WEB APPLICATION FIREWALL COMPARATIVE ANALYSIS

Performance

Author – Thomas Skybakmoen

Tested Products

Barracuda Networks Web Application Firewall 960

Citrix NetScaler AppFirewall MPX 11520

Fortinet FortiWeb 1000D

F5 Big-IP ASM 10200

Imperva SecureSphere x6500

Sangfor M5900-F-I

Environment

Web Application Firewall: Test Methodology v6.2

Overview

Implementation of web application firewall (WAF) solutions can be a complex process with multiple factors affecting the overall performance of the solution.

Each of these factors should be considered over the course of the useful life of the solution, including:

- What applications and web services will it protect?
- What is the predominant traffic mix?
- What security policy is applied?

There is usually a trade-off between *security effectiveness* and performance (capacity); a *product's security effectiveness* should be evaluated within the context of its capacity (and vice versa). This ensures that new security protections do not adversely impact capacity and security shortcuts are not taken to maintain or improve capacity.

Sizing considerations are absolutely critical, since vendor capacity claims can vary significantly from actual capacity with protection enabled. *NSS-tested capacity* is an average of all of the HTTP response-based capacity tests.

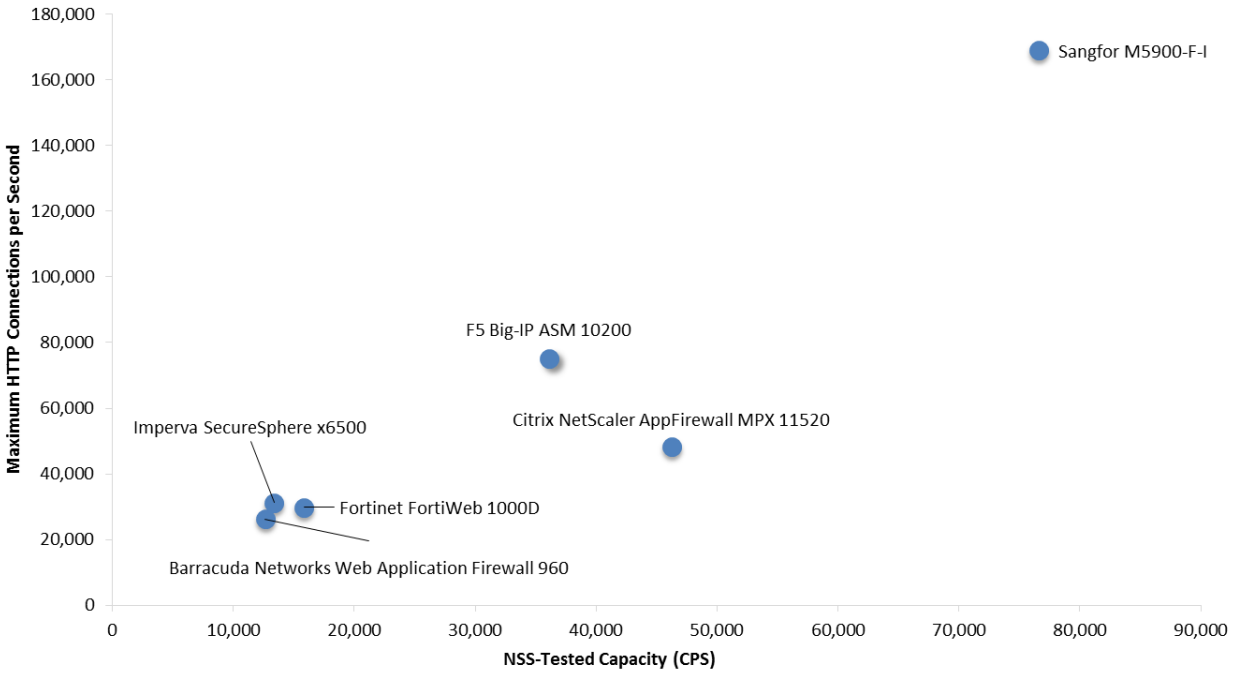

Figure 1 – Capacity and Connection Rates

Farther to the right indicates higher tested capacity. Higher up indicates higher maximum connections per second (CPS). Products with low connections/capacity ratio run the risk of exhausting connection tables before they reach their maximum potential capacity.

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

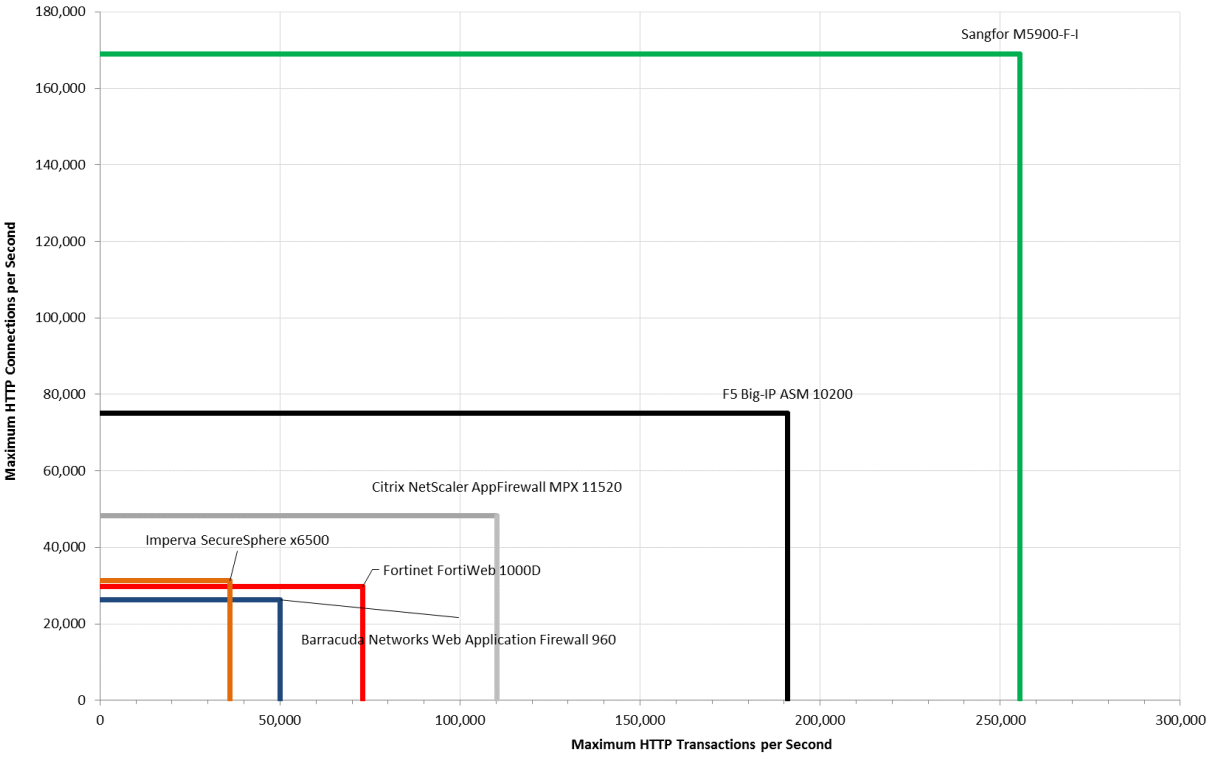

Figure 2 – Connection Dynamics

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

Table of Contents

Table of Figures

Analysis

NSS' research indicates that all enterprises tune their WAF devices. Therefore, for NSS' testing of WAF products, the devices are deployed using a tuned policy. Every effort is made to deploy policies that ensure the optimal combination of *security effectiveness* and capacity, 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 WAF *security effectiveness* and performance capabilities, based on their expected usage.

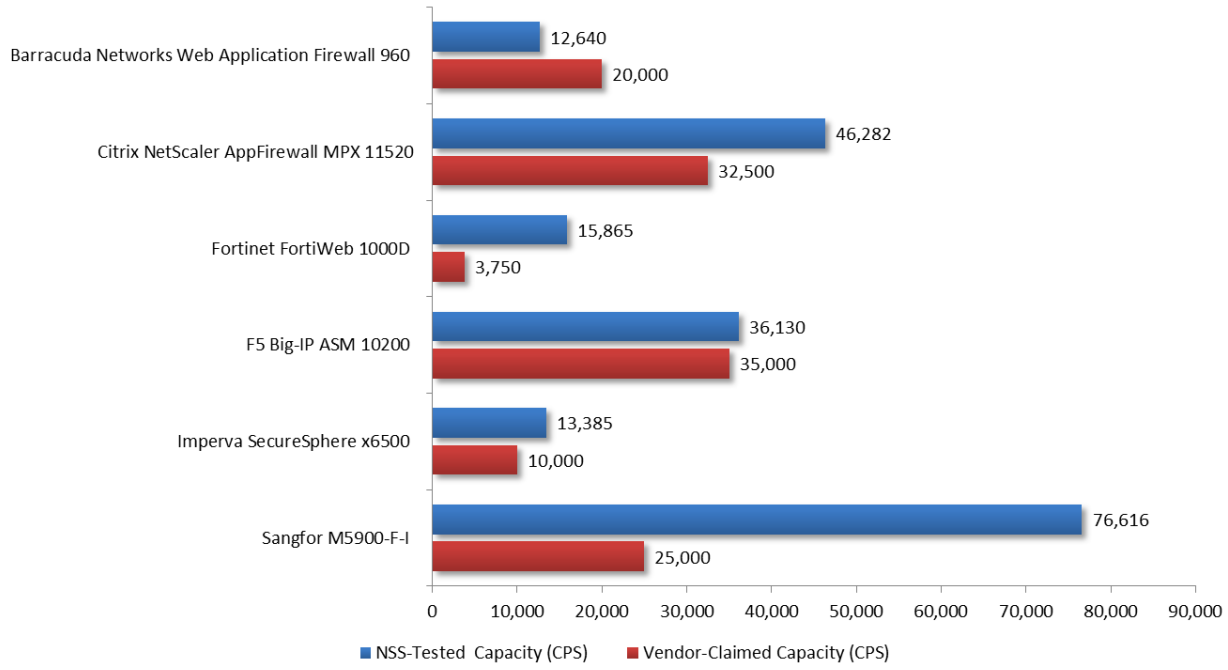

Figure 3 – Vendor-Claimed vs. NSS-Tested Capacity (CPS)

Figure 3 depicts the difference between the NSS capacity rating and the vendor capacity claims, which are often under ideal/unrealistic conditions. Where multiple figures are quoted by vendors in marketing materials, NSS selects those that relate to CPS or "with protection enabled," rather than the more optimistic UDP-only or "large packet size" performance figures often quoted.

Therefore, *NSS-tested capacity* typically is lower than that which is claimed by the vendor since it is more representative of how devices will perform in real-world deployments.

Connection Dynamics – Concurrency and Connection Rates

The aim of these tests is to stress the inspection engine and determine how it handles high volumes of 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 response time for HTTP transactions –** Latency within the WAF is causing excessive delays and increased response time to the client.
- **Unsuccessful HTTP transactions –** Normally, there should be zero unsuccessful transactions. Once these appear, it is an indication that excessive latency within the WAF is causing connections to time out.

Figure 4 depicts the key connection dynamics results from the connection dynamics tests.

Product	Maximum HTTP Connections per Second	Maximum HTTP Transactions per Second
Barracuda Networks Web Application Firewall 960	26,200	49,950
Citrix NetScaler AppFirewall MPX 11520	48,200	110,250
Fortinet FortiWeb 1000D	29,800	73,000
F5 Big-IP ASM 10200	75,000	191,000
Imperva SecureSphere x6500	31,200	36,100
Sangfor M5900-F-I	169,000	255,500

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 capacity of a system or an application. Maximum connection and transaction rates help size a device more accurately than simply examining throughput. By having knowledge of the maximum (CPS), it is possible to predict maximum capacity based on the traffic mix in a given enterprise environment. For example, if the device 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).

HTTP Connections per Second and Capacity

In-line WAF devices exhibit an inverse correlation between *security effectiveness* and capacity. The more deep-packet inspection is performed, the fewer packets can be forwarded. Furthermore, it is important to consider a real-world mix of traffic that a device will encounter.

NSS' tests aim to stress the HTTP detection engine in order to determine how the sensor copes with detecting and blocking attacks under network loads of varying average packet size and varying connections per second. By creating genuine session-based traffic with varying session lengths, the sensor is forced to track valid TCP sessions, thus ensuring a higher workload than for simple, packet-based background traffic.

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 towards HTTP traffic) at various network loads.

HTTP Connections per Second and Capacity (Capacity)

As previously stated, NSS research has found that there is usually a trade-off between *security effectiveness* and capacity. Because of this, it is important to judge a product's *security effectiveness* within the context of its capacity (and vice versa). This ensures that new security protections do not adversely impact capacity and that security shortcuts are not taken to maintain or improve capacity.

Figure 5 to Figure 9 depict the maximum capacity achieved across a range of different HTTP response sizes that may be encountered in a typical web application.

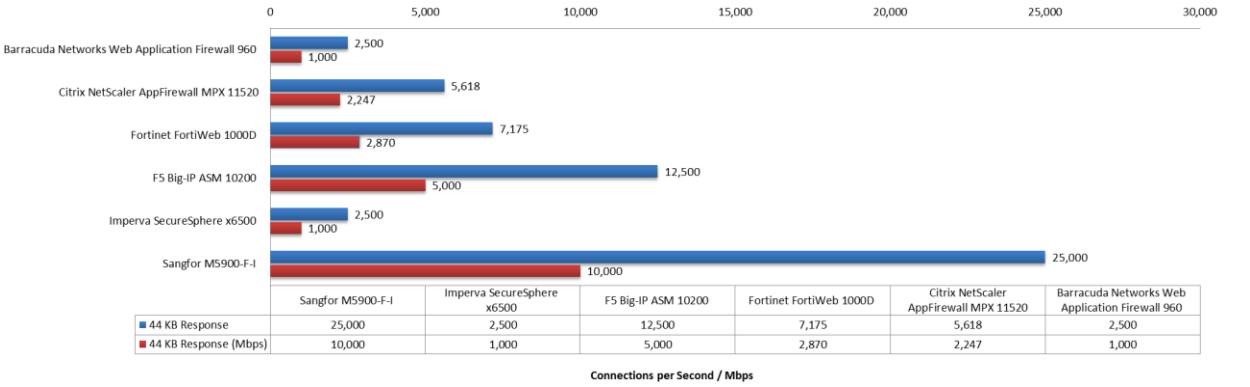

	Sangfor M5900-F-I	Imperva SecureSphere x6500	F5 Big-IP ASM 10200	Fortinet FortiWeb 1000D	Citrix NetScaler AppFirewall MPX 11520	Barracuda Networks Web Application Firewall 960
44 KB Response	25,000	2,500	12,500	7,175	5,618	2,500
44 KB Response (Mbps)	10,000	1,000	5,000	2,870	2,247	1,000

Figure 5 – Maximum Capacity per Device with 44 KB Response

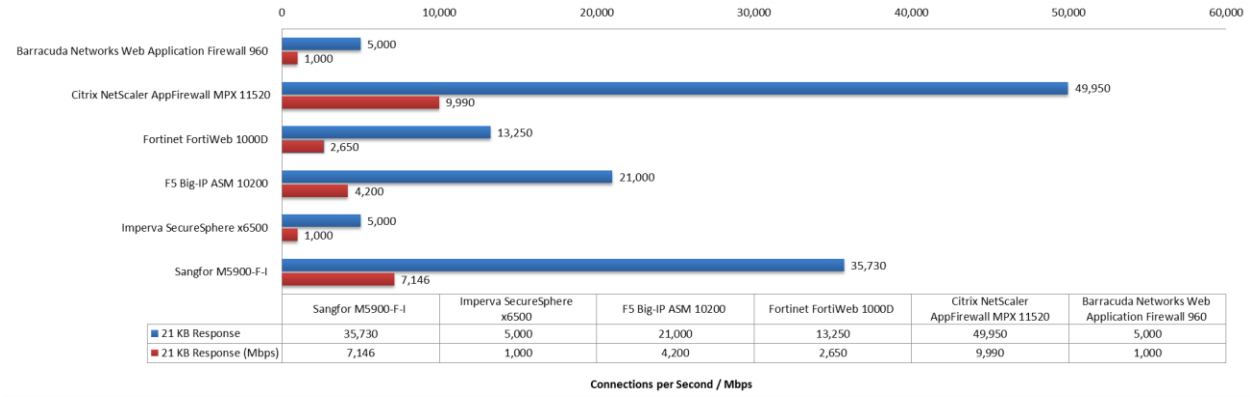

	Sangfor M5900-F-I	Imperva SecureSphere x6500	F5 Big-IP ASM 10200	Fortinet FortiWeb 1000D	Citrix NetScaler AppFirewall MPX 11520	Barracuda Networks Web Application Firewall 960
21 KB Response	35,730	5,000	21,000	13,250	49,950	5,000
21 KB Response (Mbps)	7,146	1,000	4,200	2,650	9,990	1,000

Figure 6 – Maximum Capacity per Device with 21 KB Response

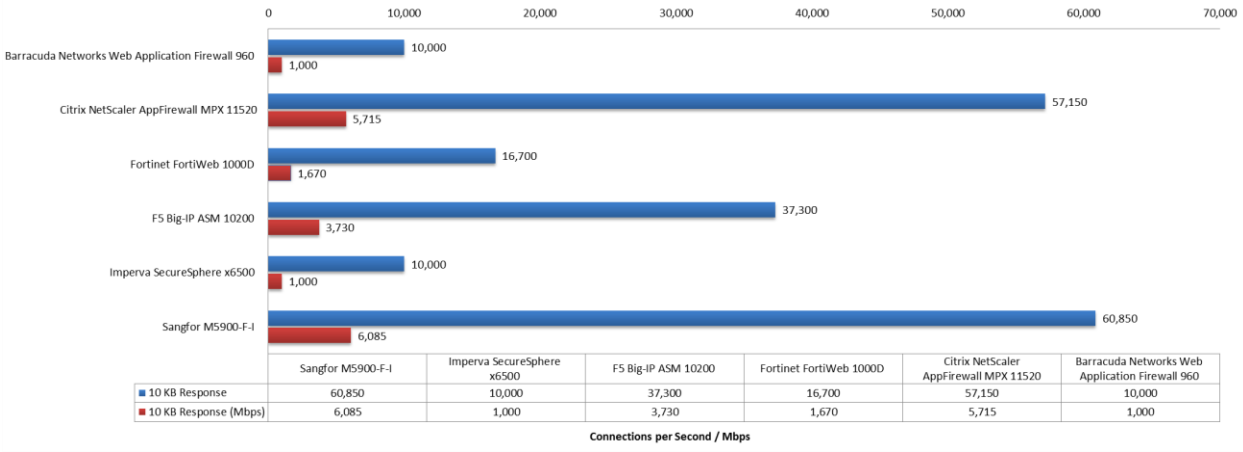

	Sangfor M5900-F-I	Imperva SecureSphere x6500	F5 Big-IP ASM 10200	Fortinet FortiWeb 1000D	Citrix NetScaler AppFirewall MPX 11520	Barracuda Networks Web Application Firewall 960
10 KB Response	60,850	10,000	37,300	16,700	57,150	10,000
10 KB Response (Mbps)	6,085	1,000	3,730	1,670	5,715	1,000

Figure 7 – Maximum Capacity per Device with 10 KB Response

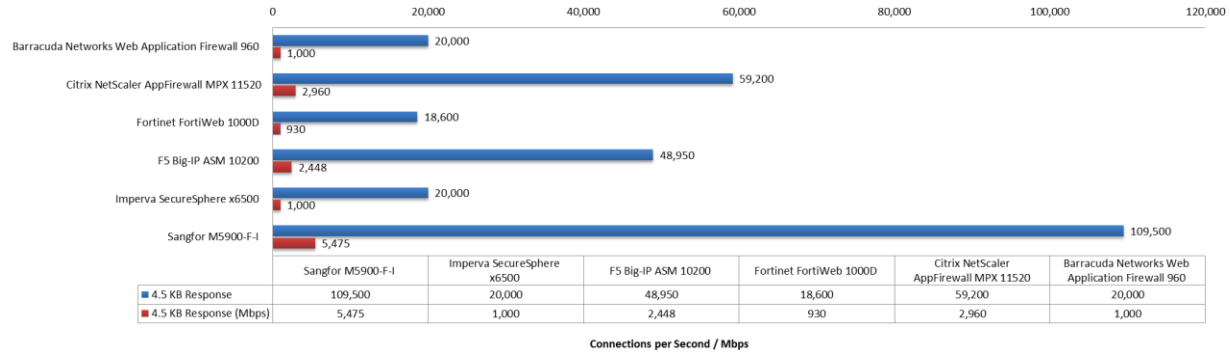

	Sangfor M5900-F-I	Imperva SecureSphere x6500	F5 Big-IP ASM 10200	Fortinet FortiWeb 1000D	Citrix NetScaler AppFirewall MPX 11520	Barracuda Networks Web Application Firewall 960
4.5 KB Response	109,500	20,000	48,950	18,600	59,200	20,000
4.5 KB Response (Mbps)	5,475	1,000	2,448	930	2,960	1,000

Figure 8 – Maximum Capacity per Device with 4.5 KB Response

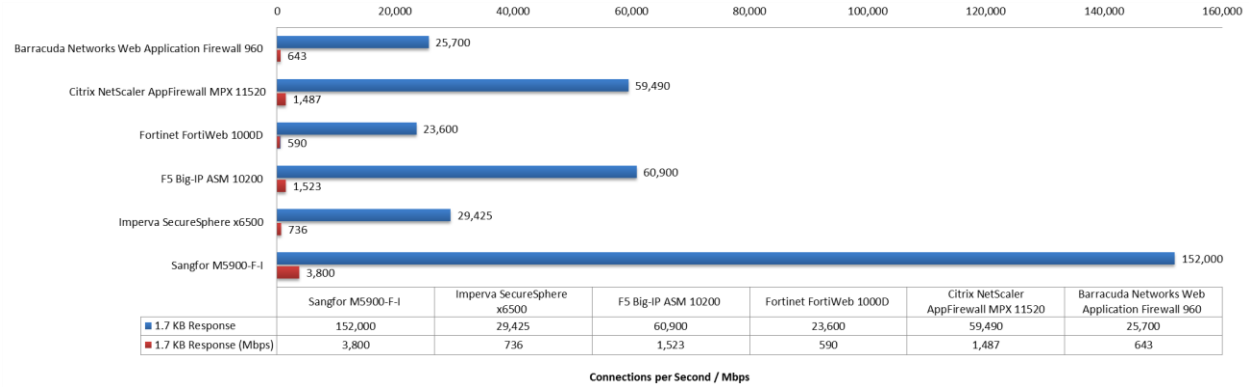

	Sangfor M5900-F-I	Imperva SecureSphere x6500	F5 Big-IP ASM 10200	Fortinet FortiWeb 1000D	Citrix NetScaler AppFirewall MPX 11520	Barracuda Networks Web Application Firewall 960
1.7 KB Response	152,000	29,425	60,900	23,600	59,490	25,700
1.7 KB Response (Mbps)	3,800	736	1,523	590	1,487	643

Figure 9 – Maximum Capacity per Device with 1.7 KB Response

Figure 10 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
Barracuda Networks Web Application Firewall 960	2,500	5,000	10,000	20,000	25,700
Citrix NetScaler AppFirewall MPX 11520	5,618	49,950	57,150	59,200	59,490
Fortinet FortiWeb 1000D	7,175	13,250	16,700	18,600	23,600
F5 Big-IP ASM 10200	12,500	21,000	37,300	48,950	60,900
Imperva SecureSphere x6500	2,500	5,000	10,000	20,000	29,425
Sangfor M5900-F-I	25,000	35,730	60,850	109,500	152,000

Figure 10 – Maximum Connection Rates per Device with Various Response Sizes

Test Methodology

Web Application Firewall: v6.2

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 Rd
Building A, Suite 200
Austin, TX 78746
info@nsslabs.com
www.nsslabs.com