

IDP Series

Concepts and Examples Guide

Release
5.1rX



Published: 2011-05-05
Revision 03

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.

IDP Series Concept and Examples Guide
Copyright © 2011, Juniper Networks, Inc.
All rights reserved.

Revision History
May 2011 —Revision 03

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The 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

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

- a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
- b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
- c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
- d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
- e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the

Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14 (ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Table of Contents

	Preface	xvii
	Objectives	xvii
	Audience	xvii
	Documentation Conventions	xvii
	Related Documentation	xix
	Requesting Technical Support	xx
Part 1	Solution Overview	
Chapter 1	IDP Series Product Overview	3
	Juniper Networks IDP Solutions	3
	IDP Series Features Overview	3
Chapter 2	IDP Series Components Overview	9
	IDP Series Operating System Overview	9
	IDP Series Multicore Architecture	9
	Auto-Recovery Feature	11
	Flow Bypass Feature	12
	Key Processes	12
	IDP Series Network Interfaces Overview	13
	Management Interface (eth0)	15
	High Availability Interface (eth1)	15
	Traffic Interfaces	15
	Internal Bypass	16
	External Bypass	17
	Interface Signaling	18
	Peer Port Modulation	18
	Centralized Management with NSM Overview	20
	J-Security Center Updates Overview	21
Part 2	Analyzing Your Network	
Chapter 3	Simulation Mode	25
	Simulation Mode Overview	25
	Topology	25
	Purpose	25
	Configuration Overview	26
	Logging	26
	Example: Getting Started with Simulation Mode	27
	Example: Using Simulation Mode to Maximize Uptime	29

Chapter 4	Profiler	31
	Profiler Overview	31
	Example: Using Profiler to Set a Baseline	33
	Example: Using Profiler to Alert You to New Hosts and Port Activity	38
	Example: Identifying Services That Use Nonstandard Ports	38
	Example: Responding to Vulnerability Announcements with Due Diligence	39
	Example: Using Profiler to Investigate Unanticipated Attacks	40
	Example: Using Profiler to Mitigate Risks from Laptops	41
Chapter 5	Security Explorer Overview	43
	NSM Security Explorer Overview	43
Chapter 6	Application Volume Tracking	45
	Application Volume Tracking Overview	45
	Example: Using NSM to Enable and View Application Volume Tracking	47
Chapter 7	Logs and Reports	53
	IDP Logs Overview	53
	Developing a Logging Strategy	58
	Developing a Log Storage Strategy	59
	Log Management Considerations	59
	Local Log Files and Directories	59
	NSM Log Collection	61
	Example: Using NSM Log Viewer Features	62
	Using Predefined Views	62
	Showing and Hiding Columns	63
	Using Filters	63
	Using Log Viewer Detail Panes	65
	Using Flags and Comments	65
	Using Custom Views	67
	Example: Packet Logging Workflow	68
	Using Packet Captures	68
	Enabling Packet Capture in Security Policy Rules	68
	Forwarding Packet Capture Logs to NSM	69
	Viewing Packet Capture Logs	70
	Using the NSM Packet Viewer	70
	Using an External Viewer to View Packet Data	71
	NSM Reports Overview	74
	IDP Reporter Overview	76
Part 3	Protecting Your Network	
Chapter 8	Security Policy Basics	79
	Understanding Non-Policy-Based Drops	79
	Understanding the Components of an IDP Security Policy	83
	Understanding the Number of Available and Installed Policies	85
	Understanding the Rule-Matching Algorithm	85

Chapter 9	Predefined Security Policies	87
	Using the Recommended Security Policy	87
	Using Other Security Policy Templates	88
Chapter 10	The IDP Rulebase	91
	Understanding the IDP Rulebase	91
	Understanding IDP Rulebase Rule Match Settings	92
	User-Role-Based Policy Feature Overview	94
	Using Application Identification	96
	Using Attack Objects	97
	Attack Objects Overview	98
	Understanding Predefined Attack Objects and Attack Object Groups	99
	Using Attack Object Groups	99
	Using Custom Attack Objects	100
	Understanding IDP Rulebase Actions	101
	Understanding IDP Rulebase Notification Options	103
	IDP Rulebase Example: User-Role-Based Policies	104
	IDP Rulebase Example: Using Application Identification	107
	IDP Rulebase Example: Specifying the Default Service	108
	IDP Rulebase Example: Using Recommended Attack Objects	109
	IDP Rulebase Example: Using Recommended Actions	110
Chapter 11	The Exempt Rulebase	113
	Understanding the Exempt Rulebase	113
	Exempt Rulebase Example: Exempting a Source Destination Pair	114
	Exempt Rulebase Example: Exempting an Attack Object	115
Chapter 12	Application Policy Enforcement Rulebase	117
	Understanding the APE Rulebase	117
	Understanding APE Rulebase Match Conditions	118
	Using Application Objects	121
	Application Objects Overview	121
	Understanding Predefined Application Objects	121
	Using Application Groups	126
	Using Custom Application Objects	127
	Understanding APE Rulebase Actions	128
	Understanding APE Rulebase Notification Options	130
	APE Rulebase Example: Using Extended Application Objects	131
	APE Rulebase Example: Aggregate and Per-Subscriber Rate Limits	136
	APE Rulebase Example: Matching Custom Application Objects	137
Chapter 13	The Backdoor Rulebase	141
	Understanding the Backdoor Rulebase	141
	Understanding Backdoor Rulebase Match Settings	143
	Understanding the Backdoor Rulebase Operation Setting	144
	Understanding Backdoor Rulebase Actions	144
	Understanding Backdoor Rulebase Notification Options	145
	Backdoor Rulebase Example: netcat	146

Chapter 14	The SYN Protector Rulebase	149
	Understanding the SYN Protector Rulebase	149
	Understanding SYN Protector Rulebase Match Settings	151
	Understanding SYN Protector Rulebase Modes	152
	Understanding SYN Protector Rulebase Notification Options	153
Chapter 15	The Traffic Anomalies Rulebase	155
	Understanding the Traffic Anomalies Rulebase	155
	Understanding Traffic Anomalies Rulebase Match Conditions	157
	Understanding Traffic Anomalies Rulebase Detection Settings	158
	Understanding Traffic Anomalies Rulebase IP Actions	159
	Understanding Traffic Anomalies Rulebase Notification Options	160
Chapter 16	The Network Honeypot Rulebase	161
	Understanding the Network Honeypot Rulebase	161
	Understanding Network Honeypot Rulebase Match Settings	162
	Understanding Network Honeypot Operation Setting	162
	Understanding Network Honeypot Rulebase IP Actions	163
	Understanding Network Honeypot Rulebase Notification Options	164
Chapter 17	Fine-Tuning a Security Policy	167
	Example: Fine-Tuning a Security Policy	167
	Fine-Tuning Security Policies Process Overview	167
	Getting Started with the Recommended Security Policy	168
	Refining Rule Matching Properties	168
	Reducing False Positives	169
	Adding Rulebases	172
Chapter 18	Additional Security Features	173
	IP Spoof Attack Prevention Overview	173
Part 4	Additional Deployment Topics	
Chapter 19	Inspection of Encapsulated and Encrypted Traffic	177
	Inspection of GRE Traffic Overview	177
	Inspection of GTP Traffic Overview	177
	Inspection of IPsec VPN Traffic Overview	178
	Inspection of MPLS Traffic Overview	178
	Inspection of SSL Traffic Overview	179
	Using the SSL Server Private Keys	179
	Using a Root Certificate Authority in SSL Forward Proxy Operations	180
	Supported SSL Specifications	181
	Example: Implementing Inspection of Outbound SSL Traffic	183
	Example: Exempting Outbound SSL Traffic from Inspection	185
Part 5	Index	
	Index	193

List of Figures

Part 1	Solution Overview	
Chapter 2	IDP Series Components Overview	9
	Figure 1: IDP Multicore Architecture	10
	Figure 2: IDP Series Network Interfaces	14
	Figure 3: Internal Bypass	17
	Figure 4: External Bypass	18
	Figure 5: Peer Port Modulation	19
	Figure 6: IDP-NSM Communication	20
Part 2	Analyzing Your Network	
Chapter 3	Simulation Mode	25
	Figure 7: Packet Processing in Simulation Mode	26
	Figure 8: NSM Log Viewer: Simulation Mode Logs	27
Chapter 4	Profiler	31
	Figure 9: NSM Network Address Object Editor	33
	Figure 10: NSM Group Object Editor	34
	Figure 11: Starting Profiler from NSM Device Manager	34
	Figure 12: NSM Profiler Configuration Tabs	35
	Figure 13: NSM Profiler Tracked Hosts Tab	35
	Figure 14: NSM Profiler Update Job Information Window	36
	Figure 15: Profiler: Network Profiler Tab	37
Chapter 5	Security Explorer Overview	43
	Figure 16: NSM Security Explorer	44
Chapter 6	Application Volume Tracking	45
	Figure 17: Profiler Settings: Enable AVT	48
	Figure 18: Profiler Viewer: Application Profiler Tab	48
	Figure 19: Profiler Viewer: Application Profiler Tab: Nested Applications	49
	Figure 20: NSM AVT Report	50
Chapter 7	Logs and Reports	53
	Figure 21: IDP Log Storage and Log Forwarding	61
	Figure 22: NSM Log Viewer: Predefined View	62
	Figure 23: NSM Log Viewer: Choose Columns	63
	Figure 24: NSM Log Viewer: Filters	64
	Figure 25: NSM Log Viewer: Filters	64
	Figure 26: Using NSM Log Viewer Attack Reference Information	65
	Figure 27: Using NSM Log Viewer Flag and Comment Features	66
	Figure 28: NSM Log Viewer: Custom View	67

	Figure 29: Notification Options: Packet Logging	69
	Figure 30: NSM Log Viewer: Has Packet Data Column	69
	Figure 31: NSM Device Configuration Editor: Report Settings	70
	Figure 32: NSM Packet Capture Viewer	71
	Figure 33: Specifying an External Viewer	72
	Figure 34: Wireshark Packet Viewer	73
Part 3	Protecting Your Network	
Chapter 8	Security Policy Basics	79
	Figure 35: Security Policy Components	83
Chapter 10	The IDP Rulebase	91
	Figure 36: IC Series Admin Console: Configuring User Roles	105
	Figure 37: ACM: Generating a One-Time Password for the Connection from the IC Series Appliance	106
	Figure 38: IC Series Admin Console: Configuring the Connection to the IDP Appliance	106
	Figure 39: IDP Rulebase: User-Role-Based Rules	107
	Figure 40: A Simplified Rule Enabled by the Application Identification Feature	108
	Figure 41: Default Service	108
	Figure 42: Recommended Attack Objects	109
	Figure 43: Recommended Action	111
Chapter 11	The Exempt Rulebase	113
	Figure 44: Exempt Rulebase Rule	114
Chapter 12	Application Policy Enforcement Rulebase	117
	Figure 45: NSM Object Manager: Predefined Application Objects	122
	Figure 46: NSM Object Manager: Predefined Application: General Tab	123
	Figure 47: NSM Object Manager: Predefined Application: Detector Tab	124
	Figure 48: NSM Object Manager: Predefined Extended Application Objects	125
	Figure 49: NSM Object Manager: Extended Application Details	126
	Figure 50: NSM Object Manager: Extended Application Member Details	126
	Figure 51: NSM Object Manager: Application Group Dialog Box	127
	Figure 52: NSM Object Manager: Custom Application Dialog Box	128
	Figure 53: NSM Object Manager: Predefined Extended Application Objects	132
	Figure 54: NSM Object Manager: Extended Application Details	133
	Figure 55: NSM Object Manager: Extended Application Details	134
	Figure 56: APE Rulebase: Using Extended Applications	134
	Figure 57: NSM Object Manager: Creating Application Groups	135
	Figure 58: NSM Object Manager: Creating Application Groups	135
	Figure 59: APE Rulebase: Using Application Groups	135
	Figure 60: APE Rulebase: User-Role-Based Rules to Support Tiered Access	136
	Figure 61: APE Rulebase: User-Role-Based Rules to Support Tiered Access	137
	Figure 62: NSM Object Manager: Custom Application Object	138
	Figure 63: NSM Object Manager: Custom Application Object	139
	Figure 64: APE Rulebase: Adding a Custom Application Object	139
	Figure 65: APE Rulebase: Rule Order	140

Chapter 13	The Backdoor Rulebase	141
	Figure 66: NSM Device Manager: Sensor Settings > Run-Time Parameters	142
	Figure 67: Backdoor Rulebase	147
Chapter 14	The SYN Protector Rulebase	149
	Figure 68: NSM Device Manager: Sensor Settings > Run-Time Parameters	150
Chapter 17	Fine-Tuning a Security Policy	167
	Figure 69: Using NSM Log Viewer Attack Reference Information	170
	Figure 70: Using NSM Log Viewer Flag and Comment Features	171
Part 4	Additional Deployment Topics	
Chapter 19	Inspection of Encapsulated and Encrypted Traffic	177
	Figure 71: SSL Inspection Using SSL Server Private Keys	180
	Figure 72: SSL Inspection Using a Root CA	180
	Figure 73: Firefox: Displaying the Server Certificate for a Website	187
	Figure 74: Internet Explorer: Displaying the Server Certificate for a Website	188

List of Tables

	Preface	xvii
	Table 1: Notice Icons	xviii
	Table 2: Text Conventions	xviii
	Table 3: Syntax Conventions	xix
	Table 4: Related IDP Series Documentation	xix
Part 1	Solution Overview	
Chapter 1	IDP Series Product Overview	3
	Table 5: IDP Series Features	3
	Table 6: Intrusion Detection Methods	6
Chapter 2	IDP Series Components Overview	9
	Table 7: IDP Multicore Architecture	10
	Table 8: Processes	12
Part 2	Analyzing Your Network	
Chapter 6	Application Volume Tracking	45
	Table 9: Application Volume Tracking Data	45
	Table 10: Application Volume Tracking Log Viewing Tools	46
	Table 11: Application Profiler Session Table	49
	Table 12: NSM: Application Volume Tracking Reports	51
Chapter 7	Logs and Reports	53
	Table 13: IDP Logging Options	53
	Table 14: NSM Log Viewer: Log Columns	54
	Table 15: IDP Local Log Directories	60
	Table 16: NSM DI/IDP Predefined Reports	74
	Table 17: NSM Profiler Predefined Reports	75
	Table 18: NSM: Application Volume Tracking Reports	75
Part 3	Protecting Your Network	
Chapter 8	Security Policy Basics	79
	Table 19: Non-Policy-Based Drops	80
	Table 20: IDP Security Policy Rulebases	83
Chapter 9	Predefined Security Policies	87
	Table 21: Recommended Security Policy Definition	87
	Table 22: IDP Security Policy Templates	88
Chapter 10	The IDP Rulebase	91

	Table 23: IDP Rulebase Match Condition Guidelines	92
	Table 24: Predefined Attack Object Groups	99
	Table 25: Recommended Action by Attack Severity	101
	Table 26: IDP Rulebase Actions	101
	Table 27: IDP Rulebase IP Actions	103
	Table 28: IDP Rulebase Notification Options	104
Chapter 12	Application Policy Enforcement Rulebase	117
	Table 29: APE Rulebase Match Condition Guidelines	119
	Table 30: IDP Rulebase Actions	129
	Table 31: APE Rulebase Notification Options	131
Chapter 13	The Backdoor Rulebase	141
	Table 32: Backdoor Detection Runtime Parameters	142
	Table 33: Backdoor Rulebase Actions	145
	Table 34: Backdoor Rulebase Notification Options	145
Chapter 14	The SYN Protector Rulebase	149
	Table 35: SYN Protector Thresholds	150
	Table 36: SYN Flood Detection Runtime Parameters	151
	Table 37: SYN Protector Rulebase Modes	152
	Table 38: SYN Protector Rulebase Notification Options	154
Chapter 15	The Traffic Anomalies Rulebase	155
	Table 39: Traffic Anomalies Rulebase Detection Settings	156
	Table 40: Traffic Signature Runtime Settings	156
	Table 41: Traffic Anomalies Rulebase IP Actions	159
	Table 42: Traffic Anomalies Rulebase Notification Options	160
Chapter 16	The Network Honeypot Rulebase	161
	Table 43: Network Honeypot Rulebase IP Actions	163
	Table 44: Network Honeypot Rulebase Notification Options	164
Chapter 17	Fine-Tuning a Security Policy	167
	Table 45: Recommended Security Policy Definition	168
	Table 46: Actions to Take To Reduce False Positives	172
Part 4	Additional Deployment Topics	
Chapter 19	Inspection of Encapsulated and Encrypted Traffic	177
	Table 47: Supported SSL Cipher Suites	181

Preface

- Objectives on page xvii
- Audience on page xvii
- Documentation Conventions on page xvii
- Related Documentation on page xix
- Requesting Technical Support on page xx

Objectives

This guide explains and provides examples of Juniper Networks IDP Series Intrusion Detection and Prevention Appliance features.

For complete procedures for implementing IDP features and monitoring security events, see the *IDP Series Administration Guide*.

For details on using Juniper Networks Network and Security Manager (NSM) user interface features, see the NSM documentation.

Audience

This guide is intended for network administrators who are familiar with TCP/IP networks and network security issues.

Documentation Conventions

This section provides all the documentation conventions that are followed in this guide. Table 1 on page xviii 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 xviii 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 xix 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</i> , <i>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 }

Related Documentation

Table 4 on page xix lists related IDP Series documentation.

Table 4: Related IDP Series Documentation

Document	Description
Release notes	Contains information about what is included in a specific product release: supported features, unsupported features, changed features, known problems, and resolved problems. If the information in the release notes differs from the information found in the documentation set, follow the release notes.
IDP Detector Engine release notes	Provides information about IDP Detector Engine releases, including new features, changed features, fixed problems, and known issues.
J-Security Center Attack Signatures	Lists predefined attack signatures developed by J-Security Center.
J-Security Center Application Signatures	Lists predefined application signatures developed by J-Security Center.
IDP Series installation guides	Describes IDP Series hardware and provides instructions for installing, configuring, updating, and servicing the device.
IDP Series Feature Documentation	A collection of topics from the <i>IDP Series Administration Guide</i> and <i>IDP Series Concepts and Examples Guide</i> , in HTML.
IDP Series Administration Guide	Provides procedures for completing IDP Series administration tasks with the Network and Security Manager (NSM) central management program; with the IDP Series device Appliance Configuration Manager (ACM); and with the IDP Series device command-line interface (CLI).
IDP Series Concepts and Examples Guide	Explains IDP Series features and provides examples of how to use the system.

Table 4: Related IDP Series Documentation (*continued*)

Document	Description
<i>IDP Series Custom Attack Objects Reference and Examples Guide</i>	Provides examples and reference information for creating custom attack objects.
<i>IDP Reporter User's Guide</i>	Describes how to use IDP Reporter, an on-box reporting platform that includes predefined reports on attack detection and application usage. You can also use IDP Reporter to schedule regular publication of reports that are of interest to you or your stakeholders.

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/customers/support/downloads/710059.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 located at <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

Solution Overview

- IDP Series Product Overview on page 3
- IDP Series Components Overview on page 9

CHAPTER 1

IDP Series Product Overview

This chapter provides an overview of the intrusion detection and prevention (IDP) standalone solution and provides a documentation map of IDP features to IDP documentation sources. It includes the following topics:

- Juniper Networks IDP Solutions on page 3
- IDP Series Features Overview on page 3

Juniper Networks IDP Solutions

Juniper Networks provides intrusion detection services and intrusion detection and prevention (IDP) technology in the following product families:

- Juniper Networks IDP Series Intrusion Detection and Prevention Appliances
- Juniper Networks ISG Series Integrated Security Gateways
- Juniper Networks SRX Series Services Gateways

This guide describes the IDP Series appliances.

Related Documentation The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- IDP Series Features Overview on page 3

IDP Series Features Overview

Table 5 on page 3 briefly describes Juniper Networks IDP Series features.

Table 5: IDP Series Features

Feature	Description	Documentation
Application-Based Management		

Table 5: IDP Series Features (*continued*)

Feature	Description	Documentation
Application identification	<p>Port-independent application identification enhances both security and manageability by eliminating the need to manually and comprehensively configure application-port mapping for the service objects and application objects used in the IDP rulebase and APE rulebase rules.</p> <p>Beginning with IDP OS Release 5.1, the application identification feature can match extended application signatures used in APE rulebase rules. <i>Extended application</i> signatures are also called <i>nested application</i> signatures. The predefined extended application signatures developed for IDP OS Release 5.1 include the most prevalent Web 2.0 applications running over HTTP.</p>	<ul style="list-style-type: none"> Using Application Identification on page 96
User-defined application signatures	If the predefined signatures do not address all of your use cases, you can use the NSM Object Manager to create custom application signatures.	<ul style="list-style-type: none"> Using Application Objects on page 121
Application policy enforcement	The application policy enforcement (APE) rulebase enables you to mark, limit, or drop traffic that matches application signatures.	<ul style="list-style-type: none"> Understanding the APE Rulebase on page 117
Application volume tracking	The application volume tracking (AVT) feature leverages Profiler functionality to collect statistics about application usage.	<ul style="list-style-type: none"> Application Volume Tracking Overview on page 45
Intrusion Detection and Prevention		
Multimethod attack detection	The IDP Series uses eight methods to detect malicious traffic.	See Table 6 on page 6 for a description of detection methods.
Zero-day protection	The IDP rulebase attack objects detect protocol usages that violate published RFCs. Protocol anomaly detection protects your network from undiscovered vulnerabilities.	<ul style="list-style-type: none"> J-Security Center Updates Overview on page 21
Protocol decoding	Juniper Networks Security Center (J-Security Center) provides a robust protocol detection engine that can decode more than 60 protocols and analyze and enforce proper usage in more than 500 contexts.	<ul style="list-style-type: none"> J-Security Center Updates Overview on page 21
Recommended security policy and predefined attack objects	<p>J-Security Center provides a robust default security policy (called Recommended) and a comprehensive set of predefined attack objects (including those flagged as Recommended for various categories of attacks).</p> <p>The J-Security Center attack database includes more than 5500 signatures for identifying anomalies, attacks, spyware, and applications.</p>	<ul style="list-style-type: none"> Using the Recommended Security Policy on page 87 Using Attack Objects on page 97

Table 5: IDP Series Features (*continued*)

Feature	Description	Documentation
User-defined security policies and attack objects	<p>If you choose, you can use the default security policy or other predefined templates as a basis for your own user-defined security policy.</p> <p>Similarly, you can use the predefined attack objects as a basis for your own user-defined attack objects.</p>	<ul style="list-style-type: none"> Understanding the Components of an IDP Security Policy on page 83 Using Attack Objects on page 97
Active response methods	<p>J-Security Center attack objects are coded with recommended actions to take on the instant session, including drop packet, drop connection, close client, close server, and close client/server. You can rely on these or set your own.</p> <p>In addition, when the IDP Series device detects an attack from a particular IP address, it can block connections from the IP address for a configurable duration of time.</p>	<ul style="list-style-type: none"> Understanding IDP Rulebase Actions on page 101
Passive response methods	The IDP Series supports several passive responses, including logging and TCP reset.	<ul style="list-style-type: none"> Understanding IDP Rulebase Notification Options on page 103
Traffic decryption and decapsulation	The IDP Series can decrypt or decapsulate traffic and then inspect the payload. We support decryption of SSL and decapsulation of GRE, GTP, IPsec ESP NULL, and MPLS traffic.	<ul style="list-style-type: none"> Inspection of SSL Traffic Overview on page 179 Inspection of GRE Traffic Overview on page 177 Inspection of GTP Traffic Overview on page 177 Inspection of IPsec VPN Traffic Overview on page 178 Inspection of MPLS Traffic Overview on page 178
Centralized Management and Logging		
Centralized management	The IDP Series is compatible with Juniper Networks Network and Security Manager (NSM).	<ul style="list-style-type: none"> Centralized Management with NSM Overview on page 20
Network profiling	The Profiler captures accurate and granular detail of your network traffic.	<ul style="list-style-type: none"> Profiler Overview on page 31
Robust logging, reporting, and notification	The IDP Series includes useful predefined log views and reports and enables you to create custom views and reports.	<ul style="list-style-type: none"> IDP Logs Overview on page 53 NSM Reports Overview on page 74 IDP Reporter Overview on page 76
Simulation Mode, Autorecovery, Bypass, and Failover		
Simulation mode	In simulation mode, the IDP Series inspects traffic according to your security policy but only simulates policy actions, generating logs of the action dictated by the security policy. You can use simulation mode when you first adopt the IDP Series solution to learn expected behavior without risk of traffic disruption.	<ul style="list-style-type: none"> Simulation Mode Overview on page 25

Table 5: IDP Series Features (*continued*)

Feature	Description	Documentation
Autorecovery	If an IDP process engine experiences failure, the IDP Series device buffers the next packets in the flow and restarts the process engine.	<ul style="list-style-type: none"> Auto-Recovery Feature on page 11
Bypass	<p>You can configure network interfaces to enter a bypass state in case of failure or graceful shutdown, or if the JNET driver encounters problems processing packets.</p> <p>The IDP Series also supports flow bypass to forward traffic when traffic exceeds IDP session capacity.</p>	<ul style="list-style-type: none"> Internal Bypass on page 16 External Bypass on page 17 Flow Bypass Feature on page 12
High availability	Feature set that operates well with third-party high availability solutions where you have deployed redundant network paths and use the failure detection features of a firewall, router, or switch to manage the cutover from the primary path to the backup path in cases of failure.	<ul style="list-style-type: none"> IDP Series Deployment Scenarios
Compatibility with Juniper Networks Access Solutions		
User role-based policies	When integrated with Juniper Networks IC Series Unified Access Control (UAC) appliance, the IDP Series appliance supports security policy rules based on UAC user roles. This feature enables you to more easily configure focused rules to implement your business security policy.	<ul style="list-style-type: none"> IDP Series Deployment Scenarios
Coordinated threat control	The IDP Series is compatible with Juniper Networks SA Series SSL VPN appliances and IC Series devices. The SA Series and IC Series devices can “subscribe” to IDP Series logs and use the logs as a basis for access rules.	<ul style="list-style-type: none"> IDP Series Deployment Scenarios

Table 6 on page 6 briefly describes IDP detection methods and provides a reference to detailed information.

Table 6: Intrusion Detection Methods

Feature	Description	Documentation
Stateful signature	The IDP rulebase attack object signatures are bound to protocol context. As a result, this detection method produces few false positives.	<ul style="list-style-type: none"> Understanding the IDP Rulebase on page 91 Using Attack Objects on page 97
Protocol anomaly	The IDP rulebase attack objects detect protocol usages that violate published RFCs. This method protects your network from undiscovered vulnerabilities.	<ul style="list-style-type: none"> Understanding the IDP Rulebase on page 91 Using Attack Objects on page 97

Table 6: Intrusion Detection Methods (*continued*)

Feature	Description	Documentation
Traffic anomaly	The Traffic Anomalies rulebase uses heuristic rules to detect unexpected traffic patterns that might indicate reconnaissance or attacks. This method blocks distributed denial-of-service (DDoS) attacks and prevents reconnaissance activities.	<ul style="list-style-type: none"> Understanding the Traffic Anomalies Rulebase on page 155
Backdoor	The Backdoor rulebase uses heuristic-based anomalous traffic patterns and packet analysis to detect Trojans and rootkits. These methods prevent proliferation of malware in case other security measures have been compromised.	<ul style="list-style-type: none"> Understanding the Backdoor Rulebase on page 141
IP spoofing	The IDP Series device checks the validity of allowed addresses inside and outside the network, permitting only authentic traffic and blocking traffic with a disguised source.	<ul style="list-style-type: none"> IP Spoof Attack Prevention Overview on page 173
Denial of service (DoS)	The SYN Protector rulebase provides two, alternative methods to prevent SYN-flood attacks.	<ul style="list-style-type: none"> Understanding the SYN Protector Rulebase on page 149
Network honeypot	The IDP Series device impersonates vulnerable ports so you can track attacker reconnaissance activity.	<ul style="list-style-type: none"> Understanding the Network Honeypot Rulebase on page 161

Related Documentation The following additional related topic is included in the *IDP Series Concepts and Examples Guide*:

- Juniper Networks IDP Solutions on page 3

CHAPTER 2

IDP Series Components Overview

This chapter provides an overview of IDP Series components. It includes the following topics:

- IDP Series Operating System Overview on page 9
- IDP Series Network Interfaces Overview on page 13
- Centralized Management with NSM Overview on page 20
- J-Security Center Updates Overview on page 21

IDP Series Operating System Overview

The following topics explain the features of the IDP Series operating system:

- IDP Series Multicore Architecture on page 9
- Auto-Recovery Feature on page 11
- Flow Bypass Feature on page 12
- Key Processes on page 12

IDP Series Multicore Architecture

The IDP Series operating system separates control plane and data plane processes, making the IDP Series devices resilient to periods of heavy load, such as a denial-of-service attack (DoS attack) or distributed denial-of-service attack (DDoS attack). In high-end platforms, control plane and data plane processes run on separate CPU, bolstering resiliency and performance. Figure 1 on page 10 shows the processes that run on each core CPU for IDP Series platforms.

Figure 1: IDP Multicore Architecture

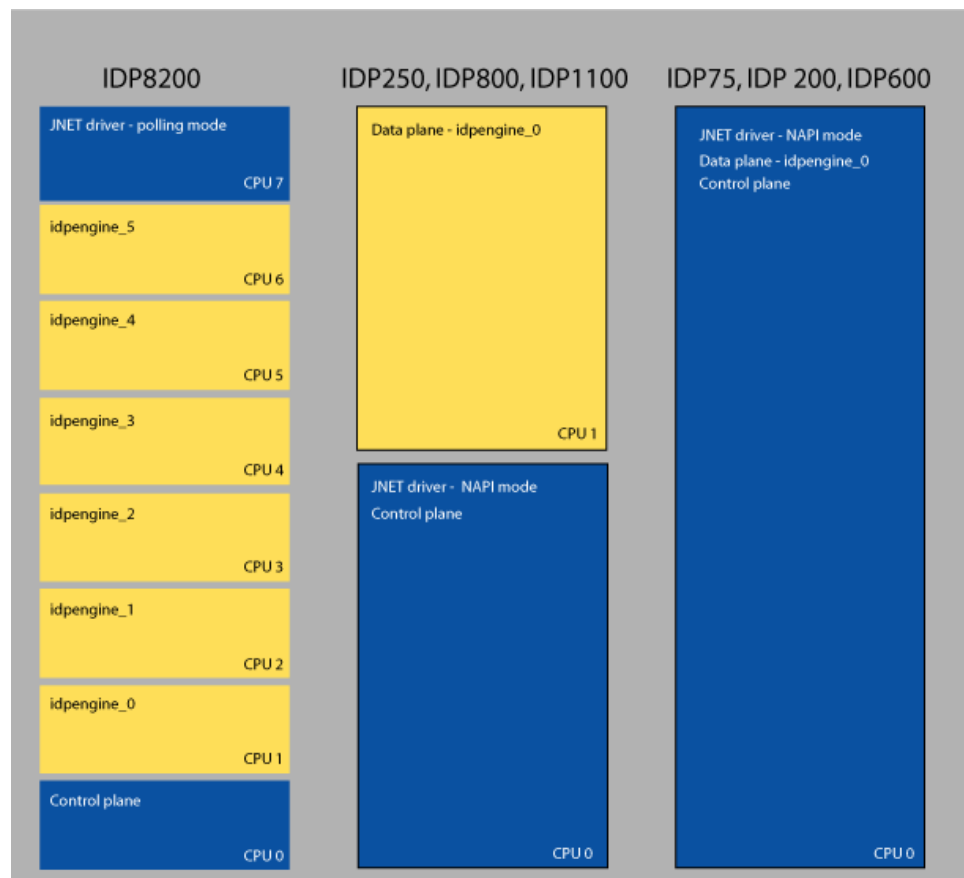


Table 7 on page 10 describes how CPUs are dedicated on IDP Series platforms.

Table 7: IDP Multicore Architecture

Platform	Multicore Architecture
IDP8200	<p>The IDP8200 appliance has eight core CPUs:</p> <ul style="list-style-type: none"> • One CPU is dedicated to control plane processes, including configuration and logging processes. • One CPU is dedicated to the JNET driver processes, which handle traffic transmission and buffering. The JNET driver runs in polling mode. • The remaining CPUs are dedicated to data plane processes, including the IDP engine processes that inspect traffic and take action against attacks. <p>NOTE: When you use the scio or sctop utilities to monitor IDP processes, you can specify the -c CPU option to display statistics related to processing on a particular IDP engine. For example, to display CPU utilization for CPU 2, the syntax is scio -c 2 idp-cpu-utilization. Without the -c option, scio displays an aggregate for all IDP engines.</p>

Table 7: IDP Multicore Architecture (*continued*)

Platform	Multicore Architecture
IDP1100, IDP800, IDP250	<p>The IDP1100, IDP800, and IDP250 appliances have two core CPUs:</p> <ul style="list-style-type: none"> One CPU is used for both control plane and JNET driver processes. The JNET driver runs in NAPI mode: when traffic is low, the JNET driver operates in interrupt mode; when traffic is high, the JNET driver switches to polling mode. The second CPU is dedicated to data plane processes.
IDP600, IDP200, IDP75	<p>The IDP600, IDP200, and IDP75 appliances have one core CPU. The JNET driver runs in NAPI mode: when traffic is low, the JNET driver operates in interrupt mode; when traffic is high, the JNET driver switches to polling mode.</p>



NOTE: If you use port monitoring utilities such as the Linux `lsof` command to view port activity, you will notice activity by host 127.0.0.1 (example below). These entries identify internal system communication. Communication on port 9101 and above and activity identifying Bacula processes is related to the internal communication between the control plane and data plane. This activity is essential to proper functioning of the IDP OS.

```
[root@idp-75-172-22-151-70 ~]# lsof -n -i
COMMAND    PID  USER  FD  TYPE DEVICE SIZE MODE NAME
idpengine  3164 root   4u  IPv4 39333      TCP *:bacula-dir (LISTEN)
idpengine  3164 root   5u  IPv4 40576      TCP
127.0.0.1:bacula-dir->127.0.0.1:39472 (ESTABLISHED)
idpengine  3164 root   6u  IPv4 39336      TCP 127.0.0.1:50000
(LISTEN)
idpengine  3164 root   7u  IPv4 302480     TCP
127.0.0.1:50000->127.0.0.1:37684 (ESTABLISHED)
idpengine  3164 root   9u  IPv4 302372     TCP
127.0.0.1:bacula-dir->127.0.0.1:59114 (ESTABLISHED)
..
```

Auto-Recovery Feature

The auto-recovery feature monitors the status of individual IDP engines. In the event an IDP engine is terminated, the auto-recovery daemon logs the event, attempts to restart the IDP engine, and logs recovery. The auto-recovery feature is enabled by default.

The auto-recovery feature bolsters resiliency of IDP Series devices by enabling restart per IDP engine. In case of failure by an IDP engine, the other IDP engines in the data plane do not need to be restarted. If an IDP engine becomes overloaded or terminates, the JNET driver buffers packets for the active sessions while the IDP engine restarts. If the IDP engine does not restart after six attempts, the IDP Series device may enter internal bypass (if enabled).

The auto-recovery feature is complemented by the bypass under congestion feature. When the IDP engine recovers, it processes the buffered packets. If the buffer becomes large enough, it triggers bypass under congestion instead of resulting in subsequent failure.



NOTE: The auto-recovery process ensures the IDP engine restarts with the same device configuration, feature configuration, and security policy that were in place before the restart. However, because the application identification feature uses the first packets of a session as a key to determining the application, the auto-recovery process cannot reliably identify the application for buffered sessions. As a result, in processing buffered traffic, the application identification feature is unavailable and application rate limiting cannot be applied. In addition, the latest interval of application volume tracking data is discarded.

Flow Bypass Feature

The flow bypass feature prevents the IDP Series device from becoming a point of failure when the network is congested. With flow bypass enabled, when the IDP system packet receive queue reaches a rising threshold that you specify, the IDP engine marks the flow as a bypass flow and passes it through, uninspected. The IDP Series device passes through subsequent flows until the IDP system receive queue falls below a reset threshold that you also specify. On IDP8200, which has multiple IDP engines, the flow bypass feature operates per IDP engine; that is, when the IDP packet receive queue for an individual IDP engine reaches its rising threshold, the individual IDP engine takes the flow bypass action and other IDP engines continue to inspect flows.

The flow bypass feature is disabled by default. It is intended for use in networks that prioritize availability over security.

For procedures for enabling the flow bypass feature, configuring its thresholds, and monitoring the feature, see the *IDP Series Administration Guide*.

Key Processes

Table 8 on page 12 describes the key processes.

Table 8: Processes

Process	Description
agent	Manages communication between the IDP Series device and Network and Security Manager.
idpengine	Performs packet processing, including decapsulation or decryption, defragmentation, reassembly, inspection, and rule actions.
idpHMD	Generates SNMP alerts when thresholds are crossed for tracked resources on the device. Responds to SNMP poll requests. Resources are CPU, memory, hard disk space, and session count.
idpLogReader	Gathers logs generated by idpengine and stores the information in a local log database. The agent process forwards the data to NSM, where you can view records.
pkid	Inspects SSL traffic, if SSL inspection is turned on.
profiler	Gathers information about hosts and applications in your network. Profiler stores the information in the Profiler database. The agent process forwards the data to NSM, where you can view records.

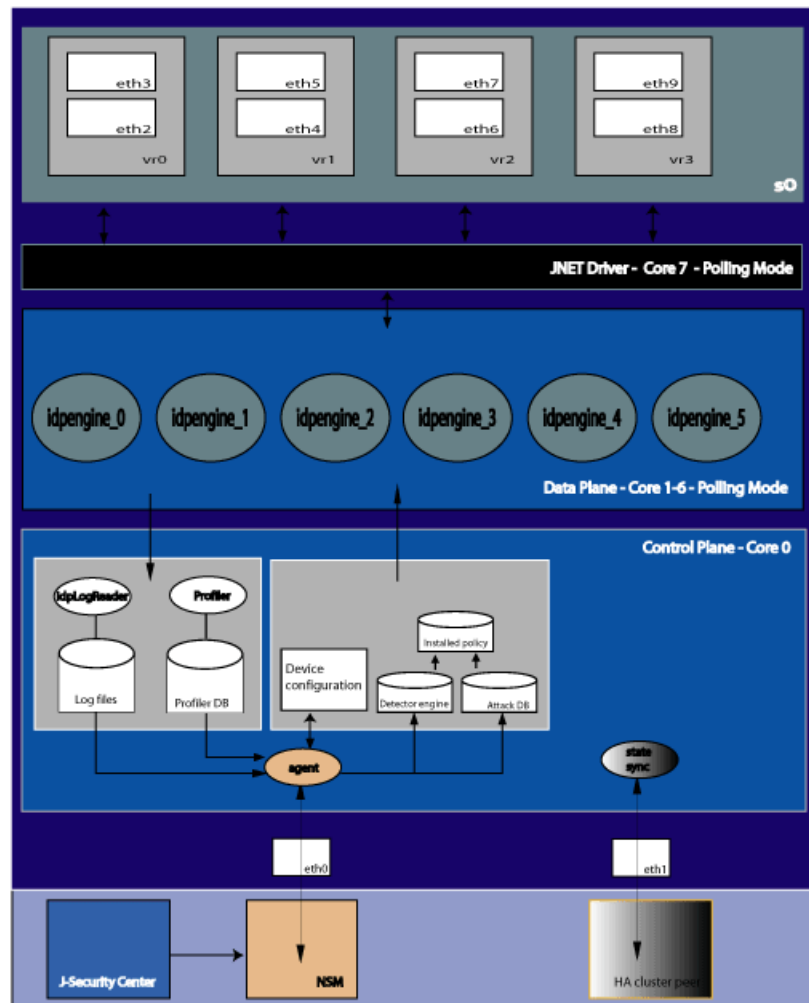
Table 8: Processes (*continued*)

Process	Description
sciod	Handles policy push, information retrieval, Profiler status, and so on.
Related Documentation	<p>The following related topic is included in the <i>IDP Series Concepts and Examples Guide</i>:</p> <ul style="list-style-type: none">• IDP Series Network Interfaces Overview on page 13 <p>The following related topics are included in the <i>IDP Series Administration Guide</i>:</p> <ul style="list-style-type: none">• Viewing Auto-Recovery Logs• Enabling the Flow Bypass Feature

IDP Series Network Interfaces Overview

In Figure 2 on page 14, eth0, eth1, eth2, eth3, and so forth are the network interfaces.

Figure 2: IDP Series Network Interfaces



The following topics explain the features of these network interfaces:

- Management Interface (eth0) on page 15
- High Availability Interface (eth1) on page 15
- Traffic Interfaces on page 15
- Internal Bypass on page 16
- External Bypass on page 17
- Interface Signaling on page 18
- Peer Port Modulation on page 18

Management Interface (eth0)

In Figure 2 on page 14, eth0 is a dedicated management interface used for communication with Network and Security Manager (NSM). The agent process is a control plane process. It manages communication between the IDP Series device and NSM. The agent process handles the following functionality:

- **Device configuration**—You set part of the active configuration with the Appliance Configuration Manager (ACM), part with the CLI, and part with NSM. The agent process pushes changes you make from NSM to the IDP Series device.
- **Security policy**—You configure policies with NSM. You push a single policy to the IDP Series device to be installed and used by the IDP process engines. The installed policy is the policy used to determine which traffic the IDP engine inspects, what to look for, and what actions to take.
- **Detector engine**—The IDP detector engine is a code base that contains the application signatures and protocol decoder definitions used by the IDP engine in packet analysis. J-Security Center periodically updates the IDP detector engine. In Figure 2 on page 14, note the process flow: first, you download updates from J-Security Center to NSM; then, you push updates from NSM to IDP Series devices.
- **Attack database**—The attack database includes the attack objects used by the IDP rulebase to match attack signatures and protocol anomalies. J-Security Center updates predefined attack object definitions as often as necessary. As with detector engine updates, you download them from J-Security Center to NSM and then push them from NSM to the IDP Series device.
- **Logging**—The IDP process engines generate logs and packet captures related to security policy and application policy enforcement rules. The Profiler generates profiling and application volume logs. The agent process sends these logs to NSM so you can use NSM monitoring features to monitor security events and application usage.

High Availability Interface (eth1)

In Figure 2 on page 14, eth1 is a dedicated high availability (HA) interface used for sync-state communication with a cluster peer in a high availability deployment.

Traffic Interfaces

In Figure 2 on page 14, eth2, eth3, and so forth are the network interfaces you connect to the network devices that route traffic in your network.

The IDP Series implements the following abstract objects to manage network interfaces:

- **Virtual circuit**—A virtual circuit corresponds with the physical interface. For example, physical interface eth2 is a virtual circuit. You use the Appliance Configuration Manager (ACM) to configure speed and duplex, as well as optional interface alias settings for each interface.
- **Virtual router**—A virtual router contains a logical pair of virtual circuits. For example, virtual router vr0 contains eth2 and eth3. In transparent mode, traffic arrives in one interface and is forwarded through the other. You use ACM to configure the deployment

mode (sniffer or transparent) and bypass options (internal, external, or off) for each virtual router. You can use the command-line interface to display information and status for each virtual router, including Address Resolution Protocol (ARP) and media access control (MAC) tables.

- **Subscriber**—A single subscriber named s0 contains all virtual routers. The subscriber maintains process and status of all traffic that flows through the device. You can use the command-line interface to view information and status maintained by subscriber s0. We test and support only configurations where the default subscriber is used.

Internal Bypass

The Internal Bypass feature is intended for deployments where a network security policy privileges availability over security. In the event of failure or graceful shutdown, traffic bypasses the IDP processing engine and is passed through the IDP Series device uninspected.

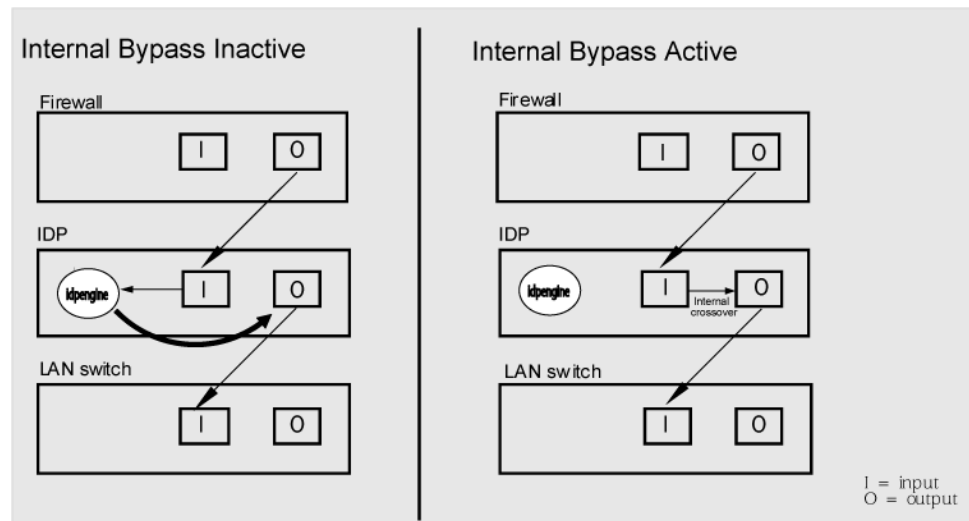
The Internal Bypass feature operates through a timing mechanism. When enabled, the timer on traffic interfaces counts down to a bypass trigger point. When the IDP Series appliance is turned on and available, it sends a reset signal to the traffic interface timer so that it does not reach the bypass trigger point. If the IDP OS encounters failure, then it fails to send the reset signal, the timer counts down to the trigger point, and the traffic interfaces enter a bypass state. If the IDP Series appliance is shut down gracefully, the traffic interfaces immediately enter bypass.

With copper NICs, the bypass mechanism joins the interfaces mechanically to form a circuit that bypasses IDP processing. Packets traverse the IDP Series device as if the path from eth2 (receiving interface) to eth3 (transmitting interface) were a crossover cable. No packet inspection or processing occurs.

With fiber NICs, the bypass mechanism uses optical relays instead of copper relays. During normal operations, the optical relays send light to the built-in optical transceivers. When bypass is triggered, the relays flip state, and the light signal is redirected to optically connect the two external ports.

Figure 3 on page 17 compares the data path when Internal Bypass is enabled but not activated with the data path when Internal Bypass is activated.

Figure 3: Internal Bypass

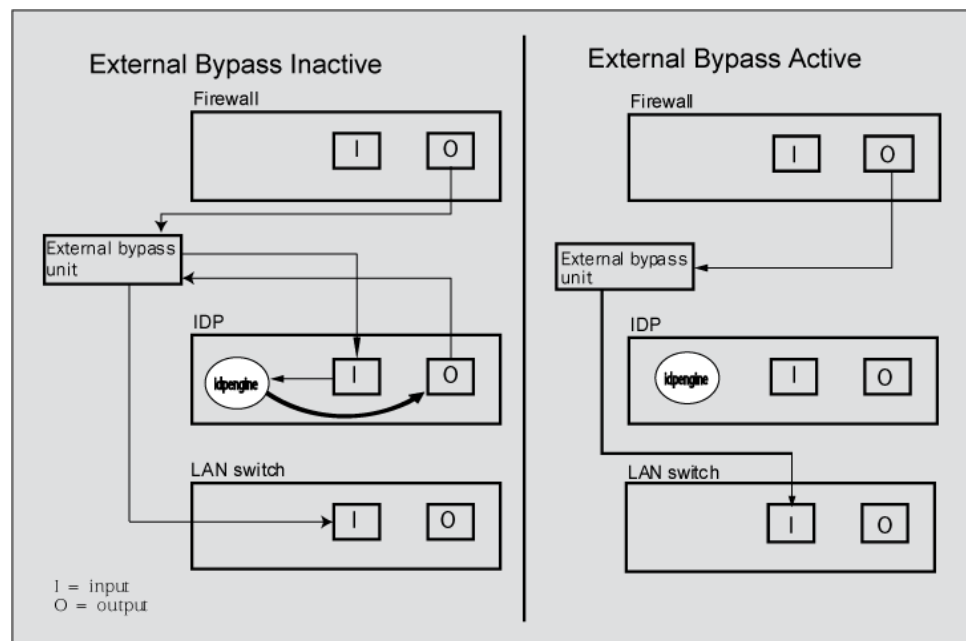


When the IDP operating system resumes healthy operations, it sends a reset signal to the traffic interfaces, and the interfaces resume normal operation.

External Bypass

The External Bypass setting supports third-party external bypass units. Deployments with external bypass units depend on the functionality of the external bypass unit to check the status of the IDP Series appliance and make the determination whether to send packets through or around the IDP Series device. Most external bypass units test for availability by sending heartbeat packets through the device. If the packets reach the expected destination, the external bypass unit allows the traffic to continue through the IDP Series appliance. If the packets fail to reach the expected destination, the external bypass unit determines the IDP Series is unavailable, so it forwards traffic around the IDP Series device. The IDP Series supports external bypass solutions by allowing the heartbeat traffic to pass through the device regardless of the Layer 2 Bypass setting. In other words, if you disable Layer 2 Bypass and enable External Bypass, most Layer 2 traffic will be dropped but the heartbeat traffic used in the external bypass deployment will be passed through. Figure 4 on page 18 compares the data path when External Bypass is enabled but not activated with the data path when External Bypass is activated.

Figure 4: External Bypass



Interface Signaling

The interface signaling feature supports high-availability deployments where there are redundant network paths, and a firewall, router, or switch chooses the active path. The interface signaling script monitors the state of the following IDP Series components:

- Traffic interfaces (eth2, eth3, and so on). In case of interface failure, the script brings down all peer interfaces so that a third-party link detection mechanism can properly detect failure.
- IDP engines (idpengine0, idpengine1, and so on). In case of IDP engine failure, the auto-recovery feature attempts to restart the IDP engine. If the IDP engine cannot be restarted after six attempts, the auto-recovery process runs an **idp.sh stop** command. The interface signaling script then brings down all traffic interfaces.

After bringing down the peer interfaces, the interface signaling script sleeps for 30 seconds to avoid link flapping issues. After 30 seconds, the script checks the state of the IDP engine or interface that had encountered the failure. When the underlying problem has been resolved and the interface is up, the interface signaling script brings up the peer interfaces.


Peer Port Modulation

The peer port modulation (PPM) feature supports deployments where routers monitor link state to make routing decisions. In these deployments, a router might be set to monitor link state on only one side of the IDP Series device. Suppose, for example, the router monitors only the IDP inbound interface. Suppose the inbound interface remains up but the outbound interface goes down. The router watching the inbound link would detect an available link and forward traffic to the IDP Series device. Traffic would be dropped

at the point of failure—the outbound link. PPM propagates a link loss state for one traffic interface to all interfaces in the IDP virtual router.

When PPM is enabled, a PPM daemon monitors the health of IDP traffic interfaces belonging to the same virtual router. If a traffic interface loses link, the PPM process turns off any associated network interfaces in the same virtual router so that other network devices detect that the virtual router is down and route around it. For example, assume you have enabled PPM and configured IDP virtual routers as shown in Figure 5 on page 19.

Figure 5: Peer Port Modulation



Suppose there is a network problem and eth3 goes down. The PPM daemon detects this and turns off the other interface in vr0: eth2. The interfaces in vr1, vr2, and vr3 are unaffected. After you fix the problem with eth3, the PPM daemon detects this, and turns on eth2.



NOTE: The PPM feature is independent of the bypass feature (NIC state setting). PPM is related to the *status of the link*, not the status of the IDP operating system. A link can be down even when the IDP operating system is healthy. Note, however, that PPM runs as a control plane process and operates only when the IDP Series device is turned on and the control plane is available. If the IDP operating system is unavailable, the PPM feature is also unavailable, regardless of the setting for the NIC state.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- IDP Series Operating System Overview on page 9
- Centralized Management with NSM Overview on page 20

The following related topics are included in the *IDP Series Administration Guide*:


- Configuring Virtual Routers (ACM Procedure)
- Tuning the JNET Driver Failure Count
- Configuring Interface Aliasing (ACM Procedure)

Centralized Management with NSM Overview

Juniper Networks Network and Security Manager (NSM) is a central management server capable of managing hundreds of IDP Series devices and other Juniper Networks devices, such as ScreenOS firewalls, SA Series devices, and IC Series devices. You typically deploy NSM in a management subnet accessible to the NSM-managed devices.

Figure 6 on page 20 illustrates the flow of information between the tiers of the central management solution: the NSM user interface, the NSM server, and IDP Series devices.

Figure 6: IDP-NSM Communication



The IDP Series configuration, security policies, attack objects, and log records are stored in NSM server databases and administered using the NSM user interface. Communication between the NSM server and IDP Series devices, and between the NSM server and the NSM user interface, is encrypted and authenticated.

For IDP Series deployments, centralized management provides the following benefits:

- Centralized management for IDP Series devices and other network devices
- Consolidated logs from different devices in a single repository
- Centralized management of enterprise security policies
- Simplified management for attack signature updates
- Role-based administration

For information about installing NSM and using NSM distributed management features, management objects (such as address objects, service objects, and templates), and navigational and display features, see the NSM documentation.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- J-Security Center Updates Overview on page 21
- Management Interface (eth0) on page 15

The following related topics are included in the *IDP Series Administration Guide*:

- NSM Device Configuration Management Task Summary
- IDP Series Logs and Reports in NSM Task Summary

J-Security Center Updates Overview

The Juniper Networks Security Center (J-Security Center) routinely makes important updates available to IDP security policy components:

- **Detector engine.** 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 IDP, whenever you upgrade, and whenever alerted to do so by Juniper Networks. You can view release notes for detector engine updates at <http://www.juniper.net/techpubs/software/management/idp/de/>.
- **Attack database.** The [attack signature database](#) stores data definitions for attack objects. Attack objects are patterns comprising stateful signatures and traffic anomalies. You specify attack objects in IDP rulebase rules.
- **Application signature database.** The [application signature database](#) stores data definitions for application objects. Application objects are patterns used to identify applications and match APE rulebase rules.

J-Security Center updates are packaged and released separately from the IDP operating system and software code base to ensure IDP products protect your network against recently discovered vulnerabilities. We recommend you schedule automatic updates for the attack database and application database. For IDP Series devices, both databases are distributed in “signature database updates”.

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



NOTE: We recommend you subscribe to the IDP Signature Updates technical bulletin to be notified when J-Security Center releases IDP detector engine updates. Go to <https://www.juniper.net/alerts/>.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Centralized Management with NSM Overview on page 20
- Using Attack Objects on page 97
- Using Application Objects on page 121

The following related topics are included in the *IDP Series Administration Guide*:

- Attack Objects Task Summary
- Application Objects Task Summary
- Loading J-Security Center Updates (NSM Procedure)

PART 2

Analyzing Your Network

- [Simulation Mode on page 25](#)
- [Profiler on page 31](#)
- [Security Explorer Overview on page 43](#)
- [Application Volume Tracking on page 45](#)
- [Logs and Reports on page 53](#)

CHAPTER 3

Simulation Mode

This chapter provides an overview of simulation mode and provides example of when you might want to enable simulation mode. It includes the following topics:

- Simulation Mode Overview on page 25
- Example: Getting Started with Simulation Mode on page 27
- Example: Using Simulation Mode to Maximize Uptime on page 29

Simulation Mode Overview

Simulation mode is not a deployment mode, but rather an operational mode. The following sections give an overview of simulation mode:

- Topology on page 25
- Purpose on page 25
- Configuration Overview on page 26
- Logging on page 26

Topology

The purpose of simulation mode is to enable you to evaluate expected results when you deploy the IDP Series device in transparent mode or sniffer mode. Therefore, in your network topology, you install and connect the IDP Series device where you intend to deploy it in transparent (in-path) or sniffer mode (out-of-path).

Purpose

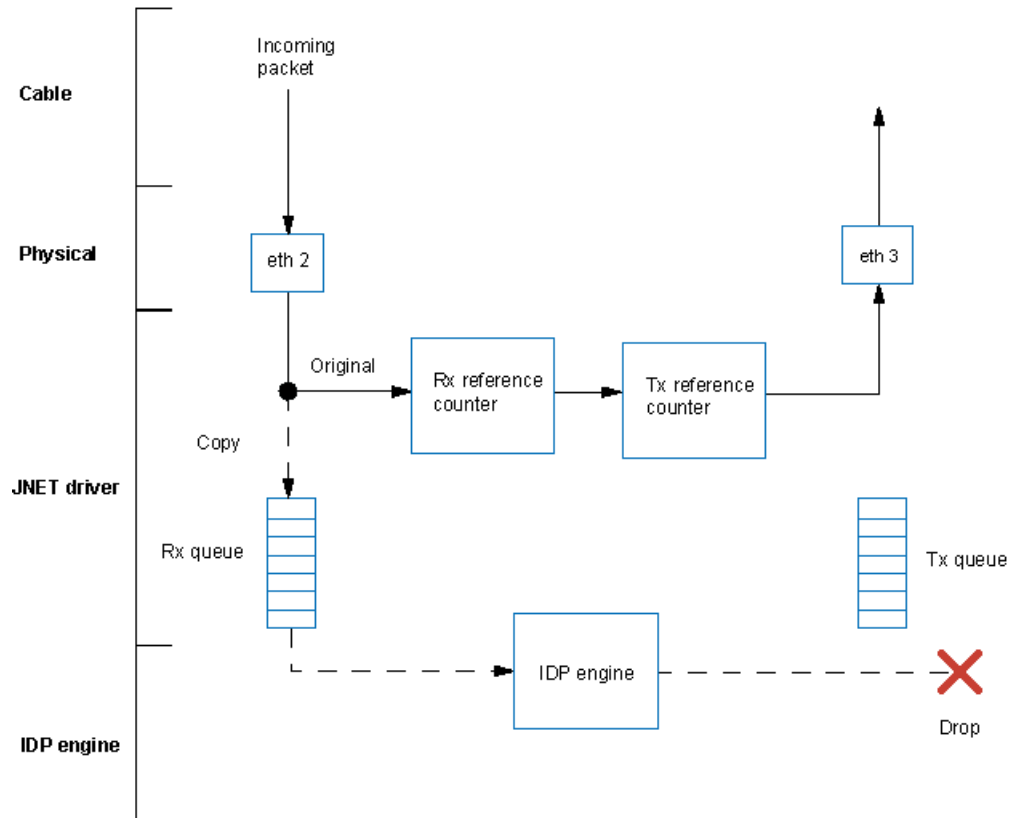
You operate an IDP Series device in simulation mode in the following situations:

- When you first deploy the IDP Series device in your network and you want to evaluate the security actions it takes without disrupting traffic.
- When you implement a new feature or change a security policy and you want to evaluate the impact without disrupting traffic.
- As a workaround to avoid traffic outages when IDP processing is resulting in crashes and other failures.

In simulation mode, when the IDP Series device receives a packet, it makes a copy. It transmits the original packet uninspected through the egress interface and enqueues

the duplicate packet into the JNET driver receive queue to be processed by the IDP engine. The IDP engine inspects the traffic against your security policy rules and implicit rules, and it generates logs when rules match. The IDP engine then drops the copy of the packet. Figure 7 on page 26 illustrates packet processing in simulation mode.

Figure 7: Packet Processing in Simulation Mode



NOTE: Because of packet queueing, when simulation mode is turned on, a few packets that are queued for processing and forwarding might be dropped. This results in retransmission depending on Layer 4 or Layer 7 behavior. When simulation mode is turned off, a few duplicate packets might be forwarded.

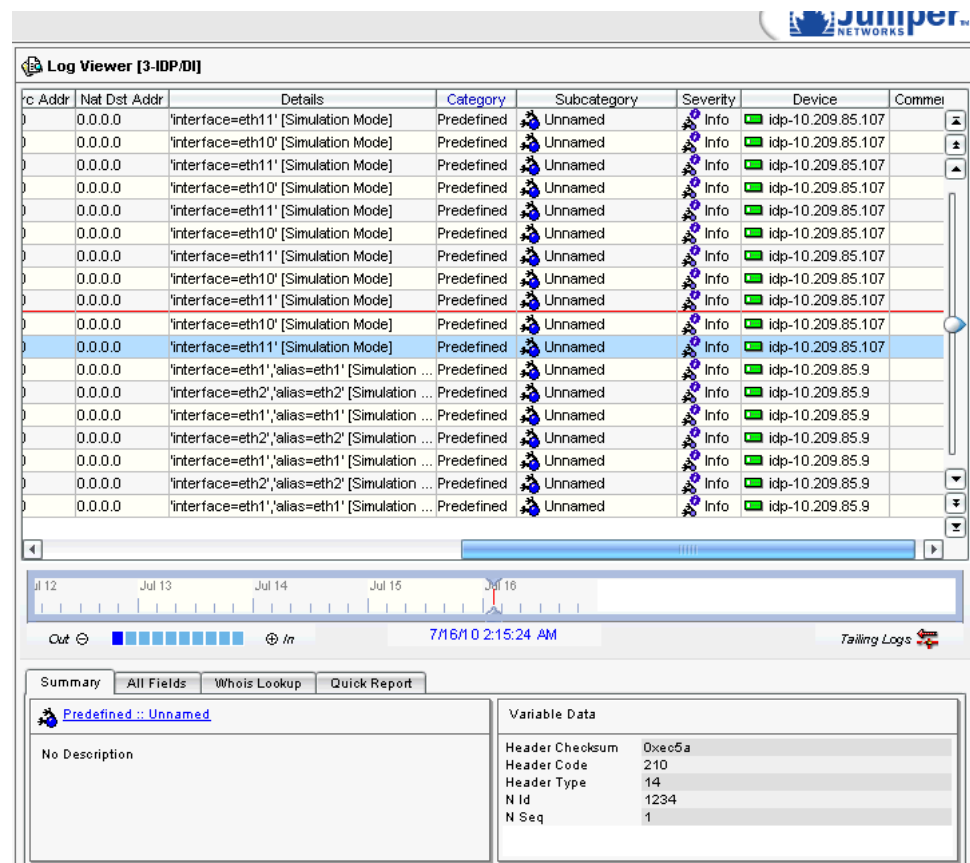
Configuration Overview

You use the CLI to enable or disable simulation mode. Simulation mode is disabled by default. You do not need to restart the IDP engine (idp.sh) or push a policy to enable or disable simulation mode.

Logging

In logs, the string [Simulation Mode] appears in the Details column, along with the details of the event. Figure 8 on page 27 shows a simulation mode log in the NSM log viewer. You can use NSM log and report filters to create log views and reports that filter for (or filter out) simulation mode logs.

Figure 8: NSM Log Viewer: Simulation Mode Logs



Related Documentation

The following related topics are included in *IDP Series Deployment Scenarios*:

- Sniffer Mode Overview
- Transparent Mode Overview

The following related topics are included in the *IDP Series Concepts and Examples Guide*

- Example: Getting Started with Simulation Mode on page 27
- Example: Using Simulation Mode to Maximize Uptime on page 29

The following related topic is included in the *IDP Series Administration Guide*

- Simulation Mode Task Summary

Example: Getting Started with Simulation Mode

The primary use case for simulation mode is when you are evaluating the effectiveness of the IDP Series device as the intrusion prevention system for your network.

Follow these basic steps to get started:

1. Read the release notes for your release. The release notes contain important release-related information about release-specific features, unsupported features, changed features, fixed issues, and known issues. The information in the release notes is more current than the information in this guide.
2. Install the IDP Series appliance, connect the management interface to your network, configure network settings, and configure virtual routers in transparent mode (in-path) or sniffer mode (out-of-path). For details, see the installation guide for your IDP Series device.
3. Upgrade IDP Series software to the latest version (if applicable).
4. Add the IDP Series device to the NSM Device Manager.
5. Update the IDP detector engine and NSM attack object database.
6. Become familiar with the default security policy (named Recommended).
7. Use the command-line interface to enable simulation mode.
8. Connect transit interfaces to the firewall and/or switch. See the installation guide for your IDP Series device.
9. Use the documentation to become familiar with the product features and user interface:
 - Use the *IDP Series Concepts and Examples Guide* to become familiar with IDP Series features.
 - Use this guide, the *IDP Series Administration Guide*, to learn the steps to implement IDP Series features and monitor security events.
 - Use the Appliance Configuration Manager (ACM) online Help for information about using ACM.
 - Use CLI man pages for syntax and parameter hints for CLI commands.
 - Use the NSM online Help for information about using the NSM user interface.
10. Run Profiler to discover the network hosts you want to protect.
11. Review logs to verify the initial deployment.
12. Fine-tune your security policy.

**Related
Documentation**

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Simulation Mode Overview on page 25
- Example: Fine-Tuning a Security Policy on page 167

The following related topic is included in the *IDP Series Administration Guide*:

- Updating IDP OS Software (NSM Procedure)
- Adding IDP Series Devices to NSM Device Manager

- Loading J-Security Center Updates (NSM Procedure)
- Developing Security Policies Task Summary
- Enabling Simulation Mode
- Profiler Task Summary
- IDP Series Logs and Reports in NSM Task Summary

Example: Using Simulation Mode to Maximize Uptime

The primary use case for simulation mode is for evaluating whether to adopt the IDP Series device as the intrusion prevention device for your network. You might also find simulation mode useful after you adopt the IDP Series as your IDP solution when you want to maximize network availability while you tune a security policy update or troubleshoot traffic outages when IDP processing results in crashes.

Suppose you discover that the device is dropping traffic and that early indicators suggest a likely false positive and that the traffic probably can be trusted. This situation might happen, for example, after an attack object database update when new attack signatures are added to a dynamic attack group that is specified in your IDP rulebase rules.

In cases like this, your choices are:

- Continue to drop traffic while you investigate.
- Change the rule action to allow traffic while you investigate. This requires you to reload the security policy with the changed rule action.
- Shut down the IDP Series device while you investigate. If you enable internal bypass, traffic passes through the device.
- Use simulation mode.

In cases like this, simulation mode is a good choice if you are an experienced IDP security administrator who suspects a false positive and are inclined to maximize uptime while you investigate. If you later conclude that it is not a false positive, you can disable simulation mode and return to active management without having to reload the security policy. You can use the logs collected during simulation mode to follow up on any subsequent security actions to take. If, on the other hand, your investigation confirms your hunch that it is a false positive, you can make iterations of modifications to your policy, load the changed policy, and observe the results. When you are satisfied with the results, you can disable simulation mode.

Simulation mode is not a good choice if you are not an experienced IDP security administrator or when you suspect a critical security risk. In these cases, we recommend that you continue to drop traffic while you investigate.

You might also switch to simulation mode on your live network when you are troubleshooting traffic outages due to IDP processing crashes. Before the IDP Series supported simulation mode, your customer support representative might have advised you to deploy the device in sniffer mode while you were waiting for a detector engine

update or service patch to resolve the root cause of a crash. With IDP OS Release 5.1 and later, simulation mode is a good choice if you want to leave the device physically in path (you do not want to reconfigure and reconnect your traffic interfaces as required for the out-of-path, sniffer mode deployment). However, in these situations, sniffer mode is a better choice if you want the device to send TCP resets to close connections when a security policy rule matches.

**Related
Documentation**

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Simulation Mode Overview on page 25
- Sniffer Mode Overview
- Example: Fine-Tuning a Security Policy on page 167

The following related topic is included in the *IDP Series Administration Guide*:

- Enabling Simulation Mode

CHAPTER 4

Profiler

This chapter provides an overview of Profiler features and examples of when and how to use Profiler. It includes the following topics:

- Profiler Overview on page 31
- Example: Using Profiler to Set a Baseline on page 33
- Example: Using Profiler to Alert You to New Hosts and Port Activity on page 38
- Example: Identifying Services That Use Nonstandard Ports on page 38
- Example: Responding to Vulnerability Announcements with Due Diligence on page 39
- Example: Using Profiler to Investigate Unanticipated Attacks on page 40
- Example: Using Profiler to Mitigate Risks from Laptops on page 41

Profiler Overview

The Profiler is a network-analysis tool that helps you learn about your internal network so you can create effective security policies and minimize unnecessary log records. The Profiler queries and correlates information from multiple IDP Series devices.

After you configure the Profiler, it automatically learns about your internal network and the elements that constitute it, including hosts, peers (which host is talking to which other host), ports (non-IP protocols, TCP/UDP ports, RPC programs), and Layer 7 data that uniquely identifies hosts, applications, commands, users, and filenames. You can use this data to investigate and analyze potential problems in the network and to resolve security incidents.

During profiling, the IDP Series device records network activity at Layer 3, Layer 4, and Layer 7 and stores this information in a searchable database called the Profiler DB. The Profiler uses session creation, session teardown, and protocol contexts to generate this database, which defines all unique activities occurring on your network. Unique activities include attempts, probes, and successful connections.

The device logs normal events only once, and it logs all unique events as often as they occur.

A *normal event* is an event that reoccurs frequently and does not change. For example, suppose Wendy holds a meeting every Tuesday at 4:00 PM in conference room A. Every meeting, she connects her laptop to the network and accesses documents on the primary

fileserver. Because the same event occurs multiple times, the device logs the event once and includes a timestamp that indicates the first and last times Wendy accessed the network from conference room A.

A *unique event* is an event that is new, unexpected, or does not match the normal traffic patterns of your network. For example, suppose that in her weekly meeting, Wendy accesses documents from a different fileserver or has a colleague lead the meeting when she is on vacation. Because the network session information differs, the device logs these activities separately from the normal Tuesday afternoon meeting.

When you configure the Profiler, you can specify:

- General settings, such as whether to collect application volume data and whether to record the OS fingerprint of network hosts
- Network and host IP addresses to track in Profiler logs
- Network and host IP addresses to exclude in Profiler logs
- Contexts to retrieve additional data
- Alerts in cases where you want to track new hosts and applications

For complete procedures on setting Profiler options, see the *IDP Series Administration Guide*.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Example: Using Profiler to Set a Baseline on page 33
- Example: Using Profiler to Alert You to New Hosts and Port Activity on page 38
- Example: Identifying Services That Use Nonstandard Ports on page 38
- Example: Responding to Vulnerability Announcements with Due Diligence on page 39
- Example: Using Profiler to Investigate Unanticipated Attacks on page 40
- Example: Using Profiler to Mitigate Risks from Laptops on page 41

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

Example: Using Profiler to Set a Baseline

A baseline is a place to start. Baseline data gives you the building blocks for your network security policy. The first time you use Profiler, the Profiler report will provide you with detailed views of the devices and applications that communicate in your network.

You use the baseline to:

- Determine which hosts to protect with security policies.
- Determine the applications that communicate over the network and therefore which services require protection in general.
- Determine specific operating systems and software versions in use and therefore which security policy attack objects are relevant and which may be exempted.
- Determine which security policy rulebases are relevant.
- Determine session contexts that can be safely ignored and those that should be monitored.

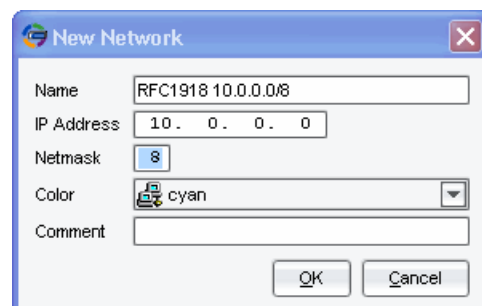
This example assumes a network that uses the private address space defined in RFC 1918: 10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16.

To discover hosts and applications in your private network:

1. Use NSM to create network address objects for each of the three private address spaces.

Figure 9 on page 33 shows the NSM network address object editor.

Figure 9: NSM Network Address Object Editor



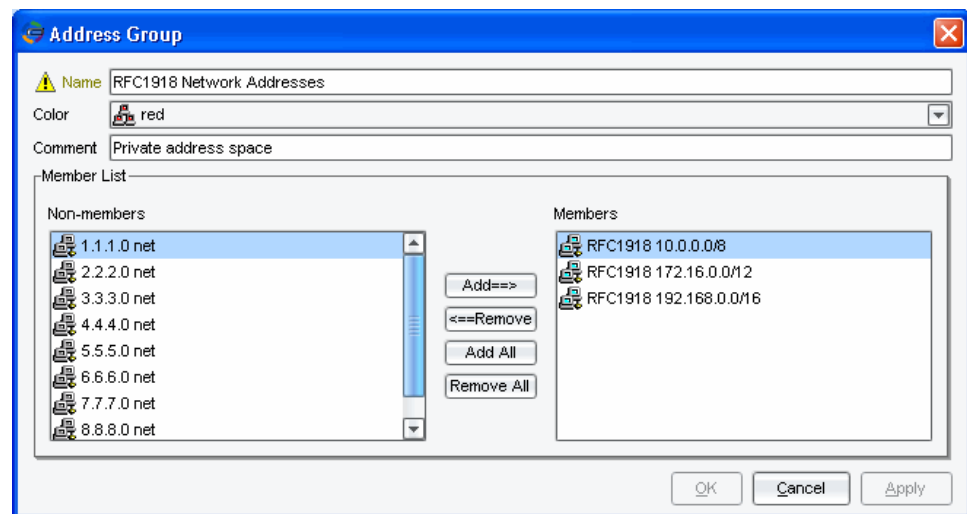
2. Create address objects for any additional networks or hosts that you are aware of.

The Profiler detects host and application information for all traffic that traverses it. If it cannot match the traffic to your network address objects, it assumes you are not tracking the host and populates the source or destination fields as **Non-tracked IP**.

3. Optionally, create a group object to contain the private address space networks.

Figure 10 on page 34 shows the NSM group object editor.

Figure 10: NSM Group Object Editor



4. In the NSM Device Manager, right-click the IDP Series device and select **IDP Profiler** > **Start Profiler**.

Figure 11 on page 34 shows how to navigate in NSM Device Manager to start Profiler.

Figure 11: Starting Profiler from NSM Device Manager

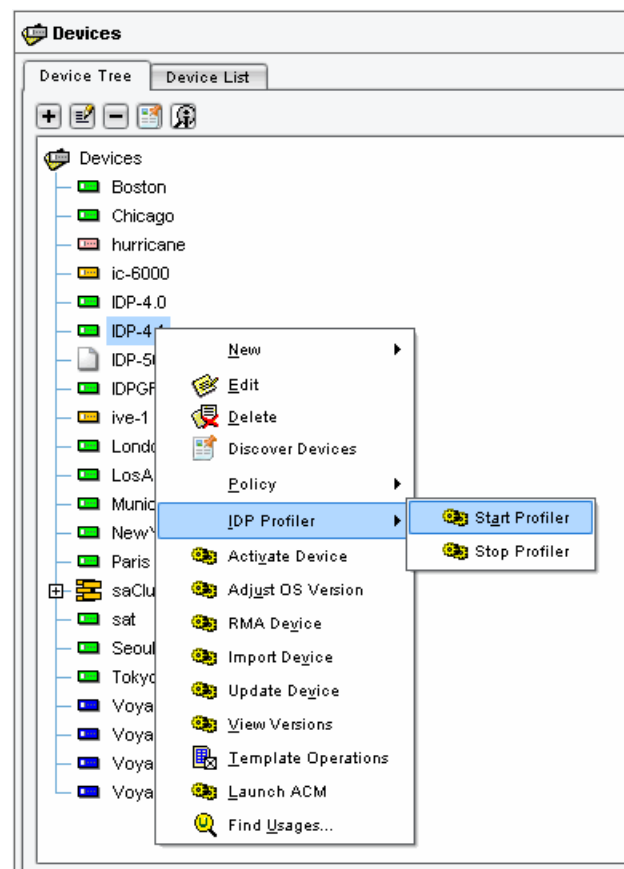
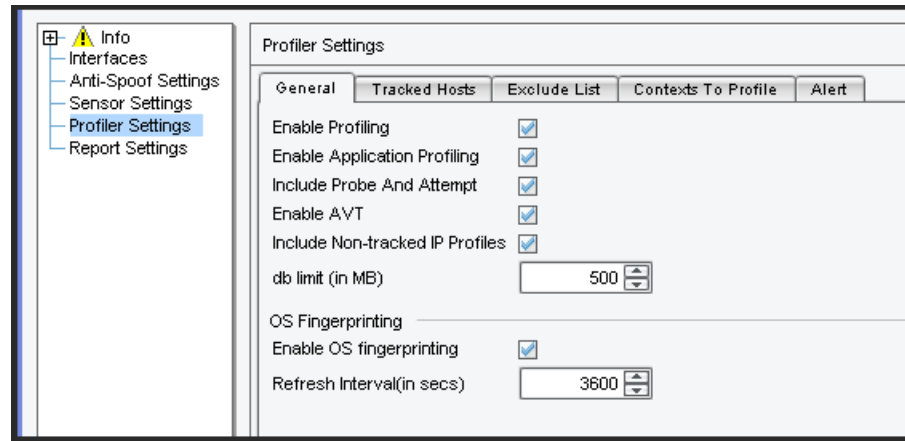


Figure 12 on page 35 shows the Profiler configuration tabs.

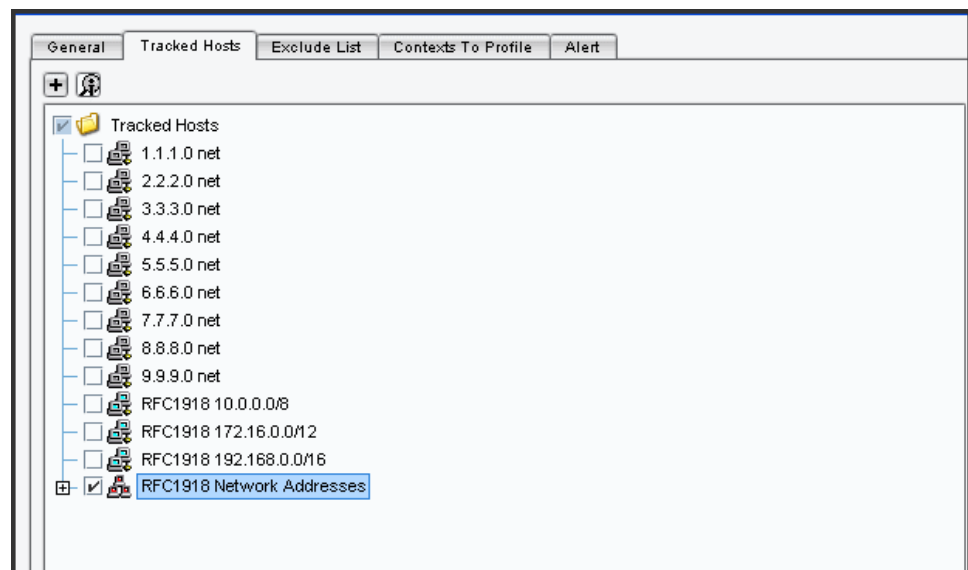
Figure 12: NSM Profiler Configuration Tabs



5. In the General tab, check the boxes to enable profiling, application profiling, OS fingerprinting, and non-tracked IP addresses.
6. Click the **Tracked Hosts** tab and add the address objects you created in Step 1.

Figure 13 on page 35 shows the Tracked Hosts tab.

Figure 13: NSM Profiler Tracked Hosts Tab



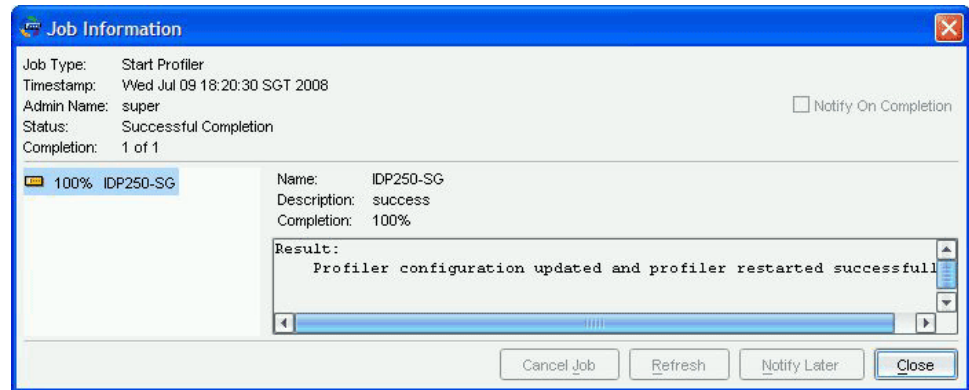
7. Click the **Contexts to Profile** tab and select all contexts.
8. Click the **Alert** tab and clear all alerts. You can use alerts after you have established your baseline but you do not need them in this initial procedure.
9. Click **Apply** to update the Profiler configuration and start the Profiler update job.

The Profiler detects network traffic that traverses the path of the IDP Series device. Consequently, it takes time to build the Profiler database. In most networks, critical

services are used frequently and you might see data in five or ten minutes. For best results, let the Profiler run for a full business day to ensure that it has had enough time to monitor all pertinent network traffic.

Figure 14 on page 36 shows the Job Information window that appears when the Profiler update job is completed.


Figure 14: NSM Profiler Update Job Information Window



10. In the NSM navigation tree, select **Investigate > Security Monitor > Profiler** to display the Profiler viewer.
11. Click the **Network Profiler** tab and examine the data gathered about hosts in your network.

Figure 15 on page 37 shows the Network Profiler tab.

Figure 15: Profiler: Network Profiler Tab



12. Optionally, use the Profiler data to create address objects and groups that you can later use when you create security policy rules.

For example, to create groups for SMTP servers, DNS servers, Windows AD servers, and, HTTP servers:

- a. Create group objects, such as SMTP, DNS, Windows AD, and HTTP.
- b. Use NSM UI features to filter and sort Profiler table rows by service.
- c. Double-click a destination entry to display the host editor, populated with data for the row you clicked.
- d. If you have created a group for the server type, such as SNMP, assign the host to the group.
- e. Click **Save**.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

Example: Using Profiler to Alert You to New Hosts and Port Activity

After you have created a baseline and installed an appropriate security policy, you can use Profiler to alert you when new hosts or applications appear in your network. You can analyze the alerts to decide whether to update your security policy.

To set alerts when Profiler detects new hosts or applications:

1. In the NSM Device Manager, right-click the IDP Series device and select **IDP Profiler > Start Profiler**.
2. Retain your baseline settings for general features, tracked hosts, and contexts.
3. Click the **Alert** tab and select options to generate alerts when Profiler detects new hosts, new protocols, or new ports.
4. Click **Apply**.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

Example: Identifying Services That Use Nonstandard Ports

Suppose you want to identify traffic that uses nonstandard ports so that you can take the appropriate security measures, such as physically removing the unauthorized network components, accounting for nonstandard ports in your existing corporate security policy, or creating rules in your security policy to restrict the traffic to specific network components.

To display a view of traffic that uses nonstandard ports:

1. In the NSM navigation tree, select **Investigate > Security Monitor > Profiler** to display the Profiler viewer.
2. Click the **Violation Viewer** tab.
3. Click the + icon that appears on the top of the right-hand window to display the New Permitted Object window.
4. For this example, name the new permitted object **Non-Standard-Ports**.
5. Right-click the Service column and select **Add Service**.
6. Select all predefined services.
7. Click **OK**.

After you have created and saved the permitted object, the object automatically becomes available in the Profiler.

8. Select the new permitted object **Non-Standard-Ports**.

The Profiler uses the object to filter the data collected from the devices. Traffic that matches the object (uses a standard service port) is filtered out, leaving only the traffic that does not match (uses a nonstandard service port).

9. Review the data for all traffic on your network that uses nonstandard service ports and take appropriate action.

**Related
Documentation**

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

Example: Responding to Vulnerability Announcements with Due Diligence

New network attacks and exploits are discovered every day. When new security patches are issued, use the Profiler to quickly identify which systems are running the affected software version, then patch them appropriately.

For large networks, it is difficult to patch everything immediately. Plan your patching process by prioritizing based on the importance of the resources. Critical, high-risk, and heavily used resources should be patched first, while less important, minimally used resources might be able to wait.

For example, suppose Microsoft announces a vulnerability in version 6.0 of the Microsoft Internet Information Services (IIS).

To quickly identify all network components running the vulnerable version:

1. In the NSM navigation tree, select **Investigate > Security Monitor > Profiler** to display the Profiler viewer.
2. Click the **Protocol Profiler** tab and review the Profiler logs keyed to protocols running on the network.
3. In the Context column, right-click a value and select **Edit Filters** to display the Context Filters dialog box.
4. Set a filter for HTTP Header Servers, for example.

The filtered view highlights the Web servers in your network. Suppose the table lists the following Web servers:

- Apache (two versions)
- Microsoft IIS, version 6.0

5. Select the Microsoft IIS 6.0 value to display your Microsoft IIS 6.0 destination server IP addresses.
6. Patch the vulnerable IIS server by using the information supplied with the Microsoft Security Bulletin.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

Example: Using Profiler to Investigate Unanticipated Attacks

Suppose your corporate security policy does not permit SQL servers on the internal network. However, during a regular Microsoft update, SQL applications are installed on a network server, without your knowledge. Because you are not aware that an SQL server is running on your network, you have not configured security policy rules to block SQL attacks.

Suppose you receive a call informing you that the SQL Slammer worm attacks and infects your network.

To investigate:

1. Create a custom TCP service object to represent Microsoft SQL (default port: TCP/1433).
2. Restart the Profiler.
3. In the NSM navigation tree, select **Investigate > Security Monitor > Profiler** to display the Profiler viewer.
4. Click the **Network Profiler** tab and review the source, destination, and service data.
5. Use a filter to display only records matching the SQL service object you created in Step 1.

The filtered view highlights the SQL servers in your network.

6. Take appropriate measures to secure the network, such as:
 - Applying patches.
 - Removing the components from your network.
 - Removing SQL server from all components.
 - Creating a rule in your security policy that drops all SQL connections between your internal network objects.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

Example: Using Profiler to Mitigate Risks from Laptops

Suppose your corporate firewall denies RPC file sharing traffic to protect sensitive corporate files from Internet users, but enables RPC file sharing on a local network for convenience.

Suppose a laptop user has a good reason to use a partner's wireless network to access the Internet. Because the laptop is configured to allow RPC, it contracts a Blaster worm from an infected user on that network. When the user returns to the office and connects the laptop to the corporate network, the worm immediately begins scanning the internal network and infecting all components that have RPC enabled.

The Profiler records all unique activity on the network, so it identifies the ICMP packet scans as a new event. If you have configured the Profiler to send alerts for new hosts, you receive an alert that a new host has joined the network. In response to the alert, you check the Profiler viewer for details on the new host, and you learn that a host in your network is scanning the entire network using ICMP, a possible sign of the Blaster worm.

To investigate:

1. Restart the Profiler.
2. In the NSM navigation tree, select **Investigate > Security Monitor > Profiler** to display the Profiler viewer.
3. Click the **Network Profiler** tab and review the source, destination, and service data.
4. Apply a filter to the Service column values so that only records matching **ICMP** are displayed.
5. Apply a second filter to the Access Type column so that only records matching **ICMP** and **probe** are displayed.
6. Apply a third filter to the Last Time column so that only records from the last two hours are displayed.

The Network Profiler displays all network components that used ICMP to probe the network in the last two hours.

Assuming the filters have cleared nonmatching records, you can now see the only IP address or IP addresses currently probing your network using ICMP. However, because

you use DHCP to dynamically assign IP addresses, you need to identify which user laptop is currently using that IP address.

7. Right-click the table cell for source properties to display the MAC address. If your enterprise maintains records matching MAC address to laptops, you can track down the specific host that is infected.

**Related
Documentation**

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary

CHAPTER 5

Security Explorer Overview

This chapter describes the NSM Security Explorer. It includes the following topic:

- NSM Security Explorer Overview on page 43

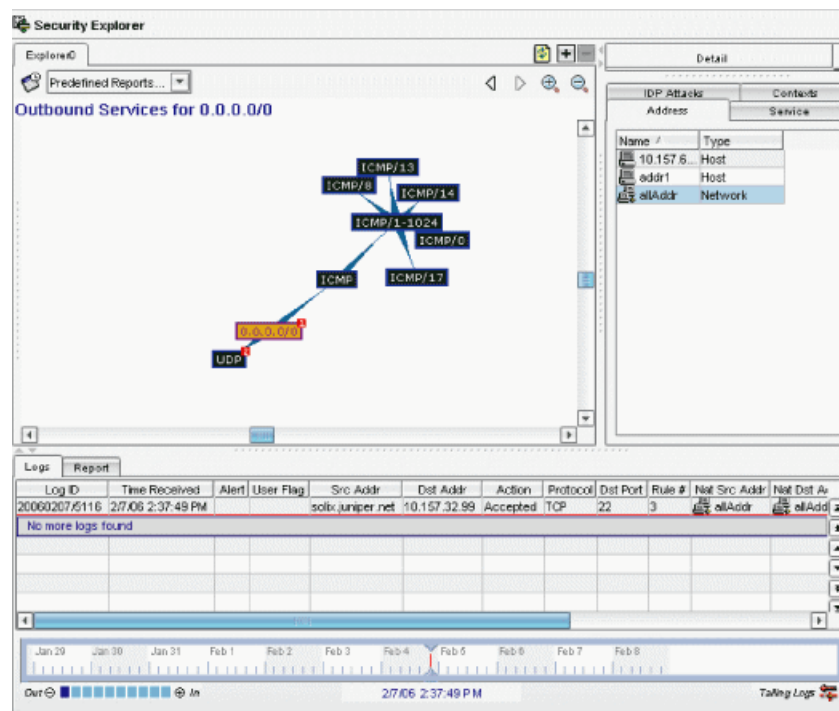
NSM Security Explorer Overview

The NSM Security Explorer is a graphical tool that enables you to visualize and correlate network behavior based on data collected in the Profiler, Log Viewer, and Report Manager. You can use the Security Explorer to:

- Get a dynamic, interactive view of your network.
- Drill down on a particular host or server and view all the different attacks, open ports, destination or peer IP addresses, and so on.
- Move between hosts and peers and trace a connection or attack.
- Toggle between different views or slices of the network, as well as explore the contextual information (logs, reports, IDP attacks, IP addresses, and so on) within the Security Explorer panel.

Figure 16 on page 44 shows the NSM Security Explorer.

Figure 16: NSM Security Explorer



For information on using the NSM Security Explorer, see the NSM documentation.

CHAPTER 6

Application Volume Tracking

This chapter explains how to use the application volume tracking (AVT) feature to gather statistics on applications used over your network. It includes the following topic:

- Application Volume Tracking Overview on page 45
- Example: Using NSM to Enable and View Application Volume Tracking on page 47

Application Volume Tracking Overview

The application volume tracking (AVT) feature uses application identification and the Profiler to collect application statistics aggregated at 15-minute and 1-hour intervals. The AVT database stores up to four sets of each interval at a time (four 15-minute intervals and four 1-hour intervals). After it has accumulated four intervals, it begins dropping the oldest interval as it collects a new one.

The AVT process writes data files to the following directories:

- `/usr/idp/device/var/stat/1hour`
- `/usr/idp/device/var/stat/15min`

The data is collected and parsed for reporting in NSM or IDP Reporter.

Table 9 on page 45 describes the columns of data in AVT records for each session.

Table 9: Application Volume Tracking Data

Data Field	Description
Session ID	Unique ID for the session.
Source IP address	IP address for the host that initiated the session.
Source port	The port number for the source host.
Destination IP address	IP address for the destination server.
Destination port	The port number of the destination host.
VLAN ID	VLAN ID (if any).

Table 9: Application Volume Tracking Data (continued)

Data Field	Description
Protocol	The IP protocol: TCP, UDP, or ICMP.
Application ID	The application identified by the application identification feature. Extended applications (also called nested applications) are reported separately from HTTP results. A 0 indicates the application was not identified.
Bytes	Throughput in bytes for sessions during the interval. AVT tracks both server-to-client and client-to-server bytes.
Packets	Number of packets for sessions during the interval. AVT tracks both server-to-client and client-to-server packets.

Table 10 on page 46 lists documentation references for AVT log viewing tools.

Table 10: Application Volume Tracking Log Viewing Tools

AVT Log Viewing Tools	Documentation
NSM Profiler Viewer > Application Profiler tab (logs)	<i>IDP Series Administration Guide</i>
NSM Report (reports)	
IDP Reporter	<i>IDP Reporter User's Guide</i>



NOTE: To avoid issues with reports, we highly recommend that you synchronize the network clocks for all devices to the same NTP server. For example, the network clocks for all IDP Series devices and NSM clients should be synchronized to the NTP server specified in the NSM configuration.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Example: Using NSM to Enable and View Application Volume Tracking on page 47
- Profiler Overview on page 31
- NSM Reports Overview on page 74
- IDP Reporter Overview on page 76

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary
- IDP Reporter Task Summary

Example: Using NSM to Enable and View Application Volume Tracking

You can use NSM to enable application volume tracking (AVT) and to view AVT logs and reports.

To enable AVT:

1. From NSM Device Manager, double-click a device and then click **Profiler Settings**.
2. Click the **General** tab.
3. Ensure **Enable AVT** is selected. This setting is enabled by default and shown in Figure 17 on page 48.
4. If you have changed settings, click **Apply**.
5. Start the Profiler:
 - a. From the NSM main menu, select **Devices > IDP Profiler > Start Profiler**.
 - b. Select the devices on which you want to start the Profiler.
 - c. Click **OK**.

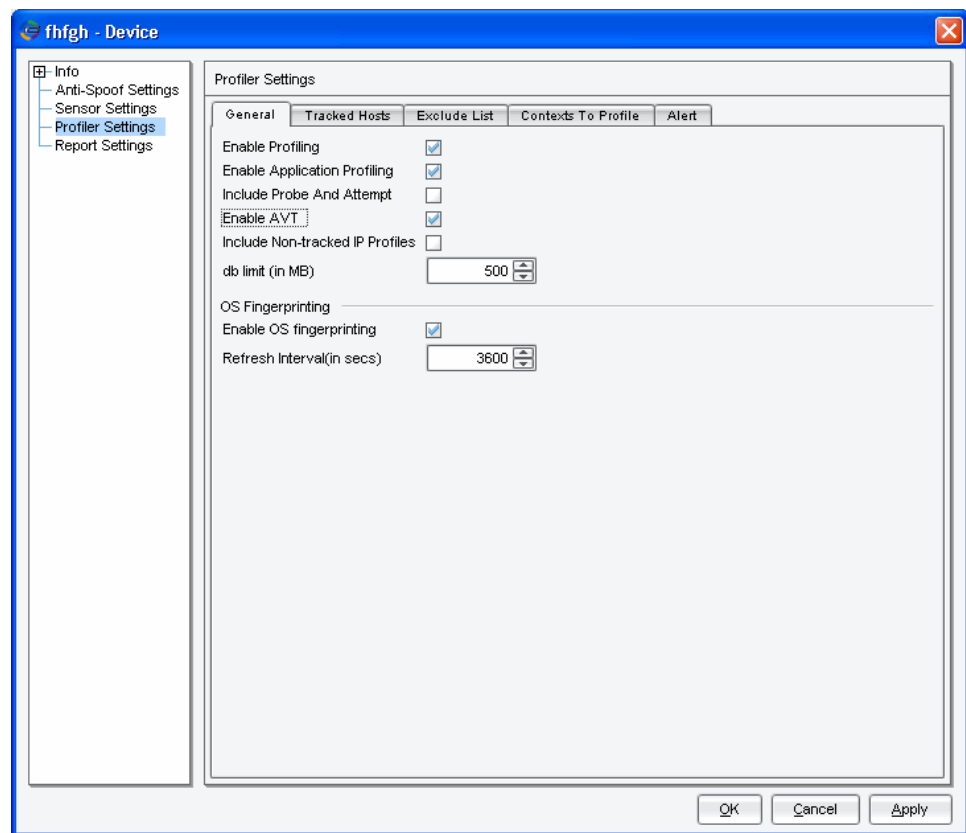


.....

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**.

.....

Figure 17: Profiler Settings: Enable AVT



To view AVT logs:

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

The Application Profiler tab displays application data. Figure 18 on page 48 shows the Application Profiler tab.

Figure 18: Profiler Viewer: Application Profiler Tab

The screenshot shows the 'Profiler' window with the 'Application Profiler' tab selected. The table displays application data for various categories and protocols.

Application Category	Bytes	Packets	Src IP	Dst IP	Dst Port	VLAN ID	Application	Byte Count	Packet Count	User	Role	First Time	Last Time	Dom	Device
Application	10.83 Mb	22,556	HTTP (1.03 Mb, 2,153)												
Unknown-application-category	9.51 Mb	19,887	Window...	128.241.220...	80	0	HTTP	222.34 Kb	329			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Web	1.03 Mb	2,153	Window...	204.160.122...	80	0	HTTP	188.69 Kb	342			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Http	1.03 Mb	2,153	Window...	204.160.122...	80	0	HTTP	150.56 Kb	295			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Remote-access	50.75 Kb	584	Window...	207.46.26.20	80	0	HTTP	122.94 Kb	312			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Misc	14.61 Kb	109	Window...	128.241.220...	80	0	HTTP	72.01 Kb	111			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Node	8.83 Kb	39	Window...	6.12.204.126	80	0	HTTP	55.79 Kb	97			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Noname	5.79 Kb	70	Window...	4.59.128.37	80	0	HTTP	52.42 Kb	76			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Enterprise	12.75 Kb	86	Window...	204.160.122...	80	0	HTTP	50.16 Kb	62			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Infrastructure	12.75 Kb	86	Window...	by2mag1204...	80	0	HTTP	30.49 Kb	111			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Dns	6.98 Kb	68	Window...	ct-in-1103 go...	80	0	HTTP	12.41 Kb	31			Mon Apr 06 ...	Mon Apr 06 ...	global	DP-NS-2...
Dhcp	5.77 Kb	18													
Encryption	2.31 Kb	24													
Ssl	2.31 Kb	24													
Messaging	880	22													
Snmp	880	22													
Peer-to-peer	678	11													
File-sharing	378	6													
Gnutella-udp	378	6													
Chat	300	5													

Extended applications, also called nested applications, are reported separately from HTTP results. Figure 19 on page 49 shows the Application Profiler tab where results for the HTTP Google Earth application are distinguished from HTTP results.

Figure 19: Profiler Viewer: Application Profiler Tab: Nested Applications

Profiler											
Protocol Profiler			Network Profiler			Violation Viewer			Application Profiler		
Application Category	Bytes	Packets	Src IP	Dst IP	Dst Port	VLAN ID	Application	Byte Count	Packet Co...	User	Role
Application	99.22 Mb	76,036	HTTP (62.14 Mb , 47,809)								
Web	99.22 Mb	76,036									
Http	62.14 Mb	47,809	vict1	att1	80	0	HTTP	62.14 Mb	47,809		
Google-earth	37.08 Mb	28,227	GOOGLE-EARTH (37.08 Mb , 28,227)								
			vict1	att1	80	0	GOOGLE-EA...	37.08 Mb	28,227		

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 11 on page 49 describes the Application Profiler session table.

Table 11: Application Profiler Session Table


Column	Description
Src IP	Source IP address of the session.
Dst IP	Destination IP address.
VLAN ID	VLAN ID (if any).
Application ID	Application. The application identified by the application identification feature. Extended applications (also called nested applications) are reported separately from HTTP results. A 0 indicates the application was not identified.
Byte count	Byte count.
Packet count	Packet count.
User	The user associated with the session.
Role	The role to which the user belongs.
First Time	Timestamp for the first time the device logged the event (within the specified time interval).

Table 11: Application Profiler Session Table (*continued*)

Column	Description
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.

The Application Profiler tab displays application data. Figure 20 on page 50 is an example of an NSM AVT report.

Figure 20: NSM AVT Report



NOTE: AVT reports are not real-time reports. On the local IDP Series device, the AVT processor writes an AVT log file at 15 minute intervals. NSM collects the interval data during its routine device log collection activity. As a result, there might be up to a 15 or 16 minute lag from the time a session is received by the IDP Series device and the display of the data in the NSM report.

To view AVT reports:

1. In the NSM navigation tree, select **Investigate > Report Manager > AVT Reports**.
2. Click the name of a predefined report to display it. Table 12 on page 51 describes the predefined AVT reports.

Table 12: 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 hour)	Applications with the highest volume in bytes in the past hour.
Top 5 Application Categories by Volume (last hour)	Application categories with the highest volume in bytes in the past hour.
Top 5 Source by Volume over Time (last hour)	Source IP addresses with the highest volume in bytes in the past hour.
Top 5 Destination by Volume over Time (last hour)	Destination IP addresses with the highest volume in bytes in the past hour.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Application Volume Tracking Overview on page 45
- NSM Reports Overview on page 74
- IDP Reporter Overview on page 76

The following related topic is included in the *IDP Series Administration Guide*:

- Profiler Task Summary
- IDP Reporter Task Summary

CHAPTER 7

Logs and Reports

This chapter describes IDP logging and reporting features. It includes the following topics:

- IDP Logs Overview on page 53
- Developing a Logging Strategy on page 58
- Developing a Log Storage Strategy on page 59
- Example: Using NSM Log Viewer Features on page 62
- Example: Packet Logging Workflow on page 68
- NSM Reports Overview on page 74
- IDP Reporter Overview on page 76

IDP Logs Overview

The IDP system generates logs for device events and security events.

Device event logs are related to the operation of the IDP Series device. IDP OS Release 5.1 supports extensive system resource instrumentation, so you can use SNMP utilities to monitor device health and load.

Security event logs are related to traffic that matches security policy rules or IP spoof attack settings.

Table 13 on page 53 summarizes options for viewing and managing logs.

Table 13: IDP Logging Options

Option	Description
Network and Security Manager (NSM)	<p>IDP Series devices automatically send device event logs to NSM. IDP Series devices send security event logs when traffic matches security policy rules for which logging has been enabled. You can use the NSM log viewer to sort through and analyze logs. If you enable packet logging for a security policy rule, you can use the NSM packet viewer to display packet data.</p> <p>For details on using NSM log utilities, see the <i>Network and Security Manager Administration Guide</i>.</p>
Syslog	<p>You can configure IDP Series devices to forward logs to a syslog server, a commonly used method for storing logs for troubleshooting or record-keeping.</p> <p>For details on configuring syslog collection, see the <i>IDP Series Administration Guide</i>.</p>

Table 13: IDP Logging Options (*continued*)

Option	Description
SNMP	<p>SNMP reporting is enabled by default. You can use SNMP methods to track the following metrics:</p> <ul style="list-style-type: none"> • CPU usage • Memory usage • Disk usage • Packet buffer usage • Session statistics • Network interface statistics • Traffic statistics <p>For details on configuring SNMP reporting, see the <i>IDP Series Administration Guide</i>.</p>
Log suppression	<p>You can configure log suppression to reduce the volume of logs. The log suppression feature eliminates multiple log entries for events with the same properties. Instead, in NSM Log Viewer, a single entry appears along with a count of all such matching events.</p> <p>For details on configuring log suppression, see the <i>IDP Series Administration Guide</i>.</p>



NOTE: To avoid issues with reports, we highly recommend that you synchronize the network clocks for all devices to the same NTP server. For example, the network clocks for all IDP Series devices and NSM clients should be synchronized to the NTP server specified in the NSM configuration.

Table 14 on page 54 describes the fields that appear in log entries.

Table 14: NSM Log Viewer: Log Columns

Column	Description
Log ID	Unique ID for the log entry, derived by combining the date and log number.
Time Received	Date and time that the management system received the log entry.
Alert	NSM-defined alert for this type of log entry. Configure alerts in policy rules.
User Flag	<p>To set a flag, right-click the log row, select Flag, and then select one of the following flags:</p> <ul style="list-style-type: none"> • High • Medium • Low • Closed • False Positive • Assigned • Investigate • Follow-up • Pending

Table 14: NSM Log Viewer: Log Columns (*continued*)

Column	Description
Src Addr	Source IP address of the packet that generated the log entry.
Dst Addr	Destination IP address of the packet that generated the log entry.
Action	<p>Action the security device performed on the packet/connection that generated this log entry:</p> <ul style="list-style-type: none"> Accepted—Did not block the packet. Closed Client—Closed the connection and sent an RST packet to the client, but did neither to the server. Closed Server—Closed the connection and sent an RST packet to the server, but did neither to the client. Closed—Closed the connection and sent an RST packet to both the client and the server. Dropped—Dropped the connection without sending an RST packet to the sender, preventing the traffic from reaching its destination. Dropped Packet—Dropped a matching packet before it could reach its destination but did not close the connection. Ignored—Matched the attack, did not take action, and ignored the remainder of the connection. <p>NOTE: Beginning in IDP OS Release 5.1, IDP logs show the action that was taken, rather than the action that was specified in the rule. For TCP events, these are the same. Close actions are not possible for UDP or ICMP packets. For UDP and ICMP events, the IDP logs show the action take—drop—instead of a close client, close server, or close client and server actions that might have been configured for the rule.</p>
Protocol	Protocol that the packet that generated the log entry used.
Dst Port	Destination port of the packet that generated the log entry.
Rule #	Security policy rule that generated the log entry.
Nat Src Addr	NAT source address of the packet that generated the log entry.
Nat Dst Addr	NAT destination address of the packet that generated the log entry.
Details	Miscellaneous string associated with log entry.

Table 14: NSM Log Viewer: Log Columns (*continued*)

Column	Description
Category	<p>Type of log entry:</p> <ul style="list-style-type: none"> • Admin • Alarm—The device generates event alarms for any security event that has a predefined severity level of emergency, critical, or alert. Additionally, the device generates traffic alarm log entries when it detects network traffic that exceeds the specified alarm threshold in a rule (the traffic alarm log entry describes the security event that triggered the alarm). • Config—A configuration change occurred on the device. • Custom—A match with a custom attack object was detected. • Implicit—An implicit rule was matched. • Info—General system information. • Predefined—A match with a predefined attack object was detected. • Profiler—Traffic matches a Profiler alert setting. • Screen—Not applicable for IDP Series devices. Generated by ScreenOS firewall devices. • Self—The device generated this log for a non-traffic related reason. • Sensor. • Traffic—Traffic matches a rule you have configured for harmless traffic. • URL Filtering—Not applicable for IDP Series devices. Generated by ScreenOS firewall devices. • User.
Subcategory	Category-specific type of log entry (examples are "Reboot" or message ID).
Severity	<p>Severity rating associated (if any) with this type of log entry:</p> <ul style="list-style-type: none"> • Not Set (the device could not determine a severity for this log entry) • Info • Device_warning_log • Minor • Major • Device_critical_log • Emergency • Error • Notice • Informational • Debug
Device	Device that generated this log entry.
Comment	User-defined comment about the log entry.
Application Name	Application associated with the current log.
Bytes In	For sessions, specifies the number of inbound bytes.
Bytes Out	For sessions, specifies the number of outbound bytes.

Table 14: NSM Log Viewer: Log Columns (*continued*)

Column	Description
Bytes Total	For sessions, specifies the combined number of inbound and outbound bytes.
Dev Domain Ver	Domain version that generated this log entry.
Device Domain	Domain for the device that generated this log entry.
Device family	Family of the device that generated this log entry.
Dst Intf	Name of the outbound interface of the packet that generated this log entry.
Dst Zone	Destination zone associated with a traffic log entry.
Elapsed Secs	For sessions, specifies how long the session lasted.
HasPacket Data	Indicates whether the log entry has associated packet data.
NAT Dst Port	The NAT destination port of the packet that generated the log entry.
NAT Src Port	The NAT source port of the packet that generated the log entry.
Packets In	For sessions, specifies the number of inbound packets.
Packets Out	For sessions, specifies the number of outbound packets.
Packets Total	For sessions, specifies the combined number of inbound and outbound packets.
Policy	Security policy that generated the log entry.
Roles	Role group associated with this log entry.
Rule Domain	The domain of the rule that generated the log entry.
Rule Domain Ver	The domain version of the rule that generated the log entry.
Rulebase	Security policy rulebase that generated the log entry.
Src Intf	Name of the inbound interface of the packet that generated this log entry.
Src Port	Source port of the packet that generated the log entry.
Src Zone	Source zone associated with a traffic log entry.
Time Generated	Date and time the device generated the log entry.
User	User associated with this log entry.

The following example shows a syslog message record:

```
[syslog@juniper.net dayId="20061012" recordId="0" timeRecv="2006/10/12 21:52:21"
timeGen="2006/10/12 21:52:21"
domain="" devDomVer2="0" device_ip="10.209.83.4" cat="Predefined"
attack="TROJAN:SUBSEVEN:SCAN"
srcZn="NULL" srcIntf="NULL" srcAddr="192.168.170.20" srcPort="63396"
natSrcAddr="NULL" natSrcPort="0"
dstZn="NULL" dstIntf="NULL" dstAddr="192.168.170.10" dstPort="27374"
natDstAddr="NULL" natDstPort="0" protocol="TCP"
ruleDomain="" ruleVer="5" policy="Policy2" rulebase="IDS" ruleNo="4" action="NONE"
severity="LOW" alert="no"
elapsedTime="0" inbytes="0" outbytes="0" totBytes="0" inPak="0" outPak="0"
totPak="0" repCount="0" packetData="no"
varEnum="31" misc="017" interface="eth2" user="NULL" app="NULL" uri="NULL"]
```

Related Documentation The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Developing a Logging Strategy on page 58
- Developing a Log Storage Strategy on page 59
- Example: Using NSM Log Viewer Features on page 62
- Example: Packet Logging Workflow on page 68
- Understanding IDP Rulebase Notification Options on page 103
- Understanding Backdoor Rulebase Notification Options on page 145
- Understanding SYN Protector Rulebase Notification Options on page 153
- Understanding Traffic Anomalies Rulebase Notification Options on page 160
- Understanding Network Honeypot Rulebase Notification Options on page 164
- IP Spoof Attack Prevention Overview on page 173

The following related topics are included in the *IDP Series Administration Guide*:

- IDP Series Logs and Reports in NSM Task Summary
- SNMP Statistic Reporting and Traps Task Summary

Developing a Logging Strategy

Intrusion prevention systems can generate hundreds of logs per hour. In order to make the best use of the security logs, you should develop strategic approaches to the following administrative tasks:

- Fine-tuning the security policy rules.

Security policy rules determine the amount of logging performed by the IDP Series device, as well as automatic actions to take on offending traffic, such as dropping the session, sending a TCP reset, blocking the IP address from future connections, and so forth. See “Example: Fine-Tuning a Security Policy” on page 167.

- Analyzing log event summaries and packet capture data.

By viewing log summaries, attack reference information, and packet data, you can verify whether the severity and actions associated with a security event are appropriate, whether refinements to your security policy are required, and whether further response actions are warranted. See “Example: Using NSM Log Viewer Features” on page 62.

- Managing log and packet storage.

Your business log management and log storage policies determine where you store IDP Series device logs and security event logs. Your IDP Series device supports local logging, central collection by NSM, and forwarding to a syslog server. See “Developing a Log Storage Strategy” on page 59.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- IDP Series Logs and Reports in NSM Task Summary

Developing a Log Storage Strategy

This topic summarizes IDP log storage and log forwarding options so you can develop a log storage strategy suitable for your business. It includes the following sections:

- Log Management Considerations on page 59
- Local Log Files and Directories on page 59
- NSM Log Collection on page 61

Log Management Considerations

An IDP Series device might generate hundreds of logs per day. Your log storage strategy depends on a number of factors:

- The nature of your business. Compliance with regulations or business agreements might determine where you collect logs or how often you retain them.
- Existing log management infrastructure. We recommend you become familiar with an use Network and Security Manager (NSM) as a central location for log analysis, but your previous investments in technology and training are also strong considerations.
- Distribution to the appropriate personnel for analysis is also a key consideration.

If your organization has not formalized a log management policy, consult the National Institute of Standards and Technology (NIST) publication, [Guide to Computer Security Log Management](#), for a treatment of the myriad considerations.

Local Log Files and Directories

Logs are stored locally on the device in subdirectories of `/usr/idp/device/var`. Log pruning occurs when a disk partition reaches 90% capacity.

Table 15: IDP Local Log Directories

Directory	Content
/usr/idp/device/var/logs	Local storage for device and security event logs before they are forwarded to NSM.
/usr/idp/device/var/pktlogs	Local storage for packet capture logs before they are forwarded to NSM.
/usr/idp/device/var/profile	Local storage for Profiler database logs before they are forwarded to NSM.
/usr/idp/device/var/sysinfo/logs	Location where system messages are written.
/usr/idp/device/var/stat/	Local storage for application volume tracking logs before they are forwarded to NSM, IDP Reporter, or Application Usage Manager.



NOTE: Although /usr/idp/device/var is a symbolic link to /var/idp/device/var, user scripts or programs created to manage files should reference the /usr/idp/device/var path.

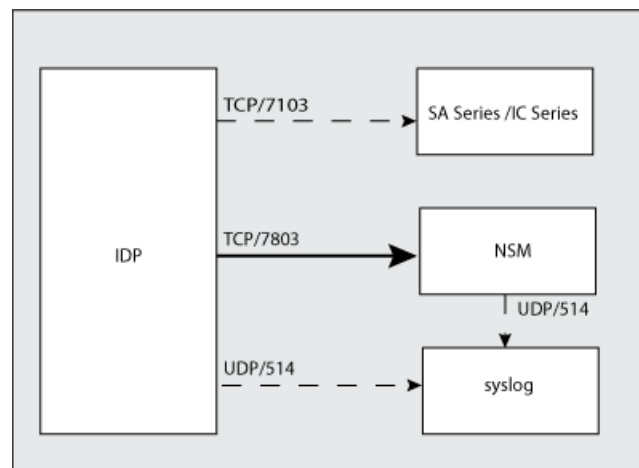
By default, logs are forwarded to NSM, which is the primary user interface for the IDP Series device.

Optionally, you can configure the IDP Series device to send copies of logs to external devices, such as:

- A syslog server, including a Juniper Networks Security Threat Response Manager (STRM) device, which reads the IDP syslog format.
- A Juniper Networks Secure Access Series or Infranet Controller Series device to inform access policies.

Figure 21 on page 61 provides a visual summary of your log forwarding options. The solid line indicates default behavior. The dashed lines indicate options you must configure to use.

Figure 21: IDP Log Storage and Log Forwarding



NOTE: In IDP OS Release 5.1, syslog protocol port are configurable. However, we recommend you use the standard protocol and port whenever feasible.

NSM Log Collection

By default, the IDP Series device sends logs to NSM where they can be displayed and analyzed with the NSM user interface. We recommend you become familiar with an use NSM as a central location for log analysis. Logs are stored on the NSM Device Server in subdirectories of `/usr/netscreen/DevSvr/var/logs`. NSM supports the following log management features:

- Command-line utilities to archive, copy, and purge logs.
- Configurable time retention policies that trigger pruning.
- Automated log management jobs based on criteria you configure, including severity, category, and so forth.
- Support for log field filters in export operations to XML, CSV, syslog, SNMP, e-mail, or script.

For complete information on NSM log management features, see Chapter 19 of the [NSM Administration Guide](#).

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- IDP Logs Overview on page 53

The following related topics are included in the *IDP Series Administration Guide*:

- IDP Series Logs and Reports in NSM Task Summary
- Connecting to the Command-Line Interface (CLI Procedure)

Example: Using NSM Log Viewer Features

The Network and Security Manager (NSM) Log Viewer includes many display features to help you sort and correlate logs so you can analyze security events. For complete information on NSM Log Viewer features, see Chapter 18 of the [NSM Administration Guide](#). The following sections are provided here to give you ideas of how to take advantage of NSM features as you develop your approach to log monitoring:

- Using Predefined Views on page 62
- Showing and Hiding Columns on page 63
- Using Filters on page 63
- Using Log Viewer Detail Panes on page 65
- Using Flags and Comments on page 65
- Using Custom Views on page 67

Using Predefined Views

Out of the box, the NSM Log Viewer includes a predefined view for DI/IDP event logs. A predefined view is a filtered view of all logs collected on the NSM device server. The DI/IDP view is filtered for events that match a predefined or custom attack object (Category field = Predefined or Category = Custom). Figure 22 on page 62 shows the DI/IDP view.

Figure 22: NSM Log Viewer: Predefined View

Log Viewer [3-IDP]

Log ID	Time Received	Alert	User	Src Addr	Dst Addr	Action	Protocol	Dst Port	Rule #	Nat Src Addr	Nat Dst Addr	Details
20090813/770	8/12/09 7:52:54 PM			1.0.0.199	2.0.0.23	Conn Dropped	TCP	1433	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/771	8/12/09 7:52:54 PM			1.0.0.22	2.0.0.115	Conn Dropped	TCP	1433	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/772	8/12/09 7:52:58 PM			1.0.0.139	2.0.0.179	Conn Dropped	TCP	12174	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/773	8/12/09 7:53:01 PM			1.0.0.148	2.0.0.248	Conn Dropped	TCP	22	1	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/774	8/12/09 7:53:01 PM			1.0.0.138	2.0.0.119	Conn Dropped	TCP	22	1	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/775	8/12/09 7:53:04 PM			1.0.0.1	2.0.0.69	Conn Dropped	TCP	80	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/776	8/12/09 7:53:04 PM			1.0.0.166	2.0.0.45	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/777	8/12/09 7:53:04 PM			1.0.0.28	2.0.0.97	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/778	8/12/09 7:53:04 PM			1.0.0.155	2.0.0.87	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/779	8/12/09 7:53:04 PM			1.0.0.43	2.0.0.238	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/780	8/12/09 7:53:04 PM			1.0.0.172	2.0.0.8	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/781	8/12/09 7:53:04 PM			1.0.0.167	2.0.0.20	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/782	8/12/09 7:53:04 PM			1.0.0.113	2.0.0.94	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/783	8/12/09 7:53:04 PM			1.0.0.230	2.0.0.49	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/784	8/12/09 7:53:04 PM			1.0.0.73	2.0.0.119	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'
20090813/785	8/12/09 7:53:04 PM			1.0.0.89	2.0.0.232	Conn Dropped	UDP	7787	2	0.0.0.0	0.0.0.0	'interface=eth2'

No more logs found

Timeline: Aug 6 Aug 7 Aug 8 Aug 9 Aug 10 Aug 11 Aug 12 Aug 13 Aug 14 Aug 15 Aug 16 Aug 17 Aug 18 Aug 19

Out 8/12/09 7:53:04 PM Tailing Logs

Summary All Fields Whois Lookup Quick Report

Predefined :: APP: Unreal Gamespy Query Protocol Buffer Overflow

References

This signature detects attempts to exploit a known vulnerability against the GameSpy query protocol supported by Unreal game engine. Attackers can crash a game server running the Unreal game engine, or execute arbitrary code with permissions of the user running the server.

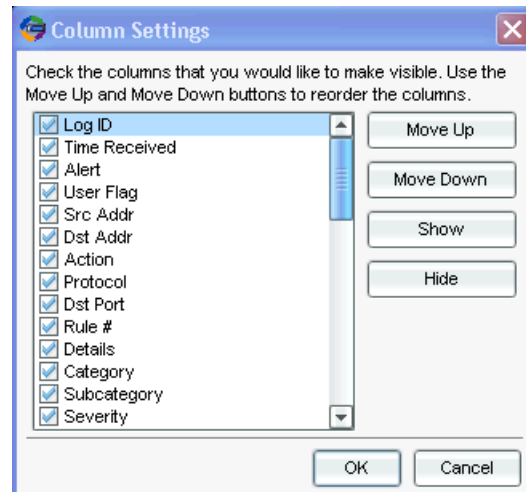
Matching Data Snippet

HEX ASCII

Showing and Hiding Columns

The default columns shown in the predefined DI/IDP view might not include all of the data fields you are interested in. To select your preferred columns and the order in which they appear, select **View > Choose Columns** and use the dialog box to organize columns according to your preference.

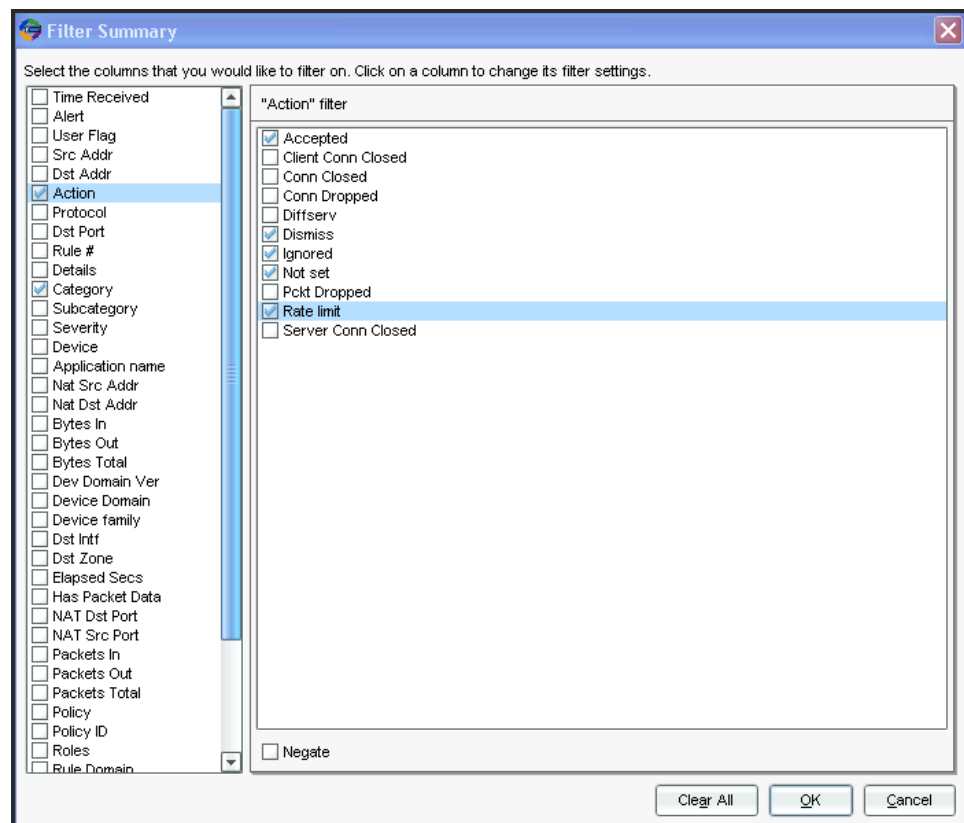
Figure 23: NSM Log Viewer: Choose Columns



Using Filters

The default DI/IDP view is filtered to display only logs where Category=Predefined or Category=Custom. To set additional filters, select **View > Filter Summary** and use the dialog box to set additional filters. In Figure 24 on page 64, filters are selected to display logs for traffic where the rule action allowed the traffic continue to the destination server. When you approach the set of logs you examine each day, you might want to start with events of high severity, where traffic continued to the destination.

Figure 24: NSM Log Viewer: Filters



You can also filter on the fly. Suppose you find a log for an attack targeting HTTP traffic. In the row for the log, you can right-click the cell containing destination port **80** and select **Filter > Only This Value** to redisplay the table with only records where destination port = 80.

Figure 25: NSM Log Viewer: Filters

Log Viewer [3-BDP.DIG]

Log ID	Time Received	Alert	User Flag	Comment	Src Addr	Dst Addr	Action	Protocol	Dst	Nat Sr	Nat Ds	Details	Category	Subcategory	
20090806/473467	8/6/09 7:13:50 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473469	8/6/09 7:13:55 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473486	8/6/09 7:14:26 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473489	8/6/09 7:14:32 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473491	8/6/09 7:14:35 AM				10.0.99	2.0.0.99	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	redefined HTTP: Missing HTTP Version	Traffic	Accept
20090806/473493	8/6/09 7:14:40 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473499	8/6/09 7:14:46 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473500	8/6/09 7:14:49 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473501	8/6/09 7:14:52 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473516	8/6/09 7:15:56 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Traffic	Accept	
20090806/473518	8/6/09 7:16:02 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473520	8/6/09 7:16:10 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473522	8/6/09 7:16:13 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473526	8/6/09 7:16:29 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473530	8/6/09 7:16:32 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473532	8/6/09 7:16:39 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473534	8/6/09 7:16:45 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473539	8/6/09 7:16:48 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473540	8/6/09 7:17:04 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473542	8/6/09 7:17:14 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473543	8/6/09 7:17:14 AM				10.0.99	2.0.0.99	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Predefined HTTP: Missing HTTP Version	Traffic	Accept
20090806/473545	8/6/09 7:17:20 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473548	8/6/09 7:17:20 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473553	8/6/09 7:17:30 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept
20090806/473557	8/6/09 7:17:36 AM				10.0.99	2.0.0.99	Rate limit	TCP	80	2	0.0.0.0	0.0.0.0	Interface-eth2	Traffic	Accept

Using Log Viewer Detail Panes

The details pane below the log table provides summary and security reference information for the attack object that triggered the log. The details pane also includes a link to WHOIS information for the source IP.

Suppose your security policy rule includes the following attack object: Predefined :: HTTP: Windows Media Services NSISLog.DLL Buffer Overflow. It generates a log when it identifies the attack pattern in traffic through the IDP Series device. Use the reference information in the details pane below the log table to learn more about the attack. You can click the hypertext linked name of the attack object in the summary tab to display reference information for the attack, as shown in Figure 26 on page 65.

Figure 26: Using NSM Log Viewer Attack Reference Information

The screenshot displays the NSM Log Viewer [3-IDP/CDI] interface. The main window contains a log table with columns: Log ID, Time Received, Alert, User Flag, Comment, Src Addr, Dest Addr, Action, Protocol, Port, Src, Dest, Details, Category, and Subcategory. The table lists multiple log entries, with one entry highlighted: Log ID 20090806416967, Time Received 8/5/09 11:13:42 PM, Alert (red triangle), User Flag (red X), Comment 'Windows 2000 SP4 only?', Src Addr 1.1.0.88, Dest Addr 1.2.0.56, Action 'Conn Dropped', Protocol 'TCP', Port '80', Src '4', Dest '0.0.0.0', Details 'Interface=eth2', Category 'Predefined', and Subcategory 'HTTP: ISS 0 WinDVD SEARCH Co...'. Below the table is a timeline view showing logs from July 30 to August 5, 2009, with a filter set to 'All Fields' and a 'Quick Report' button.

A detailed pane is open on the right, titled 'HTTP: Windows Media Services NSISLog.DLL Buffer Overflow'. It contains the following sections:

- References:**
 - <http://online.securityfocus.com/bid/3035/discussion/>
 - <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0349>
 - <http://www.kb.cert.org/vuls/id/113716>
 - <http://www.microsoft.com/technet/security/bulletin/MS03-022.mspx>
 - <http://sacunia.com/advisories/9115>
- Extended Description:**
 - Last Modified:** 2009-08-13
 - Impact:** Windows Media Services may expose IIS to remote arbitrary code execution if media logging is enabled.
 - Description:** Microsoft has reported a buffer overflow vulnerability in Windows Media Services. This is due to a problem with how the logging ISAPI extension handles incoming client requests. This could cause arbitrary code execution in IIS, which is exploitable through Media Services.
 - Technical Information:** Microsoft Media Services provides functionality for providing streaming media content to clients from IIS. It ships with a number of Microsoft Windows 2000 server releases and is also available for download for Windows NT. Microsoft has reported a buffer overflow vulnerability in Windows Media Services. This is due to a problem with how the logging ISAPI extension (nsislog.dll) handles incoming client requests. The logging facility may attempt to write excessive data to an undersized buffer when handling a malformed HTTP client request. This could trigger a denial of service or remote arbitrary code execution in IIS, which is exploitable through Media Services. The issue would occur in servers that are configured to provide logging of media requests. It is possible to exploit this issue by sending an overly long HTTP POST request to the

At the bottom of the pane, there is a 'Close' button.

Using Flags and Comments

As you work through logs, you can annotate them with flags and comments and then filter on your annotations. Figure 27 on page 66 shows a log marked as a false positive because the attack targets server versions not present in our network.

Figure 27: Using NSM Log Viewer Flag and Comment Features

Juniper Networks - NSM - global : current

Search Help

Log Viewer [3-IDP000]

Log ID	Time Received	Alert	User Flag	Comment	Src Addr	Dst Addr	Action	Protocol	Dst...	Nat Sr...	Nat De...
20090806/416941	8/5/09 11:13:33 PM				1.1.0.115	1.2.0.58	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416943	8/5/09 11:13:33 PM				1.1.0.192	1.2.0.102	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416945	8/5/09 11:13:33 PM				1.1.0.212	1.2.0.241	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416948	8/5/09 11:13:33 PM				1.1.0.248	1.2.0.132	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416949	8/5/09 11:13:36 PM				1.1.0.63	1.2.0.159	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416950	8/5/09 11:13:36 PM				1.1.0.88	1.2.0.56	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416951	8/5/09 11:13:36 PM				1.1.0.161	1.2.0.81	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416956	8/5/09 11:13:36 PM				1.1.0.231	1.2.0.243	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416957	8/5/09 11:13:36 PM				1.1.0.231	1.2.0.243	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416961	8/5/09 11:13:36 PM				1.1.0.241	1.2.0.121	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416966	8/5/09 11:13:42 PM				1.1.0.170	1.2.0.91	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416967	8/5/09 11:13:42 PM			Windows 2000 SP4 only?	1.1.0.88	1.2.0.56	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416971	8/5/09 11:13:45 PM				1.1.0.241	1.2.0.121	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416972	8/5/09 11:13:45 PM				1.1.0.23	1.2.0.139	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417097	8/5/09 11:14:59 PM				1.1.0.188	1.2.0.231	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417098	8/5/09 11:17:37 PM				1.1.0.20	1.2.0.143	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417099	8/5/09 11:17:40 PM				1.1.0.199	1.2.0.227	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417100	8/5/09 11:17:40 PM				1.1.0.114	1.2.0.190	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417101	8/5/09 11:17:40 PM				1.1.0.234	1.2.0.123	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417102	8/5/09 11:17:40 PM				1.1.0.189	1.2.0.95	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417103	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417104	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417105	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417106	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417107	8/5/09 11:17:43 PM				1.1.0.205	1.2.0.103	Conn Dropped	TCP	80	4	0.0.0.0
20090806/417108	8/5/09 11:17:43 PM				1.1.0.109	1.2.0.55	Conn Dropped	TCP	80	4	0.0.0.0

Filter Find Flag Exempt... Show Hide Log Unhide Log Goto Policy

High Medium Low Closed False Positive Assigned Investigate Follow-Up Pending Clear

Jul 30 Jul 31 Aug 1 Aug 2 Aug 3 Aug 4 Aug 5 Aug 6 Aug 7

Out In 8/5/09 11:13:45 PM

Summary All Fields Whois Lookup Quick Report

[Predefined :: HTTP: Windows Media Services NSISlog.DLL Buffer Overflow](#)

[References](#)
This signature detects attempts to exploit the buffer overflow vulnerability against Microsoft Windows Media Services, included with Microsoft Windows 2000 Server SP4. Attackers can send a maliciously crafted HTTP "POST" request to overflow the buffer and cause a denial of service or execute arbitrary code.

Matching Data Snippet

HEX

Filtered on: Category

To mark a log with a flag, right-click the cell in the Flag column and select one of the following flags:

- High (severity)
- Medium (severity)
- Low (severity)
- Closed
- False Positive
- Assigned
- Investigate
- Follow-Up
- Pending

Using Custom Views

As you become familiar with NSM Log Viewer filters, you are likely to discover views of the data you typically want to use to monitor traffic. You can save custom views. Because the custom view is based on filters, incoming log entries that match the filter criteria are automatically displayed in the view. You do not need to reapply the view to new logs.

Figure 28 on page 67 shows a custom view of columns and filters focusing on events where the IDP Series device allowed HTTP traffic to proceed to its destination.

Figure 28: NSM Log Viewer: Custom View

[illegible]

You might want to create views to help manage the following example cases:

- **Workflow**—If your team distributes responsibilities based on IDP Series device, internal servers, application, severity, or type of attack, you can create views filtered on the appropriate columns. In the same manner, you can also use the Flag or Comments columns to prioritize or delegate investigation.
- **Attackers**—Once you learn the IP address of an attacker, you can create a view filtered on Source IP to watch what the attackers activities on your network.
- **Devices**—After you deploy a new device, you can create a view filtered on the Device column to observe and validate device effectiveness.

To create a new view, select the columns you want to display and apply filters. Select **File > New View** to display a dialog box to save the view in your preferred Log Viewer folder. We recommend saving custom views in the Custom folder.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- IDP Logs Overview on page 53
- NSM Reports Overview on page 74

The following related topics are included in the *IDP Series Administration Guide*:

- IDP Series Logs and Reports in NSM Task Summary

Example: Packet Logging Workflow

This topic summarizes IDP Series packet logging basics. It includes the following sections:

- Using Packet Captures on page 68
- Enabling Packet Capture in Security Policy Rules on page 68
- Forwarding Packet Capture Logs to NSM on page 69
- Viewing Packet Capture Logs on page 70

Using Packet Captures

The IDP solution supports packet capture logging triggered by security policy rules.

You can use packet captures for a number of response activities, including:

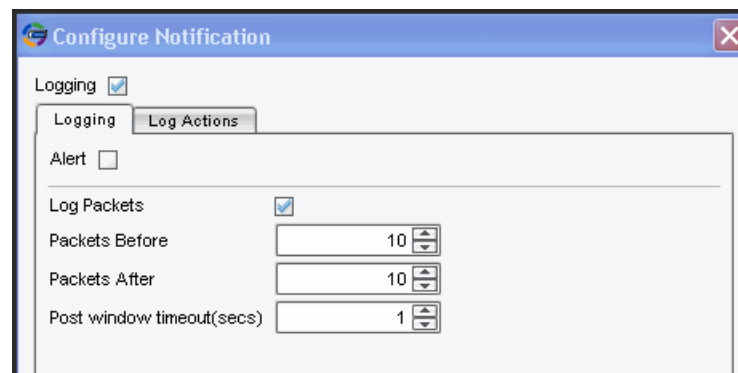
- Validation of the security policy rule and attack object. You may choose to enable packet logging to test a new attack object. Once verified, you may find packet logging for the rule unnecessary.
- Further analysis of traffic surrounding the matching event. The surrounding traffic might provide information that helps you determine whether you need to take further steps to protect the target or whether the attack should be considered a false positive.
- Reproducibility and documentation for Internet security groups, including the Juniper Networks Security Center.
- Legal evidence. Consult with your legal counsel for guidance on how local laws and rules of evidence apply if you want to use packet capture data as evidence in the prosecution of attackers.

Enabling Packet Capture in Security Policy Rules

When traffic matches a rule where packet logging is configured, the IDP Series device captures the packet that matched the rule, as well as the preceding and trailing packets (according to your configured preference).

To enable packet logging within a security policy rule, use the Security Policy editor. Right-click a cell in the Notification column and select **Configure** to display the dialog box where you can set packet logging options.

Figure 29: Notification Options: Packet Logging



In the NSM Log Viewer, logs for events where packet captures have been generated are noted by an icon in the Has Packet Data column (the last column in Figure 30 on page 69).

Figure 30: NSM Log Viewer: Has Packet Data Column

Src Addr	Dst Addr	Action	Protocol	Dst Port	Rule #	Nat Src Addr	Nat Dst Addr	Details	Category	Subcategory	Severity	Device	Comment	Has Packet Data
1.1.0.68	1.2.0.40	Conn Dropped	TCP	554	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	RTSP: Real Server Describe Overfl...	Major	DP8202		
1.1.0.28	1.2.0.26	Conn Dropped	TCP	554	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	RTSP: Real Server Transport Over...	Major	DP8202		
1.1.0.206	1.2.0.111	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	HTTP: Cisco IOS HTTP Configuratio...	Major	DP8202		
1.1.0.56	1.2.0.34	Conn Dropped	TCP	25	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SMTP: Email Domain Name	Major	DP8202		
1.1.0.56	1.2.0.34	Conn Dropped	TCP	25	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SMTP: Email Address	Major	DP8202		
1.1.0.48	1.2.0.159	Conn Dropped	TCP	25	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SMTP: Email Address	Major	DP8202		
1.1.0.80	1.2.0.175	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	HTTP: Cisco IOS HTTP Configuratio...	Major	DP8202		
1.1.0.110	1.2.0.192	Conn Dropped	TCP	25	3	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SMTP: Microsoft Exchange Mailer...	Device_critical_log	DP8202		
1.1.0.118	1.2.0.198	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: PASV/FTP: (negotio) Input Valid...	Major	DP8202		
1.1.0.230	1.2.0.123	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Username Too Long	Major	DP8202		
1.1.0.62	1.2.0.168	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Username Too Long	Major	DP8202		
1.1.0.231	1.2.0.116	Conn Dropped	TCP	25	3	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SMTP: Exchange Multiple Long Mail...	Device_critical_log	DP8202		
1.1.0.4	1.2.0.133	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Username Too Long	Major	DP8202		
1.1.0.87	1.2.0.181	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SHELLCODE: X86 NOOP (TCP)	Major	DP8202		
1.1.0.81	1.2.0.168	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Username Too Long	Major	DP8202		
1.1.0.150	1.2.0.208	Conn Dropped	TCP	25	3	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	SMTP: Sendmail Oversized Address...	Device_critical_log	DP8202		
1.1.0.231	1.2.0.116	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Username Too Long	Major	DP8202		
1.1.0.112	1.2.0.193	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Ipswitch WS_FTP Server FTP...	Major	DP8202		
1.1.0.224	1.2.0.122	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Pathname Too Long	Major	DP8202		
1.1.0.17	1.2.0.136	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Username Too Long	Major	DP8202		
1.1.0.224	1.2.0.122	Conn Dropped	TCP	21	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	FTP: Pathname Too Long	Major	DP8202		
1.1.0.223	1.2.0.239	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	HTTP: Missing HTTP Version	Major	DP8202		
1.1.0.98	1.2.0.184	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	HTTP: Microsoft WS/Who Buffer Ov...	Major	DP8202		
1.1.0.98	1.2.0.184	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	HTTP: Missing HTTP Version	Major	DP8202		
1.1.0.56	1.2.0.163	Conn Dropped	TCP	119	3	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	NNTP: XPA1 Pattern Overflow	Device_critical_log	DP8202		
1.1.0.222	1.2.0.244	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	Interface-eth2	Preddefined	HTTP: AIN WebAdmin USER Buff...	Major	DP8202		

Forwarding Packet Capture Logs to NSM

The IDP Series device writes packet captures locally to subdirectories of `/usr/idp/device/var/pktlogs/`. It forwards the packet data to NSM according to your NSM Report Settings:

- **Include packet data in log** selected. Forwards the packet capture to NSM automatically whenever it sends the corresponding event log.
- **Include packet data in log** not selected. Forwards a reference to the packet capture file to NSM automatically but forwards the packet data itself only on-demand (when an NSM user takes action to display the packet data).

Figure 31: NSM Device Configuration Editor: Report Settings

Viewing Packet Capture Logs

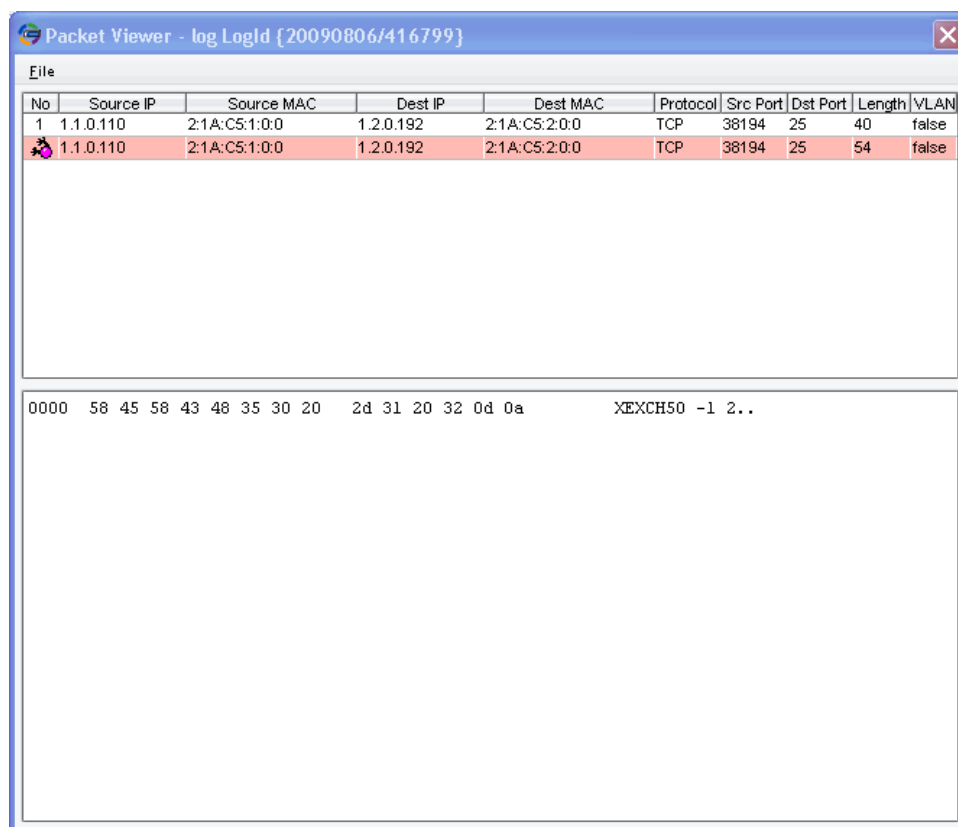
You have two options for viewing packet captures:

- Using the NSM Packet Viewer on page 70
- Using an External Viewer to View Packet Data on page 71

Using the NSM Packet Viewer

The NSM packet viewer displays the offending attack payload that triggered the alert as well as preceding and trailing packets (according to your configuration). Figure 32 on page 71 shows the NSM packet capture viewer.

Figure 32: NSM Packet Capture Viewer



To view a packet capture in the NSM packet viewer:

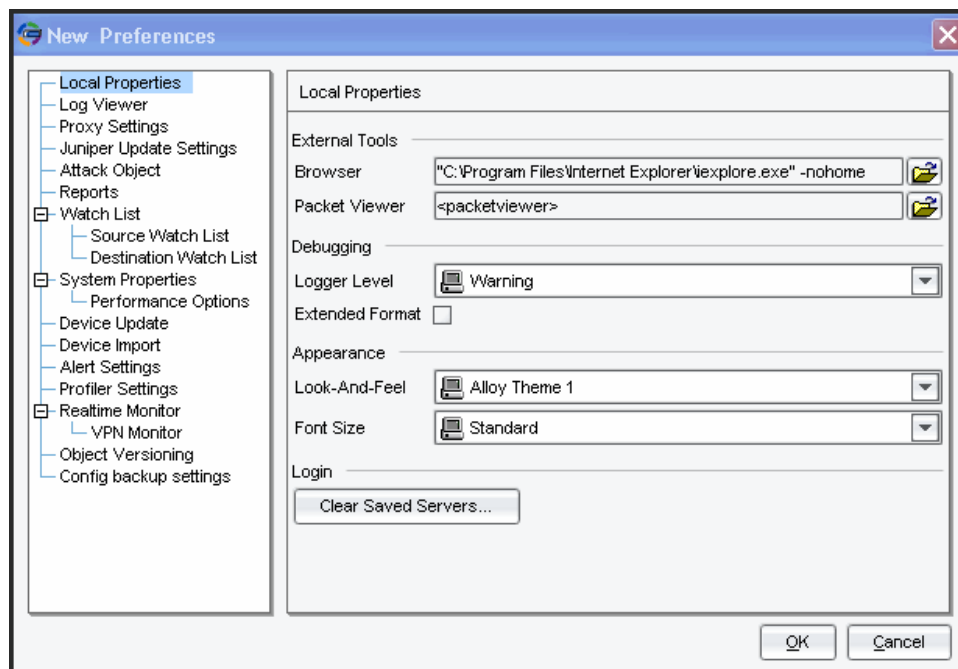
1. In the NSM navigation tree, select **Investigate > Log Viewer > Predefined > DI/IDP** to display the IDP table.
2. Select **View > Choose Columns** to display the dialog box you use to show and hide log table columns.
3. Select **Has Packet Data** to show this column.
If a security event log has packet data, an icon appears in the table cell under this column.
4. Double-click the Has Packet data icon to display the packet data in the NSM packet viewer.

Using an External Viewer to View Packet Data

You can configure NSM to launch an external viewer for packet captures.

Figure 33 on page 72 shows the NSM dialog box where you can specify the location of an external packet viewer.

Figure 33: Specifying an External Viewer

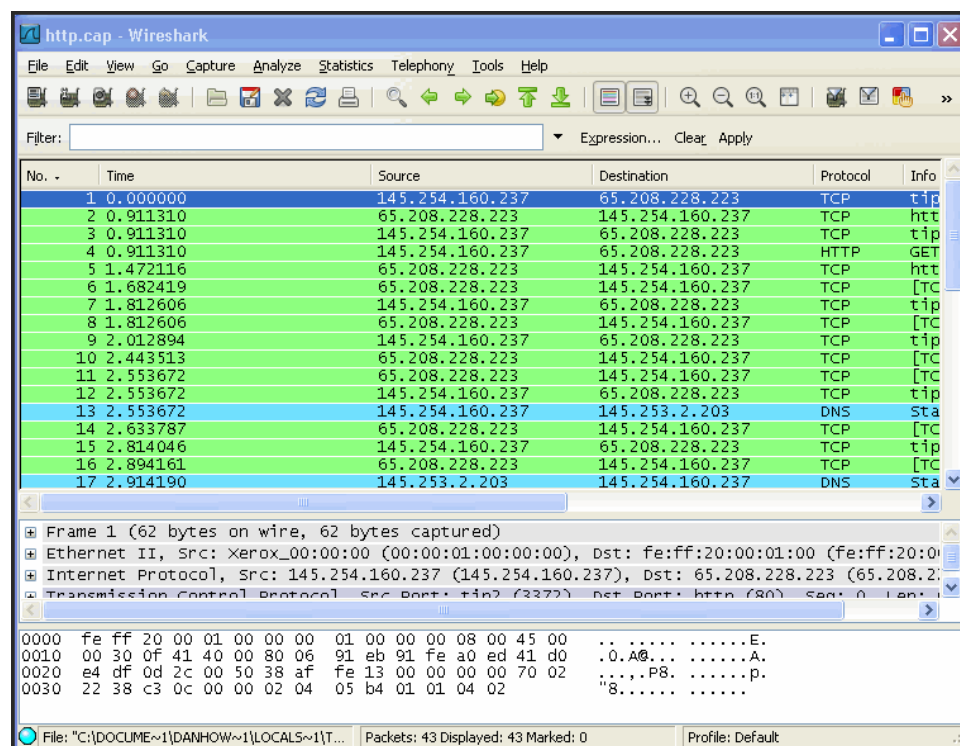


To set the location of the external viewer:

1. In NSM, select **Tools > Preferences**.
2. Select **Local Properties**.
3. Under External Tools > Packet Viewer, click the browse button and select the executable file for the external viewer (for example: C:\Program Files\Wireshark\wireshark.exe).
4. Click **OK** to close the New Preferences dialog box.

Figure 34 on page 73 shows packet data displayed in the Wireshark packet viewer.

Figure 34: Wireshark Packet Viewer



To view a packet capture in an external packet viewer:

1. In the NSM navigation tree, select **Investigate > Log Viewer > Predefined > DI/IDP** to display the IDP table.
2. Select **View > Choose Columns** to display the dialog box you use to show and hide log table columns.
3. Select **Has Packet Data** to show this column.

If a security event log has packet data, an icon appears in the table cell under this column.

4. Right-click the Has Packet data icon and select **Show > Packet Data in External Viewer**.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding IDP Rulebase Notification Options on page 103
- Using tcpdump to Capture Packets

The following related topics are included in the *IDP Series Administration Guide*:

- Specifying Rule Notification Options (NSM Procedure)
- Enabling Collection of Packet Data in NSM Logs (NSM Procedure)

NSM Reports 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 single IDP Series device through IDP Reporter.

NSM Report Manager contains a set of predefined network and security reports, including a group of deep inspection (DI) and intrusion detection and prevention (IDP) reports.

Table 16 on page 74 summarizes NSM DI/IDP predefined reports.

Table 16: 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) during the last 24 hours.
All Attacks Prevented by Severity (last 24 hours)	Number of attacks prevented by severity level (set in attack objects) during the last 24 hours.
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.

Table 16: NSM DI/IDP Predefined Reports (*continued*)

Report	Description
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 been the source of a scan during the past 7 days.
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 17 on page 75 describes Profiler predefined reports. These reports are related to activity by hosts in your network.

Table 17: NSM Profiler Predefined Reports

Report	Description
Top 10 Peers by Count	Source and destination IP addresses that appeared most frequently in the Profiler logs.
Top 10 Peers with maximum hits	Source and destination IP addresses that had the highest number of hits in the Profiler logs.

Table 18 on page 75 describes the predefined application volume tracking (AVT) reports. The reports are related to application use in your network.

Table 18: 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 hour.

Table 18: NSM: Application Volume Tracking Reports (*continued*)

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

In addition to these predefined reports, you can create custom reports based on IDP log fields. For details on creating custom reports, see the *IDP Series Administration Guide*.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- IDP Series Logs and Reports in NSM Task Summary

IDP Reporter Overview

IDP reports are representations of log data, aggregated and sorted to facilitate network and security analysis. IDP Series devices support both centralized, aggregated reporting through NSM, and on-box reporting for a singular IDP Series device through IDP Reporter.

IDP Reporter is a Java application that has been preinstalled on your IDP Series device. You can access IDP Reporter through a Web interface. Like NSM Report Manager, IDP Reporter contains predefined reports and enables you to create custom reports based on log fields.

For details on accessing and using IDP Reporter, see the *IDP Reporter User's Guide*.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- IDP Logs Overview on page 53
- Application Volume Tracking Overview on page 45

The following related topic is included in the *IDP Series Administration Guide*:

- IDP Reporter Task Summary

PART 3

Protecting Your Network

- Security Policy Basics on page 79
- Predefined Security Policies on page 87
- The IDP Rulebase on page 91
- The Exempt Rulebase on page 113
- Application Policy Enforcement Rulebase on page 117
- The Backdoor Rulebase on page 141
- The SYN Protector Rulebase on page 149
- The Traffic Anomalies Rulebase on page 155
- The Network Honeypot Rulebase on page 161
- Fine-Tuning a Security Policy on page 167
- Additional Security Features on page 173

CHAPTER 8

Security Policy Basics

This chapter provides a basic overview of what is involved in creating the security policies that protect your network. It includes the following topics:

- Understanding Non-Policy-Based Drops on page 79
- Understanding the Components of an IDP Security Policy on page 83
- Understanding the Number of Available and Installed Policies on page 85
- Understanding the Rule-Matching Algorithm on page 85

Understanding Non-Policy-Based Drops

The IDP Series device inspects traffic that traverses it and takes action according to:

- Implicit rules
- Protocol anomaly threshold settings
- Security policy rules

Table 19 on page 80 summarizes implicit rules that drop traffic. The related topics listed provide information about protocol anomaly threshold settings and IDP security policy rules.

Table 19: Non-Policy-Based Drops

Implicit Rule	Description
Layer 2 traffic (when bypass not enabled)	<p>Enabled: Dropped by default. Configurable in ACM.</p> <p>When the IDP Series device is turned on and is operating normally, the traffic interfaces process TCP/IP traffic according to implicit traffic anomaly rules and explicit security policy rules. For Layer 2 connections, the interfaces process traffic, drop it, or pass it through (uninspected), according to the following rules:</p> <ul style="list-style-type: none"> The interfaces process Address Resolution Protocol (ARP) and Internet Protocol (IPv4) traffic for inspection and process according to implicit and explicit rules. By default, the interfaces drop all other Layer 2 traffic. <p>When Layer 2 bypass is enabled, the IDP Series device passes through Layer 2 packets related to bypass and high availability deployments (such as heartbeats or Bridge Protocol Data Unit (BPDU) packets), and non-IPv4 packets and packets related to switching and routing protocols, such as IPv6, internetwork packet exchange (IPX), Cisco Discovery Protocol (CDP), and interior gateway routing protocol (IGRP), and so forth.</p> <hr/> <p>Counter: If you do not enable Layer 2 bypass, you can use the following counters to observe Layer 2 drops:</p> <pre>[root@default host ~]# scio counter get kpp grep sc_kpp_jpkt_free sc_kpp_jpkt_free 1374077</pre> <pre>[root@default host ~]# scio counter get kpp grep sc_kpp_other sc_kpp_other 305</pre>
Invalid IP header	<p>Enabled: By default.</p> <hr/> <p>Counter:</p> <pre>[root@default host ~]# scio counter get kpp grep sc_kpp_bad_ip_header sc_kpp_bad_ip_header 0</pre> <hr/> <p>Event logs: None</p> <hr/> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>
Invalid TCP header	<p>Enabled: By default.</p> <hr/> <p>Counter:</p> <pre>[root@default host ~]# scio counter get reass grep sc_reass_bad_tcp_header sc_reass_bad_tcp_header 0</pre> <hr/> <p>Event logs: Yes</p> <hr/> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>

Table 19: Non-Policy-Based Drops (*continued*)

Implicit Rule	Description
TCP checksum error	<p>Enabled: Logging for this event is disabled by default. Use the following command to enable logging for this event:</p> <pre>[root@default host ~]# scio const set sc_log_implicit_pkt_drop 1</pre> <p>Counter:</p> <pre>[root@default host ~]# scio counter get reass grep sc_reass_bad_tcp_csum sc_reass_bad_tcp_csum 0</pre> <p>Event logs: If you enable logging, event logs are generated and sent to NSM and/or a syslog server.</p> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>
UDP checksum error	<p>Enabled: The counter for this event is disabled by default. Use the following command to enable it:</p> <pre>[root@default host ~]# scio const set sc_enable_udp_csum 1</pre> <p>Counter:</p> <pre>[root@default host ~]# scio counter get flow grep sc_flow_bad_udp_csum sc_flow_bad_udp_csum 0</pre> <p>Event logs: If you enable logging, event logs are generated and sent to NSM and/or a syslog server.</p> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>
ICMP source quench	<p>Dropping ICMP source quench messages is disabled by default. Use the following command to enable it:</p> <pre>[root@default host ~]# scio const -s s0:flow set sc_icmp_drop_source_quench 1</pre> <p>Event logs: Event logs are generated and sent to NSM or a syslog server.</p> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>
TTL error	<p>Enabled: By default.</p> <p>Counter:</p> <pre>[root@default host ~]# scio counter get kpp grep sc_kpp_ttl_error sc_kpp_ttl_error 0</pre> <p>Event logs: None</p> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>

Table 19: Non-Policy-Based Drops (*continued*)

Implicit Rule	Description
Memory limit of busy packet list	<p>Enabled: By default.</p> <hr/> <p>Counter:</p> <pre>[root@defaulthost ~]# scio counter get kpp grep sc_kpp_busy_drop sc_kpp_busy_drop 0</pre> <hr/> <p>Event logs: None</p> <hr/> <p>Debug logs: None</p>
Dropped by reassembly module when per flow memory overflows	<p>Enabled: By default.</p> <hr/> <p>Counter:</p> <pre>[root@defaulthost ~]# scio counter get kpp grep sc_kpp_fdrops sc_kpp_fdrops 0</pre> <hr/> <p>Event logs: None</p> <hr/> <p>Debug logs: None</p>
Global reassembly memory overflow	<p>Enabled: By default.</p> <hr/> <p>Counter:</p> <pre>[root@defaulthost ~]# scio counter get reass grep sc_reass_ovflw_drop sc_reass_ovflw_drop 0</pre> <hr/> <p>Event logs: None</p> <hr/> <p>Debug logs: If debugging is enabled, debug logs are generated and saved to /var/idp/device/sysinfo/logs/.</p>
Transmit failure where packet has already been freed	<p>Enabled: By default.</p> <hr/> <p>Counter:</p> <pre>[root@defaulthost ~]# scio counter get kpp sc_kpp_transmit_error sc_kpp_transmit_error 0</pre> <hr/> <p>Event logs: None</p> <hr/> <p>Debug logs: None</p>

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Components of an IDP Security Policy on page 83

The following related topic is included in the *IDP Series Administration Guide*:

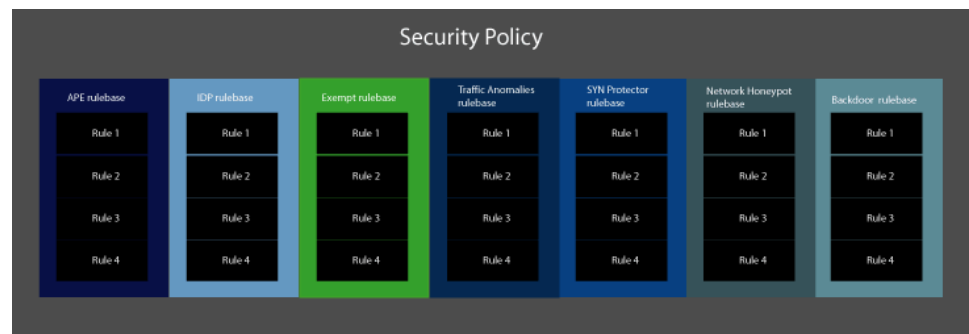
- Modifying Protocol Anomaly Thresholds

Understanding the Components of an IDP Security Policy

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

Figure 35 on page 83 illustrates the components of an IDP security policy.

Figure 35: Security Policy Components



A security policy is made up of one or more rulebases. A *rulebase* is an ordered set of rules that use a particular detection method to identify and prevent attacks.

Table 20 on page 83 describes the IDP security policy rulebases. A security policy can contain only one instance of any rulebase type.

Table 20: IDP Security Policy Rulebases

Rulebase	Description
APE rulebase	<p>Enables you to implement application policy enforcement rules. You can use APE rules to manage sessions based on application and/or user role. You can terminate matching sessions or limit bandwidth available to them.</p> <p>See “Understanding the APE Rulebase” on page 117.</p>
IDP rulebase	<p>Protects your network from attacks by using attack objects to detect known and unknown attacks. Juniper Networks Security Center (J-Security Center) provides predefined attack objects that you can use in IDP rules. You can also configure your own custom attack objects.</p> <p>See “Understanding the IDP Rulebase” on page 91.</p>
Exempt rulebase	<p>You configure rules in this rulebase to exclude known false positives or to exclude a specific source, destination, or attack object 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.</p> <p>See “Understanding the Exempt Rulebase” on page 113.</p>

Table 20: IDP Security Policy Rulebases (*continued*)

Rulebase	Description
Traffic Anomalies rulebase	<p>Protects your network from attacks by using traffic flow analysis to identify attacks that occur over multiple connections and sessions (such as scans).</p> <p>See “Understanding the Traffic Anomalies Rulebase” on page 155.</p>
SYN Protector rulebase	<p>Protects your network from SYN-floods by ensuring that the three-way handshake is performed successfully for specified TCP traffic. If your network is vulnerable to SYN-flood attacks, use the SYN-Protector rulebase to prevent it.</p> <p>See “Understanding the SYN Protector Rulebase” on page 149.</p>
Network Honeypot rulebase	<p>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.</p> <p>See “Understanding the Network Honeypot Rulebase” on page 161.</p>
Backdoor rulebase	<p>Protects your network from mechanisms installed on a host computer that facilitate 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.</p> <p>See “Understanding the Backdoor Rulebase” on page 141.</p>



NOTE: Firewall rulebases, visible in NSM, do not apply to standalone IDP Series devices.

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

- A match condition that determines which traffic to inspect
- Attack objects that determine what to look for (IDP rulebase and Exempt rulebase)
- Actions and operation modes that determine what to do when traffic matches a rule
- Notification options, including logs, alerts, and packet captures

Related Documentation

The following additional related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Number of Available and Installed Policies on page 85
- Understanding the Rule-Matching Algorithm on page 85
- Example: Fine-Tuning a Security Policy on page 167

The following additional related topic is included in the *IDP Series Administration Guide*:

- Developing Security Policies Task Summary

Understanding the Number of Available and Installed Policies

In NSM, you can create and save an unlimited number of security policies, and these policies are available to be installed on IDP Series devices.

You can install the same security policy on an unlimited number of IDP Series devices.

You can install one security policy on an IDP Series device. However, when you push a security policy update, you might observe more than one policy in place for a period after the update.

By default, the IDP system resets the flow table when you install a new policy. When the flow table is reset, existing sessions are passed through uninspected. For IDP75 and IDP200, you cannot override the default.

For high-end appliances, you can unset this default to avoid passing through sessions uninspected. Go to NSM Device Manager > Run-time Parameters and unselect **Reset flow table with policy load/unload**. If you unset this default, when you load a new policy, the IDP flow table maintains sessions belonging to the previously installed policy as well as the newly installed policy. The IDP process engine continues to use the previously installed security policy to inspect previous sessions; and uses the newly installed security policy to inspect new sessions. When the previously installed policy is no longer in use, it is unloaded and all traffic is inspected using the newly installed policy. For IDP8200 and IDP250, the IDP engine can maintain flows for as many as two security policies. For IDP1100, IDP800, and IDP600, the IDP engine can maintain flows for as many as four security policies.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Components of an IDP Security Policy on page 83

The following related topic is included in the *IDP Series Administration Guide*:

- Developing Security Policies Task Summary
- Modifying the IDP Series Device Configuration

Understanding the Rule-Matching Algorithm

The IDP process engine processes rulebases in the following order:

1. Application Policy Enforcement (APE) rulebase (terminal)
2. Traffic Anomalies rulebase (terminal)
3. SYN Protector rulebase (terminal)
4. Network Honeypot rulebase (terminal)
5. IDP rulebase (nonterminal)

6. Exempt rulebase (nonterminal)

7. Backdoor rulebase (terminal)

For terminal rulebases, the IDP rule-matching algorithm evaluates rules according to the ordered list to identify matches. As soon as the algorithm identifies a match, it applies the rule and terminates rule matching. For example, if a terminal rule 1 matches, it is applied, and rule 2 is not consulted.

For nonterminal rulebases, the IDP rule-matching algorithm also evaluates rules according to the ordered list to identify matches. However, even if it finds a match, it continues down the list to identify additional matches.

The IDP rulebase includes the option to mark a rule as a terminal rule. When a terminal rule matches source, destination, and service, IDP applies the rule and terminates rule matching. It does not matter whether the traffic matches the attack objects.

In the IDP rulebase, you can set the terminal rule flag for the following purposes:

- To disregard traffic that originates from a particular trusted source (however, an Exempt rulebase rule might be a better choice).
- To exit rule processing when you want a particular match to always trigger a particular action and no other, such as a drop connection action.
- To exit rule processing when your rule specifies precise destination addresses and precise services, and you know that the subsequent rules do not apply.

Use caution when specifying the terminal flag. You can inadvertently leave your network open to attacks by creating an inappropriate terminal rule. Be particularly careful about terminal rules using the value **Any** for both the source and destination. Terminal rules should appear near the top of the rulebase, before other rules that would match the same traffic.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Components of an IDP Security Policy on page 83
- Example: Fine-Tuning a Security Policy on page 167
- Understanding APE Rulebase Match Conditions on page 118
- Understanding IDP Rulebase Rule Match Settings on page 92
- Understanding Backdoor Rulebase Match Settings on page 143
- Understanding SYN Protector Rulebase Match Settings on page 151
- Understanding Traffic Anomalies Rulebase Match Conditions on page 157
- Understanding Network Honeypot Rulebase Match Settings on page 162

The following related topic is included in the *IDP Series Administration Guide*:

- Developing Security Policies Task Summary

CHAPTER 9

Predefined Security Policies

This chapter describes the predefined security policies that you can use to secure your network. It includes the following topics:

- Using the Recommended Security Policy on page 87
- Using Other Security Policy Templates on page 88

Using the Recommended Security Policy

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 (or customize it) to protect your network from the likeliest and most dangerous attacks.

Table 21 on page 87 summarizes the properties of the Recommended security policy.

Table 21: Recommended Security Policy Definition

Property	Value
Rulebase	IDP rulebase
Rules	Nine rules, distinguished by attack object
Source	Any
Service	Default, meaning the matching property is based on the service bindings of the attack object specified by the rule
Destination	Any
Attack objects	Recommended IP, Recommended TCP, Recommended ICMP, Recommended HTTP, Recommended SMTP, Recommended DNS, Recommended FTP, Recommended POP3, Recommended IMAP, Recommended Trojan, Recommended Virus, Recommended Worm NOTE: All of the attack objects included in the predefined policies are client-to-server attacks.
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.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the IDP Rulebase on page 91
- IDP Rulebase Example: Specifying the Default Service on page 108
- IDP Rulebase Example: Using Application Identification on page 107
- IDP Rulebase Example: Using Recommended Attack Objects on page 109
- IDP Rulebase Example: Using Recommended Actions on page 110
- Example: Fine-Tuning a Security Policy on page 167

The following related topic is included in the *IDP Series Administration Guide*:

- Developing Security Policies Task Summary

Using Other Security Policy Templates

NSM includes security policy templates you can use as the basis for a custom security policy tailored for your network. Template rules include a set of attack objects and logically associated IDP actions. If you choose to use these templates, we advise you to 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.

Table 22 on page 88 describes IDP security policy templates.

Table 22: IDP Security Policy Templates

Template	Description
all_with_logging	Includes all attack objects and enables packet logging for all rules. This policy is provided for lab use and is not recommended in production.
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.
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 set of attack groups that balances security and performance.
web_server	Protects HTTP servers from remote attacks.



NOTE: All of the attack objects included in the predefined policies are client-to-server attacks.

**Related
Documentation**

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Using the Recommended Security Policy on page 87

The following related topic is included in the *IDP Series Administration Guide*:

- Developing Security Policies Task Summary

CHAPTER 10

The IDP Rulebase

The purpose of this chapter is to explain how the IDP rulebase protects your network and to provide examples on how to leverage IDP rulebase features. It includes the following topics:

- Understanding the IDP Rulebase on page 91
- Understanding IDP Rulebase Rule Match Settings on page 92
- User-Role-Based Policy Feature Overview on page 94
- Using Application Identification on page 96
- Using Attack Objects on page 97
- Understanding IDP Rulebase Actions on page 101
- Understanding IDP Rulebase Notification Options on page 103
- IDP Rulebase Example: User-Role-Based Policies on page 104
- IDP Rulebase Example: Using Application Identification on page 107
- IDP Rulebase Example: Specifying the Default Service on page 108
- IDP Rulebase Example: Using Recommended Attack Objects on page 109
- IDP Rulebase Example: Using Recommended Actions on page 110

Understanding the IDP Rulebase

The IDP rulebase employs an attack object database to support two robust detection methods: stateful signatures and protocol anomalies.

A *stateful signature* combines an attack pattern with service, context, and other properties into a signature attack object. As a result, the IDP system does not need to expend resources inspecting huge sections of network traffic where attacks cannot possibly be, and IDP produces very few false positives.

A *protocol anomaly* is a deviation from protocol standards established by the Internet Engineering Taskforce (IETF) Request for Comment (RFC) process. Traffic that does not adhere to these standards is suspicious because most legitimate applications adhere to the standards, and anomalies can fairly be regarded as purposeful attempts to evade an intrusion detection system (IDS). IDP protocol-anomaly attack objects find traffic that deviates from IETF RFC standards.

When you create rules for the IDP rulebase, you specify:

- A source/destination/service match condition
- Attack objects
- Action
- Notification options

For complete procedures on configuring IDP rulebase rules, see the *IDP Series Administration Guide*.

Related Documentation The following additional related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding IDP Rulebase Rule Match Settings on page 92
- Using Application Identification on page 96
- Using Attack Objects on page 97
- Understanding IDP Rulebase Actions on page 101
- Understanding IDP Rulebase Notification Options on page 103
- IDP Rulebase Example: Using Application Identification on page 107
- IDP Rulebase Example: Specifying the Default Service on page 108
- IDP Rulebase Example: Using Recommended Attack Objects on page 109
- IDP Rulebase Example: Using Recommended Actions on page 110
- Example: Fine-Tuning a Security Policy on page 167

The following related topic is included in the *IDP Series Administration Guide*:

- Modifying IDP Rulebase Rules (NSM Procedure)

Understanding IDP Rulebase Rule Match Settings

The IDP engine inspects the session beginning with the first packet to determine whether the session matches a rule. If the session matches all rule settings for source, destination, service, and VLAN tag ID, the IDP system decodes the traffic and inspects the session packets for the attack objects specified in the rule. If the session matches only some of the rule settings, the rule is not a match.

Table 23 on page 92 provides guidelines for setting IDP rulebase match conditions.

Table 23: IDP Rulebase Match Condition Guidelines

Setting	Guideline
From zone/To zone	Not applicable for standalone IDP Series devices.

Table 23: IDP Rulebase Match Condition Guidelines (*continued*)

Setting	Guideline
Source	<p>Requires one of the specified source IP addresses to match the session for the rule to be applied. You can add address objects for hosts, groups, or network address ranges.</p> <p>In most cases, to detect incoming attacks that target your internal network, specify Any. Specifying Any means you are not using source as a key to your match.</p> <p>To detect traffic from spyware that has affected hosts in your internal network, specify internal network addresses as the source.</p> <p>To detect attacks between two networks you manage, specify multiple addresses. The more specific you are in defining the source and destination of an attack, the more you reduce false positives.</p> <p>NOTE: You must choose between source IP address or user role as match criteria for a rule. You cannot configure both for one rule.</p>
User Role	<p>Requires one of the specified 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.</p> <p>You must choose to configure either source IP address or user role as match criteria for a rule. User role-based rules are evaluated before IP address-based rules. If a user-role based rule matches, the rule is applied and IP address-based rules are not consulted.</p> <p>NOTE: Matching based on user role depends on integration with the Juniper Networks IC Series Unified Access Controller (UAC) appliance.</p>
Destination	<p>Requires one of the specified destination IP addresses to match the session for the rule to be applied. You can add address objects for hosts, groups, or network address ranges.</p> <p>In most cases, specify the hosts or servers you want to protect.</p> <p>Specify Any to not use destination as a key to your match. For example, it would be impossible to predict the destination IP address for traffic resulting from spyware in your internal network. Specify Any for rules that target spyware attacks.</p>
Service	<p>Requires one of the specified services to match the session for the rule to be applied. Services are Application Layer protocols that define how data is structured as it travels across the network. IDP can inspect services that use TCP, UDP, RPC, and ICMP transport layer protocols. If the application running on the destination server uses standard ports, you can select from predefined services. If the application running on the destination server uses nonstandard ports, you must create a custom service object.</p> <p>TIP: Specify Default to match the service(s) specified in the rule attack object(s). If the application identification feature is enabled, the IDP process engine identifies services even if they are running on nonstandard ports.</p> <p>If you disable application identification and specify Default, the IDP process engine assumes that standard ports are used for the service.</p> <p>NOTE: If you disable application identification and your service uses nonstandard ports, you must create custom service objects. For procedures, see the <i>IDP Series Administration Guide</i>.</p> <p>Specify Any to not use service as a key to your match.</p>
VLAN	<p>Requires one of the specified VLAN tags to match the session for the rule to be applied.</p> <p>Specify Any to not use VLAN tag as a key to your match.</p>



TIP: You can use Profiler to identify the hosts and services that are included in your network. In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM online Help.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the IDP Rulebase on page 91
- User-Role-Based Policy Feature Overview on page 94
- Using Application Identification on page 96
- IDP Rulebase Example: Specifying the Default Service on page 108
- IDP Rulebase Example: Using Application Identification on page 107

The following related topic is included in the *IDP Series Administration Guide*:

- Specifying Rule Match Conditions (NSM Procedure)

User-Role-Based Policy Feature Overview

The user role-based policy feature depends on integration with the Juniper Networks IC Series Unified Access Control (UAC) appliance. This feature requires collaboration with the UAC administrator.

The user role-based policy feature enables you to specify user roles as match criteria in IDP rulebase and application policy enforcement (APE) rulebase rules. Matching based on user role rather than IP address both simplifies and finely tunes your rules. In many networks, the IP address is dynamically assigned. To protect your network, you would have to cast a wide net for traffic sources. In most cases, you would specify a subnet mask or specify **Any** source (in the latter case, this means you really are not matching on source). For the purpose of intrusion detection and prevention, a wide net is not necessarily a bad thing: you do want to inspect any session that could potentially contain an attack. Use of role-based rules with a terminal match, however, will improve performance by providing faster matching with specific source targets and rulebase termination. In addition, you are likely to find that user role-based logs are easier to analyze because they provide visibility into the user role associated with an attack event or application usage.

UAC integration with IDP Series devices also improves end user experience authenticating to your network. In a UAC deployment, you use the Host Checker feature to quarantine users with vulnerable hosts. Instead of using a firewall to shut down access to network resources, you can use IDP security policies to enable access and inspect the traffic to guard against threats.

In the APE rulebase, role-based rules are indispensable to supporting the business cases that demand a nuanced approach to application policy enforcement. They enable you

to enforce business policies such as “Contractors, Part-Time, and Temporary employees may not use peer-to-peer filesharing applications; full-time employees may use them, but only with a limited pool of bandwidth.”

To deploy the user role-based feature:

1. Read the release notes for the IDP OS and UAC releases to verify version compatibility requirements.
2. Deploy a UAC solution for user access to the network. For details, see the *Unified Access Control Administration Guide*.
3. Use the UAC user interface to create the user roles you want to use in your security policy:
 - For security rules, you want to leverage results of the Host Checker to map users to roles that identify vulnerabilities, such as “Laptop Users,” “Unauthorized Instant Messenger Installed,” or “Windows XP Patch Required.”
 - For application policy enforcement rules, you want to map users to roles that reflect the business rule, such as Contractor, Part-Time, and Temporary.

For details, see the *Unified Access Control Administration Guide* or UAC online Help.

4. Configure communication between the UAC appliance and the IDP Series device so you can use the IDP role-based policy feature:
 - From the IDP Series side, you use the Appliance Configuration Manager (ACM) to generate a one-time password the UAC appliance can use to connect to the IDP Series device.
 - From the UAC side, you configure the connection to the IDP Series device, specifying the IP address, port 7103, and the one-time password.

For details, see the UAC online Help.

5. Use the IDP Series command-line interface to verify integration. You can use CLI commands to verify connectivity with the IC Series device and to display the user session table, which is populated by the IC Series device.

For details, see the *IDP Series Administration Guide*.

6. In NSM, create a security policy with role-based rules. The roles you specify are the roles created and managed in UAC.
7. Push the security policy from NSM to the IDP Series device.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- IDP Rulebase Example: User-Role-Based Policies on page 104
- APE Rulebase Example: Aggregate and Per-Subscriber Rate Limits on page 136

The following related topics are included in the *IDP Series Administration Guide*:

- Specifying Rule Match Conditions (NSM Procedure)

- Configuring Advanced Settings for the User-Role-Based Policy Feature
- Verifying Integration with an IC Series Unified Access Control Appliance

The following related topic is included in the *IDP Series Deployment Scenarios*:

- Deploying IDP Series with an IC Series Device to Implement User-Role-Based Security Policies

Using Application Identification

The application identification feature enables the IDP engine to detect applications running on standard or nonstandard ports. Port-independent application identification enhances both security and manageability by eliminating the need to manually and comprehensively configure application-port mapping for the service objects and application objects used in the IDP rulebase and APE rulebase rules.

The application identification feature uses application signatures provided by the Juniper Security Center team (J-Security Center) to identify the session application. Beginning with IDP OS Release 5.1, the application identification feature can match extended application signatures used in APE rulebase rules. *Extended application* signatures are also called *nested application* signatures. The predefined extended application signatures developed for IDP OS Release 5.1 include the most prevalent Web 2.0 applications running over HTTP. If your security policy includes APE rules configured to match extended application signatures, the application identification process identifies and generates the following HTTP contexts: `http-url-parsed`, `http-url-parsed-param-parsed`, `http-header-host`, and `http-header-content-type`. The application identification feature can then match application signature patterns in those contexts.

J-Security Center updates application signatures and develops new ones as necessary. Beginning with IDP OS Release 5.1, you can use NSM to browse predefined application objects, predefined extended application objects, and application groups. You can also use NSM to create custom application definitions, if needed. You cannot, however, create custom extended application definitions.

When the application identification feature identifies a new application, it caches the result (the destination address, port, protocol, and service) to reduce processing for subsequent sessions. The application cache and extended application cache are maintained separately.

When the IDP engine processes security policy rules, it examines the session, beginning with the first packet, to identify a match. To match service or application, the IDP engine first compares the session against the application identification cache to identify the application. If the session does not match the application identification cache, the IDP engine processes the session against the application signatures. If the IDP engine is still unable to determine the application, it uses the standard application protocol and port.

In IDP rulebase rules, with application identification enabled, you set the service object in rules to **Default** to allow the application identification feature to identify the correct

service. If you set service to a specific service object, application identification is not applied and the rule is processed using the service object properties.

In APE rulebase rules, with application identification enabled, you set the service object in rules to **Default** and specify rules based on application or extended application. If you disable application identification and specify a match based on application, the IDP engine uses the standard application protocol and port for the application. If the application you are interested in is not listed, you can create a custom application object to match against application properties that you define.

The application identification feature is enabled by default, and we recommend you use this feature. To support lab experimentation and troubleshooting, you can disable application identification and extended application identification, and you can tune the following settings:

- Maximum number of sessions that utilize application identification
- Maximum memory used by application identification
- Maximum memory for saving TCP or UDP packets per session

For information on tuning these parameters, see the *scio const* reference page in the *IDP Series Administration Guide*.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- IDP Rulebase Example: Using Application Identification on page 107
- Using Application Objects on page 121
- Understanding the IDP Rulebase on page 91
- J-Security Center Updates Overview on page 21

The following related topics are included in the *IDP Series Administration Guide*:

- Application Objects Task Summary
- Specifying Rule Match Conditions (NSM Procedure)
- *scio const*

Using Attack Objects

If the session matches rule settings for source, destination, service, and VLAN tag ID, the IDP engine decodes the traffic and inspects the session packets for the attack objects specified in the rule. The following topics provide guidelines for using attack objects in IDP rulebase rules:

- Attack Objects Overview on page 98
- Understanding Predefined Attack Objects and Attack Object Groups on page 99
- Using Attack Object Groups on page 99
- Using Custom Attack Objects on page 100

Attack Objects Overview

When traffic matches an IDP rulebase source/destination/service condition, the IDP engine inspects the traffic for the attack objects you specify.

A *signature attack object* detects known attacks using stateful attack signatures. An attack signature is a pattern that always exists within an attack; if the attack is present, so is the attack signature. With stateful signatures, the IDP engine can look for the specific protocol or service used to perpetrate the attack, the direction and flow of the attack, and the context in which the attack occurs. Stateful signatures produce few false positives because the context of the attack is defined, eliminating huge sections of network traffic in which the attack would not occur.

A *protocol anomaly* identifies unusual activity on the network. It detects abnormal or ambiguous messages within a connection according to the set of rules for the particular protocol being used. Protocol anomaly detection works by finding deviations from protocol standards, most often defined by RFCs and common RFC extensions. Most legitimate traffic adheres to established protocols. Traffic that does not, produces an anomaly, which may be created by attackers for specific purposes, such as evading an IPS. You cannot create protocol anomaly objects. You can specify a predefined protocol anomaly object as a component of a compound attack object.

A *compound attack object* combines multiple signatures and/or protocol anomalies into a single object. Traffic must match all of the combined signatures and/or protocol anomalies to match the compound attack object; you can specify the order in which signatures or anomalies must match. Use compound attack objects to refine your IDP policy rules, reduce false positives, and increase detection accuracy. A compound attack object enables you to be very specific about the events that need to occur before the IDP engine identifies traffic as an attack. You can use **And**, **Or**, and **Ordered and** operations to define the relationship among different attack objects within a compound attack and the order in which events occur.

Attack object definitions also include data fields to help you group and manage attack objects and use them in security policies. These data fields include category, severity, keywords, and a recommended flag.

Predefined attack objects provided by the Juniper Networks Security Center (J-Security Center) team also contain a recommended action for the IDP Series device to take against the attack session.

Custom attack objects are ones you create, if your security policy requires more or less protection, or more or less accounting than what the predefined attack objects provide.

Both predefined and custom attack objects are stored in the attack object database.

When you add attack objects to an IDP rulebase rule, you can add attack objects by group or individually.

Understanding Predefined Attack Objects and Attack Object Groups

The Juniper Networks Security Center (J-Security Center) team has developed more than 600 attack objects and these are included in the attack object database used in IDP security policies.

Table 24 on page 99 describes the attack object groups provided by the J-Security Center.

Table 24: Predefined Attack Object Groups

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. 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.
Response	Contains attack objects where the attack is detected in the server-to-client direction. This group contains a hierarchy of subgroups that includes all of the above categories.

J-Security Center updates the attack object database to provide new attack objects, to revise severities or recommendations, or to remove obsolete attack objects. We recommend you schedule routine, automatic updates.

Using Attack Object Groups

A *dynamic group* contains members that match properties you specify for the group. You use dynamic groups so that an attack database update automatically populates the group with relevant members. This 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 only contain attack objects and not attack groups. Dynamic group members can be either predefined or custom attack objects.

A *static group* is not automatically updated with new members. It contains only the attack objects or groups you have added. Use static groups when you do not want your attack group dynamically populated during NSM updates. For example, if you customize the action for predefined attack objects to meet your company's security policy guidelines, you can create one or more static groups to contain these attack objects. When you perform an NSM attack object update, your static group will not be affected.

There are two types of static groups: predefined static groups and custom static groups. Predefined static groups are categories of groups provided by default.

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

Static groups require more maintenance than dynamic groups because you must manually add or remove attack objects in a static group to manage 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.

Using Custom Attack Objects

The attack objects provided by the Juniper Networks Security Center (J-Security Center) team cover most cases for small business, enterprise, and service provider networks. Your business might encounter cases where you must modify a predefined attack object or create a new one. For example:

- You read a security advisory about a known attack and want to create an attack object that detects the malicious traffic described in that advisory.
- You need to update or improve an existing third-party signature (such as a Snort signature).
- You want to customize an existing signature or protocol anomaly attack object for your local environments. For example, you might need to customize a signature to prevent false positives generated by a specific application running on your network.
- You want to detect specific activity on your network. For example, you might want to detect abnormal traffic (possibly malicious), remote log-ins, or brute force attacks that attempt to guess usernames and passwords.

For a complete tutorial on creating custom attack objects, see the [IDP Series Custom Attack Objects Reference and Examples Guide](#).

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- J-Security Center Updates Overview on page 21
- Understanding the IDP Rulebase on page 91
- IDP Rulebase Example: Using Recommended Attack Objects on page 109
- Exempt Rulebase Example: Exempting an Attack Object on page 115

The following related topic is included in the *IDP Series Administration Guide*:

- Attack Objects Task Summary

Understanding IDP Rulebase Actions

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

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

Predefined attack objects include a recommended action. The recommended action is generally related to attack severity, but other factors are considered. Table 25 on page 101 lists the recommended actions by attack severity.

Table 25: Recommended Action by Attack Severity

Severity	Description	Recommended Action
Critical	Attacks attempt to evade an intrusion prevention system, 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 are obsolete or attempt to obtain noncritical information or scan the network.	None
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

If you choose, you can set a different action. Table 26 on page 101 describes the actions you can set for IDP rulebase rules.

Table 26: IDP Rulebase Actions

Action	Description
None	Inspects for attacks but takes no action against the connection if an attack is found.
Ignore	Does not take action and ignores the remainder of the session.

Table 26: IDP Rulebase Actions (*continued*)

Action	Description
Diffserv Marking	<p>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.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. Therefore, this action has no effect in sniffer mode.</p>
Drop Packet	<p>Drops a matching packet before it can reach its destination but does not close the connection. Use this action in rules focused on 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.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. Therefore, this action has no effect in sniffer mode.</p>
Drop Connection	<p>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.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. Therefore, this action has no effect in sniffer mode.</p>
Close Client and Server	<p>Closes the connection and sends an RST packet to both the client and the server.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and server but does not close the connection.</p>
Close Client	<p>Closes the connection to the client but not to the server.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and server but does not close the connection.</p> <p>NOTE: In VLAN tagged MPLS traffic, the Close Client action drops the connection instead of closing it.</p>
Close Server	<p>Closes the connection to the server but not to the client.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and server but does not close the connection.</p>

If the IDP engine matches an attack, it can take action not only against the current session but also against subsequent network traffic from 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.

Table 27 on page 103 describes IDP rulebase IP actions.

Table 27: IDP Rulebase IP Actions

IP Action	Description
IP Block	<p>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>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
IP Notify	Does not take any action against future traffic but logs the event or sends an alert.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the IDP Rulebase on page 91
- IDP Rulebase Example: Using Recommended Actions on page 110

The following related topic is included in the *IDP Series Administration Guide*:

- Specifying Rule Session Action (NSM Procedure)

Understanding IDP Rulebase Notification Options

You use notification features to help you manage your network, analyze your network security, validate your security policy, and capture forensic evidence of attacks. You can set notification options per rule.

The first time you design a security policy, you might be tempted to log all data for all attacks and let the policy run indefinitely. We recommend you take a more refined approach. Some attack objects are informational only, and others can generate false positives and redundant logs. If you become overloaded with data, you can miss something important. Remember that security policies that generate too many log records are hazardous to the security of your network, as you might discover an attack too late or miss a security breach entirely as a result of having to sift through hundreds of log records.

Excessive logging can also affect throughput, performance, and available disk space. A good security policy generates enough logs to fully document only the important security events on your network.

By default, logging is enabled for IDP rulebase rules. Table 28 on page 104 describes the notification options you can configure. You also have the option to disable logging.

Table 28: 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 and 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, the IDP system 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, the IDP system 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>

For complete procedures on setting IDP rulebase notification options, see the *IDP Series Administration Guide*.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the IDP Rulebase on page 91
- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- IDP Series Logs and Reports in NSM Task Summary

IDP Rulebase Example: User-Role-Based Policies

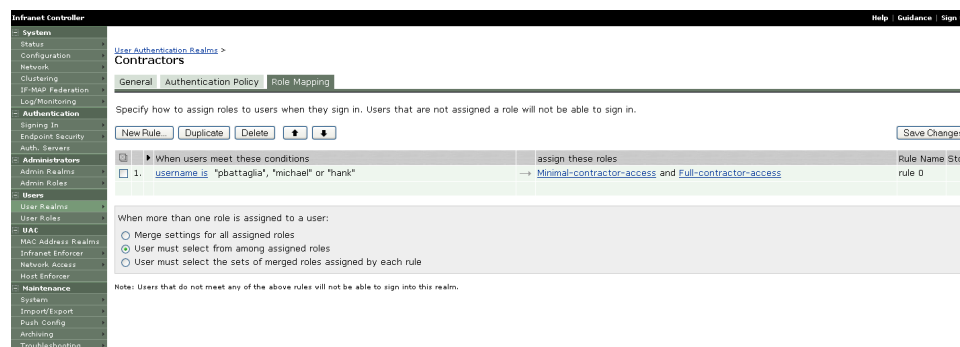
Suppose your enterprise uses Juniper Networks Unified Access Control (UAC) to authenticate access to the corporate network. When you initially rolled out the solution,

Host Checker quarantined and denied network access to many users with noncompliant systems, and you received a lot of negative feedback about end user inconvenience and lost productivity. You can ameliorate these concerns when you deploy the IDP Series device with user session signaling from UAC. When the IDP Series device is protecting your network, users who were formerly flagged for quarantine because Host Checker identified vulnerabilities do not need to be denied access. With role-based IDP security policies, you can adopt a remediation plan that allows access, and even if the vulnerability has been exploited, your network will be protected by the IDP role-focused security policy.

To deploy this solution, follow these basic steps:

1. Read the release notes for the IDP Series device and the IC Series device to verify version compatibility requirements.
2. Deploy a UAC solution for user access to the network. For details, see the *Unified Access Control Administration Guide*.
3. Use UAC to create roles you want to use in your security policy. For security rules, you want to leverage results of the Host Checker to map users with vulnerable systems to roles that identify the vulnerabilities, such as “Laptop Users,” “Unauthorized Instant Messenger Installed,” or “Windows XP Patch Required.” Figure 36 on page 105 shows the IC Series Admin Console Role Mapping page.

Figure 36: IC Series Admin Console: Configuring User Roles



For details on configuring roles and role mapping, see the *Unified Access Control Administration Guide* or UAC online Help.

4. Configure communication between the IC Series device and the IDP Series device so you can use the IDP user-role-based policy feature:
 - From the IDP Series side, you use the Appliance Configuration Manager (ACM) to generate a one-time password the IC Series device will use to connect to the IDP Series device. Figure 37 on page 106 shows the ACM page used to generate a password for the IC Series connection.

Figure 37: ACM: Generating a One-Time Password for the Connection from the IC Series Appliance

Configure NetScreen-Security Manager Communication

In this step, you can configure the Sensor to use a NetScreen-Security Manager. Configuration fields in this page are optional. Additionally, the One Time Password may be reset to authenticate communication between the Sensor and NetScreen-Security Manager. Provide the requested information below to configure the Sensor to use a NetScreen-Security Manager.

Reset One-Time Password: ☐

New One-Time Password: (password is displayed)

Device ID:

Primary NetScreen-Security Manager IP: Port No:

Secondary NetScreen-Security Manager IP: Port No:

Configure IDP IVE Server Communication

IVE one-time password (IVE OTP) may be reset by the user which is used to authenticate certificates of IDP and IVE.

Reset IVE OTP? ☐

[Next Step](#)

- From the IC Series side, you configure the connection to the IDP Series device, specifying the IP address, port 7103, and the one-time password. Figure 38 on page 106 shows the IC Series Admin Console Sensor Configuration page.

Figure 38: IC Series Admin Console: Configuring the Connection to the IDP Appliance

Sensor Configuration

Security Certificates DMI Agent Sensors

Sensor Event Policies

New Sensor Sensor Event Policies Enable Disable Reconnect Refresh

Sensor	Address	Enabled	Status	Notes

For details, see the UAC online Help.

- In NSM, configure IDP rulebase rules that inspect traffic from users with vulnerable systems. Push the security policy to the IDP Series device.

Figure 39 on page 107 shows a rule where the IDP Series device inspects traffic from vulnerable hosts for the relevant Recommended attack objects.

Figure 39: IDP Rulebase: User-Role-Based Rules

No.	ID	Source	Match	Look For	Action	IP Action	Notification
			User Role	Destination	Service	Te...	Attacks
4	1	any	Laptop Users	any	Default		(Recommended)SPYWARE
5	5	any	Unauthorized Instant Messenger Installed	any	Default		(Recommended)CHAT
6	6	any	Windows XP Patch Required	any	Default		(Recommended)Windows

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- User-Role-Based Policy Feature Overview on page 94

The following related topics are included in the *IDP Series Administration Guide*:

- Verifying Integration with an IC Series Unified Access Control Appliance
- Configuring Advanced Settings for the User-Role-Based Policy Feature

The following related topic is included in the *IDP Series Deployment Scenarios*:

- Deploying IDP Series with an IC Series Device to Implement User-Role-Based Security Policies

IDP Rulebase Example: Using Application Identification

This example demonstrates the usefulness of the application identification feature.

Suppose your corporate security policy changes, and you are charged with inspecting peer-to-peer traffic from applications such as Kazaa, Torrent, or eDonkey. To add new rules that inspect peer-to-peer traffic, you would take the following steps:

1. Analyze network traffic to identify peer-to-peer applications running in your network.
2. Research and identify the pattern for every peer-to-peer application.
3. Create signature and port definitions for every peer-to-peer application.
4. Verify the effectiveness of the signatures.
5. Repeat these steps several times for each peer-to-peer application.
6. Continually monitor the network for peer-to-peer traffic that uses nonstandard ports so you can update your signature set to inspect traffic over these ports.

Juniper Networks Security Center saves you much of this work. With predefined attack objects and application identification enabled, you can create a rule whose only elements are the predefined attack object and the service match set to **Default**.

Figure 40 on page 108 is an example of a simple rule that detects any peer-to-peer application.


Figure 40: A Simplified Rule Enabled by the Application Identification Feature

IDP

APE

+

-

	No.	Match			Look For	Action	IP Action	Notification	Comments
		Source	Destination	Service	Attacks				
	2	any	any	Default	[Recommended]P2P	Recommended	None	Logging	Client to Server P2P threats.

When the IDP engine identifies a source/destination/Default service match, it examines the session against the application signatures to determine the application, regardless of which port is used. The IDP system then decodes the traffic and inspects it for the attack objects related to that application.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Using Application Identification on page 96
- Understanding the IDP Rulebase on page 91

IDP Rulebase Example: Specifying the Default Service

This example demonstrates the usefulness of specifying the value **Default** for the service match parameter in IDP rulebase rules.

When you specify a service, you have the option to specify:

- A service object
- **Any**
- **Default**

If you specify the value **Default**, the rule gets the service parameter from the attack object. For example, if the attack object service binding specifies FTP, and you specify the value **Default** for service, the match value is FTP.

Figure 41 on page 108 is an example of a rule where the default service resolves to FTP.

Figure 41: Default Service

IDP

No.	ID	Match				Look For	Action	IP Action	Notification
		Source	Destination	Service	Terminate Ma.	Attacks			
<div>⚠</div> <div>—</div> <div>📄</div> <div>1</div>	3	<div>🌐</div> <div>any</div>	<div>🌐</div> <div>any</div>	<div>🔌</div> <div>Default</div>	<div>☐</div>	<div>🏃</div> <div>[Recommended]FTP</div>	<div>🌿</div> <div>Recommended</div>	<div>🌿</div> <div>None</div>	<div>📄</div> <div>Logging</div>



TIP: With application Identification enabled, the IDP process engine identifies services even if they are running on nonstandard ports.

If you disable application identification and specify **Default**, the IDP process engine assumes that standard ports are used for the service.



NOTE: If you do not enable application identification and your service uses nonstandard ports, you must create a custom service object. For procedures, see the NSM documentation.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- IDP Rulebase Example: Using Application Identification on page 107
- Understanding IDP Rulebase Rule Match Settings on page 92
- Understanding the IDP Rulebase on page 91

IDP Rulebase Example: Using Recommended Attack Objects

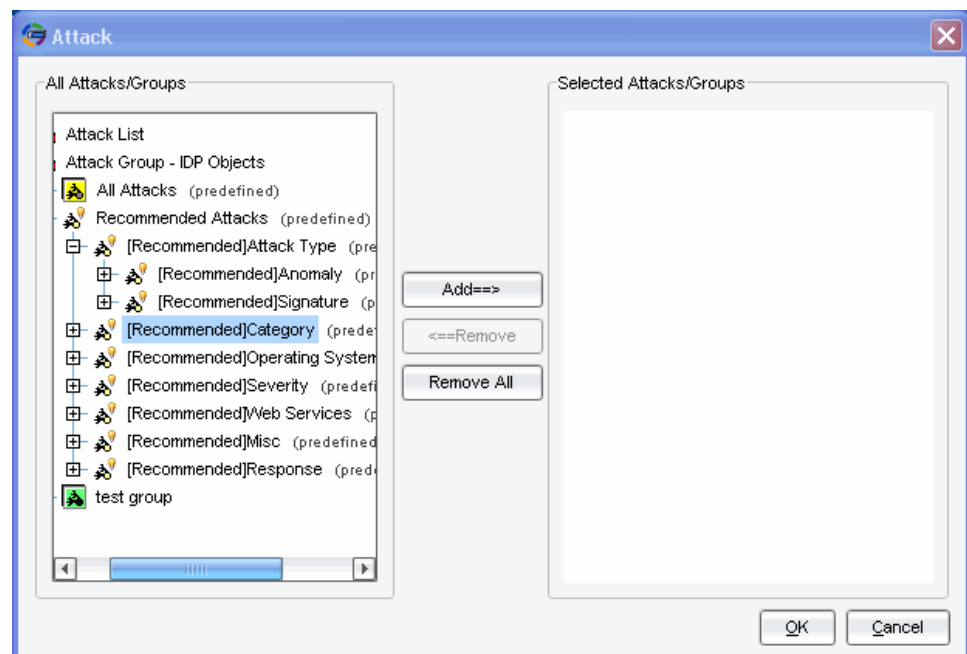
This example demonstrates the usefulness of Juniper Networks Security Center (J-Security Center) recommended attack objects.

When you add attack objects to an IDP rulebase rule, you have the option of adding:

- Predefined attack objects by group
- Recommended predefined attack objects by group
- Custom attacks

Figure 42 on page 109 shows recommended attack objects in the dialog box for adding attack objects to the IDP rulebase.

Figure 42: Recommended Attack Objects



The groups marked **Recommended** have the following special features:

- Recommended attacks have been identified and coded for their recommended purpose by J-Security Center, a world class team of security experts.
- Recommended attack groups are dynamic groups, so members are added or deleted as appropriate during NSM attack database updates.

When you get started with an IDP Series deployment, you should use the recommended attack objects and enable notification for rule matches. Later, you can turn off logging (at your discretion). If you find you need to customize attack object properties, you can make a copy of the recommended attack object and modify it with your required properties. Then you can replace the recommended attack object with the custom attack object in your IDP rulebase rule.



NOTE: If you use a recommended attack object as the basis for a custom attack object, be sure to view the original attack object from time-to-time after attack database updates. If J-Security Center makes changes to the original, you must manually propagate changes to your custom attack object.



BEST PRACTICE: Each attack object specified in an IDP rulebase rule has a performance cost. We recommend that your rules include only the attack objects that are applicable to the rule destination server and only those of a severity that concerns you. We also recommend that you create more rules with a few attack objects in each rather than fewer rules with many attack objects.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*.

- Using Attack Objects on page 97
- J-Security Center Updates Overview on page 21
- Understanding the IDP Rulebase on page 91

IDP Rulebase Example: Using Recommended Actions

This example demonstrates the usefulness of Juniper Networks Security Center (J-Security Center) recommended actions.

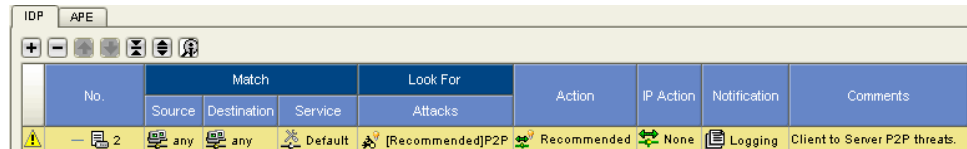
When you specify a rule action, you have the option to specify:

- No action
- A specific action
- The value **Recommended**

Recommended actions are coded in the predefined attack object by the J-Security Center team. The J-Security Center team codes a recommended action in all predefined attack objects, not just the recommended attack objects. When you use the recommended action, you leverage the experience and expertise of the J-Security Center team.

Figure 43 on page 111 shows an IDP rulebase rule with action set to **Recommended**.

Figure 43: Recommended Action



IDP		APE		Match			Look For	Action	IP Action	Notification	Comments
No.	Source	Destination	Service	Attacks							
2	any	any	Default	[Recommended]P2P		Recommended	None	Logging	Client to Server P2P threats.		

When you update the NSM attack database, any changes to recommended actions are also automatically updated.

When you get started with an IDP Series deployment, you should use the recommended actions and enable notification for rule matches. If you find these settings meet your needs, you can turn off logging (at your discretion). If you find you prefer a different action, you can specify a different action.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*.

- Understanding IDP Rulebase Actions on page 101
- Understanding the IDP Rulebase on page 91

CHAPTER 11

The Exempt Rulebase

This chapter explains how to use the Exempt rulebase to reduce logs for false positives. It includes the following topics:

- Understanding the Exempt Rulebase on page 113
- Exempt Rulebase Example: Exempting a Source Destination Pair on page 114
- Exempt Rulebase Example: Exempting an Attack Object on page 115

Understanding the Exempt Rulebase

The Exempt rulebase enhances manageability of the IDP solution by enabling you to categorically exempt traffic segments you know to be safe from processing by the IDP rulebase.

A *false positive*, also known as a false alert, is a situation in which benign traffic causes an intrusion detection system (IDS) to generate an alert. Too many false positives can degrade performance and produce oversized log files.

The IDP engine reduces false positives by using stateful signatures to detect known attacks. A stateful signature knows the pattern and location of the attack, and produces fewer false positives than regular attack signatures because it does not inspect network traffic that cannot contain the attack.

To further increase detection accuracy and reduce false positives, the IDP engine uses:

- Flow tracking to correlate multiple TCP/UDP connections into a single flow to determine the validity of the traffic.
- IP defragmentation and TCP reassembly to reconstruct fragmented traffic.
- Protocol normalization to normalize traffic to a common format for analysis.

Still, a few false positives from your IDS are normal, especially when you are testing new security policies. You can use the Exempt rulebase to manage these cases.

When you create rules for the Exempt rulebase, you specify:

- A source/destination/service match condition
- At least one attack object



NOTE: The Exempt rulebase is a non-terminal rulebase. That is, the IDP process engine processes all rules in the rulebase.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*.

- Understanding the Components of an IDP Security Policy on page 83
- Understanding the Rule-Matching Algorithm on page 85
- Understanding the IDP Rulebase on page 91
- Exempt Rulebase Example: Exempting a Source Destination Pair on page 114
- Exempt Rulebase Example: Exempting an Attack Object on page 115

The following related topics are included in the *IDP Series Administration Guide*.

- Configuring Exempt Rulebase Rules (NSM Procedure)

Exempt Rulebase Example: Exempting a Source Destination Pair

Suppose in your security policy implementation there are schedule phases where your security team probes your internal network for vulnerabilities and you want the IDP Series device to generate logs, and phases where you have put your security policy in place and now want to exclude security team traffic from the generated logs. To support these alternative phases, you can create an Exempt rulebase rule and toggle it off and on.

To create an Exempt rulebase rule:

1. Create address objects that contain the security team IP addresses and the protected servers.
2. Add the Exempt rulebase to your security policy.
3. Add a rule that specifies the source/destination match condition to exempt.
4. Add the All group of attack objects.

Figure 44 on page 114 shows an Exempt rulebase rule.

Figure 44: Exempt Rulebase Rule

No.	ID	Match		Attacks	Comments
		Source	Destination		
1	1	OurSecurity	Protected Servers	All	Security research might produce false positives from engineering servers.

To toggle the rule off, right-click it and select **Disable**.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*.

- Understanding the Exempt Rulebase on page 113

- Example: Fine-Tuning a Security Policy on page 167

The following related topic is included in the *IDP Series Administration Guide*.

- Configuring Exempt Rulebase Rules (NSM Procedure)

Exempt Rulebase Example: Exempting an Attack Object

Suppose your security policy detects HTTP Buffer Overflow: Header attacks on your internal network, but you know this can safely be ignored. You can exempt this traffic from inspection to optimize IDP Series performance and eliminate unnecessary logs.

To exempt an attack object:

1. If you have not done so already, create an address object for your internal network.
2. Add the Exempt rulebase to your security policy.
3. Add a rule that specifies a source that is the internal network and destination that is anywhere.
4. Add the relevant attack object. In this example, add HTTP Buffer Overflow: Header.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*.

- Understanding the Exempt Rulebase on page 113
- Example: Fine-Tuning a Security Policy on page 167

The following related topic is included in the *IDP Series Administration Guide*.

- Configuring Exempt Rulebase Rules (NSM Procedure)

CHAPTER 12

Application Policy Enforcement Rulebase

This chapter explains how you can use the application policy enforcement (APE) rulebase to limit bandwidth available to traffic that matches your rules. It includes the following topics:

- Understanding the APE Rulebase on page 117
- Understanding APE Rulebase Match Conditions on page 118
- Using Application Objects on page 121
- Understanding APE Rulebase Actions on page 128
- Understanding APE Rulebase Notification Options on page 130
- APE Rulebase Example: Using Extended Application Objects on page 131
- APE Rulebase Example: Aggregate and Per-Subscriber Rate Limits on page 136
- APE Rulebase Example: Matching Custom Application Objects on page 137

Understanding the APE Rulebase

The APE rulebase (application policy enforcement) leverages the application identification feature to enable you to manage network traffic based on application. APE rules match source-destination-application criteria. APE rules do not use attack objects.

You can configure rule actions to meet application policy enforcement objectives. For example:

- To use the IDP Series device like an application firewall, you can specify drop or close actions. Matching traffic is terminated at the IDP Series device.
- To set a cap on available bandwidth for disfavored applications or use of certain applications by certain users, you can specify a rate limiting action. When the limit is reached, the IDP Series device begins dropping matching traffic.
- To support deployments where you use other network equipment to implement quality-of-service (QoS) guarantees, you can specify a DiffServ marker action. If a session matches a rule, the IDP engine applies the DSCP marker to the session packets before transmitting them.

Any traffic not terminated by APE rules can be inspected subsequently by the IDP rulebase and other rulebases.

When you create rules for the APE rulebase, you specify:

- Match conditions
- An action
- Notification options

**Related
Documentation**

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding APE Rulebase Match Conditions on page 118
- Using Application Objects on page 121
- Understanding APE Rulebase Actions on page 128
- Understanding APE Rulebase Notification Options on page 130
- APE Rulebase Example: Aggregate and Per-Subscriber Rate Limits on page 136
- APE Rulebase Example: Using Extended Application Objects on page 131
- APE Rulebase Example: Matching Custom Application Objects on page 137

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring the APE Rulebase (NSM Procedure)

Understanding APE Rulebase Match Conditions

The APE rulebase is a terminal rulebase. Rules are evaluated in numerical order. The first rule to match is applied, and subsequent rules are not processed.

If an APE rule matches but the action does not drop the connection, the IDP system also processes additional rulebases to inspect for attacks. If an attack rule identifies the connection to be closed or dropped, that action is taken and the rate-limiting action is not required.

The matching tuple for APE rules includes the following elements:

- Source or user role
- Destination
- Service or the combined list of applications and extended applications
- VLAN tag

The Boolean logic of the matching tuple is as follows:

(src OR user role) AND destination AND vlan AND (service OR application list)



NOTE: You can use the **Any** wildcard to “remove” a property from the tuple. For example, if you specify **Any** for source, destination, or VLAN tag, you are creating a “traffic lane” that treats all traffic matching the specified application the same. However, **Any** has a different significance when building the service or application list. When setting service or application guidelines, be sure to follow the guidelines below.

Table 29 on page 119 provides guidelines for setting IDP rulebase match conditions.

Table 29: APE Rulebase Match Condition Guidelines

Setting	Guideline
From zone/To zone	Not applicable to IDP Series devices.
Source	<p>Requires one of the specified source IP addresses to match the session in order for the rule to be applied. You can add address objects for hosts, groups, or network address ranges.</p> <p>A rule can specify matching criteria for Source IP or user role, but not both. A policy can include rules that match on Source IP and rules that match on user role.</p> <p>NOTE: If a value for user role matches, the source parameter is not used.</p>
User Role	<p>Requires one of the specified user roles to match the session in order for the rule to be applied.</p> <p>A rule can specify matching criteria for Source IP or user role, but not both. In a rulebase, the user role-based rules are evaluated before the IP address-based rules. If a user-role based rule matches, the rule is applied and the IP address-based rules are not consulted.</p> <p>Matching based on user role depends on integration with a Juniper Networks IC Series UAC device.</p>
Destination	Requires one of the specified destination IP addresses to match the session for the rule to be applied. You can add address objects for hosts, groups, or network address ranges.
Service	<p>Requires a match of one of the specified services.</p> <p>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> <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>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 29: APE Rulebase Match Condition Guidelines (*continued*)

Setting	Guideline
Application	<p>Requires one of the specified applications to match the session for the rule to be applied.</p> <p>You use the Application and Extended Application columns to build a list of applications to match the rule. You can specify individual applications or application groups. When you add a group, you are in effect adding its members to the list. The group object itself is not evaluated. 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. If any application or member of a group matches, the rule matches.</p> <p>The predefined list of applications is populated by the application signatures included in J-Security Center signature updates. The application identification feature uses both heuristic methods and signature pattern matching to identify the application regardless of port. Port-independent application identification simplifies rule configuration and ensures that you do not miss applications that are running on nonstandard ports. For this reason, we recommend that you use the application parameter instead of the service parameter whenever possible.</p> <p>Specify Any in the Application column 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>NOTE: 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.</p>
Extended Application	<p>Requires one of the specified <i>extended applications</i> to match the session for the rule to be applied. Extended applications are also called <i>nested applications</i>. 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.</p> <p>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> <p>Specify Any in the Extended Application column 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>
VLAN	<p>Requires one of the specified VLAN IDs to match the session for the rule to be applied.</p> <p>Specifying Any effectively removes VLAN ID from the tuple.</p>



TIP: You can use Profiler to identify the destination servers and services that are included in your network. In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM online Help.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Rule-Matching Algorithm on page 85

- Understanding the APE Rulebase on page 117
- Using Application Identification on page 96
- Using Application Objects on page 121
- User-Role-Based Policy Feature Overview on page 94

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring the APE Rulebase (NSM Procedure)

Using Application Objects

You specify application objects in APE rules as a key element in the matching tuple. This topic provides an overview of application objects and includes the following sections:

- Application Objects Overview on page 121
- Understanding Predefined Application Objects on page 121
- Using Application Groups on page 126
- Using Custom Application Objects on page 127

Application Objects Overview

Application objects are also called *application signatures*. The signature comprises the Layer 7 protocol, protocol contexts, and a DFA pattern found in client-to-server and server-to-client traffic flows. An application object adds program logic to signatures, such as the capability of chaining signatures to create an ordered or unordered compound expression, a maximum number of transactions wherein the signature must occur to be a match, an order value that sets match precedence in cases where multiple signatures are identified, and a unique ID that the system uses both for logical processing and reporting. Extended application objects, also called nested applications, identify Web 2.0 applications running over HTTP.

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) makes predefined application objects and predefined extended application objects available for download to NSM during signature database updates. You use the NSM Object Manager to manage application objects. You specify application objects in APE rules as a key element in the matching tuple. You push the application signatures from NSM to your devices when you push policy updates.

Understanding Predefined Application Objects

J-Security Center makes predefined application objects and predefined extended application objects available for download to NSM during signature database updates. A complete list of application objects is maintained on the J-Security Center [website](#). We recommend that you become familiar with these objects and leverage them in your APE rules as much as possible.

Figure 45 on page 122 shows the NSM Object Manager Predefined Application Objects tab. This view displays the following properties for predefined application objects:

- Name—A unique, descriptive name.
- Application category—A classification used for sorting the list. Not unique.
- Port range—A range of ports on which the application might run. The application is identified only if the server port is within the specified range.
- Application type—A unique identifier used by the application identification feature.
- Port binding—The standard ports known to be used by the application.
- Match order—In case traffic matches protocol, port, and pattern for two or more applications, the match order determines which object is considered the match (the object with the lower match order number is considered the match).

Figure 45: NSM Object Manager: Predefined Application Objects

Application Objects					
Predefined Application Objects	Custom Application Objects	Predefined Extended Application Objects	Application Group Objects		
Name	Application Category	Port Ranges	Application Type	Port Binding	Match Order
HPOVTRACE	ENTERPRISE-INFRASTRUCTURE	TCP:5051-5053	HPOVTRACE	TCP:5051-5053	146
HSS-SSL-TCP	MISC	TCP:0-65535	HSS-SSL-TCP	...	32
HSS-SSL-UDP	MISC	UDP:0-65535	HSS-SSL-UDP	...	150
HTTP	WEB	TCP:0-65535	HTTP	TCP:80,3128,80...	99
ICA-TCP	REMOTE-ACCESS	TCP:0-65535	ICA-TCP	TCP:1494	5
ICA-UDP	REMOTE-ACCESS	UDP:0-65535	ICA	UDP:1604	51
ICQP	SCADA	TCP:102	ICQP	TCP:102	147
ICQ	PEER-TO-PEER-CHAT	TCP:0-65535	ICQ	...	22
IDENT	MISC	TCP:113	IDENT	TCP:113	68
IEC104	SCADA	TCP:2404	IEC104	TCP:2404	155
IMAP	MESSAGING	TCP:0-65535	IMAP	TCP:143	115
IPSEC-IKE-MAIN-AGGR...	ENCRYPTION	UDP:0-65535	IKE	UDP:500	25
IRC	PEER-TO-PEER-CHAT	TCP:0-65535	IRC	TCP:6666,6667...	46
JABBER	PEER-TO-PEER-CHAT	TCP:0-65535	JABBER	TCP:5222	114
JAVA-RMI	REMOTE-COMMAND	TCP:1099	JAVA-RMI	TCP:1099	188
JONDO-PROXY	WEB	TCP:0-65535	JONDO-PROXY	...	140
KADEMLIA-KAD	PEER-TO-PEER-FILE-SHARING	UDP:0-65535	KADEMLIA-KAD	...	83
KADEMLIA-OVERNET	PEER-TO-PEER-FILE-SHARING	TCP:0-65535 UDP:0-65535	KADEMLIA-OVERNET	...	79
KAZAA	PEER-TO-PEER-FILE-SHARING	TCP:0-65535 UDP:0-65535	KAZAA	...	119
KRB4	ENCRYPTION	UDP:0-65535	KRB4	...	113
KRB5	ENCRYPTION	TCP:0-65535 UDP:0-65535	KRB5	TCP:843 UDP:88	27
KUOOO	PEER-TO-PEER-FILE-SHARING	UDP:0-65535	KUOOO	UDP:7000	78
LDAP	ENTERPRISE-INFRASTRUCTURE	TCP:0-65535	LDAP	TCP:389	111
LOTUSNOTES	MESSAGING	TCP:0-65535	LOTUS-NOTES	TCP:1352	134

You can double-click the table entry to view additional details, including the signature pattern regular expression to match in client-to-server and server-to-client directions. Figure 46 on page 123 shows the general properties of the predefined application object for HTTP.

Figure 46: NSM Object Manager: Predefined Application: General Tab

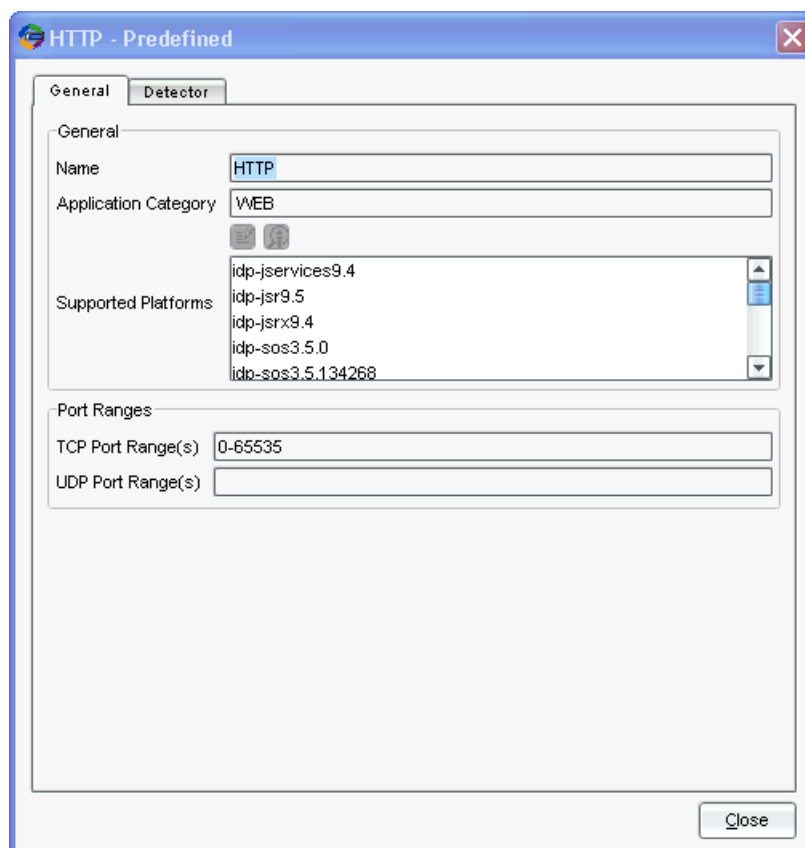


Figure 47 on page 124 shows the signature properties of the predefined application object for HTTP.

Figure 47: NSM Object Manager: Predefined Application: Detector Tab

HTTP - Predefined

General **Detector**

Port Binding

Application Type: HTTP

TCP Port Binding: 80,3128,8000,8080

UDP Port Binding:

Signature

Client-to-server

DFA Pattern: (\\OPTIONS|HEAD|GET|POST|PUT|B?DELETE|TRACE|SEARCH|B?PROPFIND|PROPPATCH|MKCOL|B?COPY|B?MOVE|LOCK|UNLOCK|CHECKOUT|

PCRE Pattern:

Server-to-client

DFA Pattern: (. *HTTP/1\\.([01])s\\.\\.?.?w<([DOCTYPE|w|\\.?.?w<([HTML|w|\\.?.?w<([xml|w|([Content-type|:|.)*

PCRE Pattern:

Minimum data length: 20

Signature Match Order: 122

Close

You can use extended application objects in APE rules if you want to treat various Web 2.0 applications running over HTTP differently. Figure 48 on page 125 shows the NSM Object Manager Predefined Extended Application Objects tab. This view displays the following properties for predefined extended application objects:

- Name—A unique, descriptive name.
- Application category—A classification used for sorting the list. Not unique.
- Extended application ID—A unique identifier. The system uses the unique ID for both logical processing and reporting.
- Application type—A unique identifier used by the application identification feature.
- L7 protocol—Only HTTP is supported.
- Chain order—Indicates whether or not the member signatures are ordered.

Figure 48: NSM Object Manager: Predefined Extended Application Objects

Application Objects

Predefined Application Objects Custom Application Objects **Predefined Extended Application Objects** Application Group Objects

Name	Application Category	Ext ID	Application Type	L7 Protocol	Chain Order
MYSPACE	SOCIAL-NETWORKING	316	MYSPACE	HTTP	No
TWITTER	SOCIAL-NETWORKING	317	TWITTER	HTTP	No
BEBO	SOCIAL-NETWORKING	321	BEBO	HTTP	No
CLASSMATES	SOCIAL-NETWORKING	322	CLASSMATES	HTTP	No
Hi5	SOCIAL-NETWORKING	329	Hi5	HTTP	No
DOOF	SOCIAL-NETWORKING	290	DOOF	HTTP	No
BLOGGER-POST	SOCIAL-NETWORKING	343	BLOGGER-POST	HTTP	No
MYSPACE-MAIL	SOCIAL-NETWORKING	351	MYSPACE-MAIL	HTTP	No
MYSPACE-CHAT	SOCIAL-NETWORKING	352	MYSPACE-CHAT	HTTP	No
MYSPACE-VIDEO	SOCIAL-NETWORKING	360	MYSPACE-VIDEO	HTTP	No
VKONTAKTE	SOCIAL-NETWORKING	501	VKONTAKTE	HTTP	No
MIKI	SOCIAL-NETWORKING	444	MIKI	HTTP	No
TIANYA	SOCIAL-NETWORKING	445	TIANYA	HTTP	No
KAXIN001	SOCIAL-NETWORKING	447	KAXIN001	HTTP	No
ODNOKLASSNIKI	SOCIAL-NETWORKING	448	ODNOKLASSNIKI	HTTP	No
RENREN	SOCIAL-NETWORKING	449	RENREN	HTTP	No
ADULTFRIENDFINDER	SOCIAL-NETWORKING	480	ADULTFRIENDFINDER	HTTP	No
TARINGA	SOCIAL-NETWORKING	481	TARINGA	HTTP	No
BADOO	SOCIAL-NETWORKING	483	BADOO	HTTP	No
NING	SOCIAL-NETWORKING	484	NING	HTTP	No
NETLOG	SOCIAL-NETWORKING	492	NETLOG	HTTP	No
HYVESDOTNL	SOCIAL-NETWORKING	506	HYVESDOTNL	HTTP	No
PLENTYOFFISH	SOCIAL-NETWORKING	508	PLENTYOFFISH	HTTP	No
NATEON	SOCIAL-NETWORKING	403	NATEON	HTTP	No
BLOGSPOT-POST	SOCIAL-NETWORKING	413	BLOGSPOT-POST	HTTP	No
PING-FM	SOCIAL-NETWORKING	509	PING-FM	HTTP	No

You can double-click the table entry to view additional details, including the matching HTTP context, signature pattern, and client-to-server or server-to-client direction. Figure 49 on page 126 shows the properties of the HTTP:Facebook-Access application object.

Figure 49: NSM Object Manager: Extended Application Details

The screenshot shows the 'FACEBOOK-ACCESS - Predefined' dialog box with the 'General' tab selected. The fields are as follows:

- Name: FACEBOOK-ACCESS
- L7 Protocol: HTTP
- Chain Order: No
- Application type: FACEBOOK-ACCESS
- Maximum Transactions: none
- Signature Match Order: 33323

Below the fields is a 'Members' section with a table:

Member /	Context	pattern	direction
m01	http-header-host	(.*)?(facebook\.com fbcdn\.net)	CTS

At the bottom right is a 'Close' button.

An application signature can include one member or more members in a compound signature. Double-click the table row entry for the member to display its details. Figure 50 on page 126 shows the properties of the HTTP:Facebook-Access application object.

Figure 50: NSM Object Manager: Extended Application Member Details

The screenshot shows the 'Signature' dialog box with the following fields:

- Member: m01
- Context: http-header-host
- pattern: (.*)?(facebook\.com|fbcdn\.net)
- direction: CTS

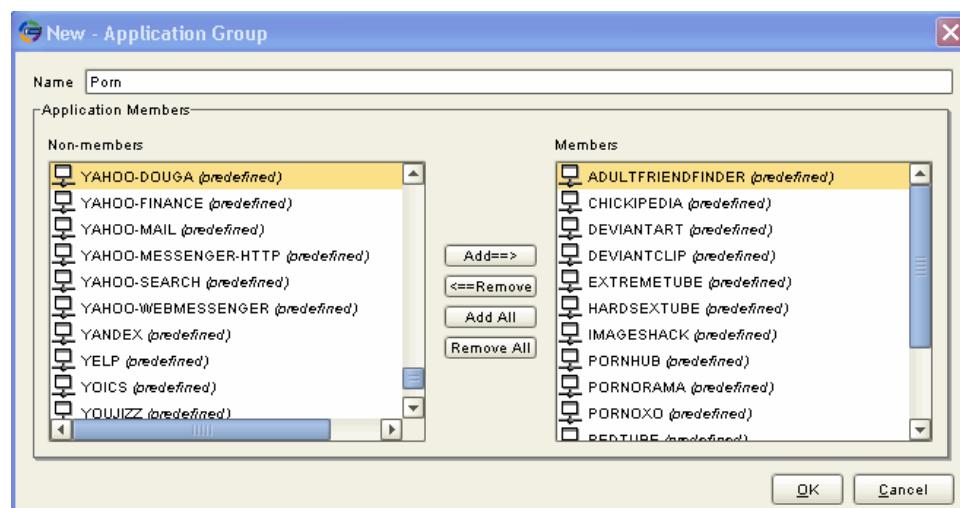
At the bottom right is a 'Close' button.

Using Application Groups

Application groups are administrative objects you can use to simplify rule configuration. A group comprises application objects that you want to treat the same—that is, you want to apply the same action to matching traffic. Figure 51 on page 127 shows the Application Group dialog box. In the Non-Members box, applications are listed first, followed by extended applications. You can nest a group within another group. In the Application

Group dialog box, the icons next to group object names and application object names differ so you can distinguish the two when you browse the lists.

Figure 51: NSM Object Manager: Application Group Dialog Box



Using Custom Application Objects

You can create rules for most business cases with the predefined application objects provided by J-Security Center. In some cases, you might need to manage traffic for applications not yet supported by J-Security Center. First, check with your Juniper Networks representative to see if a predefined application object is forthcoming. If support for your application object is not forthcoming, you can use the NSM Object Manager to define a custom application object. You can then specify that object as a match for APE rules. Figure 52 on page 128 shows the Custom Application dialog box.

Figure 52: NSM Object Manager: Custom Application Dialog Box

The screenshot shows a dialog box titled "Aspera FASP - Custom". It has two tabs: "General" and "Detector". The "General" tab is selected. Inside the "General" tab, there are three main sections: "Name" with the value "Aspera FASP", "Application Category" with the value "File-Server", and "Supported Platforms" with the value "idp5.1.0". Below these, there is a "Port Ranges" section with two fields: "TCP Port Range(s)" with the value "22,33001" and "UDP Port Range(s)" with the value "0-65535". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Apply".

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- J-Security Center Updates Overview on page 21
- Understanding the APE Rulebase on page 117
- APE Rulebase Example: Using Extended Application Objects on page 131
- APE Rulebase Example: Matching Custom Application Objects on page 137

The following related topics are included in the *IDP Series Administration Guide*:

- Application Objects Task Summary

Understanding APE Rulebase Actions

Actions are responses to sessions that match the source/destination/service or source/destination/application condition.

Table 30 on page 129 describes the actions you can specify for application policy enforcement (APE) rulebase rules.

Table 30: IDP Rulebase Actions

Action	Description
None	Does not perform rate limiting. Logs generated for traffic that match this rule display Accepted .
Drop Connection	<p>Drops the connection without sending an RST packet to the sender, thereby preventing the traffic from reaching its destination. Use this action to drop connections for traffic that is not prone to spoofing.</p> <p>Logs generated for traffic that match this rule display Drop Connection.</p> <p>NOTE: In sniffer mode, this action has no effect because the IDP Series device is not in the path of network traffic.</p>
Close Client	<p>Closes the connection to the client but not to the server.</p> <p>Logs generated for traffic that match this rule display Close Client.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and the server, but it does not close the connection.</p>
Close Server	<p>Closes the connection to the server but not to the client.</p> <p>Logs generated for traffic that match this rule display Close Server.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and the server, but it does not close the connection.</p>
Close Client and Server	<p>Closes the connection and sends an RST packet to both the client and the server.</p> <p>Logs generated for traffic that match this rule display Close.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and the server, but it does not close the connection.</p>
DiffServ Marking	<p>Assigns the DiffServ 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> <p>Logs generated for traffic that match this rule display DiffServ.</p> <p>NOTE: In sniffer mode, this action has no effect because the IDP Series device is not in the path of network traffic.</p>

Table 30: IDP Rulebase Actions (*continued*)

Action	Description
Rate Limit	<p>Rate limits set an aggregate limit for all matching sessions. If a session matches an APE rule in which a rate limit has been set, the IDP engine performs a rate-limit check. If the limit is not reached, the IDP Series device forwards the packets. If the limit is reached, the IDP Series device behaves as if no bandwidth is available: it drops packets until the aggregate bandwidth falls below the limit. When the IDP Series device drops packets, the TCP or UDP endpoints identify the packet loss and slow the transmission rate.</p> <p>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> <p>If you implement user-role-based rules, you can apply rate limiting to all users who belong to the specified role or to individual users who belong to the specified role. By default, rate limiting is applied to all users who belong to the specified role. In this case, you would configure a larger limit. You can change this setting with the command-line interface. If you change the default to enable rate limiting per user, configure a smaller limit.</p> <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 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.</p> <p>Logs generated for traffic that match this rule display Rate Limit and traffic direction (c2s or s2c).</p> <p>NOTE: In sniffer mode, this action has no effect because the IDP Series device is not in the path of network traffic.</p>
DiffServ Marking & Rate Limiting	Takes both actions described above.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the APE Rulebase on page 117

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring the APE Rulebase (NSM Procedure)
- Enabling Per-User Rate Limiting for User-Role-Based Rules

Understanding APE Rulebase Notification Options

Notification options determine whether the IDP Series device generates logs and alerts when a session matches a rule. When enabled, the IDP Series device generates a log that the client-to-server or server-to-client rate limit was reached. Logging is enabled by default. Table 31 on page 131 describes the notification options.

Table 31: APE 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. <p>You also have the option to disable logging.</p>
Related Documentation	<p>The following related topic is included in the <i>IDP Series Concepts and Examples Guide</i>:</p> <ul style="list-style-type: none"> • Understanding the APE Rulebase on page 117 <p>The following related topic is included in the <i>IDP Series Administration Guide</i>:</p> <ul style="list-style-type: none"> • Configuring the APE Rulebase (NSM Procedure)

APE Rulebase Example: Using Extended Application Objects

This example shows the usefulness of using extended application objects in APE rules when you want to treat different Web 2.0 applications differently. In IDP OS Release 5.1, the application identification feature is capable of identifying many of the most widely used Web 2.0 applications that run over HTTP. The Juniper Networks Security Center (J-Security Center) provides predefined application signatures for many Web 2.0 applications.

Let's assume you are a network administrator in the IT department of a large enterprise company. Your company executives have decided that it is in the company interest to allow employees basic access to Facebook because it will enable employees to maintain relationships and contacts that serve them professionally. However, the executives have deemed that employees should not use the corporate network to access to Facebook games or multimedia content.

To create a policy that enforces these business objectives:

1. Go to the [J-Security Center website](#) to review the list of predefined application signatures, including a number of Facebook application signatures. Make note of the signatures you want to allow and the ones you want to block.
2. Use the NSM Object Manager Application Object viewer to browse the list of predefined extended application objects. Figure 53 on page 132 shows the NSM Object Manager Application Object viewer.

Figure 53: NSM Object Manager: Predefined Extended Application Objects

Application Objects

Predefined Application Objects Custom Application Objects **Predefined Extended Application Objects** Application Group Objects

Name	Application Category	Ext ID	Application Type	L7 Protocol	Chain Ord
MYSpace	SOCIAL-NETWORKING	316	MYSpace	HTTP	No
Twitter	SOCIAL-NETWORKING	317	Twitter	HTTP	No
Bebo	SOCIAL-NETWORKING	321	Bebo	HTTP	No
Classmates	SOCIAL-NETWORKING	322	Classmates	HTTP	No
Hi5	SOCIAL-NETWORKING	329	Hi5	HTTP	No
Doof	SOCIAL-NETWORKING	290	Doof	HTTP	No
Blogger-Post	SOCIAL-NETWORKING	343	Blogger-Post	HTTP	No
MySpace-Mail	SOCIAL-NETWORKING	351	MySpace-Mail	HTTP	No
MySpace-Chat	SOCIAL-NETWORKING	352	MySpace-Chat	HTTP	No
MySpace-Video	SOCIAL-NETWORKING	360	MySpace-Video	HTTP	No
Vkontakte	SOCIAL-NETWORKING	501	Vkontakte	HTTP	No
Mixi	SOCIAL-NETWORKING	444	Mixi	HTTP	No
Tianya	SOCIAL-NETWORKING	445	Tianya	HTTP	No
Kaixin001	SOCIAL-NETWORKING	447	Kaixin001	HTTP	No
Odnoklassniki	SOCIAL-NETWORKING	448	Odnoklassniki	HTTP	No
Renren	SOCIAL-NETWORKING	449	Renren	HTTP	No
AdultFriendFinder	SOCIAL-NETWORKING	480	AdultFriendFinder	HTTP	No
Taringa	SOCIAL-NETWORKING	481	Taringa	HTTP	No
Badoo	SOCIAL-NETWORKING	483	Badoo	HTTP	No
Ning	SOCIAL-NETWORKING	484	Ning	HTTP	No
Netlog	SOCIAL-NETWORKING	492	Netlog	HTTP	No
HyvesDotNL	SOCIAL-NETWORKING	506	HyvesDotNL	HTTP	No
PlentyOfFish	SOCIAL-NETWORKING	508	PlentyOfFish	HTTP	No
Nateon	SOCIAL-NETWORKING	403	Nateon	HTTP	No
Blogspot-Post	SOCIAL-NETWORKING	413	Blogspot-Post	HTTP	No
Ping-FM	SOCIAL-NETWORKING	509	Ping-FM	HTTP	No

- Double-click a table row to display the details of the application object. You want to learn about the HTTP contexts and patterns that define the application objects so you understand what traffic would be dropped or permitted if you create rules to drop or permit them. Figure 54 on page 133 shows the properties of the HTTP:Facebook-Access application object. Note its signature is found in the HTTP header host details of an HTTP client request.

Figure 54: NSM Object Manager: Extended Application Details

The screenshot shows a window titled "FACEBOOK-ACCESS - Predefined" with a "General" tab. The window contains several configuration fields and a "Members" table.

General Tab Fields:

- Name: FACEBOOK-ACCESS
- L7 Protocol: HTTP
- Chain Order: No
- Application type: FACEBOOK-ACCESS
- Maximum Transactions: none
- Signature Match Order: 33323

Members:

Buttons: +, -, (icon)

Member /	Context	pattern	direction
m01	http-header-host	(.*)?(facebook\.com fbcdn\.net)	CTS

Close

Figure 55 on page 134 shows the properties of the HTTP:Facebook-App application object. Note its signature is found in the parsed parameters of an HTTP client request.

Figure 55: NSM Object Manager: Extended Application Details

The screenshot shows the 'FACEBOOK-APP - Predefined' window in the NSM Object Manager. The 'General' tab is active, displaying the following configuration:

- Name: FACEBOOK-APP
- L7 Protocol: HTTP
- Chain Order: No
- Application type: FACEBOOK-APP
- Maximum Transactions: none
- Signature Match Order: 32976

Below the configuration fields is the 'Members' section, which contains a table with the following data:

Member	Context	pattern	direction
m02	http-url-parsed-param-parsed	(/ap\.php\?i=.*\.\?v=app_\d+)	CTS

The window has a 'Close' button at the bottom right.

After you have familiarized yourself with the Facebook application signatures and understand how they will be identified by the application identification feature, you are ready to create APE rules that use them.

- Configure APE rules that use the extended application object in a way that meets your application usage policy objectives. Figure 56 on page 134 shows a set of rules that distinguish between “acceptable” Facebook usage and “unacceptable” Facebook usage. Rule 1 allows access to the Facebook website and Facebook mail in order to allow employees to use the popular site to keep in touch with colleagues and professional acquaintances. Rule 2 drops all other Facebook traffic.

Figure 56: APE Rulebase: Using Extended Applications

The screenshot shows the 'APE' tab in the configuration window. It displays a table of rules with the following columns: No., Source, User Role, Destination, Service, Application, Extended Application, Action, Notification, VLAN Tag, Severity, and Comments.

No.	Source	User Role	Destination	Service	Application	Extended Application	Action	Notification	VLAN Tag	Severity	Comments
1	any	any	any	Default	any	FACEBOOK:ACCESS FACEBOOK:MAIL	None	Logging	any	Default	Acceptable Facebook.
2	any	any	any	Default	any	FACEBOOK:CHAT FACEBOOK:APP FARMVILLE	Drop Connection	Logging	any	Default	Unacceptable Facebook.

- (Optional) Figure 56 on page 134 depicts APE rules where applications were added one-at-a-time. Alternatively, you can create and maintain application groups based on classifications that you choose. Figure 57 on page 135 and Figure 58 on page 135

show groups created to support the business classifications discussed in this example—acceptable and unacceptable Facebook.

Figure 57: NSM Object Manager: Creating Application Groups

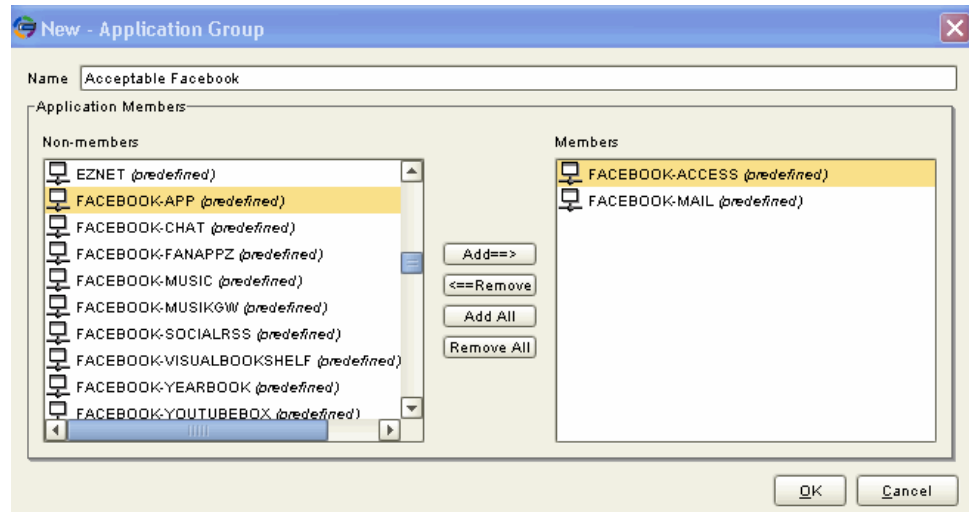


Figure 58: NSM Object Manager: Creating Application Groups

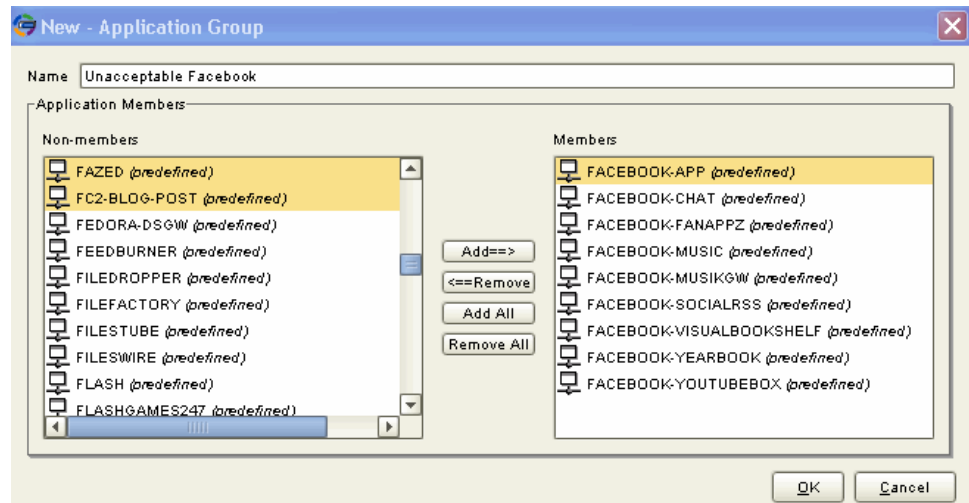


Figure 59 on page 135 shows an APE policy that uses application groups. Note that you use the Application column to add application groups to an APE rule.

Figure 59: APE Rulebase: Using Application Groups

Zone based Firewall IDP APE											
<div><div></div><div></div><div></div><div></div><div></div><div></div></div>											
No.	Match						Action	Notification	VLAN Tag	Severity	Comments
	Source	User Role	Destination	Service	Application	Extended Application					
1	any	any	any	Default	Acceptable Facebook	any	None	Logging	ANY Any	Default	Acceptable Facebook.
2	any	any	any	Default	Unacceptable Facebook	any	Drop Connection	Logging	ANY Any	Default	Unacceptable Facebook.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Using Application Objects on page 121

- Understanding the APE Rulebase on page 117

The following related topic is included in the *IDP Series Administration Guide*:

- Application Objects Task Summary
- Configuring the APE Rulebase (NSM Procedure)

APE Rulebase Example: Aggregate and Per-Subscriber Rate Limits

This example uses an Internet café business to show the usefulness of user-role-based rules. Let's assume that you own an Internet café, and you have deployed a Juniper Networks IC Series UAC solution with a firewall to create a captive portal. A captive portal redirects users to a Web page where they must enter credentials to access the WWW. As your business becomes more popular, you start hearing complaints about network performance during periods of peak use. You take great pride in your service and are particularly distressed when you hear your business customers tell you that the poor network performance is interfering with important work. You think it would be a good idea to have a separate "traffic lane" for users who need Internet access for important work.

You can deploy the IDP Series device in your network and use the application policy enforcement (APE) rulebase to create premium and economy "traffic lanes". Premium customers pay a higher rate and receive unlimited bandwidth. Economy customers pay a lower rate and are subject to a bandwidth rate limit.

To deploy this solution:

1. Deploy a Juniper Networks IC Series UAC and firewall to create a captive portal and manage user access to the Internet.
2. Use the IC Series administration console to map users to roles, including:
 - Premium—Customers who pay extra for unlimited access.
 - Economy—Customers who pay for basic service.
3. Configure communication between the IC Series device and the IDP Series device so you can use the IDP Series user-role-based policy feature.
4. Use NSM to configure APE rules.

Figure 60 on page 136 shows a set of rules that create traffic lanes for tiered access. Note that all matching parameters are set to Any except user role. This policy does not guarantee premium users a specific rate, but it conserves bandwidth for use by premium users by capping bandwidth for non-premium users.

Figure 60: APE Rulebase: User-Role-Based Rules to Support Tiered Access

No.	Match						Action	Notification	VLAN	Severity	Comments
	Source	User Role	Destin.	Service	App.	Extended App.					
1	any	Premium	any	Default	any	any	None	Log...	RVV...	Default	...
2	any	Economy	any	Default	any	any	Rate Limit (262144 kbps, 262144 kbps)	Log...	RVV...	Default	Aggregate rate limit for Economy users.

In the example above, users who belong to the Economy role have unlimited access until the aggregate of the bandwidth used by all Economy role users exceeds the rate limit. At that point, the IDP Series device begins dropping packets until the aggregate for the role falls below the rate limit. If you prefer, you can enable a system-wide option that enforces the rate limit on each user who belongs to the role. You might prefer to enforce rate limiting this way if you want to disclose to users who purchase the Economy package that their bandwidth is capped at a specific rate. You use the CLI to enable per-subscriber rate limiting, and you create an APE rule that sets the agreed upon rate. Figure 61 on page 137 shows the APE rules with a rate limit for per-subscriber enforcement. If any Economy user exceeds 256 kbps client-to-server utilization, the IDP Series device drops packets from the user's session until the rate falls below the threshold.

Figure 61: APE Rulebase: User-Role-Based Rules to Support Tiered Access

No.	Match						Action	Notification	VLAN Tag	Severity	Comments
	Source	User Role	Destination	Service	Applica.	Extended App.					
1	any	Premium	any	Default	ANY	any	None	Log...	ANY	Default	...
2	any	Economy	any	Default	ANY	any	Rate Limit (256 kbps, 512 kbps)	Log...	ANY	Default	Per subscriber rate limit.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the APE Rulebase on page 117
- User-Role-Based Policy Feature Overview on page 94

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring the APE Rulebase (NSM Procedure)
- Enabling Per-User Rate Limiting for User-Role-Based Rules

APE Rulebase Example: Matching Custom Application Objects

This example shows the usefulness of support for custom application objects.

You can create rules for most business cases with the predefined application objects provided by J-Security Center. J-Security Center makes predefined application objects and predefined extended application objects available for download to NSM during signature database updates. For a list of predefined application objects, see the J-Security Center [website](#). In some cases, you might need to manage traffic for applications not yet supported by J-Security Center. First, check with your Juniper Networks representative to see if a predefined application object is forthcoming. If support for your application object is not forthcoming, you can use the NSM Object Manager to define a custom application object. You can then specify that object as a match for APE rules.

For example, suppose you use a final “catch-all” APE rule that drops any traffic not permitted by previous rules. You learn that your organization intends to use Aspera Software’s [FASP protocol](#) instead of FTP to transport large files, so you need to explicitly permit FASP. After checking with J-Security Center to see if a predefined application object is under development, you decide to create your own. Creating a custom application is not a trivial task. To do so, you need to discover the following information:

- Protocol and port—Usually well understood through vendor documentation or network security community websites.
- Signature pattern—Sometimes you can find packet capture files (pcaps) available through network security community websites, such as [Wireshark](#) or [pcap](#). If none can be found, you might have to use a packet capture utility to create your own pcaps and discover the signature pattern.
- Relationship to other application objects—When you create an application object, you specify a signature match order. If traffic matches multiple objects, an application object with the lower signature match-order number is considered the match. You should become aware of traffic that uses the same ports or possibly the same matching pattern.

Use the predefined application naming conventions as a guide to completing the name and category properties. Figure 62 on page 138 shows basic information for FASP.

Figure 62: NSM Object Manager: Custom Application Object

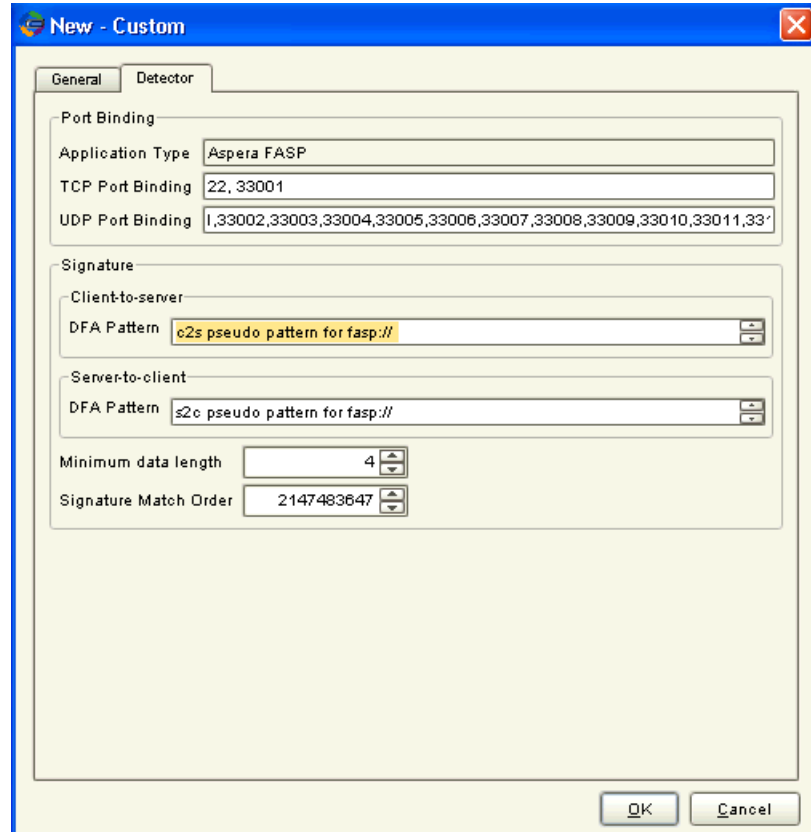
The screenshot shows a window titled "Aspera FASP - Custom" with a close button (X) in the top right corner. Inside the window, there are two tabs: "General" and "Detector". The "General" tab is selected and contains the following fields:

- Name:** Aspera FASP
- Application Category:** File-Server
- Supported Platforms:** idp5.1.0
- Port Ranges:**
 - TCP Port Range(s):** 22,33001
 - UDP Port Range(s):** 0-65535

At the bottom of the dialog, there are three buttons: "OK", "Cancel", and "Apply".

Figure 63 on page 139 shows a pseudo DFA pattern signature for FASP.

Figure 63: NSM Object Manager: Custom Application Object



After you create the custom application object, it is available to be added to APE rules. Figure 64 on page 139 shows the custom application in the APE rulebase application list.

Figure 64: APE Rulebase: Adding a Custom Application Object

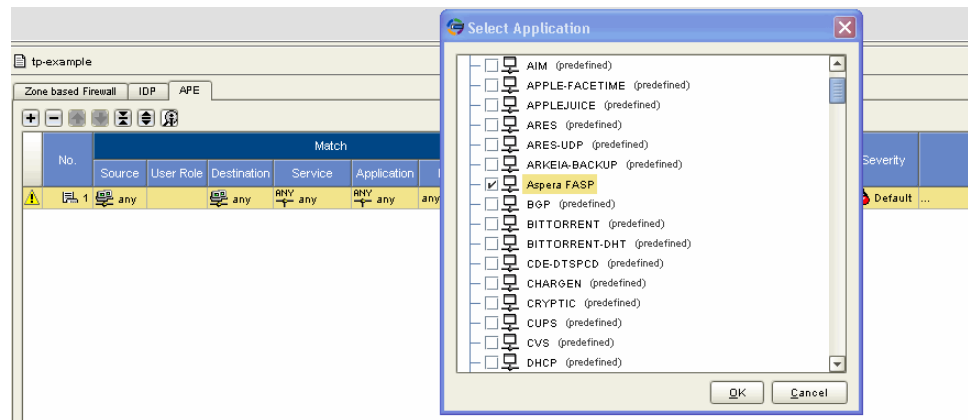


Figure 65 on page 140 shows a set of APE rules. Here, we have added the rule that permits FASP before the catch-all rule that drops all traffic not permitted by previous rules.

Figure 65: APE Rulebase: Rule Order

Zone based Firewall IDP APE										
No.	Match						Action	Notification	VLAN Tag	Severity
	Source	User Role	Destination	Service	Application	Extended Application				
1	any	any	any	Default	Aspera FASP	any	None	Logging	Any	Default
2	any	any	any	any	any	any	Drop Connection	Logging	Any	Default

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Using Application Objects on page 121

The following related topic is included in the *IDP Series Administration Guide*:

- Application Objects Task Summary

CHAPTER 13

The Backdoor Rulebase

This chapter explains how the Backdoor rulebase protects your network and provides guidelines for configuring Backdoor rulebase rules. It includes the following topics:

- Understanding the Backdoor Rulebase on page 141
- Understanding Backdoor Rulebase Match Settings on page 143
- Understanding the Backdoor Rulebase Operation Setting on page 144
- Understanding Backdoor Rulebase Actions on page 144
- Understanding Backdoor Rulebase Notification Options on page 145
- Backdoor Rulebase Example: netcat on page 146

Understanding the Backdoor Rulebase

The Backdoor rulebase detects the kind of interactive traffic produced during backdoor attacks.

A *backdoor* is a mechanism installed on a host computer that facilitates unauthorized access to the system. Attackers who have already compromised a system can install a backdoor to make future attacks easier. When attackers type commands to control a backdoor, they generate interactive traffic.

Unlike antivirus software, which scans for known backdoor files or executable files on the host system, the IDP engine detects the interactive traffic that is produced when backdoors are used. Interactive programs often transmit several short IP packets containing individual keystrokes and their echoes, reflecting the real-time actions of a user (or an attacker).

When detection is enabled, the IDP engine detects traffic that exceeds the interactive traffic thresholds you set as runtime parameters. Figure 66 on page 142 shows the backdoor detection settings in the NSM Device Manager configuration editor.

Figure 66: NSM Device Manager: Sensor Settings > Run-Time Parameters

Table 32 on page 142 shows the defaults for backdoor detection runtime parameters. You can tune these parameters if safe traffic in your network triggers false positives.

Table 32: Backdoor Detection Runtime Parameters

Parameter	Default
Minimum interval between consecutive small packets (microseconds)	20,000
Maximum interval between consecutive small packets (microseconds)	2,000,000
Byte threshold for packet sizes in a backdoor connection (bytes)	20
Minimum number of data carrying TCP packets (number)	20
Minimum percentage of back-to-back small packets (percentage)	20
Ratio of small packets to the total packets (percentage)	20

Detecting the signs of interactive traffic ensures that the IDP Series device can detect all backdoors, both known and unknown. If the IDP Series device detects interactive traffic, it can perform actions against the connection to prevent the attacker from further compromising your network.

When you create rules for the Backdoor rulebase, you specify:

- A source/destination/service match condition
- Operation
- Action
- Notification options

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Components of an IDP Security Policy on page 83

- Understanding Backdoor Rulebase Match Settings on page 143
- Understanding the Backdoor Rulebase Operation Setting on page 144
- Understanding Backdoor Rulebase Actions on page 144
- Understanding Backdoor Rulebase Notification Options on page 145
- Backdoor Rulebase Example: netcat on page 146

The following related topics are included in the *IDP Series Administration Guide*:

- Configuring Backdoor Rulebase Rules (NSM Procedure)
- Modifying the IDP Series Device Configuration

Understanding Backdoor Rulebase Match Settings

Backdoor rulebase rules are triggered when source, destination, and service for the traffic match the rule.

To detect incoming interactive traffic, set the source to **Any** and the destination to the IP address of network device you want to protect.

To detect outgoing interactive traffic, set the source to the IP address of the network device you want to protect and the destination to **Any**.

Specify not only the services on the network device you want to protect but also interactive services that can be installed and used by attackers.



NOTE: Including Telnet, SSH, RSH, NetMeeting, or VNC as services can result in false positives because these services are used to legitimately control a remote system. We recommend that you not include these services in your service list.



TIP: You can use two rules to protect a large number of servers. Configure rule 1 to match services you do not want to detect and set Operation to **Ignore**. Configure rule 2 to match any traffic and set Operation to **Detect**.



TIP: In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM documentation.



NOTE: The Backdoor rulebase is a terminal rulebase—that is, Backdoor rules are inherently terminal rules. If a Backdoor rule matches, the IDP engine does not process subsequent rules.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Rule-Matching Algorithm on page 85
- Understanding the Backdoor Rulebase on page 141

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Backdoor Rulebase Rules (NSM Procedure)

Understanding the Backdoor Rulebase Operation Setting

The Backdoor rulebase operation setting is a toggle between Ignore and Detect.

Use **Ignore** to whitelist services for accepted forms of interactive traffic, such as Telnet, SSH, RSH, NetMeeting, or VNC.

Use **Detect** for all other interactive traffic.

List the ignore rule first, followed by the detect rule.



TIP: You can use two rules to protect a large number of servers. Configure rule 1 to match services you do not want to detect and set Operation to **Ignore**. Configure rule 2 to match any traffic and set Operation to **Detect**.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Backdoor Rulebase on page 141

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Backdoor Rulebase Rules (NSM Procedure)

Understanding Backdoor Rulebase Actions

By default, Backdoor rulebase rules accept and log traffic that matches the rule. If you choose, you can set a different action. Table 33 on page 145 describes the actions you can set for Backdoor rulebase rules.

Table 33: Backdoor Rulebase Actions

Action	Description
Accept	Accepts the interactive traffic.
Drop Connection	Drops the interactive 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	<p>Closes the connection and sends an RST packet to both the client and the server.</p> <p>Logs generated for traffic that match this rule display Close.</p> <p>NOTE: In sniffer mode, the IDP Series device is not in the path of network traffic. However, if you use ACM to configure a sniffer mode reset interface, the IDP Series device can send an RST packet to both the client and server but does not close the connection.</p>
Close Client	Closes the interactive connection to the client but not to the server.
Close Server	Closes the interactive connection to the server but not to the client.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Backdoor Rulebase on page 141

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Backdoor Rulebase Rules (NSM Procedure)

Understanding Backdoor Rulebase Notification Options

By default, logging is enabled for Backdoor rulebase rules. Table 34 on page 145 describes the notification options you can configure. You also have the option to disable logging.

Table 34: Backdoor 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.

Table 34: Backdoor Rulebase Notification Options (*continued*)

Option	Description
Packet captures	<p>Viewing the packets used in an attack on your network can help you determine the extent of the attempted attack and 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, the IDP Series device 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, the IDP Series device 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>



NOTE: Backdoor rulebase notification options are the same as IDP rulebase options.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Backdoor Rulebase on page 141
- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Backdoor Rulebase Rules (NSM Procedure)

Backdoor Rulebase Example: netcat

The **netcat** utility can open connections to any port and can offer services on any port. We know that attackers use **netcat** to create and exploit backdoors.

Suppose an attacker gains access to a host in your network and installs a **netcat** utility. The following example shows a **netcat** command an attacker can run:

```
nc 10.1.1.100 4444
```

This command opens a connection to the computer at IP address 10.1.1.100 over port 4444.

Figure 67 on page 147 shows a recommended rule that would detect the interactive traffic generated by **netcat** in this case.

Figure 67: Backdoor Rulebase

No.	Match					Operation	Action
	From Zone	Source	To Zone	Destination	Service		
1	any	Web Server Group	any	any	ECHO FTP rtalk	Ignore	Accept
2	any	Web Server Group	any	any	any	Detect	Accept

Rule 1 ignores the interactive traffic you expect for interactive services in your network (Telnet, SSH, RSH, NetMeeting, and VNC). Rule 2 detects all other interactive traffic. Rule 2 detects interactive traffic that occurs over a port where there typically is not interactive traffic.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Backdoor Rulebase on page 141

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Backdoor Rulebase Rules (NSM Procedure)

The SYN Protector Rulebase

This chapter explains how the SYN Protector rulebase protects your network and provides guidelines for configuring SYN Protector rulebase rules. It includes the following topics:

- Understanding the SYN Protector Rulebase on page 149
- Understanding SYN Protector Rulebase Match Settings on page 151
- Understanding SYN Protector Rulebase Modes on page 152
- Understanding SYN Protector Rulebase Notification Options on page 153

Understanding the SYN Protector Rulebase

The SYN Protector rulebase protects your network from malicious SYN flood attacks.

A *SYN flood attack* is a type of denial-of-service (DoS) attack, where the attacker attempts to flood your server with TCP requests to overwhelm your resources.

Attackers send a SYN message from a host with a spoofed, unreachable IP address to foil the TCP three-way handshake:

- A client host sends a SYN packet to a specific port on the server.
- Next, the server sends a SYN/ACK packet to the client host. The potential connection is now in a SYN_RECV state.
- Because the system is unreachable, the server never receives an ACK or RST packet back from the client host. The potential connection remains in the SYN_RECV state, and is placed into a connection queue while it waits for an ACK or RST packet. This “half-opened” connection remains in the queue until the connection-establishment timer expires (when it will be deleted).

To exploit this vulnerability, attackers use attack programs that generate thousands of bogus SYN messages, resulting in denial of service.

When the SYN Protector rulebase is enabled, the IDP engine detects traffic that exceeds the traffic thresholds you set as runtime parameters. Figure 68 on page 150 shows the SYN protector detection settings in the NSM Device Manager configuration editor.

Figure 68: NSM Device Manager: Sensor Settings > Run-Time Parameters

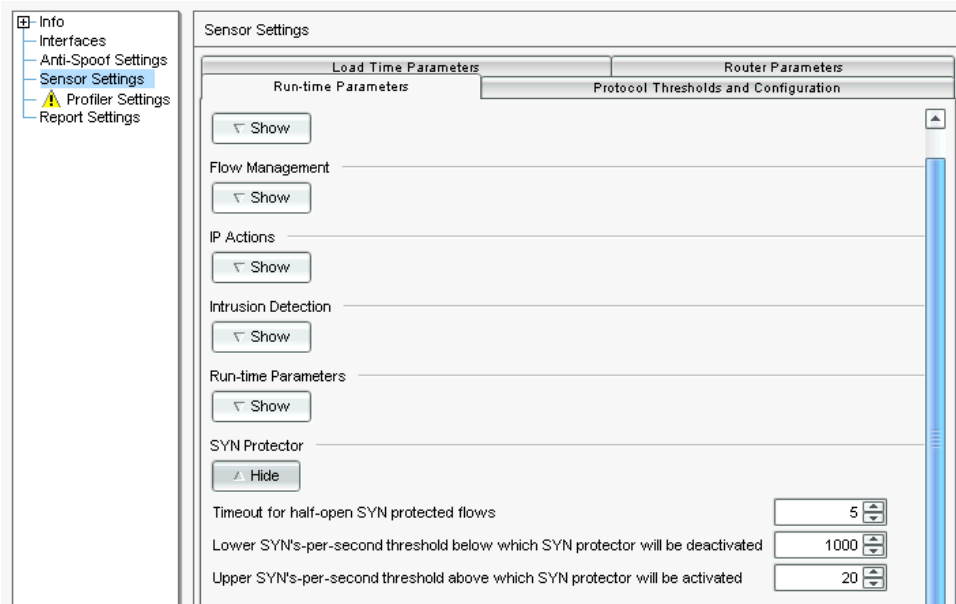


Table 35 on page 150 describes the SYN Protector thresholds.

Table 35: SYN Protector Thresholds

Setting	Description
Timeout for half-open SYN protected flows	<p>Used when SYN Protector is configured in passive mode.</p> <p>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.</p> <p>In passive mode, the IDP Series device monitors session startup. If the client does not send an ACK within the specified timeout, the IDP Series device sends a TCP reset. This setting controls the connection establishment timer, which determines the number of seconds that the IDP engine maintains a half-open SYN protected flow. The default is 5 seconds.</p>
Lower SYN's-per-second threshold below which SYN Protector will be deactivated	<p>Used to activate the SYN Protector in passive or relay mode.</p> <p>In passive mode, 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. The defaults are 1000 and 20. Using the defaults, the SYN Protector is activated when SYN's-per-second reach 1020. The SYN Protector rulebase is deactivated when the number of SYN packets per second falls below the lower SYN's-per-second threshold.</p>
Upper SYN's-per-second threshold above which SYN Protector will be activated	<p>In relay mode, the SYN Protector rulebase is activated when the number of SYN packets per second exceeds the lower SYN's-per-second threshold. The upper threshold is not used.</p>

When you create rules for the SYN Protector rulebase, you specify:

- A source/destination/service match condition

- A response mode: passive or relay
- Notification options

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding SYN Protector Rulebase Match Settings on page 151
- Understanding SYN Protector Rulebase Modes on page 152
- Understanding SYN Protector Rulebase Notification Options on page 153
- Understanding the Components of an IDP Security Policy on page 83

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring SYN Protector Rulebase Rules (NSM Procedure)

Understanding SYN Protector Rulebase Match Settings

The SYN Protector rulebase becomes active when IDP detects traffic that exceeds the thresholds you set as runtime parameters. Table 36 on page 151 shows the defaults for SYN Protector rulebase detection runtime parameters. You can tune these parameters if safe traffic in your network triggers false positives.

Table 36: SYN Flood Detection Runtime Parameters

Parameter	Default
Timeout for half-open SYN protected flows	5
Lower SYN-per-second threshold below which SYN Protector will be deactivated	1000
Upper SYN-per-second threshold above which SYN Protector will be activated.	20

In other words, using the defaults, the SYN Protector rulebase is activated when the IDP Series device counts 1020 SYN packets per second and deactivates when it falls below 1000 SYN packets per second.

When the SYN Protector rulebase is active, the IDP process engine evaluates its rules, beginning with source, destination, and service matching.

Because all TCP-IP is susceptible to a SYN flood attack, we recommend the following settings:

- Source—Any
- Destination—Servers you want to protect
- Service—TCP Any



TIP: You can use two rules to protect a large number of servers. Configure rule 1 to match servers you do not need to protect, and set Mode to **None**. Configure rule 2 to match any traffic and set Mode to **Passive** or **Relay**, as you prefer.



TIP: In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM documentation.



NOTE: The SYN Protector rulebase is a terminal rulebase—that is, SYN Protector rules are inherently terminal rules. If a SYN Protector rule matches, IDP does not process subsequent rules.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Rule-Matching Algorithm on page 85
- Understanding the SYN Protector Rulebase on page 149

The following related topics are included in the *IDP Series Administration Guide*:

- Configuring SYN Protector Rulebase Rules (NSM Procedure)
- Modifying the IDP Series Device Configuration

Understanding SYN Protector Rulebase Modes

Table 37 on page 152 summarizes SYN Protector rulebase modes.

Table 37: SYN Protector Rulebase Modes

Mode	Description
None	The IDP Series device takes no action and does not participate in the three-way handshake.
Passive	In passive mode, the IDP Series device monitors session startup. If the client does not send an ACK within a timeout period, the IDP Series device sends a TCP reset.

Table 37: SYN Protector Rulebase Modes (*continued*)

Mode	Description
Relay	<p>In relay mode, the IDP Series device acts as a relay for the connection establishment, performing the three-way handshake with the client on behalf of the server. When the IDP Series device receives the initial SYN packet, it returns a SYN/ACK packet with a SYN cookie. A SYN cookie is a 32-bit number that is put into the TCP sequence number field of a packet. If the client replies with an ACK packet with the appropriate cookie, the IDP Series device completes the three-way handshake and allows the session to become established. If the IDP Series device does not receive an appropriate ACK packet from the client, as is the case in a SYN flood attack, the IDP Series device does not establish the connection. Relay mode guarantees that the server allocates resources only to connections that are already in an established state. The relay is transparent to both the client and server.</p> <p>Relay mode has the following limitations:</p> <ul style="list-style-type: none"> • When the ACK packet from the client is lost, it can potentially lead to an unsynchronized state between client and server. • Because the IDP Series device does not save TCP options found in SYN packets, TCP extensions used for efficient transaction-oriented service (T/TCP) and Selective Acknowledgment (SACK), or protocols such as BGP, have a problem when SYN flooding is detected and the IDP Series device initiates the proxy TCP handshake. • Relay mode can be susceptible to ACK flooding because the IDP Series device must check for the validity of a cookie in the ACK messages. <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 Series device 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>



TIP: You can use two rules to protect a large number of servers. Configure rule 1 to match servers you do not need to protect, and set Mode to **None**. Configure rule 2 to match any traffic and set Mode to **Passive** or **Relay**, as you prefer.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the SYN Protector Rulebase on page 149

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring SYN Protector Rulebase Rules (NSM Procedure)

Understanding SYN Protector Rulebase Notification Options

By default, notification is not enabled for SYN Protector rulebase rules. You have the option to enable notification options. Table 38 on page 154 describes these options.

Table 38: SYN Protector 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	Packet capture is not available for SYN Protector rulebase rules.



NOTE: SYN Protector rulebase notification options are the same as IDP rulebase options, except that packet capture is not applicable.

**Related
Documentation**

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the SYN Protector Rulebase on page 149

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring SYN Protector Rulebase Rules (NSM Procedure)

The Traffic Anomalies Rulebase

This chapter explains how the Traffic Anomalies rulebase protects your network and provides guidelines for configuring Traffic Anomalies rulebase rules. It includes the following topics:

- Understanding the Traffic Anomalies Rulebase on page 155
- Understanding Traffic Anomalies Rulebase Match Conditions on page 157
- Understanding Traffic Anomalies Rulebase Detection Settings on page 158
- Understanding Traffic Anomalies Rulebase IP Actions on page 159
- Understanding Traffic Anomalies Rulebase Notification Options on page 160

Understanding the Traffic Anomalies Rulebase

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

A *traffic anomaly* is a pattern that indicates abnormal network activity. Traffic generated by automated port scanning tools trigger Traffic Anomalies rulebase rules. Attackers use automated port scanning tools to perform reconnaissance on your network. Typically, an automated port scanning tool attempts to connect to every port on a single machine (port scanning) or to connect to multiple IP addresses on a network (network scanning). Attackers do this to determine which services are allowed and responding on your network, so they can focus attacks on any vulnerabilities.

Traffic Anomalies rulebase rules count the number of ports scanned in a specified time period and use this traffic flow analysis to identify scans, as well as other attacks that occur over multiple connections and sessions. If the rule detects an attack, you can drop the connection and block the IP address that originated the attack. The IDP engine takes action against traffic that exceeds the thresholds you set.

Table 39 on page 156 summarizes Traffic Anomalies rulebase detection settings. You can tune these parameters if safe traffic in your network triggers false positives.

Table 39: Traffic Anomalies Rulebase Detection Settings

Group	Description
TCP scans, UDP Port Scans	<p>Set a port count (number of ports scanned) and the time threshold (the time period that ports are counted) in seconds.</p> <p>The default port count is 20. The default time threshold is 120 seconds. The rule is matched if the same source scans 20 TCP ports on your internal network within 120 seconds or if the same source scans 20 UDP ports on your internal network within 120 seconds.</p>
Distributed Port Scan	<p>A distributed port scan is an attack that uses multiple source IP addresses to scan ports.</p> <p>Set a port count (number of ports scanned) and the time threshold (the time period that ports are counted) in seconds.</p> <p>The default IP count is 50. The default time threshold is 120 seconds. The rule is matched if 50 IP addresses attempt to scan ports on your internal network within 120 seconds.</p>
ICMP Sweep	<p>An ICMP sweep is an attack where a single source IP pings multiple IP addresses.</p> <p>Set a port count (number of ports scanned) and the time threshold (the time period that ports are counted) in seconds.</p> <p>The default IP count is 50. The default time threshold is 120 seconds. The rule is matched if the same source IP attempts to ping 50 IP addresses within 120 seconds.</p>
Network Scan	<p>A network scan is an attack where a single source IP scans multiple IP addresses.</p> <p>Set a port count (number of ports scanned) and the time threshold (the time period that ports are counted) in seconds.</p> <p>The default IP count is 50. The default time threshold is 120 seconds. The rule is matched if the same source IP attempts to scan 50 IP addresses within 120 seconds.</p>
Session Limit	<p>Set a threshold number of sessions allowed from a single host within a second. The default is 100 sessions.</p> <p>For example, assume your internal network typically has low volume traffic. To detect a sudden increase in traffic from a specific host (which might indicate a worm), configure a rule that matches traffic over your internal network and configure a limit of 200. To block traffic that exceeds the session limit, set an IP action of IDP Block and select Source, Protocol from the Blocking Options menu.</p>

In addition, you can tune runtime parameters for Traffic Signatures. Table 40 on page 156 describes the runtime parameters associated with the Traffic Anomalies rulebase.

Table 40: Traffic Signature Runtime Settings

Setting	Description
Byte threshold for suspicious flows	Scans typically use small packets to access targets. You can exclude flows that contain large packets to reduce false positives. The default is to exclude flows where packet size exceeds 20 bytes.
Reporting frequency while scan in progress (seconds)	Specifies how frequently log messages are generated. Default is 30 seconds.

Table 40: Traffic Signature Runtime Settings (*continued*)

Setting	Description
The number of IP addresses we track for session rate	Specifies the maximum number of source IP addresses for which session rate is calculated. Default is 32,767.

When you create rules for the Traffic Anomalies rulebase, you specify:

- A source/destination/service match condition
- Detection settings
- Response options
- Notification options

For complete procedures on configuring Traffic Anomalies rulebase rules, see the *IDP Series Administration Guide*.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding Traffic Anomalies Rulebase Match Conditions on page 157
- Understanding Traffic Anomalies Rulebase Detection Settings on page 158
- Understanding Traffic Anomalies Rulebase IP Actions on page 159
- Understanding Traffic Anomalies Rulebase Notification Options on page 160
- Understanding the Components of an IDP Security Policy on page 83

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)

Understanding Traffic Anomalies Rulebase Match Conditions

We recommend the following settings for Traffic Anomalies rulebase match conditions:

- Source — Any
- Destination — IP addresses for servers you want to protect
- Service — Any (or specify specific services if you are creating an ignore list)



TIP: You can use two rules to protect a large number of servers. Configure rule 1 to match services you do not want to detect and set the detection option to **Ignore**. Configure rule 2 to match any traffic and set the detection operation to **Detect**.



TIP: In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM documentation.



NOTE: The Traffic Anomalies rulebase is a terminal rulebase—that is, Traffic Anomalies rules are inherently terminal rules. If a Traffic Anomalies rule matches, IDP does not process subsequent rules.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Rule-Matching Algorithm on page 85
- Understanding the Traffic Anomalies Rulebase on page 155

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)

Understanding Traffic Anomalies Rulebase Detection Settings

The Traffic Anomalies rulebase detection setting is a toggle detection off and on.

Specify **Ignore** to turn off traffic anomaly detection for traffic that matches the rule.

Specify **Detect** to turn on detection for traffic that matches the rule and to configure detection settings.



TIP: You can use two rules to protect a large number of servers. Configure rule 1 to match services you do not want to detect and set the detection option to **Ignore**. Configure rule 2 to match any traffic and set the detection option to **Detect**.



TIP: In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM documentation.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Traffic Anomalies Rulebase on page 155

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)

Understanding Traffic Anomalies Rulebase IP Actions

If traffic matches a traffic anomalies rule, the IDP Series device can take action against the current connection and against subsequent network traffic from 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.

Table 41 on page 159 describes Traffic Anomalies rulebase IP actions.

Table 41: Traffic Anomalies 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
IP Notify	IDP does not take any action against future traffic but logs the event or sends an alert.



NOTE: Traffic Anomalies rulebase IP actions are the same IP actions available for IDP rulebase rules.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Traffic Anomalies Rulebase on page 155

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)

Understanding Traffic Anomalies Rulebase Notification Options

By default, logging is enabled for Traffic Anomalies rulebase rules. Table 42 on page 160 describes notification options. You also have the option of disabling logging.

Table 42: Traffic Anomalies 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	Packet capture is not available for Traffic Anomalies rulebase rules.



NOTE: Traffic Anomalies rulebase notification options are the same as IDP rulebase options, except that packet capture is not applicable.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Traffic Anomalies Rulebase on page 155
- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Traffic Anomalies Rulebase Rules (NSM Procedure)

The Network Honeypot Rulebase

This chapter explains how to use the Network Honeypot rulebase to detect attacker reconnaissance activities. It includes the following topics:

- Understanding the Network Honeypot Rulebase on page 161
- Understanding Network Honeypot Rulebase Match Settings on page 162
- Understanding Network Honeypot Operation Setting on page 162
- Understanding Network Honeypot Rulebase IP Actions on page 163
- Understanding Network Honeypot Rulebase Notification Options on page 164

Understanding the Network Honeypot Rulebase

The Network Honeypot rulebase is a method to detect reconnaissance activities.

A *network honeypot* is an apparently vulnerable system that draws the attention and action of attackers. In an IDP network honeypot, the IDP Series device impersonates ports on protected servers.

When you create rules for the Network Honeypot rulebase, you specify:

- A destination/service match condition
- Operation mode
- Response options
- Notification options



NOTE: The IDP Series device 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 cannot be supported for MPLS traffic.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding Network Honeypot Rulebase Match Settings on page 162
- Understanding Network Honeypot Operation Setting on page 162
- Understanding Network Honeypot Rulebase IP Actions on page 163
- Understanding Network Honeypot Rulebase Notification Options on page 164
- Understanding the Components of an IDP Security Policy on page 83

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Network Honeypot Rulebase Rules (NSM Procedure)

Understanding Network Honeypot Rulebase Match Settings

Network Honeypot rulebase rules are triggered when a source IP address makes a connection request to the destination IP address and service specified in the rule.

We recommend you set source to **Any**; set destination and service to the server and service you want to appear to be available.



TIP: In NSM, you can create address objects and service objects to facilitate configuration. One benefit of using objects is that you can configure them once and then use them in multiple rules. For details, see the NSM documentation.



NOTE: The Network Honeypot rulebase is a terminal rulebase—that is, Network Honeypot rules are inherently terminal rules. If a Network Honeypot rule matches, IDP does not process subsequent rules.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Network Honeypot Rulebase on page 161

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Network Honeypot Rulebase Rules (NSM Procedure)

Understanding Network Honeypot Operation Setting

The Network Honeypot rulebase operation setting is used to toggle the network honeypot on and off.

Specify **Impersonate** to turn it on; specify **Ignore** to turn it off.

If the Network Honeypot rulebase is turned on, an attacker attempts to connect to an impersonated port, and the rule matches, the IDP Series device responds with a TCP SYN/ACK.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Network Honeypot Rulebase on page 161

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Network Honeypot Rulebase Rules (NSM Procedure)

Understanding Network Honeypot Rulebase IP Actions

If traffic matches a Network Honeypot rule, the IDP Series device can take action against the current connection and against subsequent network traffic from 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.

Table 43 on page 163 describes Network Honeypot rulebase IP actions.

Table 43: Network Honeypot 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 <p>NOTE: The IP Close action 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 an IP action, the IDP engine is programmed to take a server-to-client action before the traffic has reached the server. In these cases, the IDP engine does not have server-to-client MPLS label information. Therefore, the TCP reset 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>

Table 43: Network Honeypot 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.



NOTE: Network Honeypot rulebase IP actions are the same IP actions available for IDP rulebase rules.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Network Honeypot Rulebase on page 161

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Network Honeypot Rulebase Rules (NSM Procedure)

Understanding Network Honeypot Rulebase Notification Options

By default, logging is not enabled for Network Honeypot rulebase rules. You have the option to enable notification options. Table 44 on page 164 describes these options.

Table 44: Network Honeypot 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. • 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 and 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>



NOTE: Network Honeypot rulebase notification options are the same as IDP rulebase options.

**Related
Documentation**

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Understanding the Network Honeypot Rulebase on page 161
- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- Configuring Network Honeypot Rulebase Rules (NSM Procedure)

CHAPTER 17

Fine-Tuning a Security Policy

This chapter provides a suggested workflow for getting started and fine-tuning a security policy. It includes the following topic:

- Example: Fine-Tuning a Security Policy on page 167

Example: Fine-Tuning a Security Policy

This topic provides a suggested workflow for getting started and fine-tuning a security policy. It includes the following subtopics:

- Fine-Tuning Security Policies Process Overview on page 167
- Getting Started with the Recommended Security Policy on page 168
- Refining Rule Matching Properties on page 168
- Reducing False Positives on page 169
- Adding Rulebases on page 172

Fine-Tuning Security Policies Process Overview

You want to tune a security policy so that it is:

- Comprehensive—Detects all possible attacks on specific hosts in your network.
- Optimized—Each attack object specified in an IDP rulebase rule has a performance cost. In general, you want more rules with a few attack objects in each rather than fewer rules with many attack objects. In addition, we recommend that a single rule includes only the attack objects that are applicable to the rule destination server and only those of a severity that concerns you.
- Precise—Generates few false positives.
- Maintainable—As you refine your rules, you want to leverage as much of the predefined logic as possible. In your IDP rulebase rules, for example, you want to use the application identification feature, dynamic attack object groups, recommended attack objects, and recommended actions as much as possible, knowing the Juniper Networks Security Center team updates these as needed (even daily).

Fine-tuning is an iterative process. The process involves the following steps:

1. Getting Started with the Recommended Security Policy on page 168
2. Refining Rule Matching Properties on page 168
3. Reducing False Positives on page 169
4. Adding Rulebases on page 172

Getting Started with the Recommended Security Policy

When you add the IDP Series device to the NSM Device Manager, NSM automatically pushes the recommended policy to the IDP Series device. The recommended policy protects destination servers from the most frequent and severe attacks.

Table 45 on page 168 summarizes the properties of the Recommended security policy.

Table 45: Recommended Security Policy Definition

Property	Value
Rulebase	IDP rulebase
Rules	Nine rules, distinguished by attack object
Source	Any
Destination	Any
Service	Default, meaning the matching property is based on the service bindings of the attack object specified by the rule
Attack objects	Recommended IP, Recommended TCP, Recommended ICMP, Recommended HTTP, Recommended SMTP, Recommended DNS, Recommended FTP, Recommended POP3, Recommended IMAP, Recommended Trojan, Recommended Virus, Recommended Worm NOTE: All of the attack objects included in the predefined policies are client-to-server attacks.
Action	Recommended, meaning the action is specified by the attack object
Notification	Logging

Refining Rule Matching Properties

The source and destination matching parameters for template rules are set to **Any**. These broad settings provide comprehensive protection, but they entail a performance cost and might result in more logs than are necessary. We recommend you customize these settings.

Run the Profiler for several days to gather information about the hosts and applications running in your network. After several days, you should have the data you need to complete the following tasks:

- Create NSM address objects that identify groups of internal servers. When you configure rules to examine client to server traffic, you specify the internal servers as the rule's destination servers.
- Create an address object that defines your internal network. When you configure rules to examine traffic from your network to hosts on the world wide web, you can specify the internal network address object as the rule's source.
- Create NSM service objects to identify services running on internal servers. In most cases, you can specify **Default** for service so the rule uses the service relevant to the attack object. However, there might be cases where you want to specify a service object or service group.
- Identify predefined attack groups related to services (or create your own attack group, if necessary).
- Refine the IDP rulebase rule set so that each rule is focused on a single destination server (client to server traffic) or service (server to client traffic).

Reducing False Positives

A *false positive* is a log record that reflects an event on your network that you are not concerned about and no longer want to see in your logs. The IDP security policy found traffic that matched your rule, but over time you realize you do not need to track such events.

To determine whether a log is a false positive, you need to understand why the IDP Series device triggered the log and whether or not the traffic poses a real risk to the target server.

Suppose your security policy rule includes the following attack object: Predefined :: HTTP: Windows Media Services NSISlog.DLL Buffer Overflow. It generates a log when it identifies the attack pattern in traffic through the IDP Series device. Use the reference information in the details pane below the log table to learn more about the attack. You can click the hypertext linked name of the attack object in the summary tab to display reference information for the attack, as shown in Figure 69 on page 170.

Figure 69: Using NSM Log Viewer Attack Reference Information

The screenshot displays the NSM Log Viewer interface. The main window shows a table of logs with columns: Log ID, Time Received, Alert, User Flag, Comment, Src Addr, Dst Addr, Action, Protocol, Dst., Net Sr., Nat Ds., Details, Category, Subcategory, and Severity. A log entry with ID 20090806/416967 is highlighted, showing a 'Windows 2000 SP4 only?' alert. A pop-up window titled 'HTTP: Windows Media Services NSISlog.DLL Buffer Overflow' is open, displaying references, an extended description, and technical information about the vulnerability.

Log ID	Time Received	Alert	User Flag	Comment	Src Addr	Dst Addr	Action	Protocol	Dst.	Net Sr.	Nat Ds.	Details	Category	Subcategory	Severity	
20090806/416941	8/5/09 11:13:33 PM				1.1.0.115	1.2.0.58	Conn Dropped	TCP	80	4	0.0.0.0	0.0.0.0	interface=eth2	Predefined	HTTP: IIS cmd.exe Command Exec...	Major
20090806/416943	8/5/09 11:13:33 PM				1.1.0.192	1.2.0.102	Conn Dropped	TCP	80	3	0.0.0.0	0.0.0.0	interface=eth2	Predefined	HTTP: IIS5.0 WebDAV SEARCH Co...	Device
20090806/416945	8/5/09 11:13:33 PM				1.1.0.212	1.2.0.241	Conn Dropped	TCP								
20090806/416948	8/5/09 11:13:33 PM				1.1.0.248	1.2.0.132	Conn Dropped	TCP								
20090806/416949	8/5/09 11:13:36 PM				1.1.0.63	1.2.0.159	Conn Dropped	TCP								
20090806/416950	8/5/09 11:13:36 PM				1.1.0.88	1.2.0.56	Conn Dropped	TCP								
20090806/416951	8/5/09 11:13:36 PM				1.1.0.161	1.2.0.81	Conn Dropped	TCP								
20090806/416956	8/5/09 11:13:36 PM				1.1.0.231	1.2.0.243	Conn Dropped	TCP								
20090806/416957	8/5/09 11:13:36 PM				1.1.0.231	1.2.0.243	Conn Dropped	TCP								
20090806/416961	8/5/09 11:13:36 PM				1.1.0.241	1.2.0.121	Conn Dropped	TCP								
20090806/416966	8/5/09 11:13:42 PM				1.1.0.170	1.2.0.91	Conn Dropped	TCP								
20090806/416967	8/5/09 11:13:42 PM			Windows 2000 SP4 only?	1.1.0.88	1.2.0.56	Conn Dropped	TCP								
20090806/416971	8/5/09 11:13:45 PM				1.1.0.241	1.2.0.121	Conn Dropped	TCP								
20090806/416972	8/5/09 11:13:45 PM				1.1.0.23	1.2.0.139	Conn Dropped	TCP								
20090806/417097	8/5/09 11:14:59 PM				1.1.0.188	1.2.0.231	Conn Dropped	TCP								
20090806/417098	8/5/09 11:17:37 PM				1.1.0.20	1.2.0.143	Conn Dropped	TCP								
20090806/417099	8/5/09 11:17:40 PM				1.1.0.199	1.2.0.227	Conn Dropped	TCP								
20090806/417100	8/5/09 11:17:40 PM				1.1.0.114	1.2.0.190	Conn Dropped	TCP								
20090806/417101	8/5/09 11:17:40 PM				1.1.0.234	1.2.0.123	Conn Dropped	TCP								
20090806/417102	8/5/09 11:17:40 PM				1.1.0.189	1.2.0.95	Conn Dropped	TCP								
20090806/417103	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP								
20090806/417104	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP								
20090806/417105	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP								
20090806/417106	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP								
20090806/417107	8/5/09 11:17:43 PM				1.1.0.205	1.2.0.103	Conn Dropped	TCP								
20090806/417108	8/5/09 11:17:43 PM				1.1.0.109	1.2.0.55	Conn Dropped	TCP								

HTTP: Windows Media Services NSISlog.DLL Buffer Overflow

References

- <http://online.securityfocus.com/bid/8035/discussion/>
- <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-0349>
- <http://www.kb.cert.org/vuls/id/113716>
- <http://www.microsoft.com/technet/security/bulletin/MS03-022.msp>
- <http://secunia.com/advisories/9115>

Extended Description

Last Modified
2009-08-13

Impact
Windows Media Services may expose IIS to remote arbitrary code execution if media logging is enabled.

Description
Microsoft has reported a buffer overflow vulnerability in Windows Media Services. This is due to a problem with how the logging ISAPI extension handles incoming client requests. This could cause arbitrary code execution in IIS, which is exploitable through Media Services.

Technical Information
Microsoft Media Services provides functionality for providing streaming media content to clients from IIS. It ships with a number of Microsoft Windows 2000 server releases and is also available for download for Windows NT. Microsoft has reported a buffer overflow vulnerability in Windows Media Services. This is due to a problem with how the logging ISAPI extension (nsislog.dll) handles incoming client requests. The logging facility may attempt to write excessive data to an undersized buffer when handling a malformed HTTP client request. This could trigger a denial of service or remote arbitrary code execution in IIS, which is exploitable through Media Services. The issue would occur in servers that are configured to provide logging of media requests. It is possible to exploit this issue by sending an overly long HTTP POST request to the

In this example, we learn that the threat detected applies only to Microsoft Windows 2000 Server SP4. It is a false positive because all of the Windows servers in our network are Windows Server 2008. You can use the NSM Log Viewer flag and comment features to mark logs as false positives. In Figure 70 on page 171, we have marked the log ID 20090806/416967 as a false positive because the attack targets server versions not present in our network.

Figure 70: Using NSM Log Viewer Flag and Comment Features

per Networks - NSM - global : current

Search Help

Log Viewer [3-IDP001]

Log ID	Time Received	Alert	User Flag	Comment	Src Addr	Dst Addr	Action	Protocol	Dst...	Nat Sr...	Nat Ds...
20090806/416941	8/5/09 11:13:33 PM				1.1.0.115	1.2.0.58	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416943	8/5/09 11:13:33 PM				1.1.0.192	1.2.0.102	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416945	8/5/09 11:13:33 PM				1.1.0.212	1.2.0.241	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416948	8/5/09 11:13:33 PM				1.1.0.248	1.2.0.132	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416949	8/5/09 11:13:36 PM				1.1.0.63	1.2.0.159	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416950	8/5/09 11:13:36 PM				1.1.0.88	1.2.0.56	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416951	8/5/09 11:13:36 PM				1.1.0.161	1.2.0.81	Conn Dropped	TCP	80	3	0.0.0.0
20090806/416956	8/5/09 11:13:36 PM				1.1.0.231	1.2.0.243	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416957	8/5/09 11:13:36 PM				1.1.0.231	1.2.0.243	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416961	8/5/09 11:13:36 PM				1.1.0.241	1.2.0.121	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416966	8/5/09 11:13:42 PM				1.1.0.170	1.2.0.91	Conn Dropped	TCP	80	4	0.0.0.0
20090806/416967	8/5/09 11:13:42 PM			Windows 2000 SP4 only?	1.1.0.88	1.2.0.56	Cor			3	0.0.0.0
20090806/416971	8/5/09 11:13:45 PM				1.1.0.241	1.2.0.121	Cor			3	0.0.0.0
20090806/416972	8/5/09 11:13:45 PM				1.1.0.23	1.2.0.139	Cor			4	0.0.0.0
20090806/417097	8/5/09 11:14:59 PM				1.1.0.188	1.2.0.231	Cor				
20090806/417098	8/5/09 11:17:37 PM				1.1.0.20	1.2.0.143	Cor				
20090806/417099	8/5/09 11:17:40 PM				1.1.0.199	1.2.0.227	Cor				
20090806/417100	8/5/09 11:17:40 PM				1.1.0.114	1.2.0.190	Cor				
20090806/417101	8/5/09 11:17:40 PM				1.1.0.234	1.2.0.123	Cor				
20090806/417102	8/5/09 11:17:40 PM				1.1.0.189	1.2.0.95	Cor				
20090806/417103	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Cor				
20090806/417104	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Cor				
20090806/417105	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP			
20090806/417106	8/5/09 11:17:40 PM				1.1.0.48	1.2.0.155	Conn Dropped	TCP			
20090806/417107	8/5/09 11:17:43 PM				1.1.0.205	1.2.0.103	Conn Dropped	TCP			
20090806/417108	8/5/09 11:17:43 PM				1.1.0.109	1.2.0.55	Conn Dropped	TCP			

Filter Find Flag Exempt... Show Hide Log Unhide Log Goto Policy

High Medium Low Closed False Positive Assigned Investigate Follow-Up Pending Clear

Jul 30 Jul 31 Aug 1 Aug 2 Aug 3 Aug 4 Aug 5 Aug 6 Aug 7

Out In 8/5/09 11:13:45 PM

Summary All Fields Whois Lookup Quick Report

[Predefined :: HTTP: Windows Media Services NSISlog.DLL Buffer Overflow](#)

[References](#)

This signature detects attempts to exploit the buffer overflow vulnerability against Microsoft Windows Media Services, included with Microsoft Windows 2000 Server SP4. Attackers can send a maliciously crafted HTTP "POST" request to overflow the buffer and cause a denial of service or execute arbitrary code.

Matching Data Snippet

HEX

Filtered on: Category

There are a number of ways you can tune your security policy to reduce false positives. Table 46 on page 172 summarizes some basic tunings.

Table 46: Actions to Take To Reduce False Positives

Type of False Positive	Tuning Required
You trust the source.	Add an Exempt rulebase rule to “whitelist” the trusted source.
The attack applies to a hardware or software version that does not match your destination server.	<p>You have many options:</p> <ul style="list-style-type: none"> • Delete the attack from the rule. • Modify an attack group to exclude the object. • Add an Exempt rulebase rule to whitelist the non-offending attack object. • Modify rule action so traffic is stopped or permitted differently from before. • Modify the rule severity so that you can filter these events differently from before.
Your team has already patched the vulnerability detected by the attack.	Same as previous.
Upon examination, benign traffic crosses thresholds that trigger protocol anomaly events.	Use the NSM Device Manager to modify protocol anomaly thresholds.

Adding Rulebases

The IDP rulebase is the primary rulebase in an IDP security policy. When you have sufficiently tuned your IDP rulebase rules so that the security policy generates the level of logs you want, you can add additional rulebases to enable additional detection methods.

Take the same approach to tuning these additional rulebases. Instead of refining the group of attack objects that are relevant, you tune the IDP runtime parameters that set thresholds for detection mechanisms.

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Profiler Overview on page 31
- Understanding the Rule-Matching Algorithm on page 85
- Understanding IDP Rulebase Rule Match Settings on page 92
- Understanding the Components of an IDP Security Policy on page 83
- Understanding the Exempt Rulebase on page 113
- Using Attack Objects on page 97
- IDP Logs Overview on page 53

The following related topic is included in the *IDP Series Administration Guide*:

- Modifying the IDP Series Device Configuration

CHAPTER 18

Additional Security Features

This chapter describes additional security features you can use to protect your network. It includes the following topics:

- IP Spoof Attack Prevention Overview on page 173

IP Spoof Attack Prevention Overview

Every IP packet includes the destination address (where the packet is going) and the source address (where the packet came from). The routers that provide Internet communication between distant computers determine the best route for the IP packet using only the destination address and typically ignore the source address.

Attackers, who typically do not want you to know where an attack is coming from, can fake the source address of a malicious IP packet (by modifying the packet headers) so that the packet appears to come from a trusted system. The use of a fake IP address is called *IP spoofing*. You can configure the IDP system to detect these irregularities.

To detect attacks that attempt to spoof the addresses of hosts in your protected network, you can associate IDP traffic interfaces with the addresses of hosts in your protected network. IDP then detects an IP spoof attack if:

- An incoming packet uses an IP address that belongs to a network object on your internal network.
- An outgoing packet uses an IP address that does not belong to a network object on your internal network.

You can configure whether IDP drops or logs the session with a spoofed IP address.

Related Documentation

The following related topic is included in the *IDP Series Administration Guide*:

- Modifying the IDP Series Device Configuration

PART 4

Additional Deployment Topics

- [Inspection of Encapsulated and Encrypted Traffic on page 177](#)

CHAPTER 19

Inspection of Encapsulated and Encrypted Traffic

This chapter identifies support for inspection of encapsulated or encrypted traffic. It includes the following topics:

- Inspection of GRE Traffic Overview on page 177
- Inspection of GTP Traffic Overview on page 177
- Inspection of IPsec VPN Traffic Overview on page 178
- Inspection of MPLS Traffic Overview on page 178
- Inspection of SSL Traffic Overview on page 179
- Example: Implementing Inspection of Outbound SSL Traffic on page 183
- Example: Exempting Outbound SSL Traffic from Inspection on page 185

Inspection of GRE Traffic Overview

Generic Routing Encapsulation (GRE) is a tunneling protocol designed to encapsulate a wide variety of network layer packets inside IP tunneling packets. The original packet is the payload for the final packet. The protocol is used on the Internet to secure virtual private networks.

To inspect the payload of an encapsulated packet, the IDP process engine must decapsulate it. IDP Series devices support decapsulation for IP-in-GRE and PPP-in-GRE. You can configure decapsulation for one or two layers.

Related Documentation

The following related topic is included in the *IDP Series Administration Guide*:

- Enabling Inspection of GRE Traffic

Inspection of GTP Traffic Overview

GPRS Tunneling Protocol (or GTP) is an IP-based protocol used within Global System for Mobile communication (GSM) and Universal Mobile Telecommunications System (UMTS) networks.

To inspect the payload of an encapsulated traffic, the IDP process engine must decapsulate it. IDP Series devices support decapsulation for UDP GTPv0 and GTPv1. You can configure decapsulation for one or two layers.

- Related Documentation** The following related topic is included in the *IDP Series Administration Guide*:
- Enabling Inspection of GTP Traffic

Inspection of IPsec VPN Traffic Overview

Internet Protocol Security (IPsec) virtual private networks use the Encapsulated Security Payload (ESP) protocol and the NULL encryption algorithm to ensure the authenticity, integrity, and confidentiality of IP packets.

To inspect the payload of an encapsulated packet, the IDP process engine must decapsulate it. IDP Series devices support decapsulation for IPsec ESP NULL traffic. You can configure decapsulation for one or two layers.

- Related Documentation** The following related topic is included in the *IDP Series Administration Guide*:
- Enabling Inspection of IPsec VPN Traffic

Inspection of MPLS Traffic Overview

Multiprotocol Label Switching (MPLS) is an IP label switching technology that enables predetermined paths to specific destinations, called Label Switched Paths (LSPs), to be established through an inherently connectionless IP network. In MPLS networks, packets contain short labels that describe how to forward them through the network.

With MPLS decapsulation enabled, the IDP engine can inspect the IPv4 payload and pass through non-IPv4 payload. Note the following requirements and limitations:

- The IDP engine cannot decapsulate other encapsulated protocols within an MPLS frame. For example, the IDP engine cannot decapsulate the MPLS frame, find a GRE frame, decapsulate the GRE, and inspect the payload. Instead, the IDP engine passes through such traffic.
- If your traffic uses Ethernet frames larger than 1750 bytes, you must ensure the IDP default maximum frame size is sufficient (the default maximum frame size is 9014 bytes). In addition, we recommend you set the maximum transmission unit (MTU) on the switch or router connected to the IDP Series device to 1750 bytes or lower.

The IDP Series device does not participate in Label Distribution Protocol (LDP). When the IDP Series device receives the traffic, the IDP engine stores the MPLS label stacks of client-to-server or server-to-client directions. After processing the flow, the IDP Series device forwards the IP frames with the label stack it had stored when it created the flow, relying on the label switch router (LSR) to add the correct MPLS labels to the packet.

In some cases, the IDP engine is programmed to act in the server-to-client direction before it has seen and stored a server-to-client MPLS label. In effect, these connections are dropped. You might observe dropped MPLS traffic if the following rule elements apply:

- IDP rulebase – Action: Close Client (limitation applies to VLAN tagged traffic only)
- IDP rulebase – IP action: IP Close
- SYN Protector rulebase – Relay mode
- Network Honeypot rulebase – Impersonate mode

MPLS support is not enabled by default. You can use the CLI to enable MPLS support.

**Related
Documentation**

The following related topic is included in the *IDP Series Administration Guide*:

- Enabling Inspection of MPLS Traffic

Inspection of SSL Traffic Overview

Secure Sockets Layer (SSL) is a cryptographic protocol that adds security to TCP/IP communication. Several versions of the SSL and Transport Layer Security (TLS) protocols are in widespread use in applications like Web browsing, electronic mail, Internet faxing, instant messaging, and voice over IP (VoIP). SSL and TLS encrypt the Transport Layer protocol datagrams that carry the payload of these communications. While encryption is an excellent way to keep private data from prying eyes, without inspection by the IDP Series device, it also unwittingly opens a network to dangerous viruses, trojans, or network attacks.

To inspect the HTTP payload of HTTPS traffic, the IDP Series device must decrypt the HTTPS session. Your security policy can examine both the SSL session and the decrypted HTTP payload.


The following sections describe alternative methods you can use to enable SSL inspection:

- Using the SSL Server Private Keys on page 179
- Using a Root Certificate Authority in SSL Forward Proxy Operations on page 180
- Supported SSL Specifications on page 181

Using the SSL Server Private Keys

Beginning with IDP OS Release 3.2r1, we support inspection of client-to-server traffic to internal SSL servers. As shown in Figure 71 on page 180, this method depends on administrative access to the SSL server private keys.

Figure 71: SSL Inspection Using SSL Server Private Keys

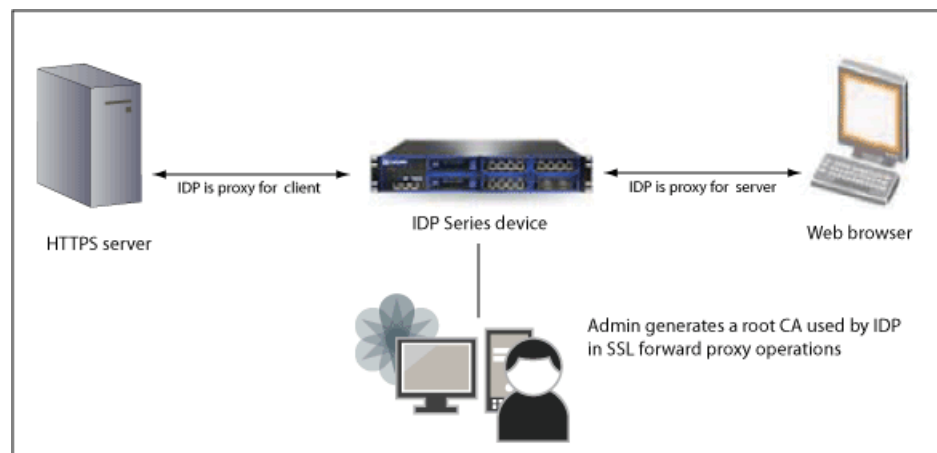


You must be able to copy the SSL server private key to the IDP SSL keystore. The IDP Series device uses the key to decrypt the inbound traffic so that it can inspect the payload. The private key must be in Privacy-Enhanced Mail (PEM) format. We have verified support for the following RSA private key lengths: 1024 bits, 2048 bits, 3072 bits, and 4096 bits.

Using a Root Certificate Authority in SSL Forward Proxy Operations

Beginning with IDP OS Release 5.0r2, we support inspection of traffic to HTTPS servers where you do not have access to the SSL private key, such as outbound traffic to the WWW. As shown in Figure 72 on page 180, this method uses a root certificate authority (CA) to proxy the SSL key negotiation. The IDP Series device inserts itself into the SSL key negotiation phase so that it can decrypt the HTTPS session and inspect the session and payload according to your security policy.

Figure 72: SSL Inspection Using a Root CA



When the special root CA is present, the IDP Series device intercepts the HTTPS connection and makes a request to the server as if it were the client; it presents to the client a CA (derived from the special root CA) as if it were the server. The IDP Series device then negotiates the key exchange, decrypts the session, inspects the payload, and re-encrypts the session as necessary before forwarding.

To ensure employee or customer privacy, you can configure a whitelist to exclude matching sessions from being processed by the SSL forward proxy feature. Traffic to destination servers on your whitelist is not intercepted and is passed through uninspected.



NOTE: When both the CA and server private keys are present, the IDP Series device uses the SSL forward proxy method to inspect HTTPS traffic.

Supported SSL Specifications

The IDP Series device supports decryption of HTTPS traffic that uses SSLv3 and TLSv1. The IDP Series device can inspect an SSLv2 header for anomalies, but it cannot decrypt and examine the HTTP payload in such sessions. In addition, the IDP Series device does not support inspection of compressed TLS traffic.

Table 47 on page 181 lists the SSL cipher suites supported by the two IDP SSL inspection methods.

Table 47: Supported SSL Cipher Suites

Cipher Suite	Decryption Using Private Keys	Decryption Using Forward Proxy
Name: TLS_RSA_WITH_NULL_MD5 Authorization: RSA Key Exchange: RSA Encryption: NULL Digest: MD5	Yes	Yes
Name: TLS_RSA_WITH_NULL_SHA Authorization: RSA Key Exchange: RSA Encryption: NULL Digest: SHA	Yes	Yes
Name: TLS_RSA_WITH_RC4_128_MD5 Authorization: RSA Key Exchange: RSA Encryption: RC4_128 Digest: MD5	Yes	Yes

Table 47: Supported SSL Cipher Suites (*continued*)

Cipher Suite	Decryption Using Private Keys	Decryption Using Forward Proxy
Name: TLS_RSA_WITH_RC4_128_SHA Authorization: RSA Key Exchange: RSA Encryption: RC4_128 Digest: SHA	Yes	Yes
Name: TLS_RSA_WITH_DES_CBC_SHA Authorization: RSA Key Exchange: RSA Encryption: DES_CBC Digest: SHA	Yes	Yes
Name: TLS_RSA_WITH_3DES_EDE_CBC_SHA Authorization: RSA Key Exchange: RSA Encryption: 3DES_EDE_CBC Digest: SHA	Yes	Yes
Name: TLS_RSA_WITH_AES_128_CBC_SHA Authorization: RSA Key Exchange: RSA Encryption: AES_128_CBC Digest: SHA	Yes	Yes
Name: TLS_RSA_WITH_AES_256_CBC_SHA Authorization: RSA Key Exchange: RSA Encryption: AES_256_CBC Digest: SHA	Yes	Yes

Related Documentation

The following related topics are included in the *IDP Series Concepts and Examples Guide*:

- Example: Implementing Inspection of Outbound SSL Traffic on page 183
- Example: Exempting Outbound SSL Traffic from Inspection on page 185

The following related topics are included in the *IDP Series Administration Guide*:

- Using the SSL Private Server Key to Enable Inspection of SSL Traffic
- Using the SSL Forward Proxy Feature to Enable Inspection of HTTPS Traffic
- Exempting HTTPS Traffic from Inspection

Example: Implementing Inspection of Outbound SSL Traffic

When users in your protected network connect to HTTPS servers on the WWW, the application activity within the encrypted sessions cannot be inspected by most intrusion prevention systems. Once an encrypted session is established with the server, the user might download a seemingly harmless file or executable that contains a trojan. If this happens, an attacker could launch an attack from within the protected network.

To protect your network against this risk, you could create a firewall policy that blocks HTTPS connections from the protected zone to the unprotected zone. To protect your network *and* support legitimate access to the WWW, you want a solution that can inspect the HTTPS traffic.

The IDP solution supports SSL inspection in two ways:

- Using server private keys. Use this method when inspecting traffic to internal servers where you have access to the server private key.
- Using the SSL forward proxy feature. Use this method when the server private key method is not practical (for example, for traffic to servers on the WWW).



NOTE: If you enable both methods, the IDP Series device performs SSL inspection using the SSL forward proxy method and does not use the server private keys.

The following procedure provides the basic steps you take to implement the SSL forward proxy feature.

To implement the SSL forward proxy feature:

1. From the IDP Series device command-line interface:

- a. Generate the root certificate authority (CA) that the IDP Series device uses to create and sign new certificates used in SSL proxy operations. The following example creates a root CA:

```
[root@default host admin]# scio ssl ca create US CA Sunnyvale 'Juniper Networks Inc.' 'SSL Inspection policy' 'Juniper IT Services' 'admin@juniper.net' 1024
```



NOTE: The system prompts the end user to install the CA you create in this step. Take care to configure an organization name that an end user is likely to accept, such as your company name.

- b. Verify the CA was added:

```
[root@default host admin]# scio ssl ca show
```

```
serial=8E0012848A2D7CCD
subject= /C=US/ST=CA/L=Sunnyvale/O=Juniper Networks Inc./OU=SSL Inspection
policy/CN=Juniper IT Services/emailAddress=admin@juniper.net
issuer= /C=US/ST=CA/L=Sunnyvale/O=Juniper Networks Inc./OU=SSL Inspection
policy/CN=Juniper IT Services/emailAddress=admin@juniper.net
notBefore=Jun 25 22:13:23 2009 GMT
notAfter=Jun 23 22:13:23 2019 GMT
```

- c. (Optional) Distribute the CA to your users and have them install the CA in their Web browser.

The following example prints to the screen the CA in PEM format. You can copy this text to a file that your users can import into their browsers.

```
[root@default host admin]# scio ssl ca export
```

```
-----BEGIN CERTIFICATE-----
MIIC1TCCAj4CCQCOABKEi18zTANBgkqhkiG9w0BAQUFADCBrijELMAkGA1UEBhMC
VVMxCzAJBgNVBAGTAKNBMRIwEAYDVQQHEw1TdW5ueXZhbGUxHjAcBgNVBAoTFUp1
bm1wZXIgaTmV0d29ya3MgSW5jLjEeMBwGA1UECxMVU1NMIEluc3B1Y3Rpb24gcG9s
aWN5MRwwGgYDVQQDEExNKDw5pcGVyIE1UIFN1cnZpY2VzMSAwHgYJKoZIhvcNAQkB
FhFhZG1pbkBgqdw5pcGVyLm5ldAeFw0wOTA2MjUyMjEzZm1wZG1wZG1wZG1wZG1w
MjNaMIGuMQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0ExEjAQBgNVBAcTCVN1bm55
dmFsZTEeMBwGA1UEChMVSnVuaXB1ciB0ZXR3b3JrcyBjb250bW50bW50bW50bW50
U0w5SW5zcGVjdG1vbiBwb2xpY3kxHDAaBgNVBAMTE0p1bm1wZXIgaTmV0d29ya3Mg
ZXIxIDAeBgkqhkiG9w0BCQEWFWFkbW1wZG1wZG1wZG1wZG1wZG1wZG1wZG1wZG1w
DQEBAAQAA4GNADCBiQKBgQDAsn2NFaXTrCpShf9sg+Ccn1rUYzPuVHTw1GUtnHHB
o/oFXeNGETggLZ/jck+L2710x3IpGd67yyHs08sXWvgC3MJukb14kqyTyguy3/E9
wkiIey8W4XzyBXRcFw2YegMc0cFExdm+C6DrAai1ddTQdgelxZ7nfIj24iiBhYYM
GQIDAQAABMA0GCSqSIB3DQEBBQUAA4GBAFTREz9DHcbohDJFqGWPjS+MDgsX9041
f/WzHXftak4ZHjOryYvVaRUyitEhMX1KvMPPQjYXf+TE2vF9yYqmoCj6710Liu2ZJ
Tw4gwy9E9p58krqvZu4F2/kVM+yEAksUIjBme1RIL6Az3kLauHvkyAbMcSFZG2b0
7Z8WbQqn3o6s
-----END CERTIFICATE-----
```

2. In NSM:

- a. Add IDP rulebase rules that target HTTPS traffic:
 - Match traffic flowing in the client-to-server direction.
 - Be sure to account for HTTPS over nonstandard ports. The application identification feature is not applied when the IDP Series device proxies an SSL session.
 - Include attack objects to inspect *both* the SSL session and the HTTP payload.



TIP: You must include at least one SSL attack object. We recommend you include the SSL: SERVER-CERT-FAILS-VALIDATION anomaly to detect invalid certificates (that is, certificates signed by an unknown CA or that cannot be validated against the issuer CA). An invalid certificate might indicate a phishing attack. You can drop or log matching sessions.

- Specify action and notification options.
- b. Push the updated security policy to the IDP Series device.
 - c. Review logs to verify the feature operates as you expect.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Inspection of SSL Traffic Overview on page 179

The following related topics are included in the *IDP Series Administration Guide*:

- Using the SSL Private Server Key to Enable Inspection of SSL Traffic
- scio ssl
- scio const

Example: Exempting Outbound SSL Traffic from Inspection

The privacy policy for your business might include cases where sessions should remain encrypted throughout. For example, suppose you have an agreement with your users that your network security infrastructure will not interfere with SSL encrypted connections to banking sites. In these cases, you can create a whitelist of destination domain names, IP addresses, and subnets you want exempted from IDP policy inspection. If a server is included in the whitelist, the IDP system does not decrypt the traffic or inspect it. Instead, this traffic is passed through the IDP Series device uninspected.



NOTE: The whitelist applies only to traffic processing based on the SSL forward proxy feature. You would not use a whitelist to exclude inspection of traffic to internal destination servers. If desired, you can use a security policy rule to exempt such traffic from inspection.

The following example shows the format of a whitelist file:

```
10.0.0.1
1.0.0.0/8
70.34.21.82
trustedsite.com
landing.trustedsearch.com
```

Each line in the whitelist file specifies the IP address or domain name for a destination server. To whitelist multiple sites with one entry, you can use an IP prefix to match address blocks and a domain suffix to include all subdomains.

The domain name in your whitelist should match the common name (CN) entry in the certificate presented by the destination server. For example, suppose the certificate for the E-Trade HTTPS server contains the following subject:

```
C=US, ST=Georgia, L=Alpharetta, O=ETRADE FINANCIAL CORPORATION, OU=Global
Information Security, CN=us.etrade.com
```

You can whitelist this site by adding the string **us.etrade.com** or the string **etrade.com** to your whitelist file.

In most cases, the CN entry in the server certificate for a website matches the server name that appears in the browser address bar. In some cases, there are differences. You can use the features of your Web browser to find the CN entry in the server certificate for the website.

Figure 73 on page 187 shows the location of the certificate details in Firefox.

Figure 73: Firefox: Displaying the Server Certificate for a Website

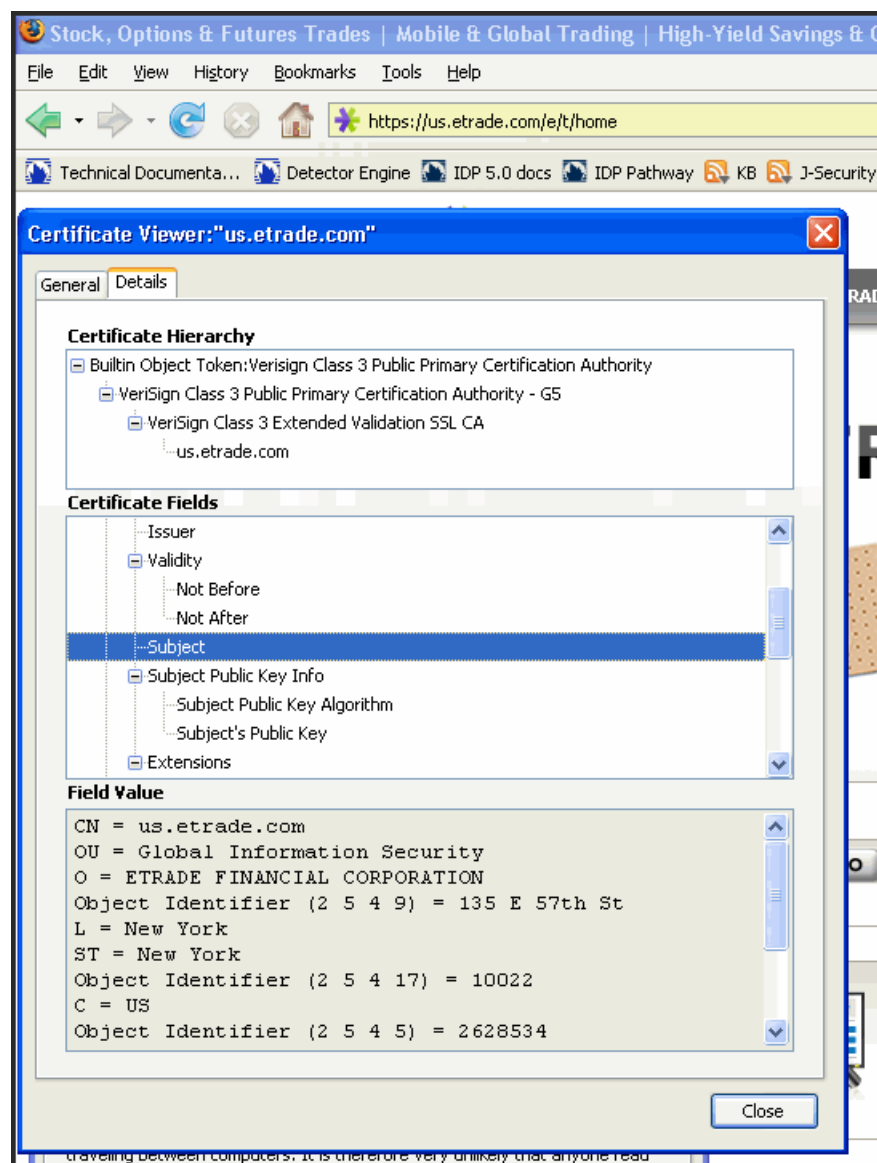
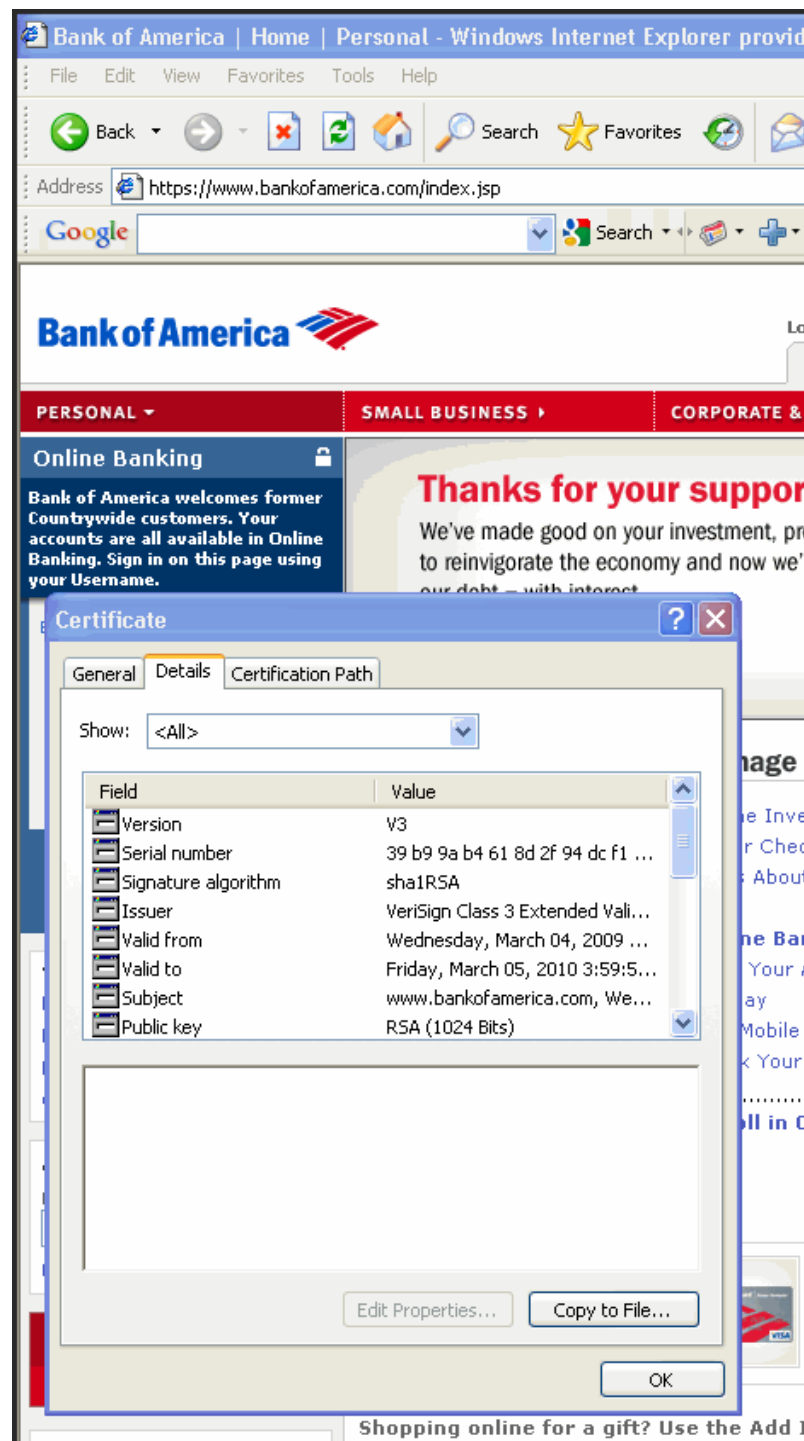


Figure 74 on page 188 shows the location of the certificate details in Internet Explorer.

Figure 74: Internet Explorer: Displaying the Server Certificate for a Website



To implement a whitelist:

1. Log into the CLI as **admin** and enter **su -** to switch to root.

2. Use an editor like vi to create a whitelist file. A whitelist file should contain the IP address prefixes and/or domain name suffixes you want to exempt from inspection. For example:

```
[root@default host admin]# vi /tmp/whitelist.txt
```

```
e-trade.com
bankofamerica.com
```

3. Run the following command to import the whitelist entries:

```
[root@default host admin]# scio ssl whitelist import /tmp/whitelist.txt
```



NOTE: To update the active whitelist, import an updated whitelist file. To clear the whitelist, import a file that contains only one empty line.

Related Documentation

The following related topic is included in the *IDP Series Concepts and Examples Guide*:

- Inspection of SSL Traffic Overview on page 179

The following related topic is included in the *IDP Series Administration Guide*:

- scio ssl

PART 5

Index

- Index on page 193

Index

A

actions.....	101
agent process.....	12
APE rulebase.....	117
application identification.....	96
application policy enforcement.....	117
application volume tracking.....	45
attack objects.....	97
auto-recovery.....	11

B

backdoor rulebase.....	141
------------------------	-----

C

control plane.....	9
customer support.....	xx
contacting JTAC.....	xx

D

data plane.....	9
DDoS attack prevention.....	155
documentation.....	xix
DoS attack prevention.....	149

E

eth0 interface.....	15
eth1 interface.....	15
eth2 interface.....	15
Exempt rulebase.....	113
extended application identification See application identification	
external bypass.....	17

F

fail close.....	16
false positives.....	169
feature list.....	3
flow bypass.....	12

G

GRE traffic.....	177
GTP traffic.....	177

H

HA interface.....	15
-------------------	----

I

IDP Reporter.....	76
IDP rulebase.....	91
idpengine process.....	12
idpHMD process.....	12
idpLogReader process.....	12
internal bypass.....	16
IP spoof attack prevention.....	173
IPsec ESP NULL traffic.....	178

J

J-Security Center.....	21
JNET driver.....	9

L

logging.....	103
logs.....	53

M

management interface.....	15
MPLS traffic.....	178
multicore architecture.....	9

N

nested application identification See application identification	
Network and Security Manager See NSM overview	
Network Honeypot rulebase.....	161
NSM overview.....	20
NSM reports.....	74

P

packet logging.....	68
---------------------	----

peer port modulation.....	18
pkid process.....	12
policies.....	83
PPM.....	18
Profiler.....	31
Application Profiler tab.....	48
process.....	12
reports.....	75

R

Recommended action.....	101, 110
Recommended attack.....	109
Recommended policy.....	87, 168

S

sciod process.....	13
security policies.....	83
support, technical See technical support	
SYN Protector rulebase.....	149

T

technical support	
contacting JTAC.....	xx
Traffic Anomalies rulebase.....	155
traffic interfaces.....	15
transparent mode.....	25

U

user-role-based policy.....	94
-----------------------------	----

V

virtual router.....	15
---------------------	----