

Q1 2021

Enterprise Firewall

AA

Overview

During an independent test of the Juniper Networks SRX4600 v18.4X3.12, the product offered excellent security, very good plain text and SSL/TLS performance/functionality, and a competitive total cost of ownership. Customer feedback was cautiously optimistic; expressing a need to see a consistent commitment to security over time before getting fully onboard.

Security Effectiveness was excellent; Juniper Networks blocked 264 out of 264 evasions, 2,320 out of 2,331 exploits, and passed all the stability and reliability tests. The Plain Text Rated Throughput was very good with 11,483 Mbps, the HTTPS Rated Throughput was very good, with 5,159 Mbps, giving Juniper a combined Rated Throughput of 7,056 Mbps.

AA | Management


AAA | Security Effectiveness

AA | SSL/TLS Functionality

A | Customer Feedback

Security Effectiveness	AAA		
	Samples Tested	Samples Blocked	Blocked %
Block Rate	2,331	2,320	99.5%
Evasions			
HTTP Evasions	150	150	100%
IP Packet Fragmentation/TCP Segmentation	89	89	100%
Combination of Evasions	25	25	100%
False Positive Testing	PASS		
Stability & Reliability	PASS		

SSL/TLS Functionality	AA
Decryption validation	Supported
Top 24 cipher support	20/24 Supported
Emergent ciphers Support for x25519 Elliptical curve specification	0/2 Supported Not Supported
Prevention of weak ciphers	4/4 Prevented
Decryption bypass policy based on Layer 3 information	Supported
Decryption bypass policy based on Layer 4 information	Supported
Certificate validation	2/2 Supported
Support for session method: ID reuse Ticket reuse	Supported Supported



Performance		Rated Throughput
Plain Text Rated Throughput	11,483 Mbps	7,056 Mbps
HTTPS (SSL/TLS) Rated Throughput	5,159 Mbps	
Single Application Flow (Mbps)		
Telephony	5,838	
Financial	995	
Email	7,740	
File Sharing	6,674	
Fileserver	10,000	
Remote Console	2,398	
Video	10,000	
Database	10,000	

Centralized Management	AA
Customer Feedback	A

3-Year Cost

\$57,297

Summary of Results

Table of Contents

SECURITY EFFECTIVENESS	3
ACCESS CONTROL	3
<i>False Positives</i>	3
EXPLOITS	3
RESISTANCE TO EVASIONS	3
PERFORMANCE.....	4
MAXIMUM CAPACITY	4
HTTP CAPACITY.....	5
APPLICATION AVERAGE RESPONSE TIME – HTTP.....	5
RAW PACKET PROCESSING PERFORMANCE (UDP THROUGHPUT)	6
SINGLE APPLICATION FLOWS.....	6
SSL/TLS.....	7
DECRYPTION VALIDATION.....	7
CIPHER SUPPORT	7
<i>Certificate Validation</i>	7
<i>TLS Session Re-use</i>	7
SSL/TLS PERFORMANCE	8
<i>Application Average Response Time: HTTPS</i>	8
STABILITY AND RELIABILITY	9
BLOCKING UNDER EXTENDED ATTACK	9
PASSING LEGITIMATE TRAFFIC UNDER EXTENDED ATTACK.....	9
BEHAVIOR OF THE STATE ENGINE.....	9
POWER FAIL	9
BACKUP/RESTORE	9
PERSISTENCE OF DATA.....	9
COST OF TESTED CONFIGURATION.....	10
INSTALLATION HOURS	10
PRICING OVER 3 YEARS.....	10
APPENDIX A – SCORECARD	11
APPENDIX B - SSL FUNCTIONALITY	13
APPENDIX C.....	14
AUTHORS	15
CONTACT INFORMATION.....	15

Security Effectiveness

An enterprise firewall is a mechanism used to protect a trusted network from an untrusted network, while allowing authorized communications to pass from one side to the other, thus facilitating secure business use of the Internet. Firewalls traditionally have been deployed to defend the network on the edge, but some enterprises have expanded their deployment to include internal segmentation.

Access Control

PASS

100%

Enterprise firewalls have undergone several stages of development, from early packet filtering and circuit relay firewalls, to application-layer (proxy-based), dynamic packet filtering firewalls, and user/application aware “next-generation” firewalls. Throughout their history, the goal has been to enforce an access control policy between two networks, and they should therefore be viewed as an implementation of policy.

Access control is the basic role of the firewall. Access control rules are configured on a firewall to permit or deny traffic from one network resource to another based on identifying criteria such as: source, destination, service, and application. A failure of any access control test would result in an overall score of zero percent for security effectiveness.

False Positives

A key to effective protection is the ability to correctly identify and allow legitimate traffic while maintaining protection against malware, exploits, and phishing attacks. False positives are any legitimate, non-malicious content/traffic that are perceived as malicious. False positive tests flex the ability of the firewall to block attacks while permitting legitimate traffic. If a device experienced false positive events, it was tuned until no further false positive events were encountered.

Exploits

Blocked

2,320/2,331(99.5%)

For a device to be eligible for security effectiveness testing, it must perform all of the tests included in the enterprise firewall test methodology with its protection against network-delivered exploitation features enabled.

We also measured the resiliency of a device by introducing previously unseen variations of a known exploit and measured the device’s effectiveness against them. Juniper Networks has since updated its software to 20.4R1.12 (Signature 3346), and now blocks all missed exploits.




Figure 1 —Exploit Block Rate

Resistance to Evasions

Blocked

264/264 (100%)

Threat actors apply evasion techniques to disguise and modify attacks in order to avoid detection by security products. Therefore, it is imperative that an enterprise firewall correctly handles evasions. If an enterprise firewall fails to detect a single form of evasion, an attack can bypass protection.

Our engineers verified that the enterprise firewall was capable of blocking exploits when subjected to numerous evasion techniques. To develop a baseline, we took several attacks that had previously been blocked. We then applied evasion techniques to those baseline samples and tested. This ensured that any misses were due to the evasions and not the underlying (baseline) attacks.

For example, we applied an HTTP protocol evasion technique to drive-by exploits where the attacker responds with a version 1.0 declaration in the status line and chunking declared in the Transfer-Encoding header, while the body is sent unchunked.




Figure 2 – Resistance to Evasions

Performance

The performance of the enterprise firewall was tested using various traffic conditions that provide metrics for real-world performance. Individual implementations will vary based on usage; however, these quantitative metrics provide a gauge as to whether a particular firewall is appropriate for a given environment.

Maximum Capacity

The use of traffic generation appliances allowed our engineers to create “real-world” traffic at multi-Gigabit speeds as a background load for the tests. The aim of these tests was 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 contained valid payload and address data. These tests provide an excellent representation of a live network at various connection/transaction rates.

Note that in all tests, final measurements were taken at the following critical “breaking points”:

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




Figure 3 – Concurrency and Connection Rates

HTTP Capacity

The goal was 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 was forced to track valid TCP sessions, thus ensuring a higher workload rather than simple packet-based background traffic. This provided a test environment that is as close to real-world conditions as possible in a lab environment, while ensuring absolute accuracy and repeatability.




Figure 4 - HTTP Capacity

Each transaction consisted of a single HTTP GET request, and there were no transaction delays (i.e., the web server responded immediately to all requests). All packets contained valid payload (a mix of binary and ASCII objects) and address data, and this test provided an excellent representation of a live network (albeit one biased towards HTTP traffic) at various network loads.

Application Average Response Time – HTTP

Test traffic was passed across the infrastructure switches and through all inline port pairs of the device simultaneously (the latency of the basic infrastructure was known and was constant throughout the tests).

Application Average Response Time – HTTP (at 95% Maximum Load)	Milliseconds
2,500 Connections per Second – 44 KB Response	4.3
5,000 Connections per Second – 21 KB Response	3.8
10,000 Connections per Second – 10 KB Response	2.7
20,000 Connections per Second – 4.5 KB Response	1.5
40,000 Connections per Second – 1.7 KB Response	1.2

Figure 5 – Average Application Response Time (Milliseconds)

Raw Packet Processing Performance (UDP Throughput)

This test used UDP packets of varying sizes generated by traffic generation appliances. 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 — was transmitted bidirectionally through each port pair of the device. Each packet contained dummy data and was 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 were verified by network monitoring tools before each test began. Multiple tests were run and averages were taken where necessary.

This traffic did not attempt to simulate any form of real-world network condition. No TCP sessions were created during this test, and there was very little for the detection engine to do. However, each vendor was required to write a signature to detect the test packets to ensure that they were being passed through the detection engine and not “fast-tracked” from the inbound port to the outbound port.




Figure 6 - Raw Packet Processing Performance (UDP Traffic)

Single Application Flows

Where previous tests provided a pure HTTP environment with varying connection rates and average packet sizes, the goal of this test was to simulate real-world single application traffic.

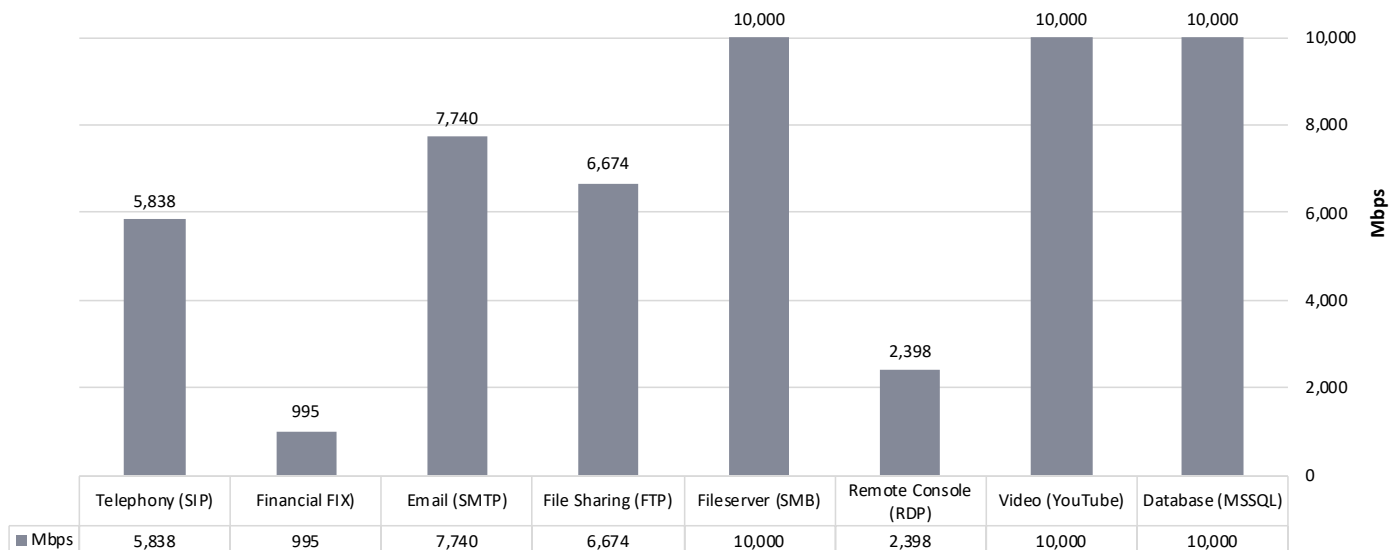


Figure 7 – Single Application Flows

SSL/TLS

Use of the Secure Sockets Layer (SSL) protocol and its current iteration, Transport Layer Security (TLS), is rising dramatically in response to an ever-increasing need for online privacy. In March 2020, data collected by Tranco on their Top 1 Million¹ shows that 60% of web traffic is being sent over HTTPS. While we believe the use of encryption is a good thing, SSL/TLS is susceptible to various security attacks at multiple levels of network communication. Attacks have been observed in the handshake protocol, record protocol, application data protocol, and Public Key Infrastructure (PKI), to name just a few.

To address the growing threat of focused attacks using the most common web protocols and applications, the capabilities of enterprise firewalls were tested to provide visibility into the SSL/TLS payloads and detect attacks concealed by encryption as well as attacks against the encryption protocols themselves.

SSL/TLS Functionality	
Decryption validation	Supported
Top 24 cipher support	20/24 Supported
Emergent ciphers Support for x25519 Elliptical curve specification	0/2 Supported Not Supported
Prevention of weak ciphers	4/4 Prevented
Decryption bypass policy based on Layer 3 information	Supported
Decryption bypass policy based on Layer 4 information	Supported
Certificate validation	2/2 Supported
Support for session method: ID reuse Ticket reuse	Supported Supported

Figure 8 – SSL/TLS Functionality

Decryption Validation

To confirm that the firewall was correctly decrypting and (if applicable) inspecting SSL/TLS traffic, a validation test was performed prior to conducting functional or performance testing. The firewall offering was expected to cover all test cases in the methodology with a single configuration.

This test used a known (previously blocked) exploit embedded in encrypted traffic and passed through the firewall. The firewall was expected to decrypt the stream, detect the exploit, and block or alert as appropriate. The purpose of this test was not to evaluate security effectiveness of the firewall, but rather to validate that it was properly decrypting and inspecting traffic.

Cipher Support

The firewall offering was expected to be capable of negotiating a wide range of commonly used SSL/TLS ciphers in order to increase the security visibility of potential threats encapsulated in real-world SSL/TLS traffic. This test covered the top 30 cipher suites as determined in the methodology. Unless otherwise specified, the functional tests used the most common key sizes for RSA (2,048 bit) and ECDSA (256 bit).

Certificate Validation

The firewall offering was expected to validate the status of all SSL/TLS certificates presented. When presented with an invalid certificate, the firewall offering must either prevent the establishment of a connection or replicate the original invalid status in the proxied/resigned certificate presented to the client, such that the client is aware of the potential risk.

TLS Session Re-use

In order to improve performance and reduce the overhead associated with conducting the full handshake for each session, the TLS protocol allows for abbreviated handshakes, which re-use previously established sessions. The two primary methods for session re-use are session IDs and session tickets. Whereas session IDs are included in the main TLS specification, session tickets are an extension of the specification, detailed in a separate RFC. Support for both of these methods was tested under this section.

¹ Tranco Top 1 Million analysis performed in March 2020, by Scott Helme (<https://scotthelme.co.uk/top-1-million-analysis-march-2020/>)

SSL/TLS Performance

The goal was to stress the HTTPS engine and determine how the device coped with network loads of varying average packet size and varying connections per second. By creating session-based traffic with varying session lengths, the device was forced to track valid TCP sessions, thus ensuring a higher workload rather than simple packet-based background traffic. Encrypting the traffic using SSL/TLS with varying algorithms forced the device to decrypt traffic prior to inspection which increased the workload further. This provided a test environment that is as close to real-world conditions as possible to achieve in a lab environment (albeit one biased towards HTTPS traffic), while ensuring accuracy and repeatability.

Tests were conducted first with one transaction per connection; a single (1) HTTP(S) GET request. Tests were then conducted with multiple transactions per connection; ten (10) HTTP(S) GET requests. There were no transaction delays (i.e., the web server responded immediately to all requests), and all packets contained valid payloads (a mix of binary and ASCII objects) and address data.

Application Average Response Time: HTTPS

Test traffic was passed across the infrastructure switches and through all inline port pairs of the device simultaneously (the latency of the basic infrastructure was known and was constant throughout the tests). The results were recorded at each response size at a load level of 90% of the maximum throughput with zero packet loss.

SSL/TLS Performance Testing		TLS ECDHE RSA W/ AES 256 GCM SHA384 (2k)	TLS ECDHE RSA W/ AES 256 GCM SHA384 (4k)	TLS ECDHE RSA W/ AES 128 GCM SHA256 (2k)	TLS ECDHE ECDSA W/ AES 128 GCM SHA256	TLS ECDHE RSA W/ AES 256 CBC SHA384 (2k)
Maximum HTTP Connections per Second		CPS				
Negligible Payload	1-byte Response	21,232	20,716	21,056	21,227	18,672
HTTPS Capacity, No Persistence		Mbps				
HTTPS 1.1 – Single GET Request	2880 KB Response	8,175	7,825	8,125	7,850	5,050
	768 KB Response	7,227	7,640	7,500	7,480	4,867
	192 KB Response	5,480	5,450	5,237	5,475	3,680
	44 KB Response	2,729	2,646	2,760	2,897	2,215
HTTP Capacity with Persistent Connections		Mbps				
HTTPS 1.1 – 10 GET Requests	288 KB Response	7,175	6,850	7,025	7,125	4,675
	76.8 KB Response	6,807	6,373	6,760	6,787	4,547
	19.2 KB Response	3,098	3,077	2,973	3,062	2,458
	4.4 KB Response	1,364	1,372	1,424	1,335	1,167
Application Average Response Time		Milliseconds				
HTTPS 1.1 – Single GET Request	2880 KB Response	136	132	137	128	257
	768 KB Response	65	74	66	63	94
	192 KB Response	22	20	17	20	28
	44 KB Response	8	6	7	8	9
HTTPS 1.1 – 10 GET Requests	288 KB Response	49	33	41	47	78
	76.8 KB Response	32	16	24	23	35
	19.2 KB Response	14	13	10	12	20
	4.4 KB Response	9	9	10	8	10

Figure 9 – SSL/TLS Performance Results

Stability and Reliability

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

The product was required to remain operational and stable throughout these tests, and to block 100% of previously blocked traffic, raising an alert for each. If any policy-forbidden traffic passes, caused by either the volume of traffic or by the product failing open for any reason, this resulted in a fail.

Blocking Under Extended Attack

The network firewall was exposed to a constant stream of policy or protocol violations over an extended period of time. The product was configured to block and alert, and thus this test provides an indication of the effectiveness of both the flow management and alert handling mechanisms.

The product was expected to remain operational and stable throughout this test and to correctly handle 100% of recognizable policy or protocol requests, raising an alert for each. If any recognizable policy violations are passed, caused by either the volume of traffic or by the product failing open for any reason, this resulted in a fail.

Behavior of the State Engine

This test determined whether the product was capable of preserving state across a large number of open connections over an extended time period. At various points throughout the test (including after the maximum had been reached), it was confirmed that the product was still capable of inspecting and blocking traffic that is in violation of the currently applied network control policy, while confirming that legitimate traffic is not blocked (perhaps as a result of exhaustion of the resources allocated to state tables). The product must be able to apply policy decisions effectively based on inspected traffic at all load levels.

- **Passing Legitimate Traffic – Normal Load:** This test ensured that the product continued to pass legitimate traffic as the number of open sessions reached 75% of the maximum determined previously in performance testing.
- **State Preservation – Maximum Exceeded:** This test determined whether the product maintained the state of pre-existing sessions as the number of open sessions exceeded the maximum determined previously in performance testing.
- **Drop Legitimate Traffic – Maximum Exceeded:** This test ensured that the product continued to drop all traffic as the number of open sessions exceeded the maximum determined previously in performance testing.

Power Fail

Power to the device is cut while passing a mixture of legitimate and disallowed traffic. Firewalls should always be configured to fail closed—no traffic should be passed once power has been cut.

Passing Legitimate Traffic under Extended Attack

This test is identical to the stability test run previously where the external interface of the product was exposed to a constant stream of policy or protocol violations over an extended period of time. The product was expected to remain operational and stable throughout this test, and to pass most or all of the legitimate traffic. If an excessive amount of legitimate traffic was blocked throughout this test, caused by either the volume of traffic or by the solution failing for any reason, this resulted in a fail.

Backup/Restore

Backing up and restoring a device's configuration is a critical component of deploying any managed device within a live network. It should be possible to export configurations and store them offline for backup purposes. Additionally, it should be possible to completely reconfigure the device using the offline configuration file(s). This includes restoring all policies and interface information in order to deploy a device.

Persistence of Data

The device should retain all configuration data, policy data, and locally logged data once it has been restored to operation following power failure.

Stability and Reliability	Result
Blocking under Extended Attack	PASS
Passing Legitimate Traffic under Extended Attack	PASS
Behavior of the State Engine under Load	
Attack Detection/Blocking – Normal Load	PASS
State Preservation – Normal Load	PASS
Pass Legitimate Traffic – Normal Load	PASS
Drop Traffic – Maximum Exceeded	PASS
Power Fail	PASS
Backup/Restore	PASS
Persistence of Data	PASS
Stability	PASS

Cost of Tested Configuration

Implementation of security solutions can be complex, with several factors affecting the overall cost of deployment, maintenance, and upkeep. Each of the following 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 10 depicts the number of hours of labor required to install each device using only local device management options. The table accurately reflects the amount of time that our engineers, with the help of vendor engineers, needed to install and configure the device to the point where it operated successfully in the test harness, passed legitimate traffic, and blocked and detected prohibited or malicious traffic. This closely mimics a typical enterprise deployment scenario for a single device. We have used a cost of USD\$75 per hour in this model.

The installation cost is based on the time that an experienced security engineer would require to perform the installation tasks described above. This approach allowed CyberRatings 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 cost figures.

Product	Installation (Hours)
Juniper Networks SRX4600 v18.4X3.12	8

Figure 10 – Sensor Installation Time (Hours)

Pricing over 3 Years

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 are for single device management and maintenance only; costs for central management solutions (CMS) may be extra.

Product	Year 1 Cost	Year 2 Cost	Year 3 Cost	3-Year Cost
Juniper Networks SRX4600 v18.4X3.12	\$40,918	\$4038+\$4,151	\$4038+\$4,151	\$57,297

Figure 11 – 3-Year Cost (US\$)

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

Appendix A – Scorecard

Description	Security Effectiveness		
False Positive Testing	PASS		
	Samples Tested	Samples Blocked	Blocked %
Block Rate	2,331	2,320	99.5%
Evasions			
HTTP Evasions	150	150	100%
IP Packet Fragmentation/TCP Segmentation	89	89	100%
Combination of Evasions	25	25	100%
Performance			
Raw Packet Processing Performance (UDP Traffic)	Mbps	Latency(μs)	
64 Byte Packets	3,940	35	
128 Byte Packets	9,010	36	
256 Byte Packets	20,000	36	
512 Byte Packets	20,000	47	
1024 Byte Packets	20,000	48	
1280 Byte Packets	20,000	65	
1514 Byte Packets	20,000	86	
Maximum Capacity	Concurrent Connections/CPS/TPS		
Max Concurrent TCP Connections	485,618		
Max TCP Connections per Second	74,500		
Max HTTP Connections per Second	95,000		
Max HTTP Transactions per Second	180,605		
HTTP Capacity with No Transaction Delays	CPS	Mbps	Average Response Time(ms)
2,500 Connections Per Second – 44Kbyte Response	24,390	9,756	4.3
5,000 Connections Per Second – 21Kbyte Response	38,710	7,742	3.8
10,000 Connections Per Second – 10Kbyte Response	55,110	5,511	2.7
20,000 Connections Per Second – 4.5Kbyte Response	65,820	3,291	1.5
40,000 Connections Per Second – 1.7Kbyte Response	79,000	1,975	1.2

Single Application Flow	Mbps
Telephony	5,838
Financial	995
Email	7,740
File Sharing	6,674
Fileserver	10,000
Remote Console	2,398
Video	10,000
Database	10,000
Stability & Reliability	
Blocking Under Extended Attack	PASS
Passing Legitimate Traffic Under Extended Attack	PASS
Behavior of The State Engine Under Load	PASS
Attack Detection/Blocking - Normal Load	PASS
State Preservation - Normal Load	PASS
Pass Legitimate Traffic - Normal Load	PASS
State Preservation - Maximum Exceeded	PASS
Drop Traffic - Maximum Exceeded	PASS
Power Fail	PASS
Backup/Restore	PASS
Persistence of Data	PASS
Stability	PASS

Appendix B - SSL Functionality

SSL/ TLS Functionality Testing	
Top 24 Cipher Suites	
TLS_DHE_RSA_WITH_AES_128_CBC_SHA	Supported
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	Supported
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	Supported
TLS_DHE_RSA_WITH_AES_256_CBC_SHA	Supported
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	Supported
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	Supported
TLS_DHE_RSA_WITH_CAMELLIA_256_CBC_SHA	Not Supported
TLS_DHE_RSA_WITH_SEED_CBC_SHA	Not Supported
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	Supported
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	Supported
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	Supported
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	Supported
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	Supported
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	Supported
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	Supported
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	Supported
TLS_RSA_WITH_AES_128_CBC_SHA	Supported
TLS_RSA_WITH_AES_128_CBC_SHA256	Supported
TLS_RSA_WITH_AES_128_GCM_SHA256	Supported
TLS_RSA_WITH_AES_256_CBC_SHA	Supported
TLS_RSA_WITH_AES_256_CBC_SHA256	Supported
TLS_RSA_WITH_AES_256_GCM_SHA384	Supported
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	Not Supported
TLS_RSA_WITH_SEED_CBC_SHA	Not Supported
Support for Emergent Cipher Suites	
TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256	Not Supported
TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256	Not Supported
Prevention of Weak Cipher Suites	
TLS_RSA_WITH_NULL_MD5	Supported
TLS_RSA_WITH_NULL_SHA	Supported
TLS_ECDH_anon_WITH_AES_128_CBC_SHA	Supported
TLS_ECDH_anon_WITH_AES_256_CBC_SHA	Supported
Inspection Bypass Exceptions	
Inspection bypass: Layer 3	Supported
Inspection bypass: Layer 4	Supported
Certificate Validation	
Revoked Certificate	Supported
Expired Certificate	Supported
TLS Session Re-use	
Session ID	Supported
Session Ticket	Supported

Appendix C

CyberRatings Classification Matrix	
RATING	DEFINITION
AAA	A product rated 'AAA' has the highest rating assigned by CyberRatings. The product's capacity to meet its commitments to consumers is extremely strong.
AA	A product rated 'AA' differs from the highest-rated products only to a small degree. The product's capacity to meet its commitments to consumers is very strong.
A	A product rated 'A' is somewhat less capable than higher-rated categories. However, the product's capacity to meet its commitments to consumers is still strong.
BBB	A product rated 'BBB' exhibits adequate stability and reliability. However, previously unseen events and use cases are more likely to negatively impact the product's capacity to meet its commitments to consumers.
	A product rated 'BB', 'B', 'CCC', 'CC', and 'C' is regarded as having significant risk characteristics. 'BB' indicates the least degree of risk and 'C' the highest. While such products will likely have some specialized capability and features, these may be outweighed by large uncertainties or major exposure to adverse conditions.
BB	A product rated 'BB' is more susceptible to failures than products that have received higher ratings. The product has the capacity to meet its commitments to consumers. However, it faces minor technical limitations that have a potential to be exposed to risks.
B	A product rated 'B' is more susceptible to failures than products rated 'BB'; however, it has the minimum capacity. Adverse conditions will likely expose the product's technical limitations that lead to an inability to meet its commitments to consumers.
CCC	A product rated 'CCC' is susceptible to failures and is dependent upon favorable conditions to perform expected functions. In the event of adverse conditions, the product is not likely to have the capacity to meet its commitments to consumers.
CC	A product rated 'CC' is highly susceptible to failures. The 'CC' rating is used when a failure has not yet occurred, but CyberRatings considers it a virtual certainty.
C	A product rated 'C' is highly susceptible to failures. The product is expected to fail under any abnormal operating conditions and does not offer a useful management systems and logging information compared with products that are rated higher.
D	A product rated 'D' is actively underperforming and failing and does not meet the use-case. The 'D' rating is used when the product is not operational without a major technical overhaul. Unless CyberRatings believes that such technical fixes will be made within a stated grace period (typically 30-90 calendar days), the 'D' rating also is an indicator that existing customers using the product have already experienced a failure and should take immediate action.

Authors

Thomas Skybakmoen, Vikram Phatak

Contact Information

CyberRatings.org
2303 Ranch Road 620 South
Suite 160, #501
Austin, TX 78734

info@cyberratings.org

www.cyberratings.org

© 2020 CyberRatings.org. All rights reserved. No part of this publication may be reproduced, copied/scanned, stored on a retrieval system, emailed or otherwise disseminated or transmitted without the express written consent of CyberRatings.org. ("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.