



NEXT GENERATION FIREWALL COMPARATIVE REPORT

Performance

JULY 17, 2018

Author – Thomas Skybakmoen

Tested Products

Barracuda Networks CloudGen Firewall F800.CCE v7.2.0

Check Point 15600 Next Generation Threat Prevention (NGTP) Appliance vR80.20

Cisco Firepower 4120 Security Appliance v6.2.2

Forcepoint NGFW 2105 Appliance v6.3.3 build 19153 (Update Package: 1056)

Fortinet FortiGate 500E V5.6.3GA build 7858

Palo Alto Networks PA-5220 PAN-OS 8.1.1

SonicWall NSa 2650 SonicOS Enhanced 6.5.0.10-73n

Sophos XG Firewall 750 SFOS v17 MR7

Versa Networks FlexVNF 16.1R1-S6

WatchGuard M670 v12.0.1.B562953

Environment

NSS Labs Next Generation Firewall Test Methodology v8.0

NSS Labs SSL/TLS Performance Test Methodology v1.3

NSS Labs Evasions Test Methodology v1.1

Overview

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

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

- 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 *NSS-Tested Throughput* and *Maximum TCP Connections per Second*. NSS-tested throughput is calculated as a weighted average of the traffic that NSS Labs expects an NGFW to experience in an enterprise environment. For more details, please see the Scorecard section in the individual test reports.




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 an 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 |
| SSL/TLS Performance | 15 |
| Single Application Flows..... | 15 |
| Test Methodology | 20 |
| Contact Information | 20 |

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 –Single Application Flow: Telephony (Mbps)..... | 15 |
| Figure 20 –Single Application Flow: Financial (Mbps) | 15 |
| Figure 21 –Single Application Flow: Email (Mbps)..... | 16 |
| Figure 22 –Single Application Flow: File Sharing (Mbps)..... | 16 |
| Figure 23 –Single Application Flow: File Server (Mbps)..... | 17 |
| Figure 24 –Single Application Flow: Remote Console (Mbps) | 17 |
| Figure 25 –Single Application Flow: Video (Mbps) | 18 |
| Figure 26 –Single Application Flow: Meeting (Mbps)..... | 18 |
| Figure 27 –Single Application Flow: Database (Mbps) | 19 |

Analysis

NSS research indicates that NGFWs 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.

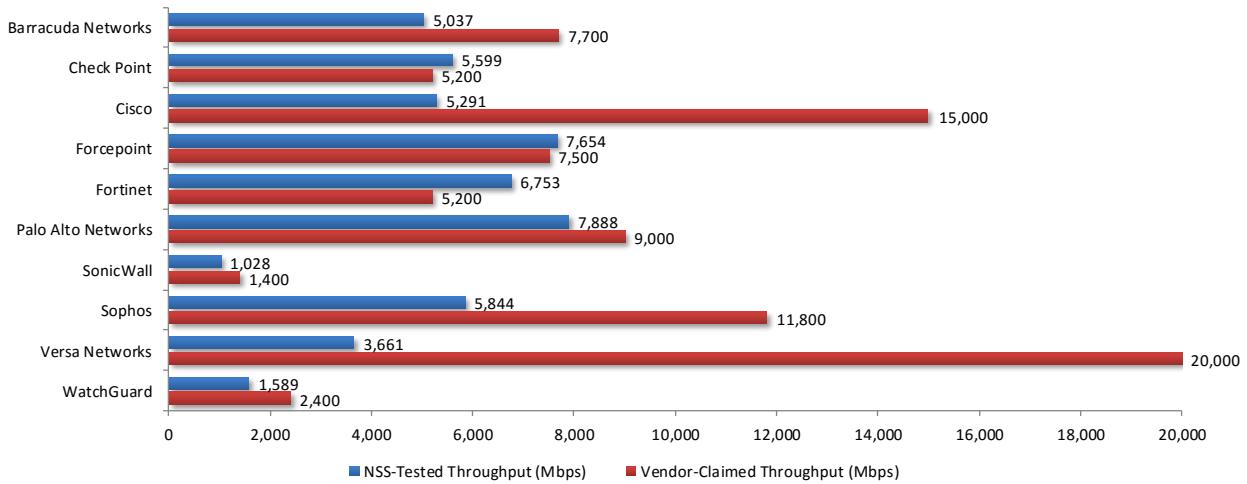


Figure 3 – Vendor-Claimed Throughput vs. NSS-Tested Throughput (Mbps)

Figure 3 depicts the difference between *NSS-Tested Throughput*¹ 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 along with variable source and destination IP addresses 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 condition. The aim of the test is to determine the raw packet processing capability of each inline port pair of the device as well as the device’s effectiveness at

¹ *NSS-Tested Throughput* is calculated as a weighted average of the traffic that NSS expects an NGFW to experience in an enterprise environment. For more details, please see the Scorecard section in the individual test reports.

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.

| Vendor | Throughput (Mbps) | | | | | |
|--------------------|-------------------|------------------|------------------|------------------|-------------------|-------------------|
| | 64-Byte Packets | 128-Byte Packets | 256-Byte Packets | 512-Byte Packets | 1024-Byte Packets | 1514-Byte Packets |
| Barracuda Networks | 1,638 | 2,688 | 4,337 | 6,131 | 8,030 | 8,979 |
| Check Point | 1,141 | 1,749 | 2,649 | 3,787 | 4,641 | 7,723 |
| Cisco | 1,341 | 2,041 | 3,238 | 4,783 | 7,230 | 8,679 |
| Forcepoint | 2,638 | 3,836 | 5,735 | 8,430 | 11,470 | 14,120 |
| Fortinet | 20,000 | 20,000 | 20,000 | 20,000 | 20,000 | 20,000 |
| Palo Alto Networks | 17,610 | 18,260 | 18,660 | 18,910 | 20,000 | 20,000 |
| SonicWall | 151 | 259 | 486 | 944 | 1,816 | 2,757 |
| Sophos | 2,261 | 3,829 | 7,705 | 13,180 | 16,670 | 16,750 |
| Versa Networks | 2,038 | 3,486 | 6,132 | 11,170 | 20,000 | 20,000 |
| WatchGuard | 740 | 1,440 | 2,637 | 4,882 | 9,027 | 12,920 |

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.

| Vendor | Latency (μ s) | | | | | |
|--------------------|--------------------|------------------|------------------|------------------|-------------------|-------------------|
| | 64-Byte Packets | 128-Byte Packets | 256-Byte Packets | 512-Byte Packets | 1024-Byte Packets | 1514-Byte Packets |
| Barracuda Networks | 76.26 | 79.07 | 80.37 | 98.33 | 77.67 | 99.99 |
| Check Point | 23.00 | 26.00 | 40.00 | 44.00 | 36.00 | 82.00 |
| Cisco | 94.72 | 108.55 | 94.51 | 92.73 | 107.96 | 111.89 |
| Forcepoint | 72.01 | 69.18 | 80.79 | 101.29 | 117.18 | 177.28 |
| Fortinet | 6.84 | 6.88 | 7.16 | 7.54 | 8.92 | 10.04 |
| Palo Alto Networks | 13.00 | 14.00 | 14.00 | 15.00 | 19.00 | 20.00 |
| SonicWall | 18.68 | 26.84 | 21.96 | 26.52 | 33.46 | 42.65 |
| Sophos | 162.11 | 166.43 | 175.19 | 174.71 | 163.00 | 145.50 |
| Versa Networks | 75.56 | 77.10 | 80.19 | 83.94 | 113.10 | 123.73 |
| WatchGuard | 44.88 | 83.65 | 86.96 | 106.71 | 125.71 | 123.52 |

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

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 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.

| Vendor | Maximum Concurrent TCP Connections | Maximum TCP CPS | Maximum HTTP CPS | Maximum HTTP Transactions per Second |
|--------------------|------------------------------------|-----------------|------------------|--------------------------------------|
| Barracuda Networks | 2,374,976 | 131,500 | 22,980 | 33,700 |
| Check Point | 5,890,170 | 46,000 | 11,960 | 119,500 |
| Cisco | 14,205,332 | 172,800 | 61,760 | 553,400 |
| Forcepoint | 30,000,000 | 90,500 | 45,320 | 100,300 |
| Fortinet | 3,000,000 | 85,000 | 69,790 | 170,500 |
| Palo Alto Networks | 2,639,904 | 38,500 | 51,000 | 71,990 |
| SonicWall | 250,000 | 10,010 | 9,971 | 34,060 |
| Sophos | 10,485,744 | 181,000 | 9,200 | 24,720 |
| Versa Networks | 999,999 | 124,900 | 28,500 | 36,410 |
| WatchGuard | 327,341 | 13,500 | 4,404 | 6,808 |

Figure 7 – Concurrency and Connection Rates (I)

In addition to overall throughput, 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. Once a device’s maximum CPS is known, 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., $(2,000/2,500) \times 1,000 \text{ Mbps} = 800 \text{ Mbps}$).

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.

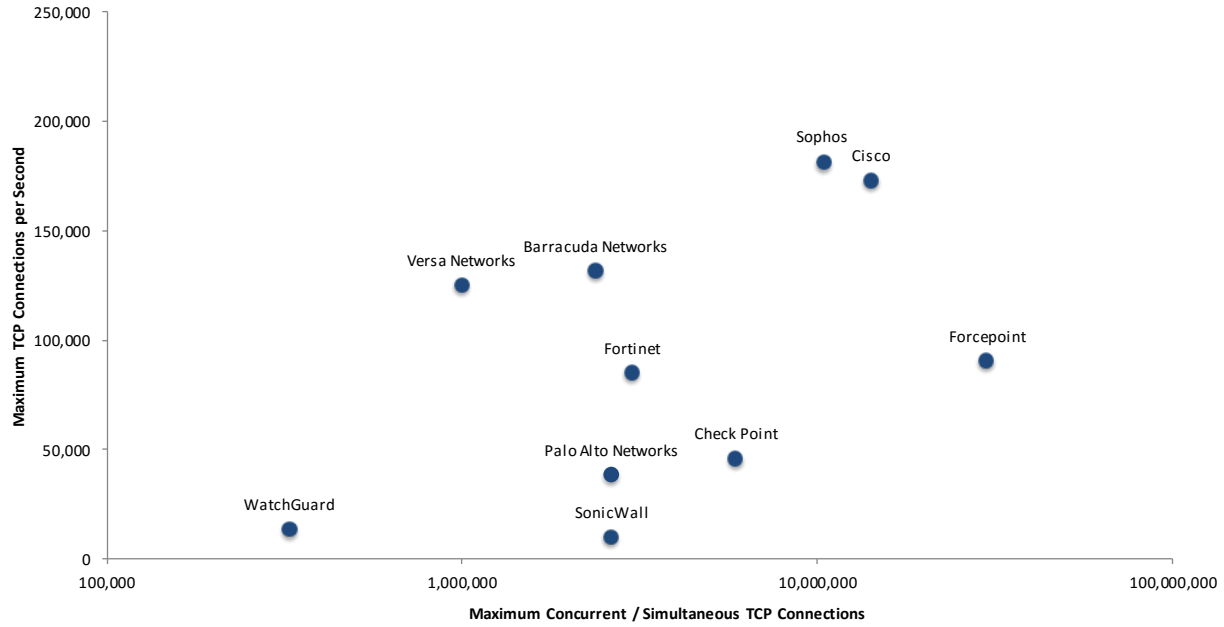


Figure 8 – Concurrency and Connection Rates (II)

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 multiple tests using genuine session-based traffic with varying session lengths, the device is forced to track valid HTTP 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).

| Vendor | 44 KB Response Size | 21 KB Response Size | 10 KB Response Size | 4.5 KB Response Size | 1.7 KB Response Size |
|--------------------|---------------------|---------------------|---------------------|----------------------|----------------------|
| Barracuda Networks | 11,270 | 12,490 | 14,800 | 18,940 | 21,540 |
| Check Point | 25,000 | 47,491 | 62,800 | 47,600 | 51,200 |
| Cisco | 25,000 | 50,000 | 54,200 | 56,060 | 58,600 |
| Forcepoint | 25,000 | 41,000 | 55,310 | 63,730 | 64,340 |
| Fortinet | 19,290 | 29,100 | 42,930 | 51,200 | 60,310 |
| Palo Alto Networks | 25,000 | 43,160 | 44,670 | 56,570 | 46,900 |
| SonicWall | 2,480 | 3,926 | 5,697 | 7,041 | 8,001 |
| Sophos | 5,994 | 8,600 | 7,800 | 9,819 | 9,972 |
| Versa Networks | 5,338 | 7,164 | 10,500 | 16,550 | 22,400 |
| WatchGuard | 2,000 | 2,850 | 3,700 | 3,647 | 4,383 |

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.

| Vendor | 44 KB Latency (ms) | 21 KB Latency (ms) | 10 KB Latency (ms) | 4.5 KB Latency (ms) | 1.7 KB Latency (ms) |
|--------------------|--------------------|--------------------|--------------------|---------------------|---------------------|
| Barracuda Networks | 2.75 | 2.08 | 1.44 | 1.12 | 0.96 |
| Check Point | 2.30 | 2.10 | 2.00 | 1.80 | 2.00 |
| Cisco | 1.60 | 1.67 | 1.24 | 1.14 | 1.07 |
| Forcepoint | 2.69 | 2.48 | 1.95 | 1.32 | 1.00 |
| Fortinet | 2.35 | 1.54 | 1.19 | 0.79 | 0.74 |
| Palo Alto Networks | 1.50 | 1.30 | 1.08 | 2.00 | 0.90 |
| SonicWall | 1.26 | 1.20 | 1.06 | 0.85 | 5.64 |
| Sophos | 22.55 | 19.60 | 19.77 | 16.13 | 17.99 |
| Versa Networks | 7.96 | 6.00 | 4.85 | 3.21 | 2.63 |
| WatchGuard | 1.73 | 3.04 | 2.50 | 1.12 | 0.61 |

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)

SSL/TLS Performance

Use of the Secure Sockets Layer (SSL) protocol and its newer iteration, Transport Layer Security (TLS), has risen in accordance with the increasing need for privacy online. Modern cybercampaigns frequently focus on attacking users through the most common web protocols and applications. NSS continues to receive inquiries from enterprise customers during their assessments of vendors that provide SSL/TLS decryption and protection technologies. For details on SSL performance, please see the SSL Performance Test Report for each device.

Single Application Flows

This test measures the performance of the device with single application flows. For details about single application flow testing, see the NSS Labs Next Generation Firewall Test Methodology, available at www.nsslabs.com.




Figure 19 –Single Application Flow: Telephony (Mbps)




Figure 20 –Single Application Flow: Financial (Mbps)




Figure 21 –Single Application Flow: Email (Mbps)




Figure 22 –Single Application Flow: File Sharing (Mbps)




Figure 23 –Single Application Flow: File Server (Mbps)




Figure 24 –Single Application Flow: Remote Console (Mbps)




Figure 25 –Single Application Flow: Video (Mbps)




Figure 26 –Single Application Flow: Meeting (Mbps)




Figure 27 –Single Application Flow: Database (Mbps)

Test Methodology

NSS Labs Next Generation Firewall Test Methodology v8.0

NSS Labs SSL/TLS Performance Test Methodology v1.3

NSS Labs Evasions Test Methodology v1.1

Contact Information

3711 South Mopac Expressway

Building 1, Suite 400

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.

© 2018 NSS Labs, Inc. All rights reserved. No part of this publication may be reproduced, copied/scanned, stored on a retrieval system, e-mailed or otherwise disseminated or transmitted without the express written consent of NSS Labs, Inc. (“us” or “we”).

Please read the disclaimer in this box because it contains important information that binds you. If you do not agree to these conditions, you should not read the rest of this report but should instead return the report immediately to us. “You” or “your” means the person who accesses this report and any entity on whose behalf he/she has obtained this report.

1. The information in this report is subject to change by us without notice, and we disclaim any obligation to update it.
2. The information in this report is believed by us to be accurate and reliable at the time of publication, but is not guaranteed. All use of and reliance on this report are at your sole risk. We are not liable or responsible for any damages, losses, or expenses of any nature whatsoever arising from any error or omission in this report.
3. NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY US. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT, ARE HEREBY DISCLAIMED AND EXCLUDED BY US. IN NO EVENT SHALL WE BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, PUNITIVE, EXEMPLARY, OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
4. This report does not constitute an endorsement, recommendation, or guarantee of any of the products (hardware or software) tested or the hardware and/or software used in testing the products. The testing does not guarantee that there are no errors or defects in the products or that the products will meet your expectations, requirements, needs, or specifications, or that they will operate without interruption.
5. This report does not imply any endorsement, sponsorship, affiliation, or verification by or with any organizations mentioned in this report.
6. All trademarks, service marks, and trade names used in this report are the trademarks, service marks, and trade names of their respective owners.