DATA CENTER INTRUSION PREVENTION SYSTEM COMPARATIVE REPORT

Performance

Authors – Thomas Skybakmoen, Keith Bormann, Morgan Dhanraj

Tested Products

Fortinet FortiGate 3000D v5.4.0, build 7184

HPE TippingPoint S7500NX v3.7.2.4252

IBM Security Network Protection XGS 7100 v5.3.2.1

Intel Security McAfee Network Security Platform NS9100 v8.2.5.120

Juniper Networks SRX5400 v12.3X48-D18

Palo Alto Networks PA-7050 v7.0.4

Environment

Data Center Intrusion Prevention System: Test Methodology v.2.0

Overview

Implementation of data center intrusion prevention system (DCIPS) solutions can be complex, with multiple factors affecting the overall performance of a solution.

The following factors should be considered over the course of the useful life of the product:

- Where will it be deployed?
- What is the predominant traffic mix?
- What security policy is applied?

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.

Sizing considerations are critical, as vendor performance claims (where protection typically is not enabled) can vary significantly from actual performance (where protection is enabled). Figure 1 depicts network-based vendors and their bandwidth performance. NSS Labs rates throughput based on the average results of "real-world" protocol mixes (enterprise perimeter, financial, education, data center, and US and EU mobile carrier) and 21 KB HTTP response-based capacity tests.

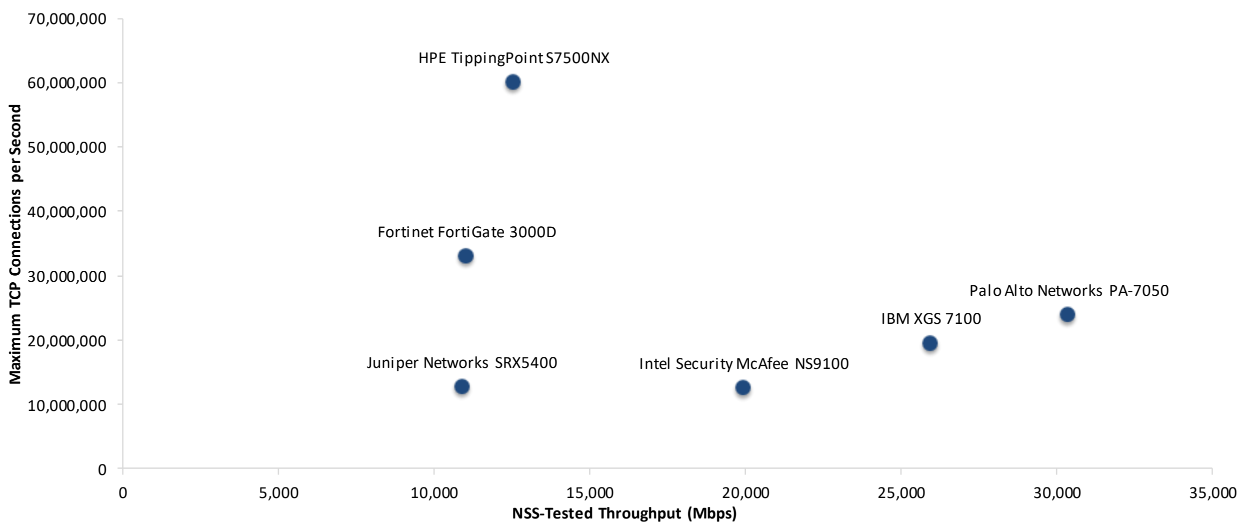

Figure 1 – Throughput and Connection Rates

Maximum TCP connections per second increases toward the top of the *y* axis. *NSS-Tested Throughput* (Mbps) increases toward the right side of the *x* axis.

Furthermore, if bypass mode is enabled, the DCIPS engine could be allowing uninspected traffic to enter the network once system resources are exhausted, and administrators would never be informed of threats in subsequent sessions.

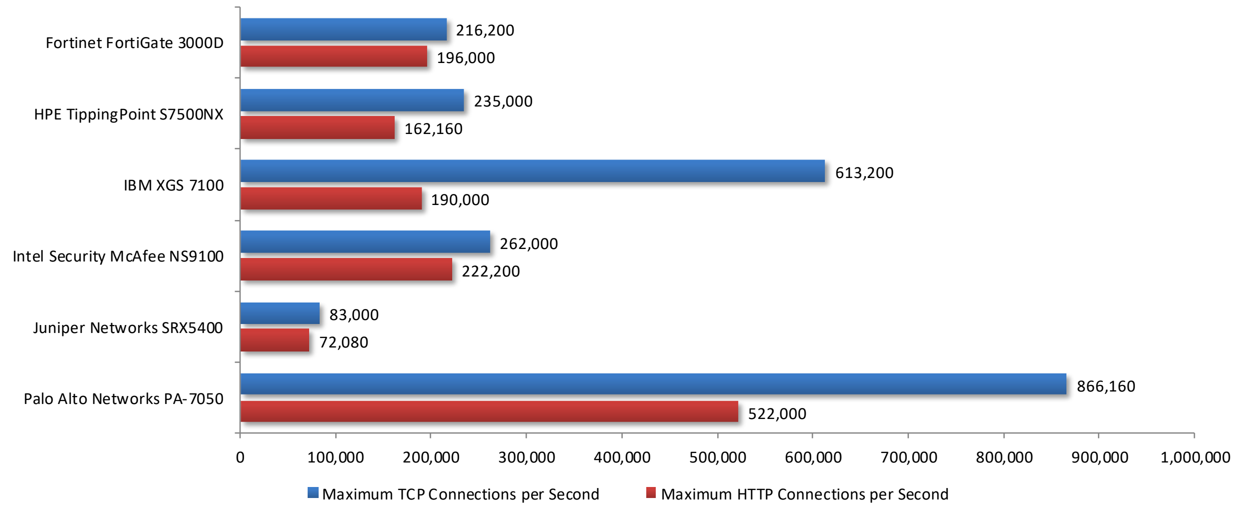

Figure 2 – Connection Dynamics

Performance is not just about raw throughput. Connection dynamics are also important and will often provide an indication of the inspection engine's effectiveness. If devices with high throughput capabilities cannot set up and tear down TCP or application-layer connections quickly enough, their maximum throughput figures can rarely be realized in a real-world deployment.

Table of Contents

Table of Figures

Analysis

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. This provides readers with the most useful information on key IPS security effectiveness and performance capabilities based on their expected usage.

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.

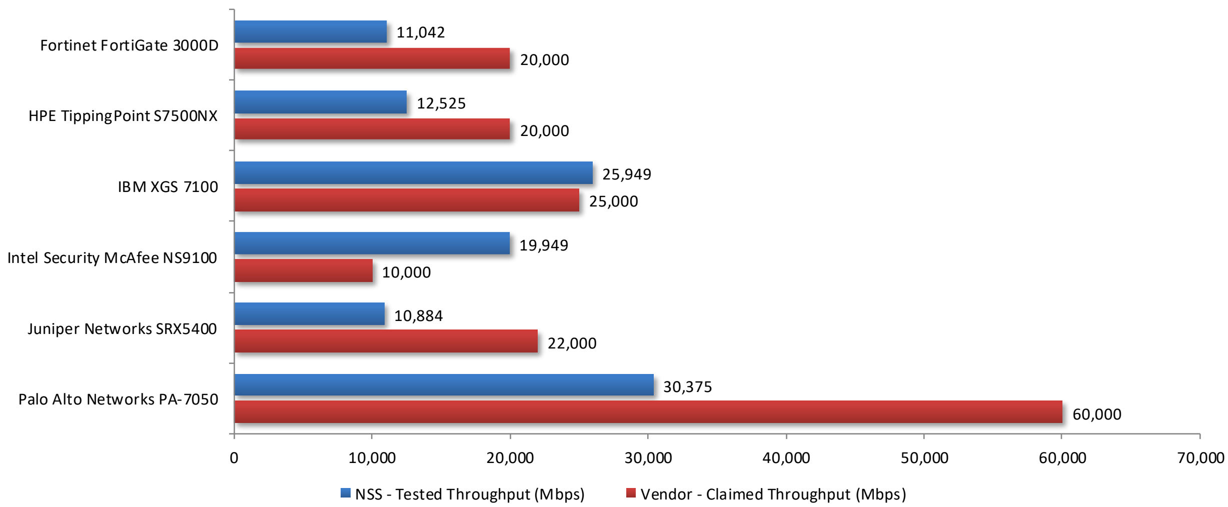

Figure 3 – Vendor-Claimed vs. NSS-Tested Throughput (Mbps)

Figure 3 depicts the difference between *NSS-Tested Throughput* and vendor performance claims, as vendor tests are often performed under ideal or unrealistic conditions. Where vendor marketing materials list throughput claims for both TCP (protection-enabled numbers) and UDP (large packet sizes), NSS selects the TCP claims, which are more realistic. Therefore, *NSS-Tested Throughput* typically is lower than vendor-claimed throughput, and often significantly so, since it more closely represents how devices will perform in real-world deployments.

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 4 depicts the results from the connection dynamics tests.

	Theoretical Maximum		Maximum Connections per Second		Maximum
Product	Concurrent TCP Connections	Concurrent TCP Connections w/Data	TCP	HTTP	HTTP Transactions per Second
Fortinet FortiGate 3000D	20,827,696	32,979,944	216,200	196,000	557,600
HPE TippingPoint S7500NX	60,000,000	60,000,000	235,000	162,160	340,000
IBM XGS 7100	19,643,120	19,556,052	613,200	190,000	179,680
Intel Security McAfee NS9100	13,948,358	12,514,912	262,000	222,200	821,400
Juniper Networks SRX5400	14,479,504	12,727,450	83,000	72,080	79,800
Palo Alto Networks PA-7050	24,080,568	23,897,100	866,160	522,000	779,640

Figure 4 – Concurrency and Connection Rates (I)

Beyond overall throughput of the device, connection dynamics can play an important role in sizing a security device that will not unduly impede the performance of a system or an application. By measuring maximum connection and transaction rates, a device can be sized more accurately than by simply examining throughput. By having knowledge of the maximum connections per second (CPS), it is possible to predict maximum throughput based on the traffic mix in a given enterprise environment. For example, if the device’s maximum HTTP CPS is 2,000, and average traffic size is 44 KB such that 2,500 CPS = 1 Gbps, then the tested device will achieve a maximum of 800 Mbps (i.e., (2,000/2,500) x 1,000 Mbps = 800 Mbps).

Maximum concurrent TCP connections and maximum TCP connections per second rates are also useful metrics when attempting to size a device accurately. Products with low connection/throughput ratios run the risk of exhausting connections before they reach their maximum potential throughput. By determining the maximum CPS, it is possible to predict when a device will fail in a given enterprise environment.

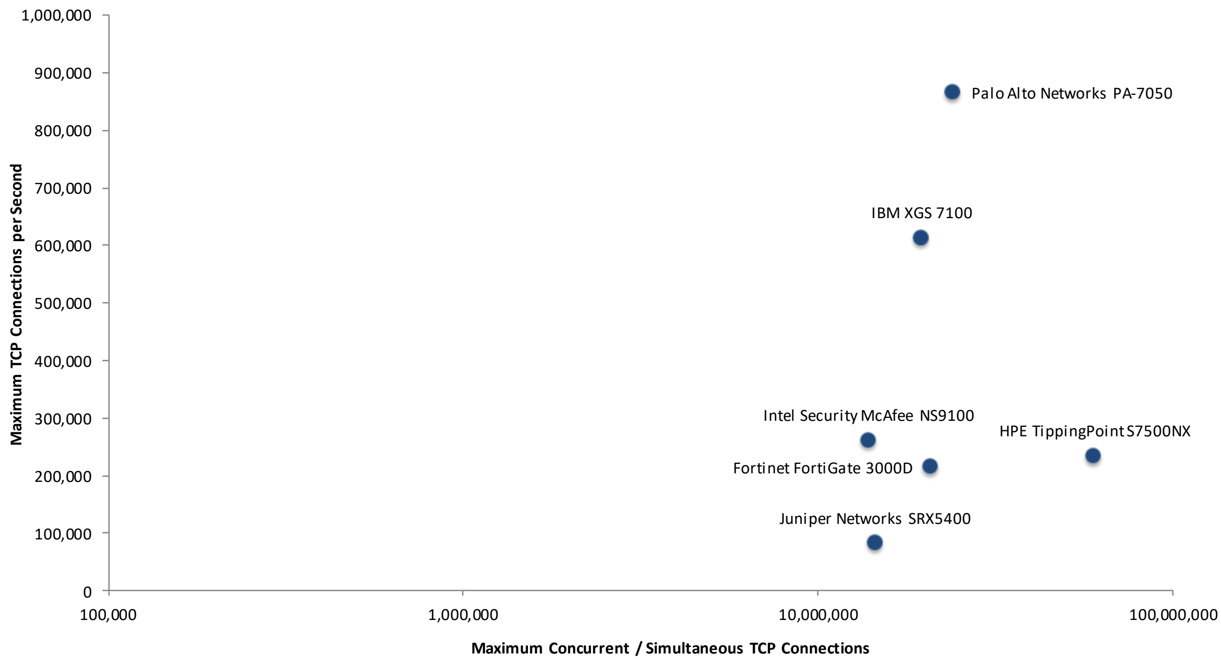

Figure 5 – Concurrency and Connection Rates (II)

The rate of maximum TCP connections per second increases toward the top of the *y* axis. The rate of concurrent/simultaneous connections increases toward the right side of the *x* axis.

HTTP Connections per Second and Capacity

Inline DCIPS devices exhibit an inverse correlation between security effectiveness and performance. The more network background traffic there is, the higher the chance of that traffic going uninspected and of malicious traffic going undetected. Furthermore, it is important to consider the “real-world” mix of traffic that a device will encounter.

The goal of these tests is to stress the HTTP detection engine and determine how the system under test(SUT) 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 SUT 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.

HTTP Connections per Second and Maximum Capacity (Throughput)

Figures 6 through 11 depict the maximum throughput achieved across a range of different HTTP response sizes that may be encountered in a typical corporate network.

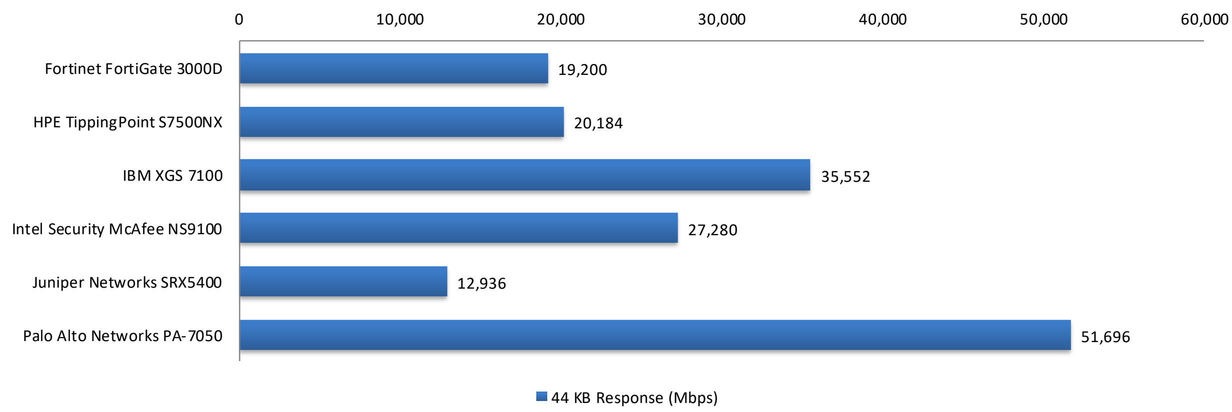

Figure 6 – Maximum Throughput per Device with 44 KB Response

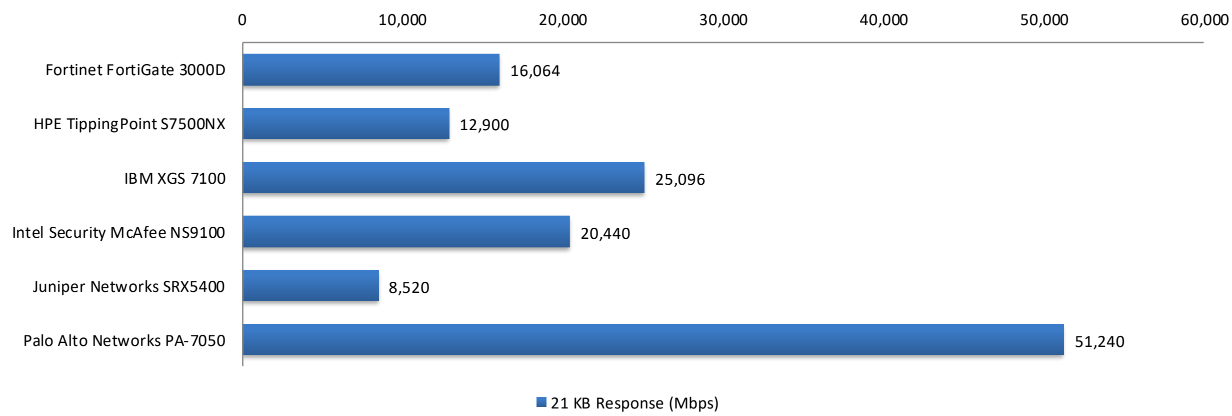

Figure 7 – Maximum Throughput per Device with 21 KB Response

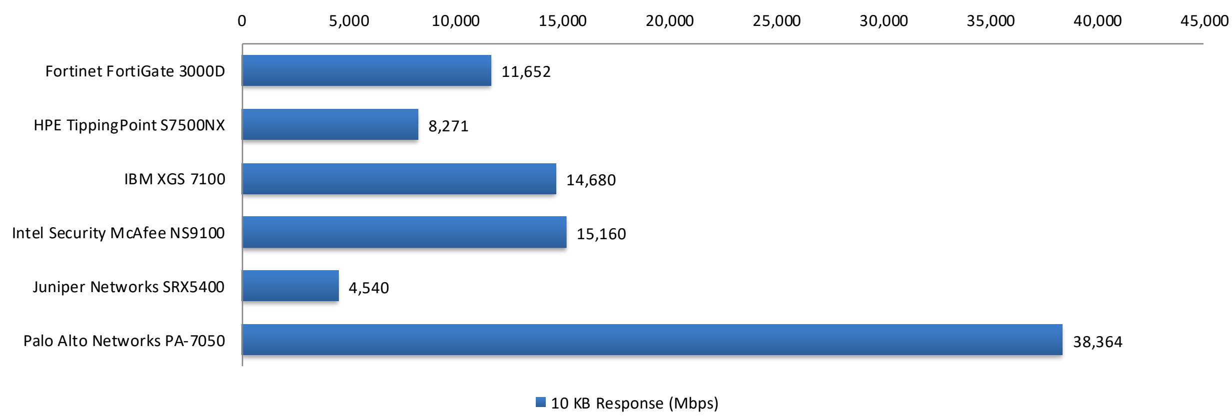

Figure 8 – Maximum Throughput per Device with 10 KB Response

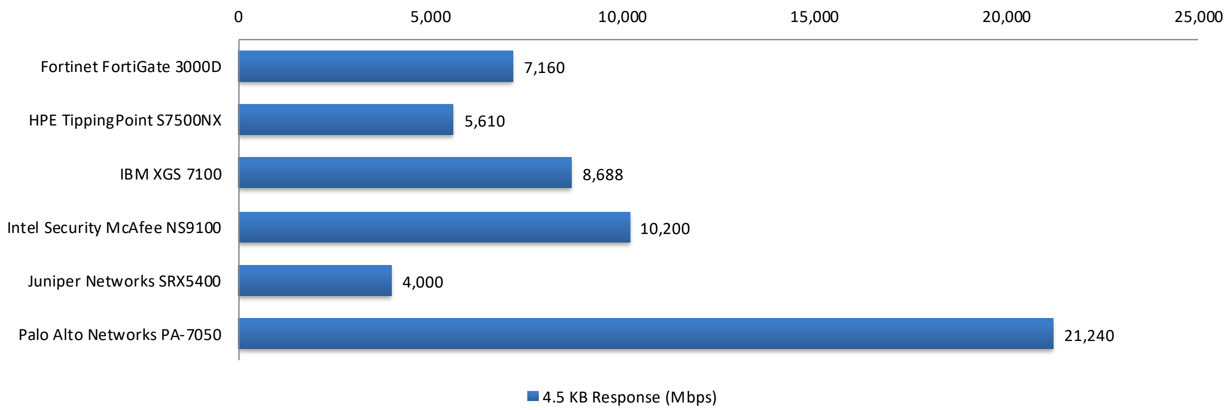

Figure 9 – Maximum Throughput per Device with 4.5 KB Response

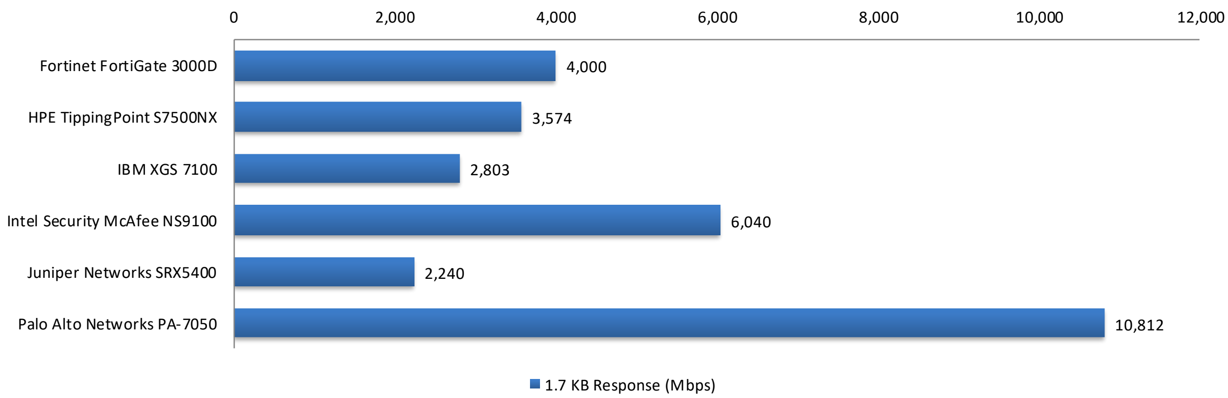

Figure 10 – Maximum Throughput per Device with 1.7 KB Response

Figure 11 depicts the maximum application layer connection rates (HTTP connections per second) achieved with different HTTP response sizes (from 44 KB down to 1.7 KB).

Product	44 KB Response	21 KB Response	10 KB Response	4.5 KB Response	1.7 KB Response
Fortinet FortiGate 3000D	48,000	80,320	116,520	143,200	160,000
HPE TippingPoint S7500NX	50,460	64,500	82,710	112,200	142,950
IBM XGS 7100	88,880	125,480	146,800	173,760	112,120
Intel Security McAfee NS9100	68,200	102,200	151,600	204,000	241,600
Juniper Networks SRX5400	32,340	42,600	45,400	80,000	89,600
Palo Alto Networks PA-7050	129,240	256,200	383,640	424,800	432,480

Figure 11 – Maximum Connection Rates per Device with Various Response Sizes

Application Average Response Time at 90% Maximum Capacity

Figure 12 depicts the average application response time (application latency, measured in milliseconds) with different packet sizes (ranging from 44 KB down to 1.7 KB) recorded at 90% of the measured maximum capacity (throughput). A lower value indicates improved application response time.

Product	44 KB Latency (ms)	21 KB Latency (ms)	10 KB Latency (ms)	4.5 KB Latency (ms)	1.7 KB Latency (ms)
Fortinet FortiGate 3000D	2.89	2.80	2.68	2.35	2.47
HPE TippingPoint S7500NX	0.99	0.75	0.18	0.02	0.01
IBM XGS 7100	0.53	0.28	0.20	0.18	0.05
Intel Security McAfee NS9100	0.96	1.24	1.53	1.58	1.58
Juniper Networks SRX5400	3.02	2.07	1.91	1.07	1.03
Palo Alto Networks PA-7050	1.13	1.27	1.10	1.08	1.04

Figure 12 – Application Latency (Milliseconds) per Device with Various Response Sizes

Real-World Traffic Mixes

For details about "real-world" traffic protocol types and percentages, see the Data Center Intrusion Prevention System Test Methodology, available at www.nsslabs.com. The aim of these tests is to measure the performance of the SUT in a "real-world" environment by introducing additional protocols and real content, while still maintaining a precisely repeatable and consistent background traffic load. In order to simulate real use cases, different protocol mixes are utilized to model different data center deployment scenarios.

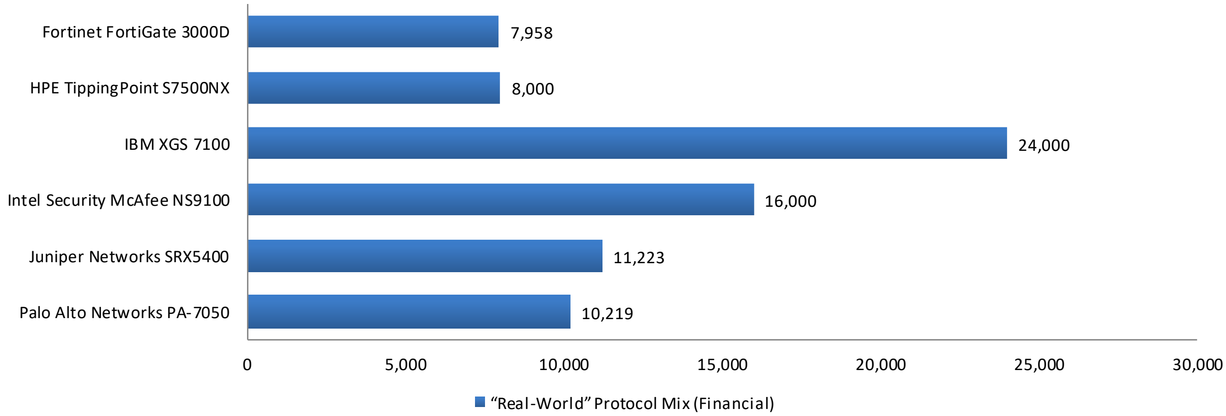

Figure 13 – "Real-World" Protocol Mix (Data Center Financial)

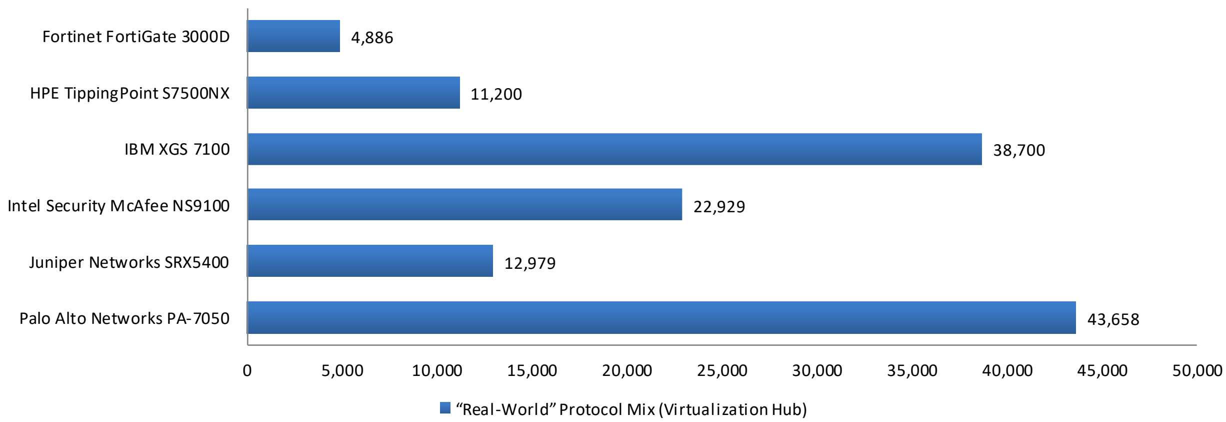

Figure 14 – "Real-World" Protocol Mix (Data Center Virtualization Hub)

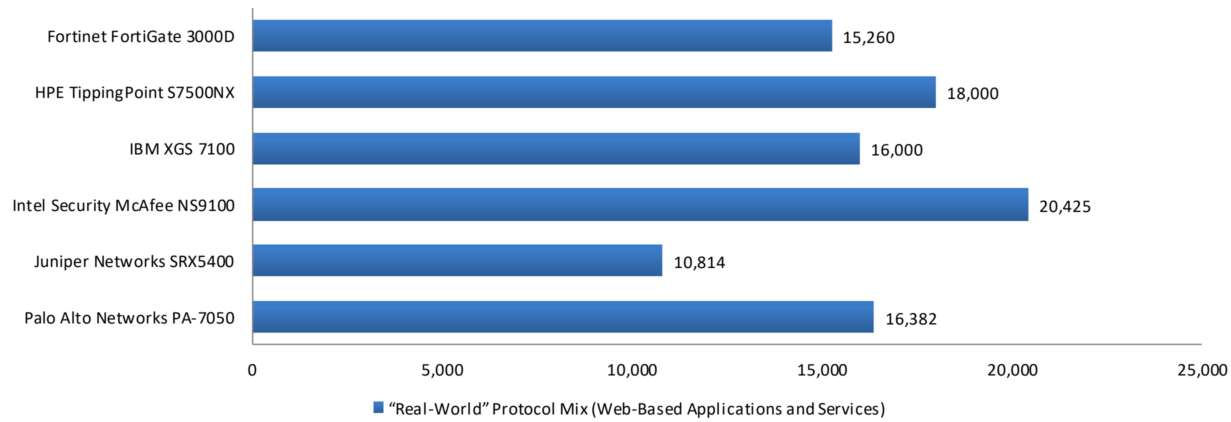

Figure 15 – "Real-World" Protocol Mix (Web-based Applications and Services)

UDP Throughput and Latency

The aim of this test is to determine the raw packet processing capability of each inline port pair of the device. The traffic does not attempt to simulate any "real-world" network condition. No TCP sessions are created during this test, and there is very little for the detection engine to do in the way of protocol analysis. However, this test is relevant because vendors are forced to perform inspection on UDP packets quickly in order to provide the highest level of network performance with the least amount of latency.

Figure 16 and Figure 17 depict the maximum UDP throughput (in megabits per second) achieved by each device using different packet sizes.

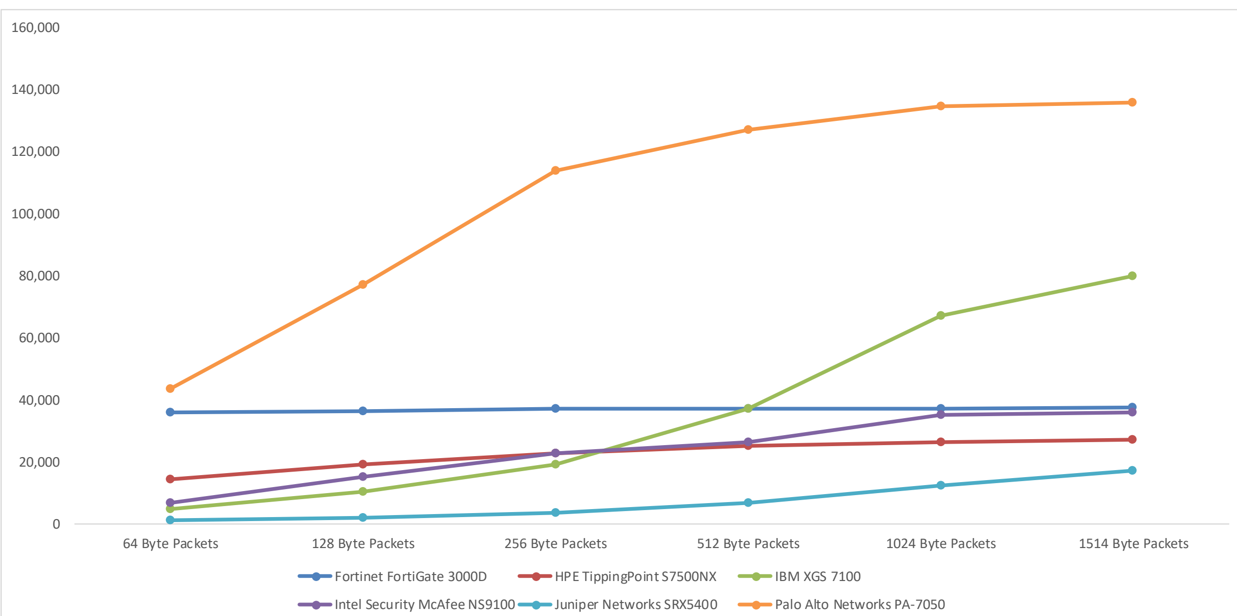

Figure 16 – UDP Throughput by Packet Size (Mbps)

The ability to provide the highest level of network performance with the least amount of latency has long been considered a minimum requirement for legacy firewalls, but it has often caused significant problems for DCIPS (and IPS) devices because of the amount of deep inspection they are expected to perform.

Product	Throughput (Mbps)					
	64-Byte Packets	128-Byte Packets	256-Byte Packets	512-Byte Packets	1024-Byte Packets	1514-Byte Packets
Fortinet FortiGate 3000D	35,840	36,432	37,012	37,112	37,320	37,520
HPE TippingPoint S7500NX	14,530	19,330	22,640	25,040	26,532	27,180
IBM XGS 7100	4,768	10,356	19,164	37,112	67,240	80,000
Intel Security McAfee NS9100	6,764	15,212	22,844	26,560	35,184	36,136
Juniper Networks SRX5400	1,182	1,986	3,684	6,682	12,378	17,070
Palo Alto Networks PA-7050	43,380	77,328	113,964	127,200	134,880	135,960

Figure 17 – UDP Throughput by Packet Size (Mbps)

Inline security devices that introduce high levels of latency lead to unacceptable response times for users, particularly where multiple security devices are placed in the data path. Figure 18 depicts the latency (in microseconds) as recorded during the UDP throughput tests at 90% of maximum load. Lower values are preferred.

Product	Latency (μs)					
	64-Byte Packets	128-Byte Packets	256-Byte Packets	512-Byte Packets	1024-Byte Packets	1514-Byte Packets
Fortinet FortiGate 3000D	3.0	3.2	3.6	4.4	5.4	6.5
HPE TippingPoint S7500NX	5.1	5.2	5.4	6.5	8.4	10.3
IBM XGS 7100	8.3	8.5	7.4	8.9	10.7	11.8
Intel Security McAfee NS9100	11.5	10.7	12.7	32.3	40.1	33.8
Juniper Networks SRX5400	78.9	81.9	84.4	85.7	93.6	96.6
Palo Alto Networks PA-7050	10.0	10.4	11.3	12.2	14.4	15.7

Figure 18 – UDP Latency by Packet Size (Microseconds [μs])

Test Methodology

Data Center Intrusion Prevention System v.2.0

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
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.