NEXT GENERATION FIREWALL COMPARATIVE REPORT

Performance

JUNE 06, 2017

Authors – Thomas Skybakmoen, Morgan Dhanraj

Tested Products

Barracuda NextGen Firewall F600.E20 Firmware Version 7.0.2

Check Point Software Technologies 15600 Next Generation Threat Prevention (NGTP) Appliance R77.20

Cisco Firepower 4110 v6.1.0.1

Forcepoint NGFW 3301 Appliance v6.1.2

Fortinet FortiGate 3200D FortiOS v5.4.4 GA Build 1117

Fortinet FortiGate 600D FortiOS v5.4.4 GA Build 1117

Juniper Networks SRX 4200 v15.1X49-D75.5

Palo Alto Networks PA-5250 PAN-OS 8.0.0

SonicWall NSA 6600 SonicOS 6.2

Sophos XG-750 Firewall v16.01

WatchGuard Firebox M4600 v11.10.7

Environment

Next Generation Firewall (NGFW) Test Methodology v7.0

Overview

Implementation of next generation firewall (NGFW) solutions can be a complex process, with multiple factors affecting the overall performance of the solution.

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

- Where will it be deployed and managed?
- What is the throughput for the target environment?
- What is the predominant traffic mix?
- Concurrency and connection rates
- 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, US mobile carrier, EU mobile carrier, and internal segmentation), and 21 KB HTTP response-based capacity tests.

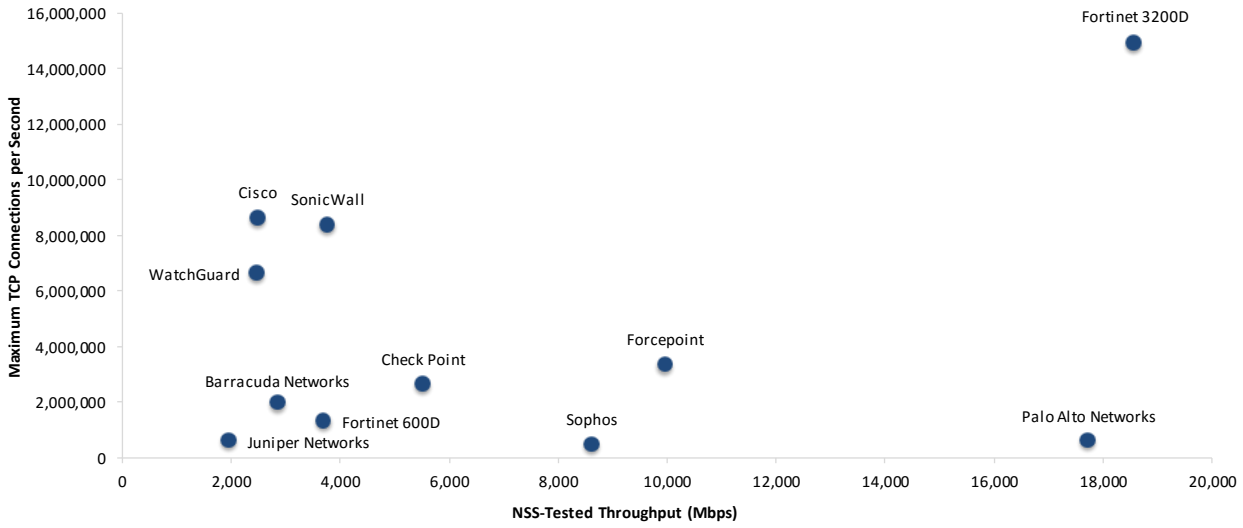

Figure 1 – Throughput and Connection Rates

Maximum TCP connections per second (CPS) 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 connection tables before they reach their maximum potential throughputs.

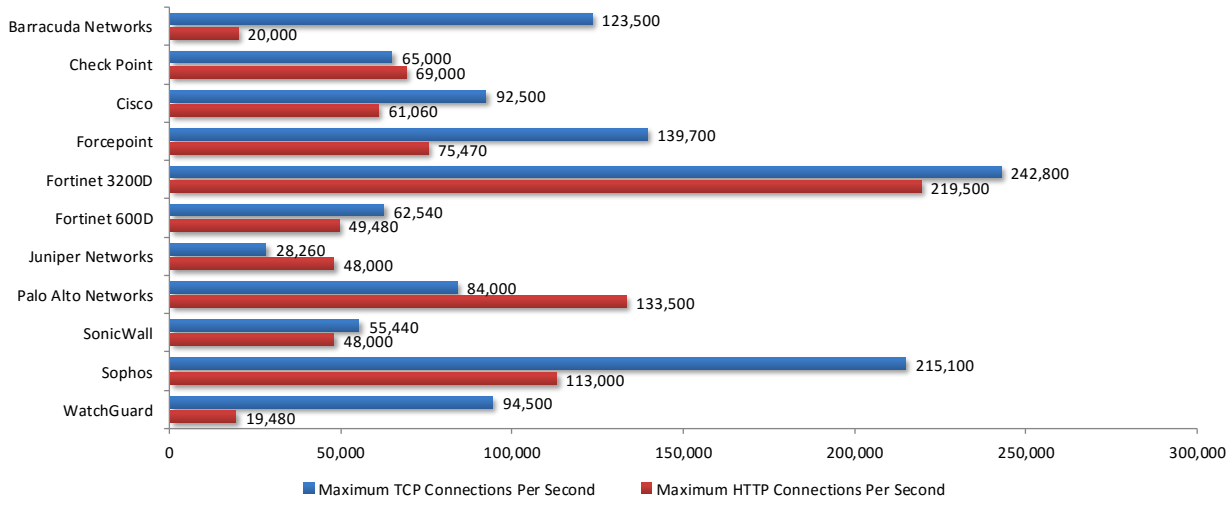

Figure 2 – Connection Dynamics

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.

Furthermore, if bypass mode is enabled, the NGFW 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.

Table of Contents

Tested Products 1

Environment 1

Overview 2

Analysis 6

UDP Throughput and Latency 6

Maximum Capacity 8

HTTP Capacity 10

Application Average Response Time at 90% Maximum Capacity 13

HTTP Capacity with HTTP Persistent Connections 13

HTTPS Capacity with HTTPS Persistent Connections 14

Real-World Traffic Mixes 16

Test Methodology 18

Contact Information 18

Table of Figures

Figure 1 – Throughput and Connection Rates 2
Figure 2 – Connection Dynamics 3
Figure 3 – Vendor-Claimed Throughput vs. NSS-Tested Throughput (Mbps) 6
Figure 4 – UDP Throughput by Packet Size (Mbps) 7
Figure 5 – UDP Throughput by Packet Size (Mbps) 7
Figure 6 – UDP Latency by Packet Size (Microseconds [µs]) 8
Figure 7 – Concurrency and Connection Rates (I) 9
Figure 8 – Concurrency and Connection Rates (II) 10
Figure 9 – Maximum Throughput per Device with 44 KB Response (Mbps) 11
Figure 10 – Maximum Throughput per Device with 21 KB Response (Mbps) 11
Figure 11 – Maximum Throughput per Device with 10 KB Response (Mbps) 11
Figure 12 – Maximum Throughput per Device with 4.5 KB Response (Mbps) 12
Figure 13 – Maximum Throughput per Device with 1.7 KB Response (Mbps) 12
Figure 14 – Maximum Connection Rates per Device with Various Response Sizes 12
Figure 15 – Application Latency (Milliseconds) per Device with Various Response Sizes 13
Figure 16 – HTTP 250 Capacity with HTTP Persistent Connections (CPS) 13
Figure 17 – HTTP 500 Capacity with HTTP Persistent Connections (CPS) 14
Figure 18 – HTTP 1000 Capacity with HTTP Persistent Connections (CPS) 14
Figure 19 – HTTPS 250 Capacity with HTTPS Persistent Connections (CPS) 14
Figure 20 – HTTPS 500 Capacity with HTTPS Persistent Connections (CPS) 15
Figure 21 – HTTPS 1000 Capacity with HTTPS Persistent Connections (CPS) 15
Figure 22 – “Real-World” Protocol Mix (Enterprise Perimeter) (Mbps) 16
Figure 23 – “Real-World” Protocol Mix (Financial) (Mbps) 16
Figure 24 – “Real-World” Protocol Mix (US Mobile Carrier) (Mbps) 17
Figure 25 – “Real-World” Protocol Mix (EU Mobile Carrier) (Mbps) 17
Figure 26 – “Real-World” Protocol Mix (Internal Segmentation) (Mbps) 17

Analysis

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

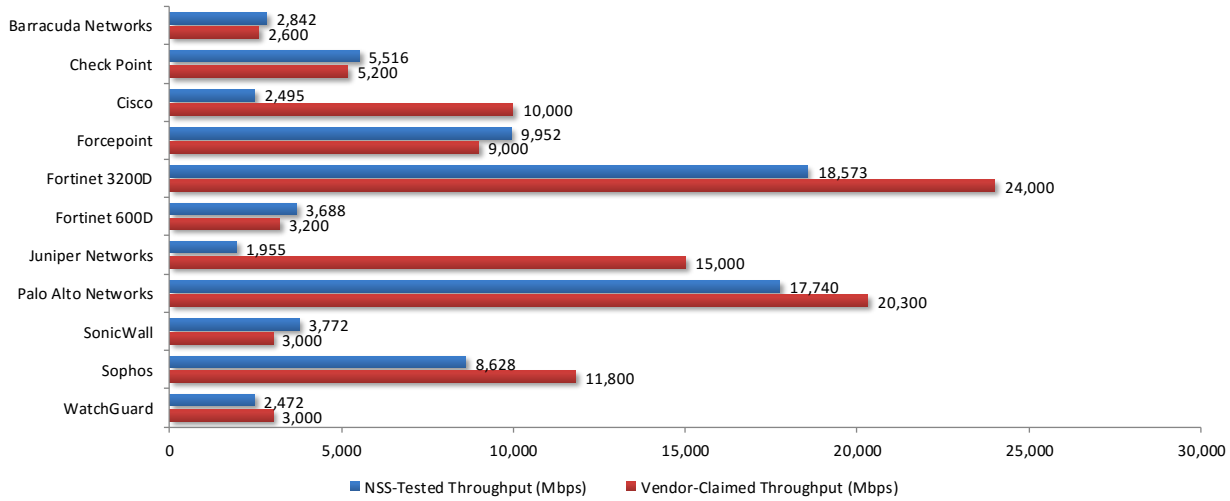

Figure 3 – Vendor-Claimed Throughput vs. NSS-Tested Throughput (Mbps)

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.

UDP Throughput and Latency

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 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. Multiple tests are run, and averages are taken where necessary.

This traffic does not attempt to simulate any real-world network conditions. No TCP sessions are created, 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 the device's effectiveness at forwarding packets quickly, in order to provide the highest level of network performance with the least amount of latency.

Figure 4 and Figure 5 depict the maximum UDP throughput (in megabits per second) achieved by each device using different packet sizes.

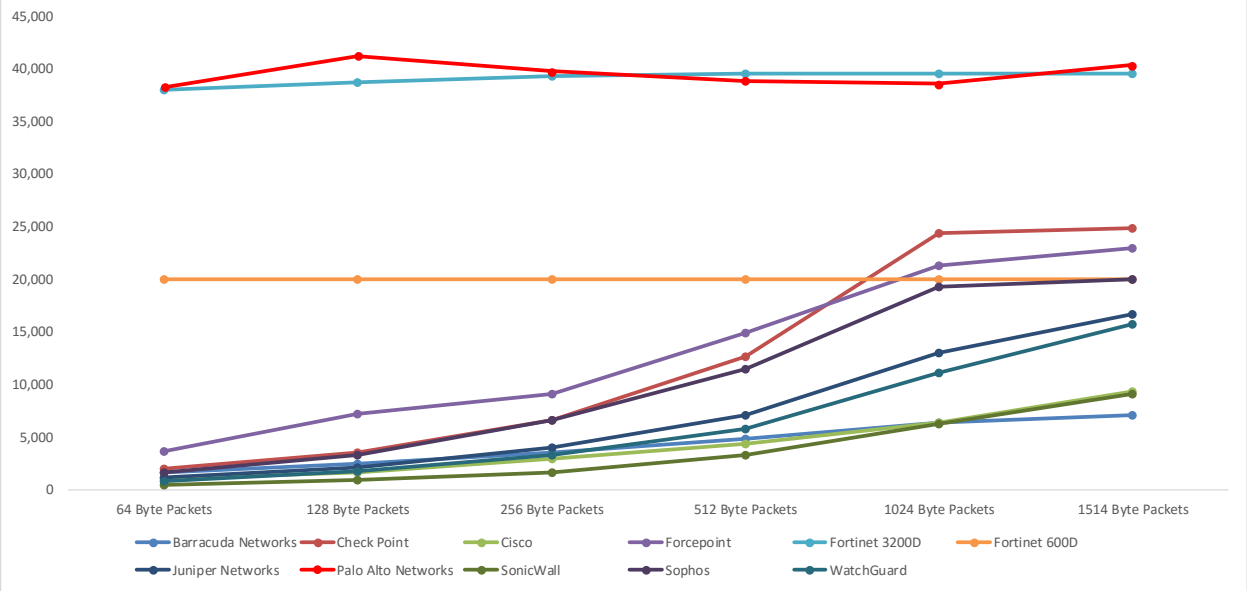

Figure 4 – UDP Throughput by Packet Size (Mbps)

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 NGFW (and IPS) devices because of the 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
Barracuda Networks	1,639	2,489	3,585	4,835	6,432	7,081
Check Point	1,977	3,476	6,571	12,660	24,440	24,840
Cisco	1,041	1,639	2,886	4,384	6,382	9,375
Forcepoint	3,635	7,169	9,127	14,860	21,350	22,960
Fortinet 3200D	38,000	38,780	39,390	39,570	39,540	39,620
Fortinet 600D	20,000	20,000	20,000	20,000	20,000	20,000
Juniper Networks	1,177	2,075	3,976	7,070	12,960	16,750
Palo Alto Networks	38,310	41,300	39,810	38,910	38,610	40,400
SonicWall	489	889	1,637	3,237	6,280	9,073
Sophos	1,637	3,278	6,581	11,520	19,350	20,000
WatchGuard	841	1,740	3,338	5,784	11,070	15,720

Figure 5 – 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 6 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
Barracuda Networks	16.0	26.0	35.0	41.0	50.0	63.0
Check Point	31.7	38.4	50.8	53.4	67.6	83.4
Cisco	208.5	111.0	98.0	77.0	61.2	172.0
Forcepoint	121.0	123.0	126.0	186.0	223.0	252.0
Fortinet 3200D	3.4	3.8	3.6	4.4	5.7	7.0
Fortinet 600D	5.0	7.8	13.0	22.6	42.0	61.0
Juniper Networks	33.0	37.7	42.0	43.1	49.2	52.1
Palo Alto Networks	9.4	9.9	10.2	10.9	12.2	12.8
SonicWall	28.8	29.1	29.3	29.9	31.0	32.0
Sophos	189.0	196.0	201.0	227.0	237.0	254.0
WatchGuard	34.4	69.0	91.0	109.0	109.0	111.0

Figure 6 – UDP Latency by Packet Size (Microseconds [µs])

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 (CPS), application layer transactions per second, and concurrent open connections. All packets contain valid payload and address data, and these tests provide an excellent representation of a live network at various connection/transaction rates.

Note that in all tests the following critical "breaking points"—where the final measurements are taken—are used:

- **Excessive concurrent TCP connections** – Latency within the NGFW is causing an unacceptable increase in open connections.
- **Excessive concurrent HTTP connections** – Latency within the NGFW is causing excessive delays and increased response time.
- **Unsuccessful HTTP transactions** – Normally, there should be zero unsuccessful transactions. Once these appear, it is an indication that excessive latency within the NGFW is causing connections to time out.

Figure 7 depicts the results from the connection dynamics tests.

Product	Theoretical Maximum		Maximum Connections per Second (CPS)		Maximum
	Concurrent TCP Connections	Concurrent TCP Connections w/Data	TCP	HTTP	HTTP Transactions per Second
Barracuda Networks	1,992,265	1,992,265	123,500	20,000	24,720
Check Point	2,797,220	2,637,271	65,000	69,000	99,540
Cisco	8,641,919	8,641,919	92,500	61,060	329,800
Forcepoint	18,120,665	3,358,100	139,700	75,470	130,900
Fortinet 3200D	15,138,012	14,964,386	242,800	219,500	567,700
Fortinet 600D	1,351,381	1,351,381	62,540	49,480	124,900
Juniper Networks	1,496,136	604,582	28,260	48,000	89,950
Palo Alto Networks	8,386,560	8,386,560	84,000	133,500	180,000
SonicWall	500,000	500,000	55,440	48,000	130,200
Sophos	10,485,744	6,668,725	215,100	113,000	204,900
WatchGuard	5,000,046	994,099	94,500	19,480	45,010

Figure 7 – Concurrency and Connection Rates (I)

In addition to overall throughput of the device, connection dynamics also 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. With knowledge of a device's maximum CPS, it is possible to predict its 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 CPS rates are also useful 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.

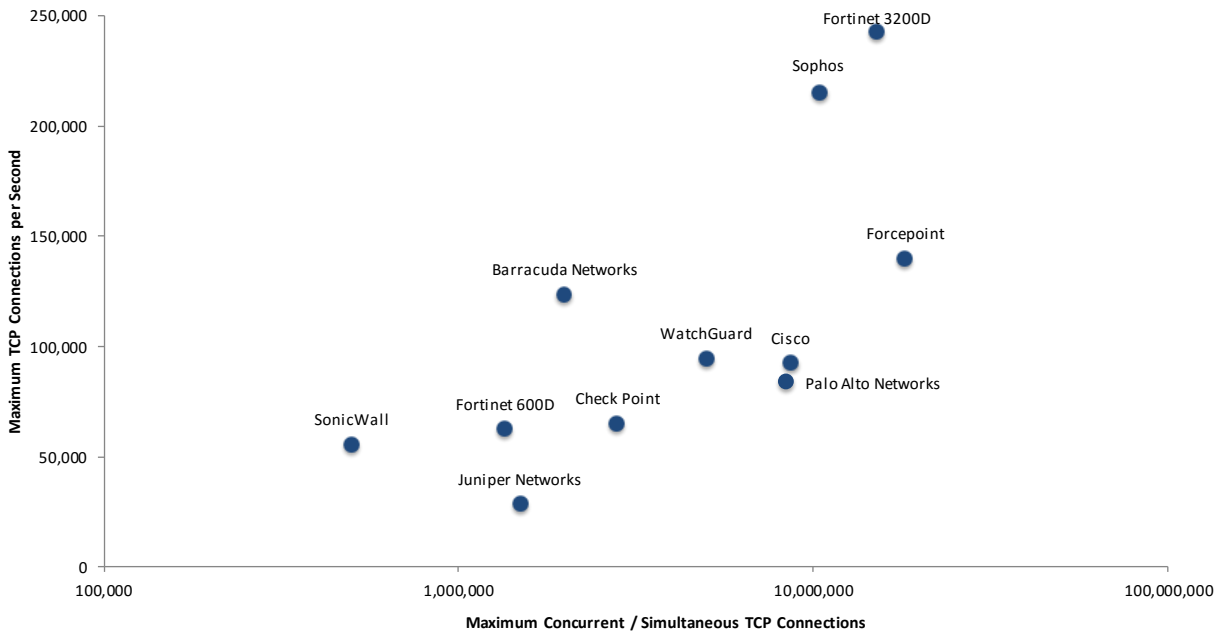

Figure 8 – Concurrency and Connection Rates (II)

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

HTTP Capacity

The aim of the HTTP capacity 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 is as close to "real-world" conditions as possible, while ensuring absolute accuracy and repeatability.

Each transaction consists of a single HTTP GET request. 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 9 through Figure 13 depict the maximum throughput achieved across a range of different HTTP response sizes that may be encountered in a typical corporate network.

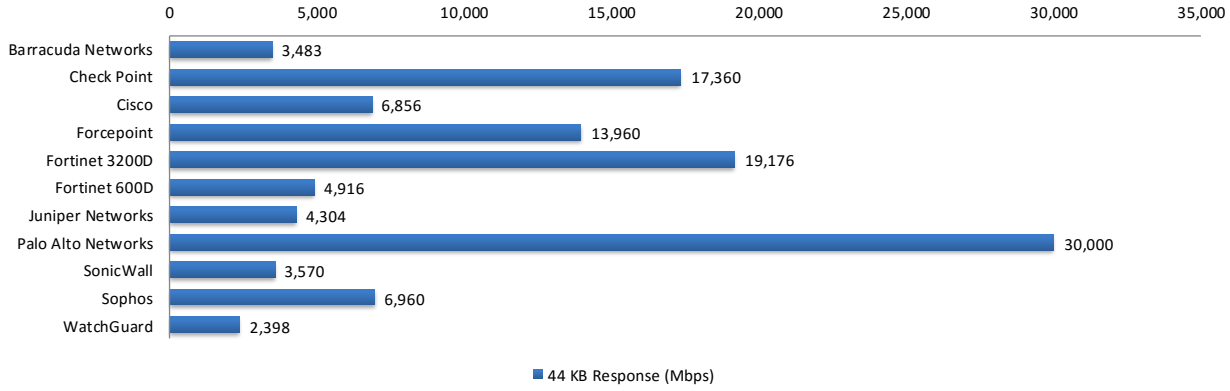

Figure 9 – Maximum Throughput per Device with 44 KB Response (Mbps)

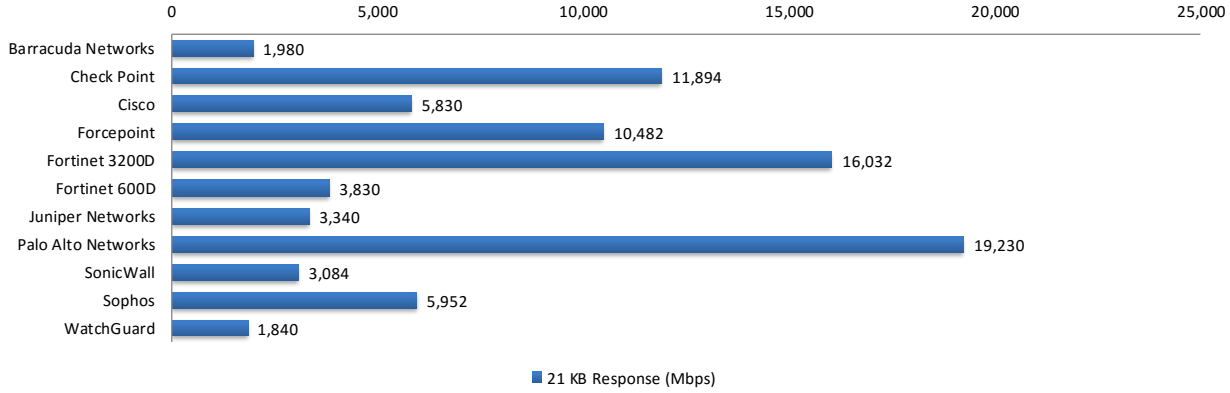

Figure 10 – Maximum Throughput per Device with 21 KB Response (Mbps)

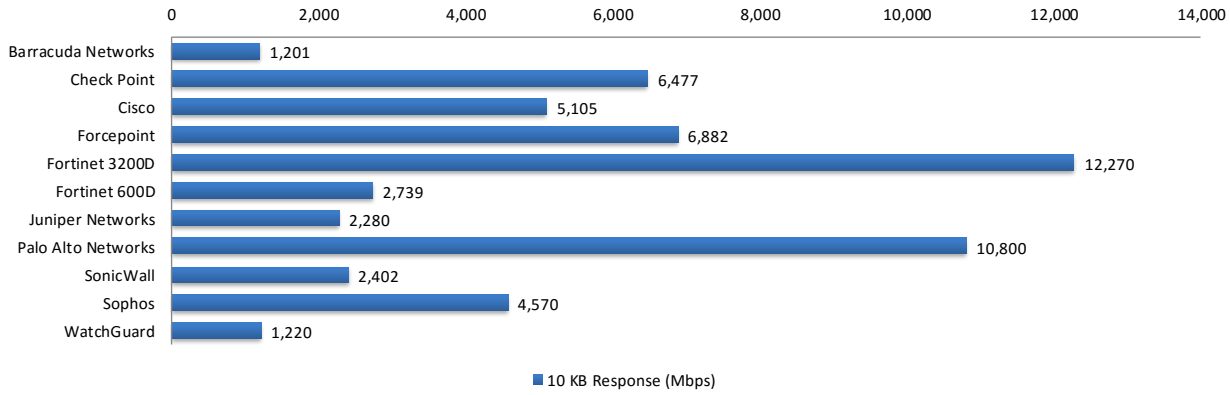

Figure 11 – Maximum Throughput per Device with 10 KB Response (Mbps)

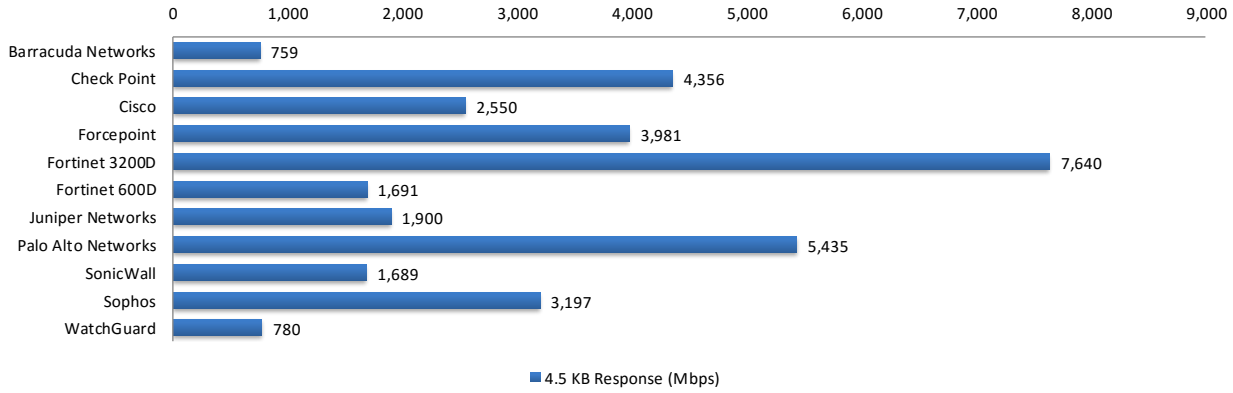

Figure 12 – Maximum Throughput per Device with 4.5 KB Response (Mbps)

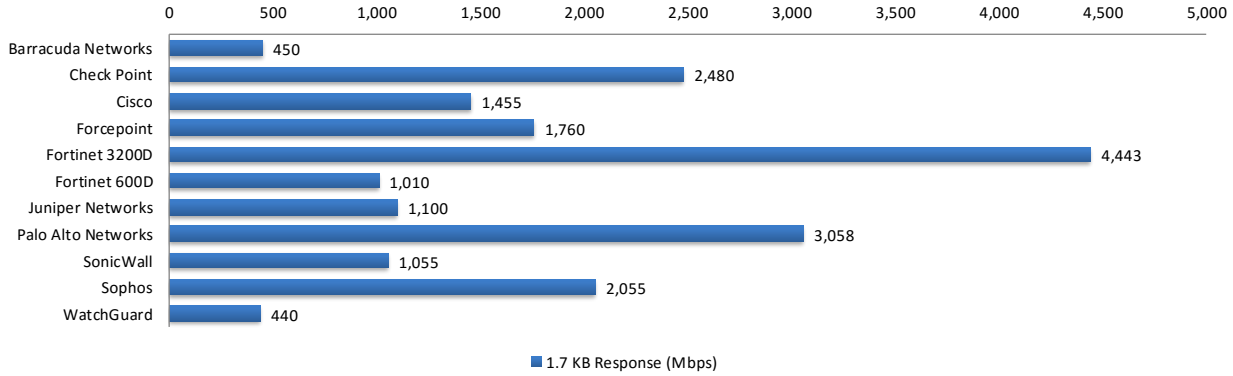

Figure 13 – Maximum Throughput per Device with 1.7 KB Response (Mbps)

Figure 14 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 Size	21 KB Response Size	10 KB Response Size	4.5 KB Response Size	1.7 KB Response Size
Barracuda Networks	8,707	9,900	12,010	15,180	18,000
Check Point	43,400	59,470	64,770	87,110	99,200
Cisco	17,140	29,150	51,050	51,000	58,200
Forcepoint	34,900	52,410	68,820	79,610	70,400
Fortinet 3200D	47,940	80,160	122,700	152,800	177,700
Fortinet 600D	12,290	19,150	27,390	33,810	40,410
Juniper Networks	10,760	16,700	22,800	38,000	44,000
Palo Alto Networks	75,000	96,150	108,000	108,700	122,300
SonicWall	8,925	15,420	24,020	33,770	42,200
Sophos	17,400	29,760	45,700	63,940	82,180
WatchGuard	5,994	9,202	12,200	15,590	17,600

Figure 14 – Maximum Connection Rates per Device with Various Response Sizes

Application Average Response Time at 90% Maximum Capacity

Figure 15 depicts the average application response time (application latency, measured in milliseconds) for 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 an 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)
Barracuda Networks	1.70	1.37	0.95	0.46	0.46
Check Point	1.40	0.71	0.22	0.24	0.20
Cisco	1.09	0.91	0.93	0.65	0.63
Forcepoint	3.89	2.75	2.14	1.40	1.38
Fortinet 3200D	2.85	3.36	4.41	5.02	5.07
Fortinet 600D	2.87	1.80	1.26	0.83	0.81
Juniper Networks	4.10	1.88	0.99	0.77	0.58
Palo Alto Networks	1.32	1.00	1.04	1.01	1.01
SonicWall	1.28	1.34	0.66	0.82	1.26
Sophos	9.69	7.45	5.22	3.33	3.12
WatchGuard	1.56	1.03	0.46	0.32	0.26

Figure 15 – Application Latency (Milliseconds) per Device with Various Response Sizes

HTTP Capacity with HTTP Persistent Connections

This test uses HTTP persistent connections, with each TCP connection containing 10 HTTP GETs and associated responses. 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 at various network loads. The stated response size is the total of all HTTP responses within a single TCP session.

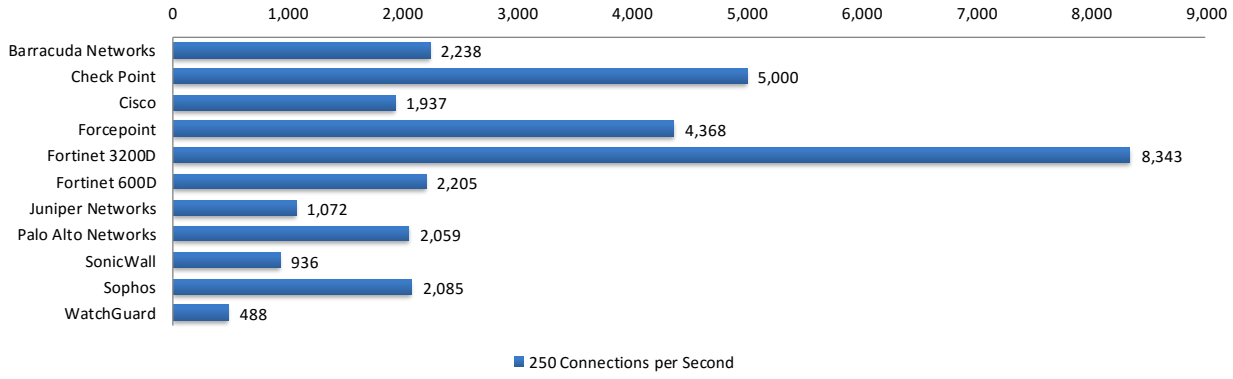

Figure 16 – HTTP 250 Capacity with HTTP Persistent Connections (CPS)

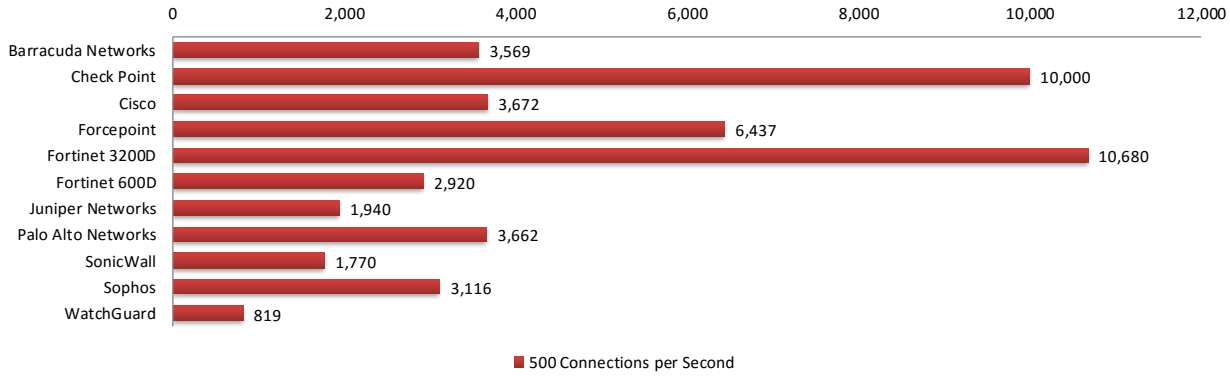

Figure 17 – HTTP 500 Capacity with HTTP Persistent Connections (CPS)

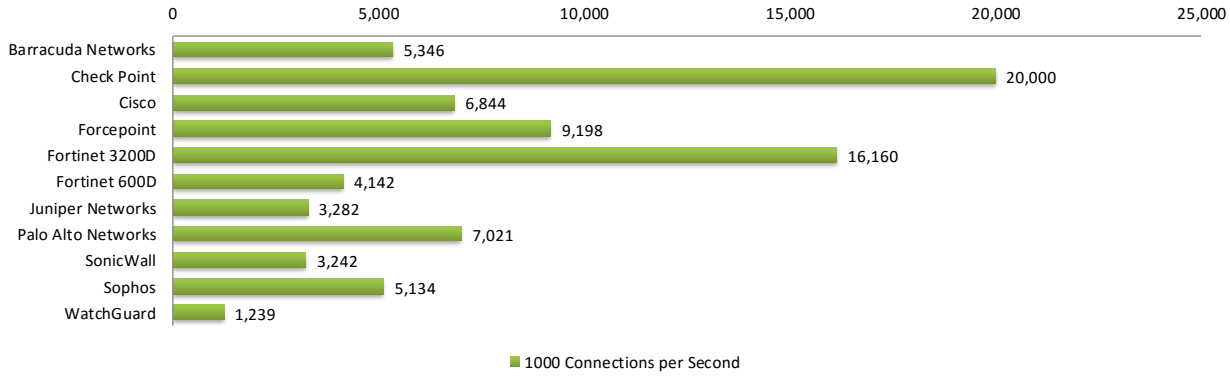

Figure 18 – HTTP 1000 Capacity with HTTP Persistent Connections (CPS)

HTTPS Capacity with HTTPS Persistent Connections

This test uses HTTPS persistent connections, with each TCP connection containing 10 HTTPS GETs and associated responses.

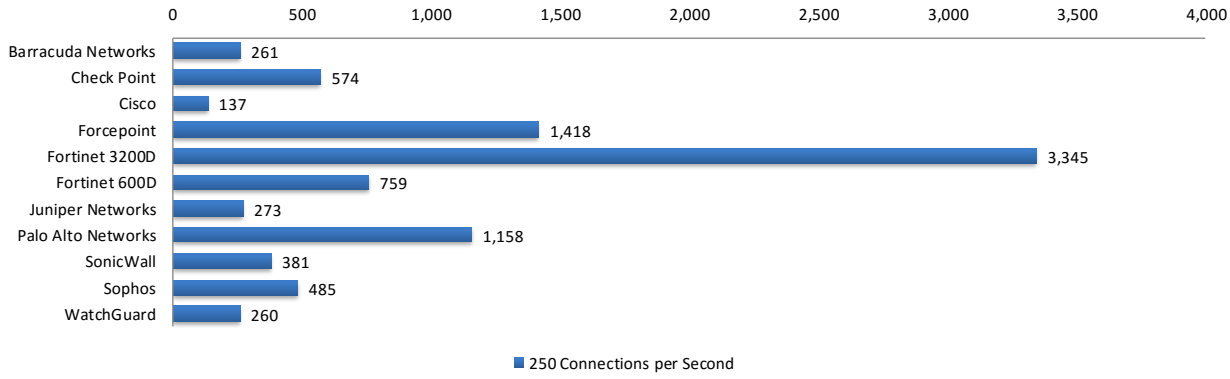

Figure 19 – HTTPS 250 Capacity with HTTPS Persistent Connections (CPS)

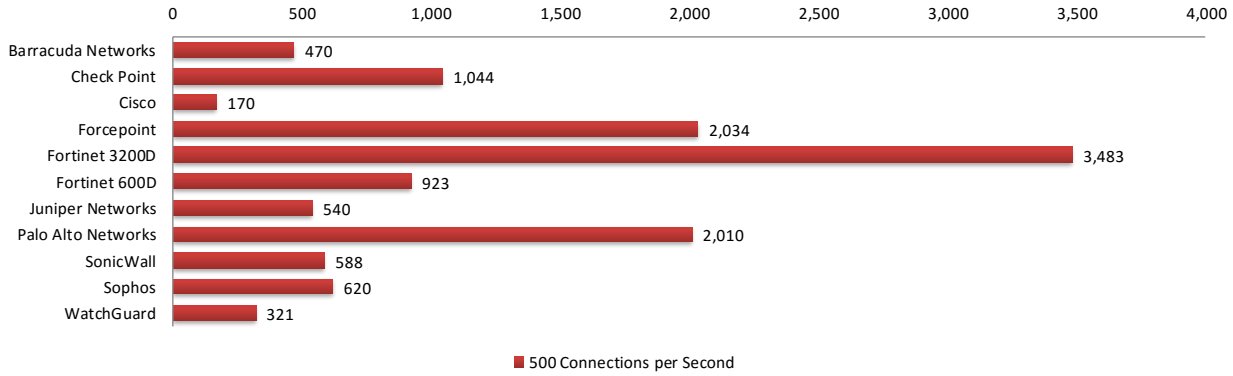

Figure 20 – HTTPS 500 Capacity with HTTPS Persistent Connections (CPS)

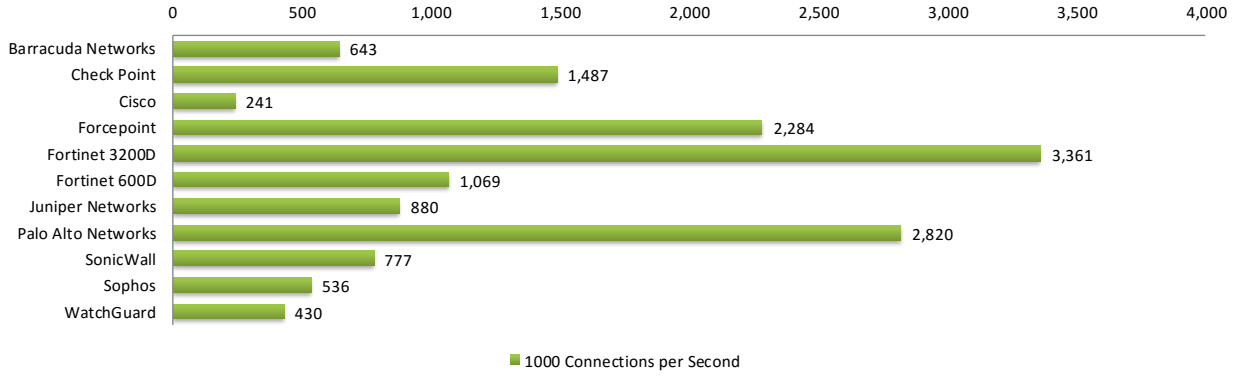

Figure 21 – HTTPS 1000 Capacity with HTTPS Persistent Connections (CPS)

Real-World Traffic Mixes

For details about real-world traffic protocol types and percentages, see the NSS Labs Next Generation Firewall Test Methodology, available at www.nsslabs.com.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. Different protocol mixes are utilized based on the intended location of the device (network core or perimeter) to reflect real use cases.

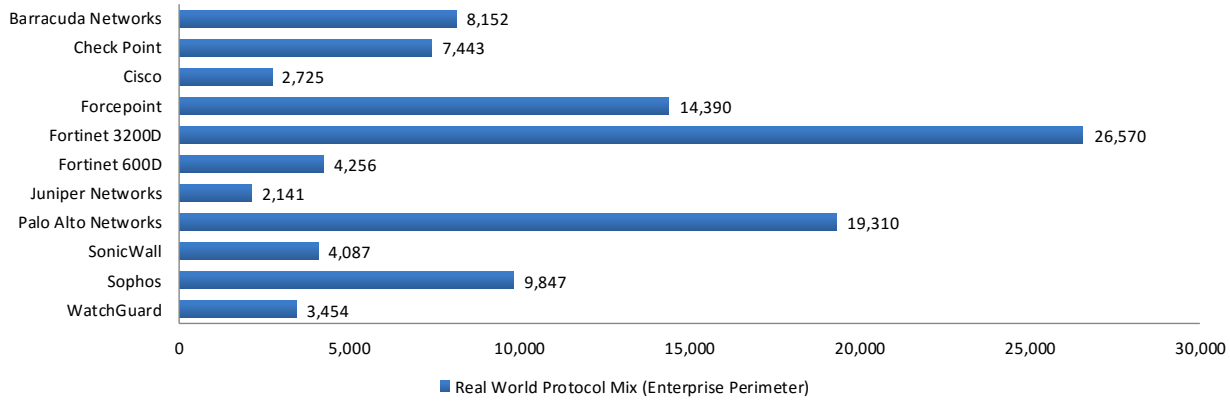

Figure 22 – "Real-World" Protocol Mix (Enterprise Perimeter) (Mbps)

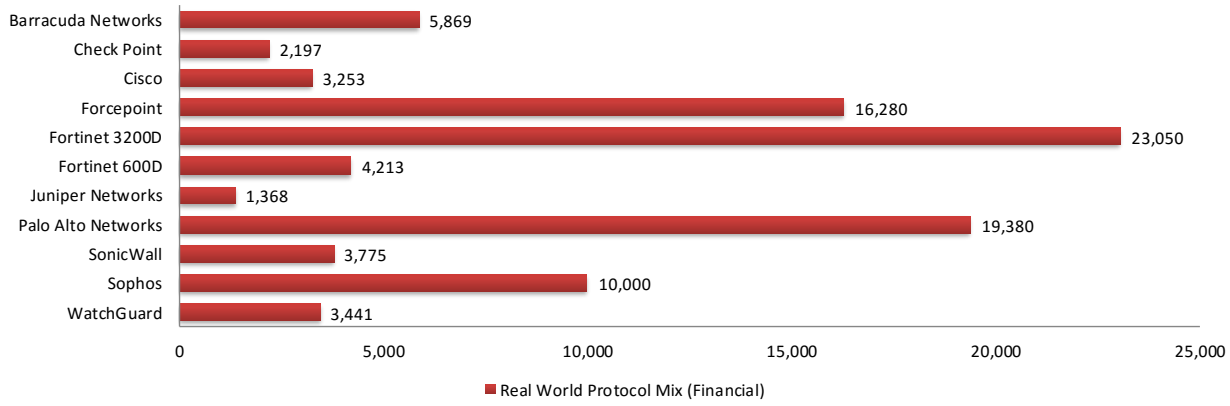

Figure 23 – "Real-World" Protocol Mix (Financial) (Mbps)

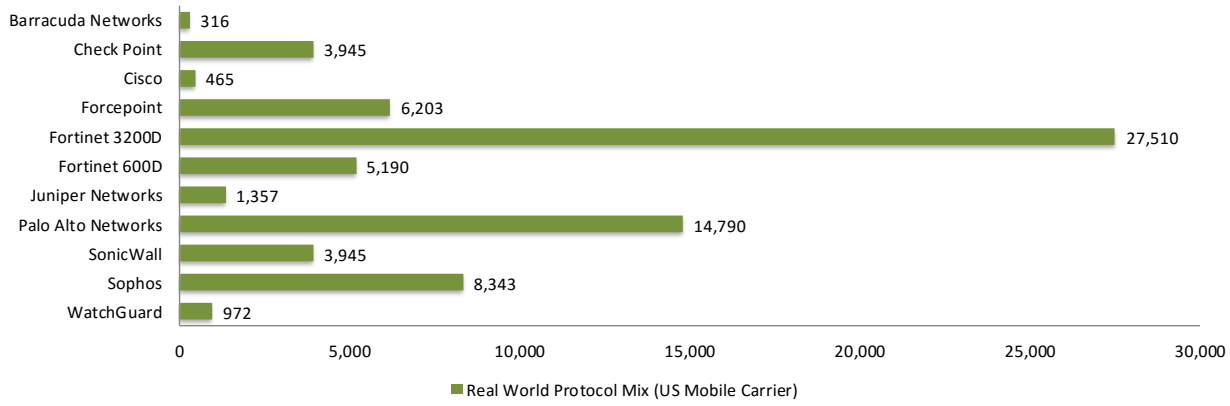

Figure 24 – "Real-World" Protocol Mix (US Mobile Carrier) (Mbps)

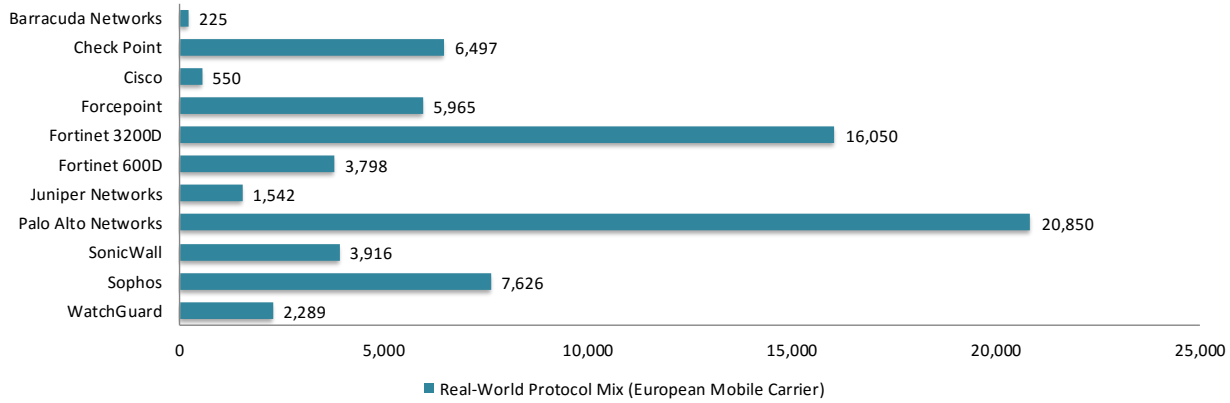

Figure 25 – "Real-World" Protocol Mix (EU Mobile Carrier) (Mbps)

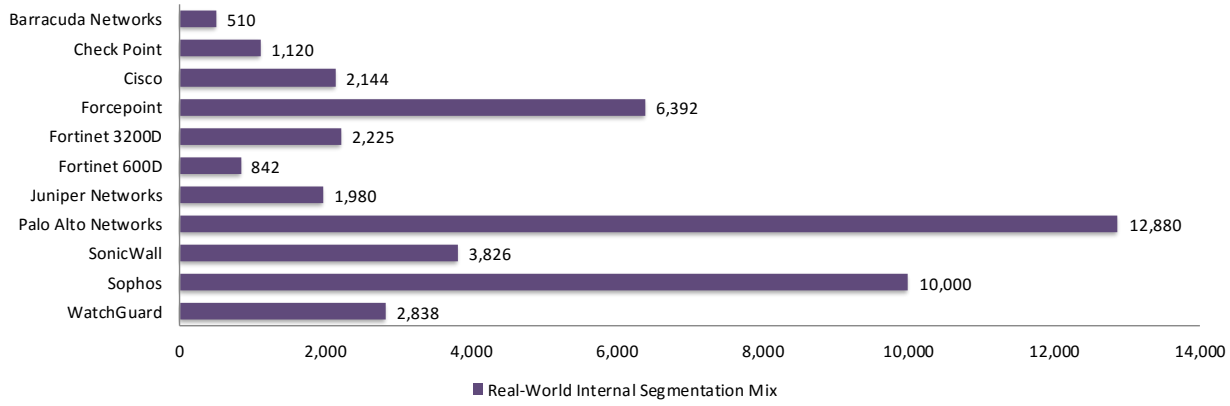

Figure 26 – "Real-World" Protocol Mix (Internal Segmentation) (Mbps)

Test Methodology

Next Generation Firewall Test Methodology Test Methodology v7.0

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

Contact Information

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