DATA CENTER FIREWALL TEST REPORT

Fortinet FortiGate 3700D FortiOS v5.4.1 GA Build 7386

APRIL 18, 2017

Author – Keith Bormann

Overview

NSS Labs performed an independent test of the Fortinet FortiGate 3700D FortiOS v5.4.1 GA Build 7386 product. The product was subjected to thorough testing at the NSS facility in Austin, Texas, based on the Data Center Firewall (DCFW) Test Methodology v2.2, which is available at www.nsslabs.com. This test was conducted free of charge and NSS did not receive any compensation in return for Fortinet's inclusion.

This report provides detailed information about this product and its security effectiveness. Additional comparative information is available at www.nsslabs.com.

Firewall devices deployed within a data center typically will be subjected to significantly higher traffic levels than a firewall or next generation firewall (NGFW) deployed at the corporate network perimeter. Furthermore, data center traffic mixes will be completely different from a typical corporate network perimeter; where perimeter devices will be 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. Figure 1 presents the overall results of the tests.

Product	Test	IPv4	IPv6
Fortinet FortiGate 3700D FortiOS v5.4.1 GA Build 7386	NSS-Tested Throughput	123,105 Mbps	127,488 Mbps
	Firewall Policy Enforcement	PASS	PASS
	Stability & Reliability	PASS	PASS

Figure 1 – Overall Test Results

The device passed all stability and reliability tests for both IPv4 and IPv6. The device also passed all firewall policy enforcement tests for both IPv4 and IPv6.

The Fortinet FortiGate 3700D is rated by NSS at 123.1 Gbps for IPv4. Fortinet rates the bidirectional throughput of this device at 160 Gbps (UDP).

The Fortinet FortiGate 3700D is rated by NSS at 127.5 Gbps for IPv6. Fortinet rates the bidirectional throughput of this device at 160 Gbps (UDP).

NSS-Tested Throughput is calculated as an average of all of the unidirectional "real-world" protocol mixes and the unidirectional 21 KB HTTP response-based capacity tests. Unidirectional tests limit the throughput of the device to 10 Gbps per port pair. Thus, an 80 Gbps device with eight 10 Gbps ports will be limited to 40 Gbps.

Table of Contents

Table of Figures

Security Effectiveness

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

Firewall Policy Enforcement

Policies are rules that are 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:

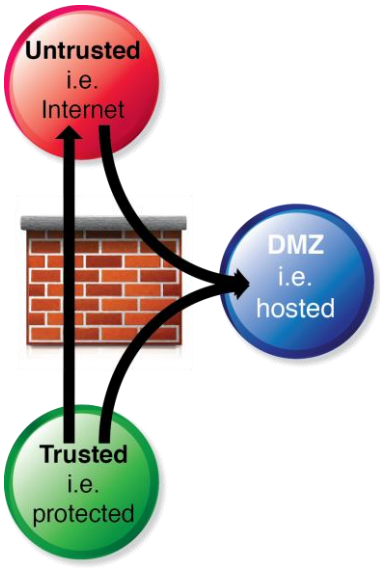

- **Untrusted –** This is typically an external network and is considered to be unknown and not secure. An example of an untrusted network would be the Internet.
- **DMZ –** This is a network that is being isolated by the firewall restricting network traffic to and from hosts contained within the isolated network.
- **Trusted –** This is typically an internal network; i.e., 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.

Figure 2 and Figure 3 depict the results for the IPv4 and IPv6 tests. If no IPv6 test was performed, results will be marked as "Not Tested."

Test Procedure	Result
Baseline Policy	PASS
Simple Policies	PASS
Complex Policies	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 (IPv4)

Test Procedure	Result
Baseline Policy	PASS
Simple Policies	PASS
Complex Policies	PASS
Static NAT	Not Tested
Dynamic/Hide NAT	Not Tested
SYN Flood Protection	Not Tested
IP Address Spoofing Protection	Not Tested
TCP Split Handshake Spoof	Not Tested

Figure 3 – Firewall Policy Enforcement (IPv6)

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.

This section depicts the performance of the firewall using various traffic conditions that provide metrics for real-world performance. Individual implementations will vary based on usage; however, these quantitative metrics provide a gauge as to whether a particular device is appropriate for a given environment. Both IPv6 and IPv4 networks can be utilized for performance testing.

Raw Packet Processing Performance (UDP Throughput)

This test uses UDP packets of varying sizes generated by traffic generation tools. 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 bi-directionally through each port pair of the device.

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.

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 device and to determine its effectiveness at forwarding packets quickly in order to provide the highest level of network performance and the lowest latency. Figure 4 depicts the results of the UDP traffic test.

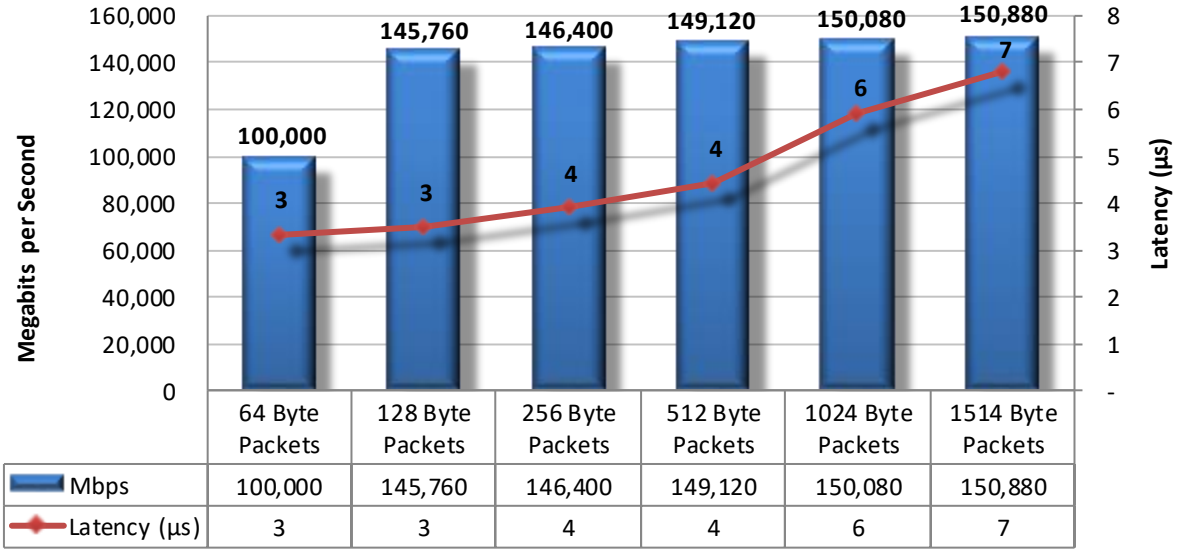

Figure 4 – Raw Packet Processing Performance (UDP Traffic)

Raw Packet Processing Performance (UDP Latency)

DCFWs 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 5 depicts UDP latency (in microseconds) as recorded during the UDP throughput tests at 90% of maximum load.

Latency – UDP	Microseconds
64-Byte Packets	3.3
128-Byte Packets	3.5
256-Byte Packets	3.9
512-Byte Packets	4.4
1024-Byte Packets	5.9
1514-Byte Packets	6.8

Figure 5 – UDP Latency in Microseconds

Maximum Capacity

The use of traffic generation equipment allows NSS engineers to create true "real-world" traffic at multi-gigabit speeds as a background load for the tests.

The purpose of these tests is to stress the inspection engine and determine how it handles 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 DCFW is causing an unacceptable increase in open connections.
- **Excessive response time for HTTP connections** – Latency within the DCFW 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 DCFW is causing connections to time out.

Figure 6 and Figure 7 depict the results of the maximum capacity tests for IPv4 and IPv6.

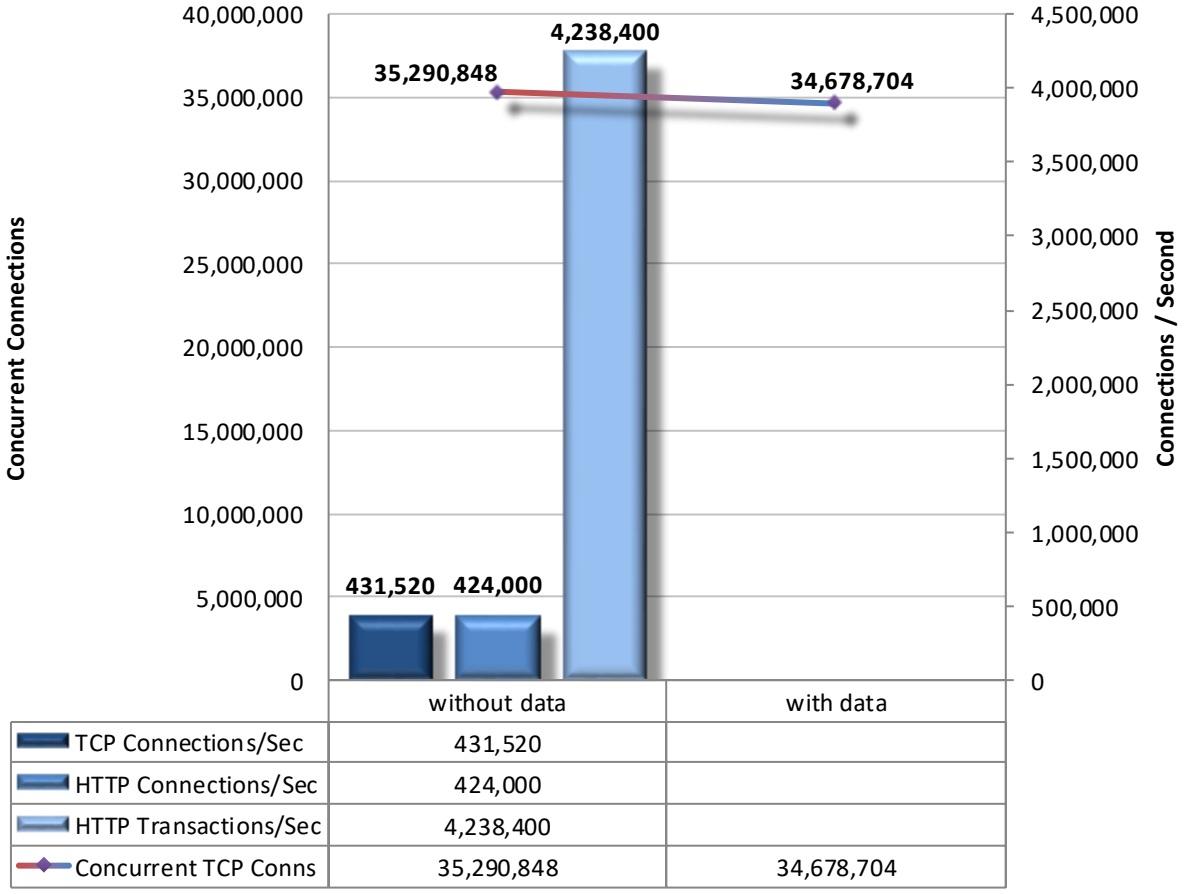

Figure 6 – Concurrency and Connection Rates (IPv4)

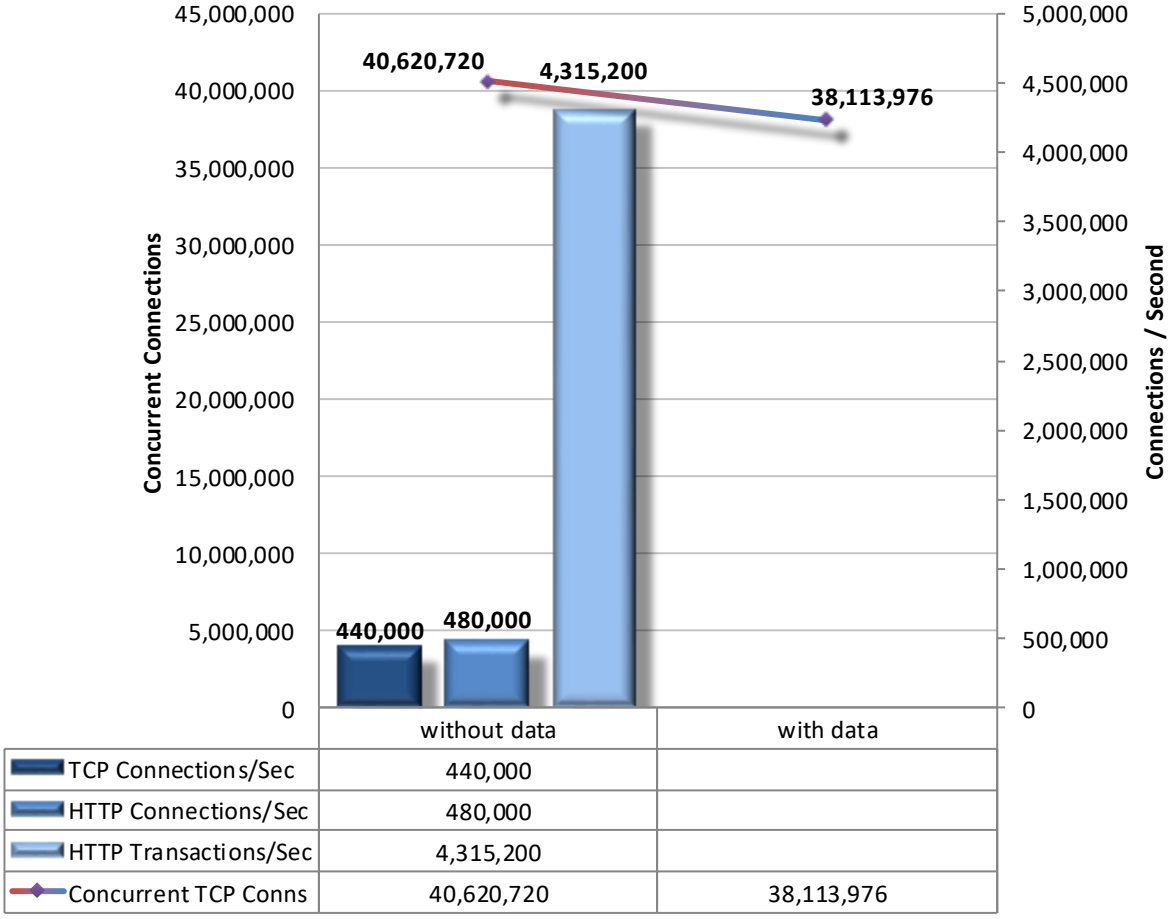

Figure 7 – Concurrency and Connection Rates (IPv6)

HTTP Capacity with No Transaction Delays

The aim of these tests is to stress the HTTP detection engine and determine how the device 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 device 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 simulates real-world HTTP transactions in the lab, 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, and this test provides an excellent representation of a live network (albeit one biased towards HTTP traffic) at various network loads.

Figure 8 and Figure 9 depict the results of the IPv4 and IPv6 tests for the HTTP capacity with no transaction delays tests.

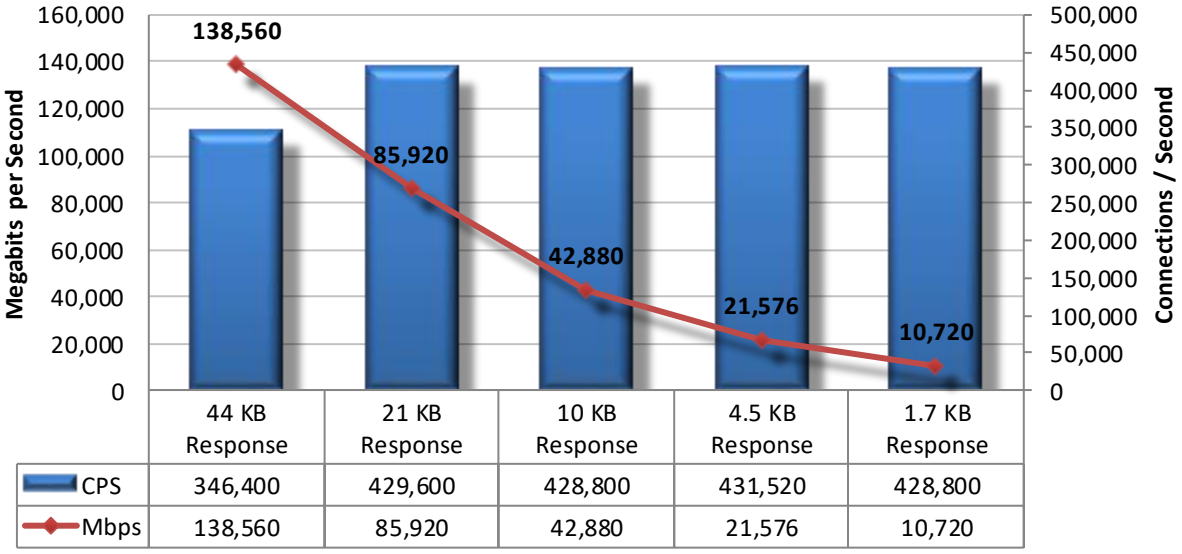

	44 KB Response	21 KB Response	10 KB Response	4.5 KB Response	1.7 KB Response
CPS	346,400	429,600	428,800	431,520	428,800
Mbps	138,560	85,920	42,880	21,576	10,720

Figure 8 – HTTP Capacity with No Transaction Delays (IPv4)

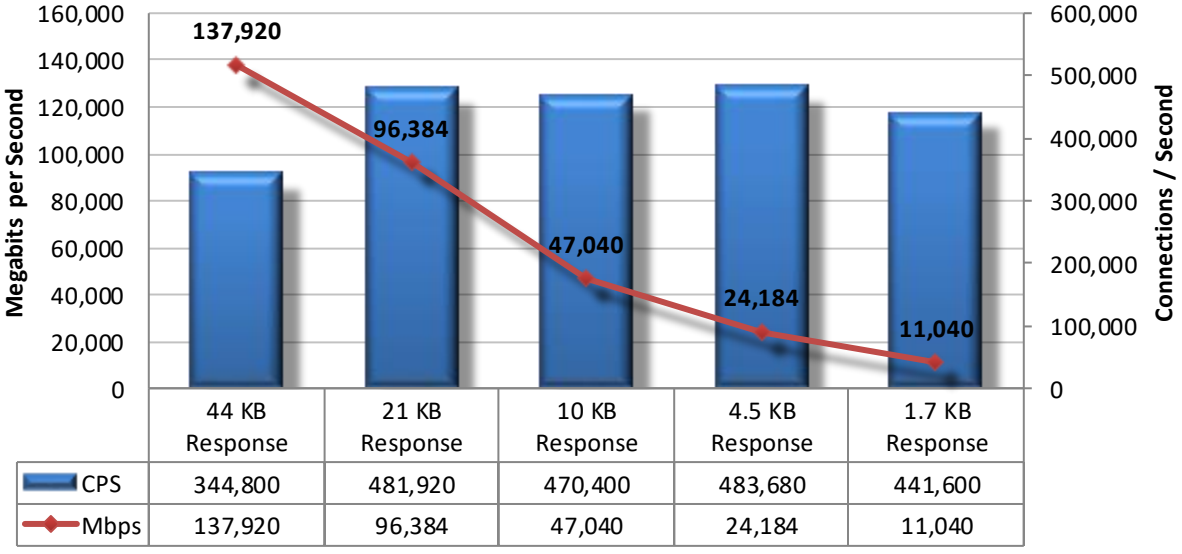

	44 KB Response	21 KB Response	10 KB Response	4.5 KB Response	1.7 KB Response
CPS	344,800	481,920	470,400	483,680	441,600
Mbps	137,920	96,384	47,040	24,184	11,040

Figure 9 – HTTP Capacity with No Transaction Delays (IPv6)

Application Average Response Time – HTTP

Test traffic is passed across the infrastructure switches and through all inline port pairs of the device simultaneously (the latency of the basic infrastructure is known and is constant throughout the tests). The results are recorded at each response size (44 KB, 21 KB, 10 KB, 4.5 KB, and 1.7 KB HTTP responses), at a load level of 90% of the maximum throughput with zero packet loss as previously determined in the HTTP capacity with no transaction delays test.

Figure 10 and Figure 11 depict the results from the IPv4 and IPv6 tests for average application response time.

Application Average Response Time – HTTP (at 90% Maximum Load)	Milliseconds
2,500 Connections per Second – 44 KB Response	0.391
5,000 Connections per Second – 21 KB Response	0.077
10,000 Connections per Second – 10 KB Response	0.004
20,000 Connections per Second – 4.5 KB Response	0.021
40,000 Connections per Second – 1.7 KB Response	0.015

Figure 10 – Average Application Response Time for IPv4 (Milliseconds)

Application Average Response Time – HTTP (at 90% Maximum Load)	Milliseconds
2,500 Connections per Second – 44 KB Response	0.284
5,000 Connections per Second – 21 KB Response	0.088
10,000 Connections per Second – 10 KB Response	0.042
20,000 Connections per Second – 4.5 KB Response	0.027
40,000 Connections per Second – 1.7 KB Response	0.026

Figure 11 – Average Application Response Time for IPv6 (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 device to utilize additional resources to track those connections.

Figure 12 and Figure 13 depict the results for the IPv4 and IPv6 tests for HTTP capacity with transaction delays.

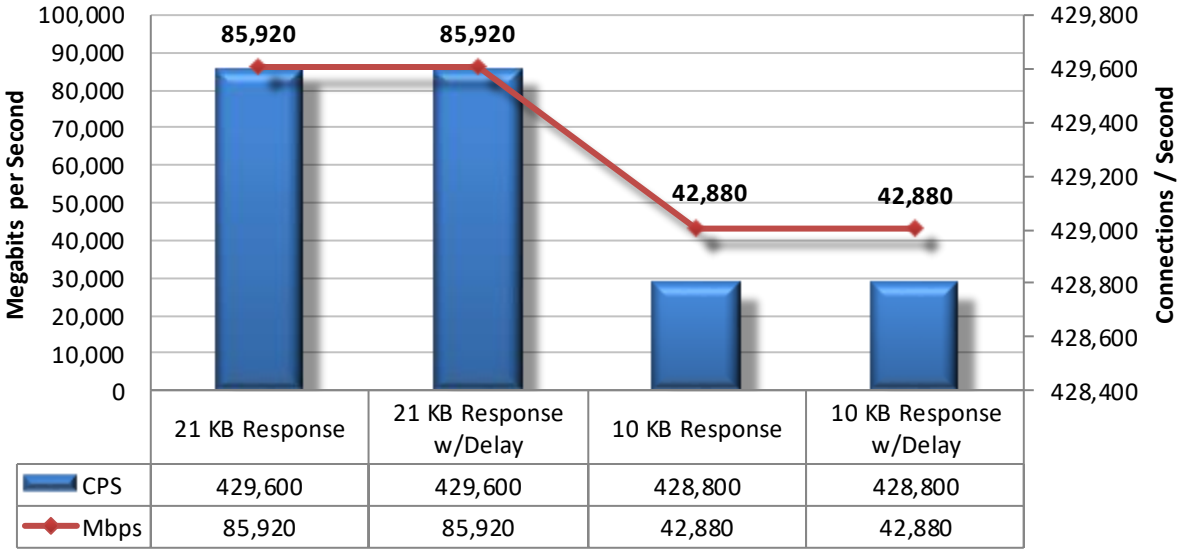

Figure 12 – HTTP Capacity with Transaction Delays (IPv4)

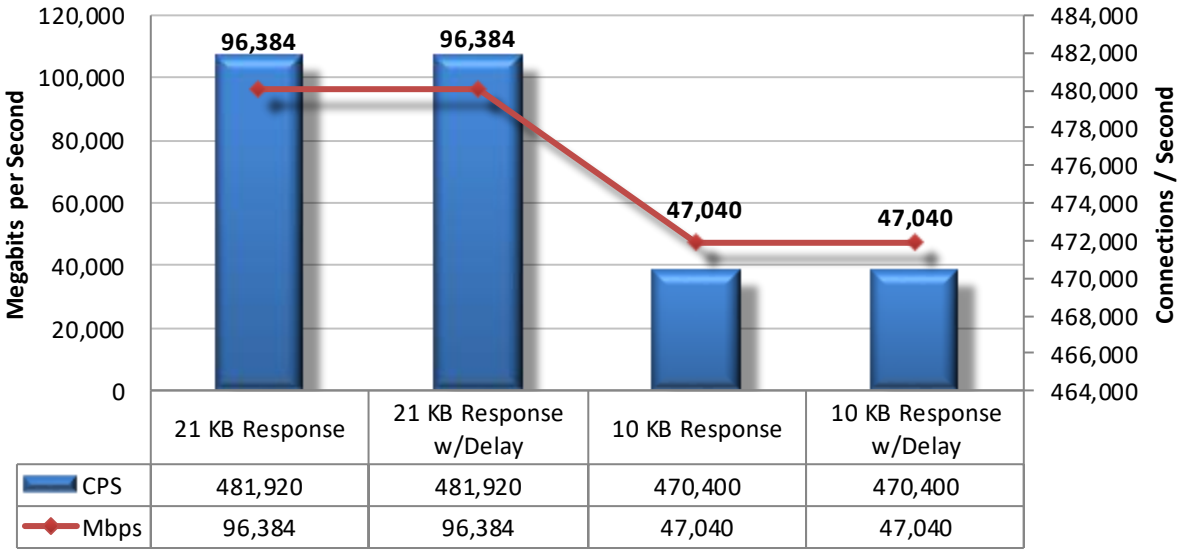

Figure 13 – HTTP Capacity with Transaction Delays (IPv6)

Real-World Traffic Mixes

This test measures the performance of the device in a "real-world" environment by introducing additional protocols and real content, while still maintaining a precisely repeatable and consistent background traffic load. For details about real-world traffic protocol types and percentages, see the NSS Labs Data Center Firewall Test Methodology, available at www.nsslabs.com.

Figure 14 and Figure 15 depict the results for the IPv4 and IPv6 tests for real-world traffic mixes.

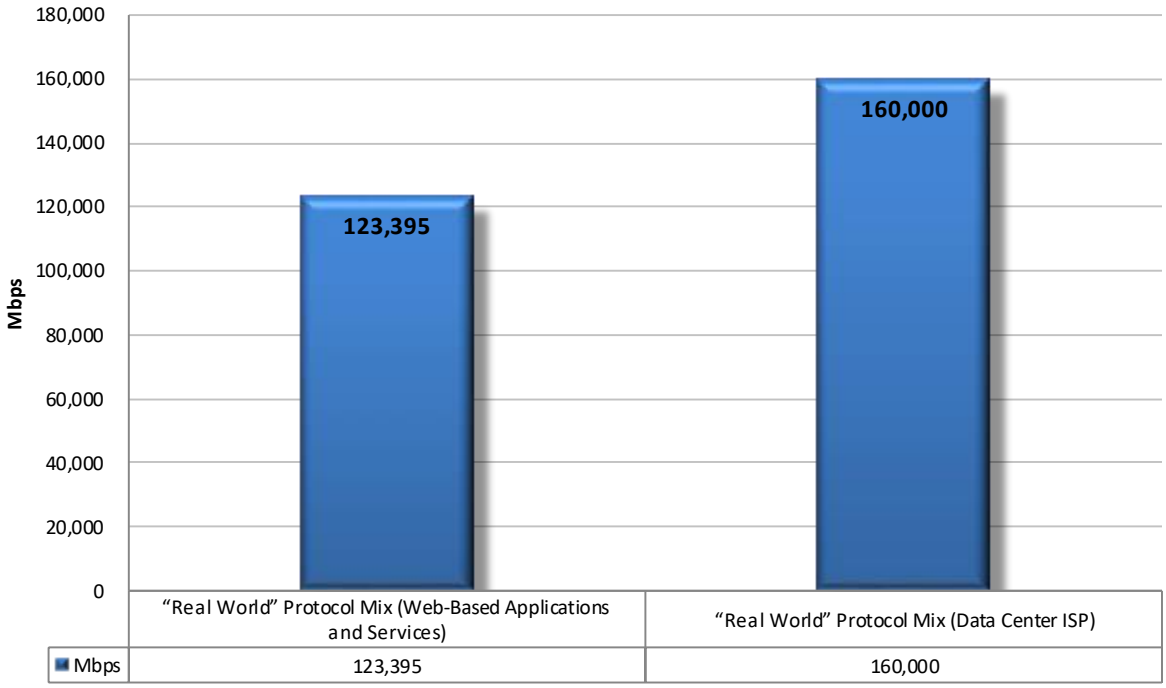

Figure 14 – Real-World Traffic Mixes (IPv4)

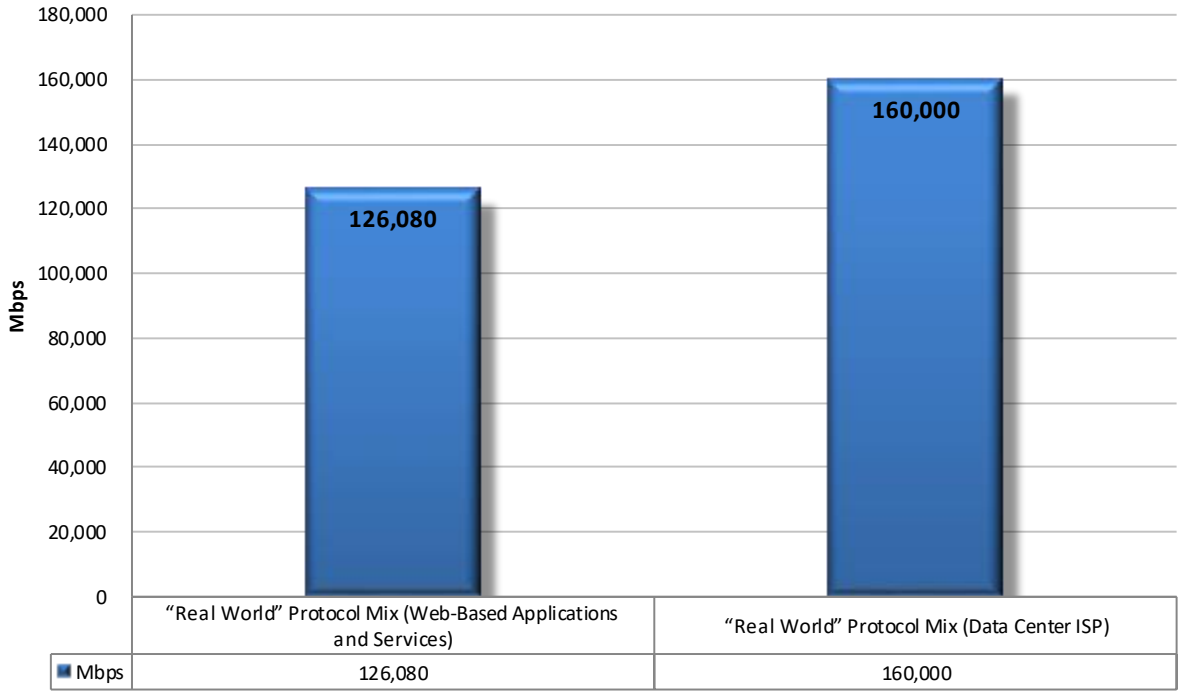

Figure 15 – Real-World Traffic Mixes (IPv6)

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 device 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 device failing open for any reason, the device will fail the test. Figure 16 and Figure 17 depict the results of the IPv4 and IPv6 tests for stability and reliability.

Stability and Reliability	Result
Blocking under Extended Attack	PASS
Passing Legitimate Traffic under Extended Attack	PASS
Protocol Fuzzing and Mutation	PASS
Power Fail	PASS
Persistence of Data	PASS

Figure 16 – Stability and Reliability Results (IPv4)

Stability and Reliability	Result
Blocking under Extended Attack	Not Tested
Passing Legitimate Traffic under Extended Attack	Not Tested
Protocol Fuzzing and Mutation	Not Tested
Power Fail	PASS
Persistence of Data	PASS

Figure 17 – Stability and Reliability Results (IPv6)

High Availability (HA)

High availability (HA) is important to many enterprise customers. This table represents the vendor's HA feature set. If no HA offering was submitted for NSS to validate, all results in this section will be marked as "Not Tested." Figure 18 depicts the results for the high availability tests for IPv4.

Description	Results
Failover – Legitimate Traffic	PASS
Stateful Operation	PASS
Active/Passive Configuration	PASS

Figure 18 – High Availability Results (IPv4)

Total Cost of Ownership (TCO)

Organizations should be concerned with the ongoing amortized cost of operating security products. This section evaluates the costs associated with the purchase, installation, and ongoing management of the device, including:

- **Product Purchase** – The cost of acquisition
- **Product Maintenance** – The fees paid to the vendor (including software and hardware support, maintenance, and updates)
- **Installation** – The time required to take the device out of the box, configure it, deploy it into the network, apply updates and patches, perform 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 firmware updates

For TCO analysis, refer to the TCO Comparative Report, which is available at www.nsslabs.com.

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 device 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)
Fortinet FortiGate 3700D FortiOS v5.4.1 GA Build 7386	8

Figure 19 – Sensor Installation Time (Hours)

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 Price	Maintenance /Year	Year 1 Cost	Year 2 Cost	Year 3 Cost	3-Year TCO
Fortinet FortiGate 3700D FortiOS v5.4.1 GA Build 7386	$80,000	$17,500	$98,100	$17,500	$17,500	$133,100

Figure 20 – 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.

Appendix A: Product Scorecard

Fortinet FortiGate 3700D	IPv4	IPv6
Security Effectiveness		
Firewall Policy Enforcement		
Baseline Policy	PASS	PASS
Simple Policies	PASS	PASS
Complex Policies	PASS	PASS
Static NAT (Network Address Translation)	PASS	Not Tested
Dynamic/Hide NAT	PASS	Not Tested
SYN Flood	PASS	Not Tested
IP Address Spoofing	PASS	Not Tested
TCP Split Handshake	PASS	Not Tested
Performance		
Raw Packet Processing Performance (UDP Traffic)	Mbps	Mbps
64 Byte Packets	100,000	Not Tested
128 Byte Packets	145,760	Not Tested
256 Byte Packets	146,400	Not Tested
512 Byte Packets	149,120	Not Tested
1024 Byte Packets	150,080	Not Tested
1514 Byte Packets	150,880	Not Tested
Latency – UDP	Microseconds	Microseconds
64 Byte Packets	3.3	Not Tested
128 Byte Packets	3.5	Not Tested
256 Byte Packets	3.9	Not Tested
512 Byte Packets	4.4	Not Tested
1024 Byte Packets	5.9	Not Tested
1514 Byte Packets	6.8	Not Tested
Maximum Capacity		
Theoretical Max. Concurrent TCP Connections	35,290,848	40,620,720
Theoretical Max. Concurrent TCP Connections w/Data	34,678,704	38,113,976
Maximum TCP Connections per Second	431,520	440,000
Maximum HTTP Connections per Second	424,000	480,000
Maximum HTTP Transactions per Second	4,238,400	4,315,200
HTTP Capacity with No Transaction Delays		
25,000 Connections per Second – 44 Kbyte Response	346,400	344,800
50,000 Connections per Second – 21 Kbyte Response	429,600	481,920
100,000 Connections per Second – 10 Kbyte Response	428,800	470,400
200,000 Connections per Second – 4.5 Kbyte Response	431,520	483,680
400,000 Connections per Second – 1.7 Kbyte Response	428,800	441,600
Application Average Response Time – HTTP (at 90% Max Load)	Milliseconds	Milliseconds
25,000 Connections per Second – 44 Kbyte Response	0.391	0.284
50,000 Connections per Second – 21 Kbyte Response	0.077	0.088
100,000 Connections per Second – 10 Kbyte Response	0.004	0.042
200,000 Connections per Second – 4.5 Kbyte Response	0.021	0.027

400,000 Connections pPer Second – 1.7 Kbyte Response	0.015	0.026
HTTP CPS & Capacity with Transaction Delays		
21 Kbyte Response with Delay	429,600	481,920
10 Kbyte Response with Delay	428,800	470,400
"Real-World" Traffic	Mbps	Mbps
"Real-World"Protocol Mix (Web-Based Applications and Services)	123,395	126,080
"Real-World" Protocol Mix (Data Center ISP)	160,000	160,000
Stability & Reliability		
Blocking Under Extended Attack	PASS	Not Tested
Passing Legitimate Traffic Under Extended Attack	PASS	Not Tested
Protocol Fuzzing & Mutation	PASS	Not Tested
Power Fail	PASS	PASS
Persistence of Data	PASS	PASS
High Availablility (HA)	PASS	Not Tested
Failover – Legitimate Traffic	PASS	Not Tested
Stateful Operation	PASS	Not Tested
Total Cost of Ownership		
Ease of Use		
Initial Setup (Hours)	8	
Time Required for Upkeep (Hours per Year)	Contact NSS Labs	
Time Required to Tune (Hours per Year)	Contact NSS Labs	
Expected Costs		
Initial Purchase (hardware as tested)	$80,000	
Installation Labor Cost (@$75/hr)	$600	
Annual Cost of Maintenance & Support (hardware/software)	$17,500	
Annual Cost of Updates (IPS/AV/etc.)	$0	
Initial Purchase (centralized management system)	Contact NSS Labs	
Annual Cost of Maintenance & Support (centralized management system)	Contact NSS Labs	
Management Labor Cost (per Year @$75/hr)	Contact NSS Labs	
Tuning Labor Cost (per Year @$75/hr)	Contact NSS Labs	
Total Cost of Ownership		
Year 1	$98,100	
Year 2	$17,500	
Year 3	$17,500	
3-Year Total Cost of Ownership	$133,100	

Figure 21 – Scorecard

Test Methodology

Data Center Firewall (DCFW) Test Methodology v2.2

A copy of the test methodology is available 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