



SOFTWARE-DEFINED WIDE AREA NETWORK TEST REPORT

Fortinet FortiGate 61E v6.0.4 GA Build 0231

June 19, 2019

Author – Ahmed Basheer

Fortinet FortiGate 61E V6.0.4 GA Build0231		
Summary	In Q1 2019 NSS Labs performed an independent test of the Fortinet FortiGate 61E v6.0.4 GA Build 0231. NSS has created three use cases to represent the most common reasons why enterprises deploy software-defined wide area network (SD-WAN) products: Manageability & Cost , Performance , and Security .	
Manageability & Cost	The ideal SD-WAN offers shorter deployment requirements, centrally managed networks, and a low total cost of ownership (TCO).	
	Zero-Touch Provisioning	
	Time to create configuration	03:31 minutes
	Time to deploy configuration	01:59 minutes
	Centralized Management System (CMS)	
	(Cloud)	
	Impact to network when CMS is offline	VoIP Mean Opinion Score (MOS): 4.41 Video MOS: 4.53
	Total Cost of Ownership	
	NSS-Tested model (3-Year TCO)	\$4,120
	TCO per Mbps	\$1
Performance	An ideal SD-WAN provides high Quality of Experience (QoE) for VoIP and Video, high throughput for VPN and cloud, and resilient WAN links and devices.	
	WAN Impairments	
	VoIP Mean Opinion Score (MOS)	4.33
	Video Mean Opinion Score (MOS)	4.00
	WAN Performance	
	NSS-Tested VPN Throughput	893 Mbps
	Split Tunnel (software as a service (SaaS) connectivity)	PASS
	High Availability (HA)	
	Power Fail	VoIP MOS: 4.39 Video MOS: 4.49
WAN Link Failover	VoIP MOS: 4.41 Video MOS: 4.53	
Security	Secure SD-WANs also provide built-in protection against network-delivered exploitation, which can reduce cost and complexity. The ideal SD-WAN will provide a high block rate.	
	Security	
	Built-in protection against network-delivered exploitation?	Yes
	Protection against network-delivered exploitation capabilities tested?	No
	Block Rate	NA
The product was subjected to thorough testing based on the Software-Defined Wide Area Network (SD-WAN) Test Methodology v2.0 (available at www.nsslabs.com). As with any NSS Labs group test, the test described in this report was conducted free of charge.		

¹ Performance testing was conducted between Branch 1 and the Headquarters (HQ) site over two established tunnels and was limited to 1,092 Mbps, as described in the NSS Labs Software-Defined Wide Area Network (SD-WAN) Test Methodology 2.0.

Table of Contents

SD-WAN Test Architecture	5
Quality of Experience (QoE).....	6
Mean Opinion Score (MOS).....	6
SD-WAN Use Cases	7
Use Case: Manageability & Cost	8
Zero-Touch Provisioning.....	8
Centralized Management.....	9
Total Cost of Ownership.....	10
TCO per Megabits per Second.....	10
Use Case: Performance	11
WAN Impairments.....	11
<i>Dynamic Path Selection and Path Conditioning</i>	11
<i>Quality of Service (QoS), Link Saturation, and Congestion</i>	11
<i>Application-Aware Traffic Steering</i>	11
<i>Mean Opinion Score (MOS) for VoIP and Video</i>	12
WAN Performance.....	13
<i>NSS-Tested VPN Throughput</i>	14
High Availability (HA).....	14
<i>Power Fail</i>	14
<i>WAN Link Failure</i>	15
Appendix A: Total Cost of Ownership	16
TCO Calculations.....	16
<i>Installation Hours</i>	16
<i>Total Cost of Ownership for Tested SD-WAN</i>	17
Appendix B: WAN Performance	18
Raw Packet Processing Performance (UDP Throughput and Latency).....	18
Maximum Capacity.....	19
HTTP Capacity.....	20
HTTP Capacity with HTTP Persistent Connections.....	21
Single Application Flows.....	22
Appendix C: Scorecard	23
Test Methodology	25
Contact Information	25

Table of Figures

Figure 1 – NSS Labs SD-WAN v2.0 Test Architecture	5
Figure 2 – Zero-Touch Provisioning: Time Measured	8
Figure 3 – CMS Offline (VoIP)	9
Figure 4 – CMS Offline (Video)	9
Figure 5 – 3-Year TCO for Tested Configuration (US\$)	10
Figure 6 – TCO per Mbps Formula	10
Figure 7 – TCO per Mbps	10
Figure 8 – WAN Impairments (VoIP).....	12
Figure 9 – WAN Impairments (Video).....	13
Figure 10 – Vendor-Claimed Throughput vs. NSS-Tested VPN Throughput (Mbps).....	14
Figure 11 – High Availability (VoIP).....	15
Figure 12 – High Availability (Video).....	15
Figure 13 – Installation Time (Hours)Total Cost of Ownership.....	16
Figure 14 – TCO for 1 Headquarters and 2 Branches with Branch 1 in HA Mode	17
Figure 15 – Raw Packet Processing Performance (UDP Traffic)	18
Figure 16 – Concurrency and Connection Rates.....	19
Figure 17 – HTTP Capacity	20
Figure 18 – HTTP Capacity HTTP with Persistent Connections	21
Figure 19 – Single Application Throughput over VPN.....	22
Figure 20 – Detailed Scorecard.....	24

SD-WAN Test Architecture

The SD-WAN test architecture was modeled with three use cases in mind: Manageability & Cost, Performance, and Security. The test architecture (see Figure 1) included two branch offices (Branch 1 and Branch 2) and one headquarters data center (HQ DC). Each branch location had two WAN links: a multiprotocol label switching (MPLS) link and an Internet service provider (ISP) connection. Additionally, Branch 1 was configured with a dedicated tunnel to a SaaS server and was set up to accommodate resilience or HA.

The ability of Branch 1 to forward traffic to the SaaS server was verified in a split tunnel test. This traffic should not impact traffic being forwarded between the two remaining WAN links. The WAN environment was provisioned with behavioral characteristics similar to those typically encountered over normal WAN link states, and the test harness baseline was recorded to ensure consistent behavior. The SD-WAN was deployed, and each test case was measured against the baseline. All tests were performed across VPN links established according to the use-case topology. The traffic flows used in this test were a mix of real-time, interactive, and bulk traffic. The MPLS link was set to 100 Mbps and the ISP link was set to 1 Gbps (maximum achievable throughput in this setup for the products tested was 1,092 Mbps without including overhead).

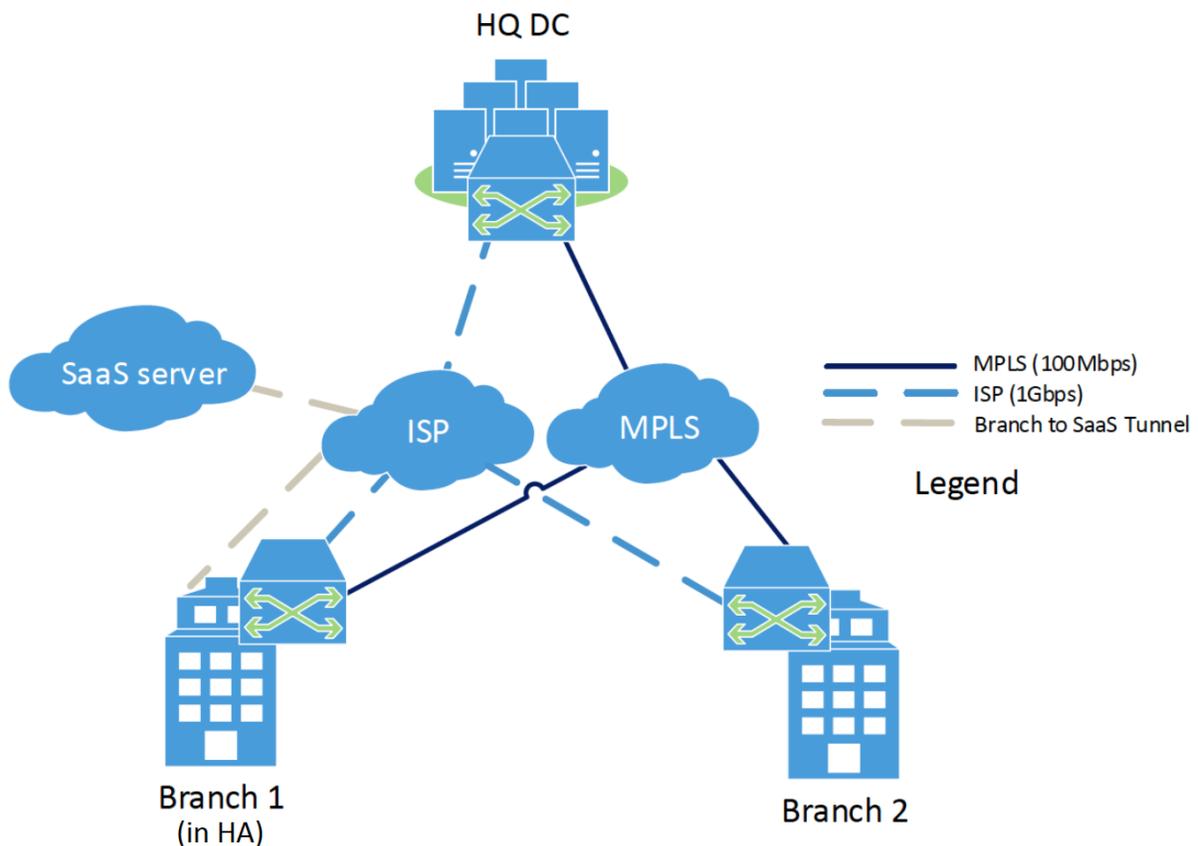


Figure 1 – NSS Labs SD-WAN v2.0 Test Architecture

Quality of Experience (QoE)

While throughput is important in an SD-WAN, so is the user's QoE. A critical function of any SD-WAN is the correct identification and routing of traffic based on policy prioritization (autonomous or configured), which is influenced by network performance characteristics (e.g., congestion, packet loss, latency, packet delay variation, etc.). Link impairment tests subject connected links to testing that represents real-world conditions encountered by enterprises today. Congestion, packet loss, latency, and packet delay variation are all commonly encountered on public links.

Mean Opinion Score (MOS)

The MOS is used to calculate the QoE that enterprises can expect when deploying SD-WAN products. Relative (video) MOS is an estimated perceptual quality score that considers the effects of codec; the impact of IP impairments (such as packet loss) on the group of pictures (GoP) structure and video content; and the effectiveness of loss concealment methods. The encoding specifications for video codec are used as guidelines and conformance, and vendors are free to design encoders to improve video quality and reduce the number of transmission bits. Simply put, MOS for video (relative MOS) can vary based on different advancements in the video estimation or encoding techniques. In the video used for the test, the maximum achievable MOS was 4.53.

VoIP (real-time protocol [RTP]) MOS, on the other hand, measures the mean opinion score for VoIP calls based on the speech codec being used. The setup used a G711 codec, which produces a maximum achievable MOS of 4.41 for an excellent VoIP call.

Any MOS below 3.5 represents a significantly degraded voice call or video stream. NSS considers a MOS below 3.4 as failing to meet the use case. In each test case, background traffic was introduced to populate links with sufficient activity so as to represent typical enterprise network communications. These measurements provide guidance on the behavior of sensitive applications in a network that uses SD-WANs.

SD-WAN Use Cases

Enterprise network traffic requirements continue to grow as applications, content, and hardware evolve. Reliable packet delivery between branch offices and headquarters is critical for business continuity. By utilizing common VPN capabilities and by separating data and control planes within software-defined networks (SDNs), SD-WANs enable enterprises to leverage high-bandwidth, consumer-grade links (often links without guaranteed performance) for business-class services at a lower cost than traditional dedicated links. Enterprises are adopting SD-WANs for their branch office network needs, capitalizing on the visibility, scalability, performance, and control benefits the technology provides.

Reasons for SD-WAN deployments often vary greatly between industries and enterprises. In order to represent the most common reasons enterprises deploy SD-WAN, NSS has created three use cases:

- Manageability & Cost
- Performance
- Security

NSS tested protection against network-delivered exploitation capabilities for vendors offering built-in security functionality. The Barracuda product, however, was tested with protection against network-delivered exploitation capabilities disabled since the SD-WAN methodology v2.0 states that security testing is optional.

Detailed information for each of the use cases can be found in the relevant use case appendix.

Use Case: Manageability & Cost

This use case is made up of three components:

- Zero-touch provisioning (ZTP)
- Centralized management
- Total cost of ownership (TCO)

Zero-Touch Provisioning

Legacy network technologies such as MPLS are highly inflexible. In some cases, deploying a new site or location can take over 100 days, drastically impacting business expansion and productivity.

- **Reasons for Adoption:** SD-WANs are faster to configure and deploy than legacy network solutions
- **Expected Outcome:** Simplified configurations and shorter deployment times can facilitate business expansion and increase productivity, which reduces cost
- **What Was Tested:** ZTP was measured to assess how quickly the SD-WAN could be set up and deployed

A configuration was created for the SD-WAN deployment using the CMS. Branch 2 was connected to the CMS and then added to the network. NSS measured the following:

- Time to create configuration
- Time to deploy configuration

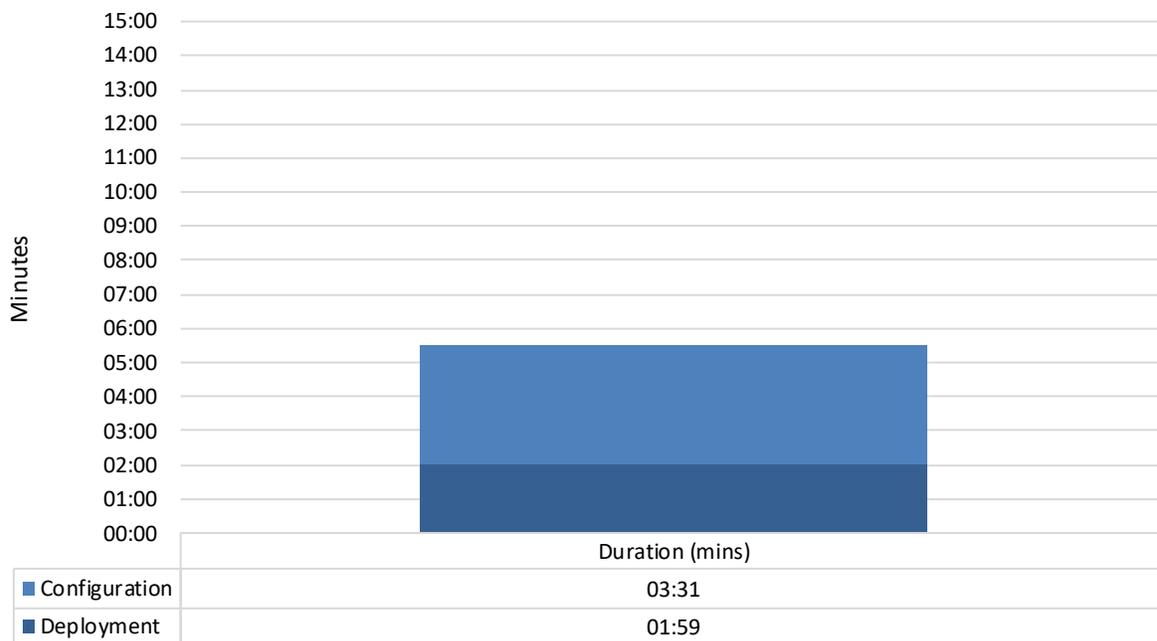


Figure 2 – Zero-Touch Provisioning: Time Measured

The *Time to Create Configuration* metric depicts the time it took to create a new configuration or clone from an existing configuration, apply updates where necessary, and add a new branch to the existing network topology. This includes creation of a template configuration, updating IP addresses for management interfaces, creating VPN

tunnels for the two WAN links (MPLS and ISP) between the new site and existing sites, and setting up thresholds and traffic policies to allow or deny traffic.

The *Time to Deploy Configuration* metric depicts the time it took to deploy the configuration. This includes connecting to the CMS, selecting the appropriate configuration to be deployed, validating for errors/issues, and provisioning it for a desired site.

Centralized Management

SD-WANs offer a CMS that provides enterprises with the ability to centrally manage and configure devices via a graphical user interface (GUI). CMS capabilities include but are not limited to monitoring, reporting, configuration changes, and the ability to update software on the SD-WAN. The ability of the network to operate autonomously in the event of a loss of connection to the central management application or controller, however, is a concern.

- **Reasons for Adoption:** Network continues to operate autonomously even if connectivity with CMS is lost
- **Expected Outcome:** Prevents network outage when the CMS is offline, which avoids revenue loss and impact to partners and customers
- **What Was Tested:** CMS outage was tested to verify if service interruption occurred as a result of loss of connectivity with CMS

During traffic flow, the CMS was disconnected or made unreachable. Such an event should not impact established links and ideally should not affect the MOS for VoIP or Video. A score of 4.41/4.53 represents a maximum achievable MOS for VoIP or Video, while any score below 3.5 represents a significantly degraded voice call or video stream.

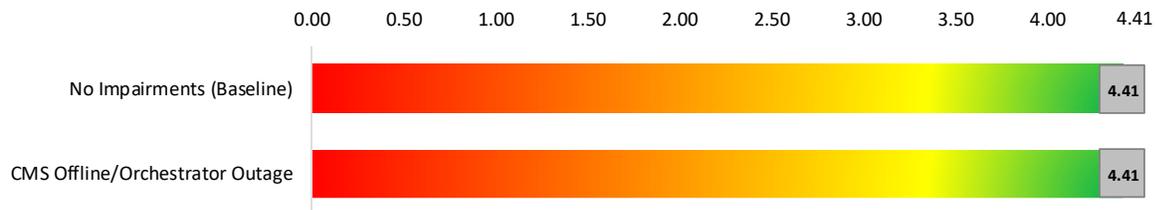


Figure 3 – CMS Offline (VoIP)

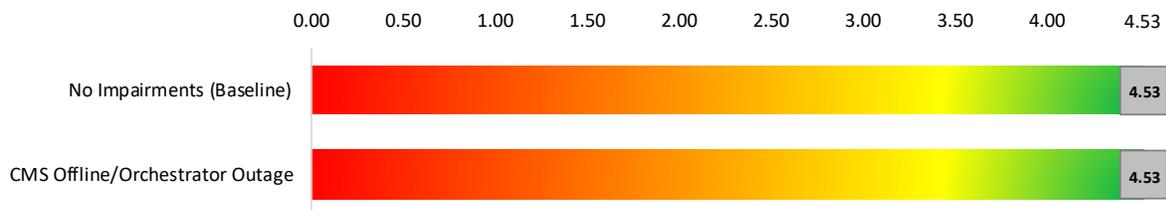


Figure 4 – CMS Offline (Video)

Total Cost of Ownership

One of the major factors motivating enterprises to move to SD-WAN technology is that it helps businesses reduce cost per throughput compared to MPLS by efficiently managing unreliable WAN links such as public Internet (ISP) or broadband. A global SD-WAN, provided as a fully managed service, will not only ensure that companies save on network maintenance and management costs, but will also free up IT resources so more strategic business initiatives can be focused on.

- **Reasons for Adoption:** SD-WANs are expected to reduce cost per throughput compared to MPLS networks
- **Expected Outcome:** Better performance-to-cost ratio than MPLS
- **What Was Tested:** Vendors were asked to submit costs for the following configuration: NSS-Tested SD-WAN (1 HQ DC +2 Branches, 1 Branch with HA).

Figure 5 depicts the TCO for the tested configuration.

For additional details on the test architecture, please reference Figure 1. For additional details on how the TCO was calculated, please see Figure 14.

Product	3-Year TCO
Fortinet FortiGate 61E v6.0.4 GA Build 0231	\$4,120

Figure 5 – 3-Year TCO for Tested Configuration (US\$)

TCO per Megabits per Second

To normalize data and account for wide-ranging differences between SD-WANs and MPLS networks, enterprises can use NSS’ *TCO per Megabits per Second (Mbps)* formula. To calculate *TCO per Mbps*, NSS used the *NSS-Tested VPN Throughput* metric (see the section on *NSS-Tested VPN Throughput* for more details) and the 3-Year TCO for the tested configuration.

The *TCO per Mbps* metric provides guidance as to whether the price for an SD-WAN is higher or lower than an alternative such as MPLS. A high price could indicate a premium based on features, bandwidth, service level agreements (SLA), brand recognition, or level of customer service. Conversely, a high price could also be a penalty for purchasing an underperforming product.

$$TCO \text{ per Mbps} = \frac{3\text{-Year TCO} / \text{NSS-VPN Tested Throughput (Mbps)}}{\text{Number of Systems purchased}}$$

Figure 6 – TCO per Mbps Formula

Product	TCO per Mbps
Fortinet FortiGate 61E v6.0.4 GA Build 0231	\$1

Figure 7 – TCO per Mbps

Use Case: Performance

This use case is made up of three components:

- WAN impairments
- WAN performance
- HA

WAN Impairments

MPLS technology is widely known for its performance, but it comes with some legacy network shortfalls, such as long deployment timelines, high cost, and lack of optimized cloud access support. The public Internet is known to provide poor application performance due to high packet loss and fluctuating latencies. Today, enterprise networks rely on the public Internet (in the absence of MPLS) for connectivity. However, enterprises expect SD-WANs, which utilize both MPLS and the public Internet, to overcome these drawbacks and provide fast and reliable application performance.

- **Reasons for Adoption:** Overcome performance issues found on the public Internet due to packet loss, packet delay variation, and high latency
- **Expected Impact:** Reduced network congestion, link saturation, packet loss, latency, etc.
- **What Was Tested:** SD-WAN tested with various network impairments seen in real-world networks

Dynamic Path Selection and Path Conditioning

This test was conducted to determine how long it took for traffic to move to an available link when preconfigured impairments were applied. To limit any visible user impact, an SD-WAN should support path decisions based on the conditions that exist on those links. The time to select a new path was measured, as was any impact to applications.

SD-WANs employ various techniques to condition WAN links in order to ensure the reliability of data transmission. Some SD-WANs employ packet duplication, forward error correction, bonding, or load balancing. An SD-WAN should identify the best path and guarantee priority policies (application, protocol, or other configured guidance) over known good links with other traffic transmitted as best effort.

Quality of Service (QoS), Link Saturation, and Congestion

Global quality of service (QoS) awareness can prevent congestion during the last mile of data delivery; therefore, the aim of this test was to ensure reliable use of bandwidth by the controller in the SD-WAN. QoS is important for business-critical applications such as VoIP and video. These applications must be prioritized if a link has bad performance indicators. This test measured QoS using voice traffic and video stream. The test included MOS for video and call measurements for VoIP. An SD-WAN should manage traffic according to configured QoS classification settings.

Application-Aware Traffic Steering

This test was conducted to verify how the SD-WAN directs various application traffic flows for applications other than video and VoIP. Behavior was observed and recorded to establish whether VoIP/video and data were sent over the same link once impairments were applied and to determine which application took precedence.

Mean Opinion Score (MOS) for VoIP and Video²

A score of 4.41/4.53 represents the maximum achievable MOS value for VoIP/video, while any score below 3.5 represents a significantly degraded voice call/video stream. Figure 8 depicts the weighted average impairments that NSS expects an SD-WAN to experience in an enterprise environment.

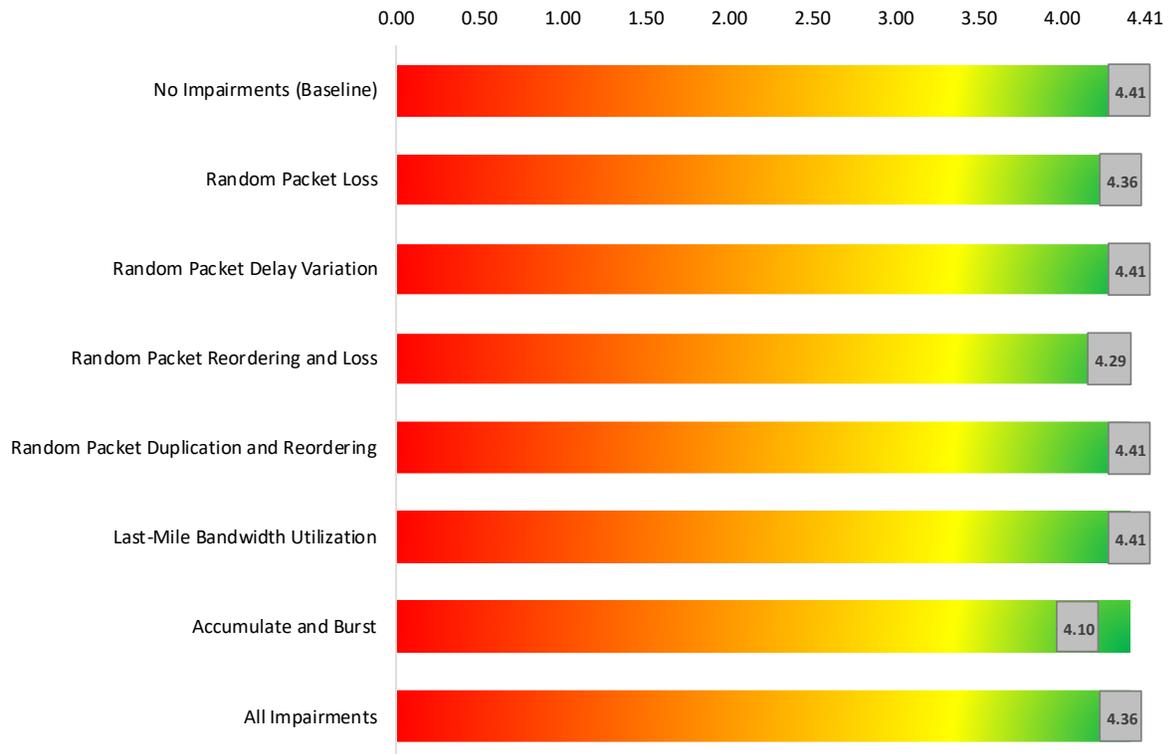


Figure 8 – WAN Impairments (VoIP)

² ITU-T P.800.1: <http://handle.itu.int/11.1002/1000/12972>

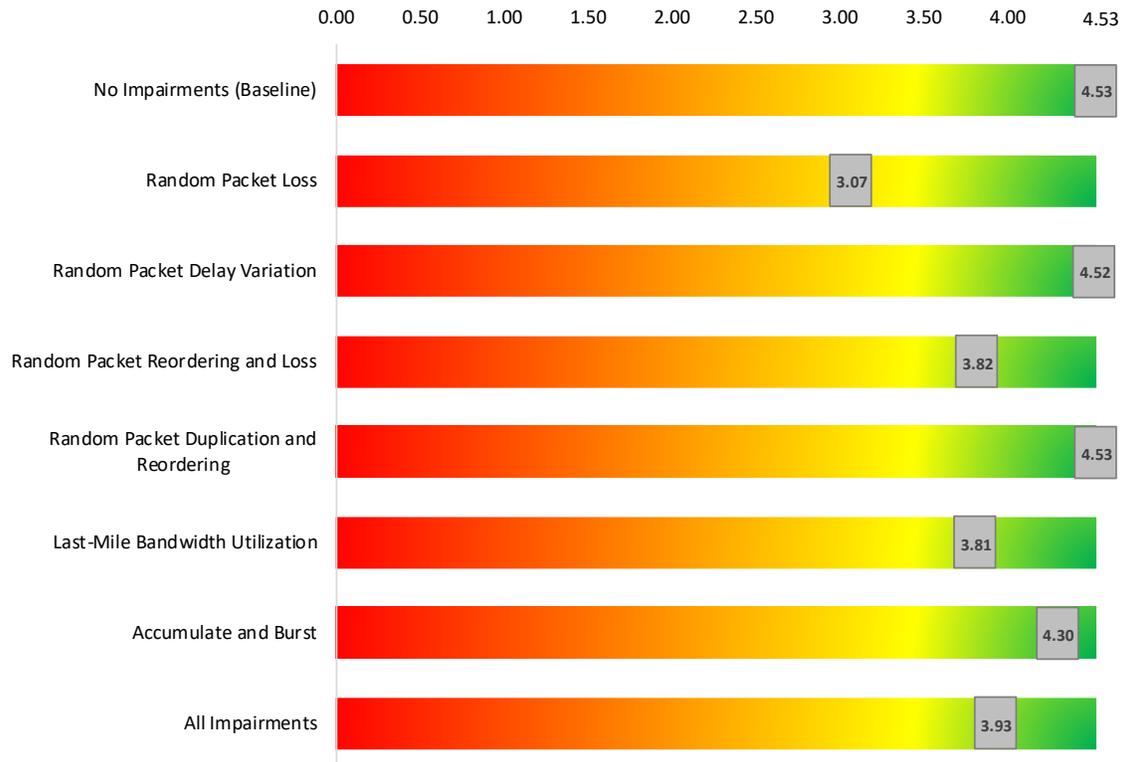


Figure 9 – WAN Impairments (Video)

WAN Performance

A growing demand for high-performance networks and a shift toward cloud-based applications has made it essential for SD-WANs to provide better performance and features that support cloud-based applications such as cloud gateways, split tunneling, etc. Split tunneling enables branch traffic to be routed directly to the public Internet rather than being routed through the HQ. This means latency at the branches is reduced, and there is less network traffic to process at HQ.

- **Reasons for Adoption:** A growing dependency on cloud-based applications means quick, direct access to the cloud is required
- **Expected Impact:** Faster Internet access, better connectivity to the cloud, more efficient traffic handling
- **What Was Tested:** Tested split tunnel to access SaaS server; measured throughput for various applications

NSS-Tested VPN Throughput

Figure 10 depicts the difference between *NSS-Tested VPN Throughput* and vendor performance claims. Vendor tests are often performed under ideal or unrealistic conditions and this provides a basis for comparison. 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 VPN Throughput* typically is lower than vendor-claimed throughput.

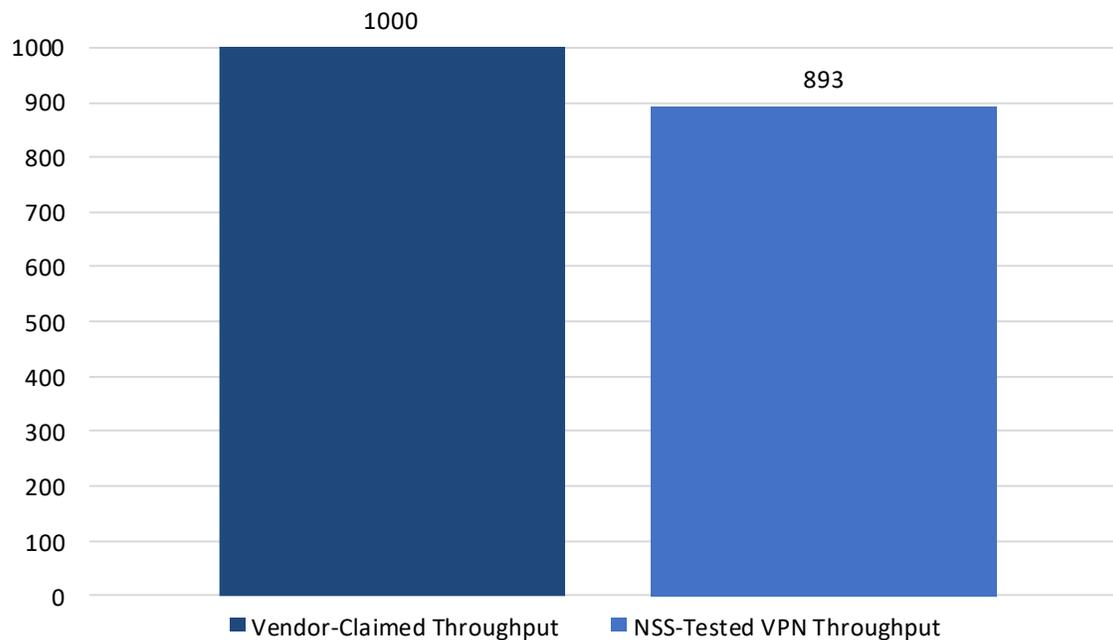


Figure 10 – Vendor-Claimed Throughput vs. NSS-Tested VPN Throughput (Mbps)

NSS-Tested VPN Throughput is calculated as a weighted average of the traffic that NSS expects an SD-WAN to experience in an enterprise environment. For additional analysis on performance, please refer to Appendix B: WAN Performance (Additional Analysis).

High Availability (HA)

HA is necessary for network stability. While enterprises often see this as an unnecessary cost rather than an essential component, businesses, their partners, and their customers require and expect it in their networks. HA provides redundancy and in turn a more stable and consistent networking experience.

- **Reasons for Adoption:** Businesses require HA for network and application redundancy/stability
- **Expected Impact:** HA can prevent downtime in a network, which impacts productivity and revenue
- **What Was Tested:** Power fail and WAN link failover

Power Fail

In this test, the SD-WAN was subjected to a device power loss event for Branch 1 (configured with HA). Branch-to-HQ communication was expected to persist with the help of the redundant device on Branch 1. NSS used the persistence of data test to validate whether or not the device was able to retain all configuration data, policy data, and locally logged data once it was restored to operation following power failure.

WAN Link Failure

In this test, an established link (ISP) between sites was interrupted, and the SD-WAN was observed to determine whether it was handling stateful session in a manner that was transparent to users. At the point of failure, routed link traffic should be redirected without loss or interruption to the applications using the link based on prioritization schema. A WAN must be able to operate resiliently in spite of link outages. Sessions or applications should be able to continue without interruption and there should be no noticeable user impact.

In a redundant network, the SD-WAN deployed in HA mode is expected to provide uninterrupted network service during any system failures while delivering an acceptable user experience. A score of 4.41/4.53 represents a maximum achievable MOS for VoIP/video, while any score below 3.5 represents a significantly degraded voice call/video stream.

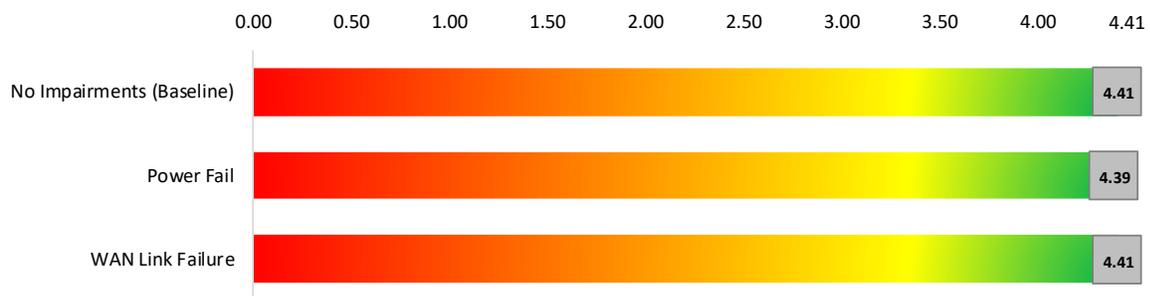


Figure 11 – High Availability (VoIP)

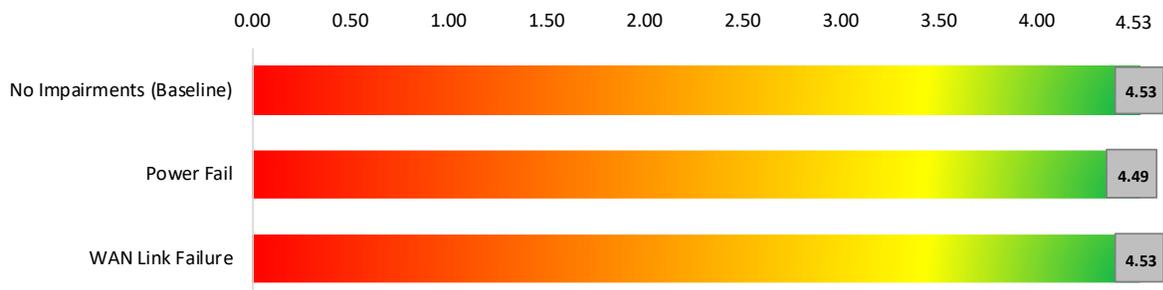


Figure 12 – High Availability (Video)

Appendix A: Total Cost of Ownership

This section offers additional information TCO, including the method for calculating the TCO of the NSS-Tested configuration.

Implementation of infrastructure and security solutions can be complex, with several factors affecting the overall cost of deployment, maintenance, and upkeep. All of the following should be considered over the course of the useful life of the SD-WAN:

- **Product Purchase** – The cost of acquisition
- **Operational Benefits** – The zero-touch provisioning concept for SD-WAN cites considerably reduced deployment requirements, specifically regarding configuration and tuning; for example, time to add a new site is measured in hours rather than days or weeks. These reduced configuration requirements contribute to operational savings for the enterprise.
- **ROI Assessment** – There are savings associated with moving from high-cost, service-assured links (e.g., MPLS) to commercial broadband. There is value both in aggregating multiple low-cost links to support demand as well as in the ease of deployment and recurring service cost reductions that are associated with moving from expensive, service-assured links to less expensive options.
- **Product Maintenance** – The fees paid to the vendor, including software and hardware support, maintenance, and other updates
- **Installation** – The time required to take the product 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 the vendor, including hardware, software, and other updates

TCO Calculations

Installation Hours

Figure 13 depicts the number of hours of labor required to install the configuration depicted in Figure 1. It reflects the amount of time that NSS engineers, with the help of vendor engineers, needed to install and configure the SD-WANs to the point where they operated successfully in the test harness, passed legitimate traffic, and if applicable, blocked and detected any prohibited or malicious traffic. Installation cost is based on the time that an experienced engineer would require to perform the tasks described. This approach allows NSS to hold constant the talent cost and measure only the difference in time required for installation. Readers should substitute their own costs to obtain accurate TCO figures.

Product	Installation (Hours)
Fortinet FortiGate 61E v6.0.4 GA Build 0231	8

Figure 13 – Installation Time (Hours)Total Cost of Ownership

Total Cost of Ownership for Tested SD-WAN

Product	Year 1 Cost	Year 2 Cost	Year 3 Cost	3-Year TCO
Fortinet FortiGate 61E v6.0.4 GA Build 0231	\$3,608	\$256	\$256	\$4,120

Figure 14 – TCO for 1 Headquarters and 2 Branches with Branch 1 in HA Mode

Calculations are based on vendor-provided pricing information. Where possible, the 24/7 maintenance and support option with 24-hour replacement is utilized to calculate maintenance/support fees, since this is the option typically selected by enterprise customers. Prices are for SD-WAN devices and maintenance only (as shown in Figure 14); costs for CMS may be extra. In detail:

- **Year 1 Cost** is calculated by adding installation costs (US\$100 per hour fully loaded labor x installation time) + purchase price of devices + first-year maintenance/support fees
- **Year 2 Cost** consists only of maintenance/support fees
- **Year 3 Cost** consists only of maintenance/support fees

For additional TCO analysis and details, refer to the TCO Comparative Report.

Appendix B: WAN Performance

This section provides additional performance tests that are not included in the *NSS-Tested VPN Throughput* average.

Raw Packet Processing Performance (UDP Throughput and Latency)

This test used 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 was transmitted bidirectionally across the WAN links. The percentage load and frames per second (fps) figures across the WAN links were verified by network monitoring tools before each test began. Multiple tests were run and averages were taken where necessary.

The goal of this test was to determine the raw packet processing capability of each inline port pair of the SD-WAN, as well as its effectiveness at forwarding packets quickly in order to provide the highest level of network performance and with the lowest latency. SD-WANs that introduce high levels of latency lead to unacceptable response times for users, especially where multiple security devices are placed in the data path. Every performance test was conducted over the established tunnels between HQ DC and Branch 1. This test helps in understanding the true throughput values based on packet sizes and to demonstrate how good results may be achieved simply by using large packets.

Larger packet sizes may be fragmented due to encryption and tunneling across WAN links. Some solutions will reassemble, and others will pass IP fragments to the end station. Figure 15 depicts UDP throughput (in Mbps) and latency (in milliseconds) as recorded during the UDP throughput tests at 90% of maximum load.

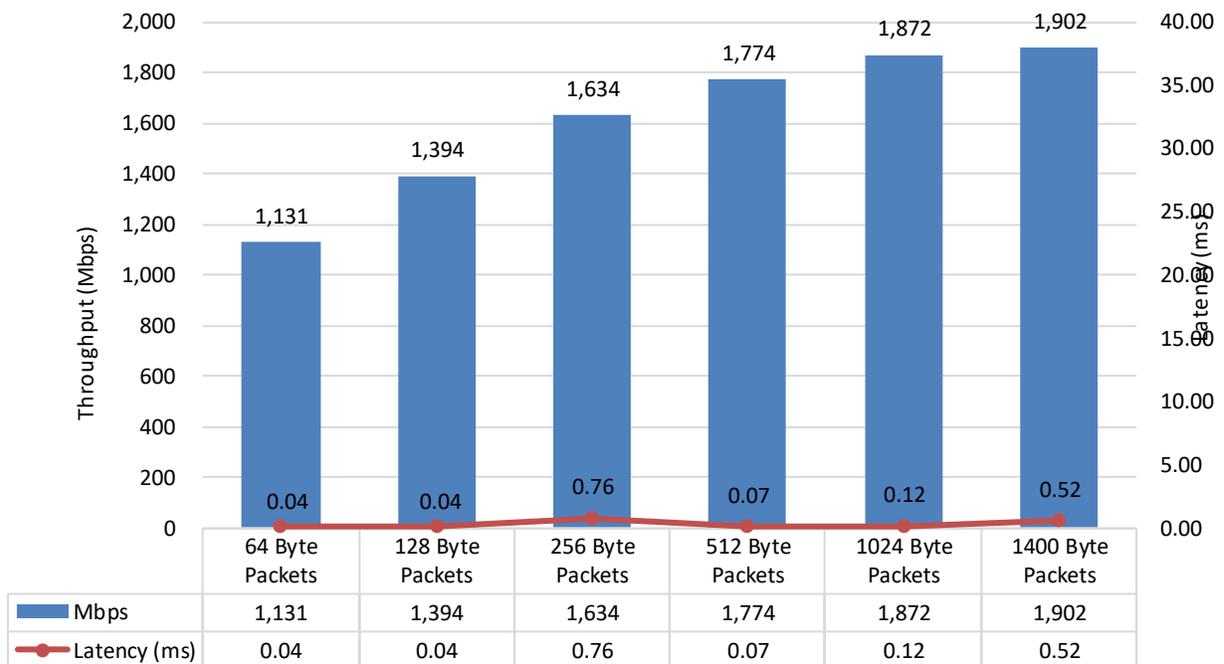


Figure 15 – Raw Packet Processing Performance (UDP Traffic)

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. Where applicable, 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 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 was causing an unacceptable increase in open connections.
- **Excessive concurrent HTTP connections** – Latency was causing excessive delays and increased response time.
- **Unsuccessful HTTP transactions** – Normally, there should be zero unsuccessful transactions. Once these appeared, it was an indication of excessive latency, which caused connections to time out.

Maximum Capacity	Connections per Second (CPS)
Max Concurrent TCP Connections	599,497
Max TCP Connections per Second	20,920
Max HTTP Connections per Second	20,620
Max HTTP Transactions per Second	176,100

Figure 16 – Concurrency and Connection Rates

HTTP Capacity

The aim of the HTTP capacity tests was to stress the HTTP detection engine and determine how the SD-WAN 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 SD-WAN was forced to track valid HTTP sessions, thus ensuring a higher workload than for simple packet-based background traffic.

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.

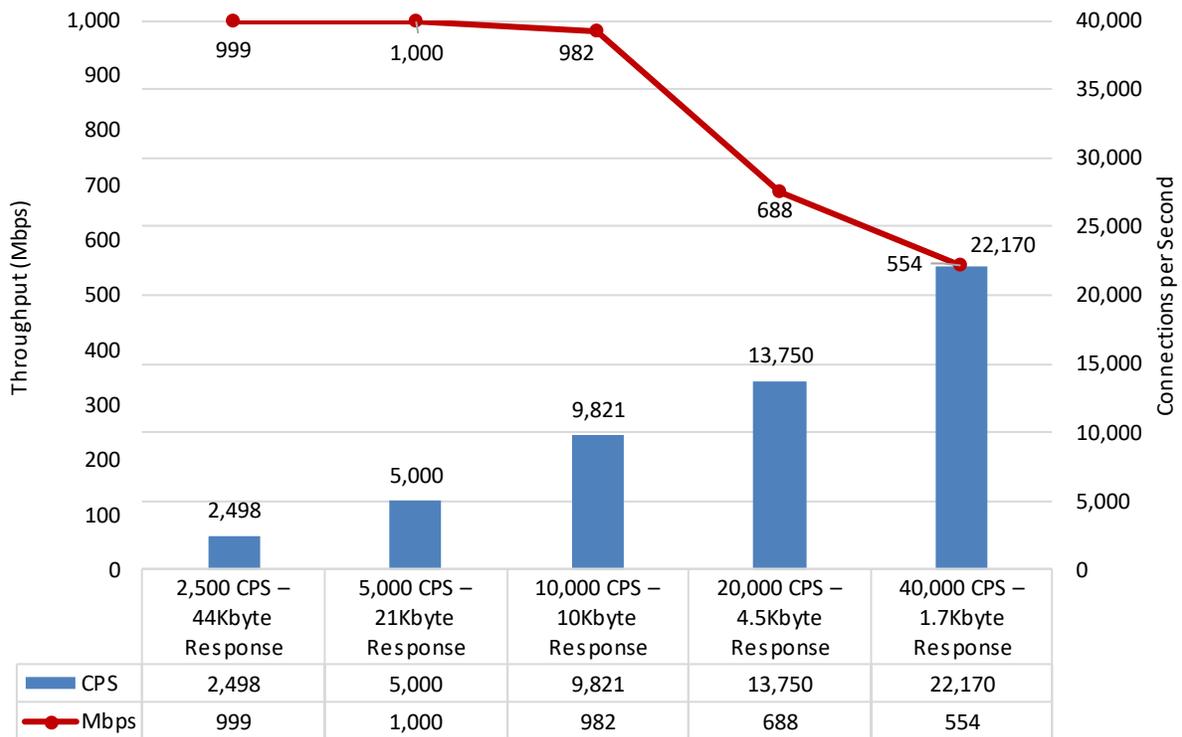


Figure 17 – HTTP Capacity

HTTP Capacity with HTTP Persistent Connections

This test used HTTP persistent connections, with each TCP connection containing 10 HTTP GETs and associated responses. 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 at various network loads. The stated response size was the total of all HTTP responses within a single TCP session.

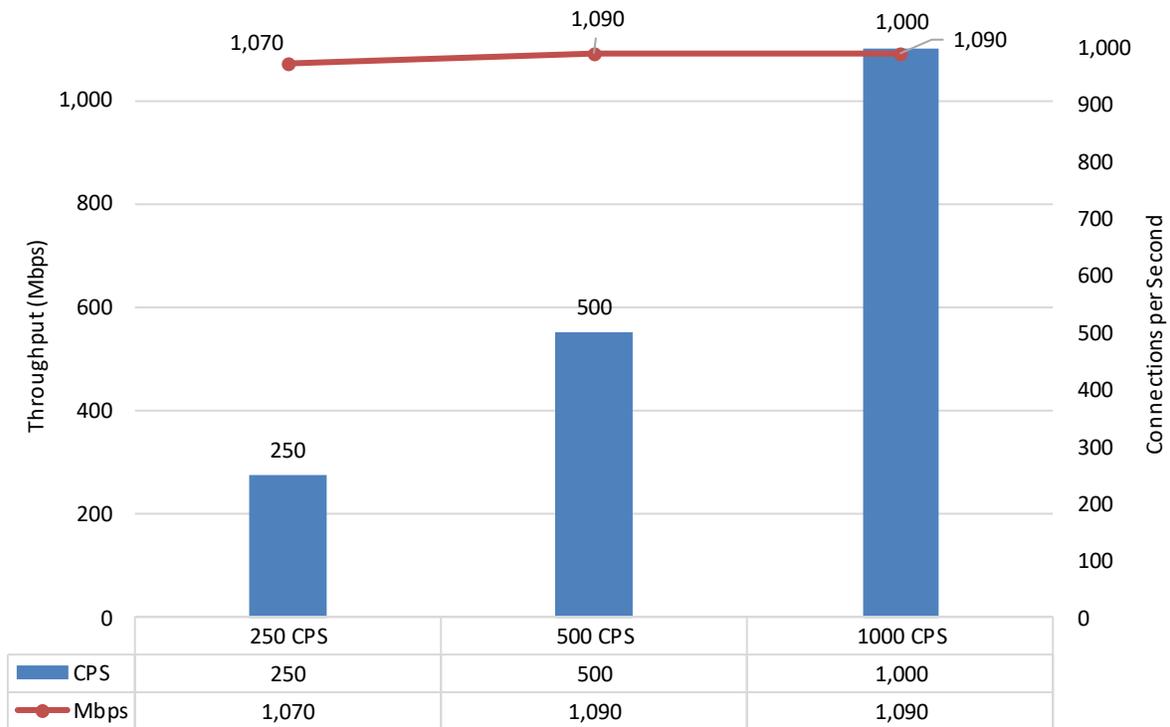


Figure 18 – HTTP Capacity HTTP with Persistent Connections

Appendix C: Scorecard

Description		Result		
Remote Initial Configuration				
Feature	Description	Duration (s)		
Zero-Touch Provisioning (ZTP)	Time to Create Configuration	211		
	Time to Deploy Configuration	119		
WAN Impairment				
Impairment Test	WAN Impairment	FTP Connection Latency (ms)	VoIP MOS	Video MOS
Baseline	No impairments	0.0	4.41	4.53
Dynamic Path Selection & Path Conditioning	Random Packet Loss	6.6	4.36	3.07
	Random PDV	9.6	4.41	4.52
	Random Packet Reordering & Loss	4.4	4.29	3.82
	Random Packet Duplication & Reordering	0.0	4.41	4.53
Quality of Service, Link Saturation & Congestion	Last-Mile Bandwidth Utilization	0.10	4.41	3.81
	Accumulate & Burst	136.4	4.10	4.30
	All impairments	2.3	4.36	3.93
High Availability	Power Fail	0.0	4.39	4.49
	WAN Link Failure (Blackout)	0.1	4.41	4.53
Centralized Management	CMS Offline/Orchestrator Outage	0.0	4.41	4.53
Performance				
Raw Packet Processing Performance (UDP Traffic)		Weighting for NSS-Tested VPN Throughput		Mbps
64-Byte Packets		0%		1,131
128-Byte Packets		1%		1,394
256-Byte Packets		1%		1,634
512-Byte Packets		1%		1,774
1024-Byte Packets		3%		1,872
1400-Byte Packets		3%		1,902
Latency – UDP				Microseconds
64-Byte Packets				0.04
128-Byte Packets				0.04
256-Byte Packets				0.76
512-Byte Packets				0.07
1024-Byte Packets				0.12
1400-Byte Packets				0.52
Maximum Capacity				CPS
Theoretical Max. Concurrent TCP Connections				599,497
Maximum TCP Connections per Second				20,920
Maximum HTTP Connections per Second				20,620
Maximum HTTP Transactions per Second				176,100

HTTP Capacity	Weighting for NSS-Tested VPN Throughput	CPS
2,500 Connections per Second – 44 KB Response	8%	2,498
5,000 Connections per Second – 21 KB Response	8%	5,000
10,000 Connections per Second – 10 KB Response	7%	9,821
20,000 Connections per Second – 4.5 KB Response	7%	13,750
40,000 Connections per Second – 1.7 KB Response	4%	22,170
Application Average Response Time – HTTP (at 90% Max Load)		Milliseconds
2,500 Connections per Second – 44 KB Response		8.29
5,000 Connections per Second – 21 KB Response		2.40
10,000 Connections per Second – 10 KB Response		0.85
20,000 Connections per Second – 4.5 KB Response		0.09
40,000 Connections per Second – 1.7 KB Response		0.08
HTTP Capacity with HTTP Persistent Connections		CPS
250 Connections per Second		250
500 Connections per Second		500
1000 Connections per Second		1,000
Single Application Flows	Weighting for NSS-Tested VPN Throughput	Mbps
Telephony	17%	865
Email	12%	892
File Sharing	7%	843
File Server	0%	832
Remote Console	1%	544
Database	0	1,000
Cloud – Storage	1%	815
Cloud – Mail	16%	590
Cloud – Productivity	3%	414
Total Cost of Ownership		
Ease of Use		
Initial Setup (Hours)		8
Expected Costs		US\$
Initial Purchase (hardware as tested)		\$2,552
Installation Labor Cost (@\$100/hr)		\$800
Annual Cost of Maintenance and Support (hardware/software)		\$256
Total Cost of Ownership		US\$
Year 1		\$3,608
Year 2		\$256
Year 3		\$256
3-Year Total Cost of Ownership		\$4,120

Figure 20 – Detailed Scorecard

Test Methodology

NSS Labs Software-Defined Wide Area Networking (SD-WAN) Test Methodology v2.0

A copy of the test methodology is available at www.nsslabs.com.

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.

© 2019 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.