



NETWORK INTRUSION PREVENTION SYSTEMS INDIVIDUAL PRODUCT TEST RESULTS

Palo Alto Networks PA-4020



METHODOLOGY VERSION: 6.0
AUGUST 2010

Licensed to: Palo Alto Networks
To receive a licensed copy or report misuse,
please contact NSS Labs at +1 (512) 961-5300 or advisor@nsslabs.com.

© 2010 NSS Labs, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the authors.

Please note that access to or use of this report is conditioned on the following:

1. The information in this report is subject to change by NSS Labs without notice.
2. The information in this report is believed by NSS Labs to be accurate and reliable at the time of publication, but is not guaranteed. All use of and reliance on this report are at the reader's sole risk. NSS Labs is not liable or responsible for any damages, losses, or expenses arising from any error or omission in this report.
3. NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY THE NSS LABS. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT ARE DISCLAIMED AND EXCLUDED BY NSS LABS. IN NO EVENT SHALL NSS LABS BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
4. This report does not constitute an endorsement, recommendation, or guarantee of any of the products (hardware or software) tested or the hardware and software used in testing the products. The testing does not guarantee that there are no errors or defects in the products or that the products will meet the reader's expectations, requirements, needs, or specifications, or that they will operate without interruption.
5. This report does not imply any endorsement, sponsorship, affiliation, or verification by or with any organizations mentioned in this report.
6. All trademarks, service marks, and trade names used in this report are the trademarks, service marks, and trade names of their respective owners.

CONTACT INFORMATION

NSS Labs, Inc.

P.O. Box 130573

Carlsbad, CA 92013 USA

+1 (512) 961-5300

info@nsslabs.com

www.nsslabs.com

TABLE OF CONTENTS

1	<i>Introduction</i>	1
2	<i>Security Effectiveness</i>	2
2.1	Coverage by Attack Vector	2
2.2	Coverage by Impact Type	3
2.3	Attack Leakage	3
2.4	Resistance to Evasion	4
3	<i>Performance</i>	5
3.1	Real-World Traffic Mixes	5
3.2	Connection Dynamics – Concurrency and Connection Rates	6
3.3	HTTP Connections per Second and Capacity	7
3.4	HTTP Connections per Second and Capacity with Delays	8
3.5	UDP Throughput	9
4	<i>Total Cost of Ownership</i>	10
4.1	Labor per Product (in Hours)	10
4.2	Purchase Price and Total Cost of Ownership	10
4.3	Value: Cost per Mbps and Exploit Blocked – Tuned Policy	11
5	<i>Detailed Product Scorecard</i>	12
6	<i>Appendix B: Special Thanks</i>	18

TABLE OF FIGURES

Figure 1: Coverage by Attack Vector – Default vs. Tuned Configurations	2
Figure 2: Product Coverage by Impact – Default vs. Tuned Configurations	3
Figure 3: Real-World Traffic Mixes – Default vs. Tuned Configurations	5
Figure 4: Concurrency and Connection Rates – Default vs. Tuned Configurations	6
Figure 5: HTTP Connections per Second and Capacity – Default vs. Tuned Configurations....	7
Figure 6: HTTP Connections per Second and Capacity (with/without Delays)	8
Figure 7: UDP Throughput – Default vs. Tuned Configurations	9

1 INTRODUCTION

During Q3 2010, NSS Labs performed an independent group test of network intrusion prevention systems (IPS) currently on the market. Each product was subjected to thorough testing at the NSS Labs facility in Austin, Texas, based on methodology v6.0 available on www.nsslabs.com. IPS vendors were invited to submit their products to NSS Labs free of charge and we did not receive any compensation in return for their participation.

While the Network IPS Group Test Report provides comparative information about those products, this Individual Test Report provides further detailed information not available elsewhere.

NSS Labs evaluated the products configured with the default, “out-of-the-box” settings, then again as optimally tuned by the vendor prior to testing to provide readers with a range of information on key IPS security effectiveness and performance dimensions.

As part of this test, Palo Alto Networks submitted the **PA-4020**

NSS Labs’ Rating: **Recommend**

Product	Effectiveness	Throughput
Palo Alto Networks PA-4020	93.4%	2,259 Mbps

Security effectiveness was excellent. Using the default policy, the PA-4020 blocked 56.6% of attacks. After rapid tuning that consisted of changing just three settings within the policy by a Palo Alto Networks engineer, the effectiveness improved by 36.8% to 93.4%. In addition, the Palo Alto Networks PA-4020 correctly identified 100% of our evasion attempts without error.

The product successfully passed 2.2 Gbps of inspected traffic. NSS Labs rates throughput based upon tuned settings—averaging out the results from tests 6.6.1, 6.6.2, and 6.4.2: “Real World” Protocol Mix (Perimeter), “Real World” Protocol Mix (Core), and 21 KB HTTP Response respectively.

Palo Alto Networks’ Management interface was flexible and excellent for reporting. Tuning and maintenance is simple and well-thought out.

For medium throughput multi-gigabit environments, the Palo Alto Networks PA-4020 provides an outstanding 3-year TCO (including labor).

2 SECURITY EFFECTIVENESS

To show the range of expectations a user should have, NSS Labs evaluated the products configured with the default, “out-of-the-box” settings, then again as optimally tuned by the vendor prior to testing.

Live Exploit Testing: NSS Labs’ security effectiveness testing leverages deep expertise of our engineers utilizing multiple commercial, open source and proprietary tools as appropriate. With 1,179 live exploits, this is the industry’s most comprehensive test to date. We retired 92 attacks and added 112 new exploits compared to our Q4 2009 test set of 1,159. Most notable, all of the live exploits and payloads in our test have been validated in our lab such that:

- a reverse shell is returned
- a bind shell is opened on the target allowing the attacker to execute arbitrary commands
- a malicious payload installed
- a system is rendered unresponsive
- etc.

Configuration	Total Number of Exploits Run	Total Number Blocked	Block Percentage
Default Configuration	1,179	666	56.4%
Tuned Configuration	1,179	1101	93.4%

2.1 COVERAGE BY ATTACK VECTOR

Because a failure to block attacks could result in significant compromise and impact to critical business systems, network IPS should be evaluated against a broad set of exploits. Exploits can be categorized into two groups: attacker-initiated and target initiated. Attacker-initiated exploits are threats executed remotely against a vulnerable application and/or operating system by an individual while target-initiated exploits are initiated by the vulnerable target. In target-initiated exploits, the attacker has little or no control as to when the threat is executed.

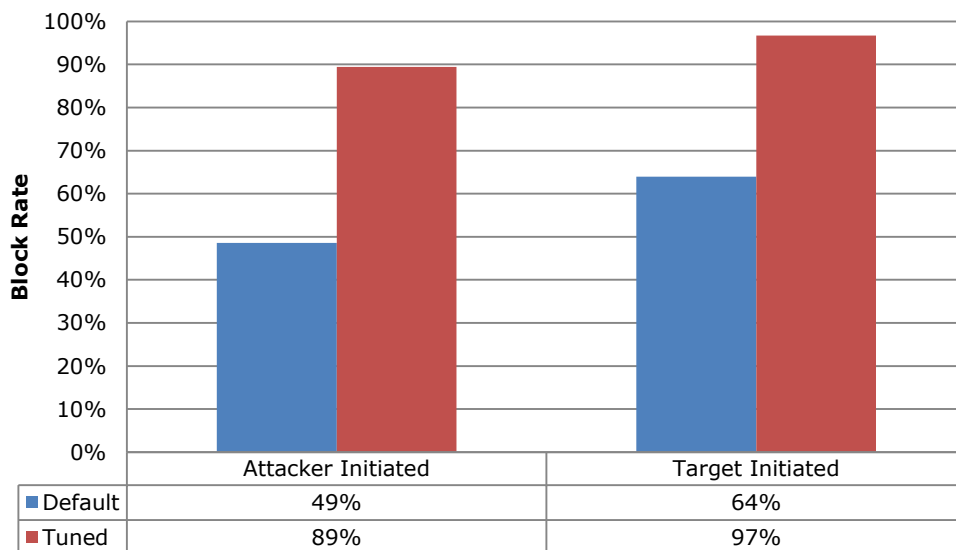


Figure 1: Coverage by Attack Vector – Default vs. Tuned Configurations

2.2 COVERAGE BY IMPACT TYPE

The most serious exploits are those which result in a remote system compromise, providing the attacker with the ability to execute arbitrary system-level commands. Most exploits in this class are “weaponized” and offer the attacker a fully interactive remote shell on the target client or server.

Slightly less serious are attacks that result in an individual service compromise, but not arbitrary system-level command execution. Typical attacks in this category include service-specific attacks—such as SQL injection—that enable an attacker to execute arbitrary SQL commands within the database service. These attacks are somewhat isolated to the service and do not immediately result in full system-level access to the operating system and all services. However, using additional localized system attacks, it may be possible for the attacker to escalate from the service level to the system level.

Finally, there are the attacks (often target initiated) which result in a system or service-level fault that crashes the targeted service or application and requires administrative action to restart the service or reboot the system. These attacks do not enable the attacker to execute arbitrary commands. Still, the resulting impact to the business could be severe, as the attacker could crash a protected system or service.

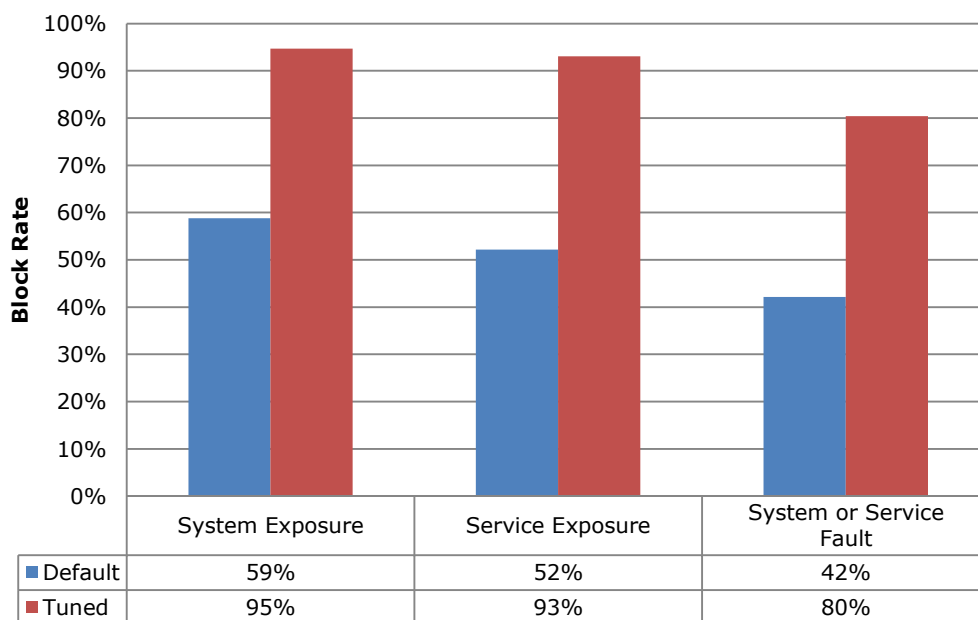


Figure 2: Product Coverage by Impact – Default vs. Tuned Configurations

2.3 ATTACK LEAKAGE

All NIPS devices have to make the choice whether to risk denying legitimate traffic or allowing malicious traffic once they run low on resources. By default, the Palo Alto Networks PA-4020 will drop new connections when resources (such as state table memory) are low, or when traffic loads exceed the device capacity. This will theoretically block legitimate traffic, but maintain state on existing connections (preventing evasion).

2.4 RESISTANCE TO EVASION

Description	IP Packet Fragmentation	TCP Stream Segmentation	RPC Fragmentation	URL Obfuscation	HTML Evasion	FTP Evasion	TOTAL
Palo Alto Networks PA-4020	✓	✓	✓	✓	✓	✓	✓

Resistance to known evasion techniques was perfect, with the Palo Alto Networks PA-4020 achieving a 100% score across the board in all related tests. *IP fragmentation*, *TCP stream segmentation*, *RPC fragmentation*, *URL obfuscation*, *HTML Evasion* and *FTP evasion* all failed to trick the product into ignoring valid attacks. Not only were the fragmented and obfuscated attacks blocked successfully, but all of them were also decoded accurately.

3 PERFORMANCE

There is frequently a trade-off between security effectiveness and performance. Because of this trade-off, it is important to judge a product’s security effectiveness within the context of its performance (and *vice versa*). This ensures that new security protections do not adversely impact performance and security shortcuts are not taken to maintain or improve performance.

3.1 REAL-WORLD TRAFFIC MIXES

The aim of this test is to measure the performance of the device under test in a “real world” environment by introducing additional protocols and real content, while still maintaining a precisely repeatable and consistent background traffic load. Different protocol mixes are utilized based on the location of the device under test to reflect real use cases. For details about real world traffic protocol types and percentages, see the NSS Labs IPS Test Methodology, available at www.nsslabs.com.

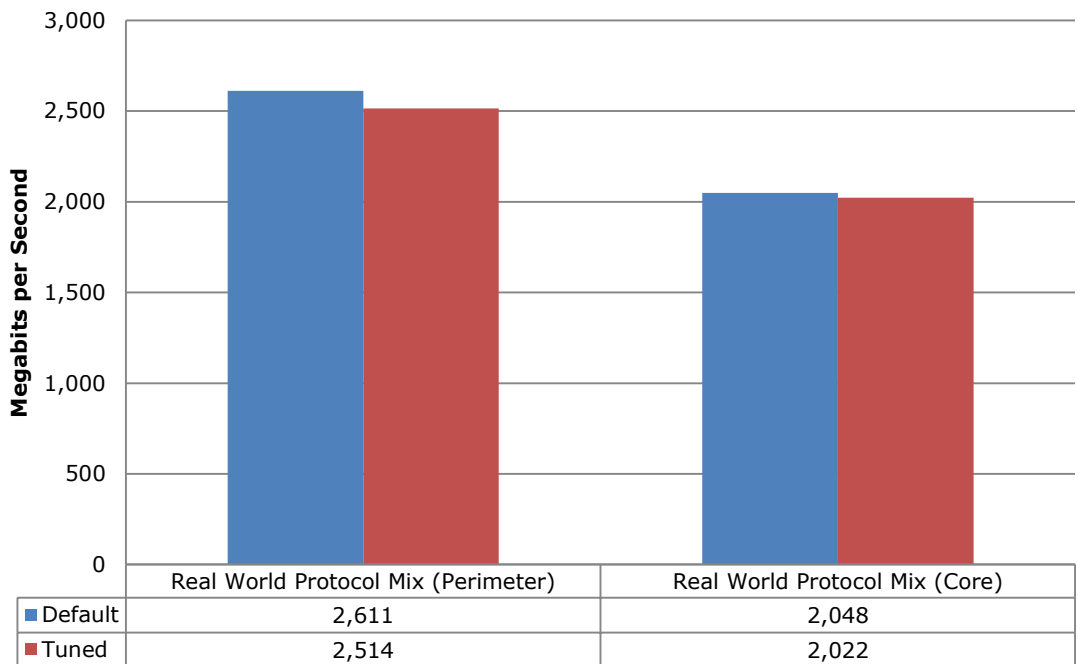


Figure 3: Real-World Traffic Mixes – Default vs. Tuned Configurations

3.2 CONNECTION DYNAMICS – CONCURRENCY AND CONNECTION RATES

The aim of these tests is to stress the detection engine and determine how the sensor copes with large numbers of TCP connections per second, application layer transactions per second, and concurrent open connections. All packets contain valid payload and address data and these tests provide an excellent representation of a live network at various connection/transaction rates.

Note that in all tests, the following critical “breaking points”—where the final measurements are taken—are used:

Excessive concurrent TCP connections - latency within the IPS is causing unacceptable increase in open connections on the server-side.

Excessive response time for HTTP transactions/SMTP sessions - latency within the IPS is causing excessive delays and increased response time to the client.

Unsuccessful HTTP transactions/SMTP sessions – normally, there should be zero unsuccessful transactions. Once these appear, it is an indication that excessive latency within the IPS is causing connections to time out.




Figure 4: Concurrency and Connection Rates – Default vs. Tuned Configurations

3.3 HTTP CONNECTIONS PER SECOND AND CAPACITY

These tests aim to stress the HTTP detection engine in order to determine how the sensor copes with detecting and blocking exploits under network loads of varying average packet size and varying connections per second. By creating genuine session-based traffic with varying session lengths, the sensor is forced to track valid TCP sessions, thus ensuring a higher workload than for simple packet-based background traffic.

Each transaction consists of a single HTTP GET request and there are no transaction delays (i.e. the web server responds immediately to all requests). All packets contain valid payload (a mix of binary and ASCII objects) and address data. This test provides an excellent representation of a live network (albeit one biased towards HTTP traffic) at various network loads.




Figure 5: HTTP Connections per Second and Capacity – Default vs. Tuned Configurations

3.4 HTTP CONNECTIONS PER SECOND AND CAPACITY WITH DELAYS

Typical user behavior introduces delays in between requests and reponses, e.g. as users read web pages and decide which links to click next. This next set of tests is identical to the previous set except that these include a 10-second delay in the server response for each transaction. This has the effect of maintaining a high number of open connections throughout the test, thus forcing the sensor to utilize additional resources to track those connections.

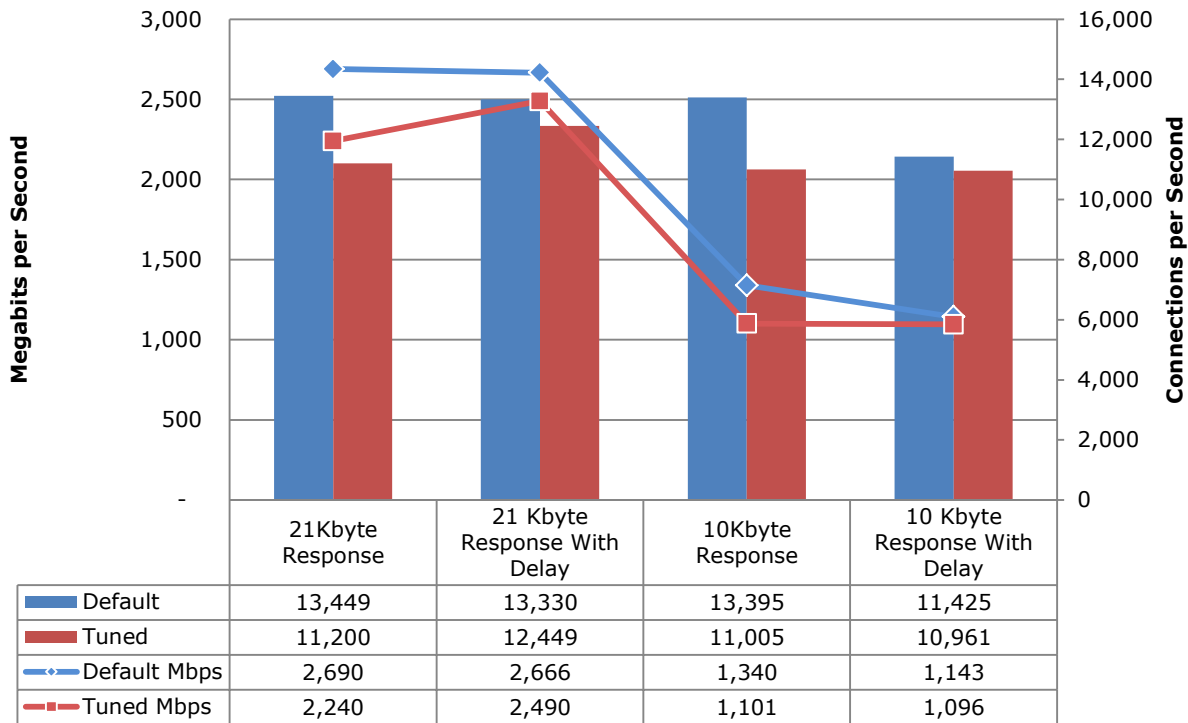


Figure 6: HTTP Connections per Second and Capacity (with/without Delays)

3.5 UDP THROUGHPUT

The aim of this test is purely to determine the raw packet processing capability of each in-line port pair of the IPS. It is not real world, and can be misleading. It is included here primarily for legacy purposes.

This traffic does not attempt to simulate any form of “real-world” network condition. No TCP sessions are created during this test, and there is very little for the detection engine to do in the way of protocol analysis (although each vendor will be required to write a signature to detect the test packets to ensure that they are being passed through the detection engine and not “fast-tracked” from the inbound to outbound port).

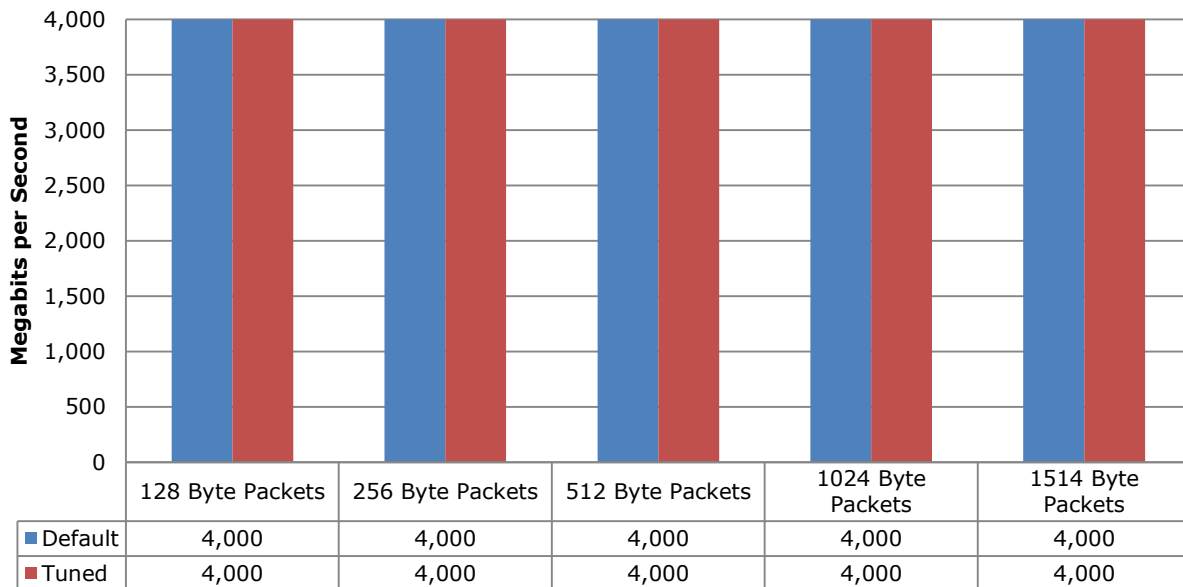


Figure 7: UDP Throughput – Default vs. Tuned Configurations

4 TOTAL COST OF OWNERSHIP

IPS solutions can be complex projects with several factors affecting the overall cost of deployment, maintenance and upkeep. All of these should be considered over the course of the useful life of the solution.

- Product Purchase – the cost of acquisition.
- Product Maintenance – the fees paid to the vendor.
- Installation – the time required to take the device out of the box, configure it, put it into the network, apply updates and patches, initial tuning, and set up desired logging and reporting.
- Upkeep – the time required to apply periodic updates and patches from vendors, including hardware, software, and protection (signature/filter/rules) updates.
- Tuning – the time required to configure the policy such that the best possible protection is applied while reducing or eliminating false alarms and false positives.

4.1 LABOR PER PRODUCT (IN HOURS)

This table estimates the annual labor required to maintain each device. Since vendors sent their very best engineers to tune, NSS Labs' assumptions are based upon the time required by a highly experienced security engineer (\$75 per hour fully loaded). This allowed us to hold the talent cost variable constant and measure only the difference in time required to tune.

Product	Installation (Hrs)	Upkeep / Year (Hrs)	Tuning / Year (Hrs)
Palo Alto Networks PA-4020	8	25	25

4.2 PURCHASE PRICE AND TOTAL COST OF OWNERSHIP

Each vendor provided pricing information. When possible, we selected the 24/7 maintenance and support option with 24-hour replacement as this is the option most organizations will select.

Product	Purchase	Maintenance / year	1 Year TCO	2 Year TCO	3 Year TCO
Palo Alto Networks PA-4020	\$35,000	\$11,200	\$50,550	\$65,500	\$80,450

- Year One TCO was determined by multiplying the Labor Rate (\$75 per hour fully loaded) x (Installation + Upkeep + Tuning) and then adding the Purchase Price + Maintenance.
- Year Two TCO was determined by multiplying the Labor Rate (\$75 per hour fully loaded) x (Upkeep + Tuning) and then adding Year One TCO.
- Year Three TCO was determined by multiplying the Labor Rate (\$75 per hour fully loaded) x (Upkeep + Tuning) and then adding Year Two TCO.

4.3 VALUE: COST PER MBPS AND EXPLOIT BLOCKED – TUNED POLICY

There is a clear difference between price and value. The least expensive product does not necessarily offer the greatest value if it blocks fewer exploits than competitors. The best value is a product with a low TCO and high level of secure throughput (security effectiveness x performance).

The following table illustrates the relative cost per unit of work performed: Mbps-Protected

Product	Protection	Throughput	3 Year TCO	Price / Mbps-Protected
Palo Alto Networks PA-4020	93.4%	2,259	\$80,450	\$38

Price per Protected Mbps was calculated by taking the Three-Year TCO and dividing it by the product of Protection x Throughput. $\text{Three-Year TCO} / (\text{Protection} \times \text{Throughput}) = \text{Price/Mbps-Protected}$.

5 DETAILED PRODUCT SCORECARD

The following chart depicts the status of each test with quantitative results where applicable. A separate product Exposure Report details specific vulnerabilities that are not protected.

Test ID	Description	Default	Tuned
5.1	Detection Engine		
5.1.1	System Exposure	59%	95%
5.1.2	Service Exposure	52%	93%
5.1.3	System or Service Fault	42%	80%
5.2	Threat Vectors		
5.2.1	Attacker Initiated	49%	89%
5.2.2	Target Initiated	64%	97%
5.2.3	Network	57%	93%
5.2.4	Local	Not for NIPS	Not for NIPS
5.3	Target Type		
5.3.1	Web Server	* See VulnerabilityScope	
5.3.2	Web Browser	*	*
5.3.3	ActiveX	*	*
5.3.4	JavaScript	*	*
5.3.5	Browser Plug-ins / Add-ons	*	*
5.4	Coverage by Result		
5.4.1	Arbitrary Code Execution	*	*
5.4.2	Buffer Overflow	*	*
5.4.3	Code Injection	*	*
5.4.4	Cross site script	*	*
5.4.5	Directory Traversal	*	*
5.4.6	Privilege Escalation	*	*
5.5	Coverage by Vendor		
5.5.1	3Com	*	*
5.5.2	Adobe	*	*
5.5.3	Alt-N	*	*
5.5.4	Apache	*	*
5.5.5	Apple	*	*
5.5.6	Atrium	*	*
5.5.7	Avast	*	*
5.5.8	BEA	*	*
5.5.9	BitDefender	*	*
5.5.10	Borland	*	*
5.5.11	CA	*	*
5.5.12	Cisco	*	*
5.5.13	Citrix	*	*

Licensed to Palo Alto Networks

Network Intrusion Prevention Systems Individual Product Test Results
 © 2010 NSS Labs, Inc. All rights reserved.

Test ID	Description	Default	Tuned
5.5.14	ClamAV	*	*
5.5.15	EMC	*	*
5.5.16	Facebook	*	*
5.5.17	GNU	*	*
5.5.18	Google	*	*
5.5.19	HP	*	*
5.5.20	IBM	*	*
5.5.21	IPSwitch	*	*
5.5.22	ISC	*	*
5.5.23	Kaspersky	*	*
5.5.24	LanDesk	*	*
5.5.25	lighttpd	*	*
5.5.26	Linux	*	*
5.5.27	Macromedia	*	*
5.5.28	MacroVision	*	*
5.5.29	Mailenable	*	*
5.5.30	McAfee	*	*
5.5.31	Mercury	*	*
5.5.32	Microsoft	*	*
5.5.33	MIT	*	*
5.5.34	Mozilla	*	*
5.5.35	Mplayer	*	*
5.5.36	Multiple Vendors	*	*
5.5.37	MySQL	*	*
5.5.38	NOD32	*	*
5.5.39	Novell	*	*
5.5.40	Nullsoft	*	*
5.5.41	OpenLDAP	*	*
5.5.42	OpenOffice	*	*
5.5.43	OpenSSH	*	*
5.5.44	OpenSSL	*	*
5.5.45	Oracle	*	*
5.5.46	Other Misc	*	*
5.5.47	Panda	*	*
5.5.48	RealNetworks	*	*
5.5.49	Samba	*	*
5.5.50	SAP	*	*
5.5.51	Snort	*	*
5.5.52	Sophos	*	*
5.5.53	SpamAssassin	*	*

Licensed to Palo Alto Networks

Network Intrusion Prevention Systems Individual Product Test Results
 © 2010 NSS Labs, Inc. All rights reserved.

Test ID	Description	Default	Tuned
5.5.54	Squid	*	*
5.5.55	Sun Microsystems	*	*
5.5.56	Symantec	*	*
5.5.57	Trend Micro	*	*
5.5.58	Trillian	*	*
5.5.59	UltraVNC	*	*
5.5.60	Veritas	*	*
5.5.61	VideoLan	*	*
5.5.62	VMWare	*	*
5.5.63	WinAmp	*	*
5.5.64	WinFTP	*	*
5.5.65	Winzip	*	*
5.5.66	Yahoo	*	*
5.6	Evasion		
5.6.1	Evasion	100%	100%
5.7	Packet Fragmentation		
5.7.1	Ordered 8 byte fragments	100%	100%
5.7.2	Ordered 24 byte fragments	100%	100%
5.7.3	Out of order 8 byte fragments	100%	100%
5.7.4	Ordered 8 byte fragments, duplicate last packet	100%	100%
5.7.5	Out of order 8 byte fragments, duplicate last packet	100%	100%
5.7.6	Ordered 8 byte fragments, reorder fragments in reverse	100%	100%
5.7.7	Ordered 16 byte frags, fragment overlap (favor new)	100%	100%
5.7.8	Ordered 16 byte frags, fragment overlap (favor old)	100%	100%
5.7.9	Out of order 8 byte fragments, interleaved duplicate packets scheduled for later delivery	100%	100%
5.8	Stream Segmentation		
5.8.1	Ordered 1 byte segments, interleaved duplicate segments with invalid TCP checksums	100%	100%
5.8.2	Ordered 1 byte segments, interleaved duplicate segments with null TCP control flags	100%	100%
5.8.3	Ordered 1 byte segs, interleaved duplicate segments with requests to resync sequence numbers mid-stream	100%	100%
5.8.4	Ordered 1 byte segments, duplicate last packet	100%	100%
5.8.5	Ordered 2 byte segments, segment overlap (favor new)	100%	100%
5.8.6	Ordered 1 byte segments, interleaved duplicate segments with out-of-window sequence numbers	100%	100%
5.8.7	Out of order 1 byte segments	100%	100%
5.8.8	Out of order 1 byte segments, interleaved duplicate segments with faked retransmits	100%	100%
5.8.9	Ordered 1 byte segments, segment overlap (favor new)	100%	100%
5.8.10	Out of order 1 byte segs, PAWS elimination (interleaved dup segs with older TCP timestamp options)	100%	100%
5.8.11	Ordered 16 byte segs, seg overlap (favor new (Unix))	100%	100%

Licensed to Palo Alto Networks

Test ID	Description	Default	Tuned
5.9	RPC Fragmentation		
5.9.1	One-byte fragmentation (ONC)	100%	100%
5.9.2	Two-byte fragmentation (ONC)	100%	100%
5.9.3	All fragments, including Last Fragment (LF) will be sent in one TCP segment (ONC)	100%	100%
5.9.4	All frags except Last Fragment (LF) will be sent in one TCP segment. LF will be sent in separate TCP seg (ONC)	100%	100%
5.9.5	One RPC fragment will be sent per TCP segment (ONC)	100%	100%
5.9.6	One LF split over more than one TCP segment. In this case no RPC fragmentation is performed (ONC)	100%	100%
5.9.7	Canvas Reference Implementation Level 1 (MS)	100%	100%
5.9.8	Canvas Reference Implementation Level 2 (MS)	100%	100%
5.9.9	Canvas Reference Implementation Level 3 (MS)	100%	100%
5.9.10	Canvas Reference Implementation Level 4 (MS)	100%	100%
5.9.11	Canvas Reference Implementation Level 5 (MS)	100%	100%
5.9.12	Canvas Reference Implementation Level 6 (MS)	100%	100%
5.9.13	Canvas Reference Implementation Level 7 (MS)	100%	100%
5.9.14	Canvas Reference Implementation Level 8 (MS)	100%	100%
5.9.15	Canvas Reference Implementation Level 9 (MS)	100%	100%
5.9.16	Canvas Reference Implementation Level 10 (MS)	100%	100%
5.1	URL Obfuscation		
5.10.1	URL encoding - Level 1 (minimal)	100%	100%
5.10.2	URL encoding - Level 2	100%	100%
5.10.3	URL encoding - Level 3	100%	100%
5.10.4	URL encoding - Level 4	100%	100%
5.10.5	URL encoding - Level 5	100%	100%
5.10.6	URL encoding - Level 6	100%	100%
5.10.7	URL encoding - Level 7	100%	100%
5.10.8	URL encoding - Level 8 (extreme)	100%	100%
5.10.9	Premature URL ending	100%	100%
5.10.10	Long URL	100%	100%
5.10.11	Fake parameter	100%	100%
5.10.12	TAB separation	100%	100%
5.10.13	Case sensitivity	100%	100%
5.10.14	Windows \ delimiter	100%	100%
5.10.15	Session splicing	100%	100%
5.11	HTML Obfuscation		
5.11.1	UTF-16 character set encoding (big-endian)	100%	100%
5.11.2	UTF-16 character set encoding (little-endian)	100%	100%
5.11.3	UTF-32 character set encoding (big-endian)	100%	100%
5.11.4	UTF-32 character set encoding (little-endian)	100%	100%

Test ID	Description	Default	Tuned
5.11.5	UTF-7 character set encoding	100%	100%
5.11.6	Chunked encoding (random chunk size)	100%	100%
5.11.7	Chunked encoding (fixed chunk size)	100%	100%
5.11.8	Chunked encoding (chaffing)	100%	100%
5.11.9	Compression (Deflate)	100%	100%
5.11.10	Compression (Gzip)	100%	100%
5.11.11	Base-64 Encoding	100%	100%
5.11.12	Base-64 Encoding (shifting 1 bit)	100%	100%
5.11.13	Base-64 Encoding (shifting 2 bits)	100%	100%
5.11.14	Base-64 Encoding (chaffing)	100%	100%
5.11.15	Combination UTF-7 + Gzip	100%	100%
5.12	FTP Evasion		
5.12.1	Inserting spaces in FTP command lines	100%	100%
5.12.2	Inserting non-text Telnet opcodes - Level 1 (minimal)	100%	100%
5.12.3	Inserting non-text Telnet opcodes - Level 2	100%	100%
5.12.4	Inserting non-text Telnet opcodes - Level 3	100%	100%
5.12.5	Inserting non-text Telnet opcodes - Level 4	100%	100%
5.12.6	Inserting non-text Telnet opcodes - Level 5	100%	100%
5.12.7	Inserting non-text Telnet opcodes - Level 6	100%	100%
5.12.8	Inserting non-text Telnet opcodes - Level 7	100%	100%
5.12.9	Inserting non-text Telnet opcodes - Level 8 (extreme)	100%	100%
6	NIPS Performance		
6.1	Raw Packet Processing Performance (UDP Traffic)	Mbps	Mbps
6.1.1	128 Byte Packets	4,000	4,000
6.1.2	256 Byte Packets	4,000	4,000
6.1.3	512 Byte Packets	4,000	4,000
6.1.4	1024 Byte Packets	4,000	4,000
6.1.5	1514 Byte Packets	4,000	4,000
6.2	Maximum Capacity		
6.2.1	Concurrent TCP Connections	527,048	527,048
6.2.2	Concurrent TCP Conns w/Data	508,065	508,065
6.2.3	Stateful Protection at Max Concurrent Connections	PASS	PASS
6.2.4	TCP Connections/Second	2,912 ¹	2,908
6.2.5	HTTP Connections/Second	16,000	15,176
6.2.6	HTTP Transactions/Second	24,668	24,361
6.3	Behavior Of The State Engine Under Load		
6.3.1	Attack Detection/Blocking - Normal Load	100%	100%

¹ TCP Connections per Second results were low due to app-id functionality that recursively examined the traffic to try and determine the application. Since the test was of TCP connections, there was no application associated with the traffic.

Test ID	Description	Default	Tuned
6.3.2	State Preservation - Normal Load	100%	100%
6.3.3	Pass Legitimate Traffic - Normal Load	100%	100%
6.3.4	Attack Detection/Blocking - Maximum Exceeded	100%	100%
6.3.5	State Preservation - Maximum Exceeded	100%	100%
6.3.6	Pass Legitimate Traffic - Maximum Exceeded	100%	100%
6.4	HTTP Capacity With No Transaction Delays	CPS	CPS
6.4.1	44Kbyte Response	10,999	10,804
6.4.2	21Kbyte Response	13,449	11,200
6.4.3	10Kbyte Response	13,395	11,005
6.4.4	4.5Kbyte Response	12,403	12,000
6.4.5	1.7Kbyte Response	13,208	12,404
6.5	HTTP Capacity With Transaction Delays	CPS	CPS
6.5.1	21 Kbyte Response With Delay	13,330	12,449
6.5.2	10 Kbyte Response With Delay	11,425	10,961
6.6	"Real World" Traffic	Mbps	Mbps
6.6.1	Real World Protocol Mix (Perimeter)	2,611	2,514
6.6.2	Real World Protocol Mix (Core)	2,048	2,022
7	Management & Configuration Costs		
7.1	Ease of Use		
7.1.1	Initial Setup (Hours)	8	8
7.1.2	Time Required for Upkeep (Hours per Year)	25	25
7.1.3	Time Required to Tune (Hours per Year)	0	25
7.2	Expected Costs		
7.2.1	Initial Purchase	35,000	35,000
7.2.2	Ongoing Maintenance & Support (Annual)	11,200	11,200
7.2.3	Installation Labor Cost (@\$75/hr)	\$600	\$600
7.2.4	Management Labor Cost (per Year @\$75/hr)	\$1,875	\$1,875
7.2.5	Tuning Labor Cost (per Year @\$75/hr)	\$0	\$1,875
7.3	Total Cost of Ownership		
7.3.1	Year 1	\$48,675	\$50,550
7.3.2	Year 2	\$13,075	\$14,950
7.3.3	Year 3	\$13,075	\$14,950
7.3.4	3 Year Total Cost of Ownership	\$74,825	\$80,450

6 APPENDIX B: SPECIAL THANKS

Special thanks go to our test infrastructure partners who provide much of the equipment, software, and support that make this testing possible:



Find it before they do.™

