87
	Viewing Predefined Application Objects (NSM Procedure)	88
	Viewing Predefined Extended Application Objects (NSM Procedure)	88
	Creating a Custom Application (NSM Procedure)	89
	Creating Application Groups (NSM Procedure)	90

Chapter 7	Configuring SNMP and Syslog Settings	91
	Configuring an SNMP Agent (NSM Procedure)	91
	Configuring Syslog Collection (NSM Procedure)	92
Chapter 8	Configuring Anti-Spoof Settings	95
	Configuring Antispoof Settings in Intrusion Detection and Prevention Devices (NSM Procedure)	95
	Example: Applying Antispoof to a Web Server and Database Server (NSM Procedure)	96
Chapter 9	Configuring Intrusion Detection and Prevention Device Settings	99
	Configuring Load-Time Parameters (NSM Procedure)	99
	Configuring Run-Time Parameters (NSM Procedure)	101
	Configuring Router Parameters (NSM Procedure)	106
	Configuring Protocol Handling (NSM Procedure)	107
Chapter 10	Configuring Additional Intrusion Detection and Prevention Features	123
	Configuring Additional Intrusion Detection and Prevention Features Overview	123
	Enabling Intrusion Detection and Prevention Processing of Encrypted and Encapsulated Traffic (NSM Procedure)	123
	Enabling SSL Decryption	124
	Enabling GRE Decapsulation	124
	Enabling GTP Decapsulation	125
Part 3	Managing Intrusion Detection and Prevention Devices	
Chapter 11	Managing Security Policies in Intrusion Detection and Prevention Devices	129
	Assigning a Security Policy in an Intrusion Detection and Prevention Device (NSM Procedure)	129
	Validating a Security Policy (NSM Procedure)	130
	Troubleshooting Security Policy Validation Errors (NSM Procedure)	130
	Pushing Security Policy Updates to an IDP Device (NSM Procedure)	131
	Troubleshooting Configuration Push Errors (NSM Procedure)	133
	Disabling Rules (NSM Procedure)	134
	Exporting Security Policies (NSM Procedure)	134
Chapter 12	Managing Profiler Settings in Intrusion Detection and Prevention Devices	135
	Managing Profiler Settings	135
	Updating Profiler Settings	135
	Starting the Profiler	135
	Stopping the Profiler	135

Part 4	Monitoring Intrusion Detection and Prevention Devices	
Chapter 13	Working with NSM Logs and Reports	139
	NSM Logs and Reports Overview	139
	Viewing Logs	139
	IDP Logs	140
	Using NSM Log Investigator	140
	Using NSM Audit Log Viewer	140
	Viewing Device Status	142
	Viewing NSM Predefined Reports	145
	Creating NSM Custom Reports	147
	Configuring Log Suppression	149
Chapter 14	Working with Intrusion Detection and Prevention Reporter Reports	151
	Intrusion Detection and Prevention Reporter Overview	151
Part 5	Index	
	Index	155

About This Guide

- [Objectives on page vii](#)
- [Audience on page vii](#)
- [Conventions on page vii](#)
- [List of Technical Publications on page ix](#)
- [Requesting Technical Support on page x](#)

Objectives

Network and Security Manager (NSM) is a software application that centralizes control and management of your Juniper Networks devices. With NSM, Juniper Networks delivers integrated, policy-based security and network management for all security devices.

Intrusion Detection and Prevention (IDP) series uses eight detection methods to detect malicious network traffic. It is able to drop attacks to prevent damage to your network and can operate inline as a forwarding gateway, directly in the path of traffic coming and going on your network.

This guide provides the various steps to configure and manage IDP devices using NSM. This guide also helps in understanding of how to configure basic and advanced NSM functionality, including adding new devices, deploying new device configurations, updating device firmware, viewing log information, and monitoring the status of IDP devices.

Audience

This guide is intended for the system administrators who are responsible for configuring IDP devices.

Conventions

This section provides all the documentation conventions that are followed in this guide. [Table 1 on page viii](#) defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 2 on page viii defines text conventions used in this guide.

Table 2: Text Conventions

Convention	Description	Examples
Bold typeface like this	<ul style="list-style-type: none"> Represents commands and keywords in text. Represents keywords Represents UI elements 	<ul style="list-style-type: none"> Issue the clock source command. Specify the keyword exp-msg. Click User Objects
Bold typeface like this	Represents text that the user must type.	user input
<code>fixed-width font</code>	Represents information as displayed on the terminal screen.	<pre>host1# show ip ospf Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an area Border Router (ABR)</pre>
Key names linked with a plus (+) sign	Indicates that you must press two or more keys simultaneously.	Ctrl + d
<i>Italics</i>	<ul style="list-style-type: none"> Emphasizes words Identifies variables 	<ul style="list-style-type: none"> The product supports two levels of access, <i>user</i> and <i>privileged</i>. <i>clusterID</i>, <i>ipAddress</i>.
The angle bracket (>)	Indicates navigation paths through the UI by clicking menu options and links.	Object Manager > User Objects > Local Objects

Table 3 on page ix defines syntax conventions used in this guide.

Table 3: Syntax Conventions

Convention	Description	Examples
Words in plain text	Represent keywords	terminal length
Words in italics	Represent variables	<i>mask, accessListName</i>
Words separated by the pipe () symbol	Represent a choice to select one keyword or variable to the left or right of this symbol. The keyword or variable can be optional or required.	diagnostic line
Words enclosed in brackets ([])	Represent optional keywords or variables.	[internal external]
Words enclosed in brackets followed by and asterisk ([]*)	Represent optional keywords or variables that can be entered more than once.	[level1 level2 11]*
Words enclosed in braces ({ })	Represent required keywords or variables.	{ permit deny } { in out } { clusterId ipAddress }

List of Technical Publications

This section provides the list of the documentations required for any additional information.

Table 4: Network and Security Manager and IDP Device Publications

Network and Security Manager Installation Guide	Details the steps to install the NSM management system on a single server or on separate servers. It also includes information on how to install and run the NSM user interface. This guide is intended for IT administrators responsible for the installation and/or upgrade to NSM.
Network and Security Manager Administration Guide	<p>Describes how to use and configure key management features in the NSM. It provides conceptual information, suggested workflows, and examples where applicable. This guide is best used in conjunction with the NSM Online Help, which provides step-by-step instructions for performing management tasks in the NSM UI.</p> <p>This guide is intended for application administrators or those individuals responsible for owning the server and security infrastructure and configuring the product for multi-user systems. It is also intended for device configuration administrators, firewall and VPN administrators, and network security operation center administrators.</p>
Network and Security Manager Configuring Firewall/VPN Devices Guide	Describes NSM features that relate to device configuration and management. It also explains how to configure basic and advanced NSM functionality, including deploying new device configurations, managing Security Policies and VPNs, and general device administration.
Network and Security Manager Online Help	Provides task-oriented procedures describing how to perform basic tasks in the NSM user interface. It also includes a brief overview of the NSM system and a description of the GUI elements.
IDP Installation Guide	Details the physical features of Juniper Networks Intrusion Detection and Prevention (IDP) series. It also explains how to install, configure, update/reimage, and service the IDP system.

Table 4: Network and Security Manager and IDP Device Publications (*continued*)

IDP Concepts & Examples Guide	Details about the Juniper Networks Intrusion Detection and Prevention (IDP) series that uses multiple methods to detect and prevent network attacks. IDP is designed to reduce false positives to ensure that only actual malicious traffic is detected and stopped.
IDP Reporter User's Guide	Details about the IDP Reporter that enables you to analyze your enterprise network thoroughly so you can assess attacks, attackers, and resource utilization.
IDP ACM Online Help	Details about how to complete the IDP QuickStart and ACM Wizard which is available through the IDP Appliance Configuration Manager (ACM) as context-sensitive online help.
IDP Detector Engine Release Notes	Details about IDP Detector Engine features and resolved issues in the recent releases. It also helps you to decide to update the IDP Detector Engine version in your deployment.

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://www.juniper.net/alerts/>

- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Getting Started

- [Intrusion Detection and Prevention Device and NSM Installation Overview on page 3](#)
- [Understanding Intrusion Detection and Prevention Device Configuration and Integration Overview on page 5](#)

CHAPTER 1

Intrusion Detection and Prevention Device and NSM Installation Overview

- [Intrusion Detection and Prevention Device Installation Overview on page 3](#)
- [NSM Installation Overview on page 3](#)

Intrusion Detection and Prevention Device Installation Overview

The Intrusion Detection and Prevention (IDP) series consists of hardware and software components. You can install the IDP device and start configuring your system using the following steps:

1. Decide on the physical location of the device.
2. Install the device into your equipment rack.
3. Connect power cables and power on.
4. Perform some initial configuration steps.
5. Install the device license key.

See the installation documentation for your IDP model to install, configure, update, and service a Juniper Networks IDP device.

Related Documentation

- [NSM Installation Overview on page 3](#)
- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)

NSM Installation Overview

NSM is a software application that enables you to integrate and centralize management of your Juniper Networks environment. You need to install two main software components to run NSM: the NSM management system and the NSM user interface (UI).

See the *Network Security Manager Installation Guide* for the steps to install the NSM management system on a single server or on separate servers. It also includes information on how to install and run the NSM user interface. The *Network Security Manager Installation Guide* is intended for IT administrators responsible for installing or upgrading to the NSM.

- Related Documentation**
- [Intrusion Detection and Prevention Device Installation Overview on page 3](#)
 - [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)

CHAPTER 2

Understanding Intrusion Detection and Prevention Device Configuration and Integration Overview

- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)
- [Intrusion Detection and Prevention Services and Device Configurations Supported in NSM on page 6](#)
- [Adding Intrusion Detection and Prevention Devices in NSM Overview on page 8](#)
- [Adding Intrusion Detection and Prevention Clusters in NSM Overview on page 8](#)
- [Using Templates and Configuration Groups in NSM Overview on page 8](#)

NSM and Intrusion Detection and Prevention Device Management Overview

NSM is the Juniper Networks network management tool that allows distributed administration of network appliances. You can use the NSM application to centralize status monitoring, logging, and reporting, and to administer IDP Series configurations.

IDP technology detects and stops attacks when deployed inline to your network. Unlike intrusion detection service (IDS), IDP uses multiple methods to detect attacks against your network and to prevent attackers from gaining access and damaging your system. IDP drops malicious packets or connections before the attacks enter your network. IDP is designed to reduce false positives and ensure that only actual malicious traffic is detected and stopped. You can also deploy IDP as a passive sniffer, similar to a traditional IDS, but with greater accuracy and manageability.

NSM is the sole means for configuring and managing IDP on the ISG1000, ISG2000, and standalone IDP Sensors running IDP 4.x. Standalone IDP sensors running IDP 3.x and earlier are managed using the IDP management server and UI.

The ISG1000 and ISG2000 security modules have an optional component installed that provides IDP functionality. If you have purchased an ISG1000 or ISG2000 device that does not have IDP capability, you can upgrade the device to be an IDP-capable system by replacing the memory chip in the CPU. You install up to three security modules and install the Advanced and IDP license keys for IDP.

With NSM, you can manage most of the parameters that you can configure through the IDP admin console. The configuration screens rendered through NSM are similar to the screens in the IDP admin console. NSM incorporates a broad configuration management framework that allows co-management using other methods.

After you have completed installation, follow these steps to get started with managing an IDP device with NSM:

1. Add the IDP device to NSM. When you first add the IDP device to NSM in first instance, NSM pushes the policy named Recommended to the device.
2. Update the IDP detector engine and attack object database.
3. Update software version (if necessary).
4. Run the Profiler.
5. Examine the logs.
6. Create address objects for IDP rulebase rules.
7. Optionally, configure additional rulebases.
8. If adding this device changes your plan to distribute administrative responsibility, create NSM users with the access privileges.

An administrator (a user of NSM or IDP) has a specific level of permission. You can create multiple administrators with specific roles to control access to the devices in each domain.

**Related
Documentation**

- [Intrusion Detection and Prevention Services and Device Configurations Supported in NSM on page 6](#)
- [Adding Intrusion Detection and Prevention Devices in NSM Overview on page 8](#)

Intrusion Detection and Prevention Services and Device Configurations Supported in NSM

The Intrusion Detection and Prevention (IDP) device supports the following services in NSM:

- Inventory management service—NSM enables upgrading license and management of the IDP hardware details. Adding or deleting licenses or upgrading or downgrading software are not supported.
- Status monitoring service—Allows the IDP device's status to be obtained, including name, domain, OS version, synchronization status, connection details, current alarms, CPU, memory, and swap.
- Logging service—Allows the IDP device's logs to be obtained in a time-generated order. Logging configuration details that are set on the IDP device will apply to NSM.
- Packaging log files or debug files for remote analysis
- Managing interface settings such as setting IP addresses, settings IDP device host and network information, interoperability with NSM, Infranet Controllers, Secure Access

devices, settings deployment mode, enabling layer 2 processing, and so on. For more information see the *ACM online Help*.

The following device configurations are not supported:

- Editing licensing information, although licenses can be viewed
- Rebooting the IDP device

On standalone IDP sensors and ISG security module settings inspects the following protocols using [Table 5 on page 7](#).

Table 5: Intrusion Detection and Prevention: Supported Protocols

AIM	HTTP	Oracle	SMTP
CHARGEN	ICMP	POP3	SNMP/Trap
DHCP	IDENT	PortMapper	SQL Mon
DISCARD	IKE	RADIUS	SSH
DNS	IMAP	Rexec	SSL
ECHO	IRC	rlogin	Syslog
FINGER	LDAP	SunRPC	TELNET
FTP	LPR	Rsh	TFTP
GNUTELLA	MSN	RTSP	VNC
GOPHER	MSRPC	NBNAME	WHOIS
GRE*	MS-SQL	NFS	Yahoo Messenger
H.225**	GTP	NNTOP	
	NTP	Rusers	
		SMB	

* GRE inspection are supported only for IP (protocol 0x0800) and PPP for CDMA A10 channel (protocol 0x8881). PPP is a Layer 2 protocol, which can carry any Layer 3 protocols. Within PPP, IDP inspects IP and Van Jacobson compressed TCP.

** Standalone IDP only.

Related Documentation

- [Adding Intrusion Detection and Prevention Clusters in NSM Overview on page 8](#)
- [Using Templates and Configuration Groups in NSM Overview on page 8](#)
- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)

Adding Intrusion Detection and Prevention Devices in NSM Overview

Before NSM can manage IDP devices, you must first add the IDP devices to the management system using the NSM UI. To add an IDP device, you create an object in the UI that represents the physical device, and then create a connection between the UI object and the device so that their information is linked. When you make a change to the UI device object, you can push that information to the real device so the two remain synchronized. You can add a single IDP device at a time or add multiple IDP devices all at once.

For complete details on adding IDP devices, see the *Network and Security Manager Administration Guide*.

- Related Documentation**
- [Adding Intrusion Detection and Prevention Clusters in NSM Overview on page 8](#)
 - [Using Templates and Configuration Groups in NSM Overview on page 8](#)

Adding Intrusion Detection and Prevention Clusters in NSM Overview

In IDP, maximum of two clusters join together to ensure continued network uptime. The device configurations are synchronized, meaning all cluster members share the same configuration settings, enabling an IDP device to handle traffic for another if one device fails.

Adding a cluster is a two-stage process:

- Add the cluster device object.
- Add the members of the cluster to the cluster device object.

For complete details on adding IDP clusters, see the *Network and Security Manager Administration Guide*.

- Related Documentation**
- [Using Templates and Configuration Groups in NSM Overview on page 8](#)
 - [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)

Using Templates and Configuration Groups in NSM Overview

Use templates to define an IDP device configuration and then reuse that configuration information across multiple IDP devices. In a template, you need to define only those configuration parameters that you want to set; you do not need to specify a complete device configuration.

Templates provide these benefits:

- You can configure parameter values for an IDP device by referring to one or more templates when configuring the device.

- When you change a parameter value in a template and save the template, the value also changes for all the IDP device configurations that refer to that template, unless specifically overridden in the device object.

For complete details on using device templates and configuration groups, see the *Network and Security Manager Administration Guide*.

**Related
Documentation**

- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)
- [Intrusion Detection and Prevention Services and Device Configurations Supported in NSM on page 6](#)

PART 2

Configuring Intrusion Detection and Prevention Devices

- [Configuring Profiler Settings on page 13](#)
- [Configuring Security Policies on page 31](#)
- [Working with Attack Objects on page 63](#)
- [Working with Application Objects on page 87](#)
- [Configuring SNMP and Syslog Settings on page 91](#)
- [Configuring Anti-Spoof Settings on page 95](#)
- [Configuring Intrusion Detection and Prevention Device Settings on page 99](#)
- [Configuring Additional Intrusion Detection and Prevention Features on page 123](#)

CHAPTER 3

Configuring Profiler Settings

Before configuring security, you must first enable and set up the Profiler. The Profiler is a network analysis tool that helps you learn about your internal network, enabling you to create effective security policies and minimize unnecessary log records. After you configure the Profiler, it automatically learns about your internal network and the elements that comprise it, including hosts, peers (communication between two hosts), ports (non-IP protocols, TCP/UDP ports, RPC programs), and Layer-7 data that uniquely identifies hosts, applications, commands, users, and filenames.

The Profiler is supported in all IDP modes and HA configurations, and also queries and correlates information from multiple devices. For details on analyzing your network, see the *Network and Security Manager Administration Guide*. This chapter provides information on setting up the Profiler and configuring antivirus settings, including antispam and Web filtering.

- [Configuring Profiler Options \(NSM Procedure\) on page 13](#)
- [Viewing Profiler Logs \(NSM Procedure\) on page 20](#)
- [Modifying Profiler Settings \(NSM Procedure\) on page 25](#)
- [Configuring Profiler Database Preferences \(NSM Procedure\) on page 26](#)
- [Displaying Profiler Database Information \(NSM Procedure\) on page 28](#)
- [Querying the Profiler Database \(NSM Procedure\) on page 28](#)
- [Purging the Profiler Database \(NSM Procedure\) on page 28](#)

Configuring Profiler Options (NSM Procedure)

Profiler option settings are valid for standalone IDP sensors only. For more information, see the *NSM online Help*.

To configure the Profiler on a given IDP sensor, open the Device window and select **Profiler Settings**.

You configure Profiler options to enable Profiler features, set network addresses and applications subject to profiling, and set alerts.

Setting Up the Profiler

Using the Profiler involves the following steps:

- Collecting specific information about your internal network
- Starting the Profiler to enable your device to begin collecting data
- Customizing Profiler preferences

You configure your device to collect specific information and compile it into the Profiler database.

Configuring the Profiler

You can configure the Profiler using the Profiler settings available on the device settings in the Device Manager. Using the Device Manager, double-click to access a device managed in NSM, and click **Profiler Settings**.

The Profile Configuration dialog box appears with the General tab selected. Once you select the device for profiling, you can configure the options for the device to collect data from your internal network.

The following topics describe the steps to configure Profiler options:

- [Specifying General Options on page 14](#)
- [Specifying Tracked Hosts on page 16](#)
- [Specifying Context Targets on page 17](#)
- [Specifying Alert Options on page 18](#)

Specifying General Options

In this tab, indicate whether you want to enable Application Profiling and Probe and Attempt and whether Non-tracked IP Profiles will be included in the profiling. Also indicate the size of the Profiler database and whether to enable OS fingerprinting.

You configure Profiler general options to enable Profiler features.

OS fingerprinting passively detects the operating system of an end-host by analyzing TCP handshake packets. To ensure that this works, you need to verify that OS fingerprinting is first enabled on the profiled device. After you have configured the Profiler with the tracked hosts and contexts, you must update the device.

OS fingerprinting works only for packets that contain a full-fledged TCP connection, that is the TCP connection should have a SYN, SYN/ACK, and a FIN connection. OS fingerprinting only works for operating systems that are supported on the device. A list of the supported operating systems is available on the device in a file called **fingerprints.set** at the following location:

```
/usr/idp/device/cfg/fingerprints.set
```

Configuring Network Objects

The first part of configuring the Profiler is to inform the device which network objects you want the device to profile. When you start the Profiler, the device begins collecting data from the selected hosts.

To specify Profiler general options:

1. From Device Manager, double-click a device and then click **Profiler Settings**.
2. Click the **General** tab.
3. Configure Profiler general options using [Table 6 on page 15](#).
4. Click **Apply**.

Table 6: Profiler Settings: General Tab

Setting	Description
Enable Profiling	Enables the Profiler.
Enable Application Profiling	Enables the Profiler to collect and track application data. This setting can be started when you disable it in the profiler setting.
Enable Application Volume Tracking	Enables the Profiler to perform application volume tracking.
Include Probe and Attempt	Enables the Profiler to collect and track specific probes and attempts.
Include Non-tracked IP Profiles	Enables the Profiler to collect and track data from external hosts.
db limit (in MB)	Specifies maximum database size for the Profiler on each device. By default, the maximum database size is set to 3GB.
Enable OS fingerprinting	Enables the Profiler to perform passive OS fingerprinting to determine the operating system of an end host. OS fingerprinting detects the operating system of an end host by analyzing TCP handshake packets. The OS fingerprinting process depends on an established TCP connection (one that has a SYN and a SYN/ACK). The OS fingerprinting process is capable of detecting the operating systems listed in <code>/usr/idp/device/cfg/fingerprints.set</code> .
Refresh Interval(in secs)	Specifies the time interval (in seconds) that the Profiler refreshes OS fingerprinting. By default, the Profiler refreshes OS fingerprinting data every 3600 seconds (60 minutes).



NOTE: If you change Profiler settings, you must push a configuration update to the device before the new settings take effect. From the Device Manager, right-click the device, select **Update Device**, select the **Restart IDP Profiler After Device Update** checkbox, and click **OK**.

Specifying Tracked Hosts

Select the known hosts you want to track in the Tracked Hosts tab. Select **Object Manager > Address Objects** to add entries to the hosts list.

In the Tracked Hosts tab, select the Network Objects that represent your internal hosts. The device collects detailed information about traffic that passes between internal hosts, and then groups traffic that does not match an internal host in a special IP: 73.78.69.84. Communication between an internal host and an external host is recorded only once. For example, the device records internal host A communicating to www.yahoo.com and www.cnn.com as one entry in the Profiler database. You can select unlimited internal network objects. You can also use the Exclude List tab to select the Network Objects that represent internal hosts that you do not want to include in IDP profiling. You might want to exclude a host from the Profiler if you select a group of network objects in the Tracked Host tab. Also, you might want to exclude specific members of that group.

You configure Profiler tracked host and excluded host settings to determine the network segments where the Profiler gathers data.

To configure the tracked host and exclude lists:

1. From Device Manager, double-click a device and then click **Profiler Settings**.
2. Click the **Tracked Hosts** tab.
3. Click the + icon and then select **Add Host > Add Network** or **Add Group**. A dialog box appears where you create your tracked hosts list.
4. Configure Profiler tracked host settings using [Table 7 on page 16](#).

Table 7: Profiler Tracked Hosts/Exclude List Dialog Boxes

Setting	Description
Add Host	Name—Enter the name of the host.
	Color—Select any color from the drop-down list.
	Comment—Enter any additional comments.
	IP/IP Address—Enter the IP address when you select IP.
	Domain/Domain name—Enter the domain name when you select domain name.
	Resolve—Resolve the domain name with the IP and vice versa.

Table 7: Profiler Tracked Hosts/Exclude List Dialog Boxes (*continued*)

Setting	Description
Add Network	Name—Enter the name of the host.
	IP Address—Enter the IP address of the network.
	Use Wildcard Mask—Enable this feature if you want to use wildcard mask.
	Netmask—Enter the netmask for the IP.
	Color—Select any color from the drop-down list.
	Comment—Enter any additional comments.
Add Group	Name—Enter the name of the group.
	Color—Select any color from the drop-down list.
	Comment—Enter any additional comments.
	Member List—Add or remove the members from the non-members list.

- Click the **Exclude List** tab.
- Click the + icon and then select **Add Host > Add Network** or **Add Group**. A dialog box appears where you create your exclude list.
[Table 7 on page 16](#) describes these dialog box settings.
- Configure Profiler settings using [Table 7 on page 16](#).
- Click **Apply**.



NOTE: If you change Profiler settings, you must push a configuration update to the device before the new settings take effect. From the Device Manager, right-click the device, select **Update Device**, select the **Restart IDP Profiler After Device Update** check box, and click **OK**.

Specifying Context Targets

Select the contexts you want to profile in the Context Targets tab. Next, determine which contexts you want the device to record. In the Contexts to Profile tab, the context list includes only the contexts that can clearly identify a host, a user, and/or an application. When you start the Profiler, the device begins collecting data on traffic that matches the selected contexts. For example, To track FTP logins, usernames, and commands, select the FTP contexts in the Contexts to Profile tab. After the Profiler is started, the device begins collecting information about FTP logins, usernames, and commands, enabling

you to quickly identify the users using FTP on your network and the actions they perform over that protocol.

When you first configure the Profiler, select all contexts. This enables the device to collect data about every context on your network, giving you a complete view of your network traffic. Later, when you have analyzed your traffic, you can eliminate contexts that you know will not be used on your network.

Select **Profile Context** to include context information. If you clear **Profile Context**, IDP profile data only includes high-level traffic data such as source, destination, and service. If you want Profiler information to include context values and network probes (for example, port scans), also configure the Profiler to include probes and attempts.

You configure Profiler context settings to determine whether Profiler logs include not only host and application data but also data pulled from application contexts. For example, if you specify context targets for FTP usernames, the Profiler logs will include the username specified for the FTP connection in addition to the hostname and service (FTP).

To specify Profiler context targets:

1. From Device Manager, double-click a device and then click **Profiler Settings**.
2. Click the **Contexts To Profile** tab.
3. Browse and select from the predefined list of contexts.
4. Click **Apply**.



NOTE: If you change Profiler settings, you must push a configuration update to the device before the new settings take effect. From the Device Manager, right-click the device, select **Update Device**, check **Restart IDP Profiler After Device Update**, and click **OK**.

Specifying Alert Options

Indicate which profiler events you want to generate alerts for in the Alert Options tab. Use this tab to configure the Profiler to indicate the appearance of a new host, protocol, or port on your internal network. When you select **New Host Detected**, **New Protocol Detected**, or **New Port Detected**, the device generates a specific log record, such as PROFILER_NEW_HOST, in the Profiler Logs section of the Log Viewer when the device discovers a new host, protocol, or port.

If you are configuring the Profiler for the first time, do not enable the new host, protocol, or port alerts. As the Profiler runs, the device views all network components as new, which can generate unnecessary log records. After the Profiler has learned about your network and has established a baseline of network activity, you should reconfigure the device to record new hosts, protocols, or ports discovered on your internal network. For details, see the *Network and Security Manager Administration Guide*.

Select the **Database Limit Exceeded** alert to indicate when you have reached the maximum limit of the database size. You can configure the maximum limit of the Profiler database using the `dbLimit` parameter in the General tab of the Profiler Configuration dialog box. The default is 500 MB; the minimum-maximum range is 0 to 500 MB. After a device reaches this limit, it begins purging the database. For example, a network host performs the normal connections required for Internet connectivity (SMTP, POP3, HTTP, and so on). If the host is infected by a worm, it begins making outbound connections on an arbitrary port. The device logs the unique event and generates `PROFILER_NEW_PROTO` and `PROFILER_NEW_PORT` log records. The system immediately e-mails these log records to the security administrator, who can investigate the worm and take action to contain it.

Repeat the configuration process for each device in your network. When you have configured all devices on your network, you are ready to start the Profiler.

You configure Profiler alert options to determine whether you receive alerts when Profiler detects new hosts, protocols, or ports in use.

If you are configuring the Profiler for the first time, do not enable the new host, protocol, or port alerts. As the Profiler runs, the device views all network components as new, which can generate unnecessary log records. After the Profiler has learned about your network and has established a baseline of network activity, you should reconfigure the device to record new hosts, protocols, or ports discovered on your internal network.

To specify Profiler alert options:

1. From Device Manager, double-click a device and then click **Profiler Settings**.
2. Click the **Alert** tab.
3. Configure alert settings using [Table 8 on page 19](#).
4. Click **Apply**.
5. Click **OK**.

Table 8: Profiler Alert Tab

Setting	Description
New Host Detected	Sends an alert when Profiler detects a new host.
New Protocol Detected	Sends an alert when Profiler detects a new protocol. New Protocol detection alerts are used only for Layer 3 protocols.
New Port Detected	Sends an alert when Profiler detects a new port.
Database Limit Exceeded	Sends an alert to indicate the maximum database size has been reached. After a device reaches this limit, it begins purging the database.



NOTE: If you change Profiler settings, you must push a configuration update to the device before the new settings take effect. From the Device Manager, right-click the device, select **Update Device**, select the **Restart IDP Profiler After Device Update** checkbox, and click **OK**.

**Related
Documentation**

- [Configuring Profiler Database Preferences \(NSM Procedure\) on page 26](#)
- [Querying the Profiler Database \(NSM Procedure\) on page 28](#)
- [Purging the Profiler Database \(NSM Procedure\) on page 28](#)
- [Viewing Profiler Logs \(NSM Procedure\) on page 20](#)

Viewing Profiler Logs (NSM Procedure)

The Profiler Viewer contains multiple tabs with different views of Profiler data. The following topics describe these views:

- [Application Profiler on page 20](#)
- [Protocol Profiler on page 22](#)
- [Network Profiler on page 23](#)
- [Violation Viewer on page 25](#)

Application Profiler

The Application Profiler tab displays Application Volume Tracking (AVT) data. The Application Profiler tab is a table of information such as the NSM Log Viewer which enables you to view and analyze dynamic application (Layer-7) traffic within a specific context.

The Application Profiler view is divided into two sections:

- In the left pane, the Application Profiler tab displays a hierarchical tree of application categories. Applications are grouped by common functionality. For example, Peer-to-Peer applications include Chat and File Sharing applications. Under Chat, you can display Yahoo messenger, MSN, and AIM; under File Sharing, you can display Kazaa, Bittorrent, and Gnutella.

The left pane also displays aggregate statistics for volume (bytes) and packet count for the application category, application group, or application you select in the tree.

- In the right pane, the Application Profiler tab displays tables of session logs related to the application category or application you select in the left pane.

[Table 9 on page 21](#) describes Application Profiler session table.

Table 9: Application Profiler Session

Column	Description
Src IP	Source IP address of the traffic profiled.
Dst IP	Destination IP address of the traffic profiled.
User	The user associated with the traffic profiled.
Role	The role group to which the user that is associated with the traffic profiled belongs.
Context	All contexts of traffic that the devices selected in the Device table recorded.
Value	When you select a context, the values that your devices recorded for a selected context.
Src MAC	Source MAC addresses of traffic profiled.
Dst MAC	Destination MAC addresses of traffic profiled.
Src OUI	Source OUIs of traffic profiled.
	NOTE: The Organizationally Unique Identifier (OUI) value is a mapping of the first three bytes of the MAC address and the organization that owns the block of MACs. You can obtain a list of OUIs at http://standards.ieee.org/regauth/oui/oui.txt .
Dst OUI	Destination OUIs of traffic profiled.
Src OS Name	Operating-system version running on the source IP of the traffic profiled.
Dst OS Name	Operating-system version running on the destination IP of the traffic profiled.
Hits	Number of occurrences that match the traffic profiled.
First Time	Timestamp for the first time the device logged the event (within the specified time interval).
Last Time	Timestamp for the last time the device logged the event (within the specified time interval).
Domain	Domain in which the device is managed in NSM.
Device	Device that profiled the data displayed.

By default, the Application Profiler view contains only the data collected during the configured time interval.

To display the Application Profiler view:

1. Navigate to **Investigate > Security Monitor > Profiler**.
2. Click the **Application Profiler** tab.



TIP: You can jump from the Application Profiler tab to the APE rulebase editor by right-clicking an application in the left pane and selecting a policy editor option. For information about using NSM features to sort, filter, and drill down on records, see the *NSM online help*.

Protocol Profiler

The Protocol Profiler tab displays information about applications that are running on your network.

Table 10 on page 22 describes the protocol profiler data.

Table 10: Protocol Profiler Data

Column	Description
Src IP	Source IP address of the session. NOTE: Profiler tracks all traffic through the IDP appliance, including traffic for hosts not in your tracked hosts list. It records a value of 73.78.69.84 for the IP address for hosts not defined in the Tracked Hosts tab, such as external hosts you would not know and therefore could not configure.
Dst IP	Destination IP address. NOTE: Communication between an internal host and an external host is recorded only once. For example, the device records internal host A communicating to http://ca.yahoo.com and http://edition.cnn.com as one entry in the Profiler DB.
User	The user associated with the session.
Role	The role to which the user belongs.
Context	Matching contexts.
Value	Value retrieved from matching context.
Src MAC	Source MAC addresses.
Dst MAC	Destination MAC addresses.
Src OUI	Source OUI. NOTE: The Organizationally Unique Identifier (OUI) value is a mapping of the first three bytes of the MAC address and the organization that owns the block of MACs. You can obtain a list of OUIs at http://standards.ieee.org/regauth/oui/oui.txt .
Dst OUI	Destination OUI.
Src OS Name	Operating-system version running on the source IP.
Dst OS Name	Operating-system version running on the destination IP.

Table 10: Protocol Profiler Data (*continued*)

Column	Description
Hits	Number of occurrences that match the session.
First Time	Timestamp for the first time the device logged the event (within the specified time interval).
Last Time	Timestamp for the last time the device logged the event (within the specified time interval).
Domain	NSM domain.
Device	Device through which the session was forwarded.

By default, the Protocol Profiler tab contains only the data collected during the configured time interval.

To display the Protocol Profiler tab:

1. In the NSM navigation tree, select **Investigate > Security Monitor > Profiler**.
2. Click the **Protocol Profiler** tab.



TIP: For information about using NSM features to sort, filter, and drill down in records, see the *NSM online Help*.

Network Profiler

The Network Profiler view is a table of information such as the NSM Log Viewer which enables you to view and analyze data related to static traffic (Layer-3, Layer-4, and RPC protocols, ports, and program numbers) within the context of data corresponding to peer, host, and operating system.

Table 11 on page 23 describes Network Profiler data.

Table 11: Network Profiler Data

Column	Description
Src IP	Source IP address of the traffic profiled.
Dst IP	Destination IP address of the traffic profiled.
User	The user associated with the traffic profiled.
Role	The role group to which the user that is associated with the traffic profiled belongs.
Service	All services of traffic profiled.

Table 11: Network Profiler Data (*continued*)

Column	Description
Access Type	Type of the traffic profiled: <ul style="list-style-type: none"> Access indicates a successful connection, during which the device recorded valid requests and responses from the server to a client. Attempt indicates a request that did not receive a reply. The device recorded a packet from a client to a server, but never saw a reply. Probe indicates a request that does not expect a reply. For non-TCP sessions, the device recorded an ICMP error; for TCP sessions, the device recorded a SYN packet from the client followed by a RST from the server.
Src MAC	Source MAC addresses of traffic profiled.
Dst MAC	Destination MAC addresses of traffic profiled.
Src OUI	Source OUIs of traffic profiled. NOTE: OUI stands for Organizationally Unique Identifier. This value is a mapping of the first three bytes of the MAC address and the organization that owns the block of MACs. You can obtain a list of OUIs at http://standards.ieee.org/regauth/oui/oui.txt .
Dst OUI	Destination OUIs of traffic profiled.
Src OS Name	Operating-system version running on the source IP of the traffic profiled.
Dst OS Name	Operating-system version running on the destination IP of the traffic profiled.
Hits	Number of occurrences that match the traffic profiled.
First Time	Timestamp for the first time the device logged the event (within the specified time interval).
Last Time	Timestamp for the last time the device logged the event (within the specified time interval).
Domain	Domain in which the device is managed in NSM.
Device	Device that profiled the data displayed.

To display the Network Profiler view:

1. Navigate to **Investigate > Security Monitor > Profiler**.
2. Click the **Network Profiler** tab.



TIP: For information about using NSM features to sort, filter, and drill down on records, see the *NSM online Help*.

Violation Viewer

The Violation Viewer is a filtered view of the Network Profiler view. In the Violation Viewer, you configure permitted objects. Permitted objects are filtered from the display, allowing you to focus on unpermitted traffic.

To configure permitted objects:

1. Navigate to **Investigate > Security Monitor > Profiler**.
2. Click the **Violation Viewer** tab.
3. Click on the + icon that appears on the top of the right-hand window to display the New Permitted Object dialog box.
4. Type a name for the permitted object.
5. Within the New Permitted Object dialog box, click the + icon to add a rule to match source, destination, and services values for the permitted object.
6. To change the source, destination, or service value from **Any**, right-click the value and select **Add Source**, **Add Destination**, or **Add Service**.
7. Use the selection controls to select the desired address objects or service objects and click **OK**.
8. Click **OK** to save the permitted object.

The object appears in the filters windows within the Violation Viewer tab.

9. Select the object and click **Apply** to hide all matching objects from the Violation Viewer.



TIP: For information about using additional NSM features to sort, filter, and drill down on records, see the *NSM online Help*.

Related Documentation

- [Configuring Profiler Options \(NSM Procedure\) on page 13](#)
- [Displaying Profiler Database Information \(NSM Procedure\) on page 28](#)
- [Querying the Profiler Database \(NSM Procedure\) on page 28](#)

Modifying Profiler Settings (NSM Procedure)

You can use the Profiler Settings dialog box to modify Profiler database settings and default settings for application volume tracking reports.

To modify profiler database and application volume tracking settings:

1. From the NSM main menu, select **Tools > Preferences**. The New Preferences dialog box is displayed.
2. Click **Profiler Settings**.

3. Modify settings as described in [Table 12 on page 26](#).
4. Click **OK**.

Table 12: Profiler Settings

Setting	Description
Purge profiler database if size exceeds (in MB)	Removes the profiler database for the selected size in MB. The default value is 1000 MB.
Max profiler database size after purging (in MB)	Specifies the maximum size of the purged profiler database. The default value is 750 MB.
Profiler query timeout (in seconds)	Specifies the timeout entry for a profiler query. The default value is 120 seconds.
Hour of day to perform db optimization (local time)	Specifies the time to perform the database optimization. The default value is 00:00 GMT.
Number of sessions to display per application	<p>Determines the number of sessions displayed in the Application Profiler application volume tracking session tables.</p> <p>The default value is 10 sessions. You can specify from 5 to 10,000 sessions.</p>
Hours of session data to display from present time	<p>Determines the hours of application volume tracking data displayed in the Application Profiler tab session tables.</p> <p>The default value is 1 hour. You can specify from 1 to 24 hours.</p> <p>This setting is also a data retention policy. By default, data older than 1 hour is deleted. If your change to 12 hours, data older than 12 hours is deleted.</p>

- Related Documentation**
- [Configuring Profiler Options \(NSM Procedure\) on page 13](#)
 - [Querying the Profiler Database \(NSM Procedure\) on page 28](#)

Configuring Profiler Database Preferences (NSM Procedure)

ScreenOS 6.2 supports application volume tracking (AVT), a feature that enables NSM to track network bandwidth usage on a per-application basis. The security device sends the NSM server periodic update messages containing details about port activity. NSM listens and processes these periodic update messages and maintains a cumulative count for each port. NSM displays this count on the console. NSM provides reports about application volume tracking. The AVT feature has the following limitations:

- Periodic updates maintained per port for each active session can slightly affect CPU performance.
- Accuracy of AVT data is dependent on communication with the NSM server. NSM, however, lacks a mechanism to ensure that periodic updates sent by AVT from ScreenOS are received, which may result in a lag between traffic instances and reporting of those instances. NSM maintains a cumulative count for all traffic on each port regardless of session, node, or protocol. The count displayed is a total across all

sessions. Because updates are periodic, the currently displayed number of bytes in NSM may be inaccurate until the next update.

- You must use NSM to view the enhanced logging provided by AVT.

For more details on AVT, see the *Network and Security Manager Administration Guide*.

Use the Profiler Settings under the Tools menu to configure the Profiler preferences mentioned in [Table 12 on page 26](#). You can use the Profiler Settings dialog box to modify Profiler database settings and default settings for application volume tracking reports.

Data discovered by Profiler is stored in a database located on the NSM GUI server.

To modify profiler database preferences and application volume tracking settings:

1. From the NSM main menu, select **Tools > Preferences**. The New Preferences dialog box is displayed.
2. Click **Profiler Settings**.
3. Modify settings as described in [Table 12 on page 26](#).
4. Click **OK**.

Table 13: Profiler Database Preferences

Setting	Description
Purge profiler database if size exceeds (in MB)	NSM purges the profiler database size if it exceeds 4GB (4000 MB) by default.
Max profiler database size after purging (in MB)	If the database size exceeds its maximum limit, NSM purges the Profiler database size until the size reaches 3 GB (3000 MB) by default.
Profiler query timeout (in seconds)	Specifies the timeout entry for a profiler query. The default value is 120 seconds.
Hour of day to perform db optimization (local time)	Specifies the time to perform the database optimization. The default value is 00:00 GMT.
Number of sessions to display per application	Determines the number of application volume tracking sessions displayed in the Application Profiler tab session tables. The default value is 10 sessions. You can specify from 5 to 10,000 sessions.
Hours of session data to display from present time	Determines the hours of application volume tracking data displayed in the Application Profiler tab session tables. The default value is 1 hour. You can specify from 1 to 24 hours. This setting is also a data retention policy. By default, data older than 1 hour is deleted. If your change to 12 hours, data older than 12 hours is deleted.

- Related Documentation**
- [Configuring Profiler Options \(NSM Procedure\) on page 13](#)
 - [Displaying Profiler Database Information \(NSM Procedure\) on page 28](#)

- [Viewing Profiler Logs \(NSM Procedure\) on page 20](#)

Displaying Profiler Database Information (NSM Procedure)

Purpose	Data discovered by Profiler is stored in a database located on the NSM GUI server. Use the steps in this procedure to display information about the Profiler database.
Action	To display Profiler database information: <ol style="list-style-type: none">1. In the NSM Navigation tree select Investigate > Security Monitor > Profiler.2. Click the Show DB Information icon in the upper right corner to view specific details about the Profiler database, including the database size.
Related Documentation	<ul style="list-style-type: none">• Configuring Profiler Options (NSM Procedure) on page 13• Configuring Profiler Database Preferences (NSM Procedure) on page 26• Querying the Profiler Database (NSM Procedure) on page 28

Querying the Profiler Database (NSM Procedure)

Purpose	Data discovered by Profiler is stored in a database located on the NSM GUI server. Use the steps in this procedure to query the Profiler database.
Action	To query records in the database: <ol style="list-style-type: none">1. Log into the NSM GUI server as the Postgres SQL user. By default, the Postgres SQL user is netSCREEN.2. Navigate to the directory where the Profiler DB is located: /usr/local/nsmpsqr/bin.3. Run any Postgres SQL command. For example, you can type the following command: <pre>./psql -d profilerDb</pre>
Related Documentation	<ul style="list-style-type: none">• Configuring Profiler Database Preferences (NSM Procedure) on page 26• Configuring Profiler Options (NSM Procedure) on page 13• Displaying Profiler Database Information (NSM Procedure) on page 28• Purging the Profiler Database (NSM Procedure) on page 28

Purging the Profiler Database (NSM Procedure)

Data discovered by Profiler is stored in a database located on the NSM GUI server. When the database reaches a maximum size (4 GB by default), it begins purging records (oldest first) automatically. The Profiler stops purging records when it reaches a certain set minimum size (3 GB by default).

Use the steps in this procedure to purge the Profiler database, if needed.

To change automatic purge settings, from the NSM main menu select **Tools > Preferences** and modify the Profiler database settings.

To manually purge the database:

1. In the NSM Navigation tree, select **Investigate > Security Monitor > Profiler**.
2. Click the **Clear All DB** icon in the upper right corner.

**Related
Documentation**

- [Configuring Profiler Options \(NSM Procedure\) on page 13](#)
- [Configuring Profiler Database Preferences \(NSM Procedure\) on page 26](#)
- [Displaying Profiler Database Information \(NSM Procedure\) on page 28](#)
- [Querying the Profiler Database \(NSM Procedure\) on page 28](#)

CHAPTER 4

Configuring Security Policies

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Configuring Predefined Security Policies \(NSM Procedure\) on page 33](#)
- [Creating a New Security Policy \(NSM Procedure\) on page 34](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)
- [Configuring Exempt Rulebase Rules \(NSM Procedure\) on page 45](#)
- [Configuring Backdoor Rulebase Rules \(NSM Procedure\) on page 47](#)
- [Configuring SYN Protector Rulebase Rules \(NSM Procedure\) on page 49](#)
- [Configuring Traffic Anomalies Rulebase Rules \(NSM Procedure\) on page 51](#)
- [Configuring Network Honeypot Rulebase Rules \(NSM Procedure\) on page 54](#)
- [Configuring Application Rulebase Rules \(NSM Procedure\) on page 57](#)

Intrusion Detection and Prevention Devices and Security Policies Overview

An IDP security policy defines how the IDP device handles network traffic. It allows you to enforce various attack detection and prevention techniques on traffic that traverses your network.

For a detailed explanation of security policy features and components, and for examples, see the *IDP Concepts & Examples Guide*.

To create an effective security policy, follow these basic steps:

1. Run the New Policy wizard to create a new security policy object. The new security policy can be based on a predefined template.
2. Use the Security Policy editor to add one or more rulebases.

A *rulebase* is an ordered set of rules that use a particular detection method to identify and prevent attacks.

[Table 14 on page 32](#) describes the IDP security policy rulebases. A security policy can contain only one instance of any rulebase type.

Table 14: IDP Security Policy Rulebases

Rulebase	Description
Application Rulebase	Enables you to limit bandwidth for specified users and applications and thus helps to manage network traffic. APE rules do not use attack objects..
IDP Rulebase	Protects your network from attacks by using attack objects to detect known and unknown attacks. Juniper Networks provides predefined attack objects that you can use in IDP rules. You can also configure your own custom attack objects.
Exempt Rulebase	You configure rules in this rulebase to exclude known false positives or to exclude a specific source, destination, or source/destination pair from matching an IDP rule. If traffic matches a rule in the IDP rulebase, IDP attempts to match the traffic against the Exempt rulebase before performing the action specified.
Backdoor Rulebase	Protects your network from mechanisms installed on a host computer that facilitates unauthorized access to the system. Attackers who have already compromised a system typically install backdoors (such as Trojans) to make future attacks easier. When attackers send and retrieve information to and from the backdoor program (as when typing commands), they generate interactive traffic that IDP can detect.
SYN Protector Rulebase	Protects your network from SYN-floods by ensuring that the three-way handshake is performed successfully for specified TCP traffic. If you know that your network is vulnerable to a SYN-flood, use the SYN-Protector rulebase to prevent it.
Traffic Anomalies Rulebase	Protects your network from attacks by using traffic flow analysis to identify attacks that occur over multiple connections and sessions (such as scans).
Network Honeypot Rulebase	Protects your network by impersonating open ports on existing servers on your network, alerting you to attackers performing port scans and other information-gathering activities.

3. Within rulebases, configure rules.

Rules are instructions that provide context to detection methods. Rules specify:

- A source/destination/service match condition that determines which traffic to inspect
- Attack objects that determine what to look for (IDP rulebase and Exempt rulebase)
- Actions that determine what to do when an attack is detected
- Notification options, including logs, alerts, and packet captures

Each rulebase can contain up to 40,000 rules.

4. Fine-tune your security policy as you learn more about your network and security requirements and IDP capabilities.

Related Documentation

- [Configuring Predefined Security Policies \(NSM Procedure\) on page 33](#)
- [Creating a New Security Policy \(NSM Procedure\) on page 34](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)

- [Troubleshooting Security Policy Validation Errors \(NSM Procedure\)](#) on page 130

Configuring Predefined Security Policies (NSM Procedure)

The highly respected Juniper Networks Security Center team (J-Security Center) provides the default IDP security policy—named Recommended. We advise that you use this policy to protect your network from the likeliest and most dangerous attacks.

[Table 15 on page 33](#) summarizes the properties of the Recommended security policy.

Table 15: Recommended Security Policy Definition

Property	Value
Rulebase	IDP Rulebase
Rules	9 rules, distinguished by attack object
Traffic source	Any
Service	Default, meaning the matching property is based on the service bindings of the attack object specified by the rule
Destination	Any
Attacks	Recommended IP, Recommended TCP, Recommended ICMP, Recommended HTTP, Recommended SMTP, Recommended DNS, Recommended FTP, Recommended POP3, Recommended IMAP, Recommended Trojan, Recommended Virus, Recommended Worm
Action	Recommended, meaning the action is specified by the attack object
Notification	Logging

If you prefer, you can copy this security policy and use it as a template for a custom security policy tailored for your network.

[Table 16 on page 33](#) describes other IDP security policy templates.

Table 16: IDP Security Policy Templates

Template	Description
all_with_logging	Includes all attack objects and enables packet logging for all rules.
all_without_logging	Includes all attack objects but does not enable packet logging.
dmz_services	Protects a typical DMZ environment.
dns_server	Protects DNS services.
file_server	Protects file sharing services, such as SMB, NFS, FTP, and others.

Table 16: IDP Security Policy Templates (*continued*)

Template	Description
getting_started	Contains very open rules. Useful in controlled lab environments, but should not be deployed on heavy traffic live networks.
idp_default	Contains a good blend of security and performance.
web_server	Protects HTTP servers from remote attacks.

If you use these templates, we advise you customize them for your deployment. At a minimum, you should change the destination IP setting from **Any** to the IP addresses for specific servers you want to protect. For more information, see the *IDP Concepts & Examples guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Creating a New Security Policy \(NSM Procedure\) on page 34](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)

Creating a New Security Policy (NSM Procedure)

You use the security policy wizard to create a new security policy. The security policies you create with the wizard must have a new name but can be based on existing policies or templates.

To create a new security policy:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select **File > New Policy** to display the New Policy wizard.
3. On the first page, complete the settings and then click **Next**. [Table 17 on page 34](#) describes page one fields.

Table 17: New Policy Wizard: Page One

Setting	Description
Name	A string to identify the policy.
Comments	Text to further identify the policy. In the security policy list, you can sort on comments.

4. On the second page, complete the settings and then click **Next**. [Table 18 on page 35](#) describes page two settings.

Table 18: New Policy Wizard: Page Two

Setting	Description
Create new Policy for	<p>Select this option to create a new security policy.</p> <p>If you select this option, the wizard displays the following set of device types:</p> <ul style="list-style-type: none"> • Firewall/VPN • Firewall/VPN with IDP • Standalone IDP <p>Select Standalone IDP.</p>
Use Existing Policy	<p>Use this option to assign an existing policy to one or more IDP devices.</p> <p>If you select this option, the wizard displays a drop-down list of existing policies.</p> <p>Select a policy from the list.</p> <p>NOTE: This procedure involves creating a new policy. For this procedure, do not select Use Existing Policy.</p>

5. On the next pages, complete pre-configuration options. [Table 19 on page 35](#) describes your choices. Click **Next** to advance through the pages.

Table 19: New Policy Wizard: Pre-configuration Options

Setting	Description
Use Predefined Policy Template	<p>Select this option to create a new security policy based on a predefined template.</p> <p>If you select this option, the wizard displays a drop-down list of predefined templates.</p> <p>Select one and click Next.</p>
Configure IDP Policy	<p>Select this option and complete the rule properties on the next page to generate a policy with the following features:</p> <ul style="list-style-type: none"> • IDP rulebase • Multiple rules matching any source, any destination, and default services • Multiple rules are distinguished by the attack object severity group, action, and notification option you configure in the next wizard page.
Empty Policy	<p>Select this option to create an empty policy that you can later modify.</p>

6. On the next to last page, select IDP devices for which you are designing this policy. Then click **Next**.

7. Click **Finish** to save the policy.

The new policy appears in the security policy list. After you have created a security policy, you can add rules to the new policy. Rules include IPv6, VPN, and also VPN link. For more information, see the *IDP Concepts & Examples guide*

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Configuring Predefined Security Policies \(NSM Procedure\) on page 33](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)

Modifying IDP Rulebase Rules (NSM Procedure)

This procedure assumes you have used the New Policy wizard to create a basic policy that you can modify.

The primary IDP security policy rulebase is the IDP rulebase. The IDP rulebase enables the IDP process engine to inspect matching traffic for signs of an attack.

For background on and examples of IDP rulebase rules, see the *IDP Concepts & Examples Guide*.

To modify IDP rulebase rules:

1. In the NSM navigation tree, select **Configure > Policy Manager > Security Policies**.
2. Select the security policy you want to edit.
3. In the security policy pane, select **IDP** tab to display the IDP rulebase table.
4. To add, delete, copy, or reorder rules, right-click the table cell for the rule number and make your selection.
5. To modify the property of a rule, right-click the table cell for the property and make your selection. [Table 20 on page 36](#) lists the rule properties you can modify and provides references documentation for these properties.

Table 20: IDP Rulebase Rule Properties

Property	Reference
ID	Identification number of the IDP rules that you add.
Match	You can select the zone from which the source sends traffic to the destination zone.
Look For	You can select the attacks that you want add IDP to match in the monitored traffic.
Action	Specifies the action you want IDP to perform against the current connection.
IP Action	Specifies the action you want IDP to perform against future connections that use the same IP address.
Notification	You can choose none, or enable logging and select the logging options that are appropriate for your network.
VLAN Tag	Specifies the VLAN tags you want to match in applying the rule.

Table 20: IDP Rulebase Rule Properties (*continued*)

Property	Reference
Severity	You can use the default severity settings of the selected attack objects, or you can choose a specific severity for your rule.
Install On	Specifies the selected source and destination zone that are available on the security device.
Optional Fields	Specifies the optional fields that you can configure in the rule.
Comments	Describes any additional comments about the rule.

Following are the updates that you can perform on an IDP rulebase rule:

- [Specifying Rule Match Conditions on page 37](#)
- [Specifying IDP Rulebase Attack Objects on page 38](#)
- [Specifying Rule Session Action on page 39](#)
- [Specifying Rule IP Action on page 41](#)
- [Specifying Rule Notification Options on page 42](#)
- [Specifying Rule VLAN Matches on page 42](#)
- [Specifying Rule Targets on page 43](#)
- [Specifying Rule Severity on page 43](#)
- [Specifying Rule Optional Fields on page 44](#)
- [Specifying Rule Comments on page 44](#)

Specifying Rule Match Conditions

To specify rule match conditions, right-click the table cell and select your setting.

[Table 21 on page 37](#) describes match condition columns for IDP rulebase rules.

Table 21: IDP Rulebase Match Condition Settings

Column	Description
From zone / To zone	Not applicable for standalone IDP devices.
Source	<p>Select Address—Display the Select Address dialog box where you can select address objects for traffic sources.</p> <hr/> <p>Any—Matches any source of traffic. To guard against incoming attacks, you typically specify Any.</p> <hr/> <p>Negate—Matches any except those specified.</p> <p>To use address negation:</p> <ol style="list-style-type: none"> 1. Add the address object. 2. Right-click the address object and select Negate.

Table 21: IDP Rulebase Match Condition Settings (*continued*)

Column	Description
User Role	<p>Select User Role—Displays the Select User Role dialog box where you can select or configure user role matches.</p> <p>If a value for User Role matches, the Source parameter is not consulted.</p> <p>User role-based rules are evaluated before IP source rules. If a user role matches, and if the other match criteria are met, the rule is applied and IP address-based rules are not consulted.</p> <p>NOTE: Matching based on user role depends on integration with Juniper Networks Infranet Controllers.</p>
Destination	<p>Select Address—Display the Select Address dialog box where you can select address objects for destination servers.</p> <p>Any—Matches any destination address.</p> <p>Negate—Specifies any except those specified.</p> <p>To use address negation:</p> <ol style="list-style-type: none"> 1. Add the address object. 2. Right-click the address object and select Negate.
Service	<p>Default—Matches the service(s) specified in the rule attack object(s).</p> <p>If you have enabled the Application Identification (AI) feature, the IDP process engine identifies services even if they are running on nonstandard ports.</p> <p>If you have not enabled AI and specify Default, the IDP process engine assumes that standard ports are used for the service.</p> <p>NOTE: If you do not enable AI and your service uses nonstandard ports, you must create a custom service objects.</p> <p>Any—Matches any service.</p> <p>Select Service—Display the Select Service dialog box where you can select predefined or custom service objects.</p>
Terminate	<p>Enable or Disable—Marks the rule a terminal rule (or clears the mark). If a session matches a terminal rule, the IDP process engine does not load any subsequent rules. It takes action, if any, according to the terminal rule.</p>

Specifying IDP Rulebase Attack Objects

To add attack objects:

1. Right-click the table cell for attacks and select **Select Attacks**.
2. In the All Attacks/Groups box, expand **Attack Groups**.

3. To add attack objects recommended by Juniper Networks Security Center (J-Security Center), expand **Recommended Attacks**, browse groups, and select groups or individual attack objects.
4. To add other predefined attack objects, expand **All Attacks**, browse groups, and select groups or individual attack objects.
5. To add attack objects that belong to custom groups, expand the node for the custom group, browse subgroups, and select groups or individual attack objects.
6. To add custom attack objects that do not belong to groups, expand **Attack List** and select from custom attack objects.
7. Click **OK**.

[Table 22 on page 39](#) describes the attack object group hierarchy for recommended and predefined attack objects provided by J-Security Center.

Table 22: Attack Object Group Hierarchy

Group	Contents
Attack Type	Contains two subgroups: anomaly and signature. Within each subgroup, attack objects are grouped by severity.
Category	Contains subgroups based on category. Within each category, attack objects are grouped by severity.
Operating System	Contains the following subgroups: BSD, Linux, Solaris, and Windows. Within each operating system, attack objects are grouped by services and severity.
Severity	Contains the following subgroups: Critical, Major, Minor, Warning, Info. Within each severity, attack objects are grouped by category. NOTE: Our severity rating is not based on CVSS (Common Vulnerability Scoring System). We do include data from Bugtraq (Symantec) and CVE (Common Vulnerabilities and Exposures).
Web Services	Contains subgroups based on Web services. Within services, attacked objects are grouped by severity.
Miscellaneous	Contains attack objects that have a significant affect on IDP performance.

Specifying Rule Session Action

Actions are responses to sessions that match the source/destination condition and attack object pattern. Actions protects your network from attacks.

If a packet triggers multiple rule actions, the IDP device takes the most severe action. For example, if the rules dictate that a packet will receive a DiffServ marking and be dropped, and then the packet will be dropped.

To specify a rule action, right-click the table cell and select your setting.

[Table 23 on page 40](#) describes the actions you can set for IDP rulebase rules.

Table 23: IDP Rulebase Actions

Action	Description
Recommended	Predefined attack objects include a recommended action. The recommended action is related to severity. Table 24 on page 40 lists the recommended actions by severity.
None	IDP inspects for attacks but takes no action against the connection if an attack is found.
Ignore	IDP does not inspect for attacks and ignores the connection.
Diffserv Marking	IDP assigns the indicated service-differentiation value to the packet, and then passes it on normally. Set the service-differentiation value in the dialog box that appears when you select this action in the rulebase. NOTE: The marking has no effect in sniffer mode.
Drop Packet	IDP drops a matching packet before it can reach its destination but does not close the connection. Use this action to drop packets for attacks in traffic that is prone to spoofing, such as UDP traffic. Dropping a connection for such traffic could result in a DoS that prevents you from receiving traffic from a legitimate source address.
Drop Connection	IDP drops the connection without sending an RST packet to the sender, preventing the traffic from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.
Close Client and Server	IDP closes the connection and sends an RST packet to both the client and the server. If IDP is in sniffer mode, IDP sends an RST packet to both the client and server but does not close the connection.
Close Client	IDP closes the connection to the client but not to the server.
Close Server	IDP closes the connection to the server but not to the client.

[Table 24 on page 40](#) describes the logic applied to the value Recommended, a setting coded in predefined attack objects provided by Juniper Networks Security Center.

Table 24: IDP Rulebase Actions: Recommended Actions by Severity

Severity	Description	Recommended Action
Critical	Attacks attempt to evade an IPS, crash a machine, or gain system-level privileges.	Drop Packet, Drop Connection
Major	Attacks attempt to crash a service, perform a denial of service, install or use a Trojan, or gain user-level access to a host.	Drop Packet, Drop Connection
Minor	Attacks attempt to obtain critical information through directory traversal or information leaks.	None
Warning	Attacks attempt to obtain noncritical information or scan the network. They can also be obsolete attacks (but probably harmless) traffic.	None

Table 24: IDP Rulebase Actions: Recommended Actions by Severity (*continued*)

Severity	Description	Recommended Action
Info	Attacks are normal, harmless traffic containing URLs, DNS lookup failures, and SNMP public community strings. You can use informational attack objects to obtain information about your network.	None



NOTE: Our severity rating is not based on CVSS (Common Vulnerability Scoring System). We do include data from Bugtraq (Symantec) and CVE (Common Vulnerabilities and Exposures).

Specifying Rule IP Action

If the IDP device matches an attack, it can take action not only against the current session but also against future network traffic that uses the same IP address. Such actions are called *IP actions*. By default, the specified IP action is permanent (timeout = 0). If you prefer, you can set a timeout.

To specify an IP action, right-click the table cell and configure options.

Table 25 on page 41 describes IDP rulebase IP actions.

Table 25: IDP Rulebase IP Actions

IP Action	Description
IP Block	<p>IDP blocks the matching connection and future connections that match combinations of the following properties you specify:</p> <ul style="list-style-type: none"> • Source IP address • Source subnet • Protocol • Destination IP Address • Destination Subnet • Destination Port • From Zone
IP Close	<p>IDP closes the matching connection and future connections that match combinations of the following properties you specify:</p> <ul style="list-style-type: none"> • Source IP address • Source subnet • Protocol • Destination IP Address • Destination Subnet • Destination Port • From Zone

Table 25: IDP Rulebase IP Actions (*continued*)

IP Action	Description
IP Notify	IDP does not take any action against future traffic but logs the event or sends an alert.

Specifying Rule Notification Options

Notification options determine how events that match the rule are logged.

To specify notification options, right-click the table cell and configure options.

[Table 26 on page 42](#) describes IDP rulebase notification options.

Table 26: IDP Rulebase Notification Options

Option	Description
Event logs and alerts	<p>You can enable the following delivery and handling options for logs:</p> <ul style="list-style-type: none"> • Send to NSM Log Viewer • Send to NSM Log Viewer and flag as an alert • Send to an e-mail address list • Send to syslog • Send to SNMP trap • Save in XML format • Save in CVS format • Process with a script
Packet captures	<p>Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack, its purpose, whether or not the attack was successful, and any possible damage to your network.</p> <p>If multiple rules with packet capture enabled match the same attack, IDP captures the maximum specified number of packets. For example, you configure rule 1 to capture 10 packets before and after the attack, and you configure rule 2 to capture 5 packets before and after the attack. If both rules match the same attack, IDP attempts to capture 10 packets before and after the attack.</p> <p>You can capture up to 256 packets before the event and 256 packets after the event.</p> <p>NOTE: If necessary, you can improve performance by logging only the packets received after the attack.</p>

Specifying Rule VLAN Matches

If you deploy an IDP device in a virtual local area network (VLAN), you can specify VLAN tags for traffic in IDP rulebase rules.

Normally, rules match source, destination, and service. If your rule specifies a VLAN tag, then the rule must also match the VLAN tag.

To specify that rules match a VLAN tag, right-click the table cell and configure your setting.

Table 27 on page 43 describes VLAN tag settings.

Table 27: IDP Rulebase VLAN Tag Settings

Option	Description
None	Matches only traffic that has no VLAN tag.
Any	Matches traffic with any or no VLAN tag (default).
Select VLAN Tags	Displays the Select VLAN Tags dialog box where you can set a single VLAN tag or a range of VLAN tags.
Delete VLAN Tags	Displays a dialog box that prompts you to confirm you want to delete the VLAN tag match setting.

Specifying Rule Targets

By default, IDP security policy rules can be applied to any IDP device. If you desire, you can specify that the rule applies to only specified IDP devices.

To specify that the rule only applies to specified devices, right-click the table cell and select **Select Target** to display the Select Targeted Devices dialog box, where you can select the specify devices on which the rule is to be applied.

Specifying Rule Severity

Severity is a rating of the danger posed by the threat the rule is designed to prevent.

To specify a rule severity, right-click the table cell and select a severity.

Table 28 on page 43 describes rule severity settings.

Table 28: IDP Rulebase Severity

Severity	Description
Default	Select Default to inherit severity from that specified in the attack object.
Critical	Attacks that attempt to evade an IPS, crash a machine, or gain system-level privileges. We recommend that you drop the packets or drop the connection for such attacks.
Major	Attacks that attempt to crash a service, perform a denial of service, install or use a Trojan, or gain user-level access to a host. We recommend that you drop the packets or drop the connection for such attacks.
Minor	Attacks that attempt to obtain critical information through directory traversal or information leaks. We recommend that you log such attacks.
Warning	Attacks that attempt to obtain noncritical information or scan the network. They can also be obsolete attacks (but probably harmless) traffic. We recommend that you log such attacks.

Table 28: IDP Rulebase Severity (*continued*)

Severity	Description
Info	<p>Attacks that are normal, harmless traffic containing URLs, DNS lookup failures, and SNMP public community strings. You can use informational attack objects to obtain information about your network.</p> <p>We recommend that you log such attacks.</p>



NOTE: Our severity rating is not based on CVSS (Common Vulnerability Scoring System). We do include data from Bugtraq (Symantec) and CVE (Common Vulnerabilities and Exposures).

Specifying Rule Optional Fields

Optional fields are user-defined name-value pairs you can configure if you want to be able to sort rules based on these fields. Optional fields do not affect the functionality of the security policy rule.

To specify optional fields, right-click the table cell and select **Edit Options** to display the Select Policy Custom Options dialog box, where you can configure name-value pairs.

Specifying Rule Comments

Comments are notations about the rule. Comments do not affect the functionality of the security policy rule.

To specify comments, right-click the table cell and select **Edit Comments** to display the Edit Comments dialog box, where you can enter a comment up to 1024 characters in length.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Configuring Predefined Security Policies \(NSM Procedure\) on page 33](#)
- [Configuring Exempt Rulebase Rules \(NSM Procedure\) on page 45](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)

Configuring Exempt Rulebase Rules (NSM Procedure)

The exempt rulebase contains rules that prevent rules in the Intrusion Detection and Prevention (IDP) rulebase from matching specific source or destination pairs for specific attack objects.

The exempt rulebase works in conjunction with the IDP rulebase. Before you can create exempt rules, you must first create rules in the IDP rulebase. If traffic matches a rule in the IDP rulebase, the IDP sensor attempts to match the traffic against the exempt rulebase before performing the specified action or creating a log record for the event.



NOTE: The exempt rulebase is a non-terminal rulebase. The IDP device checks all rules in the exempt rulebase and executes all matches.

To configure an exempt rulebase rule:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select and double-click the security policy for which you want to add an exempt rulebase rule.
3. Click **New** in the upper right corner of the policy viewer and select **Add Exempt Rulebase**.
4. Click the **New** button within the rules viewer to add a rule.
5. Modify the property of the rule by right-clicking the table cell for the property and making your modifications.
6. Configure or modify the rule using the settings described in [Table 29 on page 45](#).

Table 29: Exempt Rulebase Rule Properties

Option	Function	Your Action
No	Specifies if you want to add, delete, copy, or reorder rules.	Right-click the table cell for the rule number and make your required modifications.
Match > From Zone	Specifies the zone from where the source sends traffic.	<p>Select one or more zones for the source zone, or you can specify any for all source zones.</p> <p>NOTE: The selected zone must be available on the security device specified in the Install On column.</p>
Match > Source	Specifies the address object that is the source of the traffic.	<p>Select any to monitor network traffic originating from any IP address.</p> <p>NOTE: You can also negate one or more address objects to specify all sources except the excluded object.</p>

Table 29: Exempt Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Match > To Zone	Specifies the destination zone.	Select the destination zone.
Match > Destination	Specifies the address object that is the destination of the traffic, typically a server or other device on your network.	Select the destination object. NOTE: You can also negate one or more address objects to specify all destinations except the excluded object.
Attacks	Specifies the attack(s) you want the IDP to exempt for the specified source or destination addresses.	Select the attack objects or groups. NOTE: You must include at least one attack object in an exempt rule.
VLAN Tag	Specifies that you can configure a rule to only apply to messages in certain VLANs.	Set a value by selecting any of the following options: <ul style="list-style-type: none"> • Any—This rule is applied to messages in any VLAN and to messages without a VLAN tag. • None—This rule is applied only to messages that do not have a VLAN tag. • Select VLAN Tags—This rule specifies which VLAN tags the rule applies to.
Install On	Specifies the security devices or templates that receive and use this rule.	Select the target security device. NOTE: You can also select multiple security devices on which to install the rule.
Comments	Specifies any miscellaneous comment about the rule's purpose.	Enter any additional comments about the rule.

For more information, see the *IDP Concepts & Examples guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Creating a New Security Policy \(NSM Procedure\) on page 34](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)
- [Configuring Backdoor Rulebase Rules \(NSM Procedure\) on page 47](#)

Configuring Backdoor Rulebase Rules (NSM Procedure)

The backdoor rulebase detects if there exists any interactive traffic introduced during backdoor attacks.

To configure a backdoor rulebase rule:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select and double-click the security policy to which you want to add the backdoor rulebase rule.
3. Click **New** in the upper right corner of the policy viewer and select **Add Backdoor Rulebase**.
4. Click the **New** button within the rules viewer to add a rule.
5. Modify the property of the rule by right-clicking the table cell for the property and making your modifications.
6. Configure or modify the rule using the settings described in [Table 30 on page 47](#).

Table 30: Backdoor Rulebase Rule Properties

Option	Function	Your Action
No	Specifies if you want to add, delete, copy, or reorder rules.	Right-click the table cell for the rule number and make your required modifications.
Match > Source	Specifies the address object that is the source of the traffic.	Select any to monitor network traffic originating from any IP address. NOTE: You can also negate one or more address objects to specify all sources except the excluded object.
Match > Destination	Specifies the address object that is the destination of the traffic, typically a server or other device on your network.	Select the destination object. NOTE: You can also negate one or more address objects to specify all destinations except the excluded object.
Match > Service	Specifies service objects in rules to service an attack to access your network.	Set a service by selecting any of the following options: <ul style="list-style-type: none"> • Any—Sets any service. • Default—Accepts the service specified by the attack object. • Select Service—Chooses specific services from the list of defined service objects.

Table 30: Backdoor Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Operation	Specifies whether to detect or ignore the backdoor traffic.	Select either Detect or Ignore .
Action	Specifies an action of the IDP to detect any interactive traffic.	Select any type of action.
Notification	Allows you to create log records with attack information that you can view real-time in the Log Viewer.	Select Configure to create log records.
VLAN Tag	Specifies that you can configure a rule to only apply to messages in certain VLANs.	Set a value by selecting any of the following options: <ul style="list-style-type: none"> • Any—This rule is applied to messages in any VLAN and to messages without a VLAN tag. • None—This rule is applied only to messages that do not have a VLAN tag. • Select VLAN Tags—This rule specifies which VLAN tags the rule applies to.
Severity	Specifies if you can override the inherent attack severity on a per-rule basis within the IDP rulebase.	Set the severity to Default , Info , Warning , Minor , Major , or Critical . <i>NOTE:</i> This column only appears when you view the Security Policy in Expanded Mode.
Install On	Specifies the security devices or templates that receive and use this rule.	Select the target security device. <i>NOTE:</i> You can also select multiple security devices on which to install the rule.
Comments	Specifies any miscellaneous comment about the rule's purpose.	Enter any additional comments about the rule.

For more information, see the *IDP Concepts & Examples guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)
- [Configuring SYN Protector Rulebase Rules \(NSM Procedure\) on page 49](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)

Configuring SYN Protector Rulebase Rules (NSM Procedure)

The SYN protector rulebase protects your network from malicious SYN-flood attacks.

To configure a SYN protector rulebase rule:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select and double-click the security policy to which you want to add the SYN protector rulebase rule.
3. Click **New** in the upper right corner of the policy viewer and select **Add SYN Protector Rulebase**.
4. Click the **New** button within the rules viewer to add a rule.
5. Modify the property of the rule by right-clicking the table cell for the property and making your modifications.
6. Configure or modify the rule using the settings described in [Table 31 on page 49](#).

Table 31: SYN Protector Rulebase Rule Properties

Option	Function	Your Action
No	Specifies if you want to add, delete, copy, or reorder rules.	Right-click the table cell for the rule number and make your required modifications.
Match > Source	Specifies the address object that is the source of the traffic.	Select any to monitor network traffic originating from any IP address. NOTE: You can also negate one or more address objects to specify all sources except the excluded object.
Match > Destination	Specifies the address object that is the destination of the traffic, typically a server or other device on your network.	Select the destination object. NOTE: You can also negate one or more address objects to specify all destinations except the excluded object.
Match > Service	Specifies service objects in rules to service an attack to access your network.	Set a service by selecting any of the available options. NOTE: We recommend that you do not change the default value, TCP-ANY .

Table 31: SYN Protector Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Mode	Specifies the mode that indicates how IDP handles TCP traffic.	<p>Select any of the following options:</p> <ul style="list-style-type: none"> • None—Specifies that IDP takes no action and does not participate in the three-way handshake. • Relay—Specifies that IDP acts as the middleman or relay, for the connection establishment, performing the three-way handshake with the client host on behalf of the server. <p>NOTE: Relay mode might not work as expected for MPLS traffic. When the IDP engine processes MPLS traffic, it stores the MPLS label information for traffic in each direction. In the case of traffic that matches SYN Protector rules in relay mode, the IDP appliance is programmed to send a SYN-ACK before the traffic has reached the server. In these cases, the IDP engine does not have server-to-client MPLS label information. Therefore, the SYN-ACK packet does not include an MPLS label. Some MPLS routers can add packets without a label to an existing MPLS tunnel; others drop such packets.</p> <ul style="list-style-type: none"> • Passive—Specifies that IDP handles the transfer of packets between the client host and the server, but does not actively prevent the connection from being established.
Notification	<p>Allows you to create log records with attack information that you can view real-time in the Log Viewer.</p> <p>NOTE: For more critical attacks, you can also set an alert flag to appear in the log record.</p>	<p>Select Configure to create log records.</p> <p>NOTE: The Configure menu option does not appear if the Mode column is set to None.</p> <ul style="list-style-type: none"> • Select Logging to have a log record created each time the rule is matched. • Select Alert to have an alert flag placed in the Alert column of the Log Viewer for the matching log record. • In the Log Actions tab, select desired log actions, if any.

Table 31: SYN Protector Rulebase Rule Properties (*continued*)

Option	Function	Your Action
VLAN Tag	Specifies that you can configure a rule to only apply to messages in certain VLANs.	Set a value by selecting any of the following options: <ul style="list-style-type: none"> • Any—This rule is applied to messages in any VLAN and to messages without a VLAN tag. • None—This rule is applied only to messages that do not have a VLAN tag. • Select VLAN Tags—This rule specifies which VLAN tags the rule applies to.
Severity	Specifies if you can override the inherent attack severity on a per-rule basis within the IDP rulebase.	Set the severity to Default, Info, Warning, Minor, Major, or Critical . <i>NOTE:</i> This column only appears when you view the Security Policy in Expanded Mode.
Install On	Specifies the security devices or templates that receive and use this rule.	Select the target security device. <i>NOTE:</i> You can also select multiple security devices on which to install the rule.
Comments	Specifies any miscellaneous comment about the rule's purpose.	Enter any additional comments about the rule.

For more information, see the *IDP Concepts & Examples guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)
- [Configuring Traffic Anomalies Rulebase Rules \(NSM Procedure\) on page 51](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)

Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)

The traffic anomalies rulebase employs a traffic flow analysis method to detect attacks that occur over multiple connections and sessions (such as scans).

To configure a traffic anomalies rulebase rule:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select and double-click the security policy to which you want to add the traffic anomalies rulebase rule.

3. Click **New** in the upper right corner of the policy viewer and select **Add Traffic Anomalies Rulebase**.
4. Click the **New** button within the rules viewer to add a rule.
5. Modify the property of the rule by right-clicking the table cell for the property and making your modifications.
6. Configure or modify the rule using the settings described in [Table 32 on page 52](#).

Table 32: Traffic Anomalies Rulebase Rule Properties

Option	Function	Your Action
No	Specifies if you want to add, delete, copy, or reorder rules.	Right-click the table cell for the rule number and make your required modifications.
Match > Source	Specifies the address object that is the source of the traffic.	<p>Select any to monitor network traffic originating from any IP address.</p> <p>NOTE: You can also negate one or more address objects to specify all sources except the excluded object.</p>
Match > Destination	Specifies the address object that is the destination of the traffic, typically a server or other device on your network.	<p>Select the destination object.</p> <p>NOTE: You can also negate one or more address objects to specify all destinations except the excluded object.</p>
Match > Service	Specifies service objects in rules to service an attack to access your network.	<p>Set a service by selecting any of the available options.</p> <p>NOTE: We recommend that you do not change the default value, TCP-ANY.</p>
Traffic Anomaly	Specifies how IDP is to treat the matching traffic.	<p>Select any of the following options:</p> <ul style="list-style-type: none"> • Ignore—IDP ignores this traffic. This option excludes traffic from trusted sources that might be falsely construed as a scan. • Detect—IDP matches this traffic and takes the IP action that you have set. <p>When you select this option, the Traffic Anomalies dialog box appears. Select the scans or sweep you want to detect and enter values for Port Count and Time Threshold (in seconds) or Session Count.</p>

Table 32: Traffic Anomalies Rulebase Rule Properties (*continued*)

Option	Function	Your Action
IP Action	Allows you to log, drop, or close the current connection for each attack that matches a rule.	<p>Select Configure to do any one of the following actions:</p> <ul style="list-style-type: none"> • Enabled—Enables IP actions. • Action—Specifies the action you want the IDP to take. • Block—Specifies which parameters IDP will use to close or block further connections from the drop down list. • Logging—Specifies the log action for a matching event. • Timeout (sec)—Specifies the number of seconds that this action remains in effect on IDP after a traffic match.
Notification	<p>Allows you to create log records with attack information that you can view real-time in the Log Viewer.</p> <p>NOTE: For more critical attacks, you can also set an alert flag to appear in the log record.</p>	<p>Select Configure to create log records.</p> <p>NOTE: The Configure menu option does not appear if the Mode column is set to None.</p> <ul style="list-style-type: none"> • Select Logging to have a log record created each time the rule is matched. • Select Alert to have an alert flag placed in the Alert column of the Log Viewer for the matching log record. • In the Log Actions tab, select desired log actions, if any.
VLAN Tag	Specifies that you can configure a rule to only apply to messages in certain VLANs.	<p>Set a value by selecting any of the following options:</p> <ul style="list-style-type: none"> • Any—This rule is applied to messages in any VLAN and to messages without a VLAN tag. • None—This rule is applied only to messages that do not have a VLAN tag. • Select VLAN Tags—Specifies which VLAN tags the rule applies to.
Severity	Specifies if you can override the inherent attack severity on a per-rule basis within the IDP rulebase.	<p>Set the severity to Default, Info, Warning, Minor, or Critical.</p> <p>NOTE: This column only appears when you view the Security Policy in Expanded Mode.</p>

Table 32: Traffic Anomalies Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Install On	Specifies the security devices or templates that receive and use this rule.	Select the target security device. NOTE: You can also select multiple security devices on which to install the rule.
Comments	Specifies any miscellaneous comment about the rule's purpose.	Enter any additional comments about the rule.

For more information, see the *IDP Concepts & Examples* guide.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview](#) on page 31
- [Modifying IDP Rulebase Rules \(NSM Procedure\)](#) on page 36
- [Configuring Network Honeypot Rulebase Rules \(NSM Procedure\)](#) on page 54
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\)](#) on page 129

Configuring Network Honeypot Rulebase Rules (NSM Procedure)

The network honeypot rulebase is a method to detect investigation activities.

To configure a network honeypot rulebase rule:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select and double-click the security policy to which you want to add the network honeypot rulebase rule.
3. Click **New** in the upper right corner of the policy viewer and select **Add Network Honeypot Rulebase**.
4. Click the **New** button within the rules viewer to add a rule.
5. Modify the property of the rule by right-clicking the table cell for the property and making your modifications.
6. Configure or modify the rule using the settings described in [Table 33](#) on page 54.

Table 33: Network Honeypot Rulebase Rule Properties

Option	Function	Your Action
No	Specifies if you want to add, delete, copy, or reorder rules.	Right-click the table cell for the rule number and make your required modifications.
Source Address	Specifies the address object that is the source of the traffic.	Select any source address or group.

Table 33: Network Honeypot Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Impersonate > Destination	Specifies the address object that is the destination of the traffic, typically a server or other device on your network.	<p>Select the destination object.</p> <p>NOTE: You can also negate one or more address objects to specify all destinations except the excluded object.</p>
Impersonate > Service	Specifies the services running on your network.	Select the services you want to monitor.
Operation	Specifies whether or not IDP fakes open ports.	<p>Select any of the following options:</p> <ul style="list-style-type: none"> • Ignore—This option allows free passage on your network when creating rules for trusted traffic. • Impersonate—IDP creates a fake port open to the public based on the destination IP addresses and service you selected.
IP Action	Allows you to log, drop, or close the current connection for each attack that matches a rule.	<p>Select Configure to do any one of the following actions:</p> <ul style="list-style-type: none"> • Enabled—Enable IP actions. • Action—Specifies the action you want the IDP to take. • Block—Specifies which parameters IDP will use to close or block further connections from the drop-down list. • Logging—Specifies the log action for a matching event. • Timeout (sec)—Specifies the number of seconds that this action remains in effect on IDP after a traffic match.
Notification	<p>Allows you to create log records with attack information that you can view real-time in the Log Viewer.</p> <p>NOTE: For more critical attacks, you can also set an alert flag to appear in the log record.</p>	<p>Select Configure to create log records.</p> <p>NOTE: The Configure menu option does not appear if the Mode column is set to None.</p> <ul style="list-style-type: none"> • Select Logging to have a log record created each time the rule is matched. • Select Alert to have an alert flag placed in the Alert column of the Log Viewer for the matching log record. • In the Log Actions tab, select desired log actions, if any.

Table 33: Network Honeypot Rulebase Rule Properties (*continued*)

Option	Function	Your Action
VLAN Tag	Specifies that you can configure a rule to only apply to messages in certain VLANs.	Set a value by selecting any of the following options: <ul style="list-style-type: none"> • Any—This rule is applied to messages in any VLAN and to messages without a VLAN tag. • None—This rule is applied only to messages that do not have a VLAN tag. • Select VLAN Tags—This rule specifies which VLAN tags the rule applies to.
Severity	Specifies if you can override the inherent attack severity on a per-rule basis within the IDP rulebase.	Set the severity to Default , Info , Warning , Minor , Major , or Critical . <i>NOTE:</i> This column only appears when you view the Security Policy in Expanded Mode.
Install On	Specifies the security devices or templates that receive and use this rule.	Select the target security device. <i>NOTE:</i> You can also select multiple security devices on which to install the rule.
Comments	Specifies any miscellaneous comment about the rule's purpose.	Enter any additional comments about the rule.



NOTE: The IDP drops MPLS traffic that matches a Network Honeypot rule. When the IDP engine processes MPLS traffic, it stores the MPLS label information. It stores separate labels for client-to-server and server-to-client communication. In the case of traffic that matches Network Honeypot rules, there is no genuine server-to-client communication, so the IDP engine does not have server-to-client MPLS label information. Therefore, the impersonation operation is not supported.

For more information, see the *IDP Concepts & Examples guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)
- [Validating a Security Policy \(NSM Procedure\) on page 130](#)

Configuring Application Rulebase Rules (NSM Procedure)

The Application Policy Enforcement (APE) rulebase enables you to limit bandwidth for specified users and/or applications. You can configure APE rules to detect network traffic based on application signatures. The user can define custom application signatures to be used in the APE rules. The APE rulebase enables actions based on an application-centered matching tuple.

To configure an APE rulebase rule:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. Select and double-click the security policy to which you want to add the APE rulebase rule.
3. Click New in the upper right corner of the policy viewer and select **Add Application Rulebase**.
4. Click the **New** button within the rules viewer to add a rule.
5. Modify the property of the rule by right-clicking the table cell for the property and making your modifications.
6. Configure or modify the rule using the settings described in [Table 34 on page 57](#).
7. Click **OK** to save your changes.

Table 34: APE Rulebase Rule Properties

Option	Function	Your Action
No.	Specifies if you want to add, delete, copy, or reorder rules.	Right-click the table cell for the rule number and make your required modifications.
Match > Source	Specifies the address object that is the source of the traffic.	Select any to monitor network traffic originating from any IP address. <i>NOTE:</i> For guidelines on specifying match parameters, see the <i>IDP Concepts and Examples Guide</i> .
Match > User Role	Specifies the user roles to match the session for the rule to be applied. If a value for User Role matches, the Source parameter is not consulted. Matching based on user role depends on integration with a compatible Juniper Networks IC Series Unified Access Control appliance.	Right-click the table cell to select user roles.
Match > Destination	Specifies the address object that is the destination of the traffic, typically a server or other device on your network.	Select the destination object. <i>NOTE:</i> You can also negate one or more address objects to specify all destinations except the excluded object.

Table 34: APE Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Match > Service	Requires a match of one of the specified services. A single rule can match a service object definition or an application list, but not both. We recommend you create rules that match an application list whenever possible. Matching based on application uses the application identification feature, which can identify the application regardless of port. We support rules that match service object definitions for cases where there is not a suitable application object.	<p>Right-click the table cell and select any one of the required options.</p> <p>If you specify named values for both service and application, only the application value is used.</p> <p>If your rule includes application or extended application objects, specify Default for the service parameter.</p> <p>If you do not want to match on service or application list, specify Any for all three (service, application, and extended application).</p> <p>If there are no suitable application objects, create a rule that uses the service object and set the application and extended application columns to Any.</p> <p>NOTE: If the service uses standard ports, you can select from predefined services. If the service uses nonstandard ports, you can create a custom service object. The IDP engine can inspect services that use TCP, UDP, RPC, and ICMP transport layer protocols.</p>

Table 34: APE Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Match > Application	Requires one of the specified applications to match the session for the rule to be applied. The predefined list of applications is populated by the application identification feature. The application identification feature identifies the application regardless of port. Port-independent application identification simplifies rule configuration and ensures that you do not miss applications running on nonstandard ports. Hence it is recommended to use the application parameter instead of the service parameter whenever possible.	<p>Right-click the table cell and make your required modifications.</p> <p>If you specify named values for both service and application, only the application value is used.</p> <p>Specify Any when creating a service-based rule or when creating an application-based rule where the application list consists only of extended application objects.</p> <p>You can use the Shared Objects for Policy viewer (located below the rule editor) to browse application objects and explore object properties. You can create custom application objects.</p> <p>NOTE: To apply an APE action to all traffic matching source and destination parameters, set both the service parameter and the application parameter to Any.</p> <p>Extended application matching is more granular than application matching. Do not select HTTP in the application column if you also plan to specify extended application objects in the same rule. If you specify HTTP and HTTP:Facebook, for example, the rule matches HTTP or HTTP:Facebook. The result is indistinguishable from a rule matching only HTTP. We recommend you list rules targeting Extended Applications before a rule targeting HTTP.</p>
Match > Extended Application	Requires one of the specified extended applications to match the session for the rule to be applied. Extended applications are also called nested applications. The Juniper Networks Security Center (J-Security Center) provides predefined application signatures for many Web 2.0 applications running over HTTP. Matching on these signatures depends on the application identification feature, which is enabled by default. You use the Application and Extended Application columns to build a list of applications to match the rule. The list is evaluated as a Boolean OR, so if one of the application or extended application objects specified in the rule is identified, the “service or application” component of the tuple matches.	<p>Right-click the table cell and make your required modifications.</p> <p>Specify Any when you are creating a service-based rule or when you are creating an application-based rule where the application list consists only of application objects.</p> <p>NOTE: You can use the Shared Objects for Policy viewer (located below the rule editor) to browse extended application objects and explore object properties. You cannot create custom extended application objects.</p>

Table 34: APE Rulebase Rule Properties (*continued*)

Option	Function	Your Action
Action	Specifies which actions to perform against attacks that match rules in your security policy.	

Table 34: APE Rulebase Rule Properties (*continued*)

Option	Function	Your Action
		<p>Right-click the table cell and select any one of the following options:</p> <ul style="list-style-type: none"> • None — IDP takes no action against the connection. • Drop Packet — IDP drops a matching packet before it can reach its destination but does not close the connection. • Drop Connection — IDP drops the connection without sending an RST packet to the sender, preventing the traffic from reaching its destination. • Close Client — IDP closes the connection to the client and not to the server. • Close Server — IDP closes the connection to the server and not to the client. • Close Client and Server — IDP closes the connection and sends a RST packet to both the client and the server. • Diffserv Marking — Assigns the differentiated service value you specify to the packet. This action is useful when your network has a class of service (CoS) design, and you want to use the IDP Series device to rewrite the CoS code point based on APE rules processing. The CoS rules you have implemented for the next devices in the network path ultimately determine the effect on the transmission rate. <p>NOTE: In sniffer mode, this action has no effect because the IDP Series device is not in the path of network traffic.</p> <ul style="list-style-type: none"> • Rate Limiting — IDP enforces a rate limit for all current sessions that match the rule (separate limits for client-to-server and server-to-client traffic). If the limit has not been reached, IDP forwards the packets. If the limit has been reached, IDP behaves as if no bandwidth is available. The rate limits that are best suited for your business case depend on the bandwidth for your links. If you have a 1-Gbps link and want no more than 10% available to peer-to-peer traffic, the sum of the rate limits you specify for all peer-to-peer rules must be less than 102.4 Mbps (in each direction). <p>You configure separate rate limits for client-to-server and server-to-client directions. For peer-to-peer traffic, we recommend that you set the same rate for each direction.</p> <p>NOTE: For TFTP traffic, all traffic is considered client-to-server traffic. A TFTP server responds</p>

Table 34: APE Rulebase Rule Properties (*continued*)

Option	Function	Your Action
		<p>to get requests by establishing an ephemeral port from which to send the reply. In this case, both directions appear to the IDP Series device as client-to-server flows. We recommend you set the same rate for each direction. In sniffer mode, this action has no effect because the IDP Series device is not in the path of network traffic</p> <ul style="list-style-type: none"> • Diffserv Marking & Rate Limiting — Takes both actions as described for Diffserv Marking and Rate Limiting.
Notification	Specifies logging options. Packet capture is not applicable for APE rulebase rules.	Right-click the table cell and select Configure to display a dialog box where you can configure logging options.
VLAN Tag	Specifies rules to traffic on certain VLANs. Normally, for a rule to take effect, it must match the packet source, destination, service, and attack objects. If the VLAN cell is populated with a value other than any, then the rule will also consider the packet's VLAN tag when determining a match.	Right-click the table cell to assign a VLAN object to a rule or to set the VLAN tag value to none.
Install On	Specifies target IDP devices for the rule. By default, IDP security policy rules can be applied to any IDP device.	Right-click the table cell and select Select Target to display a dialog box to specify the IDP devices to which the rule can be installed.
Comments	Adds notations about the rule. This setting is optional and does not affect the functionality of the security policy rule.	Right-click the table cell and select Edit Comments to display a dialog box where you can make notations about the rule.

You can verify the APE rulebase functionality in your lab and view APE related statistics in the Command-Line Interface (CLI). It is recommended that you retain defaults for APE rulebase. By default:

- IDP does not limit the rate of sessions that do not match APE rules. Rate limiting is done by service based till application is identified in the session i.e. default services running on the port.
- When the application identification feature fails to identify the application, IDP does not try to match the rule but instead applies the default rate limit (if any). You can modify this so that in cases where application identification fails, IDP attempts to match the session to the standard protocol and port for the application.

For more information, see the *IDP Concepts & Examples guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)

CHAPTER 5

Working with Attack Objects

- [Attack Objects in Intrusion Detection and Prevention Security Policies Overview on page 63](#)
- [Loading J-Security-Center Updates \(NSM Procedure\) on page 64](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\) on page 65](#)
- [Working with Attack Groups \(NSM Procedure\) on page 66](#)
- [Creating a Signature Attack Object \(NSM Procedure\) on page 69](#)
- [Creating a Compound Attack Object \(NSM Procedure\) on page 81](#)
- [Verifying the Attack Object Database Version \(NSM Procedure\) on page 84](#)
- [Updating the IDP Detector Engine \(NSM Procedure\) on page 85](#)

Attack Objects in Intrusion Detection and Prevention Security Policies Overview

You use the NSM Object Manager to manage NSM administrative objects, including the attack objects used in IDP security policies.

For more explanation of attack objects and examples, see the *Network and Security Manager Administration Guide*.

For details on how to use NSM Object Manager features to copy objects, find references to objects, search for unused objects, or configure object versions, see the *NSM online Help*.

IDP administration using attack objects can include the following tasks related to attack objects:

- Updating IDP detector engine and the NSM attack database
- Viewing predefined attack object definitions
- Viewing attack object groups
- Updating predefined IDP attack objects and groups

Related Documentation

- [Working with Attack Groups \(NSM Procedure\) on page 66](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\)](#)
- [Loading J-Security-Center Updates \(NSM Procedure\) on page 64](#)

Loading J-Security-Center Updates (NSM Procedure)

The Juniper Networks Security Center (J-Security Center) routinely makes important updates available to IDP security policy components, including updates to the IDP detector engine and NSM attack database.

The IDP detector engine is a dynamic protocol decoder that includes support for decoding more than 60 protocols and more than 500 service contexts. You should update IDP detector engine when you first install the IDP device, whenever you upgrade, and whenever alerted to do so by Juniper Networks.

The NSM attack database stores data definitions for the attack objects that are key components of IDP security policies. Updates can include new attack objects, revised severity settings, or removed attack objects. You should schedule daily updates to the NSM attack database.

After you have completed the update, any new attack objects are available in the security policy editor. If you use dynamic groups to your IDP rulebase rules and a new attack object belongs to the dynamic group, the rule automatically inherits the new attacks.

[Table 35 on page 64](#) provides procedures for updating IDP detector engine and the NSM attack database.

Table 35: IDP Detector Engine and NSM Attack Database Update Procedures

Task	Procedure
To download IDP detector engine and NSM attack database updates to the NSM GUI server	<p>From the NSM main menu, select Tools > View/Update NSM attack database and complete the wizard steps.</p> <p>NOTE: The default URL from which to obtain updates is https://services.netscreen.com/restricted/sigupdates/nsm-updates/NSM-SecurityUpdateInfo.dat. If you encounter connection errors, ensure this setting has not been inadvertently changed.</p> <ol style="list-style-type: none"> 1. From the NSM main menu, select Tools > Preferences. 2. Click Attack Object. 3. Click Restore Defaults. NSM restores the URL in the Download URL for ScreenOS Devices text box. 4. Click OK.
To push an IDP detector engine update from the NSM GUI server to IDP devices	<p>From the NSM main menu, select Devices > IDP Detector Engine > Load IDP Detector Engine for ScreenOS and complete the wizard steps.</p> <p>NOTE: Updating the IDP detector engine on a device does not require a reboot of the device.</p>

Table 35: IDP Detector Engine and NSM Attack Database Update Procedures (*continued*)

Task	Procedure
To push predefined attack object updates from the NSM GUI server to IDP devices	<ol style="list-style-type: none"> 1. From the NSM main menu, select Devices > Configuration > Update Device Config. 2. Select the devices that you want to push configuration updates to and to set update job options on. 3. Click OK. <p>NOTE: Only the attack objects that are used in IDP rules for the device are pushed from the GUI server to the device.</p>
To schedule regular updates	<ol style="list-style-type: none"> 1. Log in to the NSM GUI server command line. 2. Change directory to <code>/usr/netscreen/GuiSvr/utils</code>. 3. Create a shell script called <code>attackupdates.sh</code> with the following contents: <ul style="list-style-type: none"> • Set the <code>NSMUSER</code> environment variable with an NSM domain/user pair. The command for setting environment variables depends on your OS. Example: <pre>export NSMUSER=domain/user</pre> • Set the <code>NSMPASSWD</code> environment variable with an NSM password. The command for setting environment variables depends on your OS and shell. Example: <pre>export NSMPASSWD=password</pre> • Specify a <code>guiSvrCli</code> command string. Example: <pre>/usr/netscreen/GuiSvr/utils/guiSvrCli.sh --update-attacks --post-action --update-devices --skip</pre> 4. Make the script executable by the user associated with the cron job: <pre>chmod 700 attackupdates.sh</pre> 5. Run the crontab editor: <pre>crontab -e</pre> 6. Add an entry for the shell script: <pre>minutes_after_hour hour * * * /usr/netscreen/GuiSvr/utils/attackupdates.sh</pre> <p>During the update, the <code>guiSvrCli</code> utility updates the attack object database, then performs the post actions. After updating and executing actions, the system generates an exit status code of 0 (no errors) or 1 (errors).</p>

Related Documentation

- [Attack Objects in Intrusion Detection and Prevention Security Policies Overview on page 63](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\) on page 65](#)
- [Working with Attack Groups \(NSM Procedure\) on page 66](#)

Viewing Predefined Attack Objects (NSM Procedure)

Purpose Juniper Networks Security Center (J-Security Center) develops predefined attack objects and attack object groups for IDP rulebase rules.

In most cases, the predefined attack objects are the only attack objects you need to protect your network.

The predefined attack object list in the NSM Object Manager provides the following summary of each attack object:

- Name of the attack object
- Severity of the attack: critical, major, minor, warning, info
- Category
- Keywords
- Common Vulnerabilities and Exposures database (CVE) number
- Security Focus Bugtraq database number

Action To view predefined attack object details:

1. In the Object Manager, click **Attack Objects > IDP Objects** to display the IDP Objects dialog box.
2. Click either the **Predefined Attacks** or **Predefined Attack Groups** tab to view the predefined attack object list.
3. Double-click the table row entry for the attack object to display its details.



NOTE: You cannot create, edit, or delete predefined attack objects.

**Related
Documentation**

- [Attack Objects in Intrusion Detection and Prevention Security Policies Overview on page 63](#)
- [Working with Attack Groups \(NSM Procedure\) on page 66](#)
- [Loading J-Security-Center Updates \(NSM Procedure\) on page 64](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\)](#)

Working with Attack Groups (NSM Procedure)

NSM groups are administrative objects that facilitate configuration and monitoring tasks. You can add attack groups or individual attack objects to IDP rulebase rules and Exempt rulebase rules.

- [Creating Dynamic Groups on page 67](#)
- [Creating Static Groups on page 68](#)

Creating Dynamic Groups

A dynamic group contains attack objects that are automatically added or deleted based on specified criteria for the group. The NSM Object Manager includes predefined dynamic groups that work with recommended attack objects, predefined attack objects, the recommended security policy, and predefined policy templates.

When you run an NSM attack database update job, the process automatically performs the following tasks:

- For all new attack objects, compares the predefined attributes of each attack object to each dynamic group criteria and adds the attack objects that match.
- For all updated attack objects, removes attack objects that no longer meet their dynamic group criteria.
- Reviews updated attack objects to determine if they now meet any other dynamic group criteria, and adds them to those groups if necessary.
- For all deleted attack objects, removes the attack objects from their dynamic groups.

Use of dynamic groups eliminates the need to review each new signature to determine if you need to use it in your existing security policy.

A predefined or custom dynamic group can contain only attack objects and not attack groups. Dynamic group members can be either predefined or custom attack objects.

To create a custom dynamic group:

1. In Object Manager, select **Attack Objects > IDP Objects** to display the IDP Objects dialog box.
2. Click **Custom Attack Groups** tab, then click the + icon and select **Add Dynamic Group** to display the New Dynamic Group dialog box.
3. Enter a name and description for the static group. Select a color for the group icon.
4. In the Filters tab, click + icon and add select filters that determine which attack objects should be in the group using [Table 36 on page 67](#).
5. Click **Members** tab to view the attack objects now belonging to the group.
6. Click **OK** to save your settings.

Table 36: Dynamic Attack Group Filters

Filter	Description
Add Products Filter	Filters attack objects based on the application that is vulnerable to the attack.
Add Service Filter	Filter to add attack objects based on the attack service, such as FTP, HTTP, NetBios, and so on.
Add Direction Filter	Filter to add predefined attacks to the dynamic group based on the direction specified in the attacks.

Table 36: Dynamic Attack Group Filters (*continued*)

Filter	Description
Add Severity Filter	Filters attack objects based on attack severity. NOTE: All predefined attack objects are assigned a severity level by Juniper Networks. However, you can edit this setting to match the needs of your network.
Add Category Filter	Filters attack objects based on category.
Add Last Modified Filter	Filters attack objects based on their last modification date.
Add Attack Type Filter	Filter to add attack objects based on the type of attack object (signature or protocol anomaly).
Add False Positives Filter	Filter to add attack objects based on the frequency that the attack produces a false positive on your network.
Add Recommended Filter	Filters attack objects based on whether they have been marked Recommended.
Add Performance (Detection Complexity) Filter	Filter to add attack objects based on the performance level that is vulnerable to the attack. NOTE: Detection Complexity filter only supports IDP standalone device.

Creating Static Groups

A static group contains a specific, finite set of attack objects or groups. There are two types of static groups: predefined static groups and custom static groups.

A custom static group can include the same members as a predefined static group (predefined attack objects, predefined static groups, and predefined dynamic groups), plus the following members:

- Custom attack objects
- Custom dynamic groups
- Other custom static groups

Use static groups to define a specific set of attacks to which you know your network is vulnerable, or to group custom attack objects. For example, you might want to create a group for a specific set of informational attack objects that keep you aware of what is happening on your network.

Static groups require more maintenance than dynamic groups because you must manually add or remove attack objects in a static group to change the members. However, you can include a dynamic group within a static group to automatically update some attack objects. For example, the predefined attack object group Operating System is a static group that contains four predefined static groups: BSD, Linux, Solaris, and Windows. The BSD group contains the predefined dynamic group BSD-Services-Critical, to which attack objects can be added during an attack database update.

To create a custom static group:

1. In Object Manager, select **Attack Objects > IDP Objects** to display the IDP Objects dialog box.
2. Click the **Custom Attack Groups** tab, then click the + icon and select **Add Static Group** to display the New Static Group dialog box.
3. Enter a name and description for the static group.
4. Select a color for the group icon.
5. Select the attack or group from the Attacks/Group list and click **Add**.
6. Click **OK**.

Related Documentation

- [Attack Objects in Intrusion Detection and Prevention Security Policies Overview on page 63](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\)](#)
- [Verifying the Attack Object Database Version \(NSM Procedure\) on page 84](#)

Creating a Signature Attack Object (NSM Procedure)

A signature attack object is a pattern you want the system to detect. You use a DFA expression to represent the pattern. All of the other signature properties you can set (such as service or protocol context, direction, and other constraints) are provided so you can optimize performance of the system in detecting the pattern and eliminate false positives. In general, you want to tune settings of a signature attack object so that the system looks for it in every context where it might occur and in no other context.

To configure a signature attack object:

1. In the Object Manager, select **Attack Objects > IDP Objects**.
2. Click the **Custom Attacks** tab.
3. Click the + icon to display the Custom Attack dialog box.
4. Configure attack object settings. [Table 37 on page 69](#) provides guidelines for completing the settings.

Table 37: Custom Attack Dialog Box: General Tab Settings

Setting	Description
Name	The name displayed in the UI. TIP: Include the protocol the attack uses as part of the attack name.
Description	(Optional) Information about the attack. Although a description is optional when you create a new attack object, it can help you remember important information about the attack. For examples, view the attack descriptions for predefined attacks.

Table 37: Custom Attack Dialog Box: General Tab Settings (*continued*)

Setting	Description
Severity	Info, Warning, Minor, Major, or Critical. Critical attacks are attempts to crash your server or gain control of your network. Informational attacks are the least dangerous and typically are used by network administrators to discover holes in their own security system.
Category	A predefined or new category.
Keywords	Unique identifiers that can be used to search and sort log records.
Recommended	Indicates that this attack object is among your highest-risk set of attack objects. Later, when you add this attack object to dynamic groups, you can specify whether to include only recommended attack objects.
Attack Versions	Skip this for now.
Detection Performance	Select High , Medium , Low , or Not Defined .

- Configure additional attack details on the Extended tab. [Table 38 on page 70](#) provides guidelines for completing the settings.

Table 38: Custom Attack Dialog Box: Extended Tab Settings

Setting	Description
Primary URL	Up to three URLs (primary, secondary, tertiary) to external references you used to research the attack.
Secondary URL	
Tertiary URL	
CVE	The Common Vulnerabilities and Exposures (CVE) ID that the attack object addresses. CVE is a standardized list of vulnerabilities and other information security exposures. The CVE number is an alphanumeric code, such as CVE-2209.
BugTraq	The BugTraq ID number that the attack object addresses. BugTraq is a moderated mailing list that discusses and announces computer security vulnerabilities. The BugTraq ID number is a three-digit code, such as 831 or 120.
Impact	Information about the impact of a successful attack, including information about system crashes and access granted to the attacker.
Description	Additional information.
Tech Info	Information about the vulnerability, the commands used to execute the attack, which files are attacked, registry edits, and other low-level information.
Patches	Any patches available from the product vendor, as well as information about how to prevent the attack.

- Click the **General** tab.

7. Under Attack Versions, click the + icon to display the New Attack wizard.
8. On the Target Platform and Type page, select a device platform and attack type.
[Table 39 on page 71](#) describes the attack types.

Table 39: Attack Object Types

Type	Description
Signature	<p>Uses a stateful attack signature (a pattern that always exists within a specific section of the attack) to detect known attacks.</p> <p>Stateful signature attack objects also include the protocol or service used to perpetrate the attack and the context in which the attack occurs.</p> <p>If you know the exact attack signature, the protocol, and the attack context used for a known attack, select this option.</p>
Compound Attack	<p>Detects attacks that use multiple methods to exploit a vulnerability. This object combines multiple signatures or protocol anomalies into a single attack object, forcing traffic to match all combined signatures or anomalies within the compound attack object before traffic is identified as an attack.</p> <p>By combining and even specifying the order in which signatures or anomalies must match, you can be very specific about the events that must place before the IDP engine identifies traffic as an attack.</p> <p>If you need to detect an attack that uses several benign activities to attack your network, or if you want to enforce a specific sequence of events to occur before the attack is considered malicious, select this option.</p>

9. Select **Signature** and click **Next**.
10. On the Custom Attack – General Properties page, configure constraints and other settings. [Table 40 on page 71](#) provides guidelines for completing the settings.

Table 40: Custom Attack – General Properties

Property	Description
Info	
False Positives	<p>Select the frequency that the attack object produces a false positive on your network: Unknown, Rarely, Occasionally, Frequently.</p> <p>Typically, you do not initially know the frequency of false positives. You can update this setting as your observations change.</p>
Service Binding	

Table 40: Custom Attack – General Properties (*continued*)

Protocol Type	<p>Service—If you were able to determine the service through your research, select Service. Later in the wizard, you can specify a service context.</p>
	<p>IP—If you are not sure of the service but you know IP details, select IP and specify a protocol type number.</p>
	<p>TCP, UDP, or ICMP—If you do not know the service context but you know protocol details, select the protocol.</p>
	<p>For TCP and UDP protocol types, specify the port ranges.</p>
	<p>RPC—If you are detecting threats over remote procedure call (RPC) protocol, select this option and specify the program ID.</p>
	<p>RPC is used by distributed processing applications to handle interaction between processes remotely. When a client makes a remote procedure call to an RPC server, the server replies with a remote program. Each remote program uses a different program number.</p>
	<p>IPv6 or ICMPv6—Do not select these options. IDP Series devices do not support inspection of IPv6.</p>
	<p>Any—If you are unsure of the correct service, select Any to match the signature in all services. Matching any service essentially turns off service binding and has a significant performance impact. Specify Any when you know that attacks are using multiple services to attack your network.</p>
	<p>NOTE: You must select a service binding other than Any if you want to select a context for the attack.</p>

Table 40: Custom Attack – General Properties (*continued*)

Time Binding	
Enable	Time binding attributes track how many times a signature is repeated. By configuring the scope and count of an attack, you can detect a sequence of the same attacks over a period of time (one minute) across sessions. This method is useful for detecting brute force attacks that attempt to guess authentication credentials or overwhelm system capacity to handle data.
Scope	<p>Select the scope within which the count occurs:</p> <ul style="list-style-type: none"> • Source—Detects the signature in traffic from the source IP address for the specified number of times, regardless of the destination IP address. • Destination—Detects the signature in traffic from the destination IP address for the specified number of times, regardless of the source IP address. • Peer—Detects the signature in traffic between source and destination IP addresses of the sessions for the specified number of times.
Count/Min	<p>Enter the number of times per minute that the signature must be detected within the specified scope before the device identifies the traffic as a match.</p> <p>The minute timer starts when the signature first matches the event. If the signature matches the same event for the specific count or higher within the next 60 seconds, the traffic is a match.</p> <p>The system increments the count each time it detects the signature, either regardless of port (application identification) or according to your port binding settings. For example, when the system detects the signature on TCP/80 and then on TCP/8080, the count is 2.</p>

Table 40: Custom Attack – General Properties (*continued*)

Constraints	
Within Bytes Constraint	<p>Use this constraint to require that the pattern be found within a byte range:</p> <ul style="list-style-type: none"> • Lower limit—Specify the beginning of the range. • Upper limit—Specify the end of the range. • Start point—Your selection must be consistent with your pattern context setting. For example, if you configured one of the service contexts, select Context. If you configured one of the packet contexts, select Packet. If you configured one of the stream contexts, select Stream. <p>In NSM, it is possible to select a start point that is inconsistent with the pattern context setting. For example, the NSM object editor allows you to configure a pattern context http-variable and then set a within bytes start point that is start-of-packet. However, the within bytes match logic will be resolved to the start point you should have selected: context.</p> <p>Inspection for this object terminates when the range limit is reached.</p> <p>Example: If you know a threat can be identified either completely within the first 20 bytes of the http-variable context or not identified at all, you set the context to http-variable and use the within-bytes constraint to terminate inspection after bytes 1-20 of the generated http-variable context are processed.</p> <p>You can set multiple constraints. The constraints are evaluated as a Boolean OR.</p> <p>Example: You configure two start-of-stream constraints with byte ranges of 20-40 and 80-100. The constraint rules out matches unless found within either byte range.</p>
Within Packets Constraint	<p>Use this constraint to require that the pattern be found completely within a packet range:</p> <ul style="list-style-type: none"> • Lower limit—Specify the beginning of the range. • Upper limit—Specify the end of the range. <p>Inspection (for this object) terminates when the range limit is reached.</p> <p>Example: If you know a threat can be identified either in the first 2 packets or not identified at all, you set a stream context and use the within packets constraint to terminate inspection after 2 packets.</p>
Context Check	<p>Use this constraint to require the matching context be of a specified byte length to be a hit:</p> <ul style="list-style-type: none"> • Constraint—Select length. • Comparison operator—Select =, !, >, or <. • Operand—Select a byte length. <p>Example: You can use the context check constraint as a tuning device to limit processing for harmless traffic. For example, if you know that a certain class of attack, like a buffer overflow attack, always has an unusually large byte length in a given context, you can use this constraint to ignore contexts of normal length. If you set the FTP username context length requirement to be > 18, you would only see signature hits if the FTP username context is longer than 18 bytes.</p> <p>You can specify multiple constraints. For example, if you add a < 25 constraint to the previous example, you would only see hits if the username context is between 18 and 25 bytes.</p>

Click **Next**.

11. On the Custom Attack – Attack Pattern page, configure pattern settings. [Table 41 on page 75](#) provides guidelines for completing the settings.

Table 41: Custom Attack – Attack Pattern

Setting	Description
Pattern	A DFA expression. The following rows summarize DFA syntax conventions. For detailed information, consult a standard source on programming with regular expressions.

Table 41: Custom Attack – Attack Pattern (*continued*)

Setting	Description
\B.0.1..00\B	<p>Bit-level matching for binary protocols. The length of the bitmask must be in multiples of 8.</p> <p>The first \B denotes the start of the bitmask. The last \B denotes the end of the bitmask.</p> <p>The decimal (.) indicates the bit can be either 0 or 1.</p> <p>A 0 or 1 indicates the bit at that position must be 0, or must be 1.</p>
\0 <octal_number>	For a direct binary match.
\X<hexadecimal-number>\X	For a direct binary match.
\[<character-set>\]	For case-insensitive matches.
.	To match any symbol.
*	To match 0 or more symbols.
+	To match 1 or more symbols.
?	To match 0 or 1 symbol.
()	Grouping of expressions.
	<p>Alternation. Typically used with ().</p> <p>Example: The following expression matches dog or cat: (dog cat).</p>
[]	<p>Character class. Any explicit value within the bracket at the position matches.</p> <p>Example: [Dd]ay matches Day and day.</p>
[<start>--<end>]	<p>Character range. Any value within the range (denoted with a hyphen). You can mix character class and a hexadecimal range.</p> <p>Example: [AaBbCcDdEeFf0-9].</p>
[^<start>--<end>]	<p>Negation of character range.</p> <p>Example: [^Dd]ay matches Hay and ray, but not Day or day.</p> <p>NOTE: To negate an entire signature pattern, select the Negate option under the pattern text box.</p>
\u<string>\u	Unicode insensitive matches.
\s	Whitespace.

Table 41: Custom Attack – Attack Pattern (*continued*)

Setting	Description
\	Use a backslash to escape special characters so that they are matched and not processed as regular expression operators.
Negate	Negates the attack pattern.

Table 41: Custom Attack – Attack Pattern (*continued*)

Setting	Description
Context	<p>Binds pattern matching to a context.</p> <p>For known services, such as HTTP, select the service in the first box, and select the HTTP context you discovered with scio ccap, such as HTTP POST Parsed Param, in the second box.</p> <p>If you were unable to discover the context, select Other in the first box, and select one of the following contexts in the second box:</p> <ul style="list-style-type: none"> • Packet—Detects the pattern in any packet. • First Packet—Inspects only the first packet of a stream. When the flow direction is set to any, the detector engine checks the first packet of both the server-to-client (STC) and client-to-server (CTS) flows. Less processing means greater performance. If you know that the pattern appears in the first packet of a session, select First Packet. • First Data Packet—Inspection ends after the first packet of a stream. Select this option to detect the attack in only the first data packet of a stream. If you know that the pattern appears in the first data packet of a stream, select First Data Packet. • Stream 256—Reassembles packets and searches for a pattern match within the first 256 bytes of a traffic stream. Stream 256 is often the best choice for non-UDP attacks. When the flow direction is set to any, the detector engine checks the first 256 bytes of both the STC and CTS flows. If you know that the pattern will appear in the first 256 bytes of a session, select Stream 256. • Stream 8K—Like Stream 256 except reassembles packets and searches for a pattern match within the first 8192 bytes of a traffic stream. • Stream 1K—Like Stream 256 except reassembles packets and searches for a pattern match within the first 1024 bytes of a traffic stream. • Line—Detects a pattern within a specific line. Use this context for line-oriented applications or protocols (such as FTP). • Stream—Reassembles packets and extracts the data to search for a pattern match. However, the IDP engine does not recognize packet boundaries for stream contexts, so data for multiple packets is combined. Select this option only when no other context option contains the attack. <p>NOTE: If you select a line, stream, or service context, you do not configure match criteria for IP settings and protocol header fields.</p>
Direction	<p>Select the direction in which to detect the pattern:</p> <ul style="list-style-type: none"> • Client to Server—Detects the pattern only in client-to-server traffic. • Server to Client—Detects the pattern only in server-to-client traffic. • Any—Detects the pattern in either direction. <p>The session initiator is considered the client, even if that source IP is a server.</p>
Flow	<p>Select the flow in which to detect the attack:</p> <ul style="list-style-type: none"> • Control—Detects the pattern in the initial connection that is established to issue commands, requests, and so on. Ninety-nine percent of signatures use control. • Auxiliary—Detects the pattern in the response connection that is established intermittently to transfer requested data. This option supports a small number of protocols, such as PTP. • Both—Detects the pattern in the initial and response connections. <p>TIP: Using a single flow (instead of Both) improves performance and increases detection accuracy.</p>

Click **Next** to display the Custom Attack – IP Settings and Header Matches page. [Table 42 on page 79](#) provides guidelines for completing the settings.

12. If you have selected a line, stream, stream 256, or service context, do not configure match criteria for IP settings and protocol header fields. Click **Finish**.

If you are using a packet context, you can refine matching by adding criteria for IP flags and packet headers, as described in the following tables.



TIP: If you are unsure of the IP flags and IP fields you want to match, leave all fields blank. If no values are set, the IDP engine attempts to match the signature for all header contents.

Table 42: Custom Attack – IP Settings and Header Matches Page

Setting	Description
IP Version	Select IPv4 . IDP Series devices do not support inspection of IPv6.
Type of Service	Service type. Common service types are: <ul style="list-style-type: none"> • 0000 Default • 0001 Minimize Cost • 0002 Maximize Reliability • 0003 Maximize Throughput • 0004 Minimize Delay • 0005 Maximize Security
Packet Length	Number of bytes in the packet, including all header fields and the data payload.
ID	Unique value used by the destination system to reassemble a fragmented packet.
Time-to-live	Time-to-live (TTL) value of the packet. This value represents the number of routers the packet can pass through. Each router that processes the packet decrements the TTL by 1; when the TTL reaches 0, the packet is discarded.
Protocol	Protocol used in the attack.
Source	IP address of the attacking device.
Destination	P address of the attack target.
RB	Reserved bit. This bit is not used.
MF	More fragments. When set (1), this option indicates that the packet contains more fragments. When unset (0), it indicates that no more fragments remain.
DF	Don't fragment. When set (1), this option indicates that the packet cannot be fragmented for transmission.

[Table 43 on page 80](#) provides guidelines for completing the settings.

Table 43: Custom Attack Object: TCP Packet Header Fields

Setting	Description
Source Port	Port number on the attacking device.
Destination Port	Port number of the attack target.
Sequence Number	Sequence number of the packet. This number identifies the location of the data in relation to the entire data sequence.
ACK Number	ACK number of the packet. This number identifies the next sequence number; the ACK flag must be set to activate this field.
Header Length	Number of bytes in the TCP header.
Window Size	Number of bytes in the TCP window size.
Data Length	Number of bytes in the data payload. For SYN, ACK, and FIN packets, this field should be empty.
Urgent Pointer	Data in the packet is urgent; the URG flag must be set to activate this field.
URG Bit	When set, the urgent flag indicates that the packet data is urgent.
ACK Bit	Acknowledgment flag. When set, acknowledges receipt of a packet.
PSH Bit	Push flag. When set, indicates that the receiver should push all data in the current sequence to the destination application (identified by the port number) without waiting for the remaining packets in the sequence.
RST Bit	Reset flag. When set, resets the TCP connection, discarding all packets in an existing sequence.
FIN Bit	Final flag. When set, indicates that the packet transfer is complete and the connection can be closed.
R1 Bit, R2 Bit	Reserved bit. Unused.

[Table 44 on page 80](#) provides guidelines for completing the settings.

Table 44: Custom Attack Object: UDP Header Fields

Setting	Description
Source Port	Port number on the attacking device.
Destination Port	Port number of the attack target.
Data Length	Number of bytes in the data payload.

[Table 45 on page 81](#) provides guidelines for completing the settings.

Table 45: Custom Attack Object: ICMP Packet Header Fields

Setting	Description
ICMP	
ICMP Type	Primary code that identifies the function of the request or reply.
ICMP Code	Secondary code that identifies the function of the request or reply within a given type.
Sequence Number	Sequence number of the packet. This number identifies the location of the request/reply in relation to the entire sequence.
ICMP ID	Identification number, which is a unique value used by the destination system to associate requests and replies.
Data length	Number of bytes in the data payload.



NOTE: ICMPv6 header fields are not applicable. IDP Series devices do not support inspection of IPv6.

13. Click **Finish**.

Related Documentation

- [Verifying the Attack Object Database Version \(NSM Procedure\) on page 84](#)

Creating a Compound Attack Object (NSM Procedure)

Use compound attack objects in cases where:

- Attacks use multiple methods to exploit a vulnerability and, inspected independently, the individual contexts appear benign.
- Matching multiple contexts reduces false positives.
- Coupling a signature with a protocol anomaly reduces false positives.

To configure a compound attack object:

1. Configure general attack object properties and reference information as described for signature attack objects.
On the Target Platform and Type page, select a target platform, select **Compound Attack**, and click **Next**.
2. On the Custom Attack – General Properties page, configure the settings described in [Table 46 on page 82](#).

Table 46: Custom Attack – General Properties

Property	Description
False Positives	Same guidelines as for signature attack objects.
Service Binding	
Time Binding	

Click **Next**.

- On the Compound Members page, specify compound attack parameters and add members. [Table 47 on page 82](#) provides guidelines for completing the settings.

Table 47: Compound Attack Parameters

Setting	Description
Scope	<p>Select one of the following:</p> <ul style="list-style-type: none"> Session—Allows multiple matches for the object within the same session. Transaction—Matches the object across multiple transactions that occur within the same session.
Reset	<p>Enable to detect multiple occurrences of the attack object in the same session. Disable to log multiple occurrences as one.</p>
Boolean Expression	<p>Type a Boolean expression using the following Boolean operators:</p> <ul style="list-style-type: none"> OR—If either of the member name patterns match, the expression matches. AND—If both of the member name patterns match, the expression matches. It does not matter which order the members appear in. OAND—If both member name patterns match, and if they appear in the same order as in the Boolean expression, the expression matches. <p>For example, the Boolean expression ((s1 OAND s2) OR (s1 OAND s3)) AND (s4 AND s5) would match an attack that contains s1 followed by either s2 or s3, and that also contains s4 and s5 in any location.</p>
Add member	<p>Click the + icon, select Signature or Protocol Anomaly, and complete the configuration details.</p> <p>For signature members, specify the same contextual information as you do for a signature attack object.</p> <p>For protocol anomaly members, select from a list of predefined protocol anomalies.</p> <p>BEST PRACTICE: Our signature team uses the following naming convention for members: m01, m02, m03, and so on. We recommend you use this same naming convention.</p>

Table 47: Compound Attack Parameters (*continued*)

Setting	Description
Context Check	<p>Use this constraint to require the matching context be of a specified byte length to be a hit:</p> <ul style="list-style-type: none"> • Constraint—Select length. • Comparison operator—Select =, !, >, or <. • Operand—Select a byte length. <p>Example: You can use the context check constraint as a tuning device to limit processing for harmless traffic. For example, if you know that a certain class of attack, like a buffer overflow attack, always has an unusually large byte length in a given context, you can use this constraint to ignore contexts of normal length. If you set the FTP username context length requirement to be > 18, you see signature hits only when the FTP username context is longer than 18 bytes.</p> <p>You can specify multiple constraints. For example, if you add a < 25 constraint to the previous example, you see hits only when the username context is between 18 and 25 bytes.</p>
Match within same context	<p>Use this constraint to require selected signature members to be found in the same context instance (in any order). You can select up to 32 signature members.</p> <p>Protocol anomaly members are not selectable and are not a component of this constraint.</p> <p>Example: You design a compound attack with service context ftp-filename, and you enable this restraint. The pattern for member 1 is test; the pattern for member 2 is hello. A user opens an FTP session and requests files test.txt and hello.txt. Each file transfer occurs in its own context, not within the same context instance, so the FTP session does not trigger this attack object. Instead, consider what happens when the user requests a file named test-hello.txt. In this case, both members are found in a single context instance, so the FTP session is a match.</p>
Within Bytes Constraint	<p>NOTE: IDP OS Release 5.1 does not support the within bytes constraint for compound attack objects.</p>
Within Packets Constraint	<p>Use this constraint to require that the pattern be found completely within a packet range of a stream:</p> <ul style="list-style-type: none"> • Lower limit—Specify the beginning of the range. • Upper limit—Specify the end of the range. • Member—Select one or two members. You cannot configure a relationship for more than two members. <p>If you set a packet constraint for one member, the program logic counts packets beginning with the start-of-stream. The member must be found completely within the packet range indicated.</p> <p>If you select two members and apply a packet constraint to them, the program logic counts the first match as packet 0. If you specify a range of 1-2 with member 1 and member 2, the second pattern must occur within one or two packets from the packet containing the first match.</p> <p>Specifying 0-1 requires the pattern to appear in the same packet or within one packet from the first match. Order does not matter unless you use an Boolean ordered AND to specify the order in which the patterns must appear.</p> <p>Inspection for this object terminates when the range limit has been reached.</p>

4. Click **Finish**.



NOTE: For more information on custom attack objects references and examples, see *IDP Series Custom Attack Object Reference and Examples Guide*.

Related Documentation

- [Creating a Signature Attack Object \(NSM Procedure\) on page 69](#)

Verifying the Attack Object Database Version (NSM Procedure)

Purpose New attack objects are added to the attack object database server frequently; downloading these updates and installing them on your managed devices regularly ensures that your network is protected against the latest threats. As new attack objects are added to the attack object database server, the version number of the database increments by 1. When you download a version of the attack object database from the server, NSM stores the version number of that database.

Action Automatic Verification

The management system uses the database version number to detect and notify you when the stored attack object database on the GUI server is:

- Older than the most recent database available from the attack object database server.
- Newer than the attack object database currently installed on your ScreenOS 5.1 and later managed devices.

When NSM detects that the managed device contains an older attack object database version than the one stored on the GUI server, the UI displays a warning for that device, indicating that you should update the attack object database on the device.

Manual Verification

You can also check to see if the attack object database on the server is more recent than the one on the security device.

To check the attack object database version:

1. From the Device Manager, select **Security Updates > Check Attack Database Server Version**. The Check Attack Database Server Version dialog box appears.
2. Select the devices or group of devices to be checked.
3. Click **OK**. The Job Information window displays the status of the version check.

Related Documentation

- [Attack Objects in Intrusion Detection and Prevention Security Policies Overview on page 63](#)
- [Loading J-Security-Center Updates \(NSM Procedure\) on page 64](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\)](#)

Updating the IDP Detector Engine (NSM Procedure)

The IDP detector engine is dynamically changeable firmware that runs on ISG security devices running ScreenOS 5.0.0-IDP1, standalone IDP sensors, J-series devices, and SRX-series devices. Automatic updates to the IDP detector engine occur when you:

- Upgrade security device firmware.
- Load a new detector engine—New detector engines may be downloaded with normal attack object updates. You must load the new detector engine onto the device.



NOTE: You cannot downgrade the IDP detector engine version on the device.

To update the IDP detector engine for an IDP device:

1. From the Device Manager, select **Security Updates > Update ScreenOS/IDP Device Detector**. The Load IDP Detector Engine wizard starts.
2. Click **Next**, and then follow the instructions in the wizard to update the IDP detector engine on the selected device.



NOTE: Updating the IDP detector engine on a device does not require a reboot of the device.

For more information, see *Network and Security Manager Administration Guide*.

Related Documentation

- [Attack Objects in Intrusion Detection and Prevention Security Policies Overview on page 63](#)
- [Working with Attack Groups \(NSM Procedure\) on page 66](#)
- [Loading J-Security-Center Updates \(NSM Procedure\) on page 64](#)
- [Viewing Predefined Attack Objects \(NSM Procedure\)](#)

CHAPTER 6

Working with Application Objects

- [Application Objects in Intrusion Detection and Prevention Security Policies Overview on page 87](#)
- [Viewing Predefined Application Objects \(NSM Procedure\) on page 88](#)
- [Viewing Predefined Extended Application Objects \(NSM Procedure\) on page 88](#)
- [Creating a Custom Application \(NSM Procedure\) on page 89](#)
- [Creating Application Groups \(NSM Procedure\) on page 90](#)

Application Objects in Intrusion Detection and Prevention Security Policies Overview

Application objects are stored in the NSM database application signature table (also referred to as the appsig table) and extended application signature table (also referred to as the extappsig table). Juniper Networks Security Center (J-Security Center) develops predefined application signatures and makes them available for download to NSM during updates to the signature database. You use the NSM Object Manager to manage application objects. When you push security policy updates, the application signatures are pushed to the IDP Series device. During traffic inspection, the application identification feature matches traffic to the application objects that are specified in APE rules.

IDP Series administration can include the following tasks related to application objects:

- Updating the application signature table
- Viewing predefined application object definitions
- Creating custom application objects
- Working with application object groups

For information about how to use NSM Object Manager search features, see the *NSM online Help*.

Related Documentation

- [Viewing Predefined Application Objects \(NSM Procedure\) on page 88](#)
- [Viewing Predefined Extended Application Objects \(NSM Procedure\) on page 88](#)
- [Creating a Custom Application \(NSM Procedure\) on page 89](#)

Viewing Predefined Application Objects (NSM Procedure)

Purpose In most cases, the predefined application objects and predefined extended application objects developed by the Juniper Networks Security Center (J-Security Center) are the only ones you need to create APE rules that meet your business objectives. J-Security Center maintains a list of predefined application objects on its [website](#). We recommend that you become familiar with the predefined objects and leverage them in your APE rules as much as possible.

You can also use the NSM Object Manager to view a sortable table of predefined application objects.

Action To view the table that lists predefined application objects:

1. In the Object Manager, select **Application Objects**.
2. Click the **Predefined Application Objects** tab.
3. Click a column heading to sort the table by the column property. Double-click the table row entry for the application object to display additional details.



NOTE: You cannot edit or delete predefined application objects. You can create custom application objects.

- Related Documentation**
- [Viewing Predefined Extended Application Objects \(NSM Procedure\) on page 88](#)
 - [Creating a Custom Application \(NSM Procedure\) on page 89](#)
 - [Creating Application Groups \(NSM Procedure\) on page 90](#)

Viewing Predefined Extended Application Objects (NSM Procedure)

Purpose Extended application objects, also called nested applications, identify Web 2.0 applications running over HTTP. Extended application objects are predefined objects developed by the Juniper Networks Security Center (J-Security Center) and distributed during NSM signature database updates. You can use extended application objects in APE rules to treat various Web 2.0 applications running over HTTP differently.

Action To view table listings of predefined extended application objects:

1. In the Object Manager, select **Application Objects**.
2. Click **Predefined Extended Application Objects** tab.
3. Double-click the table row entry to display additional details.



NOTE: You cannot edit or delete predefined extended application objects. You cannot create custom extended application objects.

Related Documentation

- [Viewing Predefined Application Objects \(NSM Procedure\) on page 88](#)
- [Creating a Custom Application \(NSM Procedure\) on page 89](#)
- [Creating Application Groups \(NSM Procedure\) on page 90](#)

Creating a Custom Application (NSM Procedure)

You use the NSM Object Manager to create a custom application.

To create a custom application object:

1. In the NSM Object Manager, select **Application Objects**.
2. Click the **Custom Application Objects** tab.
3. Click the + icon to display the New Custom Application dialog box.
4. Configure custom application properties, as described in [Table 48 on page 89](#).
5. Click **OK** to save the object.

Table 48: NSM Object Manager: Custom Application Objects

Tab	Property	Configuration Guidelines
General	Name	Specify a descriptive name. Use the conventions of the predefined application object names as a model.
	Application Category	Specify an application category. Use the same categories as the predefined application objects, or specify a new category if needed.
	Supported Platforms	Click the edit icon to display the selection box. Then select the platforms you plan to test against.
	Port Ranges	Specify the range of TCP and UDP ports where the application might run. The application is identified only if the server port is within the specified range.

Table 48: NSM Object Manager: Custom Application Objects (*continued*)

Tab	Property	Configuration Guidelines
Detector	Port Binding	(Optional) Specify the standard ports on which the application usually runs.
	Signature	Specify a pattern match for client-to-server and server-to-client directions.
	Minimum data length	Specify a minimum data length to examine to match this pattern.
	Signature Match Order	Specify a signature match order. Order numbers are relative to each other. In cases where traffic matches multiple objects, an application object with the lower signature match-order number is considered the match. Be sure to examine all applications that might have the same protocol, port, and pattern; and then select a relative match order suited for the results you expect.

**Related
Documentation**

- [Creating Application Groups \(NSM Procedure\) on page 90](#)
- [Application Objects in Intrusion Detection and Prevention Security Policies Overview on page 87](#)

Creating Application Groups (NSM Procedure)

You use the NSM Object Manager to create application groups.

To create an application group:

1. In the NSM Object Manager, select **Application Objects**.
2. Click the **Application Group Objects** tab.
3. Click the + icon to display the New Application Group dialog box.
4. Give the group a name. Then use the selector controls to add or remove members to or from the group.
5. Click **OK** to save the object.

**Related
Documentation**

- [Creating a Custom Application \(NSM Procedure\) on page 89](#)
- [Application Objects in Intrusion Detection and Prevention Security Policies Overview on page 87](#)

CHAPTER 7

Configuring SNMP and Syslog Settings

Use Global Settings to specify syslog and SNMP servers. In device templates, find these settings under IDP Device Settings. In an individual device, find these settings under Global Settings.

This chapter includes the following topics:

- [Configuring an SNMP Agent \(NSM Procedure\) on page 91](#)
- [Configuring Syslog Collection \(NSM Procedure\) on page 92](#)

Configuring an SNMP Agent (NSM Procedure)

The IDP sensor creates and sends a log entry whenever one of the following statistics reaches 90%. It also sends a log when the value drops below 90%. Logs are sent no more than once a minute. The log entry specifies the value of the setting at the time the log is sent.

- **CPU Usage**—Log entry generated when CPU usage reaches 90%.
- **Hard Disk Usage**—Log entry generated when disk space reaches 90%.
- **Memory Usage**—Log entry generated when memory usage reaches 90%.
- **Session Count**—Log entry generated when session count reaches 90% of the total possible session count. (Maximum total session count for each device is specified on that device's product sheet.)

These logs are generated and sent to NSM automatically. However, you can also set the sensor to send the entries to your SNMP server. To do so, enable SNMP and specify the server's community and IP address, and then load the new configuration onto the sensor.

You configure an SNMP agent if you want to send device event logs to an SNMP server.

You have the option of configuring an SNMP agent for NSM (if you want to send the NSM collection to SNMP) or configuring an SNMP agent for each IDP device.

To configure an SNMP agent for NSM, see the NSM online Help.

To configure an SNMP agent for a single IDP device:

1. In the NSM Device Manager, double-click the IDP device to display the device configuration editor.
2. Click **Report Settings**.
3. Add or modify the settings in the SNMP Settings grid as described in the [Table 49 on page 92](#).
4. Click **OK**.

Table 49: IDP Configuration: SNMP Settings

Setting	Description
Enable SNMP	Enables forwarding to a network management system that reads SNMP.
SNMP Read Only Community	Specifies a string. The SNMP read-only community string resembles a password used for the exchange between the IDP device and the network management system.
SNMP Manager IP	Specifies the IP address of the SNMP server.
SNMP Contact	Specifies an e-mail address for the IDP administrator contact to be included in SNMP communications. If the network management system encounters a problem with the SNMP communication, it can use the contact information to follow up.
SNMP Location	Specifies a location of the IDP device to be included in SNMP communications.
New SNMP allowed hosts	Specifies the network/host IP address and the network/host netmask for the agent.

- Related Documentation**
- [Configuring Syslog Collection \(NSM Procedure\) on page 92](#)
 - [NSM Logs and Reports Overview on page 139](#)

Configuring Syslog Collection (NSM Procedure)

You configure syslog settings if you want to forward a copy of IDP logs to a syslog server.

You have the option of configuring NSM to forward a copy of its log collection to a syslog server or configuring syslog settings for each IDP device.

To have all IDP logs sent to a Syslog server, select the check box and specify the syslog server IP address. Then load the new configuration onto the sensor.

To configure syslog forwarding for NSM, see the NSM online Help.

To configure syslog forwarding for a single IDP device:

1. In the NSM Device Manager, double-click the IDP device to display the device configuration editor.
2. Click **Report Settings**.

3. Select **Enable Syslog**.
4. Enter the syslog server IP address in **Syslog Server IP**.
5. Set the syslog server port in **Syslog Server Port**. Port 514 is the industry standard and is used as the default.
6. Select the protocol in **Protocol**. UDP is the industry standard and is used as the default.
7. Specify whether to forward packet logs to the syslog server.
8. Click **OK**.

**Related
Documentation**

- [NSM Logs and Reports Overview on page 139](#)
- [Viewing Logs on page 139](#)
- [Configuring an SNMP Agent \(NSM Procedure\) on page 91](#)

CHAPTER 8

Configuring Anti-Spoof Settings

- [Configuring Antispoof Settings in Intrusion Detection and Prevention Devices \(NSM Procedure\)](#) on page 95
- [Example: Applying Antispoof to a Web Server and Database Server \(NSM Procedure\)](#) on page 96

Configuring Antispoof Settings in Intrusion Detection and Prevention Devices (NSM Procedure)

Antispoof settings are valid for standalone IDP sensors only. You can assign address objects to specific interfaces on your sensor. You can set the sensor to log or drop any connections that do not match the permitted address objects for that interface.

In addition, you can set the sensor to check incoming IP addresses against the permitted address objects for other interfaces. If the sensor detects an IP address entering the wrong interface, it can log or drop that connection.

To configure antispoof settings:

1. In NSM Device Manager, double-click the IDP device you want to configure antispoof settings. The device configuration editor appears.
2. Click **Anti-Spoof Settings**.
3. Click **New** to display the Anti-Spoof Settings dialog box.
4. Configure antispoof settings using [Table 50 on page 95](#).
5. Click **OK**.

Table 50: IDP Device Configuration: Anti-Spoof Settings

Setting	Description
Interface Name	Select a forwarding interface to configure.
Logging	Enable logging for spoofed IP address.
Alarm	Enable alerts for spoofed IP addresses.
Check Other Interfaces	Indicate whether the device should check the status of other interfaces when determining spoofing.

Table 50: IDP Device Configuration: Anti-Spoof Settings (*continued*)

Setting	Description
Action	Specify the action for the IDP device to take: None or Drop Packet .
Network Objects	Browse and select the address objects you associate with the selected interface.

Related Documentation

- [Configuring Additional Intrusion Detection and Prevention Features Overview on page 123](#)
- [Adding Intrusion Detection and Prevention Devices in NSM Overview on page 8](#)
- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)

Example: Applying Antispoof to a Web Server and Database Server (NSM Procedure)

To apply antispoof settings to a Web server and a database server:

1. Add your Web server and database server to the list of address objects.
2. Connect the Web server to the Sensor through eth2. Connect the database server to the Sensor through eth4.
3. Open the device in Device Manager.
4. Click **Anti-Spoof Settings**.
5. Click **New** to add a new entry to the antispoof table. In the dialog box that opens, configure the following settings:
 - a. Select **eth4** as the forwarding interface for the database server.
 - b. Check both the **Logging** and **Alert** check boxes because your database server is important.
 - c. Select **None** from the Action list.
 - d. Select your database server from the list of address objects.
 - e. Click **OK**.
 - f. Click **New** again to add your Web server.
 - g. Select **eth2** as the interface.
 - h. Select the **Logging** check box.
 - i. Select the **Check other interfaces** check box.

If this check box is selected, the sensor compares each IP address to the list of addresses known to be assigned to other interfaces. In other words, if the database server IP address appears at this interface, you want the sensor to let you know.
 - j. Select **None** from the Action list. You just want to log this event.
 - k. Select the Web server as the address object assigned to this interface.

**Related
Documentation**

- [Configuring Antispoof Settings in Intrusion Detection and Prevention Devices \(NSM Procedure\) on page 95](#)
- [Intrusion Detection and Prevention Services and Device Configurations Supported in NSM on page 6](#)

CHAPTER 9

Configuring Intrusion Detection and Prevention Device Settings

The IDP SM and sensor settings specify how the security module(s) on the ISG Series devices and IDP sensors handle traffic. When you add IDP, default values for all security module parameters are used. As you fine-tune a security policy to fit network traffic, you may want to edit these default values. If you make changes to the default settings, the changes only affect that device to which the security module settings apply. For detailed information on fields, refer to the *NSM online Help* or the *IDP Concepts and Examples Guide*.

- [Configuring Load-Time Parameters \(NSM Procedure\) on page 99](#)
- [Configuring Run-Time Parameters \(NSM Procedure\) on page 101](#)
- [Configuring Router Parameters \(NSM Procedure\) on page 106](#)
- [Configuring Protocol Handling \(NSM Procedure\) on page 107](#)

Configuring Load-Time Parameters (NSM Procedure)

Load-time parameters include options for tuning IDP performance. In general, you modify these settings only if you encounter performance issues. These options control the security module functions when it first powers on.

To configure load-time parameters:

1. In NSM Device Manager, double-click the IDP device for which you want to configure load-time parameters. The device configuration editor appears.
2. Click **Sensor Settings**.
3. Click the **Load Time Parameters** tab.
4. Configure load-time parameters using [Table 51 on page 100](#).
5. Click **Apply**.
6. Click **OK**.

Table 51: IDP Device Configuration: Load Time Parameters

Setting	Description
Flow table size (requires sensor restart)	For improved IDP performance, set the flow table size to limit the size of the connection table. This setting should reflect the maximum number of concurrent flows you expect to have at any one time. A TCP connection has about two flows per session, and a UDP connection has about three flows per session. The default setting is 100,000 concurrent flows. If you change this value, you have to restart the IDP device.
Enable log suppression	Log suppression reduces the number of logs displayed in the Log Viewer by displaying a single record for multiple occurrences of the same event. NOTE: If the reporting interval is set too high, log suppression can negatively impact IDP performance.
Include destination IP's while performing log suppression	When log suppression is enabled, multiple occurrences of events with the same source IP, service, and matching attack object generate a single log record with a count of occurrences. If you enable this option, log suppression combines log records for events with the same destination IP.
Number of log occurrences after which log suppression begins	This number represents the number of identical log records received before suppression starts. The default is 1 (meaning log suppression begins with the first redundancy).
Maximum number of logs that log suppression can operate on	When log suppression is enabled, IDP must cache log records so that it can identify when multiple occurrences of the same event occur. This number represents the number of log records in the IDP management server that IDP tracks for log suppression. The default is 16,384 log records.
Time (seconds) after which suppressed logs will be reported	When log suppression is enabled, the IDP device maintains a count of multiple occurrences of the same event. This number represents the number of seconds that pass before IDP reports a single log entry containing the count of occurrences. The default is 10 seconds.
Enable application identification	The application identification feature is used to detect the session application regardless of port. We recommend you disable this feature only when troubleshooting.
Maximum number of Application Identification sessions	Specifies the maximum number of sessions where application identification is in use. The default is 100,000. Valid values are 0 - 200,000. We recommend you tune this setting only if you encounter issues.
Enable policy sharing	This option allows two CPUs on a security module to share a policy. This enables the policy with all attacks to withhold maximum memory. Also the memory usage increases while the attacks database grows.

Related Documentation

- [Pushing Security Policy Updates to an IDP Device \(NSM Procedure\) on page 131](#)
- [Troubleshooting Configuration Push Errors \(NSM Procedure\) on page 133](#)
- [Configuring Run-Time Parameters \(NSM Procedure\) on page 101](#)

Configuring Run-Time Parameters (NSM Procedure)

Run-time parameters include options for tuning IDP detection methods. In general, you modify these settings only if you encounter false positives or performance issues. These options control the security module operations.

To configure run-time parameters:

1. In NSM Device Manager, double-click the IDP device for which you want to configure run-time parameters. The device configuration editor appears
2. Click **Sensor Settings**.
3. Click the **Run-time Parameters** tab.
4. Configure run-time settings using [Table 52 on page 101](#).
5. Click **Apply**.
6. Click **OK**.

Table 52: IDP Device Configuration: Run-Time Parameters

Setting	Description
Backdoor Detection	<p>Minimum interval between consecutive small packets (microseconds) / Maximum interval between consecutive small packets (microseconds)—Controls the minimum and maximum intervals (in microseconds) between the arrival of two consecutive small packets in suspected interactive traffic. If the IDP device sees small packets arrive in less than the minimum or more than the maximum number of microseconds, it does not consider the traffic to be interactive.</p> <p>The defaults are 20,000 and 2,00,00,000. This means that consecutive small packets must arrive within 20,000 to 2,00,00,000 microseconds to be considered interactive.</p> <hr/> <p>Byte threshold for packet sizes in a backdoor connection—Controls the maximum number of bytes a TCP packet must contain before the IDP device uses the packet for backdoor detection heuristics. The default is 20 bytes.</p> <hr/> <p>Minimum number of data carrying TCP packets—Controls the minimum number of data-carrying TCP packets in suspected interactive traffic. The default is 20 packets.</p> <hr/> <p>Minimum percentage of back-to-back small packets—Controls the minimum percentage of consecutive small packets in suspected interactive traffic. If the IDP device sees less than this percentage, it does not report a backdoor event. The default is 20%.</p> <hr/> <p>Ratio of small packets to the total packets (percentage)—Controls the minimum percentage of small packets that the IDP device uses for backdoor detection heuristics. If the IDP device sees less than this minimum, it does not report a backdoor event. The default is 20%.</p>

Table 52: IDP Device Configuration: Run-Time Parameters (*continued*)

Setting	Description
Flow Management	Timeout (seconds) for non-UDP/TCP/ICMP flows —Each connection through the security module typically has two non-UDP/TCP/ICMP flows, one in each direction. If IDP does not see flow activity for the specified timeout, it removes the idle flow from the flow table. The default is 30 seconds.
	Timeout (seconds) for UDP flows —Each connection through the security module typically has two UDP flows, one in each direction. If IDP does not see flow activity for the specified timeout, it removes the idle flow from the flow table. The default is 30 seconds.
	Timeout (seconds) for TCP flows —Each connection through the security module typically has two TCP flows, one in each direction. If IDP does not see flow activity for the specified timeout, it removes the idle flow from the flow table. The default is 30 seconds.
	Timeout (seconds) for ICMP flows —Each connection through the security module typically has two ICMP flows, one in each direction. If IDP does not see flow activity for the specified timeout, it removes the idle flow from the flow table. The default is 30 seconds.
	Maximum TCP Sessions —Controls the maximum number of TCP sessions that IDP maintains. If IDP reaches the maximum, it drops all new sessions and writes a SESSION_LIMIT_EXCEEDED log. Defaults vary according to model.
	Maximum UDP Sessions —Controls the maximum number of UDP sessions that IDP maintains. If IDP reaches the maximum, it drops all new sessions and writes a SESSION_LIMIT_EXCEEDED log. Defaults vary according to model.
	Maximum ICMP Sessions —Controls the maximum number of ICMP sessions that IDP maintains. If IDP reaches the maximum, it drops all new sessions and writes a SESSION_LIMIT_EXCEEDED log. Defaults vary according to model.
	Maximum IP (non-UDP/TCP/ICMP) sessions —Controls the maximum number of IP sessions that IDP maintains. If IDP reaches the maximum, it drops all new sessions and writes a SESSION_LIMIT_EXCEEDED log. Defaults vary according to model.
	Reset flow table with policy load/unload —Enables IDP to reset the flow table each time you load or unload a security policy. If you do not enable this option, IDP maintains the flow table until all flows referencing that security policy go away. This setting is enabled by default. We recommend that you keep this setting enabled to preserve memory.
IP Actions	Log flow related errors —Enables logging for flow-related errors. This setting is not enabled by default.
	Reset block table with policy load/unload —Allows the IDP device to reset the block table. The block table maintains the state of active IP actions each time a security policy loads or unloads. This setting is enabled by default.

Table 52: IDP Device Configuration: Run-Time Parameters (*continued*)

Setting	Description
Intrusion Detection	<p>Buffer Overflow emulator—Turns on buffer overflow emulation.</p> <hr/> <p>Attack matches per packet when Signature Hierarchy (0 to disable) take effect—Sets the threshold for activating Signature Hierarchy calculations.</p> <p>Common attack can be composed of several known vulnerabilities. Each vulnerability has an attack object, and each would generate a separate log entry if the signature hierarchy feature were disabled.</p> <p>For example, for a policy with critical, high, medium, low, and info attacks and logging enabled, a single detection of HTTP:IIS:COMMAND-EXEC attack generates the following logs:</p> <ul style="list-style-type: none"> • HTTP:IIS:COMMAND-EXEC [wininnt/system32/cmd.exe] (medium) • HTTP:WIN-CMD:WIN-CMD-EXE [cmd.exe] (medium) • HTTP:REQERR:REQ-MALFORMED-URL [anomaly for %xx] (medium) • HTTP:DIR:TRAVERSE-DIRECTORY (anomaly for ../) (medium) • HTTP:REQERR:REQ-LONG-UTF8CODE (anomaly for oe) (medium) • TCP:AUDIT:BAD-SYN-NONSYN (info) • HTTP:AUDIT:URL (info) • TCP:AUDIT:BAD-SYN-NONSYN (info) <p>If the number of attacks in a packet exceeds the set value, then IDP examines its signature hierarchy to see if some attacks are actually part of a larger attack. If so, then only the parent attack is displayed in the logs. In this example, if the value was set to 9 or lower, then only a log for HTTP:IIS:COMMAND-EXEC would be generated.</p> <p>An attack in the signature hierarchy may have multiple parents or multiple children. If a child attack is part of two discovered parents, IDP takes action based on the parent with the highest severity.</p> <p>Specify 0 to disable.</p>

Table 52: IDP Device Configuration: Run-Time Parameters (*continued*)

Setting	Description
Run-time Parameters	<p>Enable Per subscriber rate limit—If you implement user-role-based rules, you can apply rate limiting to all users who belong to the specified role or to each user who belongs to the specified role. By default, rate limiting is applied to all users who belong to the specified role.</p> <p>Select the check box to change the default to per user enforcement.</p> <p>NOTE: User-role-based policies require integration with an IC Series UAC deployment.</p> <hr/> <p>RPC program timeout (seconds)—IDP performs a stateful inspection of all RPC messages on port 111, then builds a table of program-to-port mapping for each RPC server that it finds on the network. This setting indicates how long an entry in the table is maintained. The default is 300 seconds.</p> <hr/> <p>RPC transaction timeout (Seconds)—All RPC messages (port 111) are based on a request/response protocol. When the IDP receives a request, it adds the request to a request table. If IDP does not receive an RPC reply in the specified timeout, the RPC entry times out. The default is 5 seconds.</p> <hr/> <p>Exempt management server flows—Exempts NSM connections from IDP processing. This setting is enabled by default.</p> <hr/> <p>Fragment timeout (seconds)—Controls when IDP drops an incomplete fragment chain because one or more fragments did not arrive. If IDP does not receive missing fragments in the specified timeout, it generates a log (FRAGMENT_TIME_EXCEEDED). The default is 5 seconds.</p> <hr/> <p>Minimum fragment size (bytes)—IDP drops all IP fragments less than the specified size (bytes). The default is 0 bytes (no fragments are dropped).</p> <hr/> <p>Maximum fragments per IP datagram—An IP datagram can be broken into many fragments which, when assembled, should not exceed 64 K. Because IP fragment processing is CPU and memory intensive, this setting controls the size of the IP fragment chain. If the number of fragments in a chain exceeds this number, IDP drops the entire fragment chain. The default is 65,535 bytes.</p> <hr/> <p>Maximum concurrent fragments in queue—IDP can perform pseudo reassembly of IP fragment chains. This setting controls the maximum number of reassembled fragment chains. Once this limit is reached, IDP drops all new IP fragment chains and generates a log (TOO_MANY_FRAGMENTS). If your network produces a large number of IP fragments, such as those produced by Network File System (NFS), increase the number of fragments per chain to eliminate unnecessary logs. The default is 16 fragments.</p> <hr/> <p>Log fragment related errors—Logs fragment related errors. This setting is not enabled by default.</p> <hr/> <p>Enable GRE decapsulation support—Enables IDP to decode generic routing encapsulation (GRE) tunnels where IP-in-GRE or PPP-in-GRE encapsulation is used. This allows IDP to inspect the packet in its original form. GRE decapsulation is not enabled by default.</p> <hr/> <p>Enable GTP decapsulation support—Enables GPRS Tunneling Protocol (GTP) decapsulation. IDP supports decapsulation of UDP GTPv0 and GTPv1 only. GTP decapsulation is not enabled by default.</p>

Table 52: IDP Device Configuration: Run-Time Parameters (*continued*)

Setting	Description
SYN Protector	<p>Enable SSL decryption support—Enables SSL inspection. SSL decryption is not enabled by default.</p> <p>Timeout for half-open SYN protected flows—A half-open SYN flow occurs during the TCP three-way handshake, after the client has sent a SYN/ACK packet to the server. The half-open connection is now in the SYN_RECV state, and is placed into a connection queue while it waits for an ACK or RST packet. The connection remains in the queue until the connection-establishment timeout expires and the half-open connection is deleted. This setting controls the connection establishment timer, which determines the number of seconds that the security module maintains a half-open SYN protected flow. The default is 5 seconds.</p>
TCP Reassembler	<p>Lower SYN's-per-second threshold below which SYN Protector will be deactivated / Upper SYN's-per-second threshold above which SYN Protector will be activated—The SYN Protector rulebase is activated when the number of SYN packets per second is greater than the sum of the lower SYN's-per-second threshold and the upper SYN's-per-second threshold.</p> <p>The SYN Protector rulebase is deactivated when the number of SYN packets per second is less than the lower SYN's-per-second threshold.</p> <p>The defaults are 1000 and 20. The SYN Protector is activated when SYN's-per-second reach 1020 and deactivated when SYN's-per-second fall below 1000.</p> <p>Ignore packets in TCP flows where a SYN hasn't been seen (recommended)—The absence of SYN flags in TCP flows is suspect, yet still a very common occurrence. IDP can ignore packets within TCP flows that do not yet contain a SYN flag. This is enabled by default.</p> <p>Close flows as soon as a FIN is seen—Enables when a TCP connection closes, IDP sees a FIN packet from each side of the connection followed by an ACK packet from each side of the connection. However, TCP does not guarantee delivery of the final ACK.</p> <p>Enables IDP to quickly close a TCP connection after receiving a FIN packet. When enabled, IDP maintains a connection waiting for a final ACK for 5 seconds, then closes the connection. This is enabled by default and recommended.</p> <p>Timeout for connected, idle TCP flows (seconds)—Controls the number of seconds that IDP maintains connected (but idle) TCP flows. The default is 3600 seconds.</p> <p>Timeout for closed TCP flows (seconds)—Controls when IDP sees a RST packet or FIN/FIN+ACK packets on a TCP connection, it closes the connection flows. IDP drops any further packets for the closed flow, but does not delete existing, closed flows from the flow table. Controls the number of seconds that closed TCP flows are maintained in the flow table. The default is 5 seconds.</p> <p>Timeout for CLOSE-WAIT/LAST-ACK TCP flows (seconds)—Controls when a TCP connection closes, IDP sees a FIN packet from each side of the connection followed by an ACK packet from each side of the connection. However, TCP does not guarantee delivery of the final ACK.</p> <p>Controls the number of seconds a connection is maintained while waiting for the final ACK.</p> <p>To improve IDP performance during heavy loads, decrease the timeout—this reduces the size of the flow table by closing connections sooner. The default is 120 seconds.</p>

Table 52: IDP Device Configuration: Run-Time Parameters (*continued*)

Setting	Description
Traffic Signatures	<p>Byte threshold for suspicious flows—Specifies a threshold for what IDP considers a small packet.</p> <p>A scan typically uses small packets to access its targets. You can exclude suspicious flows that contain large packets to prevent false positives when detecting scans.</p> <p>If IDP sees more than this maximum, it does not consider the connection to be a scan. The default is 20 bytes.</p> <p>Reporting frequency when scan is in progress—Controls how often IDP generates "in progress" logs for a stealthy scan.</p> <p>Attackers can perform blatant scans very quickly, mapping your network in just a few seconds, but these scans typically trigger IDSes and leave evidence behind. Stealthy scans are performed over much longer time periods, lasting hours, days, or even weeks, making them more difficult to detect. The default is 30 seconds.</p> <p>The number of IP tracked for session rate—Controls the number of IP addresses tracked by the session rate counter. If IDP sees more addresses than the maximum, it does not track the additional IP addresses. The default is 32,767 IP addresses.</p>
Related Documentation	<ul style="list-style-type: none"> • Updating the IDP Detector Engine (NSM Procedure) on page 85 • Configuring SYN Protector Rulebase Rules (NSM Procedure) on page 49 • Configuring Router Parameters (NSM Procedure) on page 106

Configuring Router Parameters (NSM Procedure)

Router parameters control how the security module handles address resolution protocol (ARP) requests/replies and media access control (MAC) address issues. These settings apply to proxy-ARP and bridge mode deployments. These options control packet handling for specific protocols. Use these options to control IDP Sensor routing, if applicable.

To configure router parameters:

1. In NSM Device Manager, double-click the IDP device for which you want to configure router parameters. The device configuration editor appears.
2. Click **Sensor Settings**.
3. Click the **Router Parameters** tab.
4. Configure the router parameters using [Table 53 on page 107](#).
5. Click **Apply**.
6. Click **OK**.

Table 53: IDP Device Configuration: Router Parameter Settings

Setting	Description
ARP timeout (seconds)	When the virtual router is in proxy-ARP mode, this setting controls how long an ARP entry is maintained in the virtual router. If IDP does not receive an ARP reply before the timeout expires, the ARP entry times out. The default is 3600 seconds.
ARP proxy timeout (seconds)	In proxy-ARP mode, IDP sends out proxy ARPs on all interfaces except the one on which an ARP request was received. This setting indicates how long the original ARP entry is maintained in the virtual router if IDP does not receive an ARP reply through that interface. The default is 20 seconds.
Log ARP attacks	When selected, IDP detects and logs all spoofed ARP requests/replies and other ARP anomalies. This setting is enabled by default.
MAC timeout (seconds)	When the virtual router is in bridge mode, this setting controls how long a MAC entry is maintained in the virtual router. The default is 3600 seconds.
MAC proxy timeout (seconds)	In bridge mode, IDP performs MAC discovery if the target MAC address is not in its MAC table. This setting controls how long the entry is maintained in the virtual router until a reply comes back. The default is 20 seconds.

Related Documentation

- [Configuring Protocol Handling \(NSM Procedure\) on page 107](#)
- [Configuring Load-Time Parameters \(NSM Procedure\) on page 99](#)

Configuring Protocol Handling (NSM Procedure)

The protocol anomaly detection methods identify traffic that deviates from RFC specifications. In general, you modify protocol thresholds and configuration settings only if you encounter false positives or performance issues.

To tune protocol anomaly detection thresholds:

1. In NSM Device Manager, double-click the IDP device that you want to modify. The device configuration editor appears.
2. Click **Sensor Settings**.
3. Click the **Protocol Thresholds and Configuration** tab.
4. Configure the protocol thresholds using [Table 54 on page 108](#).
5. Click **Apply**.
6. Click **OK**.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings

Setting	Description
AIM	Maximum header length —Raises a protocol anomaly if IDP detects a header containing more bytes than the specified maximum. The default is 10,000 bytes.
	Maximum type-length-value length —Raises a protocol anomaly if IDP detects an AIM/ICQ type-length-value (TLV) containing more bytes than the specified maximum. A TLV is a tuple used for passing typed information to the protocol. The default is 8000 bytes.
	Maximum inter-client-message-block length —Raises a protocol anomaly if IDP detects an AIM/ICQ inter-client-message-block (ICMB) containing more bytes than the specified maximum. The default is 2000 bytes.
	Maximum filename length —Raises a protocol anomaly if IDP detects an AIM/ICQ file name containing more bytes than the specified maximum. The default is 10,000 bytes.
DHCP	Check to see if the source port of client's packets is 68 —Raises a protocol anomaly if IDP detects DHCP traffic that originates from a port other than 68. This setting is not enabled by default.
DNS	Report unknown DNS parameters (high noise) —Detects and reports unknown DNS parameters. You must also configure an IDP rulebase rule to detect DNS anomalies. This setting is not enabled by default.
	Report unexpected DNS parameters (high noise) —Detects and reports unexpected DNS parameters. This setting is not enabled by default. You must also configure an IDP rulebase rule to detect DNS anomalies.
	Maximum length of a DNS UDP packet —Raises a protocol anomaly if IDP detects a DNS UDP packet containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum size of a NXT resource record —Raises a protocol anomaly if IDP detects an NXT resource record in a DNS request or response message of a greater size. The default is 4096 bytes. This setting tunes the following protocol anomaly attack object: DNS_BIND_NXT_OVERFLOW (key is DNS:OVERFLOW:NXT-OVERFLOW).
	Maximum time of a dns cache —Controls the maximum amount of time for a DNS query and reply. The default is 60 seconds.
	Maximum number of logs in a session —Controls the maximum number of DNS queries kept to match a reply. The default is 1000 queries.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
FTP	Maximum Line length —Raises a protocol anomaly if IDP detects an FTP username containing more bytes than the specified maximum. The default is 32 bytes.
	Maximum Username length —Raises a protocol anomaly if IDP detects an FTP password containing more bytes than the specified maximum. The default is 64 bytes.
	Maximum Password length —Raises a protocol anomaly if IDP detects an FTP pathname containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum Pathname length —Raises a protocol anomaly if IDP detects an FTP pathname containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum Sitestring length —Raises a protocol anomaly if IDP detects an FTP sitestring containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum number of login failures per-minute —Raises a protocol anomaly if IDP detects more FTP login failures in one minute than the specified maximum. The default is 4 FTP login failures per minute.
GNUTELLA	Maximum TTL hops —Raises a protocol anomaly if IDP detects a number of TTL hops that is higher than the specified maximum. The default is 8 TTL hops.
	Maximum Line length —Raises a protocol anomaly if IDP detects, in a Gnutella connection, a line that contains more bytes than the specified maximum. The default is 2048 bytes.
	Maximum Query size —Raises a protocol anomaly if IDP detects a Gnutella client query that contains more bytes than the specified maximum. The default is 256 bytes.
GOPHER	Maximum line length —Raises a protocol anomaly if IDP detects, in a Gopher server-to-client connection, a line sent by a Gopher server to a client that contains more bytes than the specified maximum. The default is 512 bytes.
	Maximum hostname length —Raises a protocol anomaly if IDP detects, in a Gopher server-to-client connection, a hostname that contains more bytes than the specified maximum. The default is 64 bytes.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
HTTP	Maximum Request length —Raises a protocol anomaly if IDP detects an HTTP request that contains more bytes than the specified maximum. The default is 8192 bytes.
	Maximum Header length —Raises a protocol anomaly if IDP detects an HTTP header that contains more bytes than the specified maximum. The default is 8192 bytes.
	Maximum Cookie length —Raises a protocol anomaly if IDP detects a cookie that contains more bytes than the specified maximum. The default is 8192 bytes. Cookies that exceed the cookie length setting can match the protocol anomaly "r;HTTP-HEADER-OVERFLOW" and produce unnecessary log records. If you are getting too many log records for the HTTP-HEADER-OVERFLOW protocol anomaly, increase the maximum cookie length.
	Maximum Authorization length —Raises a protocol anomaly if IDP detects an HTTP header authorization line that contains more bytes than the specified maximum. The default is 512 bytes. Use this setting to tune results from the Auth Overflow attack object (key is HTTP:OVERFLOW:AUTH-OVFLW).
	Maximum Content-type length —Raises a protocol anomaly if IDP detects an HTTP header content-type that contains more bytes than the specified maximum. The default is 512 bytes.
	Maximum User-agent length —Raises a protocol anomaly if IDP detects an HTTP header user-agent that contains more bytes than the specified maximum. The default is 256 bytes.
	Maximum Host length —Raises a protocol anomaly if IDP detects an HTTP header host that contains more bytes than the specified maximum. The default is 64 bytes.
	Maximum Referrer length —Raises a protocol anomaly if IDP detects an HTTP header referrer that contains more bytes than the specified maximum. The default is 8192 bytes.
	Use alternate ports as http service —If selected, the security module watches for HTTP traffic on the following ports in addition to tcp/80: 7001; 8000; 8001; 8100; 8200; 8080; 8888; 9080. This setting is enabled by default.
	Maximum number of login failures per-minute —Raises a protocol anomaly if IDP detects, between a unique pair of hosts, more login failures than the specified maximum. The default is 4 HTTP authentication failures per minute. This setting tunes the BRUTE_FORCE attack object.
	Maximum number of 301/403/404 or 405 errors per-minute —Raises a protocol anomaly if IDP detects, between a unique pair of hosts, more 301/403/404/405 errors than the specified maximum. The default is 16 HTTP errors per minute.
ICMP	Maximum Packets per second to trigger a flood —Raises a protocol anomaly if IDP detects more ICMP packets than the specified maximum. The default is 250 packets per second.
	Minimum time interval (in seconds) between packets —Raises a protocol anomaly if IDP detects ICMP packets that have less than the specified minimum time interval between them. The default is 1 second. Use this setting to tune the Flood attack object (ICMP:EXPLOIT:FLOOD).

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
IDENT	<p>Maximum requests per session—Raises a protocol anomaly if IDP detects more IDENT (identification protocol) requests than the specified maximum. The default is 1 request per session.</p> <p>This setting tunes the Too Many Requests attack object (key is IDENT:OVERFLOW:REQUEST-NUM).</p>
	<p>Maximum Request length—Raises a protocol anomaly if IDP detects an IDENT request containing more bytes than the specified maximum. The default is 15 bytes.</p> <p>This setting tunes the Request Too Long attack object (key is IDENT:OVERFLOW:REQUEST).</p>
	<p>Maximum Reply length—Raises a protocol anomaly if IDP detects an IDENT reply containing more bytes than the specified maximum. The default is 128 bytes.</p> <p>This setting tunes the Reply Too Long attack object (key is IDENT:OVERFLOW:REPLY).</p>
IKE	<p>Maximum number of payloads in an IKE message—Raises a protocol anomaly if IDP detects an IKE message with a higher number of payloads. The default is 57 payloads.</p> <p>This setting tunes detection with the TOO-MANY-PAYLOADS attack object (key is IKE:MALFORMED:2MANY-PAYLOAD).</p>
IMAP	<p>Maximum Line length—Raises a protocol anomaly if IDP detects an IMAP line containing more bytes than the maximum. The default is 2048 bytes.</p>
	<p>Maximum Username length—Raises a protocol anomaly if IDP detects an IMAP username containing more bytes than the maximum. The default is 64 bytes.</p>
	<p>Maximum Password length—Raises a protocol anomaly if IDP detects an IMAP password containing more bytes than the specified maximum. The default is 64 bytes.</p>
	<p>Maximum Mailbox length—Raises a protocol anomaly if IDP detects an IMAP mailbox containing more than the maximum. The default is 64 bytes.</p>
	<p>Maximum Reference length—Raises a protocol anomaly if IDP detects an IMAP reference containing more bytes than the specified maximum. The default is 64 bytes.</p>
	<p>Maximum Flag length—Raises a protocol anomaly if IDP detects an IMAP flag containing more bytes than the specified maximum. The default is 64 bytes.</p>
	<p>Maximum Literal length—Raises a protocol anomaly if IDP detects a literal with more octets than the specified maximum. In IMAP4 protocol, a string can be in one of two forms: literal and quoted. As defined in RFC 2060 4.3, a literal is a sequence of zero or more octets (including CR and LF), prefix-quoted with an octet count in the form of an open brace ("{"), the number of octets, close brace ("}"), and CRLF. Valid range is 1 to 1,67,77,215. The default is 65,535 bytes.</p> <p>This setting tunes detection with the imap_literal_length_overflow attack object (key is IMAP:OVERFLOW:LIT_LENGTH_OFLOW).</p>
	<p>Maximum number of login failures per-minute—Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the maximum. The default is 4 IMAP login failures per minute.</p>

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
IRC	Maximum Password length —Raises a protocol anomaly if IDP detects an Internet Relay Chat (IRC) password containing more bytes than the specified maximum. The default is 16 bytes.
	Maximum Username length —Raises a protocol anomaly if IDP detects an IRC username containing more bytes than the specified maximum. The default is 16 bytes.
	Maximum Channel length —Raises a protocol anomaly if IDP detects an IRC channel name containing more bytes than the specified maximum. The default is 64 bytes.
	Maximum Nickname length —Raises a protocol anomaly if IDP detects an IRC nickname containing more bytes than the specified maximum. The default is 16 bytes.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
LDAP	Maximum length of Integer representation in BER encoding —Raises a protocol anomaly if IDP detects an integer field of the LDAP BER containing more bytes than the specified maximum. The default is 4 bytes.
	Maximum number of left zeros for tag in BER encoding —Raises a protocol anomaly if IDP detects more left zeros in any tag in LDAP BER encoding than the specified maximum. The default is 4 left zeros.
	Maximum value of any LDAP tag in BER encoding —Raises a protocol anomaly if IDP detects a value for a tag that can be seen in the LDAP BER encoding that is greater than the specified maximum. LDAP tags are represented using 1 byte, with the top 3 bits reserved. The default is 31.
	Maximum number of left zeros for length in BER encoding —Raises a protocol anomaly if IDP detects more left zeros in any length field in LDAP BER encoding than the specified maximum. The default is 64 left zeros.
	Maximum number of search results requested by LDAP client —Raises a protocol anomaly if IDP detects an LDAP client request for more matching entries than the specified maximum. The default is 0 (indicating no limit).
	Maximum timelimit for search result requested by LDAP client —Raises a protocol anomaly if IDP detects a time limit greater than the specified maximum. The time limit is the number of seconds before a client request times out waiting for a response from the server. The default is 0 (indicating no limit).
	Maximum length of an LDAP Attribute Descriptor —Raises a protocol anomaly if IDP detects a length of an attribute descriptor field in an LDAP message containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum length of an LDAP Distinguished Name —Raises a protocol anomaly if IDP detects a length of a distinguished name field in the LDAP message containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum value of Message id in any LDAP Message —Raises a protocol anomaly if IDP detects a message ID greater than the specified maximum. The default is 2,14,74,83,647.
	Maximum length of an LDAP message —Raises a protocol anomaly if IDP detects a LDAP message that will be processed by the LDAP subsystem larger than the specified maximum. The default is 8100 bytes.
	This setting tunes the MESSAGE_TOO_LONG attack object. If IDP raises this anomaly, it logs the event and skips the message.
	Maximum number of nested operators in an LDAP search request —Raises a protocol anomaly if IDP detects a number of nested levels allowed in an LDAP search request filter argument greater than the specified maximum. The default is 8 nested operators.
	Maximum Number of login failures per-minute —Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the maximum. The default is 4 LDAP login failures per minute.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
LPR	<p>Maximum Sub-command length in RECEIVE-JOB Command—Raises a protocol anomaly if IDP detects in an Line Printer Protocol (LPR) control file a sub command line containing more bytes than the specified maximum. LPR is a TCP-based print server protocol used by line printer daemons (client and server) to communicate over networks. An LPR client uses the LPR protocol to send a print command to an LPR server (a line printer) at TCP/515. After the print command is received by the server, the client can issue subcommands to the server and send control and data files. Control files tell the line printer which functions to perform when printing the file; data files carry the payload. The default is 256 bytes.</p> <p>Maximum Reply length from server—Raises a protocol anomaly if IDP detects an LPR control filename containing more bytes than the specified maximum. The default is 64 bytes.</p> <p>Maximum Control filename length—Raises a protocol anomaly if IDP detects an LPR control filename containing more bytes than the specified maximum. The default is 64 bytes.</p> <p>Maximum Data filename length—Raises a protocol anomaly if IDP detects a data filename containing more bytes than the specified maximum. The default is 64 bytes.</p> <p>Maximum Control file size—Raises a protocol anomaly if IDP detects an LPR control file size greater than the specified maximum. The default is 1024 bytes.</p> <p>Maximum Data file size—Raises a protocol anomaly if IDP detects an LPR data file size greater than the specified maximum. The default is 64 bytes.</p> <p>Maximum Banner string length—Raises a protocol anomaly if IDP detects an LPR banner string containing more bytes than the specified maximum. A banner string is typically the filename of the print job. The default is 32 bytes.</p> <p>Maximum E-mail length—Raises a protocol anomaly if IDP detects an LPR control file e-mail address containing more bytes than the specified maximum. After the file has printed, it is sent to the e-mail address specified in the control file. The default is 32 bytes.</p> <p>Maximum Symbolic link length—Raises a protocol anomaly if IDP detects in an LPR control file a symbolic link containing more bytes than the specified maximum. A symbolic link is a file that points to another file (entry) in a UNIX file system, but does not contain the data in the target file. When the LPR protocol receives a symbolic link command in a control file, it records the symbolic link data for the print job filename to prevent directory entry changes from reprinting the file. The default maximum is 128 bytes.</p> <p>Maximum font length—Raises a protocol anomaly if IDP detects in an LPR control file a font name containing more bytes than the specified maximum. The default is 64 bytes.</p> <p>Maximum filename length for format related sub commands—Raises a protocol anomaly if IDP detects in an LPR control file a format-related file name containing more bytes than the specified maximum. The default is 32 bytes.</p>

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
MSN	Maximum Username length —Raises a protocol anomaly if IDP detects an MSN (Microsoft Instant Messaging) username containing more bytes than the specified maximum. The default is 84 bytes.
	Maximum Display name length —Raises a protocol anomaly if IDP detects an MSN display name containing more bytes than the specified maximum. The default is 128 bytes.
	Maximum Group name length —Raises a protocol anomaly if IDP detects an MSN group name containing more bytes than the specified maximum. The default is 84 bytes.
	Maximum User state length —Raises a protocol anomaly if IDP detects an MSN user state containing more bytes than the specified maximum. A user state is a three-letter code that indicates the status of the user's connection (online, offline, idle, and so on). The default is 10 bytes.
	Maximum Phone number length —Raises a protocol anomaly if IDP detects a phone number containing more bytes than the specified maximum. The default is 20 bytes.
	Maximum Length of IP:port —Raises a protocol anomaly if IDP detects an IP:port parameter containing more bytes than the specified maximum. An IP:port parameter indicates the IP address and port number of the MSN server for a switchboard session. The default is 30 bytes.
	Maximum URL length —Raises a protocol anomaly if IDP detects a URL containing more bytes than the specified maximum. The default is 1024 bytes.
MSRPC	Maximum fragment length in MSRPC message —Raises a protocol anomaly if IDP detects an MSRPC (Microsoft Remote Procedure Call) message with a fragment length greater than the specified maximum. The default is 8192.
	Maximum tower data length in endpoint mapper messages —Raises a protocol anomaly if IDP detects an endpoint mapper message with a tower data length greater than the specified maximum. The default is 8192.
	Maximum number of entries in an insert message —Raises a protocol anomaly if IDP detects an MSRPC insert message with more entries than the specified maximum. The default is 100 entries.
NFS	Maximum Name length —Raises a protocol anomaly if IDP detects an NFS packet name containing more bytes than the specified maximum. The default is 256 bytes.
	Maximum Path length —Raises a protocol anomaly if IDP detects an NFS packet pathname containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum buffer length for read/write —Raises a protocol anomaly if IDP detects an NFS read/writer buffer larger than the specified maximum. The default is 32,768 bytes.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
NTP	Minimum time (in seconds) between two requests —Raises a protocol anomaly if IDP detects the time between two client-to-server NTP requests is greater than the specified maximum. Valid values range from 64 to 1024 seconds. The default is 0 seconds (which turns the feature off).
	Maximum length for NTPv3 message —Raises a protocol anomaly if IDP detects an NTPv3 message containing more bytes than the specified maximum. The default is 68 bytes.
	Maximum length for NTPv4 message —Raises a protocol anomaly if IDP detects an NTPv4 message containing more bytes than the specified maximum. The default is 68 bytes.
	Maximum stratum value for any NTP peer —Raises a protocol anomaly if IDP detects a stratum value larger than the specified maximum. The default is 15 bytes.
	Maximum time since last update of Reference clock —Raises a protocol anomaly if IDP detects that the NTP reference clock has not been updated in more time than the specified maximum. The default is 86,400 seconds.
	Match timestamps on NTP request and response —Enables IDP to perform timestamp matching on client requests and server responses. With this setting enabled, IDP expects the server response original timestamp to match the client request transmit timestamp; otherwise IDP considers the packet a possible protocol anomaly. This setting is enabled by default.
	Maximum Authorization field length in NTP control message —Raises a protocol anomaly if IDP detects that the length of the Authentication fields in an NTP control message is larger than the specified maximum. The default is 20 bytes.
	Maximum length of any NTP control variable —Raises a protocol anomaly if IDP detects that the length of NTP control data variable name is larger than the specified maximum. The default is 128 bytes.
	Maximum length of any NTP variable value —Raises a protocol anomaly if IDP detects that the length of any NTP control data variable value is larger than the specified maximum. The default is 255 bytes.
	Maximum length of buffer to store between control packets —NTP control messages can be split across multiple UDP packets. This setting is the maximum number of characters that IDP stores in memory to ensure continuity from one packet to the other. The default is 255 bytes.
	Maximum time for an NTP Symmetric passive association to dissolve —A symmetric passive association between two NTP peers must be dissolved after sending one reply. This setting is the time in seconds after which IDP considers such an association as expired. The default is 900 seconds.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
POP3	Maximum Line length —Raises a protocol anomaly if IDP detects a POP3 line containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum Username length —Raises a protocol anomaly if IDP detects a POP3 username containing more bytes than the specified maximum. The default is 64 bytes.
	Maximum Password length —Raises a protocol anomaly if IDP detects a POP3 password containing more bytes than the specified maximum. The default is 64 bytes.
	Maximum APOP length —Raises a protocol anomaly if IDP detects an APOP containing more bytes than the specified maximum. The default is 100 bytes.
	Maximum message number —Raises a protocol anomaly if IDP detects a POP3 message number that is higher than the specified maximum. The default is 10,00,000.
	Maximum number of login failures per-minute —Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the specified maximum. The default is 4 POP3 login failures per minute.
RADIUS	Maximum number of authenticated failures per-minute —Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the specified maximum. The default is 4 RADIUS login failures per minute.
SIP	Max-Forwards threshold —Raises a protocol anomaly if IDP detects maximum number of thresholds.
SMB	Maximum registry key length —Raises a protocol anomaly if IDP detects an SMB registry key containing more bytes than the specified maximum. The default is 8192 bytes.
	Maximum number of login failures per-minute —Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the specified maximum. The default is 4 SMB login failures per minute.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
SMTP	Maximum Number of mail recipients —Raises a protocol anomaly if IDP detects an SMTP message containing more recipients than the specified maximum. The default is 100 recipients.
	Maximum Username length in RCPT and MAIL —Raises a protocol anomaly if IDP detects an SMTP message with a username containing more bytes than the specified maximum. The default is 256 bytes.
	Maximum Domain name length in RCPT and MAIL —Raises a protocol anomaly if IDP detects an SMTP message with a domain name containing more bytes than the specified maximum. The default is 64 bytes.
	Maximum Path length in RCPT and MAIL —Raises a protocol anomaly if IDP detects an SMTP message with a pathname containing more bytes than the specified maximum. The default is 256 bytes.
	Maximum Command line length (before DATA) —Raises a protocol anomaly if IDP detects an SMTP message with a command-line entry containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum Reply line length from server (default) —Raises a protocol anomaly if IDP detects an SMTP message with a reply line from the server containing more bytes than the specified maximum. The default is 512 bytes.
	Maximum Text line length (after DATA) —Raises a protocol anomaly if IDP detects an SMTP text line containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum number of nested mime multi-part attachments —Raises a protocol anomaly if IDP detects more nested attachments than the specified maximum. The default is 4 nested mime multi-part attachments.
	Maximum number of base-64 bytes to decode —Raises a protocol anomaly if IDP detects more bytes of encoded mime data than the specified maximum. The default is 64 bytes.
	Maximum length of the value for content-type's name attribute —Raises a protocol anomaly if IDP detects a name attribute in the content-type header containing more bytes than the specified maximum. The default is 128 bytes.
SYSLOG	Maximum length of the value for the content-disposition's filename attribute —Raises a protocol anomaly if IDP detects a filename attribute in the content-disposition header containing more bytes than the specified maximum. The default is 128 bytes.
	Look for email headers in message data —Controls whether IDP looks for e-mail headers in the message data, which can occur when a bounced email contains an attachment. This setting is not enabled by default.
TELNET	Maximum number of login failures per-minute —Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the specified maximum. The default is 4 TELNET login failures per minute.
TFTP	Maximum Filename length —Raises a protocol anomaly if IDP detects a filename containing more bytes than the specified maximum. The default is 128 bytes.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
VNC	Maximum Reason string length —Raises a protocol anomaly if IDP detects a VNC (Virtual Network Computing) reason string length greater than the specified maximum. A reason string contains the text that describes why a connection between a VNC server and client failed. The default is 512 bytes.
	Maximum Display name length —Raises a protocol anomaly if IDP detects a VNC display name containing more bytes than the specified maximum. The default is 128 bytes.
	Maximum cut text length —Raises a protocol anomaly if IDP detects a VNC cut text buffer containing more bytes than the specified maximum. The default is 4096 bytes.
	Verify message after the initial handshake —Enables the security module to verify VNC connections after the initial handshake. This setting is not enabled by default.
	Maximum number of login failures per-minute —Raises a BRUTE_FORCE protocol anomaly if IDP detects more login failures than the specified maximum. The default is 4 VNC login failures per minute.
WHOIS	Maximum Request length —Raises a protocol anomaly if IDP detects a WHOIS request containing more bytes than the specified maximum. The default is 128 bytes.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
YMSG	Maximum Message length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger message with a header that indicates more bytes for the total message than the specified maximum. The default is 8192 bytes.
	Maximum Username length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger ID containing more bytes than the specified maximum. The default is 84 bytes.
	Maximum Groupname length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger group name containing more bytes than the specified maximum. The default is 84 bytes.
	Maximum Crypt length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger encrypted password containing more bytes than the specified maximum. The default is 124 bytes.
	Maximum Instant message length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger message containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum Activity string length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger activity data type containing more bytes than the specified maximum. The default is 8000 bytes.
	Maximum Challenge length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger challenge containing more bytes than the specified maximum. The default is 15 bytes.
	Maximum Cookie length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger cookie containing more bytes than the specified maximum. The default is 84 bytes.
	Maximum URL length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger Web Name containing more bytes than the specified maximum. The default is 400 bytes.
	Maximum Conference message length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger join conference message containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum Conference name length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger conference name containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum E-mail length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger new e-mail alert containing an e-mail that has more bytes than the specified maximum. The default is 84 bytes.
	Maximum E-mail subject length —Raises a protocol anomaly if IDP detects an Yahoo! Messenger subject line containing more bytes than the specified maximum. The default is 128 bytes.
	This setting tunes the Mail Subject Overflow attack object (key is CHAT:YIM:OVERFLOW:MAIL-SUBJECT).
	Maximum Filename length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger file transfer containing a filename that has more bytes than the specified maximum. The default is 1000 bytes.
	Maximum Chatroom name length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger chat room name containing more bytes than the specified maximum. The default is 1024 bytes.
	Maximum Chatroom message length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger chat room message containing more bytes than the specified maximum. The default is 2000 bytes.

Table 54: IDP Device Configuration: Protocol Thresholds and Configuration Settings (*continued*)

Setting	Description
	Maximum buddy list length —Raises a protocol anomaly if IDP detects a Yahoo! Messenger buddy list containing more bytes than the specified maximum. The default is 8000 bytes.
	Maximum webcam key length —Raises a protocol anomaly if IDP detects an Yahoo! Messenger Webcam key containing more bytes than the specified maximum. The default is 124 bytes.

- Related Documentation**
- [Updating the IDP Detector Engine \(NSM Procedure\) on page 85](#)
 - [Configuring Traffic Anomalies Rulebase Rules \(NSM Procedure\) on page 51](#)

CHAPTER 10

Configuring Additional Intrusion Detection and Prevention Features

- [Configuring Additional Intrusion Detection and Prevention Features Overview on page 123](#)
- [Enabling Intrusion Detection and Prevention Processing of Encrypted and Encapsulated Traffic \(NSM Procedure\) on page 123](#)

Configuring Additional Intrusion Detection and Prevention Features Overview

You can configure additional IDP features in NSM, including enabling IDP processing of encrypted and encapsulated traffic. For more information on configuring the following features, see the *IDP ACM online Help*.

- Traffic interfaces
- IDP in advanced deployment modes
- Enable Spanning Tree Protocol
- High availability deployments
- Interoperability with Secure Access devices and Infranet Controllers

Related Documentation

- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)
- [Configuring Antispoof Settings in Intrusion Detection and Prevention Devices \(NSM Procedure\) on page 95](#)
- [Enabling Intrusion Detection and Prevention Processing of Encrypted and Encapsulated Traffic \(NSM Procedure\) on page 123](#)

Enabling Intrusion Detection and Prevention Processing of Encrypted and Encapsulated Traffic (NSM Procedure)

You can enable IDP processing of encrypted and encapsulated traffic through NSM.

1. [Enabling SSL Decryption on page 124](#)
2. [Enabling GRE Decapsulation on page 124](#)
3. [Enabling GTP Decapsulation on page 125](#)

Enabling SSL Decryption

You can enable inspection of SSL traffic by first adding keys for the target SSL servers to the IDP keystore and then enabling the IDP SSL decryption feature.

For an overview of the IDP SSL decryption feature and lists of supported encryption algorithms and SSL ciphers, see the *IDP Concepts & Examples Guide*.

To add keys for target SSL servers to the IDP keystore:

1. Use SCP or FTP to copy your private key file to the IDP device. IDP does not run an FTP server, so you have to initiate the FTP session from the IDP device.
2. Add the key to the IDP keystore.
3. Retrieve the key ID from the IDP keystore.
4. Add any other servers that use the same key.

To enable SSL decryption:

1. In the NSM Device Manager, double-click the IDP device to display the device configuration editor.
2. Click **Sensor Settings**.
3. Click the **Run-Time Parameters** tab.
4. Expand the **Run-Time Parameters** group.
5. Select **Enable SSL decryption support**.
6. Click **OK**.

Enabling GRE Decapsulation

To enable GRE decapsulation:

1. In the NSM Device Manager, double-click the IDP device to display the device configuration editor.
2. Click **Sensor Settings**.
3. Click the **Run-Time Parameters** tab.
4. Expand the **Run-Time Parameters** group.
5. Select **Enable GRE decapsulation support**.
6. Click **OK**.

Enabling GTP Decapsulation

To enable GTP decapsulation:

1. In the NSM Device Manager, double-click the IDP device to display the device configuration editor.
2. Click **Sensor Settings**.
3. Click the **Run-Time Parameters** tab.
4. Expand the **Run-Time Parameters** group.
5. Select **Enable GTP decapsulation support**.
6. Click **OK**.

Related Documentation

- [Configuring Additional Intrusion Detection and Prevention Features Overview on page 123](#)
- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)

PART 3

Managing Intrusion Detection and Prevention Devices

- [Managing Security Policies in Intrusion Detection and Prevention Devices on page 129](#)
- [Managing Profiler Settings in Intrusion Detection and Prevention Devices on page 135](#)

CHAPTER 11

Managing Security Policies in Intrusion Detection and Prevention Devices

- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)
- [Validating a Security Policy \(NSM Procedure\) on page 130](#)
- [Troubleshooting Security Policy Validation Errors \(NSM Procedure\) on page 130](#)
- [Pushing Security Policy Updates to an IDP Device \(NSM Procedure\) on page 131](#)
- [Troubleshooting Configuration Push Errors \(NSM Procedure\) on page 133](#)
- [Disabling Rules \(NSM Procedure\) on page 134](#)
- [Exporting Security Policies \(NSM Procedure\) on page 134](#)

Assigning a Security Policy in an Intrusion Detection and Prevention Device (NSM Procedure)

After you have created the necessary firewall and IDP rules within the security policy, you must perform the following steps to assign a security policy.

To assign an existing policy to the IDP device:

1. In the NSM navigation tree, select **Policy Manager > Security Policies**.
2. In Device Manager, right-click the IDP device and select **Policy > Assign Policy**.
3. From the Security Policy Name list, select the security policy you just created.
4. Click **OK**.

For more information, see the *IDP Concepts & Examples Guide* or *Network and Security Manager Administration Guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Validating a Security Policy \(NSM Procedure\) on page 130](#)
- [Troubleshooting Security Policy Validation Errors \(NSM Procedure\) on page 130](#)

Validating a Security Policy (NSM Procedure)

Validating a security policy can identify potential problems before you install it.

To validate a security policy:

1. In the navigation tree, select **Device Manager**. The Device manager appears.
2. Select **Validate > Validate IDP Policy** and select the device. A Job Manager window displays job information and progress.
3. Click **OK**.

For more information, see either the *IDP Concepts & Examples Guide* or the *Network and Security Manager Administration Guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)
- [Troubleshooting Security Policy Validation Errors \(NSM Procedure\) on page 130](#)

Troubleshooting Security Policy Validation Errors (NSM Procedure)

Problem If NSM identifies a problem in the policy during policy validation, it displays information about the problem at the bottom of the selected rulebase. For example, if you included a non-IDP capable security device in the Install On column of an IDP rule, policy validation displays an error message. You can validate those errors and troubleshoot them.

[Table 55 on page 130](#) describes security policy validation errors and how to resolve them.

Table 55: Troubleshooting: Security Policy Validation Errors

Error	Description
Rule Duplication	<p>Rule appears more than once.</p> <p>To resolve this problem, delete the duplicate.</p>
Rule Shadowing	<p>Rule shadowing occurs when two rules are designed to detect the same attack, and the first rule is either a terminal match rule or contains a more severe action than the second rule. In these cases, the second rule will never be applied.</p> <p>To resolve this problem, modify or delete one of the rules.</p>
Protocol Mismatches	<p>Protocol mismatches occur when a service object that is specified in the Service column of the security policy uses a different protocol from that specified by the default service binding of the attack object for that rule. Remember that the service binding specifies the service and port that the attack uses. Because two different protocols are specified, IDP cannot match attacks for the attack object.</p> <p>To resolve this problem, set Service to Default.</p>

Table 55: Troubleshooting: Security Policy Validation Errors (*continued*)

Error	Description
Any-Any-None Rules	<p>Any-Any-None rules are rules that specify any for the source and destination and none for attacks. Because IDP must log all packets for all connections, this rule can cause severe IDP performance penalties.</p> <p>To resolve this problem, specify network objects for the destination and attack objects for the attacks.</p>
Any-Any-One Rules	<p>Any-Any-One rules are rules that specify any for the source and destination and a single attack object for attacks. Because IDP must look at all network traffic, this rule can cause severe IDP performance penalties.</p> <p>To resolve this problem, specify network objects for the destination.</p>
Unsupported Options	<p>Rule contains options that are not supported on the target device.</p> <p>To resolve this problem, upgrade the target device or remove the option from the rule.</p>

**Related
Documentation**

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)
- [Validating a Security Policy \(NSM Procedure\) on page 130](#)

Pushing Security Policy Updates to an IDP Device (NSM Procedure)

You must run a device configuration update job (also called *pushing* an update) in the following cases:

- After you have revised the security policy assigned to an IDP device. The configuration changes you make in NSM do not affect the IDP device until you have successfully pushed the configuration to the IDP device.
- If you have deleted the device from NSM and reinstall it. In these cases, the IDP device does not retain the previous security policy assignment.
- If you use the NSM Device Manager to change IDP device settings.

To push configuration updates to multiple IDP devices:

1. Select **Devices > Configuration > Update Device Config** to display the Update Devices Options dialog box.
2. Select the devices that you want to push configuration updates to and to set update job options on. [Table 56 on page 132](#) describes devices update job options.
3. Click **OK**.

Table 56: Devices Update Job Options

Tab	Description
General	Run Summarize Delta Config —Summarizes and runs the delta change in the configuration.
Netconf	Lock configuration during update —Locks configuration while updating device configuration.
	Update to candidate config first before commit to running config —Updates the configuration before committing.
	Use confirmed commit —Enables commit confirmed.
	Rollback candidate config to running config in error —Rollbacks when there is error generated during the configuration.
	Discard uncommitted changes when exclusive lock is available —Discards any uncommitted changes during exclusive lock.
ScreenOS and IDP	Show unconnected devices —Lists all devices that are not connected.
	Update when device connects —Updates configuration when the devices are connected.
	Firewall Device Options —Not applicable.
	Standalone IDP device options —Includes the following option: <ul style="list-style-type: none"> • Restart IDP Profiler after Device Update—Restarts the Profiler.
	ISG Device Options —Not applicable.

To push an update to a specific, single device:

1. In Device Manager, right-click the device that you want to push the update to and select **Update Device** to display the Update Device Options dialog box.
2. Set update job options using [Table 57 on page 132](#).
3. Click **OK**.

Table 57: Device Update Job Options

Option	Description
Update When Device Connects	Updates the device whenever there exist a connection between the devices.
Restart IDP Profiler After Device Update	Restarts the profiler when the device gets updated.
Update IDP Rulebase Only	Updates IDP rulebase only.
Don't Show This Dialog	Does not allow this dialog box to appear again.

For more information, see the *IDP Concepts & Examples Guide*.

- Related Documentation**
- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)
 - [Troubleshooting Configuration Push Errors \(NSM Procedure\) on page 133](#)

Troubleshooting Configuration Push Errors (NSM Procedure)

Problem [Table 58 on page 133](#) provides tips for troubleshooting errors related to NSM configuration push jobs.

Table 58: Troubleshooting: Configuration Push Errors

Error	Description
Timeout	<p>The default timeout for IDP policy is 2400000 milliseconds (40 minutes).</p> <p>When you first push a policy to a newly deployed IDP device, NSM must send a lot of information (mostly attack definitions). In some cases, the update job can time out before it completes.</p> <p>To modify the timeout setting:</p> <ol style="list-style-type: none"> 1. On the NSM Device Server, open the following file in a text editor: <code>/usr/netscreen/DevSvr/var/devSvr.cfg</code> 2. Modify the following setting: <code>devSvrDirectiveHandler.idpPolicyPush.timeout 2400000</code>
The following attacks/groups cannot be updated. Not supported for version.	<p>Different versions of IDP use different detector engines. Not all attack objects are valid for all versions of the detector engine. IDP indicates which attack objects in the security policy were not valid for the loaded detector engine and, therefore, not loaded.</p> <p>This message is for information purposes only and does not indicate a problem with the IDP device or the policy.</p>
No firewall rules can be updated for device in assigned policy policyName.	<p>You try to load a policy that contains a firewall rulebase onto a standalone IDP device.</p> <p>This message just means that IDP cannot process the firewall rulebase. The IDP rulebases are still processed normally, assuming no other errors.</p>
Rule #: Packet logging with any/any rule has serious performance implications.	<p>Setting the rule to log packets causes IDP to save packets until it is sure that they will not be needed for a log entry. A rule that has any in the Source IP column and any in the Destination IP column examines all traffic. So, IDP has to save a lot of packets all the time, which impacts performance.</p>
Policy has not changed and hence will not be updated.	<p>For performance reasons, IDP does not spend resources recompiling a security policy that has not changed.</p>
Failed to update device. Failed to compile policy.	<p>Something has gone wrong with the policy compilation. Other error messages may indicate why.</p>
No license for idp.	<p>The device does not have a valid license. Unlicensed devices do not accept policy uploads.</p>

- Related Documentation**
- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)
 - [Pushing Security Policy Updates to an IDP Device \(NSM Procedure\) on page 131](#)

Disabling Rules (NSM Procedure)

You can disable a rule without deleting it in cases where you run tuning experiments, troubleshoot an issue, or otherwise need to make a quick or temporary modification.

To disable a rule, right-click inside the No. column (the first column) of the rule and select **Disable**. The rule remains in the rulebase, but a gray diagonal stripe indicates that it has been disabled. While the rule is disabled, NSM does not install the rule on any devices.

To enable a rule, right-click inside the No. column (the first column) of the rule and select **Disable** again to clear the check box. You can disable rule groups using the same method. For more information, see the *Network and Security Manager Administration Guide*.



NOTE: You cannot disable an entire security policy or a rulebase.

Related Documentation

- [NSM and Intrusion Detection and Prevention Device Management Overview on page 5](#)
- [Modifying IDP Rulebase Rules \(NSM Procedure\) on page 36](#)

Exporting Security Policies (NSM Procedure)

You can export a security policy rulebase to an HTML file.

To export a security policy to an HTML file:

1. Select the security policy and select **File > Export Policy** to display the Export Policy dialog box.
2. Select the rulebases you want to export.
3. Select a directory in which to save the exported file.
4. Click **Export** to complete the operation.

Each export creates a new directory. The default directory name is *policyname_YYMMDD_HHMMSS*. The export process puts each rulebase in a separate HTML file in that directory.

Use an HTML browser to view the exported file. For more information, see the *Network and Security Manager Administration Guide*.

Related Documentation

- [Intrusion Detection and Prevention Devices and Security Policies Overview on page 31](#)
- [Assigning a Security Policy in an Intrusion Detection and Prevention Device \(NSM Procedure\) on page 129](#)

CHAPTER 12

Managing Profiler Settings in Intrusion Detection and Prevention Devices

- [Managing Profiler Settings on page 135](#)

Managing Profiler Settings

- [Updating Profiler Settings on page 135](#)
- [Starting the Profiler on page 135](#)
- [Stopping the Profiler on page 135](#)

Updating Profiler Settings

After you have finished configuring settings on the Profiler, you must update those settings on the device. You can do this in the Device Manager by right-clicking the device and selecting **Update Device**. The Device Update Options window appears and prompts you to restart IDP Profiler after the device updates. Click **OK** to confirm.

The Job Information window appears indicating the status of the update. After this is finished, the device begins collecting data for the Profiler database.

Starting the Profiler

To manually start the Profiler, use the Devices menu, and select **IDP Profiler > Start Profiler**. In the Start Profiler dialog box, select the devices you want to use for profiling, and then click **OK**. Alternatively, you can right-click any device from the Device Manager, and select **IDP Profiler > Start Profiler**.

As your devices begin profiling your internal network, they gather information about your network hosts, their peers, ports, and Layer-7 data.

Stopping the Profiler

To manually stop the Profiler, use the Devices menu, and select **IDP Profiler > Stop Profiler**. In the Stop Profiler dialog box, select the appropriate devices, and then click **OK**. Alternatively, you can right-click any device from the Device Manager, and select **IDP Profiler > Stop Profiler**.

Related Documentation

- [Configuring Profiler Options \(NSM Procedure\) on page 13](#)

- [Modifying Profiler Settings \(NSM Procedure\) on page 25](#)

PART 4

Monitoring Intrusion Detection and Prevention Devices

- [Working with NSM Logs and Reports on page 139](#)
- [Working with Intrusion Detection and Prevention Reporter Reports on page 151](#)

CHAPTER 13

Working with NSM Logs and Reports

- [NSM Logs and Reports Overview on page 139](#)
- [Viewing Logs on page 139](#)
- [Viewing Device Status on page 142](#)
- [Viewing NSM Predefined Reports on page 145](#)
- [Creating NSM Custom Reports on page 147](#)
- [Configuring Log Suppression on page 149](#)

NSM Logs and Reports Overview

IDP devices generate logs about device status based on built-in criteria and about security events based on the security policy notification settings. These logs are automatically sent to the NSM GUI server and can be viewed in the NSM log viewer.

IDP administration includes the following log-related tasks:

- Viewing device status, logs, and reports
- Configuring log suppression, if you want to reduce the number of identical log files

Related Documentation

- [Viewing Logs on page 139](#)
- [Configuring Syslog Collection \(NSM Procedure\) on page 92](#)

Viewing Logs

NSM logging tools provide a high-level view of the activity on your network, enabling you to view summaries as well as detailed information. You can choose to view log entries for an event that occurs across domains. This section includes the following primary sections:

1. [IDP Logs on page 140](#)
2. [Using NSM Log Investigator on page 140](#)
3. [Using NSM Audit Log Viewer on page 140](#)

IDP Logs

NSM collects logs from managed IDP devices and stores them in a central log database. You can use NSM to view, manipulate, and export logs.

Table 59 on page 140 provides a reference of log views.

Table 59: Log Viewing Options

Log Views	Description
NSM Log Viewer / Log Investigator	Logs based on notification options you set for security policy rules. Logs related to device events, such as changes in the state of a traffic interface.
NSM Log Viewer / Log Investigator NSM Security Monitor	Logs produced by the Profiler feature.
NSM Audit Log Viewer	Logs generated by NSM related to the use of NSM to manage the IDP device.
statview utility	Logs produced by the application volume tracking (AVT) feature.

Using NSM Log Investigator

Purpose You use the NSM Log Investigator to analyze aggregations of logs and drill down based on properties of interest.

Action To display logs in NSM Log Investigator, select **Investigate > Log Investigator**.



TIP: For details on using NSM to modify aggregation or display options, see the NSM online Help.

Using NSM Audit Log Viewer

Purpose You use the NSM Audit Log Viewer to track the administrative changes made to a managed device. Log-entry details include the administrator that performed the change, when the change occurred, and the job results.

Action To display the NSM Audit Log Viewer table, select **Investigate > Audit Log Viewer**.

Table 60 on page 141 describes the columns in the Audit Log Viewer table.

Table 60: NSM Audit Log Viewer Table

Column	Description
Time Generated	The time the object was changed. The Audit Log Viewer displays log entries in order of time generated by Greenwich Mean Time (GMT).
Admin Name	The name of the NSM administrator who changed the object.
Admin Login Domain	The name of the domain (global or subdomain) that contains the changed object.
Authorization Status	The final access-control status of activities is either success or failure.
Command	The command applied to the object or system, for example, sys_logout or modify.
Targets	For changes made to a device configuration or object, the Audit Log Viewer displays the object type, an object name, and object domain.
Devices	For changes made to a device, the Audit Log Viewer displays the device name, object type, and device domain. For changes made to the management system, such as administrator login or logout, the Audit Log Viewer does not display target or device data.
Miscellaneous	Additional information that is not displayed in other audit log columns.

To display details of a configuration change, such as a changed IP address or renamed device, select the audit log entry for that change in the Audit Log table and view details in the Target View table, which appears below the Audit Log Viewer table.

[Table 61 on page 141](#) describes the Target View table.

Table 61: NSM Audit Log Viewer: Target View Table

Column	Description
Target Name	To see additional details for an target view entry, double-click the entry. NSM displays the configuration screen that the change was made in and marks the changed field with a solid green triangle.
Table	To set the table details for the target view entry, double-click the table. Enter or update the options.
Domain ID	Specifies the domain ID of the target view.

To display details of a non-configuration event, such as adding the device, auto-detecting a device, or rebooting a device, select the audit log entry for that change in the Audit Log table and view details in the Device View table, which is displayed below the Audit Log Viewer table.

[Table 62 on page 142](#) describes the Device View table.

Table 62: NSM Audit Log Viewer: Device View Table

Column	Description
Device Name	To see additional details for an device view entry, double-click the entry. NSM displays the Job Manager information window for the job task.
Table	To set the table details for the device view entry, double-click the table. Enter or update the options.
Domain ID	Specifies the domain ID of the device view.

Related Documentation

- [Viewing Device Status on page 142](#)
- [NSM Logs and Reports Overview on page 139](#)
- [Configuring Syslog Collection \(NSM Procedure\) on page 92](#)

Viewing Device Status

NSM keeps you up-to-date on the health of your network with the following information:

- View critical information about your devices and IDP sensors in the Device Monitor:
 - Configuration and connection status of your security devices
 - Individual device details, such as memory usage and active sessions
 - Device statistics
- View the status of your IDP clusters in the IDP Cluster Monitor.

[Table 63 on page 142](#) lists and describes device information that you can view through the Device Monitor.

Table 63: Device Status Information

Column	Description
Name	Unique name assigned to the device in NSM.
Domain	Domain in NSM in which the device is managed.
Platform	Model number of the device.
OS Version	Operating system firmware version running on the device.

Table 63: Device Status Information (*continued*)

Column	Description
Config Status	<p>Current configuration status of the device in NSM:</p> <ul style="list-style-type: none"> • None—No state has been set (does not show in Device Monitor). • Modeled—The device exists in NSM, but a connection to the device has not yet been established. • RMA—Equivalent to bringing the device into the Modeled state. RMA results from an administrator selection in the UI when a device goes down. • Waiting for 1st connect—NSM is waiting for the device to connect. You must enter a command on the device to make it connect to NSM. • Import Needed—You must import the configuration of the device into NSM. When you add a device for the first time, verify that your status indicates “Import Needed” before you attempt to import the device. During migration, this state indicates that import of the security device configuration is still required. • OS Version Adjustment Needed—The firmware version detected running on the device is different from what was previously detected in NSM. This could happen in the event that the automatic adjustment option was cleared during a change device firmware directive or an Update Device directive was issued to an IDP device with a firmware version mismatch. • Platform Mismatch—The device platform selected when adding the DMI device in NSM does not match the device itself. A device in this state cannot connect to NSM. • Device Firmware Mismatch—The OS version selected when adding a DMI device does not match the OS version running on the device itself. • Device Type Mismatch—The type of device specified when adding the device in NSM does not match the device itself. The device type might indicate whether the device is part of a vsys device, part of a cluster, or part of a virtual chassis. A device in this state cannot connect to NSM. • Detected duplicate serial number—The device has the same sequence number as another managed device. A device in this state cannot connect to NSM. • Update Needed—An update to this device is required. • Managed—The device is currently being managed by NSM. • Managed, In Sync—The physical device configuration is synced with the modeled configuration in NSM.
Config Status (continued)	<ul style="list-style-type: none"> • Managed, Device Changed—The physical device configuration is out of sync with the modeled configuration in NSM. Changes were made to the physical device configuration (the configuration on the physical device is newer than the modeled configuration). • Managed, NSM Changed—The modeled device configuration is out of sync with the physical device configuration. Changes were made to the modeled configuration (the configuration on the NSM is newer than the physical device configuration). • Managed, NSM and Device Changed—Both device configurations (physical and modeled) are out of sync with each other. Changes were made to the physical device configuration and to the modeled configuration. • Managed, Sync Pending—Completion of the Update Device directive is suspended and waiting for the device to reconnect. This state occurs only for ScreenOS devices that have the Update When Device Connects option selected during the device update.

Table 63: Device Status Information (*continued*)

Column	Description
Connection Status	<p>Connection status of the device in NSM:</p> <ul style="list-style-type: none"> Up—Device is currently connected to NSM. Down—Device is not currently connected to NSM but has connected in the past. Never Connected—Device has never connected to NSM. <p>The Device Server checks the connection status of each device every 120 seconds by default. You can change this behavior by editing the value for the <code>devDaemon.deviceHeartbeatTimeout</code> parameter in the Device Server configuration file. Refer to the <i>Network and Security Manager Installation Guide</i> for more information on editing configuration files.</p> <p>NOTE: If the network connection goes down for a period longer than six to eight minutes, the device connection will permanently time out. If this occurs and the device goes down for any reason, the device still appears as Up in the Device Monitor.</p>
Alarm	<p>Displays the current alarm status for each device in NSM:</p> <ul style="list-style-type: none"> If device has any alarms, the most severe alarm severity is displayed (either Major or Minor). None—The device has no alarms. Unknown—The device status is unknown. For example, the device might not be connected. N/A—The device's alarm is not pollable or discoverable, for example, this column shows "N/A" for ScreenOS and IDP devices. Alarm is colored: <ul style="list-style-type: none"> Red for Major. Orange for Minor. Green for Ignore, None, Unknown, or N/A.
H/W Inventory Status	<p>Displays the inventory status for hardware on the device:</p> <ul style="list-style-type: none"> In Sync—The inventory information in the NSM database is synchronized with the information on the device. Out Of Sync—The inventory information in the NSM database is not synchronized with the information on the device. N/A—The connected device is a ScreenOS or IDP device, or the device is not connected and imported.
S/W Inventory Status	<p>Displays the inventory status for software on the device:</p> <ul style="list-style-type: none"> In Sync—The inventory information in the NSM database is synchronized with the software on the device. Out Of Sync—The inventory information in the NSM database is not synchronized with the software on the device. N/A—The connected device is a ScreenOS or IDP device, or the device is not connected and imported.

Table 63: Device Status Information (*continued*)

Column	Description
License Inventory Status	<p>Displays the inventory status for software on the device:</p> <ul style="list-style-type: none"> • In Sync—The inventory information in the NSM database is synchronized with the licenses on the device. • Out Of Sync—The inventory information in the NSM database is not synchronized with the licenses on the device. • N/A—The connected device is a ScreenOS or IDP device, or the device is not connected and imported.
First Connect	The first time the security device connected to the NSM device server.
Latest Connect	The last time the security device connected to the NSM device server.
Latest Disconnect	The last time the security device disconnected from the NSM device server.

- Related Documentation**
- [Creating NSM Custom Reports on page 147](#)
 - [NSM Logs and Reports Overview on page 139](#)

Viewing NSM Predefined Reports

You can use the predefined reports to validate the effectiveness of your security policies.

[Table 64 on page 145](#) describes NSM DI/IDP predefined reports. These reports are related to attacks.

Table 64: NSM DI/IDP Predefined Reports

Report	Description
Top 100 Attacks (last 24 hours)	Those attacks that are detected most frequently within the last 24 hours.
Top 100 Attacks Prevented (last 24 hours)	Those attacks that are prevented most frequently within the last 24 hours.
Top 20 Attackers (All Attacks - last 24 hours)	IP addresses that have most frequently been the source of an attack during the last 24 hours.
Top 20 Attackers Prevented (All Attacks - last 24 hours)	IP addresses that have most frequently been prevented from attacking the network during the last 24 hours.
Top 20 Targets (last 24 hours)	IP addresses that have most frequently been the target of an attack during the last 24 hours.
Top 20 Targets Prevented (last 24 hours)	IP addresses that have most frequently prevented attacks during the last 24 hours.
All Attacks by Severity (last 24 hours)	Number of attacks by severity level (set in attack objects).

Table 64: NSM DI/IDP Predefined Reports (*continued*)

Report	Description
All Attacks Prevented by Severity (last 24 hours)	Number of attacks prevented by severity level (set in attack objects).
All Attacks Over Time (last 7 days)	All attacks detected during the last 7 days.
All Attacks Prevented Over Time (last 7 days)	All attacks prevented during the last 7 days.
All Attacks Over Time (last 30 days)	All attacks detected during the last 30 days.
All Attacks Prevented Over Time (last 30 days)	All attacks prevented during the last 30 days.
Critical Attacks (last 24 hours)	All attacks categorized as "critical" detected during the past 24 hours.
Critical Attacks Prevented (last 24 hours)	All attacks categorized as "critical" prevented during the past 24 hours.
Critical through Medium Attacks (last 24 hours)	All attacks categorized as either "critical" or "medium" detected during the past 24 hours.
Critical through Medium Attacks Prevented (last 24 hours) (last 24 hours)	All attacks categorized as either "critical" or "medium" prevented during the past 24 hours.
Top 50 Scan Sources (last 7 days)	IP addresses that have most frequently performed a scan of a managed device.
Top 50 Scan Targets (last 7 days)	IP addresses that have most frequently been the target of a scan over the last 7 days.
Profiler - New Hosts (last 7 days)	New hosts listed in the Profiler over the last 7 days.
Profiler - New Ports (last 7 days)	New ports listed in the Profiler over the last 7 days.
Profiler - New Protocols (last 7 days)	New protocols listed in the Profiler over the last 7 days.
Top IDP Rules	The total number of log entries generated by specific rules in your IDP policies. You can use the Top Rules report to identify those rules that are generating the most log events. This enables you to better optimize your rulebases by identifying those rules that are most and least effective. You can then modify or remove those rules from your security policies.

[Table 65 on page 146](#) describes Profiler predefined reports. These reports are related to activity by hosts in your network.

Table 65: NSM Profiler Predefined Reports

Report	Description
Top 10 Peers by Count	Ten source and destination IP addresses that appeared most frequently in the Profiler logs.

Table 65: NSM Profiler Predefined Reports (*continued*)

Report	Description
Top 10 Peers with maximum hits	Ten source and destination IP addresses that had the highest number of hits in the Profiler logs.

[Table 66 on page 147](#) describes the predefined application volume tracking reports. The reports are related to application use in your network.

Table 66: NSM: Application Volume Tracking Reports

Report	Description
Top 10 Applications by Volume	Applications with the highest volume in bytes in the past 24 hours.
Top 10 Application Categories by Volume	Application categories with the highest volume in bytes in the past 24 hours.
Top 5 Applications by Volume over Time (last 1 hour)	Applications with the highest volume in bytes in the past 1 hour.
Top 5 Application Categories by Volume (last 1 hour)	Application categories with the highest volume in bytes in the past 1 hour.
Top 5 Source by Volume over Time (last 1 hour)	Source IP addresses with the highest volume in bytes in the past 1 hour.
Top 5 Destination by Volume over Time (last 1 hour)	Destination IP addresses with the highest volume in bytes in the past 1 hour.

- Related Documentation**
- [Creating NSM Custom Reports on page 147](#)
 - [NSM Logs and Reports Overview on page 139](#)

Creating NSM Custom Reports

Purpose You use custom reports if you require a view of data not covered by predefined reports.

Action To create a custom report:

1. In the NSM navigation tree, select **Investigate > Report Manager**.
2. Select a pre-defined report with data similar to what you ultimately want to save.
3. Select **File > Save As**.
4. Use pre-defined report as a template and example, complete the configuration options, and click **OK** to save the new report settings.

[Table 67 on page 148](#) describes configuration options.

Table 67: Custom Report Configuration Options

Tab	Field	Description
General	Name	Specify a name for the report as you would like it to appear in the NSM navigation tree.
	Report Title	Specify a name for the report as you would like it to appear at the top of the report.
	Type of Report	<p>Select a report type:</p> <ul style="list-style-type: none"> • Count-Based—Displays total current activity to date. For example, the Top Scan Targets report is a count-based report that displays the total number of scans currently recorded against a specified number of destination IP addresses. • Time-Based—Displays activity over time. For example, the Attacks Over Time report is a time-based report that measures the top attacks recorded in log records over a specified period. • Sum-Based—Displays the sum of the activities to date.
	Columns for Report	In reports, columns are the same as log fields.
	Time Period	<p>You can configure a report to display all available data from either a specific date and time or during a specific time interval. For example, if you suspect that your network was attacked on September 15 at 6:00 PM, you could set the Starting At Time Period Duration report field in the options on a Top Screen Attacks report to that time, then generate the report. If you are not sure of the exact date or time of the attack, but know it occurred during the past 2 days, set the Duration field in the Time Period Duration report options on a Top Screen Attacks report to two days, then generate the report.</p> <p>NOTE: The data that you can display in each report is limited by the amount of log information available.</p>
	Data point count	<p>Typically, the top 50 occurrences of each data type are displayed in each report. You can configure a report to display more or fewer data points depending upon the level of detail you need. For example, if you want to obtain a more precise view of the top occurrences of events, you would configure a lower data point count (such as 25).</p> <p>NOTE: The minimum data point count that you can configure in all reports is 5; the maximum data point count is 200.</p>
	Chart type	<p>Select from the following choices:</p> <ul style="list-style-type: none"> • Horizontal bar (default) • Pie • Line • Vertical bar
Filter	Save Report In	<p>In the first selection box, specify whether to save in the My Reports or Shared Reports node.</p> <p>In the second box, select the Others folder or type a new folder name.</p>
	Columns for Report	The columns you selected on the General tab are passed through. Select the column with the cursor to display the corresponding Filter Settings controls.
	Filter Settings	Specify filter values related to column settings.



TIP: For information on deleting custom reports, organizing report folders, exporting reports, and using the NSM guiSvrCli.sh command line utility and Linux cron utility to automate reporting jobs, see the NSM online Help.

- Related Documentation**
- [NSM Logs and Reports Overview on page 139](#)
 - [Intrusion Detection and Prevention Reporter Overview on page 151](#)

Configuring Log Suppression

You can configure log suppression if you want to reduce the number of logs displayed in the NSM log viewer. If you enable log suppression, NSM displays a single record for multiple occurrences of similar events, along with a count of all such occurrences.

To enable and configure log suppression:

1. In the NSM Device Manager, double-click the IDP device to display the configuration editor.
2. Click **Sensor Settings**.
3. Click **Load-Time Parameters**.
4. Complete the settings related to log suppression using [Table 68 on page 149](#).

Table 68: IDP Configuration: Log Suppression Settings

Setting	Description
Enable log suppression	Log suppression is enabled by default. Use this setting to turn log suppression off and on.
Include destination IPs when performing log suppression	When log suppression is enabled, multiple occurrences of events with the same source IP, service, and matching attack object generate a single log record with a count of occurrences. If you enable this option, log suppression combines log records for events with the same destination IP.
Number of log occurrences after which log suppression begins	This number represents the number of identical log records received before suppression starts. The default is 1 (meaning log suppression begins with the first redundancy).
Maximum number of logs that log suppression can operate on	When log suppression is enabled, IDP must cache log records so that it can identify when multiple occurrences of the same event occur. This number represents the number of log records in the IDP Management Server that IDP tracks for log suppression. The default is 16384 log records.
Time (seconds) after which suppressed logs will be reported	When log suppression is enabled, the IDP device maintains a count of multiple occurrences of the same event. This number represents the number of seconds that pass before IDP reports a single log entry containing the count of occurrences. The default is 10 seconds.
Maximum number of logs that can be stored	Determines the limit for logs stored on the IDP device. The default is 50,000. The minimum value is 1,000. The maximum is 65,535.

Table 68: IDP Configuration: Log Suppression Settings (*continued*)

Setting	Description
Maximum number of packet captures that can be stored	Determines the limit for packet captures stored on the IDP device. The default is 10,000. The minimum value is 100; the maximum is 1,000 packets. The maximum is 65,535.

- Related Documentation**
- [NSM Logs and Reports Overview on page 139](#)
 - [Configuring Syslog Collection \(NSM Procedure\) on page 92](#)

CHAPTER 14

Working with Intrusion Detection and Prevention Reporter Reports

- [Intrusion Detection and Prevention Reporter Overview on page 151](#)

Intrusion Detection and Prevention Reporter Overview

IDP reports are representations of log data, aggregated and sorted to facilitate network and security analysis. The standalone IDP solution supports both centralized, aggregated reporting through NSM and on-box reporting for a singular IDP device instance through IDP Reporter.

You can perform the following reporting tasks:

- View predefined reports on attacks.
- Generate reports based on criteria you specify
- Create jobs that regularly generate predefined or custom reports and e-mail them to customers or other third-party organizations.

For more information on using IDP Reporter, see the *IDP Reporter User's Guide*.

Related Documentation

- [NSM Logs and Reports Overview on page 139](#)
- [Creating NSM Custom Reports on page 147](#)

PART 5

Index

- [Index on page 155](#)

Index

A

application groups	
viewing	89, 90
application objects	
task list.....	87
viewing predefined.....	88
attack objects	
creating compound.....	81
creating signature.....	69

B

Boolean expressions in compound attack	
objects.....	82

C

compound attack objects	
Boolean expressions.....	82
creating.....	81
protocol anomalies.....	82
compound attack objects, creating.....	81
context check constraint.....	74
context, specifying.....	78
custom application objects	
viewing.....	89, 90
customer support.....	x
contacting JTAC.....	x

D

DFA expressions	
syntax.....	75
direction	
specifying.....	78

E

extended application objects	
viewing	88
viewing predefined.....	89, 90

F

flow, specifying.....	78
-----------------------	----

I

ICMP packet header matching.....	79
IP packet header matching.....	79

P

packet header matching.....	79
protocol anomalies in compound attack	
objects.....	82

S

Security.....	13
service binding.....	71
signature attack objects	
creating.....	69
signature attack objects, creating.....	69
support, technical See technical support	

T

TCP packet header matching.....	79
technical support	
contacting JTAC.....	x
time binding.....	73

U

UDP packet header matching.....	79
---------------------------------	----

W

within bytes constraint.....	74
within packets constraint.....	74

