NEXT GENERATION INTRUSION PREVENTION SYSTEM COMPARATIVE REPORT

Performance

OCTOBER 13, 2016

Authors – Thomas Skybakmoen, Chris Conrad, Tim Otto, Morgan Dhanraj

Tested Products

Check Point Software Technologies, Ltd. 13800 Next Generation Firewall Appliance vR77.20

Cisco FirePOWER 8350 v6.0.1

Forcepoint Stonesoft Next Generation Firewall 3301 v6.0.2

Fortinet FortiGate 3000D v5.4.0

IBM Security Network Protection XGS 7100 v5.3.2.1

Intel Security McAfee Network Security Platform NS9100 v8.2.5.158

Palo Alto Networks PA-7050 v7.0.4

Trend Micro TippingPoint 7500NX v3.8.4.4525

Environment

Next Generation Intrusion Prevention System Test Methodology v.2.0

Overview

Implementation of a next generation intrusion prevention system (NGIPS) product can be complex, with multiple factors affecting the overall performance of the system.

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

- 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, and education) and 21 KB HTTP response-based capacity tests.

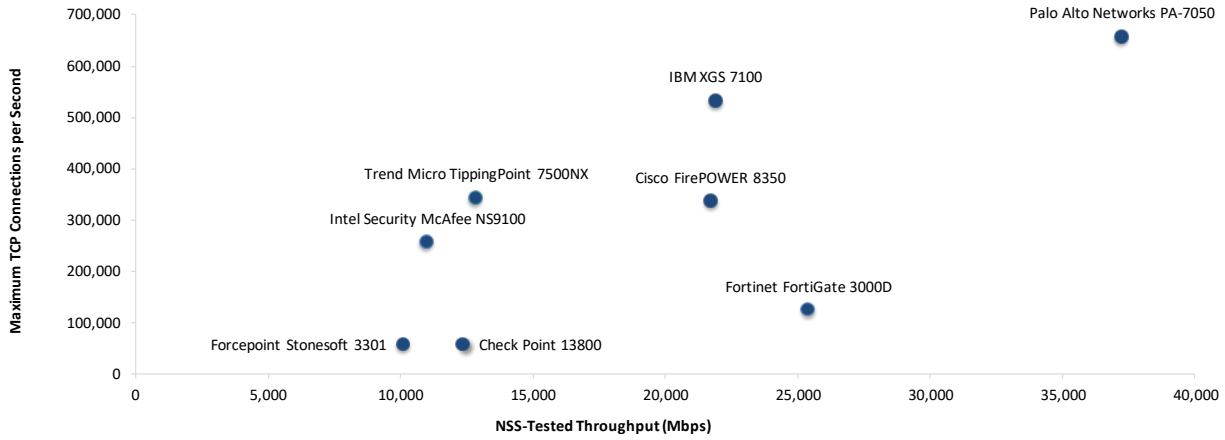

Figure 1 – Throughput and Connection Rates

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.

Products with low connection/throughput ratios run the risk of exhausting connections tables before they reach their maximum potential throughput, Furthermore, if bypass mode is enabled, the NGIPS 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.

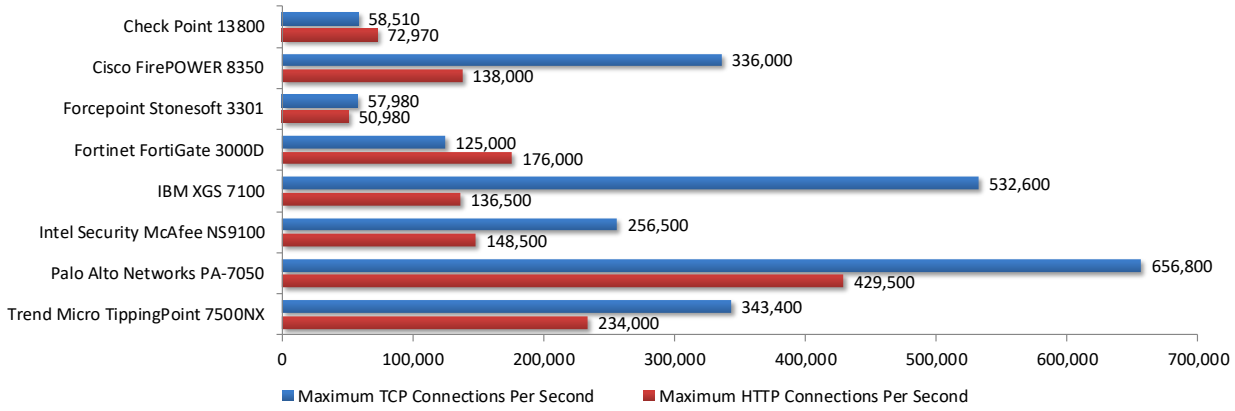

Figure 2 – Connection Dynamics

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

NGIPS products are deployed at the perimeter and within corporate networks to protect employee desktops, laptops, and PCs. NSS research has determined that the majority of enterprises do not tune their NGIPS products, but rather rely on a vendor's default/recommended policies and settings. Since there are enterprises that do tune their devices, NSS has tested NGIPS products using both tuned and vendor-recommended settings.

NSS defines a Vendor-Recommended policy as an out-of-the-box, vendor-pre-defined policy that is available to all customers. NSS defines tuning as the act of changing the device setting from the default/recommended setting to a specific configuration based on the environment it is protecting. In both cases, the signatures/filters/rules that trigger false positives are turned off, since that is what would happen in an enterprise environment.

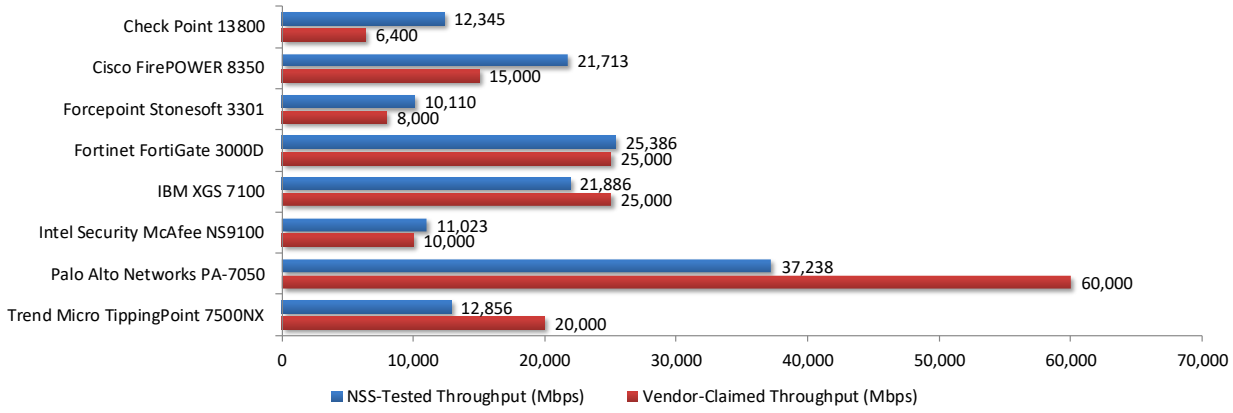

Figure 3 – Vendor-Claimed vs. NSS-Tested Throughput (Mbps)

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

Figure 4 depicts the results from the connection dynamics tests.

Product	Theoretical Maximum			Maximum Connections per Second		Maximum
	Concurrent TCP Connections	Concurrent TCP Connections w/Data	Single Connection	TCP	HTTP	HTTP Transactions per Second
Check Point 13800	2,233,218	2,256,265	788	58,510	72,970	120,100
Cisco FirePOWER 8350	60,000,000	60,000,000	1347	336,000	138,000	899,700
Forcepoint Stonesoft 3301	4,779,038	7,908,686	176	57,980	50,980	139,600
Fortinet FortiGate 3000D	17,268,580	17,363,009	1392	125,000	176,000	478,500
IBM XGS 7100	15,073,584	16,426,634	1417	532,600	136,500	150,100
Intel Security McAfee NS9100	12,447,669	12,470,033	988	256,500	148,500	717,800
Palo Alto Networks PA-7050	18,102,035	23,904,366	988.6	656,800	429,500	781,500
Trend Micro TippingPoint 7500NX	2,233,218	2,256,265	730	58,510	72,970	120,100

Figure 4 – Concurrency and Connection Rates

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.

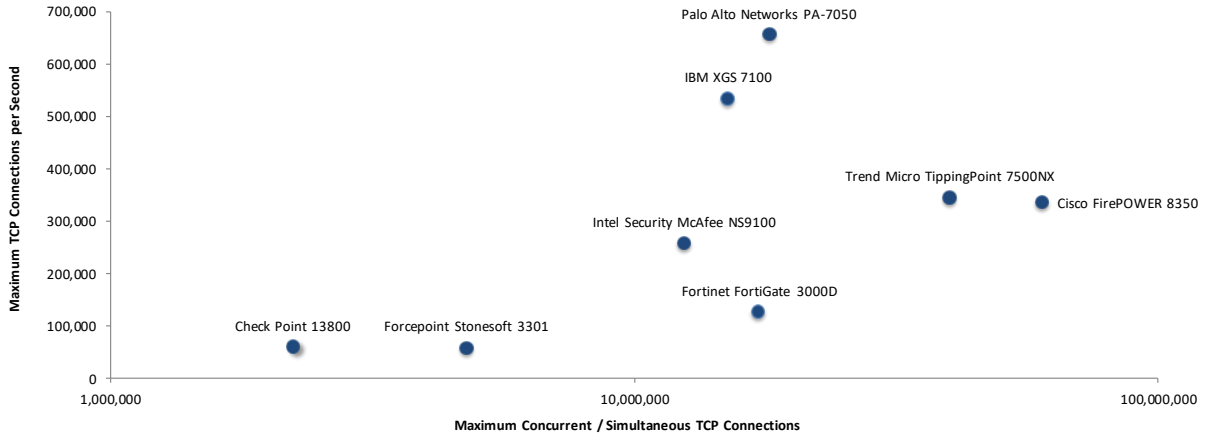

Figure 5 – Concurrency and Connection Rates (II)

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 NGIPS 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 10 depict the maximum throughput achieved across a range of different HTTP response sizes that may be encountered in a typical corporate network.

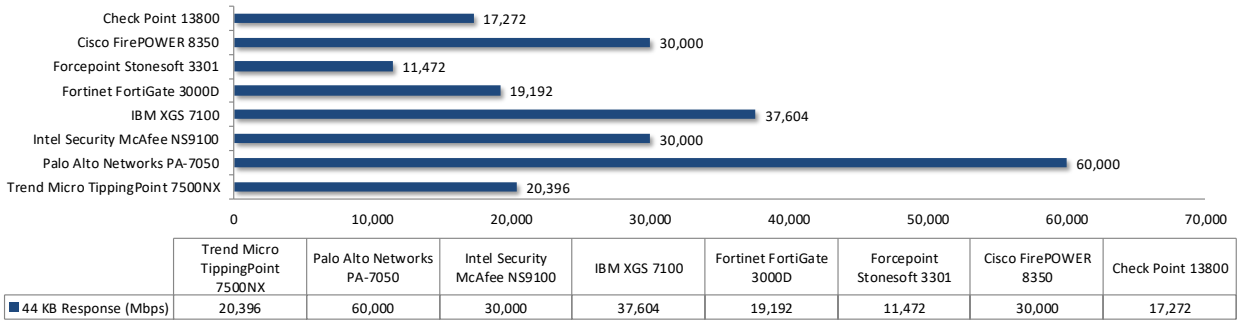

	Trend Micro TippingPoint 7500NX	Palo Alto Networks PA-7050	Intel Security McAfee NS9100	IBM XGS 7100	Fortinet FortiGate 3000D	Forcepoint Stonesoft 3301	Cisco FirePOWER 8350	Check Point 13800
44 KB Response (Mbps)	20,396	60,000	30,000	37,604	19,192	11,472	30,000	17,272

Figure 6 – Maximum Throughput per Device with 44 KB Response

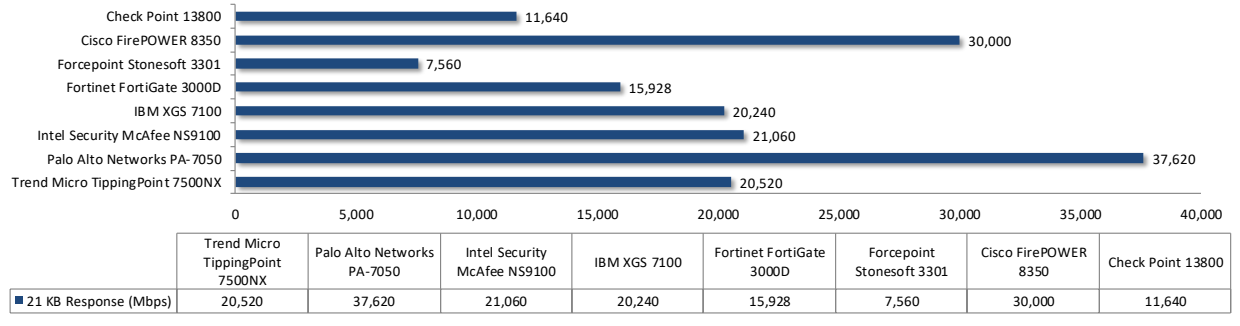

	Trend Micro TippingPoint 7500NX	Palo Alto Networks PA-7050	Intel Security McAfee NS9100	IBM XGS 7100	Fortinet FortiGate 3000D	Forcepoint Stonesoft 3301	Cisco FirePOWER 8350	Check Point 13800
21 KB Response (Mbps)	20,520	37,620	21,060	20,240	15,928	7,560	30,000	11,640

Figure 7 – Maximum Throughput per Device with 21 KB Response

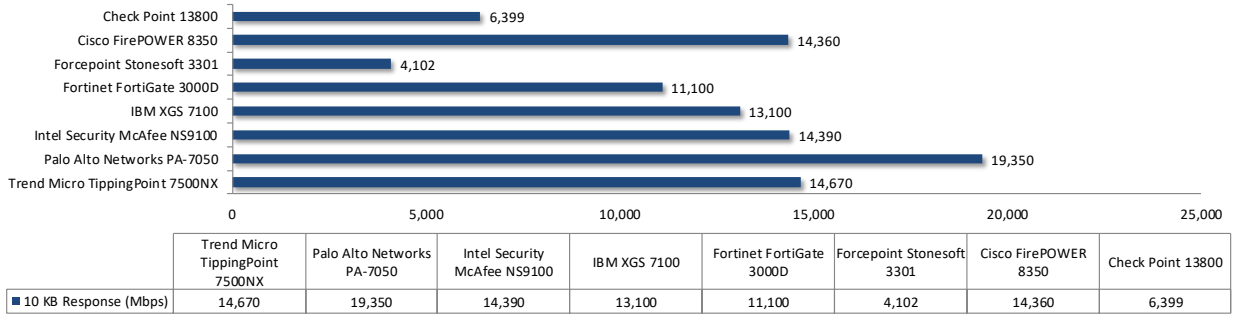

	Trend Micro TippingPoint 7500NX	Palo Alto Networks PA-7050	Intel Security McAfee NS9100	IBM XGS 7100	Fortinet FortiGate 3000D	Forcepoint Stonesoft 3301	Cisco FirePOWER 8350	Check Point 13800
10 KB Response (Mbps)	14,670	19,350	14,390	13,100	11,100	4,102	14,360	6,399

Figure 8 – Maximum Throughput per Device with 10 KB Response

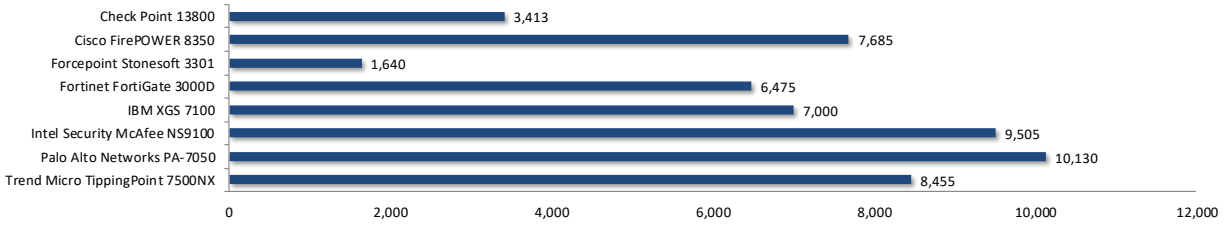

	Trend Micro TippingPoint 7500NX	Palo Alto Networks PA-7050	Intel Security McAfee NS9100	IBM XGS 7100	Fortinet FortiGate 3000D	Forcepoint Stonesoft 3301	Cisco FirePOWER 8350	Check Point 13800
4.5 KB Response (Mbps)	8,455	10,130	9,505	7,000	6,475	1,640	7,685	3,413

Figure 9 – Maximum Throughput per Device with 4.5 KB Response

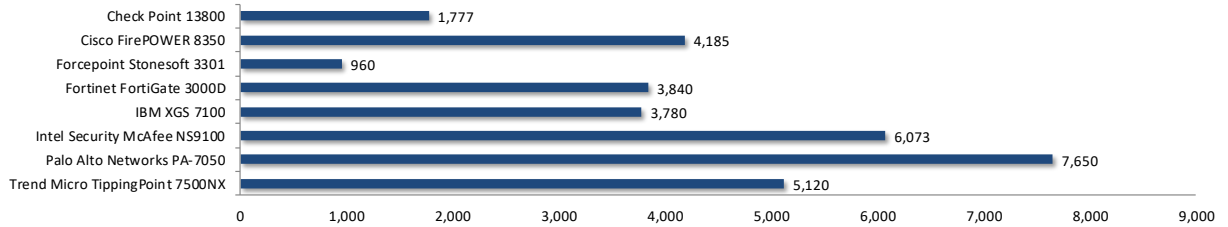

	Trend Micro TippingPoint 7500NX	Palo Alto Networks PA-7050	Intel Security McAfee NS9100	IBM XGS 7100	Fortinet FortiGate 3000D	Forcepoint Stonesoft 3301	Cisco FirePOWER 8350	Check Point 13800
1.7 KB Response (Mbps)	5,120	7,650	6,073	3,780	3,840	960	4,185	1,777

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
Check Point 13800	43,180	58,200	63,990	68,260	71,080
Cisco FirePOWER 8350	75,000	150,000	143,600	153,700	167,400
Forcepoint Stonesoft 3301	28,680	37,800	41,020	32,800	38,400
Fortinet FortiGate 3000D	47,980	79,640	111,000	129,500	153,600
IBM XGS 7100	94,010	101,200	131,000	140,000	151,200
Intel Security McAfee NS9100	75,000	105,300	143,900	190,100	242,900
Palo Alto Networks PA-7050	150,000	188,100	193,500	202,600	306,000
Trend Micro TippingPoint 7500NX	50,990	102,600	146,700	169,100	204,800

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)
Check Point 13800	1.71	1.21	0.80	0.44	0.09
Cisco FirePOWER 8350	1.33	1.62	1.63	2.36	3.31
Forcepoint Stonesoft 3301	4.11	3.66	2.38	0.49	0.40
Fortinet FortiGate 3000D	2.62	2.79	2.90	2.23	0.81
IBM XGS 7100	1.27	0.67	0.45	0.29	0.06
Intel Security McAfee NS9100	2.50	1.91	1.98	1.95	1.16
Palo Alto Networks PA-7050	2.37	1.56	1.01	1.00	0.11
Trend Micro TippingPoint 7500NX	0.64	0.49	0.23	0.02	0.00

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 Next Generation 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 deployment scenarios.

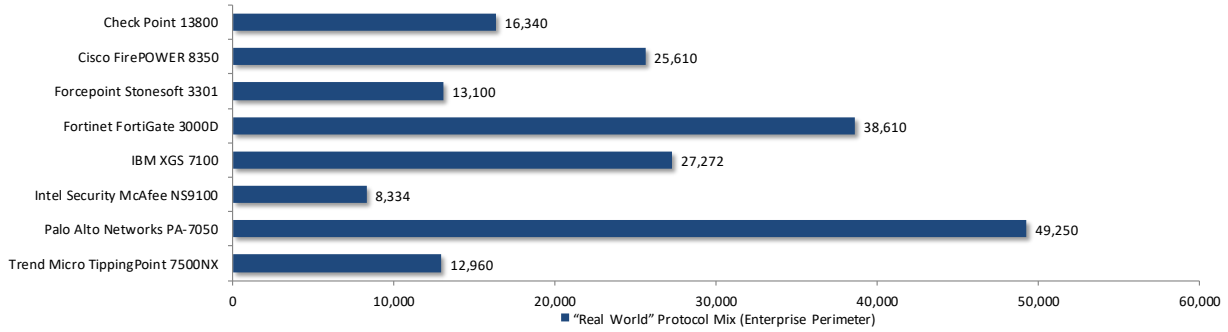

Figure 13 – "Real-World" Protocol Mix (Enterprise Perimeter)

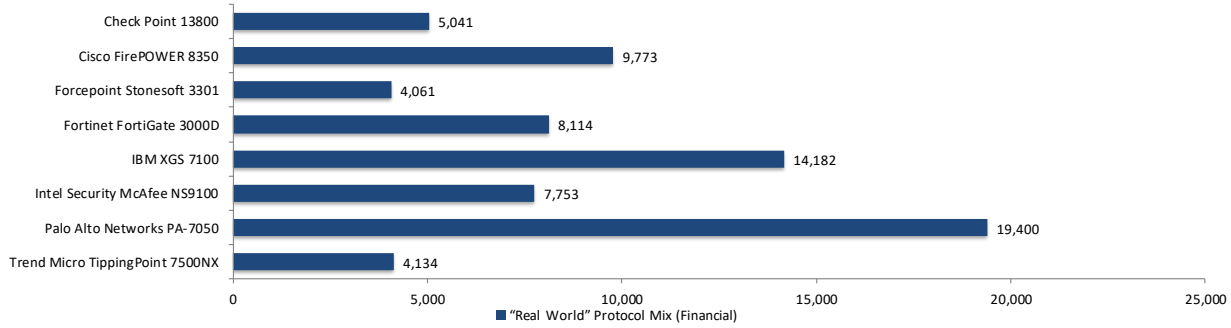

Figure 14 – "Real-World" Protocol Mix (Financial)

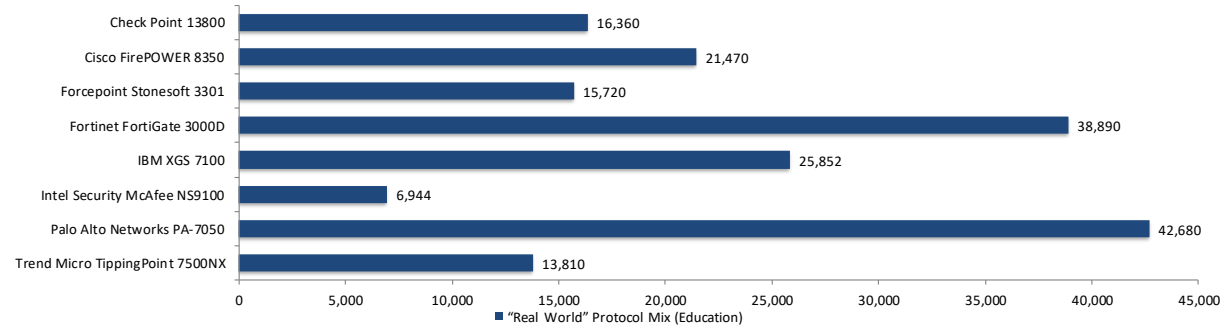

Figure 15 – "Real-World" Protocol Mix (Education)

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.

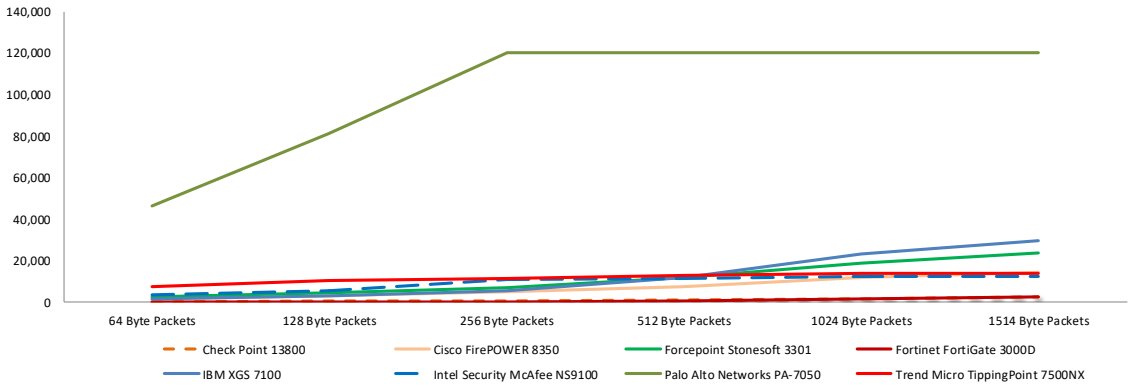

Figure 16 – UDP Throughput by Packet Size (Mbps)

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 NGIPS (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
Check Point 13800	92	275	321	871	1,616	2,511
Cisco FirePOWER 8350	1,324	2,673	4,746	7,166	11,820	14,060
Forcepoint Stonesoft 3301	2,590	4,383	6,689	11,984	18,564	23,480
Fortinet FortiGate 3000D	178	180	184	651	1,643	2,334
IBM XGS 7100	1,402	3,030	5,577	11,870	22,970	29,660
Intel Security McAfee NS9100	3,363	5,364	10,750	11,350	12,350	12,350
Palo Alto Networks PA-7050	46,080	81,180	120,000	120,000	120,000	120,000
Trend Micro TippingPoint 7500NX	7,574	10,180	11,370	12,960	13,950	13,950

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
Check Point 13800	28.0	46.0	44.0	49.0	49.0	64.0
Cisco FirePOWER 8350	86.0	99.0	99.0	104.0	111.0	124.0
Forcepoint Stonesoft 3301	233.0	245.4	255.4	285.2	297.9	304.4
Fortinet FortiGate 3000D	60.4	62.4	62.3	63.6	72.1	73.3
IBM XGS 7100	6.2	6.3	7.0	8.2	9.0	10.3
Intel Security McAfee NS9100	9.0	11.1	13.4	26.0	27.0	30.0
Palo Alto Networks PA-7050	10.2	10.6	11.7	12.7	14.6	15.7
Trend Micro TippingPoint 7500NX	5.0	5.1	5.5	6.5	8.7	11.0

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

Test Methodology

Next Generation 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