WEB APPLICATION FIREWALL PRODUCT ANALYSIS

F5 Big-IP ASM 10200 v11.4.0

Authors – Ryan Liles, Orlando Barrera

Overview

NSS Labs performed an independent test of the F5 Big-IP ASM 10200. The product was subjected to thorough testing at the NSS facility in Austin, Texas, based on the *Web Application Firewall Methodology v6.2* available at www.nsslabs.com. This test was conducted free of charge, and NSS did not receive any compensation in return for F5's participation.

While the companion Comparative Analysis Reports (CARs) on security, performance, and total cost of ownership (TCO) will provide comparative information about all tested products, this individual Product Analysis Report (PAR) provides detailed information not available elsewhere.

NSS testing has found that the majority of web application firewalls (WAFs) operate in an adaptive learning mode ("learning mode"). In this mode, a WAF learns the behavior of applications and automatically generates policy recommendations. These recommendations require review and approval before the WAF device is deployed. Periodic manual tuning may also be required.

As part of the initial WAF test setup, devices are tuned as deemed necessary 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 WAF security effectiveness and performance capabilities based upon their expected usage.

Product	Block Rate[1]	NSS-Tested Capacity
F5 Big-IP ASM 10200 V11.4.0	99.89%	36,130 CPS
Evasions	False Positives	Stability & Reliability
PASS	0.124%	PASS

Figure 1 – Overall Test Results

Using a tuned policy, the Big-IP ASM 10200 blocked 99.21% of WAF attacks. The device proved effective against all evasion techniques tested. The device also passed all stability and reliability tests. The Big-IP ASM 10200 presented a 0.124% false positive rate

The F5 Big-IP ASM 10200 is rated by NSS at 36,130 connections per second (CPS), which is in line with the vendor-claimed performance. This is a minimum rating using one transaction per connection. F5 rates this device at 35,000 CPS. NSS-tested capacity is an average of all of the HTTP response-based capacity tests. These performance numbers represent a baseline which you can use to model your environment.

[1] Block rate is defined as the number of attacks blocked under test.

Table of Contents

Table of Figures

Security Effectiveness

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

Attack Types

In order to represent accurately the protection that is likely to be achieved by a typical enterprise, NSS evaluates the DUT using the vendor's optimally tuned policy.

The NSS threat and attack suite contains thousands of publically available exploits (including multiple variants of each exploit) and a number of complex web applications that have been constructed to include known vulnerabilities and coding errors. Groups of exploits are carefully selected from this library to test based on the intended attack type as listed below. Each exploit has been validated to impact the target vulnerable host(s) by compromising either the underlying OS, the web server, or the web application itself.

Attack types:

- **URL Parameter Manipulation –** altering URL data to gain potentially protected information or access protected areas of a website.
- **Form/Hidden Field Manipulation —** constructing POST requests to access protected information or protected areas of a website, or to manipulate "fixed" data directly (such as pricing information).
- **Cookie/Session Poisoning —** manipulation of cookie or session variables to access protected information or protected areas of a website.
- **Cross-Site Scripting (XSS) —** the process of manipulating user input in such a way that, when rendered in the context of a webpage, it will be interpreted by the browser as code.
- **Directory traversal —** altering the URL to access areas of the web server that should not otherwise be accessible.
- **SQL Injection —** manipulating user input in such a way that, when processed by the database server, it will be interpreted as code, potentially providing direct access to private data.
- **Padding Oracle attacks —** altering a block-cypher cryptographic hash in such a way as to decrypt encrypted information.

Test Procedure	Results
URL Parameter Manipulation	100%
Form/Hidden Field Manipulation	100%
Cookie/Session Poisoning	100%
Cross-Site Scripting (XSS)	99.25%
Directory Traversal	100%
SQL Injection	100%
Padding Oracle Attacks	100%

Figure 2 – Attack Types

Resistance to Evasion Techniques

Evasion techniques disguise and modify attacks at the point of delivery in order to avoid detection and blocking by security products. Missing a particular type of evasion means an attacker can use an entire class of exploits for which a device is supposed to have protection, rendering it virtually useless. Many of the techniques used in this test have been widely known for years and should be considered minimum requirements for the WAF product category.

Providing exploit protection results without fully factoring in evasion can be misleading since the more *different* types of evasion that are missed – packet fragmentation reassembly, stream segmentation, URL obfuscation and normalization – the worse the situation. For example, it is better to miss all techniques in one evasion category (say, stream segmentation) than one technique in each category.

Figure 3 provides the results of the evasion tests for F5 Big-IP ASM 10200.

Test Procedure	Results
Packet Fragmentation Reassembly	PASS
Stream Segmentation	PASS
URL Obfuscation and Normalization	PASS

Figure 3 – Resistance to Evasion Results

The device proved effective against all evasion techniques tested. This resulted in an overall PASS result for F5 Big-IP ASM 10200.

False Positive Testing

The ability of the DUT to identify and allow legitimate traffic while maintaining protection against attacks and exploits is of equal importance to providing protection against malicious content. This test includes a varied sample of legitimate application traffic, which should properly be identified and allowed.

Figure 4 shows the percentage of non-malicious traffic mistakenly identified as malicious (lower score is better). F5 Big-IP ASM 10200 had a false positive rate of 0.124%.

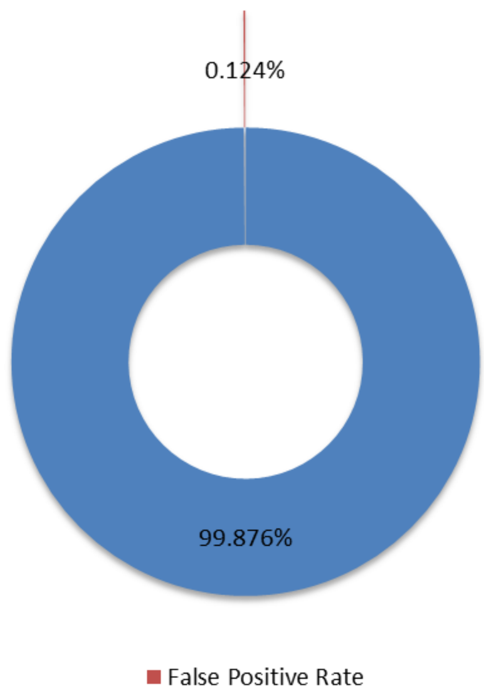

Figure 4 – False Positives

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 security shortcuts are not taken to maintain or improve performance.

Connection Dynamics – Concurrency and Connection Rates

The use of sophisticated test equipment appliances allows NSS engineers to create true "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 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 DUT 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 DUT is causing connections to time out.

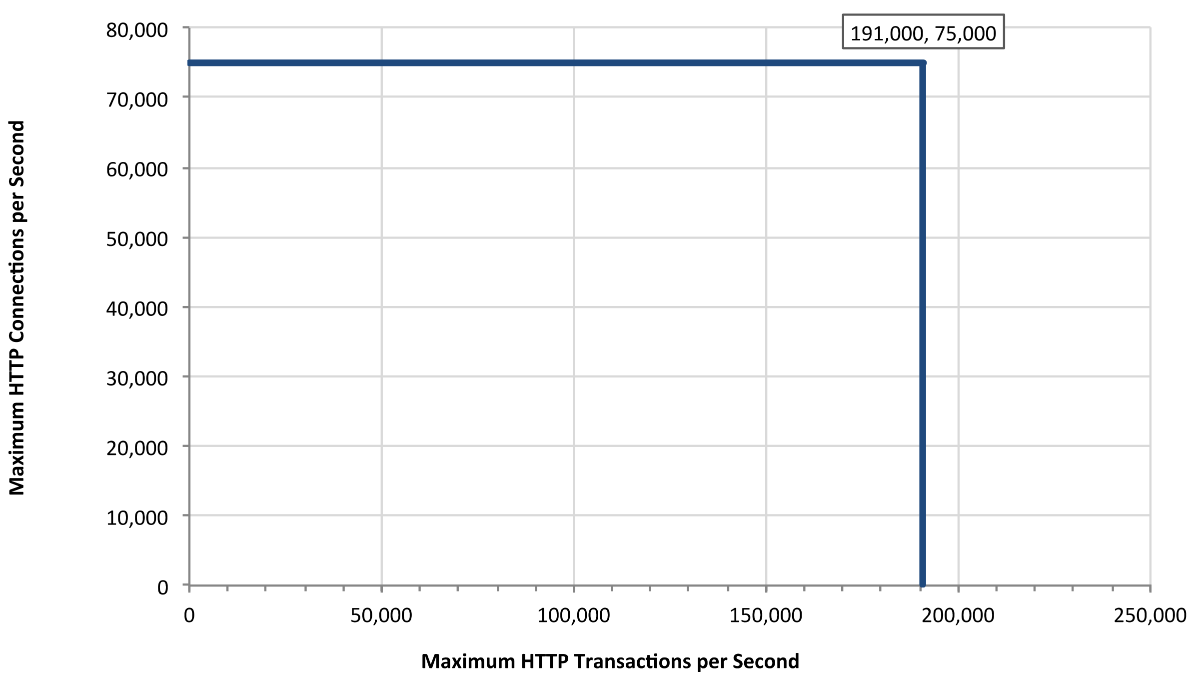

Figure 5 – Concurrency and Connection Rates

HTTP Connections Per Second and Capacity

The aim of these tests is to stress the HTTP detection engine and determine how the DUT 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 DUT 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” as it is possible to achieve in a lab environment, 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 a good benchmark representation of a live network (albeit one biased towards HTTP traffic) at various network loads. In real life scenarios, browsers may use one connection to send multiple HTTP requests resulting in enhanced throughput for the systems.

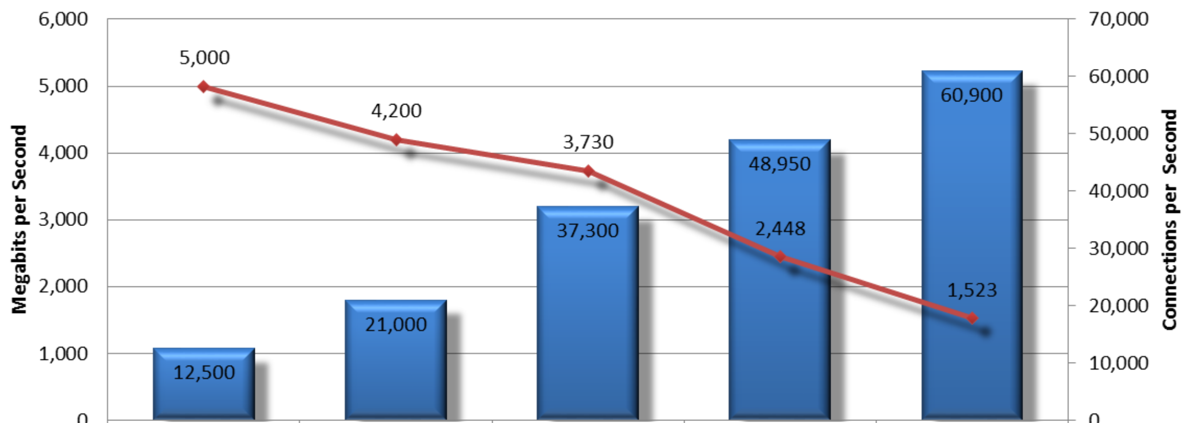

	44 KB Response	21 KB Response	10 KB Response	4.5 KB Response	1.7 KB Response
CPS	12,500	21,000	37,300	48,950	60,900
Mbps	5,000	4,200	3,730	2,448	1,523

Figure 6 – HTTP Connections per Second and Capacity

HTTP Connections Per Second and Capacity (With Delays)

Typical user behavior introduces delays between requests and responses, e.g., “think time,” as users read web pages and decide which links to click next. This next set of tests is identical to the previous set except that these include a 5-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 DUT to utilize additional resources to track those connections.

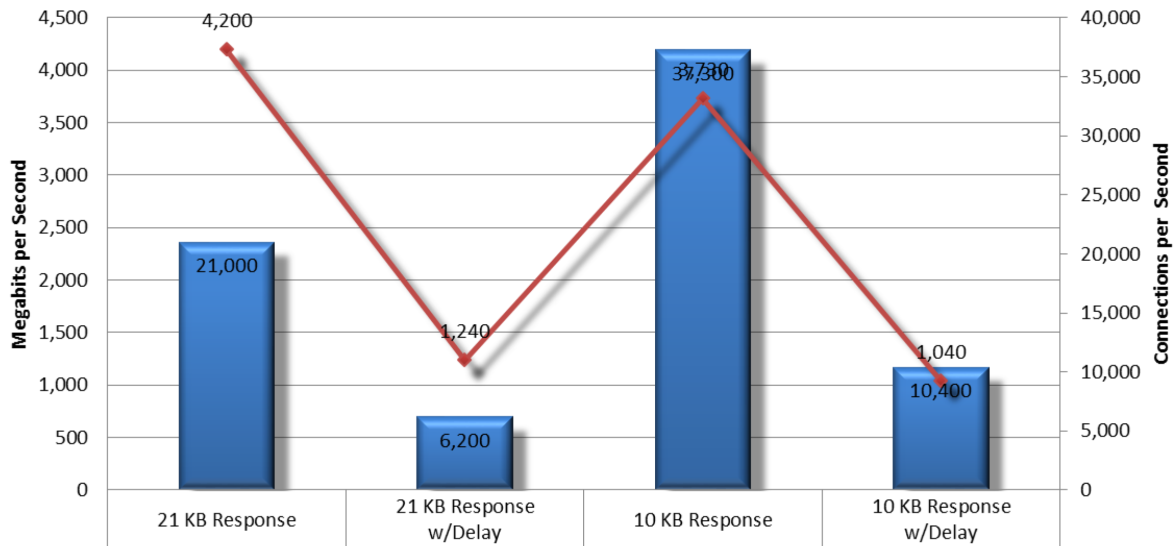

	21 KB Response	21 KB Response w/Delay	10 KB Response	10 KB Response w/Delay
CPS	21,000	6,200	37,300	10,400
Mbps	4,200	1,240	3,730	1,040

Figure 7 – HTTP Connections per Second and Capacity (With Delays)

The F5 Big-IP ASM 10200 demonstrated an unusual drop in throughput when a 5000ms delay was introduced to the test traffic.

Stability and Reliability

Long-term stability is particularly important for an in-line device, where failure can produce network outages. These tests verify the stability of the DUT along with its ability to maintain security effectiveness while under normal load and while passing malicious traffic. Products that are not able to sustain legitimate traffic (or that crash) while under hostile attack will not pass.

The Big-IP ASM 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 by either the volume of traffic or the DUT failing open for any reason, this will result in a FAIL.

Test Procedure	Result
Blocking Under Extended Attack	PASS
Passing Legitimate Traffic Under Extended Attack	PASS
Protocol Fuzzing & Mutation	
Protocol Fuzzing & Mutation – Detection Ports	PASS
Protocol Fuzzing & Mutation – Management Port	PASS
Power Fail	PASS
Redundancy	YES
Persistence of Data	PASS

Figure 8 – Stability and Reliability Results

These tests also determine the behavior of the state engine under load. All WAF devices must choose whether to risk denying legitimate traffic or allowing malicious traffic once they run low on resources. Dropping new connections when resources (such as state table memory) are low, or when traffic loads exceed the device capacity will theoretically block legitimate traffic but maintain state on existing connections (preventing attack leakage).

Management and Configuration

Security devices are complicated to deploy; essential systems such as centralized management console options, log aggregation, and event correlation/management systems further complicate the purchasing decision.

Understanding key comparison points will allow customers to model the overall impact on network service level agreements (SLAs); estimate operational resource requirements to maintain and manage the systems; and better evaluate the required skill/competencies of staff. Enterprises should include management and configuration during their evaluation, focusing on the following at a minimum:

- **General Management and Configuration –** How easy is it to install, configure, and deploy multiple devices throughout a large enterprise network?
- **Policy Handling –** How easy is it to create, edit, and deploy complicated security policies across an enterprise?
- **Alert Handling –** How accurate and timely is the alerting, and how easy is it to drill down to locate critical information needed to remediate a security problem?
- **Reporting –** How effective is the reporting capability, and how readily can it be customized?

Total Cost of Ownership (TCO)

Implementation of security solutions can be complex, with several factors affecting the overall cost of deployment, maintenance, and upkeep. All of these should be considered over the course of the useful life of the solution.

- **Product Purchase –** The cost of acquisition
- **Product Maintenance –** The fees paid to the vendor (including software and hardware support, maintenance, and other updates)
- **Installation –** The time required to take the device out of the box, configure it, put it into the network, apply updates and patches, and set up desired logging and reporting
- **Upkeep –** The time required to apply periodic updates and patches from vendors, including hardware, software, and other updates
- **Management –** Day-to-day management tasks including device configuration, policy updates, policy deployment, alert handling, and so on

For the purposes of this report, capital expenditure (capex) items are included for a single device only (the cost of acquisition and installation).

Installation (Hours)

Figure 9 provides the number of hours of labor required to install each device using local device management options only. This will reflect accurately the amount of time taken for NSS engineers, with the help of vendor engineers, to install and configure the DUT to the point where it operates successfully in the test harness, passes legitimate traffic, and blocks/detects prohibited/malicious traffic. This closely mimics a typical enterprise deployment scenario for a single device.

Costs are based upon the time that would be required by an experienced security engineer to perform the abovementioned tasks (assumed US $75 per hour for the purposes of these calculations), allowing 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)
F5 Big-IP ASM 10200 V11.4.0	8

Figure 9 – Sensor Installation Time in Hours

Purchase Price and 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 reflect single device management and maintenance only; costs for central management solutions (CMS) may be extra. For additional TCO analysis, refer to the *TCO CAR*.

Product	Purchase	Maintenance/ Year	Year 1 Cost	Year 2 Cost	Year 3 Cost	3-Year TCO
F5 Big-IP ASM 10200 V11.4.0	$84,995	$10,709	$96,304	$10,709	$10,709	$117,722

Figure 10 – 3-Year TCO

- **Year 1 Cost** is calculated by adding installation costs ($75 USD 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**.**

This provides a TCO figure consisting of hardware, installation, and maintenance costs for a single device only: costs for CMS may be extra. For additional TCO analysis, refer to the *TCO CAR*.

Value: Total Cost of Ownership Per Protected-CPS

There is a clear difference between price and value. The least expensive product does not necessarily offer the greatest value if it offers significantly lower performance than only slightly more expensive competitors. The best value is a product with a low TCO and high level of secure capacity (Block Rate x NSS-Tested Capacity).

Figure 11 depicts the relative cost per unit of work performed, described as TCO per Protected-CPS.

Product	Block Rate	NSS-Tested Capacity	3-Year TCO	TCO per Protected-CPS
F5 Big-IP ASM 10200 V11.4.0	99.89%	36,130 CPS	$117,722	$3

Figure 11 – Total Cost of Ownership per Protected-CPS

TCO per Protected-CPS was calculated by taking the 3-Year TCO and dividing it by the product of Block Rate x NSS-Tested Capacity. Therefore, 3-Year TCO / (Block Rate x NSS-Tested Capacity) = TCO per Protected-CPS.

Detailed Product Scorecard

The following chart depicts the status of each test with quantitative results where applicable.

Description	Result
Security Effectiveness	
Attack Types	99.89%
URL Parameter Manipulation	100%
Form/Hidden Field Manipulation	100%
Cookie/Session Poisoning	100%
Cross-Site Scripting (XSS)	99.25%
Directory Traversal	100%
SQL Injection	100%
Padding Oracle attacks	100%
Evasions and Attack Leakage	
Resistance to Evasion	100%
Packet Fragmentation Reassembly	100%
Ordered 8 byte fragments	100%
Ordered 16 byte fragments	100%
Ordered 24 byte fragments	100%
Ordered 32 byte fragments	100%
Out of order 8 byte fragments	100%
Ordered 8 byte fragments, duplicate last packet	100%
Out of order 8 byte fragments, duplicate last packet	100%
Ordered 8 byte fragments, reorder fragments in reverse	100%
Ordered 16 byte fragments, fragment overlap (favor new)	100%
Ordered 16 byte fragments, fragment overlap (favor old)	100%
Out of order 8 byte fragments, interleaved duplicate packets scheduled for later delivery	100%
Ordered 8 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
Ordered 16 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
Ordered 24 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%
Ordered 32 byte fragments, duplicate packet with an incrementing DWORD in the options field. The duplicate packet has random payload.	100%

Stream Segmentation	100%
Ordered 1 byte segments, interleaved duplicate segments with null TCP control flags	100%
Ordered 1 byte segments, interleaved duplicate segments with null TCP control flags	100%
Ordered 1 byte segments, interleaved duplicate segments with requests to resync sequence numbers mid-stream	100%
Ordered 1 byte segments, duplicate last packet	100%
Ordered 2 byte segments, segment overlap (favor new)	100%
Ordered 1 byte segments, interleaved duplicate segments with out-of-window sequence numbers	100%
Out of order 1 byte segments	100%
Out of order 1 byte segments, interleaved duplicate segments with faked retransmits	100%
Ordered 1 byte segments, segment overlap (favor new)	100%
Out of order 1 byte segments, PAWS elimination (interleaved duplicate segments with older TCP timestamp options)	100%
Ordered 16 byte segments, segment overlap (favor new (Unix))	100%
Ordered 32 byte segments	100%
Ordered 64 byte segments	100%
Ordered 128 byte segments	100%
Ordered 256 byte segments	100%
Ordered 512 byte segments	100%
Ordered 1024 byte segments	100%
Ordered 2048 byte segments (sending MSRPC request with exploit)	100%
Reverse Ordered 256 byte segments, segment overlap (favor new) with random data	100%
Reverse Ordered 512 byte segments, segment overlap (favor new) with random data	100%
Reverse Ordered 1024 byte segments, segment overlap (favor new) with random data	100%
Reverse Ordered 2048 byte segments, segment overlap (favor new) with random data	100%
Out of order 1024 byte segments, segment overlap (favor new) with random data, Initial TCP sequence number is set to 0xffffffff - 4294967295	100%
Out of order 2048 byte segments, segment overlap (favor new) with random data, Initial TCP sequence number is set to 0xffffffff - 4294967295	100%
URL Obfuscation And Normalization	100%
URL encoding - Level 1 (minimal)	100%
URL encoding - Level 2	100%
URL encoding - Level 3	100%
URL encoding - Level 4	100%
URL encoding - Level 5	100%
URL encoding - Level 6	100%
URL encoding - Level 7	100%
URL encoding - Level 8 (extreme)	100%
Directory Insertion	100%
Premature URL ending	100%
TAB separation	100%
Windows\delimiter	100%
Base-64 Encoding	100%
Base-64 Encoding (shifting 1 bit)	100%
Base-64 Encoding (shifting 2 bits)	100%
Base-64 Encoding (chaffing)	100%

False Positives	0.124%
Performance	
Maximum Capacity	
Maximum HTTP Connections per Second	75,000
Maximum HTTP Transactions per Second	191,000
HTTP Capacity with no Transaction Delays	
2,500 Connections per Second – 44 Kbyte Response	12,500
5,000 Connections per Second – 21 Kbyte Response	21,000
10,000 Connections per Second – 10 Kbyte Response	37,300
20,000 Connections per Second – 4.5 Kbyte Response	48,950
40,000 Connections per Second – 1.7 Kbyte Response	60,900
HTTP CPS & Capacity With Transaction Delays	
21 Kbyte Response with Delay	6,200
10 Kbyte Response with Delay	10,400
Stability & Reliability	
Blocking Under Extended Attack	PASS
Passing Legitimate Traffic Under Extended Attack	PASS
Protocol Fuzzing & Mutation	
Protocol Fuzzing & Mutation – Detection Ports	PASS
Protocol Fuzzing & Mutation – Management Port	PASS
Power Fail	PASS
Redundancy	YES
Persistence of Data	PASS
Total Cost of Ownership	
Ease of Use	
Initial Setup (Hours)	8
Time Required for Upkeep (Hours per Year)	***Contact NSS***
Expected Costs	
Initial Purchase (hardware as tested)	$84,995
Initial Purchase (enterprise management system)	***Contact NSS***
Annual Cost of Maintenance & Support (hardware/software)	$10,709
Annual Cost of Maintenance & Support (enterprise management system)	***Contact NSS***
Installation Labor Cost (@ US$75/hr)	$600.00
Management Labor Cost (per Year @ US$75/hr)	***Contact NSS***
Total Cost of Ownership (TCO)	
Year 1	$96,304
Year 2	$10,709
Year 3	$10,709
3-Year TCO	$117,722

Figure 12 – Detailed Scorecard

Test Methodology

Web Application Firewall (WAF): 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