



# DATA CENTER INTRUSION PREVENTION SYSTEM TEST REPORT

**Fortinet FortiGate 3000D** v5.4.0, build 7184

Author – Keith Bormann

## Overview

NSS Labs performed an independent test of the Fortinet FortiGate 3000D v5.4.0, build 7184. The product was subjected to thorough testing at the NSS facility in Austin, Texas, based on the Data Center Intrusion Prevention System (DCIPS) Test Methodology v2.0 available at [www.nsslabs.com](http://www.nsslabs.com). This test was conducted free of charge and NSS did not receive any compensation in return for Fortinet’s participation.

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 the majority of enterprises tune their DCIPS products. Therefore, NSS tests DCIPS products that have been optimally tuned by the vendor. Every effort is made to deploy policies that ensure the optimal combination of security effectiveness and performance, as would be the aim of a typical customer deploying the device in a live network environment.

IPS devices deployed within a data center typically are subjected to significantly higher traffic levels than are IPS or next generation firewalls (NGFWs) deployed at the corporate network perimeter. Furthermore, data center traffic mixes are significantly different from network perimeter traffic mixes. Where perimeter devices are expected to protect a wide range of end user applications, a data center device may be deployed to protect a single type of server, supporting far fewer network protocols and applications. Latency is also a concern since applications will be adversely affected if the IPS introduces delays.

Product	NSS Exploit Library Block Rate <sup>1</sup>	NSS-Tested Throughput	3-Year TCO (List Price)	3-Year TCO (Street Price)
<b>Fortinet FortiGate 3000D</b> v5.4.0, build 7184	99.9%	11.042 Gbps	US\$99,975	US\$80,100
	Evasions		Stability and Reliability	
	PASS		PASS	

**Figure 1 – Overall Test Results**

Using a tuned policy, the Fortinet FortiGate 3000D blocked 99.9% of exploits. The device proved effective against all evasion techniques tested. The device also passed all stability and reliability tests.

The Fortinet FortiGate 3000D is rated by NSS at 11.042 Gbps, which is lower than the vendor-claimed performance; Fortinet rates this device at 20 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.

<sup>1</sup> NSS Exploit Block Rate is defined as the number of exploits blocked under test

## Table of Contents

<b>Overview</b>	<b>2</b>
<b>Security Effectiveness</b>	<b>5</b>
NSS Exploit Library	5
<i>False Positive Testing</i>	5
<i>Coverage by Impact Type</i>	5
<i>Coverage by Date</i>	6
<i>Coverage by Target Vendor</i>	6
Resistance to Evasion Techniques	7
<b>Performance</b>	<b>8</b>
Maximum Capacity	8
HTTP Capacity with No Transaction Delays	10
HTTP Capacity with Transaction Delays	10
Application Average Response Time – HTTP	11
Real-World Traffic Mixes	12
Raw Packet Processing Performance (UDP Throughput)	12
Raw Packet Processing Performance (UDP Latency)	13
<b>Stability and Reliability</b>	<b>14</b>
<b>Management and Configuration</b>	<b>15</b>
<b>Total Cost of Ownership (TCO)</b>	<b>16</b>
Installation Hours	16
List Price and Total Cost of Ownership	17
Street Price and Total Cost of Ownership	17
<b>Detailed Product Scorecard</b>	<b>18</b>
<b>Test Methodology</b>	<b>24</b>
<b>Contact Information</b>	<b>24</b>

## Table of Figures

Figure 1 – Overall Test Results.....	2
Figure 2 – Number of Exploits Blocked (%).....	5
Figure 3 – Product Coverage by Date .....	6
Figure 4 – Product Coverage by Target Vendor.....	6
Figure 5 – Resistance to Evasion Results .....	7
Figure 6 – Concurrency and Connection Rates.....	9
Figure 7 – HTTP Capacity with No Transaction Delays .....	10
Figure 8 – HTTP Capacity with Transaction Delays .....	11
Figure 9 –Application Average Response Time (Milliseconds) .....	11
Figure 10 – “Real-World” Traffic Mixes .....	12
Figure 11 – Raw Packet Processing Performance (UDP Traffic) .....	13
Figure 12 – UDP Latency in Microseconds.....	13
Figure 13 – Stability and Reliability Results .....	14
Figure 14 – Sensor Installation Time (Hours).....	16
Figure 15 – List Price 3-Year TCO (US\$) .....	17
Figure 16 – Street Price 3-Year TCO (US\$).....	17
Figure 17 – Detailed Scorecard.....	23

## Security Effectiveness

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

### NSS Exploit Library

NSS’ security effectiveness testing leverages the deep expertise of our engineers who utilize multiple commercial, open-source, and proprietary tools as appropriate. With 896 server 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
<b>Fortinet FortiGate 3000D</b> v5.4.0, build 7184	896	895	99.9%

Figure 2 – Number of Exploits Blocked (%)

### False Positive Testing

The Fortinet FortiGate 3000D 5.4.0 correctly identified traffic and did not fire alerts for non-malicious content.

### 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 3 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 3 – Product Coverage by Date

### Coverage by Target Vendor

Exploits within the *NSS Exploit Library* target a wide range of protocols and applications. Figure 4 depicts the coverage offered by the Fortinet FortiGate 3000D for five of the top vendors targeted in this test. More than 50 vendors are represented in the test. Clients can contact NSS for more information about this test.




Figure 4 – 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 DCIPS 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 5 provides the results of the evasion tests for the Fortinet FortiGate 3000D.

Test Procedure	Result
IP Packet Fragmentation	PASS
Stream Segmentation	PASS
RPC Fragmentation	PASS
URL Obfuscation	PASS
FTP Evasion	PASS
Layered Evasions	
IP Fragmentation + TCP Segmentation	PASS
IP Fragmentation + MSRPC Fragmentation	PASS
TCP Segmentation + SMB / NetBIOS Evasions	PASS

Figure 5 – 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.

In addition, when considering an IPS for the data center rather than for the network perimeter, there are several key metrics that need to be adjusted. Performance metrics, while important in any IPS, become more critical in a device that is intended for data center deployment. In a data center IPS, the volume of traffic will be significantly higher than it would for a device that is intended to protect end user desktops behind the corporate network perimeter. A data center IPS also needs to support much higher data rates, as it handles traffic for potentially hundreds of thousands of users who are accessing large applications in a server farm inside the network perimeter. Connection rate and concurrent connection capacity are additional metrics that become even more critical in a data center IPS.

Traffic mix will differ significantly between a corporate network perimeter and a data center, and this can put additional load on the IPS inspection process. Stateless UDP traffic (such as that seen in a Network File System [NFS]), and long-lived transmission control protocol (TCP) connections (as would be seen in an iSCSI Storage Area Network [SAN] or backup application) are common in many data center networks. These types of applications present a continuous and heavy load to the network.

Within the data center, application traffic puts a very different load on the network than does file system traffic. Communications between users and servers, and communications between applications, database, and directory servers have very different profiles. Application traffic is connection intensive, with connections constantly being set up and torn down. An IPS that includes any form of application awareness capabilities will find significant challenges in data center deployments. Another critical concern is latency, since applications will be adversely affected if the IPS introduces delays.

## 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 DCIPS is causing an unacceptable increase in open connections.
- **Excessive concurrent HTTP connections** – Latency within the DCIPS 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 DCIPS is causing connections to time out.



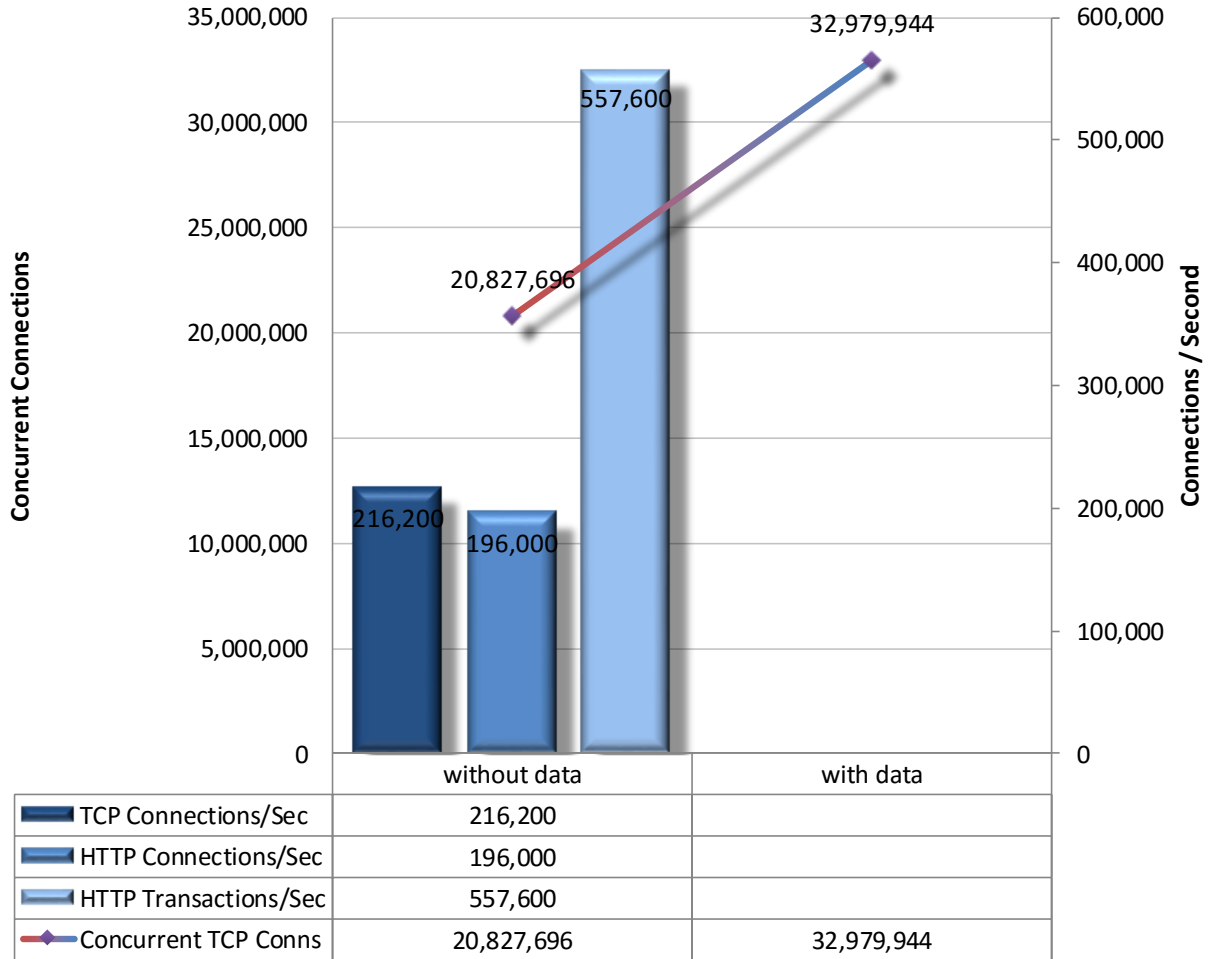


Figure 6 – Concurrency and Connection Rates

## 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 7 – HTTP Capacity with No Transaction Delays

## 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 8 – HTTP Capacity with 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	2.89
5,000 Connections per Second – 21 KB Response	2.80
10,000 Connections per Second – 10 KB Response	2.68
20,000 Connections per Second – 4.5 KB Response	2.35
40,000 Connections per Second – 1.7 KB Response	2.47

Figure 9 –Application Average Response Time (Milliseconds)

## 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 Data Center Intrusion Prevention System Test Methodology, available at [www.nsslabs.com](http://www.nsslabs.com).

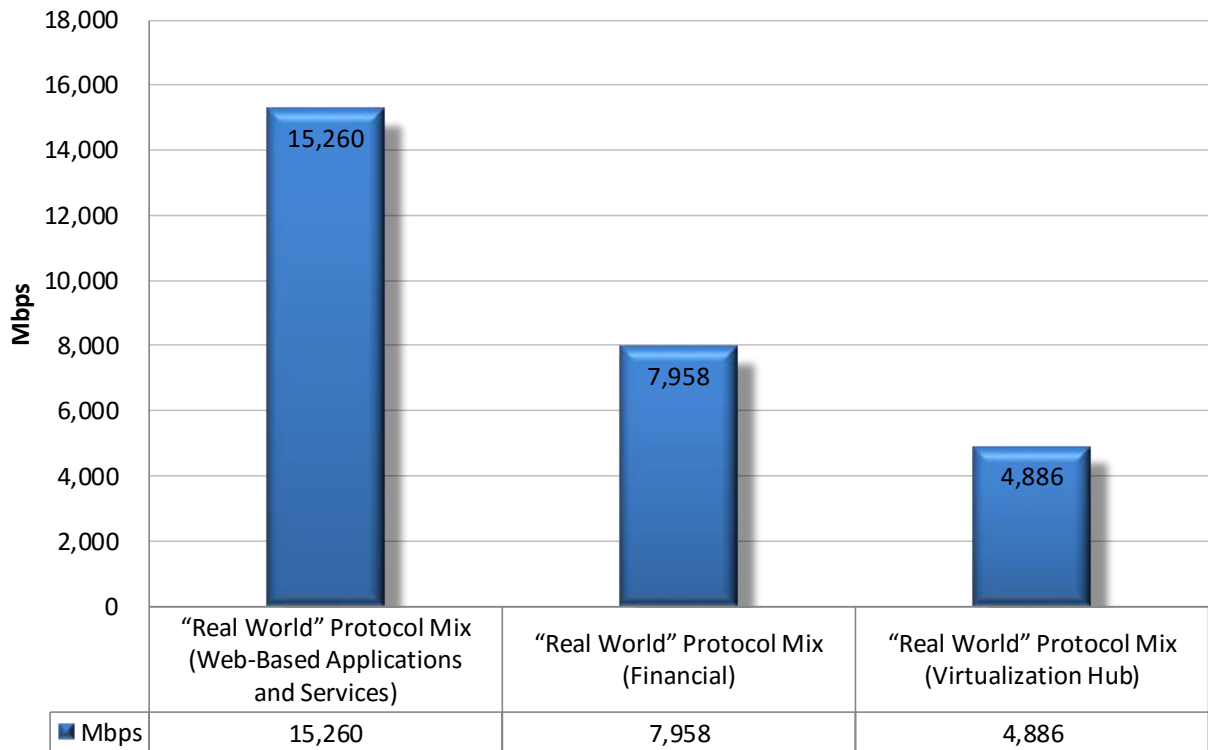


Figure 10 – “Real-World” Traffic Mixes

The Fortinet FortiGate 3000D was tested by NSS and performed below the throughput claimed by the vendor for all “real-world” traffic mixes.

## 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.

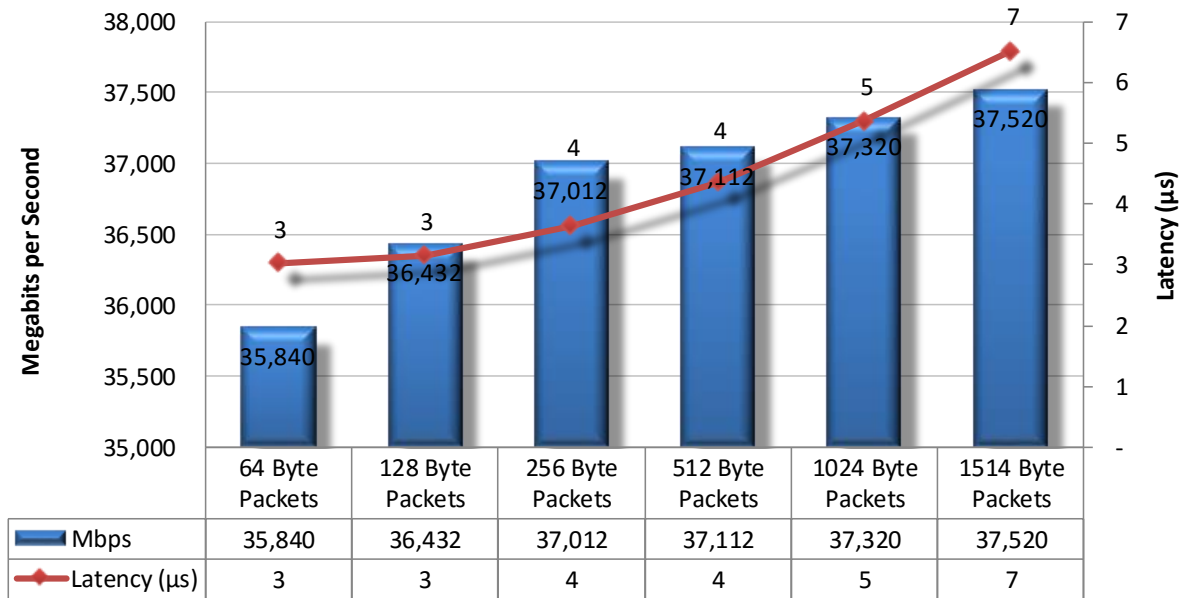


Figure 11 – Raw Packet Processing Performance (UDP Traffic)

### Raw Packet Processing Performance (UDP Latency)

DCIPS 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 12 depicts UDP latency (in microseconds) as recorded during the UDP throughput tests at 90% of maximum load.

Latency – UDP	Microseconds
64-Byte Packets	3
128-Byte Packets	3
256-Byte Packets	4
512-Byte Packets	4
1024-Byte Packets	5
1514-Byte Packets	7

Figure 12 – UDP Latency in Microseconds

## 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	
• State Preservation – Normal Load	PASS
• State Preservation – Maximum Exceeded	PASS
Protocol Fuzzing and Mutation	PASS
Power Fail	PASS
Persistence of Data	PASS

**Figure 13 – Stability and Reliability Results**

These tests also determine the behavior of the state engine under load. All DCIPS devices must choose whether to risk denying legitimate traffic or risk allowing malicious traffic once they run low on resources. A DCIPS 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 DCIPS 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)
<b>Fortinet FortiGate 3000D</b> v5.4.0, build 7184	8

Figure 14 – 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
<b>Fortinet FortiGate 3000D</b> v5.4.0, build 7184	\$60,000	\$13,125	\$73,725	\$13,125	\$13,125	\$99,975

Figure 15 – List Price 3-Year TCO (US\$)

- **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
<b>Fortinet FortiGate 3000D</b> v5.4.0, build 7184	\$48,000	\$10,500	\$59,100	\$10,500	\$10,500	\$80,100

Figure 16 – Street Price 3-Year TCO (US\$)

- **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
<b>Security Effectiveness</b>	
<b>Exploit Library</b>	
<i>NSS Exploit Library</i> Block Rate	99.9%
False Positive Testing	PASS
<b>Coverage by Impact Type</b>	
System Exposure	Contact NSS
Service Exposure	Contact NSS
System or Service Fault	Contact NSS
<b>Coverage by Date</b>	Contact NSS
Coverage by Target Vendor	Contact NSS
Coverage by Result	Contact NSS
Coverage by Target Type	Contact NSS
<b>Evasions and Attack Leakage</b>	
Resistance to Evasion	PASS
<b>IP Packet Fragmentation</b>	100%
Ordered 8 byte fragments	100%
Ordered 16 byte fragments	100%
Ordered 24 byte fragments	100%
Ordered 32 byte fragments	100%
Out of order 8 byte fragments	100%
Ordered 8 byte fragments, duplicate last packet	100%
Out of order 8 byte fragments, duplicate last packet	100%
Ordered 8 byte fragments, reorder fragments in reverse	100%
Ordered 16 byte fragments, fragment overlap (favor new)	100%
Ordered 16 byte fragments, fragment overlap (favor old)	100%
Out of order 8 byte fragments, interleaved duplicate packets scheduled for later delivery	100%
Ordered 8 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
Ordered 16 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
Ordered 24 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
Ordered 32 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
<b>TCP Stream Segmentation</b>	<b>100%</b>
Ordered 1 byte segments, interleaved duplicate segments with invalid TCP checksums	100%
Ordered 1 byte segments, interleaved duplicate segments with null TCP control flags	100%
Ordered 1 byte segments, interleaved duplicate segments with requests to resync sequence numbers mid-stream	100%
Ordered 1 byte segments, duplicate last packet	100%
Ordered 2 byte segments, segment overlap (favor new)	100%
Ordered 1 byte segments, interleaved duplicate segments with out-of-window sequence numbers	100%
Out of order 1 byte segments	100%
Out of order 1 byte segments, interleaved duplicate segments with faked retransmits	100%
Ordered 1 byte segments, segment overlap (favor new)	100%
Out of order 1 byte segments, PAWS elimination (interleaved duplicate segments with older TCP timestamp options)	100%
Ordered 16 byte segments, segment overlap (favor new (Unix))	100%
Ordered 32 byte segments	100%
Ordered 64 byte segments	100%
Ordered 128 byte segments	100%

Ordered 256 byte segments	100%
Ordered 512 byte segments	100%
Ordered 1024 byte segments	100%
Ordered 2048 byte segments (sending MSRPC request with exploit)	100%
Reverse Ordered 256 byte segments, segment overlap (favor new) with random data	100%
Reverse Ordered 512 byte segments, segment overlap (favor new) with random data	100%
Reverse Ordered 1024 byte segments, segment overlap (favor new) with random data	100%
Reverse Ordered 2048 byte segments, segment overlap (favor new) with random data	100%
Out of order 1024 byte segments, segment overlap (favor new) with random data, Initial TCP sequence number is set to 0xffffffff – 4294967295	100%
Out of order 2048 byte segments, segment overlap (favor new) with random data, Initial TCP sequence number is set to 0xffffffff – 4294967295	100%
RPC Fragmentation	100%
One-byte fragmentation (ONC)	100%
Two-byte fragmentation (ONC)	100%
All fragments, including Last Fragment (LF) will be sent in one TCP segment (ONC)	100%
All frags except Last Fragment (LF) will be sent in one TCP segment. LF will be sent in separate TCP seg (ONC)	100%
One RPC fragment will be sent per TCP segment (ONC)	100%
One LF split over more than one TCP segment. In this case no RPC fragmentation is performed (ONC)	100%
Canvas Reference Implementation Level 1 (MS)	100%
Canvas Reference Implementation Level 2 (MS)	100%
Canvas Reference Implementation Level 3 (MS)	100%
Canvas Reference Implementation Level 4 (MS)	100%
Canvas Reference Implementation Level 5 (MS)	100%
Canvas Reference Implementation Level 6 (MS)	100%
Canvas Reference Implementation Level 7 (MS)	100%
Canvas Reference Implementation Level 8 (MS)	100%
Canvas Reference Implementation Level 9 (MS)	100%
Canvas Reference Implementation Level 10 (MS)	100%
MSRPC messages are sent in the big endian byte order, 16 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
MSRPC messages are sent in the big endian byte order, 32 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
MSRPC messages are sent in the big endian byte order, 64 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
MSRPC messages are sent in the big endian byte order, 128 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
MSRPC messages are sent in the big endian byte order, 256 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
MSRPC messages are sent in the big endian byte order, 512 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
MSRPC messages are sent in the big endian byte order, 1024 MSRPC fragments are sent in the same lower layer message, MSRPC requests are fragmented to contain at most 2048 bytes of payload	100%
URL Obfuscation	100%
URL encoding – Level 1 (minimal)	100%
URL encoding – Level 2	100%
URL encoding – Level 3	100%
URL encoding – Level 4	100%
URL encoding – Level 5	100%
URL encoding – Level 6	100%
URL encoding – Level 7	100%
URL encoding – Level 8 (extreme)	100%
Directory Insertion	100%
Premature URL ending	100%
Long URL	100%

Fake parameter	100%
TAB separation	100%
Case sensitivity	100%
Windows \ delimiter	100%
Session splicing	100%
FTP Evasion	100%
Inserting spaces in FTP command lines	100%
Inserting non-text Telnet opcodes – Level 1 (minimal)	100%
Inserting non-text Telnet opcodes – Level 2	100%
Inserting non-text Telnet opcodes – Level 3	100%
Inserting non-text Telnet opcodes – Level 4	100%
Inserting non-text Telnet opcodes – Level 5	100%
Inserting non-text Telnet opcodes – Level 6	100%
Inserting non-text Telnet opcodes – Level 7	100%
Inserting non-text Telnet opcodes – Level 8 (extreme)	100%
Layered Evasions	
IP Fragmentation + TCP Segmentation	100%
Ordered 8 byte fragments + Ordered TCP segments except that the last segment comes first	100%
Ordered 24 byte fragments + Ordered TCP segments except that the last segment comes first	100%
Ordered 32 byte fragments + Ordered TCP segments except that the last segment comes first	100%
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	100%
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	100%
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	100%
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	100%
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	100%
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	100%
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	100%
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	100%
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	100%
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	100%
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	100%
IP Fragmentation + MSRPC Fragmentation	100%
Ordered 8 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a shuffled payload + MSRPC messages are sent in the big endian byte order with	100%

8 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 2048 bytes of payload.	
Ordered 16 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a shuffled payload + MSRPC messages are sent in the big endian byte order with 16 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 2048 bytes of payload.	100%
Ordered 32 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a shuffled payload + MSRPC messages are sent in the big endian byte order with 32 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 64 bytes of payload.	100%
Ordered 64 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a shuffled payload + MSRPC messages are sent in the big endian byte order with 64 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 64 bytes of payload.	100%
Ordered 128 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + MSRPC messages are sent in the big endian byte order with 1024 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 128 bytes of payload.	100%
Ordered 256 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + MSRPC messages are sent in the big endian byte order with 1024 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 256 bytes of payload.	100%
Ordered 512 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + MSRPC messages are sent in the big endian byte order with 1024 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 512 bytes of payload.	100%
Ordered 1024 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + MSRPC messages are sent in the big endian byte order with 1024 MSRPC fragments sent in the same lower layer message. MSRPC requests are fragmented to contain at most 1024 bytes of payload.	100%
<b>IP Fragmentation + SMB Evasions</b>	100%
Ordered 1024 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + SMB chaff message before real messages. The chaff is a WriteAndX message with a broken write mode flag, and has random MSRPC request-like payload	100%
Ordered 8 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + A chaffed NetBIOS message is sent before the first actual NetBIOS message. The chaff message is an unspecified NetBIOS message with MSRPC request like payload	100%
Ordered 8 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has a random payload + A chaffed NetBIOS message is sent before the first actual NetBIOS message. The chaff message is an unspecified NetBIOS message with HTTP GET request like payload	100%
<b>TCP Segmentation + SMB / NetBIOS Evasions</b>	100%
Reverse Ordered 2048 byte TCP segments, segment overlap (favor new) with random data + A chaffed NetBIOS message is sent before the first actual NetBIOS message. The chaff message is an unspecified NetBIOS message with MSRPC request like payload	100%
<b>Performance</b>	
<b>Vendor Claimed Performance</b>	
<b>Raw Packet Processing Performance (UDP Traffic)</b>	<b>Mbps</b>
64 Byte Packets	35,840
128 Byte Packets	36,432
256 Byte Packets	37,012
512 Byte Packets	37,112
1024 Byte Packets	37,320
1514 Byte Packets	37,520
<b>Latency – UDP</b>	<b>Microseconds</b>
64 Byte Packets	3

128 Byte Packets	3
256 Byte Packets	4
512 Byte Packets	4
1024 Byte Packets	5
1514 Byte Packets	7
<b>Maximum Capacity</b>	
Theoretical Max. Concurrent TCP Connections	20,827,696
Theoretical Max. Concurrent TCP Connections w/Data	32,979,944
Maximum TCP Connections per Second	216,200
Maximum HTTP Connections per Second	196,000
Maximum HTTP Transactions per Second	557,600
<b>HTTP Capacity With No Transaction Delays</b>	
25,000 Connections per Second – 44 KB Response	48,000
50,000 Connections per Second – 21 KB Response	80,320
100,000 Connections per Second – 10 KB Response	116,520
200,000 Connections per Second – 4.5 KB Response	143,200
400,000 Connections per Second – 1.7 KB Response	160,000
<b>Application Average Response Time – HTTP (at 90% Max Load)</b>	<b>Milliseconds</b>
25,000 Connections per Second – 44 KB Response	2.8939
50,000 Connections per Second – 21 KB Response	2.8017
100,000 Connections per Second – 10K KB Response	2.6800
200,000 Connections per Second – 4.5 KB Response	2.3460
400,000 Connections per Second – 1.7 KB Response	2.4680
<b>HTTP CPS &amp; Capacity With Transaction Delays</b>	
21 KB Response With Delay	80,320
10 KB Response With Delay	116,520
<b>“Real World” Traffic</b>	<b>Mbps</b>
“Real-World” Protocol Mix (Web-Based Applications and Services)	15,260
“Real-World” Protocol Mix (Financial)	7,958
“Real-World” Protocol Mix (Virtualization Hub)	4,886
<b>Stability &amp; Reliability</b>	
Blocking Under Extended Attack	PASS
Passing Legitimate Traffic Under Extended Attack	PASS
<b>Behavior Of The State Engine Under Load</b>	<b>PASS</b>
State Preservation – Normal Load	PASS
State Preservation – Maximum Exceeded	PASS
Protocol Fuzzing & Mutation	PASS
Power Fail	PASS
Persistence of Data	PASS
<b>Total Cost of Ownership (List Price)</b>	
<b>Ease of Use</b>	
Initial Setup (Hours)	8
Time Required for Upkeep (Hours per Year)	See Comparative
Time Required to Tune (Hours per Year)	See Comparative
<b>Expected Costs</b>	
Initial Purchase (hardware as tested)	\$60,000
Installation Labor Cost (@\$75/hr)	\$600
Annual Cost of Maintenance & Support (hardware/software)	\$13,125
Annual Cost of Updates (IPS/AV/etc.)	\$0
Initial Purchase (enterprise management system)	See Comparative
Annual Cost of Maintenance & Support (enterprise management system)	See Comparative
<b>Total Cost of Ownership</b>	
Year 1	\$73,725
Year 2	\$13,125
Year 3	\$13,125
3-Year Total Cost of Ownership	\$99,975

<b>Total Cost of Ownership (Street Price)</b>	
<b>Ease of Use</b>	
Initial Setup (Hours)	8
Time Required for Upkeep (Hours per Year)	See Comparative
Time Required to Tune (Hours per Year)	See Comparative
<b>Expected Costs</b>	
Initial Purchase (hardware as tested)	\$48,000
Installation Labor Cost (@\$75/hr)	\$600
Annual Cost of Maintenance & Support (hardware/software)	\$10,500
Annual Cost of Updates (IPS/AV/etc.)	\$0
Initial Purchase (enterprise management system)	See Comparative
Annual Cost of Maintenance & Support (enterprise management system)	See Comparative
<b>Total Cost of Ownership</b>	
Year 1	\$59,100
Year 2	\$10,500
Year 3	\$10,500
3-Year Total Cost of Ownership	\$80,100

Figure 17 – Detailed Scorecard

## Test Methodology

Data Center Intrusion Prevention System (DCIPS) v2.0

A copy of the test methodology is available on the NSS Labs website at [www.nsslabs.com](http://www.nsslabs.com).

## Contact Information

NSS Labs, Inc.  
206 Wild Basin Road  
Building A, Suite 200  
Austin, TX 78746 USA  
[info@nsslabs.com](mailto:info@nsslabs.com)  
[www.nsslabs.com](http://www.nsslabs.com)

This and other related documents are available at [www.nsslabs.com](http://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.