ENTERPRISE FIREWALL COMPARATIVE SECURITY VALUE MAP™ Q2 2023

Empirical data from testing is used to create CyberRatings.org Security Value Map™ (SVM). The SVM illustrates the relative value of security investment options by mapping the *Security Effectiveness* and the Total Cost *per Protected Mbps* (*Value*) of tested products.

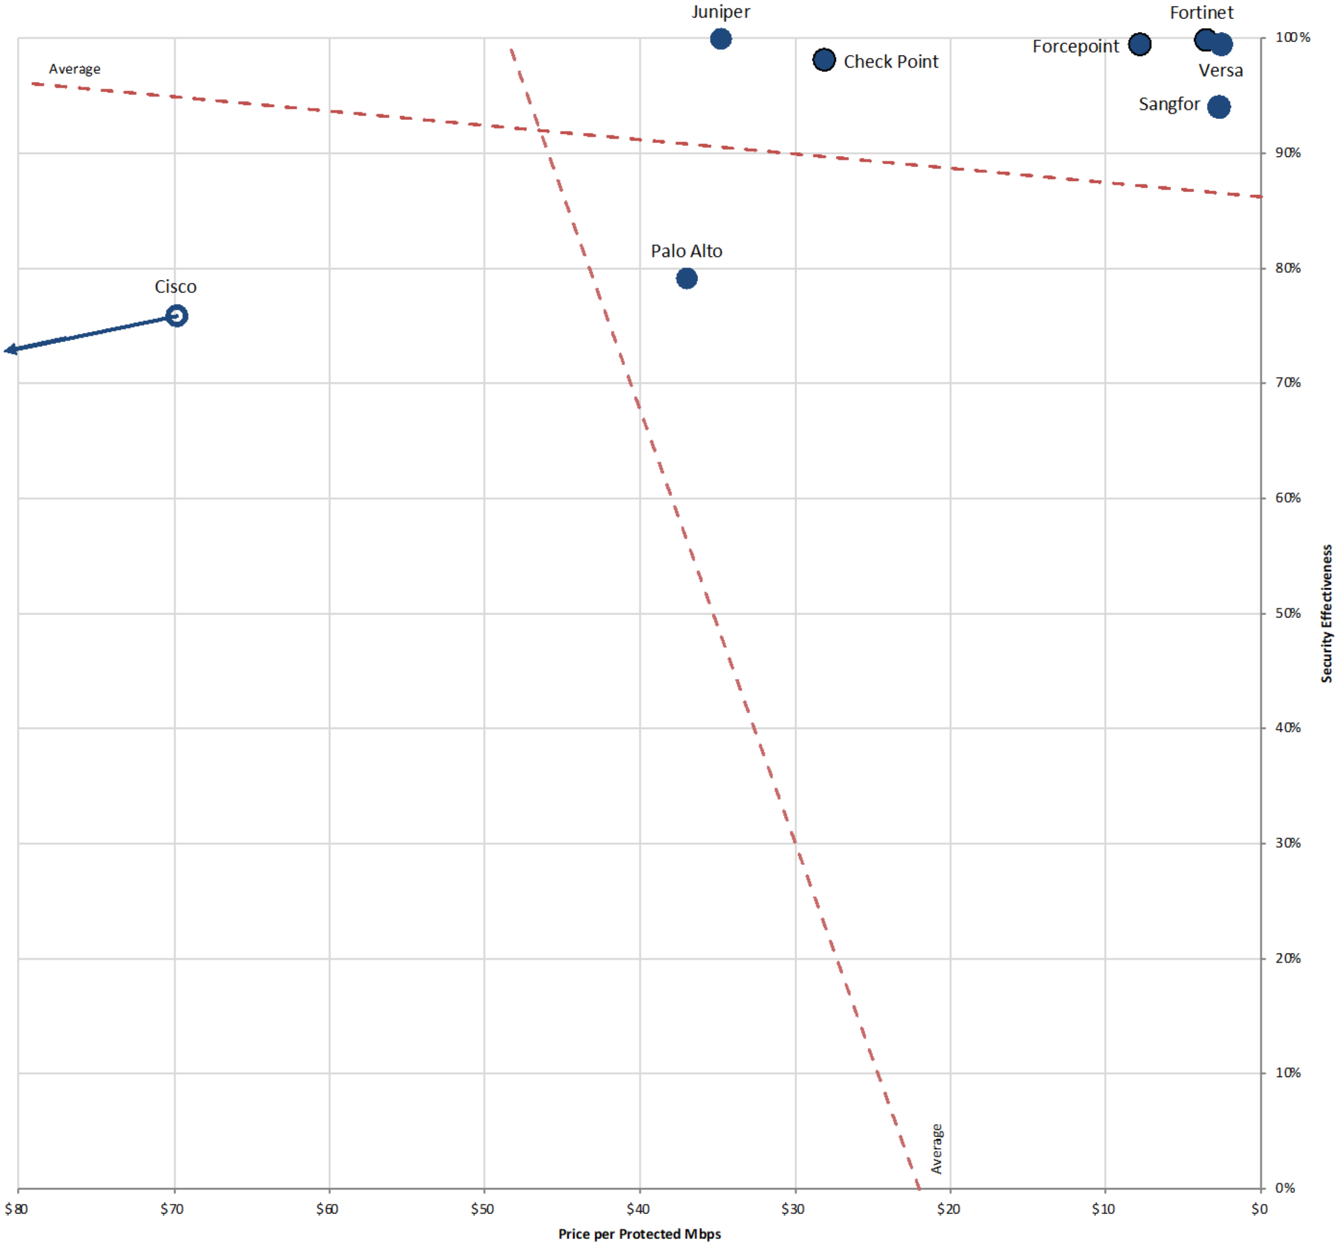

Enterprise Firewall	Rating	Security Effectiveness	Rated Throughput	MSRP + 24/7
Check Point	Recommended	98.14%	5,438	$149,970
Cisco	Caution	19.50%	1,537	$81,482
Forcepoint	Recommended	99.48%	4,235	$32,915
Fortinet	Recommended	99.88%	11,383	$40,532
Juniper Networks	Recommended	99.94%	7,572	$263,105
Palo Alto Networks	Neutral	79.15%	1,717	$50,200
Sangfor	Recommended	94.05%	5,782	$14,920
Versa Networks	Recommended	99.48%	12,160	$30,851

Ratings

The SVM provides a quick, high-level analysis of the detailed findings from our tests. Every enterprise has its own set of unique requirements and the SVM should only be a starting point. In addition to this comparative report, individual test reports are available for each product tested and can be found at www.CyberRatings.org.

The rating in the Security Value Map ™ (SVM) is determined by which section of the SVM the product falls within Recommended (top right), Security Recommended (top left), Neutral (bottom right), or Caution (bottom left). For more information on how the SVM is constructed, see the Appendix: *How to Read the SVM* section of this document.

CyberRatings.org is a 501(c)6 non-profit organization that provides independent, objective ratings of security product efficacy through our research and testing programs.

KEY FINDINGS

- Encryption matters: Roughly 80% of web traffic is encrypted. The top four cipher suites account for over 95% of HTTPS traffic.
    - Decryption is not on by default: Firewalls will not see attacks delivered via HTTPS unless configured to do so.
    - There is a performance cost when TLS/SSL is turned on. Sometimes performance is significantly different.
- When a "known good" exploit is blocked by a firewall, applying an evasion technique to that exploit is often easier for an attacker than finding a new exploit that isn't blocked by that firewall.
    - Many firewall evasion defenses are not on by default, potentially leaving customers at significant risk.
    - Most enterprises are not testing for evasions. Properly testing exploits is hard; properly testing evasions is very hard.
    - Some products have concerning gaps when it comes to evasions.
- At times, we found multiple signatures/rules for the same CVE, with some more effective than others.
    - Attempts to provide rapid coverage for vulnerabilities that are not fully understood can result in multiple exploit-specific signatures that may be inaccurate, ineffective, or prone to false positives.
    - A single poorly written signature/rule can significantly impact performance.

RECOMMENDATIONS

- Plan deployments using encrypted throughput, not just clear text.
    - Existing firewall deployments should enable TLS/HTTPS decryption features whenever possible and prioritize upgrading equipment when not.
    - Block unsupported cipher suites since traffic must be decrypted before detection can occur.
- Verify that your evasion defenses are enabled; when in doubt, check your vendor's best practices.
    - Test your evasion defenses; if you lack the capabilities, ask for help.
- Update firewalls regularly. New versions of software often have code/capabilities that signatures rely on.
- Be careful when using a vendor's test plan since their priorities might not be yours.
- Leverage CyberRatings test results to preselect vendors for a proof of concept (PoC).

Security Effectiveness

Implementing an enterprise firewall can be a complex process with multiple factors affecting overall security effectiveness. The following factors should be considered over the course of the useful life of the device:

- Deployment use cases:
  - What Operating systems and applications are to be protected?
  - What TLS/SSL cipher support is required?
  - What defensive capabilities are necessary (exploit block rate)?
- Product's ability to protect against common evasion techniques.
- Device stability and reliability

Security Effectiveness tests measured how well the enterprise firewall controlled network access, applications, and users while preventing exploits and evasions while remaining resistant to false positives.

There is a trade-off between security effectiveness and performance. Because of this trade-off, judging a product's security effectiveness within the context of its performance is essential, 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. To determine the relative security effectiveness of devices on the market and to facilitate accurate product comparisons, we use the following formula:

Security Effectiveness = Routing & access control *x* SSL/TLS Functionality *x* Threat Prevention *x* Stability & Reliability

By focusing on security effectiveness as a whole instead of on exploit block rate alone, we can factor in the ease with which defenses can be bypassed as well as the reliability of the device. Figure 1 summarizes the security results using the abovementioned formula.

Enterprise Firewall	Routing & access control	SSL/TLS Functionality	Threat Prevention		Stability & Reliability	**Security Effectiveness**
			Exploits	Evasions		
Check Point	100.0%	100.0%	99.13%	99.0%	100.0%	**98.14%**
Cisco	100.0%	25.7%	81.84%	92.7%	100.0%	**19.50%**
Forcepoint	100.0%	100.0%	99.48%	100.0%	100.0%	**99.48%**
Fortinet	100.0%	100.0%	99.88%	100.0%	100.0%	**99.88%**
Juniper Networks	100.0%	100.0%	99.94%	100.0%	100.0%	**99.94%**
Palo Alto Networks	100.0%	100.0%	91.24%	86.8%	100.0%	**79.15%**
Sangfor	100.0%	100.0%	98.26%	95.7%	100.0%	**94.05%**
Versa Networks	100.0%	100.0%	99.48%	100.0%	100.0%	**99.48%**

Figure 1 – Security Effectiveness

Routing & Access Control

ACCESS CONTROL

Access control is the primary responsibility of a firewall. Throughout its history, the goal of a firewall has been to enforce an access control policy between two networks. Rules are configured to permit or deny traffic from one network resource to another based on identifying criteria such as source IP, destination IP, source port, destination port, and protocols.

Enterprise Firewall	Unrestricted Traffic Test	Segmented Traffic Test
Check Point	Pass	Pass
Cisco	Pass	Pass
Forcepoint	Pass	Pass
Fortinet	Pass	Pass
Juniper Networks	Pass	Pass
Palo Alto Networks	Pass	Pass
Sangfor	Pass	Pass
Versa Networks	Pass	Pass

Figure 2 – Access Control

ROUTING

This test validated that the firewalls enforced security policies over various policy use cases, from simple to complex. The tests were incrementally built on a baseline consisting of a simple configuration with no policy restrictions and no content inspection – to a complex multiple-zone configuration that supports many users, networks, policies, and applications. Traffic was tested at each level of complexity to ensure specified policies were enforced.

Enterprise Firewall	Simple Policies	Complex Multi-Zone Policies
Check Point	Pass	Pass
Cisco	Pass	Pass
Forcepoint	Pass	Pass
Fortinet	Pass	Pass
Juniper Networks	Pass	Pass
Palo Alto Networks	Pass	Pass
Sangfor	Pass	Pass
Versa Networks	Pass	Pass

Figure 3 – Routing

SSL/TLS Functionality

The use of the Secure Sockets Layer (SSL) protocol and its current iteration, Transport Layer Security (TLS), is now the norm. Let's Encrypt statistics show that as of January 2023, over 77% of web traffic is being sent over HTTPS.[1]

While CyberRatings believes using encryption is good, SSL/TLS is susceptible to various security attacks at multiple levels of network communication. For example, attacks have been observed in the handshake protocol, record protocol, application data protocol, and Public Key Infrastructure (PKI). To address the growing threat of 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. The table below lists the tested SSL/TLS in order of prevalence[2] per March 2023.

DECRYPTION VALIDATION

Version	Prevalence	Cipher Suites
TLS 1.3	63.90%	TLS_AES_256_GCM_SHA384 (0x13, 0x02)
TLS 1.2	13.70%	TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xC0, 0x30)
TLS 1.2	9.90%	TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xC0, 0x2F)
TLS 1.3	7.70%	TLS_AES_128_GCM_SHA256 (0x13, 0x01)
TLS 1.2	1.30%	TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256 (0xCC, 0xA8)
TLS 1.2	1.10%	TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (0xC0, 0x28)
TLS 1.3	1.10%	TLS_CHACHA20_POLY1305_SHA256 (0x13, 0x03)
TLS 1.2	0.30%	TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256 (0xCC, 0xA9)
TLS 1.2	0.30%	TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 (0xC0, 0x2C)
TLS 1.2	0.20%	TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 (0xC0, 0x2B)

Figure 4 – SSL/TLS Functionality (I)

First, we tested how the firewall handled cipher suites known to be insecure, using null ciphers (no encryption of data) and anonymous ciphers (no authorization). Then we validated the ability to correctly decrypt and inspect SSL/TLS traffic using prohibited content previously blocked during testing. (The content was then encrypted and verified that it was still blocked.) We then tested to see if we could permit conditional bypass of decryption. This might be required to preserve privacy for regulatory or other reasons. Lastly, we tested TLS session reuse; to improve performance and reduce the overhead associated with conducting the full handshake for each session. The TLS protocol allows for abbreviated handshakes, which reuse previously established sessions. Figures 5 and 6 show the results for each product under test. For additional details, please see the Individual test reports, which are available for each product tested and can be found at www.cyberratings.org.

[1] Let's Encrypt Stats (https://letsencrypt.org/stats/)

[2] https://crawler.ninja/files/ciphers.txt

Enterprise Firewall	Decryption Validation	Top 10 Cipher Support	Prevention of Weak Ciphers	Decryption Bypass Exceptions
Check Point	Pass	10/10	Pass	Pass
Cisco	Pass	5/10	Pass	Pass
Forcepoint	Pass	10/10	Pass	Pass
Fortinet	Pass	10/10	Pass	Pass
Juniper Networks	Pass	10/10	Pass	Pass
Palo Alto Networks	Pass	10/10	Pass	Pass
Sangfor	Pass	10/10	Pass	Pass
Versa Networks	Pass	10/10	Pass	Pass

Figure 5 – SSL/TLS Functionality (II)

Session reuse is one of the mechanisms to improve TLS/SSL performance for TLS, figure 6 lists which products had either functionality as an option.

Enterprise Firewall	TLS Session Reuse - Session Tickets	TLS Session Reuse - Session IDs
Check Point	Supported	Supported
Cisco	Not Supported	Supported
Forcepoint	Supported	Supported
Fortinet	Supported	Supported
Juniper Networks	Not Supported	Supported
Palo Alto Networks	Supported	Supported
Sangfor	Not Supported	Not Supported
Versa Networks	Not Supported	Not Supported

Figure 6 – SSL/TLS Functionality (III)

Threat Prevention

A 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. The CyberRatings exploit repository contains exploits for many protocols and applications. Exploit sets for individual tests are selected based on CVSS score (how widely used is an application + what can an attacker do?), use case, and customer relevance. This has implications for the age of exploits since some applications in industrial environments are deployed and then left untouched for years. In contrast, other applications within office environments are refreshed every 5-7 years.

EXPLOIT PROTECTION

An exploit is an attack that takes advantage of a protocol, product, operating system, or application vulnerability. CyberRatings verified that the firewall could detect and block exploits while remaining resistant to false positives by attempting to send exploits through the product under test. Additionally, we verified that the malicious traffic was blocked, and all appropriate logging and notifications were performed.

Coverage by Attack Vector

Because a failure to block attacks could result in significant compromise and severely impact critical business systems, firewalls should be evaluated against a broad set of exploits. Exploits can be categorized as either client-initiated or server-initiated. Server-initiated exploits are threats executed remotely against a vulnerable application and/or operating system by an individual, while client-initiated exploits are initiated by the vulnerable target. Client-initiated exploits are the most common type of attack experienced by the end user, and the attacker has little or no control as to when the threat is executed.

Tuning

Our research indicates that the majority of enterprises tune their enterprise firewall products. Therefore, the tested firewalls were tuned for this test. In addition, 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.

False Positives

A key to effective protection is correctly identifying and allowing legitimate traffic while maintaining protection against malware, exploits, and phishing attacks. False positives are any legitimate, non-malicious content/traffic 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.

Enterprise Firewall	Client-Initiated	Server-Initiated
Check Point	677/681 (99.41%)	1032/1043 (98.95%)
Cisco	470/681 (69.02%)	941/1043 (90.22%)
Forcepoint	680/681 (99.85%)	1035/1043 (99.23%)
Fortinet	681/681 (100.00%)	1041/1043 (99.81%)
Juniper Networks	681/681 (100.00%)	1042/1043 (99.90%)
Palo Alto Networks	581/681 (85.32%)	992/1043 (95.11%)
Sangfor	677/681 (99.41%)	1017/1043 (97.51%)
Versa Networks	679/681 (99.71%)	1036/1043 (99.33%)

Figure 7 – Exploit Protection (Client & Server – Initiated)

Exploit Block Rate by Year

Our research indicates that the most significant risks are not always driven by the latest “Patch Tuesday” disclosures. Studies reveal that many older applications, operating systems, and attacks are still in circulation and remain relevant.

Vendors may retire older signatures in attempts to alleviate the performance limitations of products; however, this may result in poor coverage for older vulnerabilities and inconsistent protection across products. The figure below classifies coverage by disclosure date, as tracked by CVE numbers.

Enterprise Firewall	Check Point	Cisco	Forcepoint	Fortinet	Juniper Networks	Palo Alto Networks	Sangfor	Versa Networks
2012	100.0%	81.9%	100.0%	100.0%	100.0%	96.6%	99.2%	99.2%
2013	100.0%	89.8%	100.0%	100.0%	100.0%	76.1%	100.0%	100.0%
2014	97.7%	79.5%	100.0%	100.0%	100.0%	85.2%	100.0%	100.0%
2015	100.0%	78.8%	100.0%	100.0%	100.0%	98.0%	100.0%	100.0%
2016	100.0%	74.2%	100.0%	100.0%	100.0%	91.8%	100.0%	100.0%
2017	99.3%	85.8%	100.0%	99.6%	100.0%	93.7%	99.6%	99.6%
2018	100.0%	71.4%	95.2%	100.0%	100.0%	84.1%	100.0%	100.0%
2019	99.8%	77.6%	98.8%	99.8%	99.8%	91.7%	99.6%	99.0%
2020	96.9%	95.9%	100.0%	100.0%	100.0%	85.7%	81.6%	100.0%
2021	97.8%	86.7%	100.0%	100.0%	100.0%	90.0%	98.9%	98.9%
2022	94.3%	87.4%	100.0%	100.0%	100.0%	92.0%	93.1%	100.0%

Figure 8 – Exploit Block Rate by Year

Different vendors take different approaches to adding coverage once a vulnerability is disclosed. Attempts to provide rapid coverage for vulnerabilities that are not fully understood can result in multiple exploit-specific signatures that may be inaccurate, ineffective, or prone to false positives. Vendors that have the resources to fully research a vulnerability should be able to produce vulnerability-oriented signatures that provide coverage for all exploits written to take advantage of that flaw. This approach provides more effective coverage with fewer false positives.

Exploit Coverage for Top Vendors

Exploits within the CyberRatings exploit library target a wide range of protocols and applications. The figure below shows how the product under test offers exploits protection for ten top vendors targeted in this test.

Enterprise Firewall	Adobe	Advantech	Apache	Apple	Cisco	Microsoft	Oracle	SolarWinds	VMware	HPE
Check Point	100.0%	98.6%	100.0%	100.0%	95.2%	99.7%	98.3%	100.0%	100.0%	100.0%
Cisco	61.5%	95.7%	97.2%	63.2%	100.0%	73.3%	81.0%	95.2%	80.0%	94.7%
Forcepoint	100.0%	100.0%	100.0%	100.0%	90.5%	100.0%	100.0%	100.0%	100.0%	100.0%
Fortinet	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Juniper Networks	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Palo Alto Networks	71.9%	100.0%	98.1%	100.0%	100.0%	84.4%	95.7%	97.6%	93.3%	97.4%
Sangfor	100.0%	100.0%	100.0%	100.0%	100.0%	98.6%	98.3%	97.6%	93.3%	98.6%
Versa Networks	100.0%	100.0%	99.1%	100.0%	100.0%	99.7%	100.0%	100.0%	100.0%	100.0%

Figure 9 – Exploit Coverage for Top 10 Vendors

RESISTANCE TO EVASIONS

Threat actors apply evasion techniques to disguise and modify attacks to avoid detection by security products. Therefore, it is imperative that a firewall correctly handles evasions. An attacker can bypass protection if a firewall fails to detect a single form of evasion.

Handling evasions is hard. And to our knowledge, this was the most comprehensive evasion test performed to date. Our engineers verified that the firewall could block exploits with evasion techniques applied. To develop a baseline, we took several previously blocked attacks. We then applied evasion techniques to those baseline samples and tested them. This ensured that any misses were due to the evasions, not the baseline samples.

We adjusted scoring for evasions according to their impact: For example, TCP evasions are more impactful than HTML evasions. A TCP evasion can be applied to thousands of exploits, vs. an HTML evasion is limited to far fewer exploits.

During testing, we used multiple exploits for each evasion technique to see how each product defended against these combinations. Some products properly handled an evasion technique with all tested exploits while others handled evasions with only some of the exploits.

Client-initiated evasions								
Enterprise Firewall	Check Point	Cisco	Forcepoint	Fortinet	Juniper Networks	Palo Alto Networks	Sangfor	Versa Networks
IP Spoofing	5	5	5	5	5	5	5	5
IP Fragmentation	96	96	96	96	96	96	96	96
TCP Segmentation	264	264	264	264	264	262	264	264
Layered Evasions	16	16	16	16	16	16	16	16
HTTP Obfuscation	172	172	172	172	172	151	172	172
HTTP Compression	72	72	72	72	72	28	72	72
HTML Obfuscation	122	108	124	124	124	120	124	124
Combination Evasions	4	4	4	4	4	4	4	4
Server-initiated evasions								
TCP Split Handshake	5	5	5	5	5	5	5	5
IP Fragmentation	238	238	238	238	238	238	237	238
TCP Segmentation	458	458	458	458	458	457	456	458
Layered Evasions	28	25	28	28	28	28	28	28

Figure 10 – Evasion Coverage (Client & Server – Initiated)

Evasion techniques are means of disguising and modifying attacks to avoid detection and blocking by security products. Missing a type of evasion means a hacker can use an entire class of exploits to circumvent the security product, rendering it virtually useless. The techniques used in this test have been widely known for years and should be considered minimum requirements for the enterprise firewall product category.

Missing a single evasion technique opens holes for attackers to get through and vendors should rectify such omissions immediately. Any financially motivated hacker with basic skills will know how to take advantage of these weaknesses, and toolkits exist to assist them.

STABILITY & RELIABILITY

Long-term stability is essential for an inline device, where failure can produce network outages. These tests verified the firewalls' stability and ability to maintain security effectiveness while under normal load and passing malicious traffic. A firewall that could not sustain legitimate traffic (or that crashed) while under hostile attack would 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 results in a failure.

All the devices remained operational and stable throughout all these tests and blocked 100% of previously known malicious attacks, raising an alert for each.

Enterprise Firewall	Blocking with Minimal Load	Blocking Under Load	Attack Detection/Blocking – Normal Load	State Preservation – Normal Load
Check Point	Pass	Pass	Pass	Pass
Cisco	Pass	Pass	Pass	Pass
Forcepoint	Pass	Pass	Pass	Pass
Fortinet	Pass	Pass	Pass	Pass
Juniper Networks	Pass	Pass	Pass	Pass
Palo Alto Networks	Pass	Pass	Pass	Pass
Sangfor	Pass	Pass	Pass	Pass
Versa Networks	Pass	Pass	Pass	Pass

Figure 11 – Stability & Reliability (I)

Enterprise Firewall	Pass Legitimate Traffic – Normal Load	State Preservation – Maximum Exceeded	Drop Traffic – Maximum Exceeded	Protocol Fuzzing & Mutation
Check Point	Pass	Pass	Pass	Pass
Cisco	Pass	Pass	Pass	Pass
Forcepoint	Pass	Pass	Pass	Pass
Fortinet	Pass	Pass	Pass	Pass
Juniper Networks	Pass	Pass	Pass	Pass
Palo Alto Networks	Pass	Pass	Pass	Pass
Sangfor	Pass	Pass	Pass	Pass
Versa Networks	Pass	Pass	Pass	Pass

Figure 12 – Stability & Reliability (II)

Performance

We tested 54 performance use cases for each product to capture their performance curves. This included maximum connections and transactions per second, concurrency, throughput, and latency to see how the firewall performed under various adverse conditions. As a result, each product has achieved a rated throughput. For more tests and details, please see the individual test reports.

RATED THROUGHPUT

We measured performance with different packet sizes and payloads to capture the firewall's performance curves for UDP, HTTP, and HTTPS. The "Rated Throughput" is an average of UDP, HTTP, and HTTPS Capacity (1,000, 2,000,4,000, and 8,000 CPS), and the "Real World Application Flows" is a good benchmark for what an enterprise can expect the firewall to achieve in a typical enterprise network.

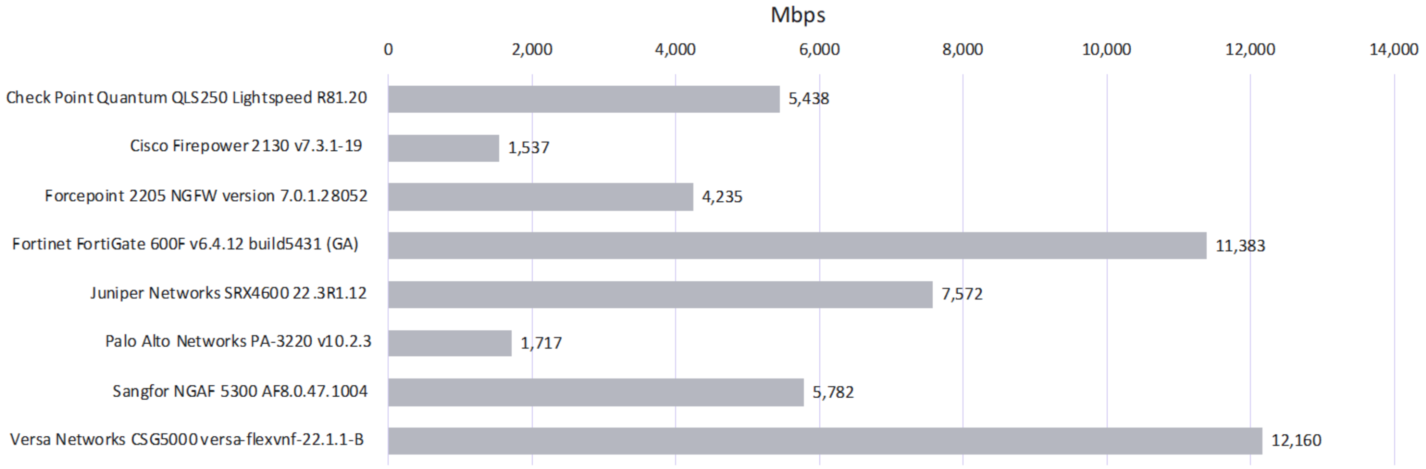

Figure 13 – Rated Throughput (Mbps)

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. 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.

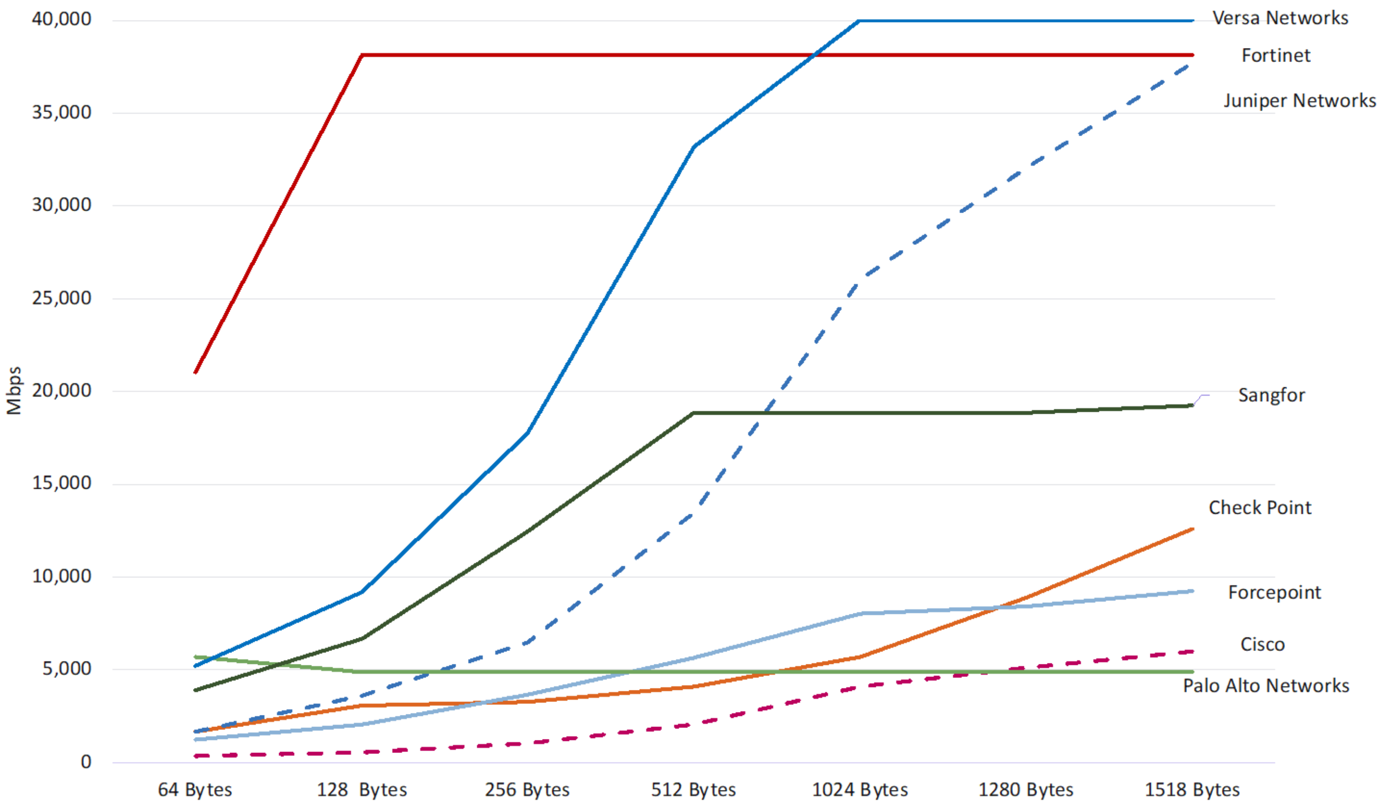

Figure 14 – UDP Throughput

MAXIMUM CAPACITY

These tests aimed 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. 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 increase in open connections.
- Excessive concurrent HTTP connections – Latency within the firewall is causing delays and increased response time.
- Unsuccessful HTTP transactions – Normally, there should be zero unsuccessful transactions. Once these appear, it indicates that firewall latency is causing connections to time out.

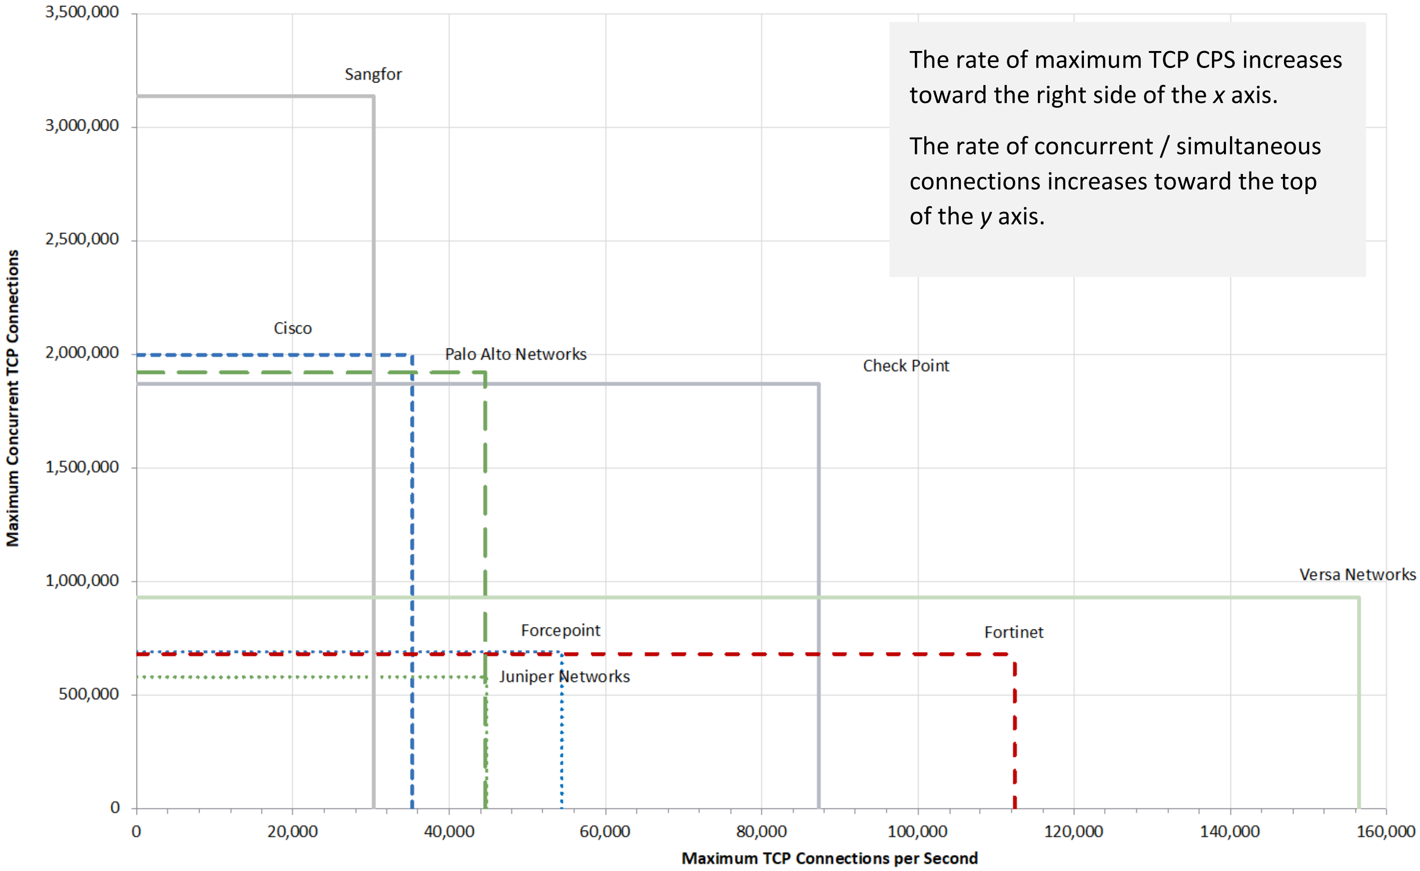

The rate of maximum TCP CPS increases toward the right side of the *x* axis.

The rate of concurrent / simultaneous connections increases toward the top of the *y* axis.

Figure 15 – Maximum TCP CPS & Concurrent/Simultaneous TCP Connections

Theoretical Maximum Concurrent TCP Connections

This test is designed to determine the device's maximum concurrent TCP connections with no data passing across the connections. This type of traffic would not typically be found on a normal network, but it provides the means to determine the maximum possible concurrent connections.

Maximum TCP Connections per Second

This test is designed to determine the maximum TCP connection rate of the device with one byte of data passing across the connections. This type of traffic would not typically be found on a normal network, but it provides the means to determine the maximum possible TCP connection rate.

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.

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. This test provided an excellent representation of a live network (albeit one biased towards HTTP traffic) at various network loads. For the application average response time, test traffic was passed across the infrastructure switches and through all inline port pairs of the device simultaneously (the basic infrastructure latency was known and constant throughout the tests).

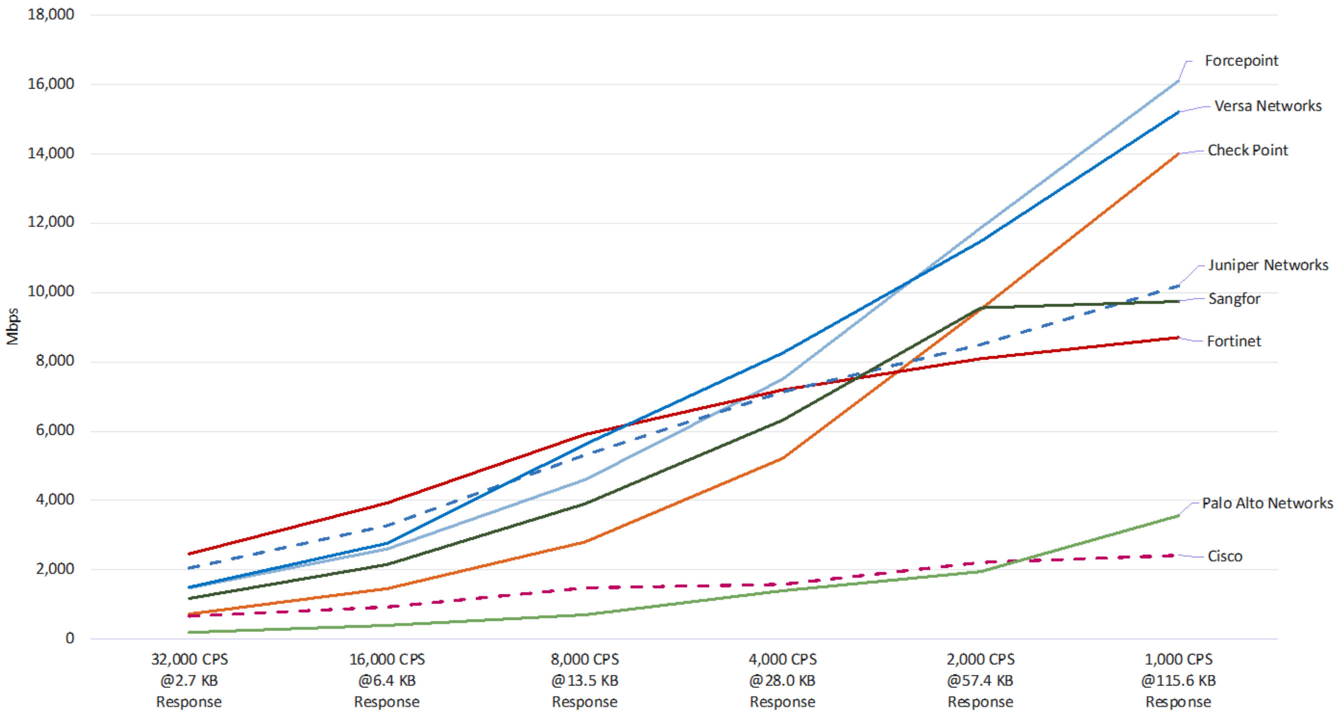

Figure 16 – HTTP Capacity

HTTPS CAPACITY

The goal was to stress the HTTPS engine and determine how the device coped with network loads of varying average packet sizes 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 than simple packet-based background traffic. Encrypting the traffic using SSL/TLS with varying algorithms forced the device to decrypt traffic before inspection, increasing the workload further.

Tests were conducted with one transaction per connection; a single (1) HTTP(S) GET request. There were no transaction delays (the webserver responded immediately to all requests), and all packets contained valid payloads (a mix of binary and ASCII objects) and address data. Testing determined the maximum rate at which the firewall could process HTTPS packets of various sizes and its efficiency at forwarding packets quickly to provide the highest level of network performance with the lowest latency. The results were recorded at a load level of 95% of the maximum throughput with zero packet loss at each response size.

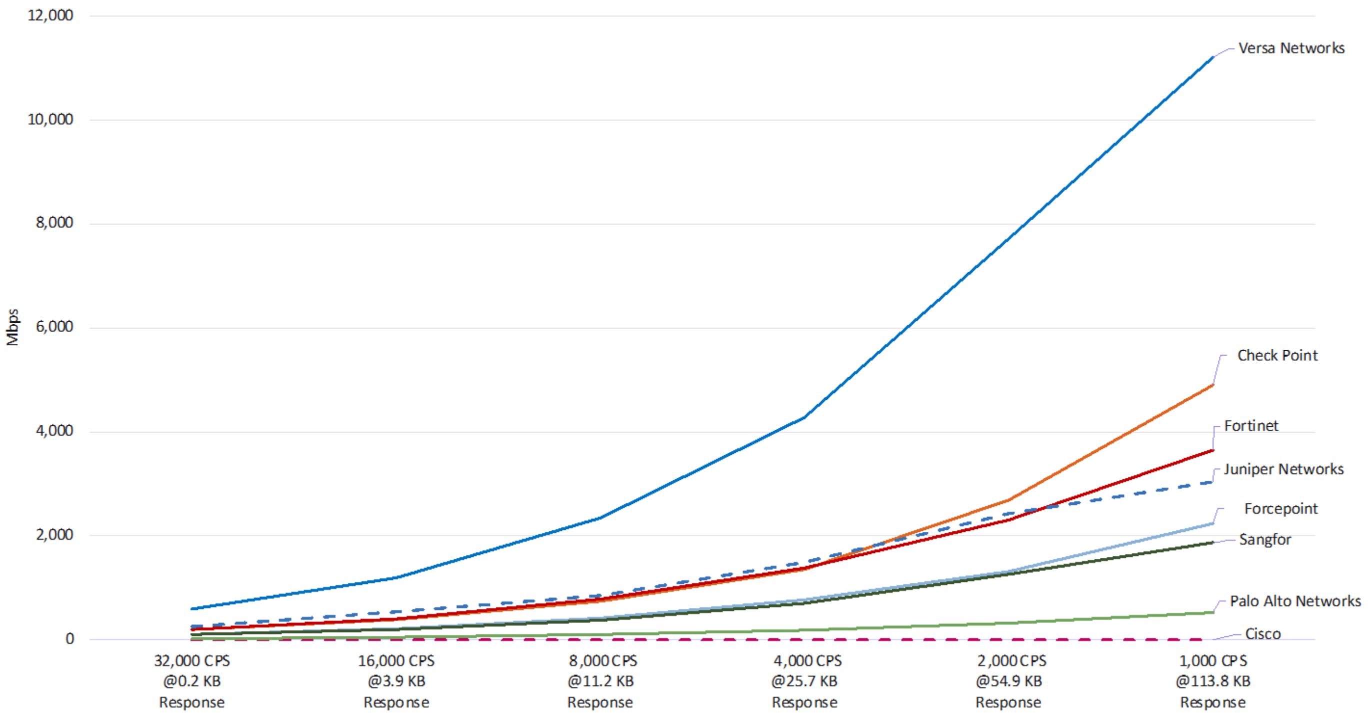

Figure 17 – HTTPS Capacity [TLS_AES_256_GCM_SHA384 (0x13, 0x02)]

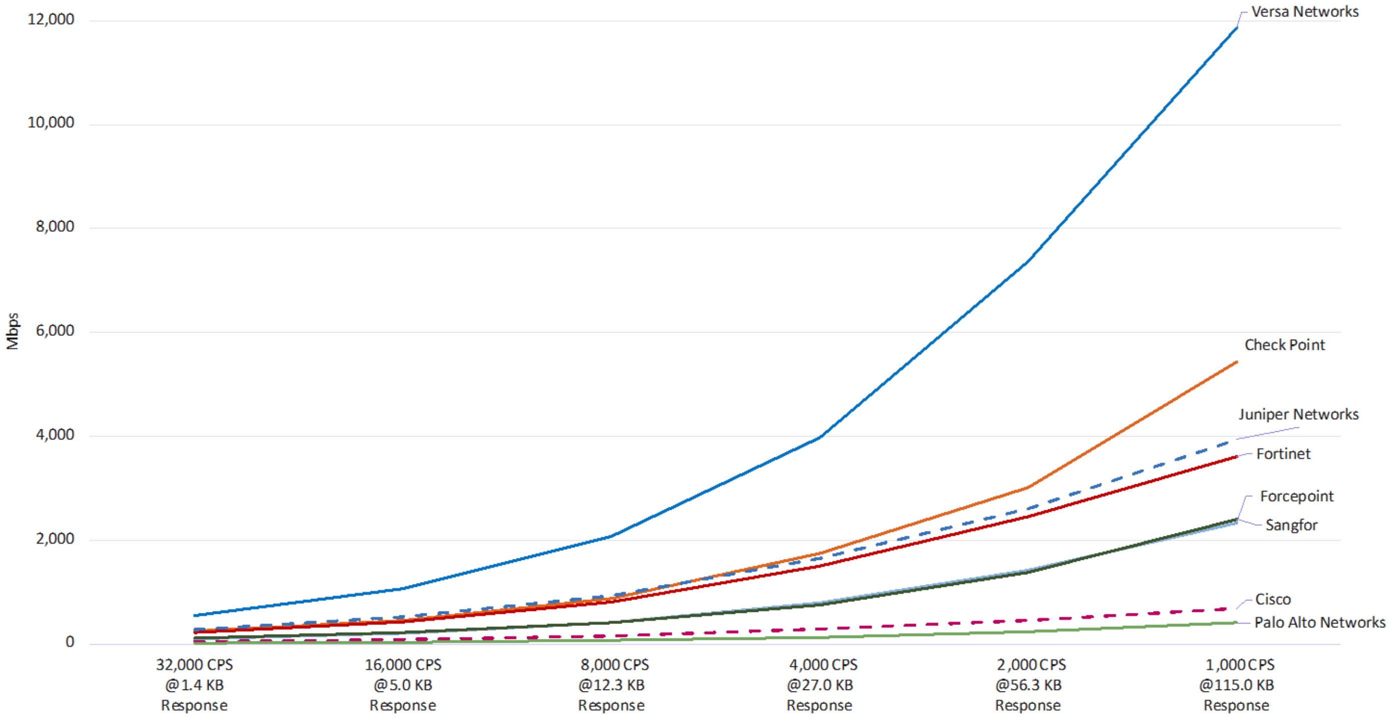

Figure 18 – HTTPS Capacity [TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (0xC0, 0x30)]

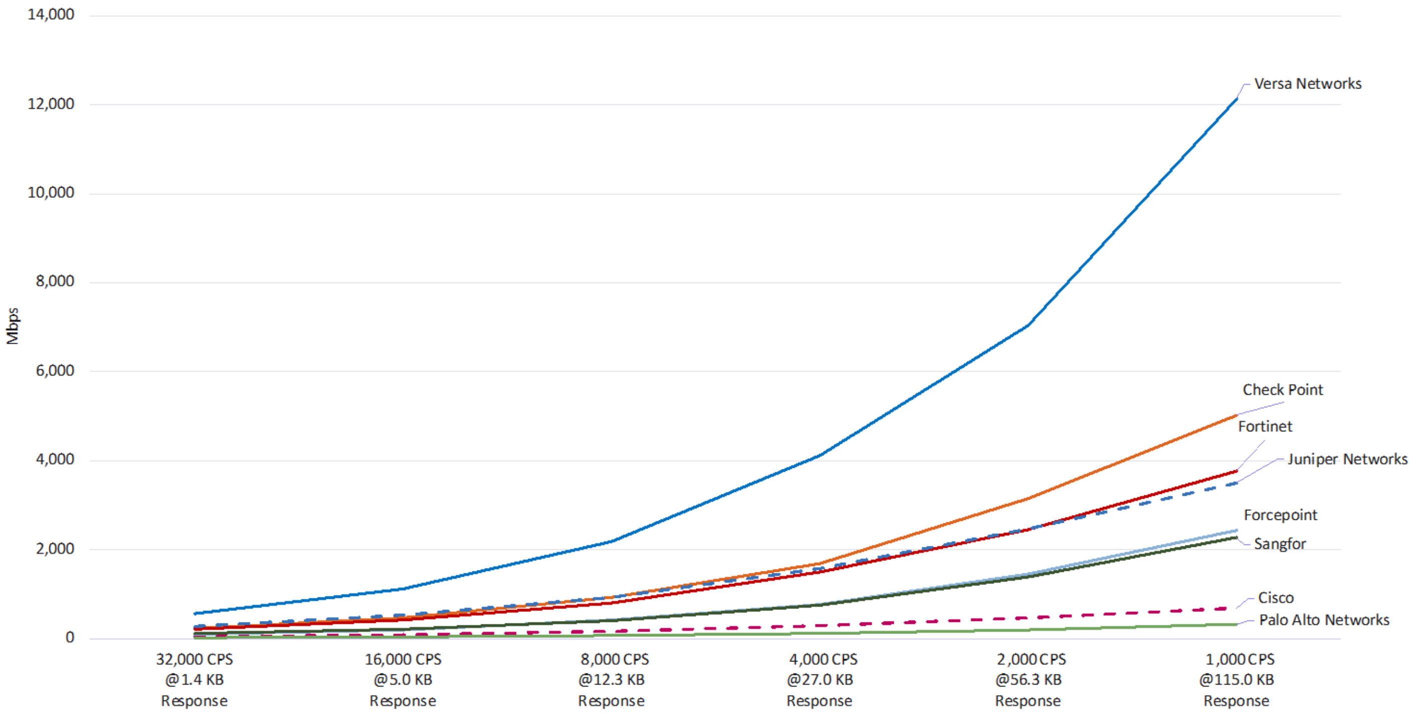

Figure 19 – HTTPS Capacity [TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xC0, 0x2F)]

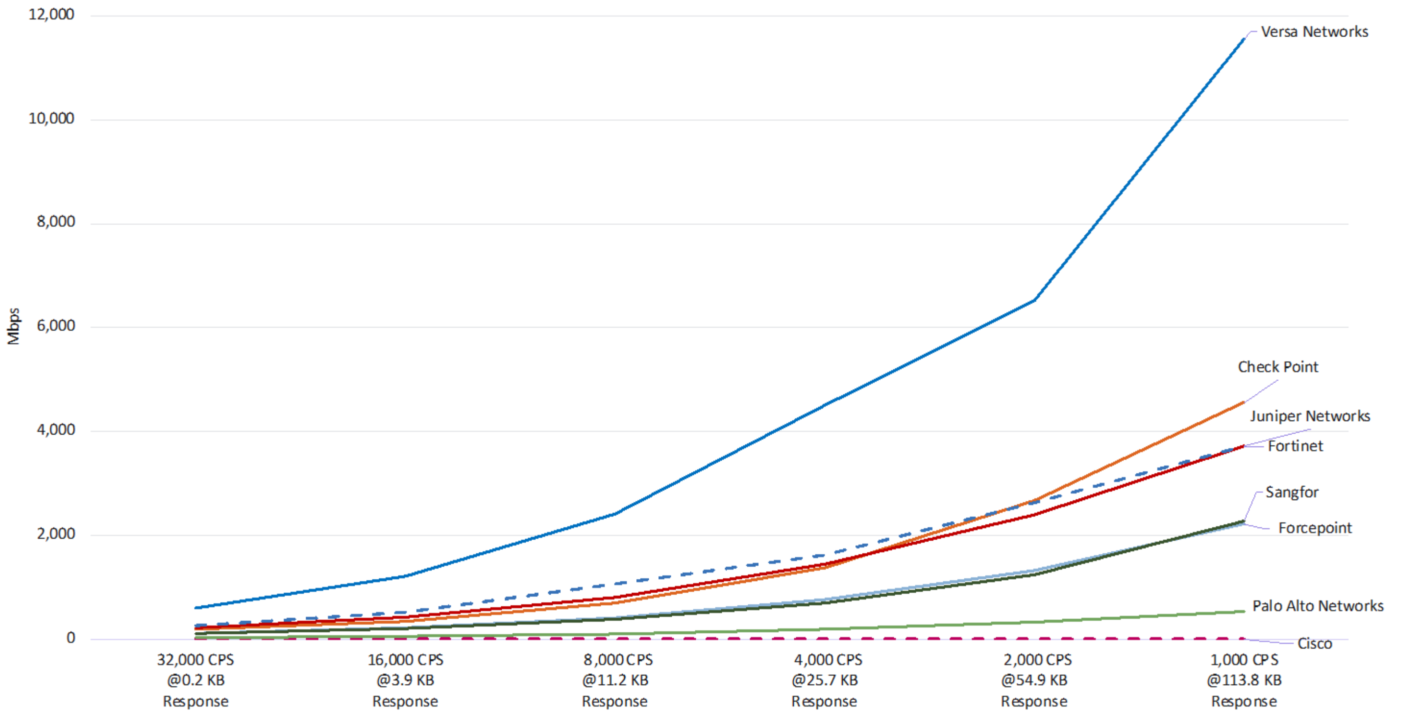

Figure 20 – HTTPS Capacity [TLS_AES_128_GCM_SHA256 (0x13, 0x01)]

"REAL-WORLD" SINGLE APPLICATION FLOWS

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

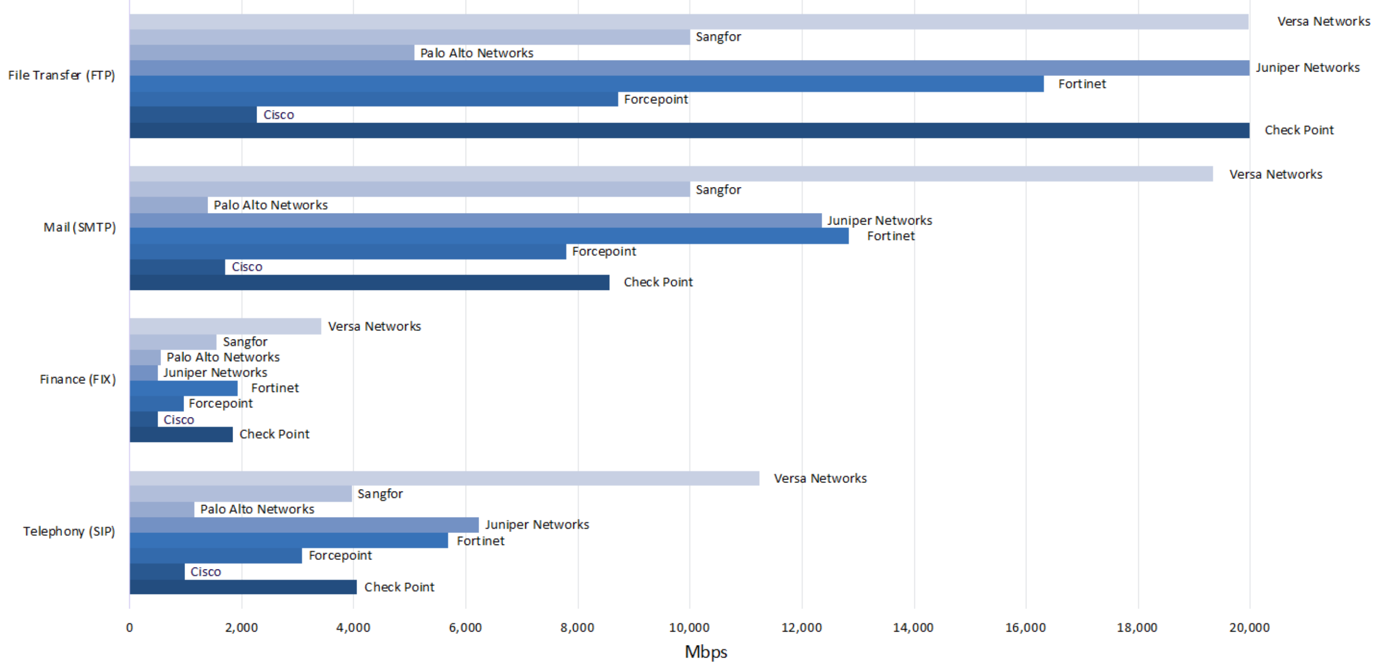

Figure 21 – "Real–World" Single Application Flows (I)

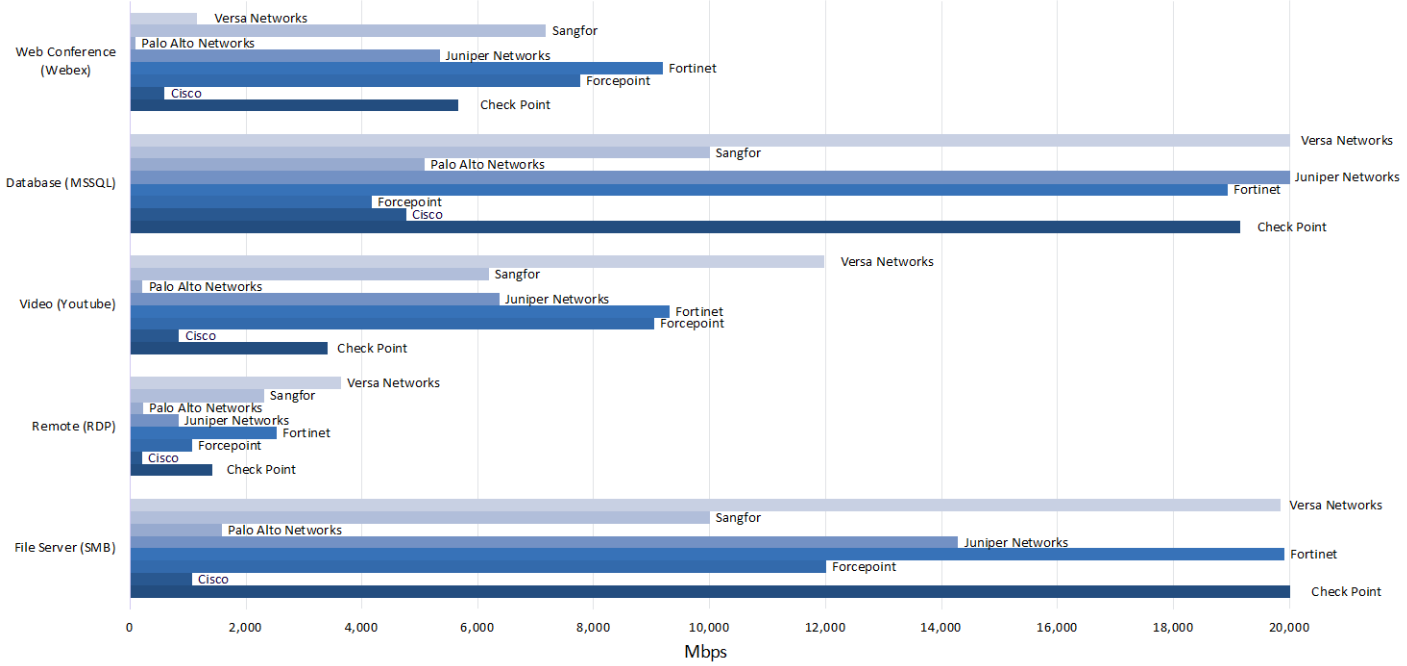

Figure 22 – "Real – World" Single Application Flows (II)

Price & Value

Now that we know how well the firewall performs from a security and performance perspective, there is almost always a budget to consider. For additional details on the cost of the product, please see the Individual test report at www.CyberRatings.org

One way to look at value is to think of it within the context of price/performance, or in this case, Price/Mbps. We have previously calculated the rated throughput for each product. Please see the performance section for more details.

Price per Mbps = Total Cost (3-Years)/ Rated Throughput (Mbps)

Using this formula, we can normalize data and account for wide-ranging price differences and performance among products.

Enterprise Firewall	Total Cost (3-Years)	Rated Throughput (Mbps)	Price per Mbps
Check Point	$149,970	5,438	$27.58
Cisco	$81,482	1,537	$53.01
Forcepoint	$32,915	4,235	$7.77
Fortinet	$40,532	11,383	$3.56
Juniper	$263,105	7,572	$34.75
Palo Alto Networks	$50,200	1,717	$29.25
Sangfor Technologies	$14,920	5,782	$2.58
Versa Networks	$30,851	12,160	$2.54

Figure 23 – Value (Price per Mbps)

Given that this is a security device, a low cost must be viewed within the context of security effectiveness. After all, an inexpensive device that only blocks 10 percent of attacks is not serving the purpose for which it was purchased; there is no value— performance matters, but not at the expense of security. Therefore, calculating a security device's value requires considering the relationship between price, performance, manageability, and security; we take the Price/Mbps and divide it by security effectiveness. Using our formula, a device that provides less security, i.e., 50%, will be twice as expensive as a device that offers 100% security. We have previously calculated the Security Effectiveness of each tested product; please see the Security Effectiveness section for more details.

Price per Protected Mbps = Total Cost (3-Years)/ (Rated Throughput (Mbps) x Security Effectiveness)

Enterprise Firewall	Total Cost (3-Years)	Rated Throughput	Price per Mbps	Security Effectiveness	Price per Protected Mbps
Check Point	$149,970	5,438	$27.58	98.14%	$28.10
Cisco	$81,482	1,537	$53.01	19.50%	$271.84
Forcepoint	$32,915	4,235	$7.77	99.48%	$7.81
Fortinet	$40,532	11,383	$3.56	99.88%	$3.57
Juniper	$263,105	7,572	$34.75	99.94%	$34.77
Palo Alto Networks	$50,200	1,717	$29.25	79.15%	$36.95
Sangfor Technologies	$14,920	5,782	$2.58	94.05%	$2.74
Versa Networks	$30,851	12,160	$2.54	99.48%	$2.55

Figure 24 – Value (Price per Protected Mbps)

How to Read the SVM

The SVM depicts the value of one enterprise firewall product using the list price. For larger deployments, some vendors offer significant discounts off the list price, and as a result, the outcome in value (X-Axis) could be very different. The Security Effectiveness will, however, remain the same for these types of attacks.

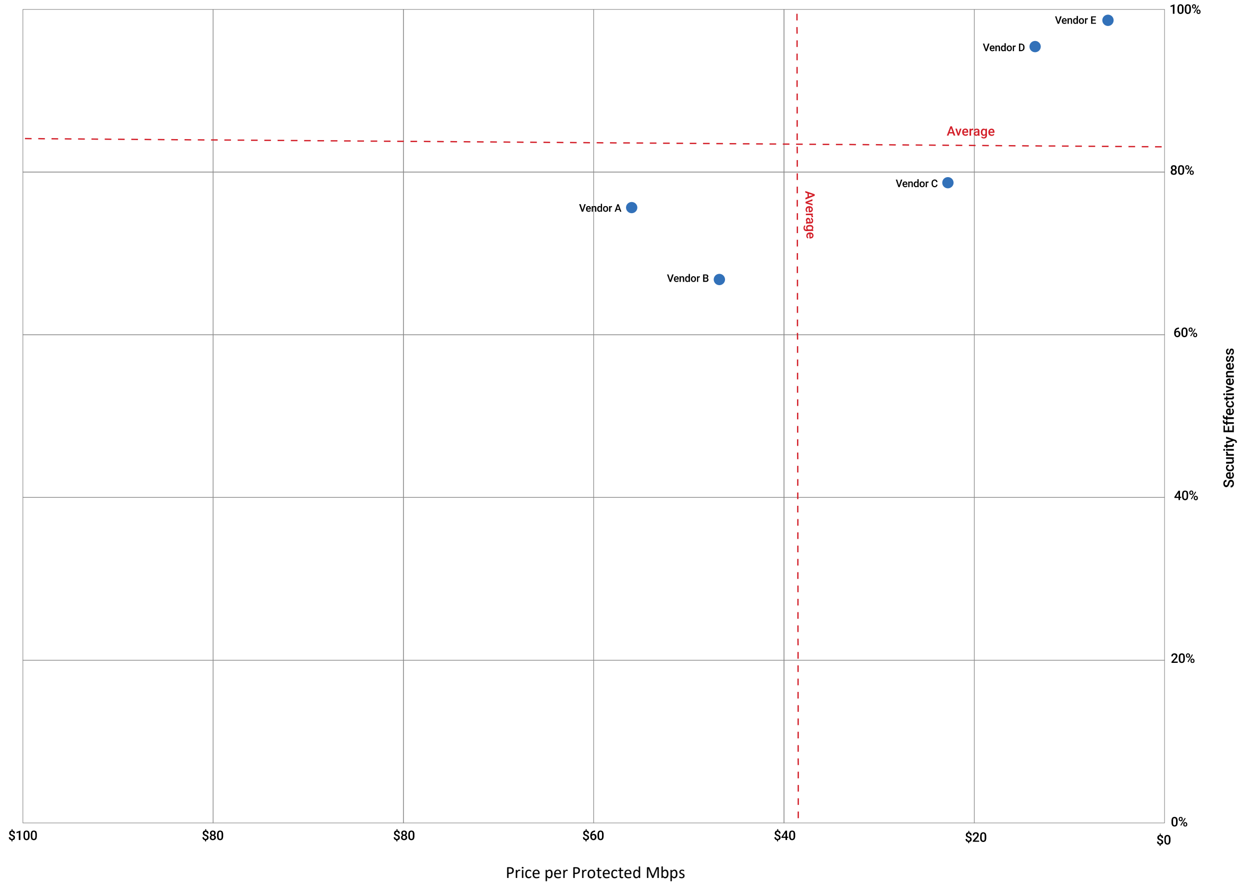

Figure 25 – Example SVM

No two security products deliver the same security effectiveness or performance, making precise comparisons extremely difficult. In order to enable value-based comparisons of enterprise firewall products on the market, we have developed a unique metric: *Price per Protected Mbps.*

The x-axis displays the *Price per Protected Mbps* in US dollars, which decreases from left to right. This metric incorporates the 3-Year cost with the *Security Effectiveness* score to provide a data point with which to compare the actual value of each product tested. The formula used is as follows: 3-Year Cost/ (*Security Effectiveness* x *Tested Throughput*). The Cost incorporates capital expenditure (capex) costs over three years, including initial acquisition and deployment costs and annual maintenance and update costs (software and hardware updates).

The y-axis displays the *Security Effectiveness* score as a percentage. *Security Effectiveness* is greater toward the top of the y-axis. Therefore, devices that are missing critical security capabilities will have a reduced *Security Effectiveness* score.

The SVM displays two dashed lines that represent the average for the *Security Effectiveness* and *Price per Protected Mbps* ratings of all the tested products. These lines divide the SVM into four unequally sized sections.

Where a product's *Security Effectiveness* and *Price per Protected Mbps* scores map on the SVM will determine which section it falls into:

- **Recommended**: Products that map into the upper-right section of the SVM score well for both *Security Effectiveness* and *Price per Protected Mbps***.** These products provide a high level of detection and value for money.
- **Caution:** Products that map into the lower-left section of the SVM offer limited value for money given their 3-Year Cost and measured *Security Effectiveness*.
- **Neutral:** Products that map into either the upper-left or lower-right sections may be good choices for organizations with specific security or budget requirements.

Neutral products in the upper-left section score above the average for *Security Effectiveness* but below the average for *Price per Protected Mbps* (*Security Recommended).* These products are suitable for environments requiring a high level of detection, albeit at a higher-than-average cost.

Conversely, *Neutral* products in the lower-right section score below the average for *Security Effectiveness* but above the average for *Price per Protected Mbps*. These products would be suitable for environments where a slightly lower level of detection is acceptable in exchange for a lower TCO.

PRODUCTS

Check Point Quantum QLS250 Lightspeed R81.20

Cisco Firepower 2130 v7.3.1-19

Forcepoint 2205 NGFW version 7.0.1.28052

Fortinet FortiGate 600F v6.4.12 build5431 (GA)

Juniper Networks SRX4600 22.3R1.12

Palo Alto Networks PA-3220 v10.2.3

Sangfor NGAF 5300 AF8.0.47.1004

Versa Networks CSG5000 versa-flexvnf-22.1.1-B

SPECIAL THANKS

We would like to issue a special thank you to Keysight for providing their CyPerf and Breaking Point tools for us to test Enterprise Firewall.

We would also like to thank TeraPackets for providing us with their Threat Replayer tool.

AUTHORS

Thomas Skybakmoen, Ahmed Basheer, Vikram Phatak

CONTACT INFORMATION

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

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.