#221100

January 2021

Commissioned by

Citrix Systems, Inc.

Citrix ADC VPX

Performance vs. F5 BIG-IP Virtual Edition

EXECUTIVE SUMMARY

Application Delivery Controllers (ADCs) are the unseen workhorses of the internet. Every back-end application that consists of two or more servers has an ADC of some type on the front end to optimize delivery to the user. Citrix ADC VPX (formerly NetScaler) is designed to enhance end-user experience by providing high performance traffic management coupled with low latency.

Citrix Systems commissioned Tolly to benchmark the performance of Citrix ADC VPX virtual appliance and compare with the F5 BIG-IP Virtual Edition (VE). Tests were run in an Amazon Web Services (AWS) environment and included tests of ADC and web application firewall (WAF) capabilities using different encryption protocols. Tests focused on measuring latency as an indicator of responsiveness, transaction throughput, and, ultimately, end-user experience. Tests measured P99 latency which measures the latency of the worst 1% of the flows.

The Citrix ADC VPX outperformed the F5 BIG-IP VE in all test scenarios having lower (better) latency with lower CPU utilization when benchmarked with comparable throughput levels. See Figure 1.

THE BOTTOM LINE

Citrix ADC VPX delivers:

1. 1/5 the latency under stress of F5 BIG-IP VE in data throughput tests using both TLS 1.3 and TLS 1.2 encryption
2. 1/2 the latency under stress of F5 BIG-IP VE in transaction tests using both TLS 1.3 and TLS 1.2 encryption
3. 10% greater transactions than F5 BIG-IP VE, translating into a potential 1.85 billion more transactions in 24 hours
4. Lower latency than F5 BIG-IP VE in all ADC + WAF tests run using both TLS 1.3 and TLS 1.2

HTTPS Request Latency with TLS 1.3 & 1.2 for *Volume Traffic Test*

P99 Results

(Lower result is better)

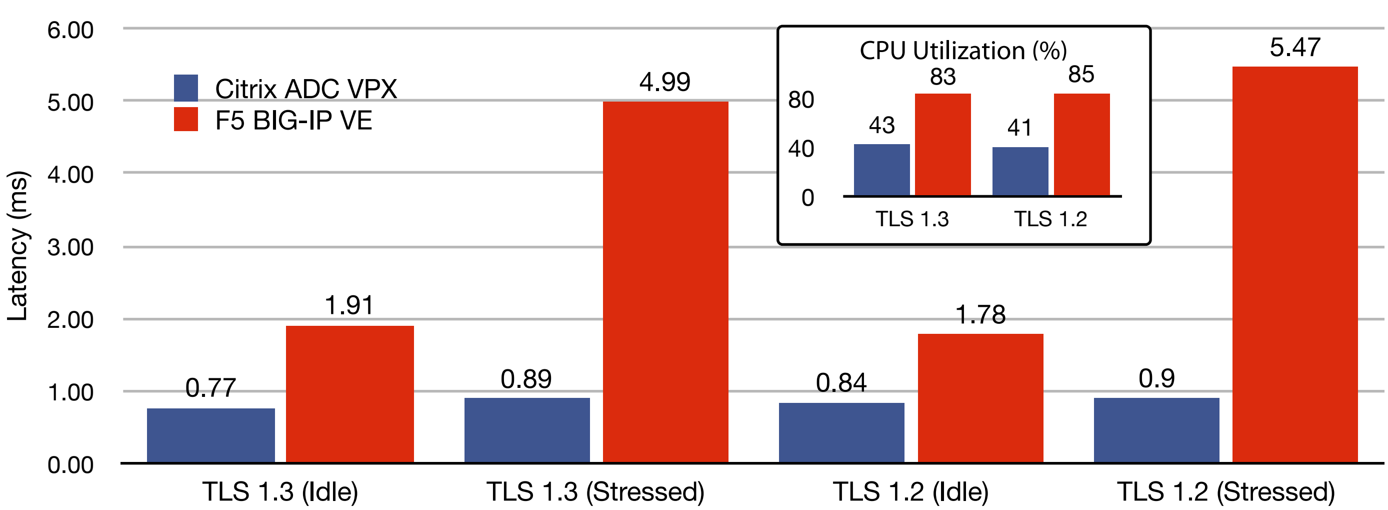

Note: Results as reported by Fortio. Spirent Avalanche Virtual traffic with 30KB response size with throughput for each ADC of ~14Gbps.

Source: Tolly, December 2020

Figure 1

HTTPS Request Latency with TLS 1.3 & 1.2 for *Transaction Test*

P99 Results

(Lower result is better)

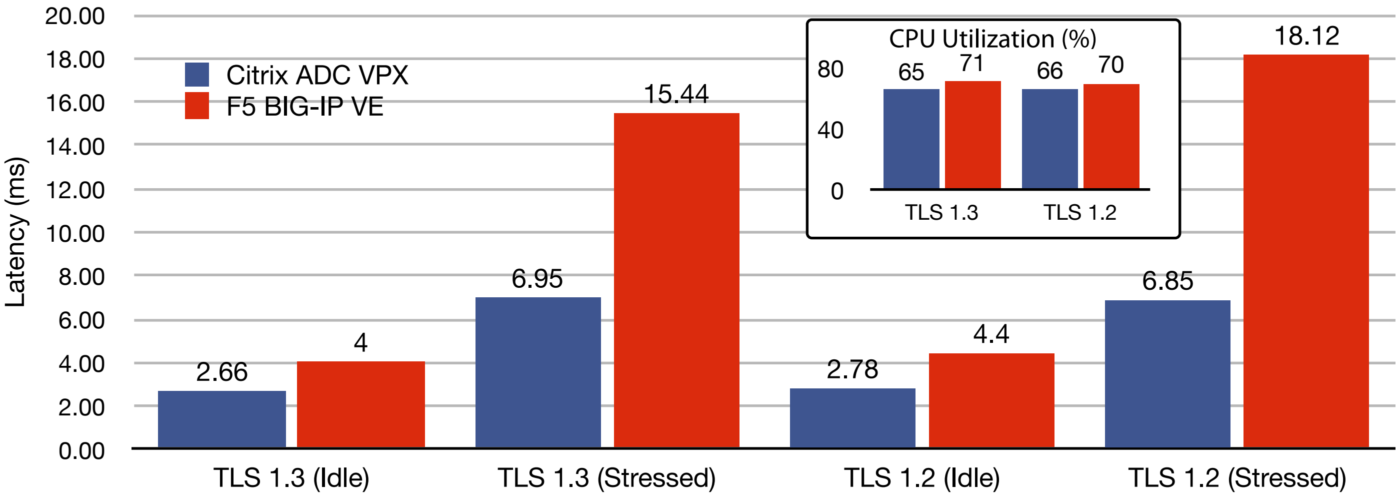

Note: Results as reported by Fortio. Spirent Avalanche Virtual traffic with 1-byte response size with throughput for each ADC of ~200,000 RPS.
Source: Tolly, December 2020

Figure 2

Key Takeaways

Tests were run first using only core ADC functionality and then with ADC + WAF functionality. Tests were conducted separately using both TLS 1.3 and TLS 1.2. Systems were benchmarked in "idle" (low traffic) and "stressed" (high traffic) states. Additionally, a test of total transactions was run where both systems were running at the same CPU usage to gauge effective throughput.

There are 36 throughput bars and 12 CPU bars in this report. While the reader is encouraged to review each and every test, the results are consistent throughout. The Citrix ADC VPX delivers better results than the F5 BIG-IP in every variant of every test.

Citrix ADC VPX latency is better (lower) in every scenario. Citrix ADC VPX CPU usage is better (lower) in every scenario. Lower latency means reduced wait time for application users. Lower CPU usage means more user traffic can be processed and potentially fewer ADCs required. Citrix ADC VPX "reduces the drag" on user experience in every scenario tested.

ADC Latency

This series of tests measured the latency of the two solutions in HTTPS volume throughput (large data request) and transaction (small data request) scenarios. See Figures 1 and 2.

In the volume traffic test, F5 BIG-IP VE latency was over 5X longer (worse) than Citrix ADC VPX for TLS 1.3 tests and over 6X longer than Citrix for TLS 1.2 tests under stressed (loaded) conditions. To achieve the same approximate throughput as Citrix,

Total HTTPS Transactions with Equivalent (65%) CPU Usage

(Higher result is better)

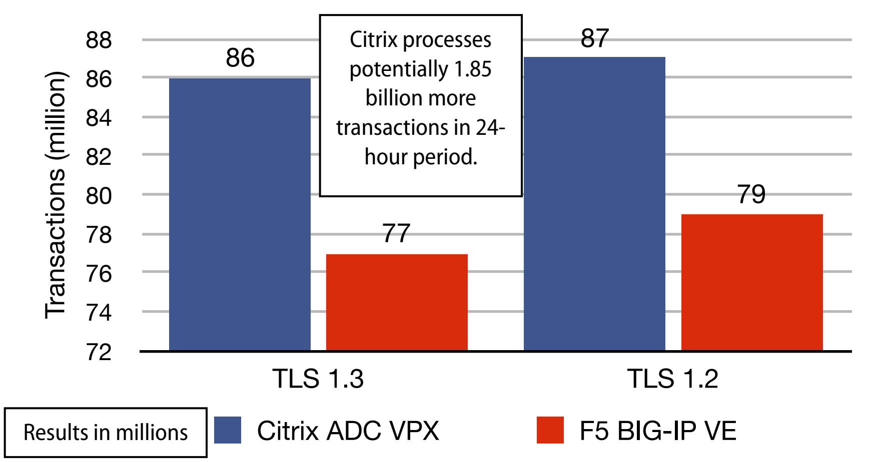

Note: Results as reported by Spirent Avalanche Virtual. Test run time was seven minutes.
Source: Tolly, December 2020

Figure 3

BIG-IP used more than 2X the CPU usage of Citrix ADC VPX.

In the transaction traffic test, F5 BIG-IP VE latency was over 2X longer (worse) than Citrix ADC VPX for TLS 1.3 tests and over 2.5X longer than Citrix for TLS 1.2 tests under load. To achieve the same approximate transaction throughput as Citrix, F5 BIG-IP VE used more CPU than Citrix.

ADC Transactions Over Time

This test focused on measuring the number of transactions processed on each virtual appliance under a moderate CPU load, measuring that over the course of a seven-minute test and then extrapolating that out to a 24-hour period to quantify the potential transaction processing delta between the two solutions. See Figure 3.

To establish the test transaction rate, engineers increased the transaction load on each system until the CPU utilization of that system was approximately 65%. This represented a system under load but was not constrained by CPU resource (i.e., was not close to 100% CPU usage.)

With the transaction rate established, the test was run for seven minutes and results recorded. With the TLS 1.3 encryption protocol, Citrix ADC completed 9 million more transactions than F5 BIG-IP VE. With TLS 1.2, Citrix ADC completed 8 million more transactions than F5 BIG-IP VE.

Over the course of a 24-hour period, the TLS 1.3 transaction results extrapolate to a potential additional 1.85 billion more transactions processed by Citrix over F5 BIG-IP VE.

ADC + WAF Latency

Many customers leverage the additional application security capabilities available in leading ADCs. With WAF enabled, the ADC scans the traffic for application security threats. Because additional scanning is conducted latency is naturally longer than with ADC-only traffic.

This series of tests measured the latency of the two solutions in the volume test with two different allocations for WAF processing. In the first scenario 100% of the traffic was inspected by the WAF. In the second scenario only 33% of the traffic was inspected. See Figures 4 and 5.

Again, the Citrix ADC VPX had better results than F5 BIG-IP VE in all test combinations. In the stressed test using TLS 1.3, Citrix latency was 2.23ms lower than F5 in the 100% WAF test and 2.79ms lower in the 33% WAF test.

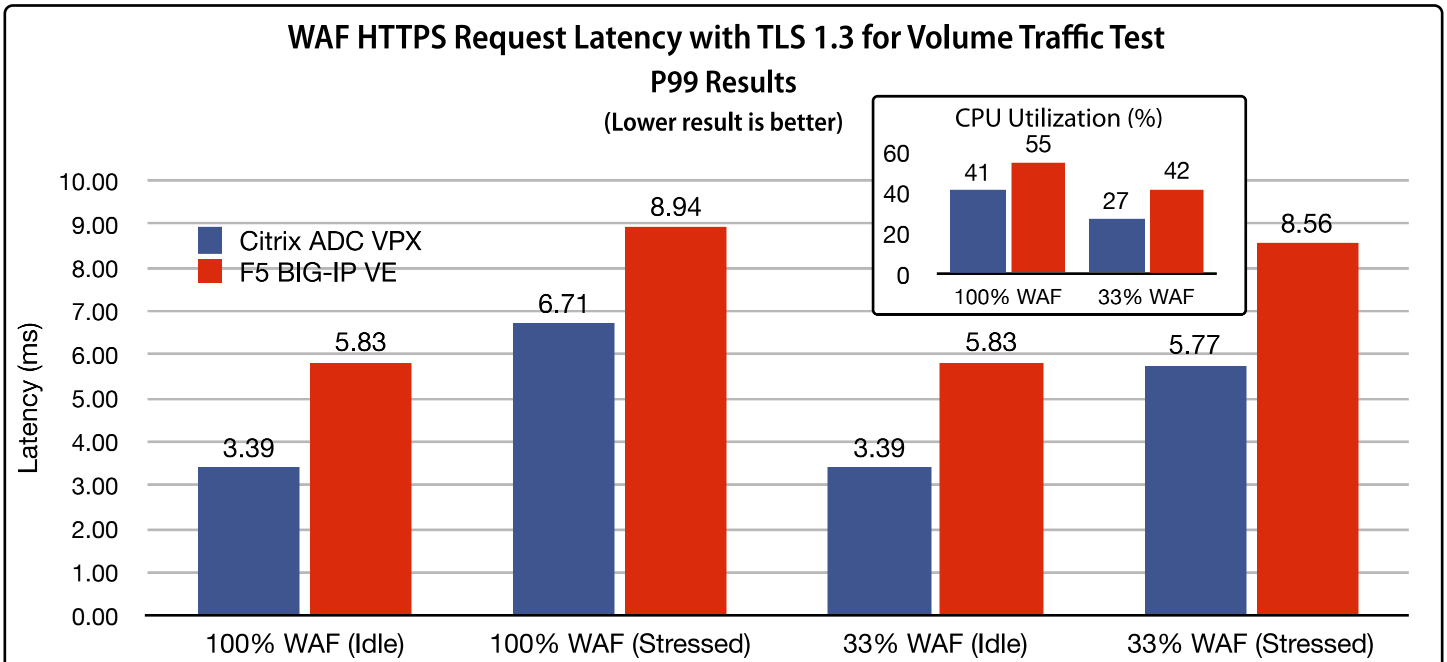

Note: Results as reported by Fortio. Spirent Avalanche Virtual traffic with 30KB response size with throughput for each ADC of ~3.8Gbps.In 100% WAF tests all traffic was inspected. In the 33% WAF tests only one third of the traffic was inspected. 80% GET/20% POST.

Source: Tolly, December 2020

Figure 4

The Citrix ADC VPX CPU usage was also lower than F5 BIG-IP VE in both tests. In the 100% WAF test F5 BIG-IP VE CPU requirements were 34% more than Citrix while in the 33% WAF test the F5 BIG-IP VE CPU requirement was nearly 60% higher.

In the stressed test using TLS 1.2 Citrix latency was 2.88ms lower than F5 BIG-IP VE in the 100% WAF test and 4.18ms lower in the 33% WAF test.

The Citrix ADC VPX CPU usage was also lower than F5 BIG-IP VE in both tests. In the 100% WAF test F5 BIG-IP VE CPU requirements were more than 20% more than Citrix while in the 33% WAF test the F5 BIG-IP VE CPU requirement was nearly 50% higher.

Test Setup & Methodology

The focus of the test was to benchmark the performance of virtualized application delivery controllers. Benchmarking focused on measuring session latency (delay) in traffic environments designed to model real-world conditions. See Tables 1-4 for details.

Testing used six vServers configured on a single ADC appliance and tests were run with web application firewall (WAF) functions enabled and disabled.

Comparable virtual appliances from Citrix Systems and F5 BIG-IP VE were tested in the Amazon Web Services (AWS) US cloud environment in early December 2020. ADC-only tests were run on c5n.4xlarge instances. Hyperthreading (HT) was disabled. ADC + WAF tests were run on c5.4xlarge instances because the F5 BIG-IP VE virtual appliance that included WAF functionality was not supported on the C5n instance.

Citrix Systems, Inc.

Citrix ADC VPX

Cloud ADC Performance

Tolly. Certified

Tested December 2020

ADC Configurations

Depending upon the use case an ADC might be configured with or without web

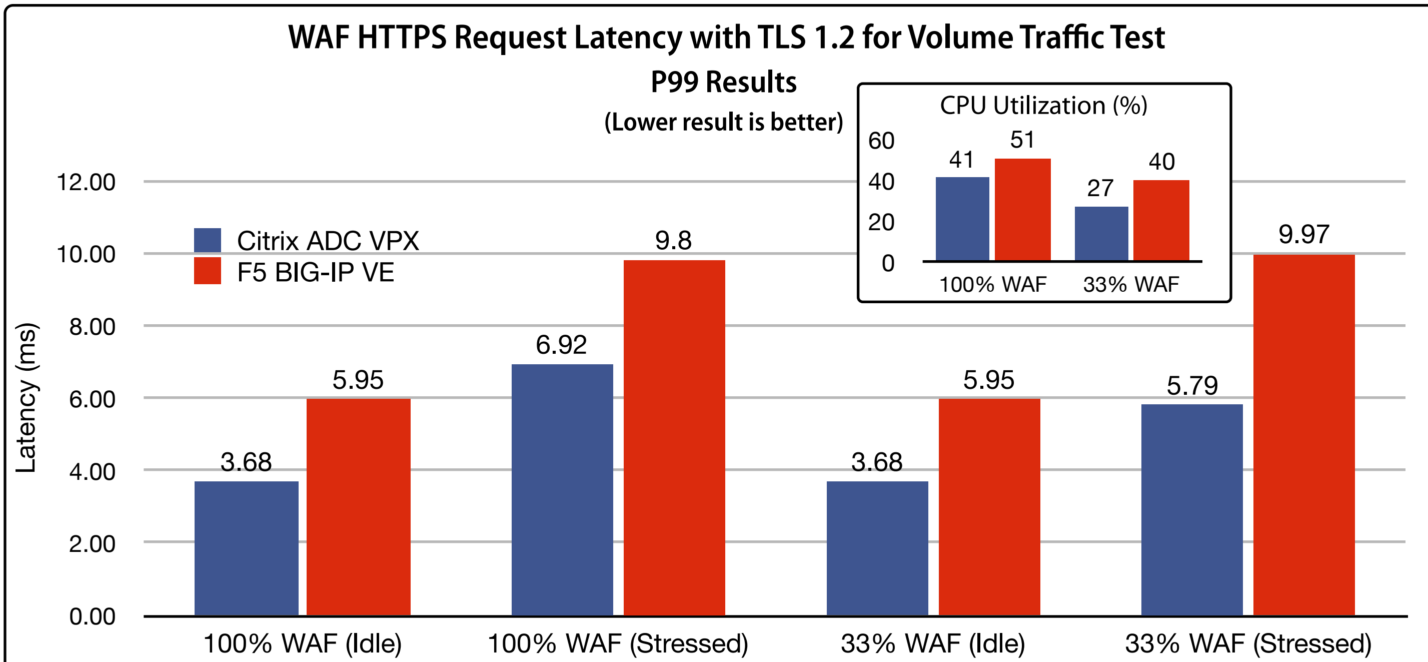

Note: Results as reported by Fortio. Spirent Avalanche Virtual traffic with 30KB response size with throughput for each ADC of ~3.8Gbps. In 100% WAF tests all traffic was inspected. In the 33% WAF tests only one third of the traffic was inspected. 80% GET/20% POST.

Source: Tolly, December 2020

Figure 5

application firewall (WAF) functions enabled on its vServers. Additionally, customers might deploy a hybrid environment.

Since almost all web traffic is now encrypted, all tests were run using encrypted sessions (detailed below).

P99 Latency

Latency (delay) through the ADC is the primary metric that was used in the test. While some tests report average latency, all latency results in this test are P99 latency. P99 latency is a more stringent

Test Details

Solutions Under Test (AWS Instances)

Vendor	ADC	Version
Citrix Systems	Citrix ADC VPX	13.0-70.7
F5	BIG-IP Virtual Edition (VE)	15.1.1.0.0.6 (15.0.1.1-0.0.3 used for ADC+WAF TLS 1.3 tests)

Note: 16vCPU license used for both DUTs. Instance had 8vCPUs

Table 1

Test Variables

ADC Configuration	Encryption Protocol - Transport Level Security (TLS)	
ADC Only	v1.2 TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	v1.3 TLS_AES_256_GCM_SHA384
ADC + WAF 33% and 100%	v1.2 TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	v1.3 TLS_AES_256_GCM_SHA384

Note: Only volume tests run with WAF. WAF traffic profile was GET 80%/POST 20%.

Table 2

Test Tools

Vendor	Solution	Function
Spirent Comm.	Avalanche Virtual	Traffic generator
Fortio (Open Source)	Fortio	Measure latency

Table 3

WAF Feature Settings

Feature	Setting
Cookies All Checks	Enabled
Evasion Techniques Detected	Enabled
File Types	Alarm/Block enabled on violations
General Settings	Alarm/block enabled if request length increases buffer size of 20 MB
Headers	Alarm/Block enabled for illegal header length
HTTP Protocol Compliance	Enabled: Multiple Host header, Bad host header, Host header contains IP address, No host Header, content length should be a positive number, CRLF characters before request start, Bad multipart parameters/ form data request parsing, Body in GET or HEAD requests, Chunked Request with Content length header, Several Content length headers
SQL Injection	Enabled
Cross-site Scripting	Enabled
Buffer Overflow	Enabled
Signatures for CVE-2019-14994	Enabled
Command Injection	Enabled

Table 4

Source: Tolly, December 2020

measurement that represents the latency of the worst 1% of all flows.

Latency was measured twice in each test scenario, when the ADC was "idle" when processing a light traffic load of 100 requests per second and again under "stressed" conditions when the ADC was handling a heavy traffic load.

Traffic Profiles: Data Throughput & Transaction Tests

It is customary to test both the volumetric (data) throughput and number of HTTPS transactions of ADCs. This allows users to understand the range of performance for a solution. For these tests, the goal was to benchmark the systems when running comparable loads.

Data Throughput Test

The data throughput uses a large request size and, therefore, generates fewer transactions per second than the small requests used in the transaction test.

For this test, a 30KB response size was used. This represents a common web object retrieval size. For ADC-only tests, load was increased until it reached approximately 14Gbps of throughput. The traffic profile was 100% "GET." For WAF tests, the target throughput was approximately 3.8Gbps. The traffic profile 80% "GET" and 20% "POST."

For the "idle" measurement traffic was set to 100 requests (transactions) per second. The "stressed" measurement was taken after the traffic had ramped up to the target rate. Tests ran for at least seven minutes.

Test Systems Summary

Vendor	Product	Web	
Fortio	**Fortio**	https://fortio.org	
Spirent Comm.	**Avalanche Virtual**	https://www.spirent.com	

Transaction Test

The transaction throughput uses a small response size and, therefore, generates more transactions per second than the large response used in the transaction test.

For this test, a GET request generated 1 byte HTTPS response. Load was increased until it reached approximately 200,000 TPS. The traffic profile was 100% "GET."

For the "idle" measurement traffic was set to 100 requests (transactions) per second. The "stressed" measurement was taken after the traffic had ramped up to the target rate. Tests ran for at least seven minutes.

Total HTTPS Transactions Over Time

This test determined the transaction processing of each ADC while running with the same CPU usage. Engineers chose 65% CPU to represent an active system where CPU is low enough not to represent a bottleneck. The test was run for seven minutes. The reported results were extrapolated to a 24-hour period to illustrate the delta in workload processing over a longer period.

Transaction tests were not run for the ADC + WAF configuration.

Encryption: Transport Layer Security 1.2 & 1.3

The amount of encrypted traffic on the web is large and growing all the time. Thus, it is important to benchmark the performance with encrypted (i.e., HTTPS) traffic in order to get an accurate reading on how an ADC will perform in the real-world. Thus, all the benchmarking performed used traffic encrypted using TLS (which is the follow-on to SSL).

Today, traffic will consist of sessions using the older TLS 1.2 protocol along with a growing number of sessions using the current TLS 1.3 protocol (which will ultimately replace TLS 1.2). Thus, separate tests were run with each to benchmark encryption protocol specific performance.

Test Procedure

ADCs under tests were configured with six vServers.

Spirent Avalanche Virtual was used to generate all traffic for the transaction/data throughput tests. Avalanche Virtual was resident in the same datacenter as the ADCs under test and simulates both the client and server sides of the connections.

Fortio was used to measure P99 latency and ran in tandem with Avalanche Virtual.

For the idle measurement, Avalanche Virtual generated 100 requests per second for several minutes. During this time, Fortio was also run and reported P99 latency.

For the stressed measurement, Avalanche Virtual was used, first, to generate a load level that would bring the two vendors SUTs to equivalent throughput levels. Engineers would then reduce the load by approximately 5%. This was done to be sure that Fortio could run and still stay within the maximum levels established for the test. Fortio was run and reported P99 latency.

About Tolly

The Tolly Group companies have been delivering world-class IT services for more than 30 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company by E-mail at sales@tolly.com, or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:
http://www.tolly.com

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is," and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com. No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.

dj-m-wt-2020-12-23-VerI