SD-WAN COMPARATIVE REPORT

Performance

AUGUST 8, 2018

Author – Thomas Skybakmoen

Tested Products

Barracuda Networks NextGen Firewall F-Series F80 v7.1.1

Citrix Systems NetScaler SD-WAN v10.0.0.207

Cradlepoint AER2200-600M v6.5.0

FatPipe Networks MPVPN/SD-WAN v9.1.2

Forcepoint NGFW 1101 vSMC 6.3.6, engine 6.3.6.19302

Fortinet FortiGate 61E v6.0.1 GA Build 5068

Talari Networks Adaptive Private Networking (APN) Software APN 7.1

Versa Networks FlexVNF v120

VMware NSX SD-WAN by VeloCloud Edge v3.2

Environment

NSS Labs Software-Defined Wide Area Network (SD-WAN) Test Methodology v1.2

Overview

Implementation of software-defined wide area network (SD-WAN) solutions can be a complex process, with multiple factors affecting the overall performance of the device.

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

The mean opinion score (MOS) is used to calculate the QoE 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. Unlike speech codecs, video codecs have no limits on a maximum possible MOS.

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 QoE 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 QoE score of 4.41 for an excellent VoIP call. Any score below 3.5 represents a significantly degraded voice call and video stream. NSS considers a score below 3.4 as failing to meet the use case. Figure 1 presents the *NSS-Tested VPN Throughput (Mbps)*, VoIP QoE, and Video QoE for each product.

Vendor	NSS–Tested VPN Throughput (Mbps)	VoIP QoE	Video QoE
Barracuda Networks	124	2.49	2.75
Citrix Systems	751	4.25	4.04
Cradlepoint	17	3.52	1.10
FatPipe Networks	447	4.31	3.85
Forcepoint	713	4.20	4.04
Fortinet	749	4.38	4.26
Talari Networks	745	4.37	4.47
Versa Networks	552	4.09	4.09
VMware	880	4.27	4.21

Figure 1 – Throughput and Quality Results

Table of Contents

Table of Figures

Analysis

The marriage of software-defined networking (SDN) benefits to wide area network (WAN) technology yields the software-defined wide area network (SD-WAN), which allows consumer-grade links (or links without assured performance) to be leveraged for business-class services. Through the use of common VPN capabilities and the separation of data and control planes within SDN, software-managed connections can be established and managed between multiple sites over any number of link types (e.g., fixed circuit, DSL, cable, mobile, MPLS, etc.) without the operational challenges of having to manage multiple links simultaneously.

NSS research indicates that SD-WANs are typically deployed with the vendor's pre-defined or recommended (i.e., "out-of-the-box") settings. The tested SD-WAN products were configured with vendor-recommended settings in order to provide readers with relevant QoE and performance based on their expected usage.

Figure 2 depicts the difference between *NSS-Tested VPN Throughput*[1] and vendor performance claims. Vendor-claimed throughput is normalized to represent the branch device and not necessarily the maximum capacity of the headquarters. Please see the individual Test Reports for details on each vendor's submitted SD-WAN configuration.

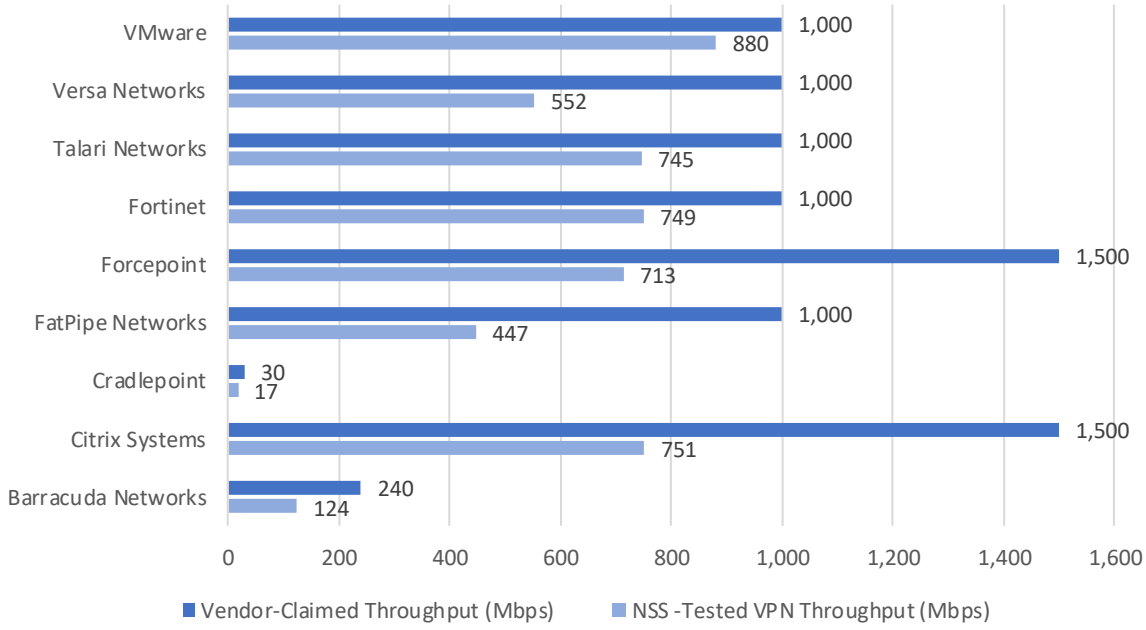

Figure 2 – Vendor-Claimed Throughput vs. NSS-Tested Throughput (Mbps)

[1] *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 more details, please see the Scorecard section in the individual Test Reports.

WAN Impairment and Link Failover

A critical function of any SD-WAN product is the identification and correct routing of traffic based on policy prioritization (autonomous or configured), which is influenced by network performance characteristics (e.g., variability, latency, jitter). Link impairment tests subject connected links to testing that represents real-world conditions encountered by enterprises today. The solution adapts the capabilities of the WAN for bandwidth, congestion, loss, latency, and jitter in real time.

In each test case, background traffic was introduced to populate links with sufficient activity as to represent typical enterprise network communications. Additionally, traffic-specific flows were introduced in order to capture accurate measurements, including RTP MOS for VoIP, relative MOS for video, and one-way delay for RTP. These measurements provide guidance as to how sensitive applications behave across a tested SD-WAN configuration.

Dynamic Path Selection with SLA Measurements

The goal of this test was 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, the 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.

Vendor	Dynamic Path Selection	
	VoIP QoE	Video QoE
Barracuda Networks	2.56	2.32
Citrix Systems	4.33	4.52
Cradlepoint	4.16	1.19
FatPipe Networks	4.26	3.55
Forcepoint	4.39	4.04
Fortinet	4.40	4.40
Talari Networks	4.41	4.38
Versa Networks	4.37	4.53
VMware	4.39	4.46

Figure 3 – Packet Delay Variation and Packet Loss (Voice and Video)

Path Conditioning and Application-Aware Steering

SD-WANs employ various techniques to condition WAN links in order to ensure reliability of data transmission. Some SD-WANs employ packet duplication, forward error correction, bonding, or load balancing. An SD-WAN product 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) is important for business-critical applications such as voice and video. These applications must be prioritized if a link has bad performance indicators. This test measures QoS using voice traffic and video stream. The test includes QoE scores for video and call measurements for VoIP (one-way delay for RTP). The SD-WAN should manage traffic according to configured QoS classification settings.

Vendor	Path Conditioning		Application-Aware Steering	
	VoIP QoE	Video QoE	VoIP QoE	Video QoE
Barracuda Networks	3.09	4.44	1.59	1.90
Citrix Systems	4.41	4.51	4.28	2.84
Cradlepoint	3.89	1.03	2.69	1.09
FatPipe Networks	4.41	4.26	4.41	3.85
Forcepoint	4.24	4.28	4.22	3.52
Fortinet	4.41	4.51	4.41	4.14
Talari Networks	4.40	4.53	4.40	4.46
Versa Networks	4.41	4.53	3.27	3.03
VMware	4.41	4.48	4.20	3.89

Figure 4 – Path Conditioning and Application-Aware Steering (Voice and Video)

Link Saturation and Congestion

As global QoS awareness can prevent congestion during the last mile of data delivery, the goal of this test was to ensure reliable use of bandwidth by the controller in the SD-WAN.

Vendor	Link Saturation and Congestion	
	VoIP QoE	Video QoE
Barracuda Networks	2.74	2.37
Citrix Systems	4.00	4.29
Cradlepoint	3.36	1.09
FatPipe Networks	4.16	3.74
Forcepoint	3.97	4.34
Fortinet	4.29	3.98
Talari Networks	4.29	4.51
Versa Networks	4.32	4.29
VMware	4.10	4.02

Figure 5 – Congestion and Saturation Impairments (Voice and Video)

Application-Aware Traffic Steering

This test verifies how the SD-WAN directs various application traffic flows for applications besides video and VoIP. Behavior was observed and recorded to establish whether voice/video and data were sent over the same link once impairments were applied and to establish which application took precedence. For a full breakdown of scores, please see the individual Test Reports. Figure 6 captures latency during steering of application-aware traffic in milliseconds. A high value for latency indicates that real-time traffic was prioritized over bulk traffic.

Vendor	No Impairments (baseline)	Failover	Dynamic Path Selection	Path Conditioning	Application-aware Steering
Barracuda Networks	752.0	526.8	403.0	449.5	330.0
Citrix Systems	1.6	10.4	36.0	3.1	4.0
Cradlepoint	771.0	3335.0	2222.5	983.5	950.5
FatPipe Networks	1.9	30.6	42.3	3.3	5.2
Forcepoint	25.4	26.3	80.6	69.9	8.1
Fortinet	0.0	56.7	41.5	0.4	0.9
Talari Networks	1.2	9.4	30.8	1.2	1.7
Versa Networks	4.0	196.5	74.8	6.8	220.0
VMware	0.6	1.8	5.6	0.8	11.3

Figure 6 – Application-Aware Traffic Steering with All Impairments (milliseconds)

Raw Packet Processing Performance (UDP Throughput)

This test uses UDP packets of varying sizes generated by test equipment. A constant stream of the appropriate packet size along with variable source and destination IP addresses is transmitted bidirectionally across the WAN links.

The percentage load and frames per second (fps) figures across the WAN links are verified by network monitoring tools before each test begins. Multiple tests are run and averages are taken where necessary.

The aim of the test is to determine the raw packet processing capability of the SD-WAN as well as its effectiveness at forwarding packets quickly in order to provide the highest level of network performance with the least amount of latency.

Figure 7 and Figure 8 depict the maximum UDP throughput (in megabits per second) achieved by each device using different packet sizes.

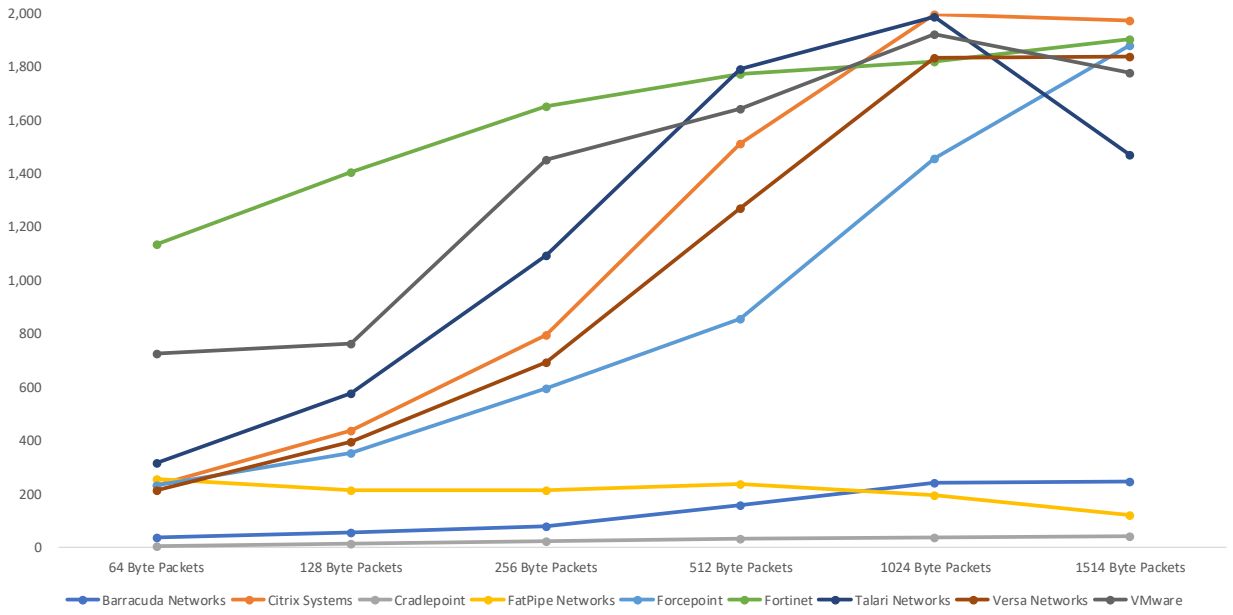

Figure 7 – UDP Throughput by Packet Size (Mbps)

Vendor	Throughput (Mbps)					
	64-Byte Packets	128-Byte Packets	256-Byte Packets	512-Byte Packets	1024-Byte Packets	1514-Byte Packets
Barracuda Networks	36	56	77	156	240	249
Citrix Systems	235	436	795	1512	2,000	1,974
Cradlepoint	4	13	24	33	37	40
FatPipe Networks	256	216	216	236	195	120
Forcepoint	235	356	596	856	1,460	1,880
Fortinet	1,135	1,405	1,654	1,774	1,820	1,905
Talari Networks	316	576	1,096	1,794	1,987	1,470
Versa Networks	216	396	696	1,271	1,833	1,842
VMware	726	764	1,455	1,642	1,925	1,780

Figure 8 – UDP Throughput by Packet Size (Mbps)

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. Figure 9 depicts UDP latency (in milliseconds) as recorded during the UDP throughput tests at 90% of maximum load. Lower values are preferred.

Vendor	Latency (ms)					
	64-Byte Packets	128-Byte Packets	256-Byte Packets	512-Byte Packets	1024-Byte Packets	1514-Byte Packets
Barracuda Networks	158	93	75.0	40.0	25	21
Citrix Systems	10	25	19.8	14.0	2.2	6.3
Cradlepoint	45	36	38.0	29.1	21.0	19.0
FatPipe Networks	1	2	1.6	2.7	4.1	21.0
Forcepoint	5	16	21.0	25.6	87.0	20.6
Fortinet	27	32	39.0	53.0	86.0	102.0
Talari Networks	2	2	10.3	43.0	75.0	3.6
Versa Networks	4	10	9.4	13.0	12.0	3.0
VMware	2	6	2.4	2.7	8.3	12.3

Figure 9 – UDP Latency by Packet Size (Milliseconds [ms])

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 is to stress the inspection engine and determine how it copes with high volumes of TCP connections per second, application-layer transactions per second, and concurrent open connections. All packets contain valid payload and address data, and these tests provide an excellent representation of a live network at various connection/transaction rates.

Note that in all tests the following critical "breaking points"—where the final measurements are taken—are used:

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

Figure 10 depicts the results from the connection dynamics tests.

Vendor	Maximum Concurrent TCP Connections	Maximum TCP CPS	Maximum HTTP CPS	Maximum HTTP Transactions per Second
Barracuda Networks	10,434	2,945	1,602	4,897
Citrix Systems	1,497,552	16,920	12,400	28,730
Cradlepoint	5,500	220	175	1,330
FatPipe Networks	1,351,908	17,200	15,990	35,620
Forcepoint	1,469,872	16,920	11,477	49,347
Fortinet	1,218,776	33,000	21,390	45,990
Talari Networks	1,487,972	17,100	13,500	62,900
Versa Networks	213,450	8,000	6,420	36,040
VMware	773,313	12,410	10,990	27,580

Figure 10 – Concurrency and Connection Rates

HTTP Capacity

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

This provides a test environment that is as close to real-world conditions as possible, while ensuring absolute accuracy and repeatability.

Each transaction consists of a single HTTP GET request. All packets contain valid payload (a mix of binary and ASCII objects) and address data. This test provides an excellent representation of a live network (albeit one biased toward HTTP traffic) at various network loads. Figure 11 through Figure 15 depict the maximum throughput achieved across a range of different HTTP response sizes that may be encountered in a typical corporate network.

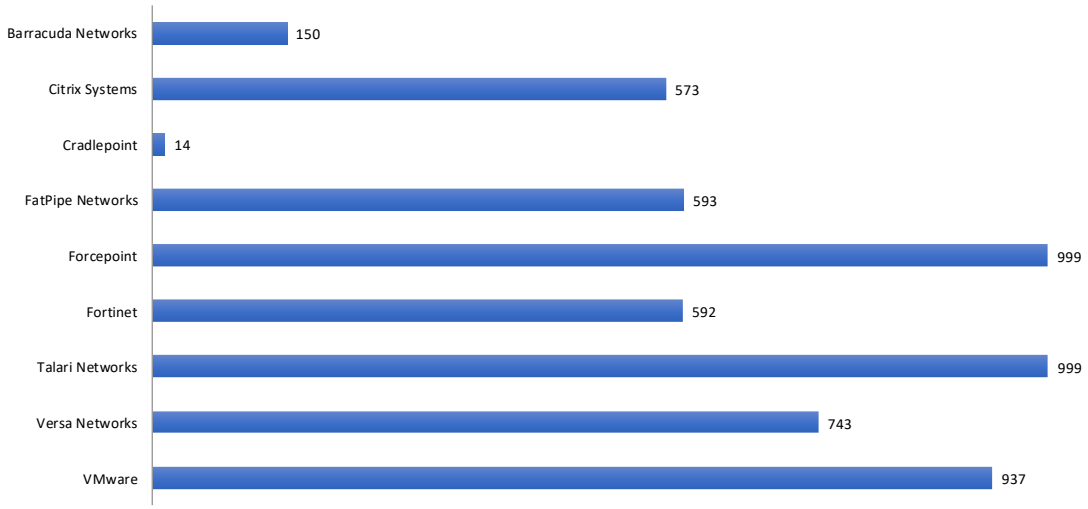

Figure 11 – Maximum Throughput per Device with 44 KB Response (Mbps)

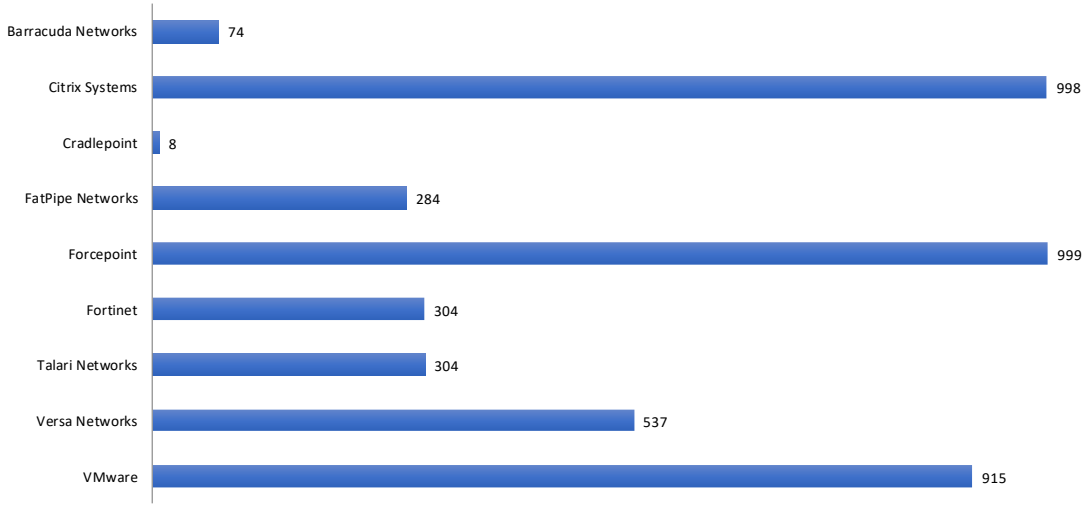

Figure 12 – Maximum Throughput per Device with 21 KB Response (Mbps)

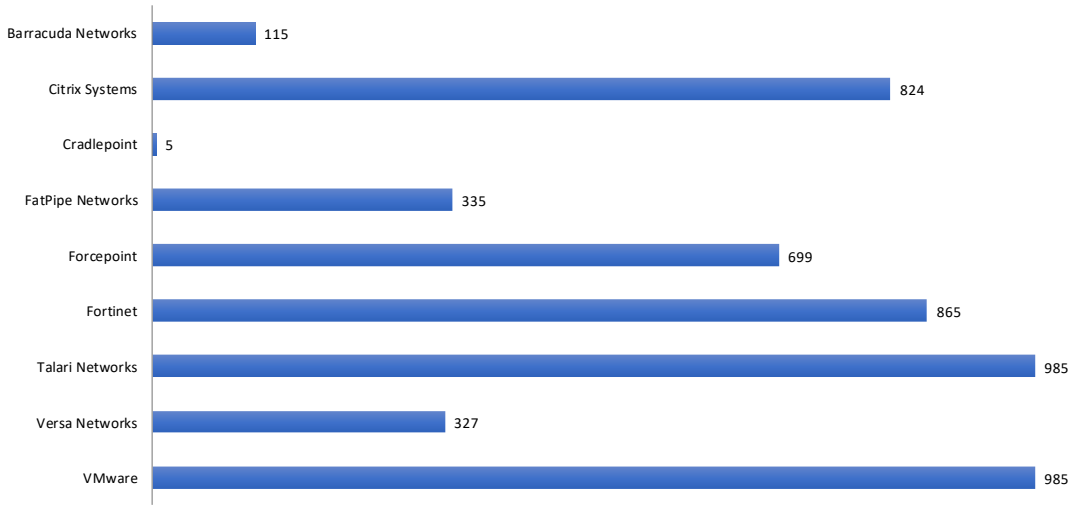

Figure 13 – Maximum Throughput per Device with 10 KB Response (Mbps)

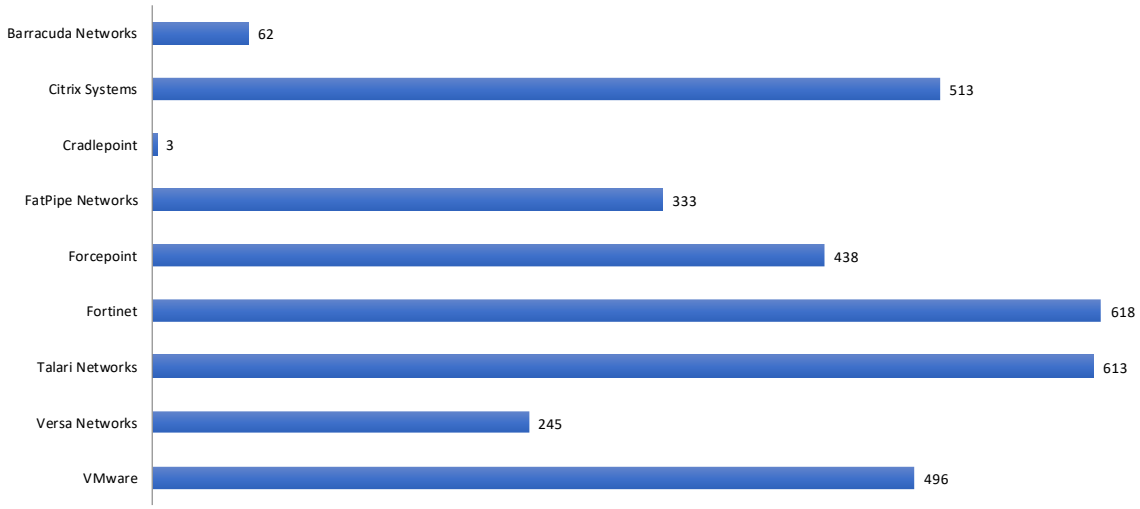

Figure 14 – Maximum Throughput per Device with 4.5 KB Response (Mbps)

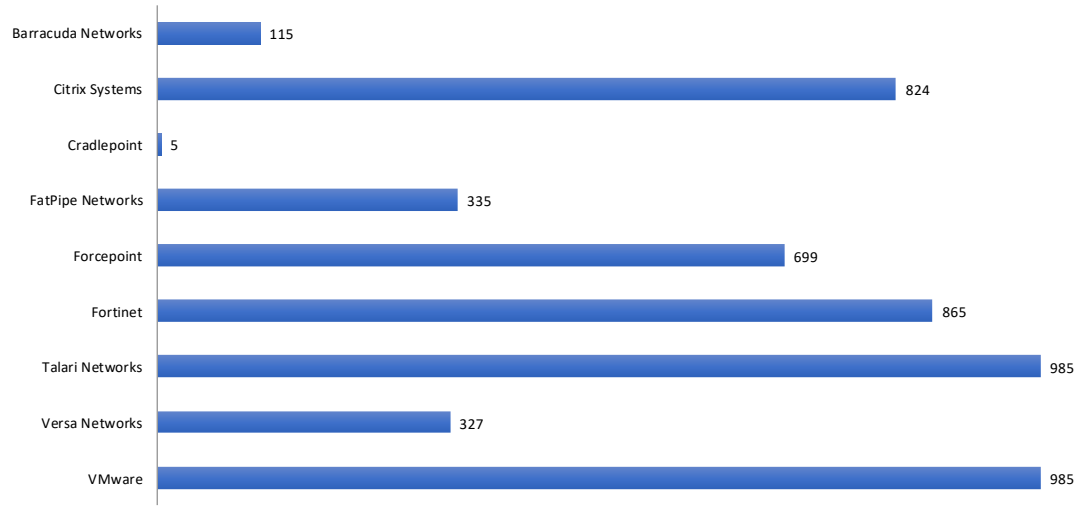

Figure 15 – Maximum Throughput per Device with 1.7 KB Response (Mbps)

Figure 16 depicts the maximum application layer connection rates (HTTP connections per second) achieved with different HTTP response sizes (from 44 KB down to 1.7 KB).

Vendor	44 KB Response Size	21 KB Response Size	10 KB Response Size	4.5 KB Response Size	1.7 KB Response Size
Barracuda Networks	375	369	1,150	1,249	1,400
Citrix Systems	1,432	4,990	8,242	10,250	15,500
Cradlepoint	35	40	45	60	78
FatPipe Networks	1,482	1,420	3,353	6,650	14,600
Forcepoint	2,497	4,995	6,993	8750	12,600
Fortinet	1,479	1,520	8,650	12,350	19,500
Talari Networks	2,498	1,522	9,850	12,250	16,000
Versa Networks	1,857	2,686	3,268	4,900	5,705
VMware	2,342	4,575	9,850	9,920	8,492

Figure 16 – Maximum Connection Rates per Device with Various Response Sizes

Application Average Response Time at 90% Maximum Capacity

Figure 17 depicts the average application response time (application latency, measured in milliseconds) for different packet sizes (ranging from 44 KB down to 1.7 KB), recorded at 90% of the measured maximum capacity (throughput). A lower value indicates an improved application response time.

Vendor	44 KB Latency (ms)	21 KB Latency (ms)	10 KB Latency (ms)	4.5 KB Latency (ms)	1.7 KB Latency (ms)
Barracuda Networks	3.1	2.0	1.9	1.0	0.0
Citrix Systems	13.6	13.9	10.5	8.0	1.8
Cradlepoint	1.2	1.5	3.6	12.0	1.1
FatPipe Networks	1.8	2.1	3.1	1.5	1.1
Forcepoint	7.9	2.7	18.8	19.8	1.3
Fortinet	2.2	1.1	2.1	1.0	0.1
Talari Networks	10.1	8.8	10.0	9.8	1.6
Versa Networks	29.0	29.6	30.0	13.0	4.4
VMware	13.2	11.2	11.9	4.6	2.2

Figure 17 – Application Latency (Milliseconds) per Device with Various Response Sizes

HTTP Capacity with HTTP Persistent Connections

This test uses HTTP persistent connections, with each TCP connection containing 10 HTTP GETs and associated responses. All packets contain valid payload (a mix of binary and ASCII objects) and address data, and this test provides an excellent representation of a live network at various network loads. The stated response size is the total of all HTTP responses within a single TCP session.

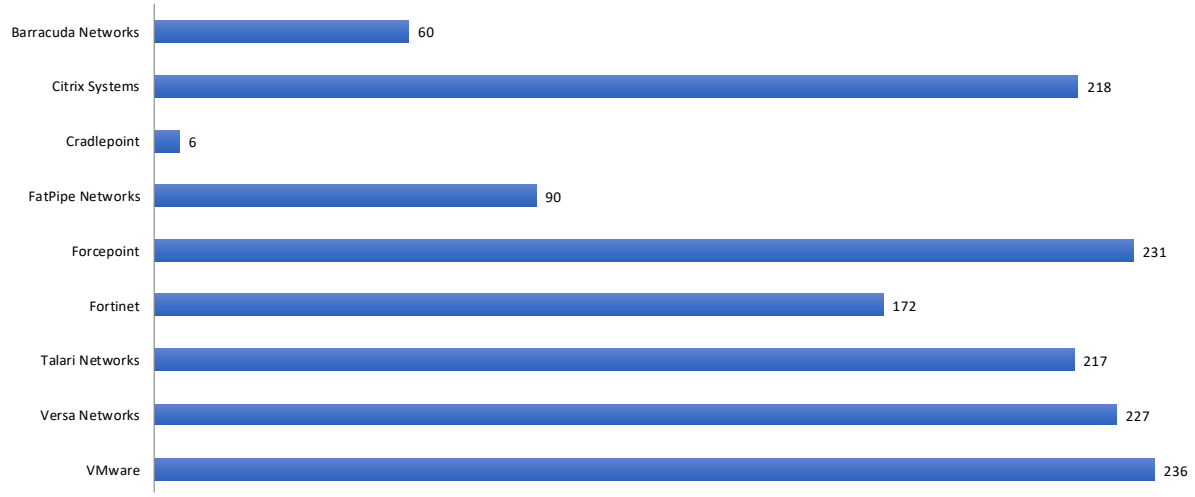

Figure 18 – HTTP 250 Capacity with HTTP Persistent Connections (CPS)

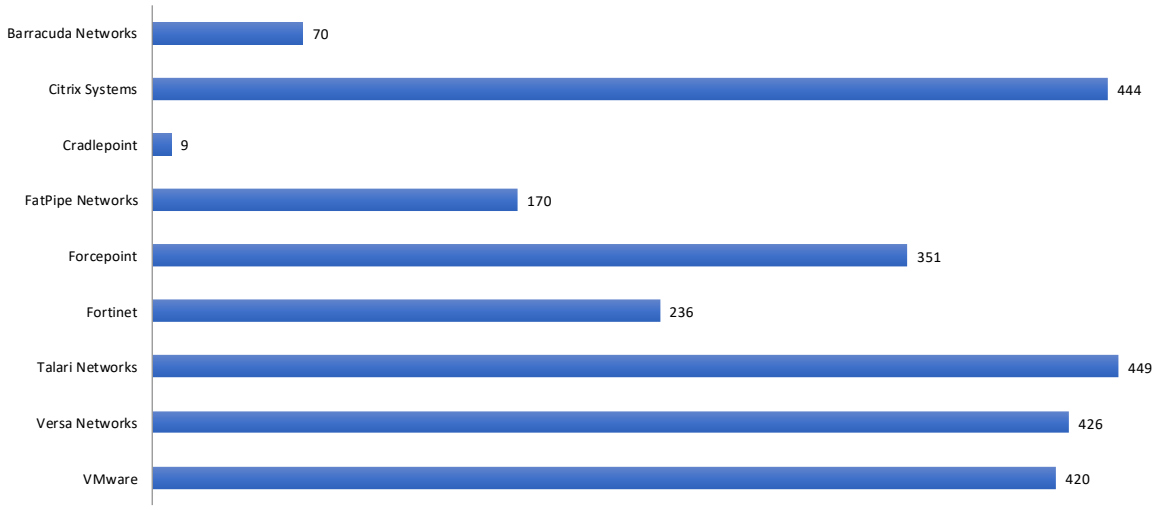

Figure 19 – HTTP 500 Capacity with HTTP Persistent Connections (CPS)

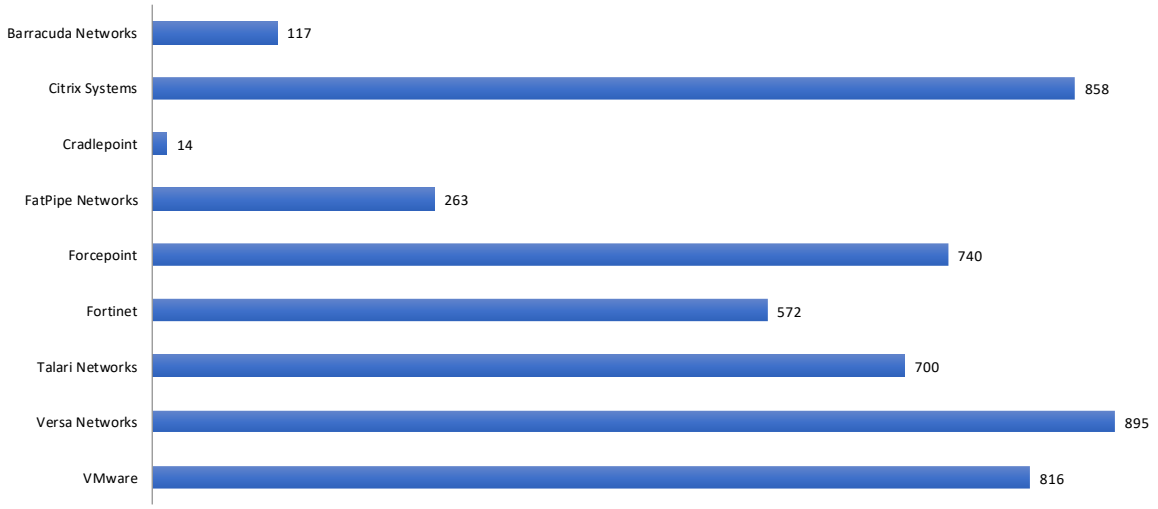

Figure 20 – HTTP 1000 Capacity with HTTP Persistent Connections (CPS)

Single Application Flows

This test measures the performance of the SD WAN using single application flows. These application flows are what NSS expects an SD-WAN product will face in an enterprise environment. Using a frame size distribution ranging from 64 to 1024 bytes, performance testing was conducted between Branch 1 and the headquarters site over two established tunnels and was limited to 1,092 Mbps, as described in the SD-WAN Test Methodology v1.2.

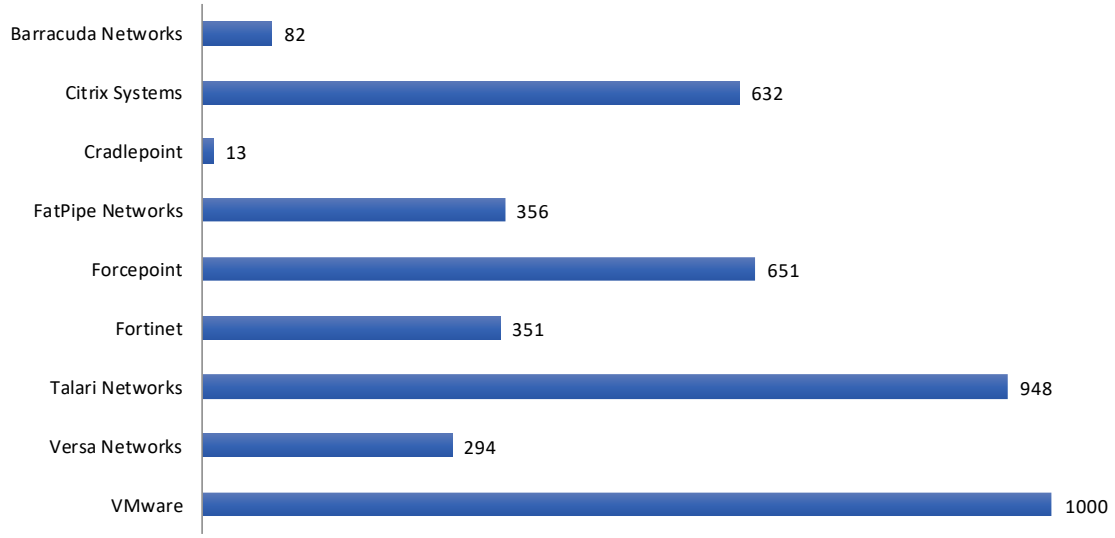

Figure 21 –Single Application Flow: Telephony (Mbps)

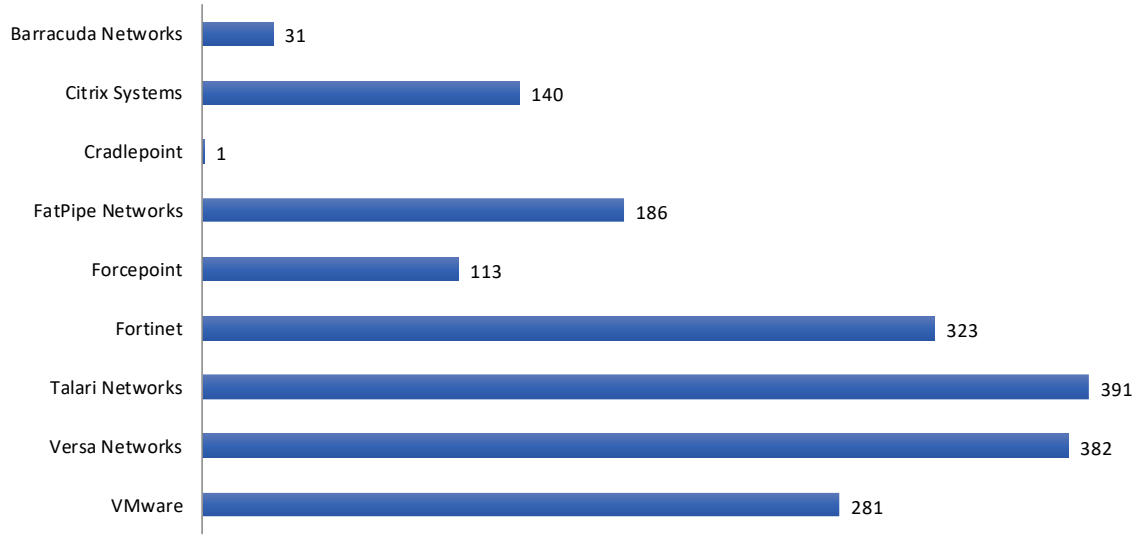

Figure 22 –Single Application Flow: Financial (Mbps)

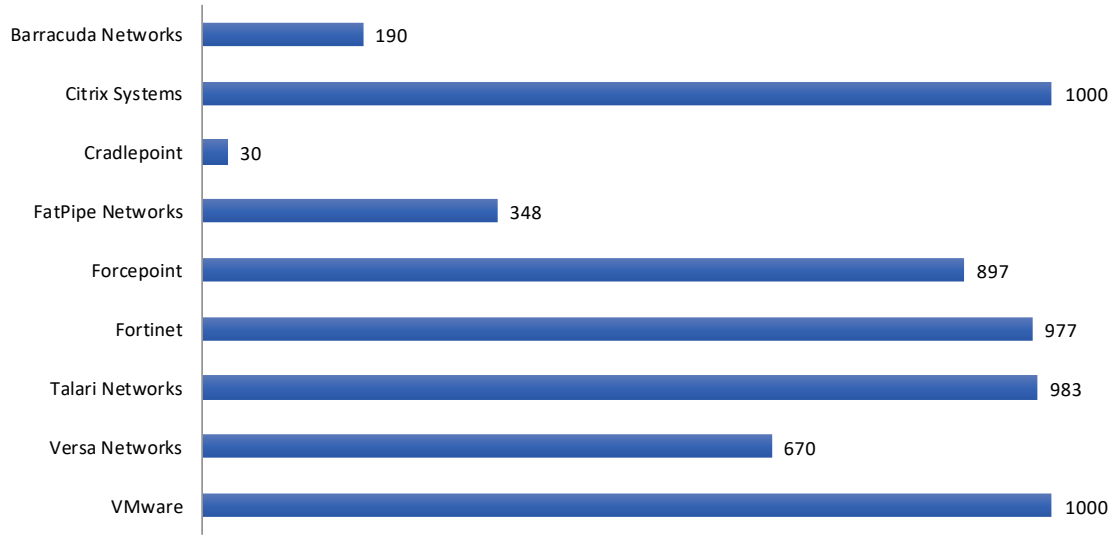

Figure 23 –Single Application Flow: Email (Mbps)

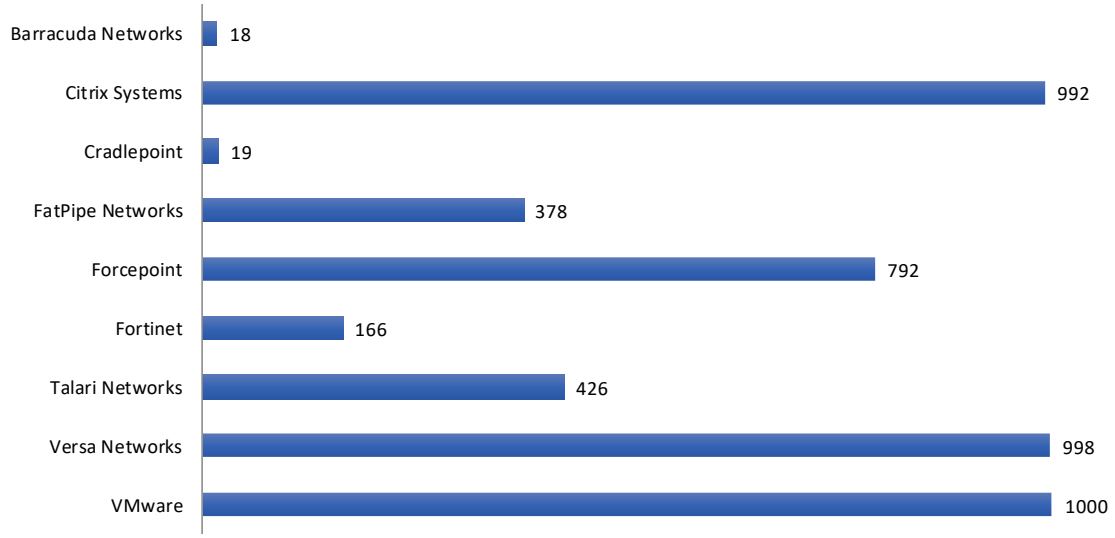

Figure 24 –Single Application Flow: File Sharing (Mbps)

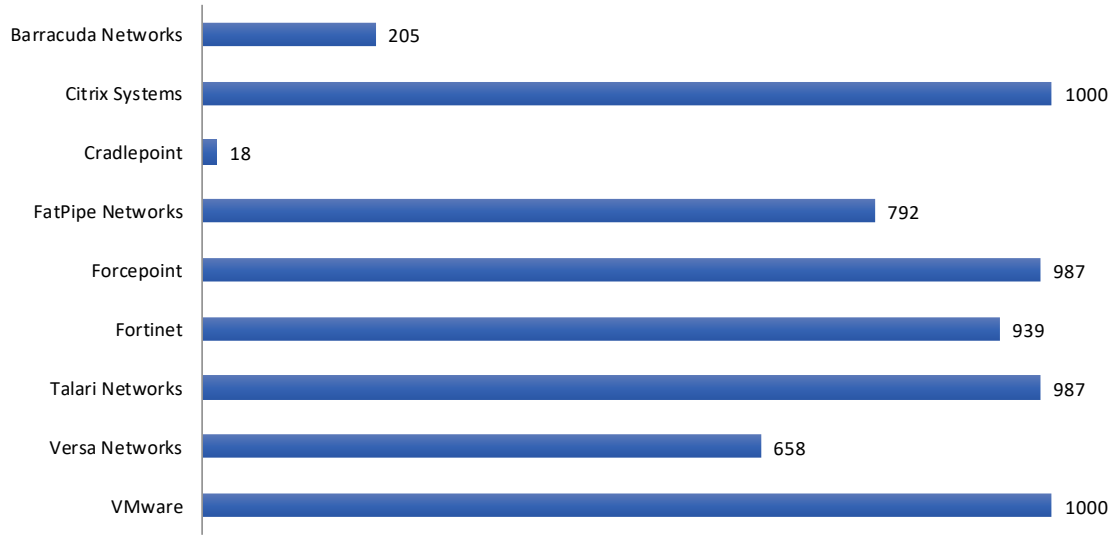

Figure 25 –Single Application Flow: File Server (Mbps)

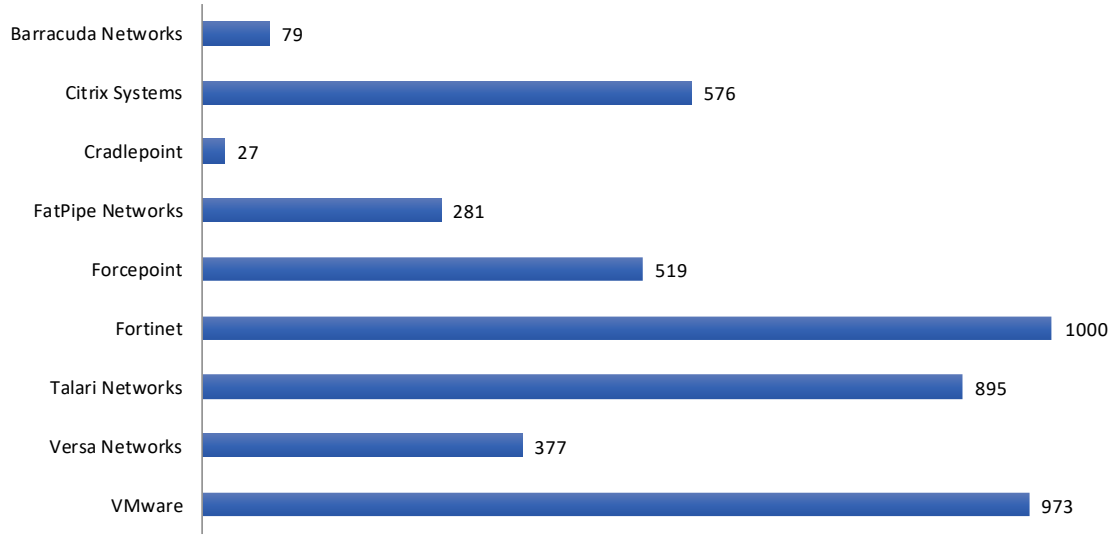

Figure 26 –Single Application Flow: Remote Console (Mbps)

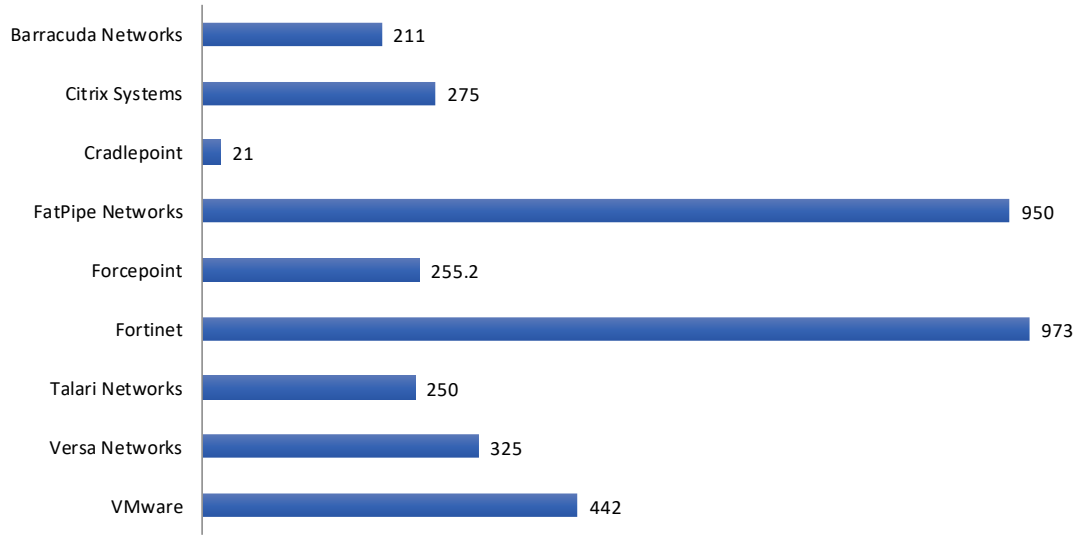

Figure 27 –Single Application Flow: Video (Mbps)

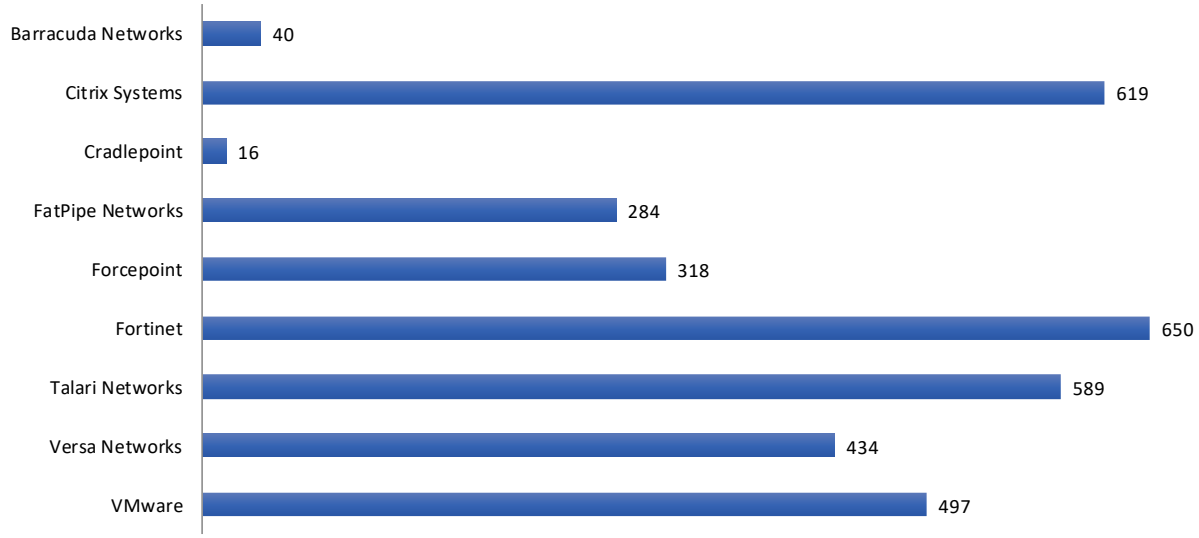

Figure 28 –Single Application Flow: Meeting (Mbps)

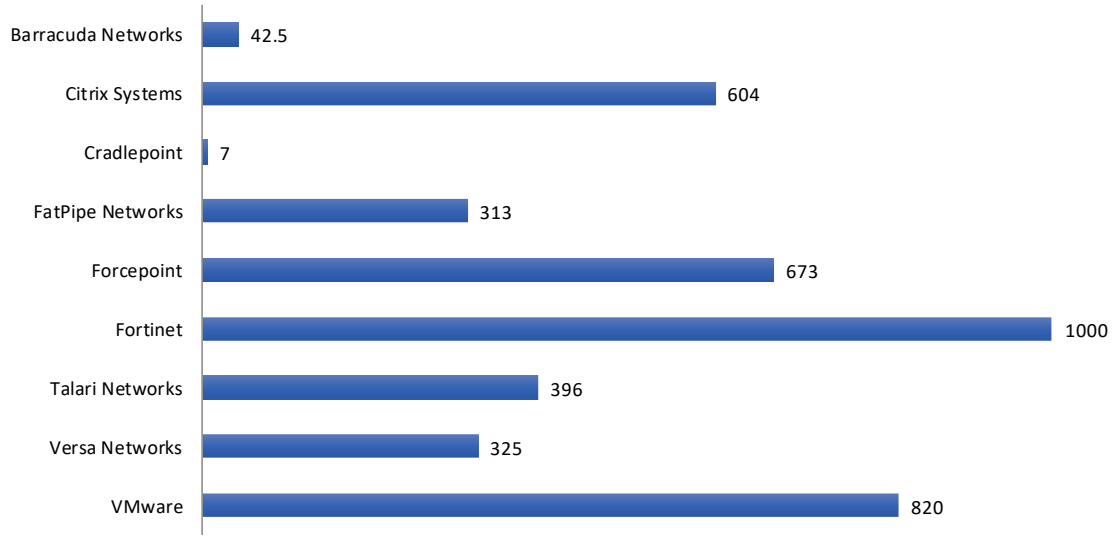

Figure 29 –Single Application Flow: Database (Mbps)

Test Methodology

Software-Defined Wide Area Network (SD-WAN) Test Methodology v1.2

Contact Information

3711 South Mopac Expressway
Building 1, Suite 400
Austin, TX 78746
info@nsslabs.com
www.nsslabs.com