



NEXT GENERATION FIREWALL TEST REPORT

Palo Alto Networks PA-7050 v6.0.11-h1

Author – Christopher Conrad

Overview

NSS Labs performed an independent test of the Palo Alto Networks PA-7050 v6.0.11-h1. The product was subjected to thorough testing at the NSS facility in Austin, Texas, based on the Next Generation Firewall (NGFW) Test Methodology v6.0 available at www.nsslabs.com. This test was conducted free of charge and NSS did not receive any compensation in return for Palo Alto Networks’ participation. For additional information on NGFW technology, refer to the NSS Analyst Brief, *What Do CIOs Need to Know About Next Generation Firewalls?*¹

While the companion Comparative Reports on security, performance, and total cost of ownership (TCO) will provide information about all tested products, this Test Report provides detailed information not available elsewhere.

NSS research indicates that NGFW devices are typically deployed to protect users rather than data center assets, and that the majority of enterprises will not separately tune intrusion prevention system (IPS) modules within their NGFWs. Therefore, during NSS testing, NGFW products are configured with the vendor’s pre-defined or recommended (i.e., “out-of-the-box”) settings in order to provide readers with relevant security effectiveness and performance dimensions based on their expected usage.

Product	NSS Exploit Library Block Rate ²	NSS-Tested Throughput	3-Year TCO (List Price)	3-Year TCO (Street Price)
Palo Alto Networks PA-7050 v6.0.11-h1	95.8%	42,324 Mbps	\$1,256,000	\$1,004,920
	Firewall Policy Enforcement	Application Control	Evasions	Stability and Reliability
	PASS	PASS	PASS	PASS

Figure 1 – Overall Test Results

Using the recommended policy, the Palo Alto Networks PA-7050 blocked 94.4% of attacks against server applications, 97.1% of attacks against client applications, and 95.8% of attacks overall. The device proved effective against all evasion techniques tested. The device also passed all stability and reliability tests.

The Palo Alto Networks PA-7050 is rated by NSS at 42,324 Mbps, which is lower than the vendor-claimed performance; Palo Alto Networks rates this device at 60 Gbps. *NSS-Tested Throughput* is calculated as an average of all of the “real-world” protocol mixes and the 21 KB HTTP response-based capacity test.

¹ What Do CIOs Need to Know About Next Generation Firewalls? NSS Labs

² Defined as the rate at which the device under test blocked exploits from the *NSS Exploit Library*. This value is a component of the overall block rate, which is reported in the NSS Labs Security Value Map™.

Table of Contents

Overview	2
Security Effectiveness	5
Firewall Policy Enforcement	5
Application Control	6
NSS Exploit Library	6
<i>False Positive Testing</i>	7
<i>Coverage by Attack Vector</i>	7
<i>Coverage by Impact Type</i>	7
<i>Coverage by Date</i>	8
<i>Coverage by Target Vendor</i>	8
Resistance to Evasion Techniques	9
Performance	10
Raw Packet Processing Performance (UDP Throughput)	10
Raw Packet Processing Performance (UDP Latency)	11
Maximum Capacity	11
HTTP Capacity with No Transaction Delays	13
Application Average Response Time – HTTP	13
HTTP Capacity with Transaction Delays	13
“Real-World” Traffic Mixes	15
Stability and Reliability	16
Management and Configuration	17
Total Cost of Ownership (TCO)	18
Installation Hours	18
List Price and Total Cost of Ownership	19
Street Price and Total Cost of Ownership	19
Detailed Product Scorecard	20
Test Methodology	25
Contact Information	25

Table of Figures

Figure 1 – Overall Test Results.....	2
Figure 2 – Firewall Policy Enforcement	5
Figure 3 – Application Control	6
Figure 4 – Number of Exploits Blocked (%).....	6
Figure 5 – Coverage by Attack Vector	7
Figure 6 – Product Coverage by Date	8
Figure 7 – Product Coverage by Target Vendor.....	8
Figure 8 – Resistance to Evasion Results	9
Figure 9 – Raw Packet Processing Performance (UDP Traffic)	10
Figure 10 – UDP Latency in Microseconds.....	11
Figure 11 – Concurrency and Connection Rates.....	12
Figure 12 – HTTP Capacity with No Transaction Delays	13
Figure 13 – Average Application Response Time (Milliseconds)	13
Figure 14 – HTTP Capacity with Transaction Delays	14
Figure 15 – “Real-World” Traffic Mixes	15
Figure 16 – Stability and Reliability Results	16
Figure 17 – Sensor Installation Time (Hours).....	18
Figure 18 – List Price 3-Year TCO	19
Figure 19 – Street Price 3-Year TCO.....	19
Figure 20 – Detailed Scorecard.....	24

Security Effectiveness

This section verifies that the device under test (DUT) is capable of enforcing the security policy effectively.

Firewall Policy Enforcement


Policies are rules configured on a firewall to permit or deny access from one network resource to another based on identifying criteria such as source, destination, and service. A term typically used to define the demarcation point of a network where policy is applied is *demilitarized zone* (DMZ). Policies are typically written to permit or deny network traffic from one or more of the following zones:

- **Untrusted** – This is typically an external network and is considered unknown and nonsecure. An example of an untrusted network would be the Internet.
- **DMZ** – This is a network that is being *isolated* by the firewall restricting network traffic to and from hosts contained within the isolated network.
- **Trusted** – This is typically an internal network; a network that is considered secure and protected.

The NSS firewall tests verify performance and the ability to enforce policy between the following:

- Trusted to Untrusted
- Untrusted to DMZ
- Trusted to DMZ

Note: Firewalls must provide at a minimum one DMZ interface in order to provide a DMZ or “transition point” between untrusted and trusted networks.



Test Procedure	Result
Baseline Policy	PASS
Simple Policy	PASS
Complex Policy	PASS
Static NAT	PASS
Dynamic/Hide NAT	PASS
SYN Flood Protection	PASS
IP Address Spoofing Protection	PASS
TCP Split Handshake Spoof	PASS

Figure 2 – Firewall Policy Enforcement

Application Control

An NGFW must provide granular control based on applications, not just ports. This capability is needed to reestablish a secure perimeter where unwanted applications are unable to tunnel over HTTP/S. As such, granular application control is a requirement of an NGFW since it enables the administrator to define security policies based on applications rather than on ports alone.

Test Procedure	Result
Block Unwanted Applications	PASS
Block Specific Actions	PASS

Figure 3 – Application Control

Our testing found that the Palo Alto Networks PA-7050 correctly enforced complex outbound and inbound policies consisting of multiple rules, objects, and applications. NSS engineers verified that the device successfully determined the correct application and took the appropriate action based on the policy.

NSS Exploit Library

NSS’ security effectiveness testing leverages the deep expertise of our engineers who utilize multiple commercial, open-source, and proprietary tools, including NSS’ network live stack test environment³ as appropriate. With 1999 exploits, this is the industry’s most comprehensive test to date. Most notably, all of the exploits and payloads in this test have been validated such that:

- A reverse shell is returned
- A bind shell is opened on the target, allowing the attacker to execute arbitrary commands
- Arbitrary code is executed
- A malicious payload is installed
- A system is rendered unresponsive
- Etc.

Product	Total Number of Exploits Run	Total Number Blocked	Block Percentage
Palo Alto Networks PA-7050 v6.0.11-h1	1,999	1,916	95.8%

Figure 4 – Number of Exploits Blocked (%)

³ See the NSS Cyber Advanced Warning System™ for more details.

False Positive Testing

The Palo Alto Networks PA-7050 v6.0.11-h1 correctly identified traffic and did not fire alerts for non-malicious content.

Coverage by Attack Vector

Because a failure to block attacks could result in significant compromise and could severely impact critical business systems, NGFWs should be evaluated against a broad set of exploits. Exploits can be categorized as either *attacker initiated* or *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. Target-initiated exploits are the most common type of attack experienced by the end user, and the attacker has little or no control as to when the threat is executed.

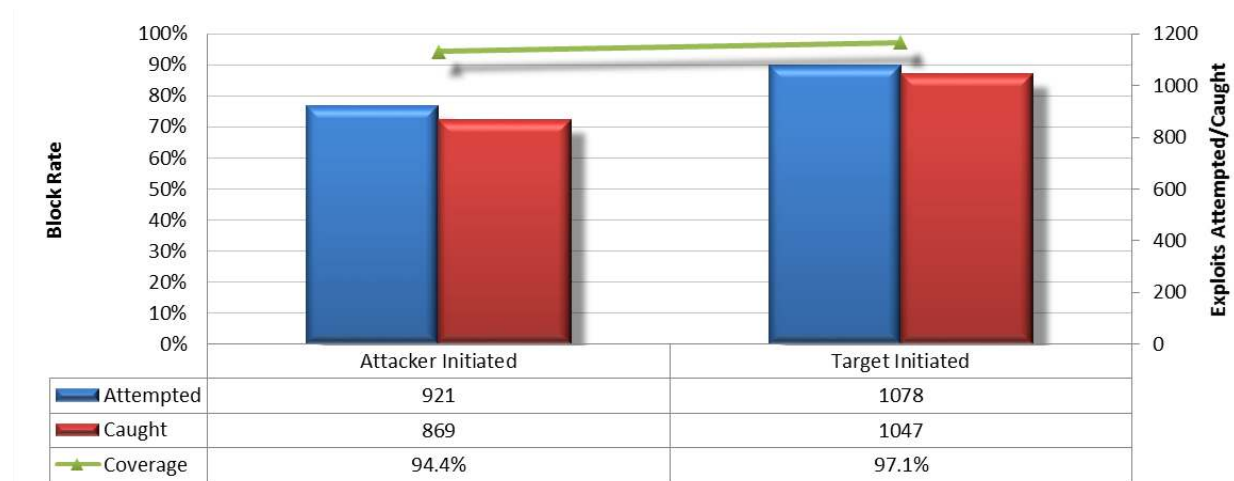


Figure 5 – Coverage by Attack Vector

Coverage by Impact Type

The most serious exploits are those that 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. Finally, there are attacks that 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. Clients can contact NSS for more information about these tests.

Coverage by Date

Figure 6 provides insight into whether or not a vendor is aging out protection signatures aggressively enough to preserve performance levels. It also reveals whether a product lags behind in protection for the most current vulnerabilities. NSS reports exploits by individual years for the past ten years. Exploits older than ten years are grouped together.




Figure 6 – Product Coverage by Date

Coverage by Target Vendor

Exploits within the *NSS Exploit Library* target a wide range of protocols and applications. Figure 7 depicts the coverage offered by the Palo Alto Networks PA-7050 for five of the top vendors targeted in this test. More than 70 vendors are represented in the test. Clients can contact NSS for more information about this test.




Figure 7 – Product Coverage by Target Vendor

Resistance to Evasion Techniques

Evasion techniques are a means of disguising and modifying attacks at the point of delivery to avoid detection and blocking by security products. Failure of a security device to correctly identify a specific type of evasion potentially allows an attacker to use an entire class of exploits for which the device is assumed to have protection. This renders the device virtually useless. Many of the techniques used in this test have been widely known for years and should be considered minimum requirements for the NGFW product category.

Providing exploit protection results without fully factoring in evasion can be misleading. The more classes of evasion that are missed (such as IP packet fragmentation, stream segmentation, RPC fragmentation, URL obfuscation, HTML obfuscation, payload encoding, and FTP evasion), the less effective the device. For example, it is better to miss all techniques in one evasion category, such as FTP evasion, than one technique in each category, which would result in a broader attack surface.

Furthermore, evasions operating at the lower layers of the network stack (IP packet fragmentation or stream segmentation) have a greater impact on security effectiveness than those operating at the upper layers (HTTP or FTP obfuscation.) Lower-level evasions will potentially impact a wider number of exploits; missing TCP segmentation, for example, is a much more serious issue than missing FTP obfuscation.

Figure 8 provides the results of the evasion tests for the Palo Alto Networks PA-7050.

Test Procedure	Result
IP Packet Fragmentation	PASS
Stream Segmentation	PASS
RPC Fragmentation	PASS
URL Obfuscation	PASS
HTML Obfuscation	PASS
Payload Encoding	PASS
FTP Evasion	PASS
IP Packet Fragmentation + TCP Segmentation	PASS

Figure 8 – Resistance to Evasion Results

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 that security shortcuts are not taken to maintain or improve performance.

Raw Packet Processing Performance (UDP Throughput)

This test uses UDP packets of varying sizes generated by test equipment. A constant stream of the appropriate packet size, with variable source and destination IP addresses transmitting from a fixed source port to a fixed destination port, is transmitted bidirectionally through each port pair of the DUT.

Each packet contains dummy data and is targeted at a valid port on a valid IP address on the target subnet. The percentage load and frames per second (fps) figures across each inline port pair are verified by network monitoring tools before each test begins. Multiple tests are run and averages are taken where necessary.

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 state engine to do. The aim of this test is to determine the raw packet processing capability of each inline port pair of the DUT, and to determine the DUT’s effectiveness at forwarding packets quickly, in order to provide the highest level of network performance and with the least amount of latency.




Figure 9 – Raw Packet Processing Performance (UDP Traffic)

Raw Packet Processing Performance (UDP Latency)

NGFWs that introduce high levels of latency lead to unacceptable response times for users, especially where multiple security devices are placed in the data path. Figure 10 depicts UDP latency (in microseconds) as recorded during the UDP throughput tests at 90% of maximum load.

Latency – UDP	Microseconds
64-Byte Packets	10
128-Byte Packets	11
256-Byte Packets	12
512-Byte Packets	13
1024-Byte Packets	15
1514-Byte Packets	17

Figure 10 – UDP Latency in Microseconds

Maximum Capacity

The use of traffic generation appliances allows NSS engineers to create “real-world” traffic at multi-Gigabit speeds as a background load for the tests. The aim of these tests is to stress the inspection engine and determine how it copes with high volumes 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 NGFW is causing an unacceptable increase in open connections.
- **Excessive concurrent HTTP connections** – Latency within the NGFW is causing excessive delays and increased response time.
- **Unsuccessful HTTP transactions** – Normally, there should be zero unsuccessful transactions. Once these appear, it is an indication that excessive latency within the NGFW is causing connections to time out.




Figure 11 – Concurrency and Connection Rates

HTTP Capacity with No Transaction Delays

The aim of these tests is to stress the HTTP detection engine and determine how the DUT copes with network loads of varying average packet size and varying connections per second. By creating genuine session-based traffic with varying session lengths, the DUT is forced to track valid TCP sessions, thus ensuring a higher workload than for simple packet-based background traffic. This provides a test environment that is as close to “real-world” conditions as possible, while ensuring absolute accuracy and repeatability.

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 toward HTTP traffic) at various network loads.



Figure 12 – HTTP Capacity with No Transaction Delays

Application Average Response Time – HTTP

Application Average Response Time – HTTP (at 90% Maximum Load)	Milliseconds
2,500 Connections per Second – 44 KB Response	3.29
5,000 Connections per Second – 21 KB Response	2.25
10,000 Connections per Second – 10 KB Response	1.28
20,000 Connections per Second – 4.5 KB Response	1.10
40,000 Connections per Second – 1.7 KB Response	1.30

Figure 13 – Average Application Response Time (Milliseconds)

HTTP Capacity with Transaction Delays

Typical user behavior introduces delays between requests and responses (for example, “think time”) as users read web pages and decide which links to click next. This group of tests is identical to the previous group except that these include a five-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.

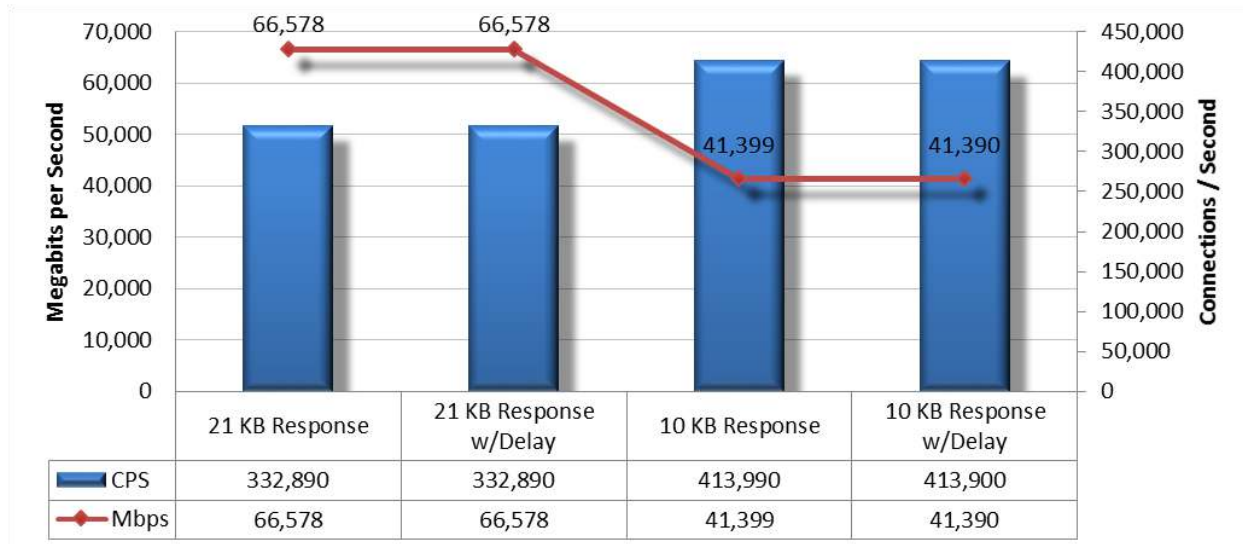


Figure 14 – HTTP Capacity with Transaction Delays

“Real-World” Traffic Mixes

This test measures 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 intended location of the device under test (network core or perimeter) to reflect real use cases. For details about “real-world” traffic protocol types and percentages, see the NSS Labs Next Generation Firewall Test Methodology, available at www.nsslabs.com.

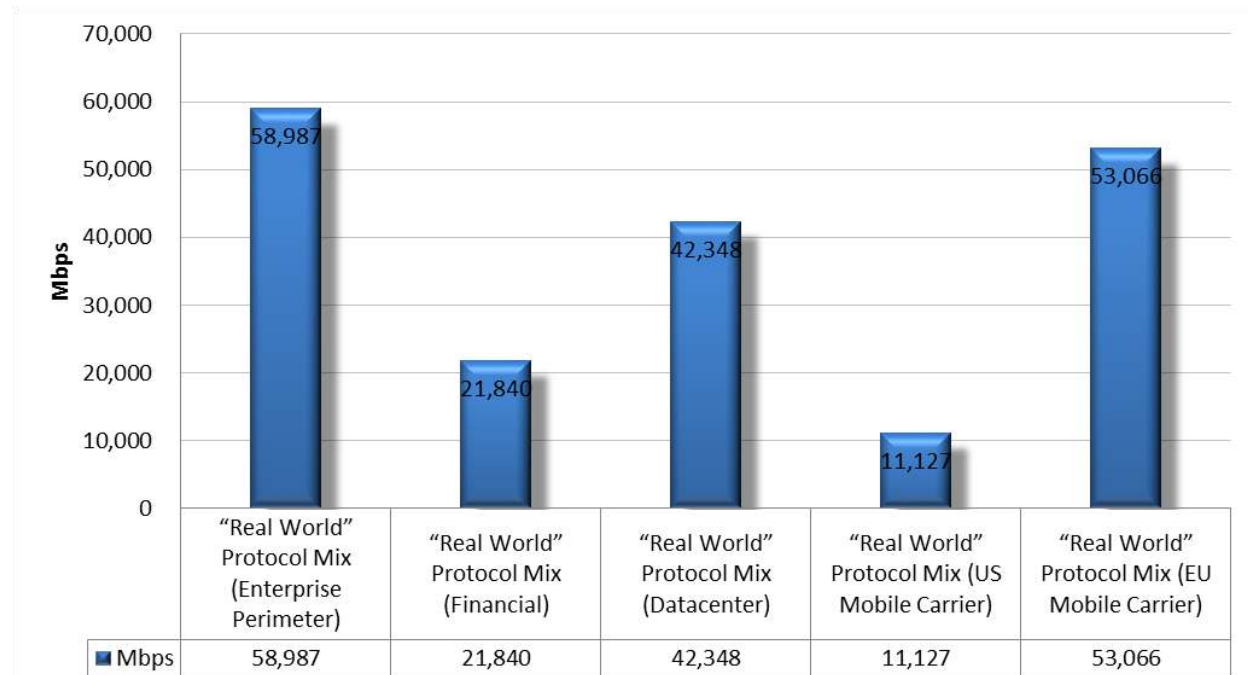


Figure 15 – “Real-World” Traffic Mixes

The Palo Alto Networks PA-7050 was tested by NSS and performed below the throughput claimed by the vendor for all “real-world” traffic mixes.

Stability and Reliability

Long-term stability is particularly important for an inline device, where failure can produce network outages. These tests verify the stability of the DUT along with its ability to maintain security effectiveness while under normal load and while passing malicious traffic. Products that cannot sustain legitimate traffic (or that crash) while under hostile attack will not pass.

The device is required to remain operational and stable throughout these tests, and to block 100% of previously blocked traffic, raising an alert for each. If any non-allowed traffic passes successfully, caused either by the volume of traffic or by the DUT failing open for any reason, the device will fail the test.

Stability and Reliability	Result
Blocking under Extended Attack	PASS
Passing Legitimate Traffic under Extended Attack	PASS
Behavior of the State Engine under Load	
<ul style="list-style-type: none"> Attack Detection/Blocking – Normal Load 	PASS
<ul style="list-style-type: none"> State Preservation – Normal Load 	PASS
<ul style="list-style-type: none"> Pass Legitimate Traffic – Normal Load 	PASS
<ul style="list-style-type: none"> State Preservation – Maximum Exceeded 	PASS
<ul style="list-style-type: none"> Drop Traffic – Maximum Exceeded 	PASS
Protocol Fuzzing and Mutation	PASS
Power Fail	PASS
Persistence of Data	PASS

Figure 16 – Stability and Reliability Results

These tests also determine the behavior of the state engine under load. All NGFW devices must choose whether to risk denying legitimate traffic or risk allowing malicious traffic once they run low on resources. An NGFW device will drop new connections when resources (such as state table memory) are low, or when traffic loads exceed its capacity. In theory, this means the NGFW will block legitimate traffic but maintain state on existing connections (and prevent attack leakage).

Management and Configuration

Security devices are complicated to deploy; essential systems such as centralized management console options, log aggregation, and event correlation/management systems further complicate the purchasing decision.

Understanding key comparison points will allow customers to model the overall impact on network service level agreements (SLAs), to estimate operational resource requirements to maintain and manage the systems, and to better evaluate the required skills/competencies of staff.

Enterprises should include management and configuration during their evaluations, focusing on the following at a minimum:

- **General Management and Configuration** – How easy is it to install and configure devices, and how easy is it to deploy multiple devices throughout a large enterprise network?
- **Policy Handling** – How easy is it to create, edit, and deploy complicated security policies across an enterprise?
- **Alert Handling** – How accurate and timely is the alerting, and how easy is it to drill down to locate critical information needed to remediate a security problem?
- **Reporting** – How effective is the reporting capability, and how readily can it be customized?

Total Cost of Ownership (TCO)

Implementation of security solutions can be complex, with several factors affecting the overall cost of deployment, maintenance, and upkeep. All of the following 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, including software and hardware support, maintenance, and other updates
- **Installation** – The time required to take the device out of the box, configure it, put it into the network, apply updates and patches, and set up desired logging and reporting
- **Upkeep** – The time required to apply periodic updates and patches from vendors, including hardware, software, and other updates
- **Management** – Day-to-day management tasks, including device configuration, policy updates, policy deployment, alert handling, and so on

For the purposes of this report, capital expenditure (capex) items are included for a single device only (the cost of acquisition and installation).

Installation Hours

This table depicts the number of hours of labor required to install each device using only local device management options. The table accurately reflects the amount of time that NSS engineers, with the help of vendor engineers, needed to install and configure the DUT to the point where it operated successfully in the test harness, passed legitimate traffic, and blocked and detected prohibited or malicious traffic. This closely mimics a typical enterprise deployment scenario for a single device.

The installation cost is based on the time that an experienced security engineer would require to perform the installation tasks described above. This approach allows NSS to hold constant the talent cost and measure only the difference in time required for installation. Readers should substitute their own costs to obtain accurate TCO figures.

Product	Installation (Hours)
Palo Alto Networks PA-7050 v6.0.11-h1	8

Figure 17 – Sensor Installation Time (Hours)

List Price and Total Cost of Ownership

Calculations are based on vendor-provided pricing information. Where possible, the 24/7 maintenance and support option with 24-hour replacement is utilized, since this is the option typically selected by enterprise customers. Prices are for single device management and maintenance only; costs for central management solutions (CMS) may be extra.

Product	Purchase	Maintenance /Year	Year 1 Cost	Year 2 Cost	Year 3 Cost	3-Year TCO
Palo Alto Networks PA-7050 v6.0.11-h1	\$1,025,000	\$76,800	\$1,102,400	\$76,800	\$76,800	\$1,256,000

Figure 18 – List Price 3-Year TCO

- **Year 1 Cost** is calculated by adding installation costs (US\$75 per hour fully loaded labor x installation time) + purchase price + first-year maintenance/support fees.
- **Year 2 Cost** consists only of maintenance/support fees.
- **Year 3 Cost** consists only of maintenance/support fees.

Street Price and Total Cost of Ownership

Calculations are based on vendor-provided pricing information. Where possible, the 24/7 maintenance and support option with 24-hour replacement is utilized, since this is the option typically selected by enterprise customers. Prices are for single device management and maintenance only; costs for CMS may be extra.

Product	Purchase	Maintenance /Year	Year 1 Cost	Year 2 Cost	Year 3 Cost	3-Year TCO
Palo Alto Networks PA-7050 v6.0.11-h1	\$820,000	\$61,440	\$882,040	\$61,440	\$61,440	\$1,004,920

Figure 19 – Street Price 3-Year TCO

- **Year 1 Cost** is calculated by adding installation costs (US\$75 per hour fully loaded labor x installation time) + purchase price + first-year maintenance/support fees.
- **Year 2 Cost** consists only of maintenance/support fees.
- **Year 3 Cost** consists only of maintenance/support fees.

For additional TCO analysis, including for the CMS, refer to the TCO Comparative Report.

Detailed Product Scorecard

The following chart depicts the status of each test with quantitative results where applicable.

Description	Result
Security Effectiveness	
Firewall Policy Enforcement	
Baseline Policy	PASS
Simple Policy	PASS
Complex Policy	PASS
Static NAT	PASS
Dynamic / Hide NAT	PASS
SYN Flood Protection	PASS
Address Spoofing Protection	PASS
TCP Split Handshake	PASS
Application Control	
Block Unwanted Applications	PASS
Block Specific Action	PASS
Intrusion Prevention	
Exploit Library	
NSS Exploit Library Block Rate	95.8%
False Positive Testing	PASS
Coverage by Attack Vector	
Attacker Initiated	94.4%
Target Initiated	97.1%
Combined Total	95.8%
Evasions and Attack Leakage	
Resistance to Evasion	
IP Packet Fragmentation	PASS
Ordered 8-byte fragments	PASS
Ordered 16-byte fragments	PASS
Ordered 24-byte fragments	PASS
Ordered 32-byte fragments	PASS
Out of order 8-byte fragments	PASS
Ordered 8-byte fragments, duplicate last packet	PASS
Out of order 8-byte fragments, duplicate last packet	PASS
Ordered 8-byte fragments, reorder fragments in reverse	PASS
Ordered 16-byte fragments, fragment overlap (favor new)	PASS
Ordered 16-byte fragments, fragment overlap (favor old)	PASS
Out of order 8-byte fragments, interleaved duplicate packets scheduled for later delivery	PASS
Ordered 8-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	PASS
Ordered 16-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	PASS
Ordered 24-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	PASS
Ordered 32-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	PASS
TCP Stream Segmentation	PASS
Ordered 1-byte segments, interleaved duplicate segments with invalid TCP checksums	PASS
Ordered 1-byte segments, interleaved duplicate segments with null TCP control flags	PASS
Ordered 1-byte segments, interleaved duplicate segments with requests to resync sequence numbers mid-stream	PASS
Ordered 1-byte segments, duplicate last packet	PASS
Ordered 2-byte segments, segment overlap (favor new)	PASS
Ordered 1-byte segments, interleaved duplicate segments with out-of-window sequence numbers	PASS

Out of order 1-byte segments	PASS
Out of order 1-byte segments, interleaved duplicate segments with faked retransmits	PASS
Ordered 1-byte segments, segment overlap (favor new)	PASS
Out of order 1-byte segments, PAWS elimination (interleaved duplicate segments with older TCP timestamp options)	PASS
Ordered 16-byte segments, segment overlap (favor new (Unix))	PASS
Ordered 32-byte segments	PASS
Ordered 64-byte segments	PASS
Ordered 128-byte segments	PASS
Ordered 256-byte segments	PASS
Ordered 512-byte segments	PASS
Ordered 1024-byte segments	PASS
Ordered 2048-byte segments (sending MSRPC request with exploit)	PASS
Reverse Ordered 256-byte segments, segment overlap (favor new) with random data	PASS
Reverse Ordered 512-byte segments, segment overlap (favor new) with random data	PASS
Reverse Ordered 1024-byte segments, segment overlap (favor new) with random data	PASS
Reverse Ordered 2048-byte segments, segment overlap (favor new) with random data	PASS
Out of order 1024-byte segments, segment overlap (favor new) with random data, Initial TCP sequence number is set to 0xffffffff - 4294967295	PASS
Out of order 2048-byte segments, segment overlap (favor new) with random data, Initial TCP sequence number is set to 0xffffffff - 4294967295	PASS
RPC Fragmentation	PASS
One-byte fragmentation (ONC)	PASS
Two-byte fragmentation (ONC)	PASS
All fragments, including Last Fragment (LF) will be sent in one TCP segment (ONC)	PASS
All frags except Last Fragment (LF) will be sent in one TCP segment. LF will be sent in separate TCP seg (ONC)	PASS
One RPC fragment will be sent per TCP segment (ONC)	PASS
One LF split over more than one TCP segment. In this case no RPC fragmentation is performed (ONC)	PASS
Canvas Reference Implementation Level 1 (MS)	PASS
Canvas Reference Implementation Level 2 (MS)	PASS
Canvas Reference Implementation Level 3 (MS)	PASS
Canvas Reference Implementation Level 4 (MS)	PASS
Canvas Reference Implementation Level 5 (MS)	PASS
Canvas Reference Implementation Level 6 (MS)	PASS
Canvas Reference Implementation Level 7 (MS)	PASS
Canvas Reference Implementation Level 8 (MS)	PASS
Canvas Reference Implementation Level 9 (MS)	PASS
Canvas Reference Implementation Level 10 (MS)	PASS
URL Obfuscation	PASS
URL encoding – Level 1 (minimal)	PASS
URL encoding – Level 2	PASS
URL encoding – Level 3	PASS
URL encoding – Level 4	PASS
URL encoding – Level 5	PASS
URL encoding – Level 6	PASS
URL encoding – Level 7	PASS
URL encoding – Level 8 (extreme)	PASS
Directory Insertion	PASS
Premature URL ending	PASS
Long URL	PASS
Fake parameter	PASS
TAB separation	PASS
Case sensitivity	PASS
Windows \ delimiter	PASS
Session splicing	PASS
HTML Obfuscation	PASS
UTF-16 character set encoding (big-endian)	PASS

UTF-16 character set encoding (little-endian)	PASS
UTF-32 character set encoding (big-endian)	PASS
UTF-32 character set encoding (little-endian)	PASS
UTF-7 character set encoding	PASS
Chunked encoding (random chunk size)	PASS
Chunked encoding (fixed chunk size)	PASS
Chunked encoding (chaffing)	PASS
Compression (Deflate)	PASS
Compression (Gzip)	PASS
Base-64 Encoding	PASS
Base-64 Encoding (shifting 1 bit)	PASS
Base-64 Encoding (shifting 2 bits)	PASS
Base-64 Encoding (chaffing)	PASS
Combination UTF-7 + Gzip	PASS
Payload Encoding	PASS
x86/call4_dword_xor	PASS
x86/fnstenv_mov	PASS
x86/jmp_call_additive	PASS
x86/shikata_ga_nai	PASS
FTP Evasion	PASS
Inserting spaces in FTP command lines	PASS
Inserting non-text Telnet opcodes – Level 1 (minimal)	PASS
Inserting non-text Telnet opcodes – Level 2	PASS
Inserting non-text Telnet opcodes – Level 3	PASS
Inserting non-text Telnet opcodes – Level 4	PASS
Inserting non-text Telnet opcodes – Level 5	PASS
Inserting non-text Telnet opcodes – Level 6	PASS
Inserting non-text Telnet opcodes – Level 7	PASS
Inserting non-text Telnet opcodes – Level 8 (extreme)	PASS
Layered Evasions	PASS
IP Fragmentation + TCP Segmentation	PASS
Ordered 8-byte fragments + Ordered TCP segments except that the last segment comes first	PASS
Ordered 24-byte fragments + Ordered TCP segments except that the last segment comes first	PASS
Ordered 32-byte fragments + Ordered TCP segments except that the last segment comes first	PASS
Ordered 8-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Reverse order TCP segments, segment overlap (favor new), Overlapping data is set to zero bytes	PASS
Ordered 16-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to zero bytes	PASS
Ordered 24-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to zero bytes	PASS
Ordered 32-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to zero bytes	PASS
Ordered 8-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random alphanumeric	PASS
Ordered 16-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random alphanumeric	PASS
Ordered 32-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random alphanumeric	PASS
Ordered 8-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random bytes	PASS

Ordered 16-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random bytes	PASS
Ordered 24-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random bytes	PASS
Ordered 32-byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload + Out of order TCP segments, segment overlap (favor new), Overlapping data is set to random bytes	PASS
Performance	
Raw Packet Processing Performance (UDP Traffic)	Mbps
64-Byte Packets	43,344
128-Byte Packets	77,364
256-Byte Packets	113,964
512-Byte Packets	127,920
1024-Byte Packets	134,880
1514-Byte Packets	144,240
Latency - UDP	Microseconds
64-Byte Packets	10
128-Byte Packets	11
256-Byte Packets	12
512-Byte Packets	13
1024-Byte Packets	15
1514-Byte Packets	17
Maximum Capacity	
Theoretical Max. Concurrent TCP Connections	23,952,345
Theoretical Max. Concurrent TCP Connections w/Data	21,637,073
Maximum TCP Connections Per Second	545,400
Maximum HTTP Connections Per Second	522,210
Maximum HTTP Transactions Per Second	839,880
HTTP Capacity With No Transaction Delays	
2,500 Connections Per Second – 44 KB Response	161,670
5,000 Connections Per Second – 21 KB Response	332,890
10,000 Connections Per Second – 10 KB Response	413,990
20,000 Connections Per Second – 4.5 KB Response	412,600
40,000 Connections Per Second – 1.7 KB Response	448,760
Application Average Response Time - HTTP (at 90% Max Load)	
2,500 Connections Per Second – 44 KB Response	3.3
5,000 Connections Per Second – 21 KB Response	2.3
10,000 Connections Per Second – 10 KB Response	1.3
20,000 Connections Per Second – 4.5 KB Response	1.1
40,000 Connections Per Second – 1.7 KB Response	1.3
HTTP CPS and Capacity With Transaction Delays	
21 KB Response With Delay	332,890
10 KB Response With Delay	413,900
“Real World” Traffic	
“Real-World” Protocol Mix (Enterprise Perimeter)	58,987
“Real-World” Protocol Mix (Financial)	21,840
“Real-World” Protocol Mix (Datacenter)	42,348
“Real-World” Protocol Mix (US Mobile Carrier)	11,127
“Real-World” Protocol Mix (EU Mobile Carrier)	53,066
Stability and Reliability	
Blocking Under Extended Attack	PASS
Passing Legitimate Traffic Under Extended Attack	PASS
Behavior Of The State Engine Under Load	
Attack Detection/Blocking – Normal Load	PASS
State Preservation – Normal Load	PASS

Pass Legitimate Traffic – Normal Load	PASS
State Preservation – Maximum Exceeded	PASS
Drop Traffic – Maximum Exceeded	PASS
Protocol Fuzzing and Mutation	PASS
Power Fail	PASS
Redundancy	YES
Persistence of Data	PASS
Total Cost of Ownership (List Price)	
Ease of Use	
Initial Setup (Hours)	8
Time Required for Upkeep (Hours per Year)	See Comparative
Time Required to Tune (Hours per Year)	See Comparative
Expected Costs	
Initial Purchase (hardware as tested)	\$1,025,000
Installation Labor Cost (@\$75/hr)	\$600
Annual Cost of Maintenance and Support (hardware/software)	\$28,800
Annual Cost of Updates (IPS/AV/etc.)	\$48,000
Initial Purchase (enterprise management system)	See Comparative
Annual Cost of Maintenance and Support (enterprise management system)	See Comparative
Total Cost of Ownership	
Year 1	\$01,102,400
Year 2	\$76,800
Year 3	\$76,800
3 Year Total Cost of Ownership	\$1,256,000
Total Cost of Ownership (Street Price)	
Ease of Use	
Initial Setup (Hours)	8
Time Required for Upkeep (Hours per Year)	See Comparative
Time Required to Tune (Hours per Year)	See Comparative
Expected Costs	
Initial Purchase (hardware as tested)	\$820,000
Installation Labor Cost (@\$75/hr)	\$600
Annual Cost of Maintenance and Support (hardware/software)	\$23,040
Annual Cost of Updates (IPS/AV/etc.)	\$38,400
Initial Purchase (enterprise management system)	See Comparative
Annual Cost of Maintenance and Support (enterprise management system)	See Comparative
Total Cost of Ownership	
Year 1	\$882,040
Year 2	\$61,440
Year 3	\$61,440
3 Year Total Cost of Ownership	\$1,004,920

Figure 20 – Detailed Scorecard

Test Methodology

Next Generation Firewall (NGFW): v6.0

A copy of the test methodology is available on the NSS Labs website at www.nsslabs.com.

Contact Information

NSS Labs, Inc.
206 Wild Basin Road
Building A, Suite 200
Austin, TX 78746 USA
info@nsslabs.com
www.nsslabs.com

This and other related documents are available at www.nsslabs.com. To receive a licensed copy or report misuse, please contact NSS Labs.

© 2016 NSS Labs, Inc. All rights reserved. No part of this publication may be reproduced, copied/scanned, stored on a retrieval system, e-mailed or otherwise disseminated or transmitted without the express written consent of NSS Labs, Inc. (“us” or “we”).

Please read the disclaimer in this box because it contains important information that binds you. If you do not agree to these conditions, you should not read the rest of this report but should instead return the report immediately to us. “You” or “your” means the person who accesses this report and any entity on whose behalf he/she has obtained this report.

1. The information in this report is subject to change by us without notice, and we disclaim any obligation to update it.
2. The information in this report is believed by us to be accurate and reliable at the time of publication, but is not guaranteed. All use of and reliance on this report are at your sole risk. We are not liable or responsible for any damages, losses, or expenses of any nature whatsoever arising from any error or omission in this report.
3. NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY US. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, ARE HEREBY DISCLAIMED AND EXCLUDED BY US. IN NO EVENT SHALL WE BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, PUNITIVE, EXEMPLARY, 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/or 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 your 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.