

Dell EMC PowerEdge R7525 Virtual Desktop Infrastructure (VDI)

Windows 10 Scalability vs. Prior-Generation Platform

Executive Summary

Virtualization of Windows 10 desktops is increasingly relied upon by organizations to provide a scalable and economical solution for large groups of corporate users. The new Dell EMC PowerEdge R7525 is built around the 2nd Gen AMD EPYC 7002 series CPU which boasts a re-architected memory structure to improve performance for VDI and other applications.

Dell Technologies commissioned Tolly to benchmark the VDI performance of the R7525 and compare that with performance of the prior generation platform. Additionally, tests were run to illustrate that 230 knowledge worker users per node could be supported on a 3-node system as well as graphical performance with GPUs. Tests focused on showing scalable VDI performance from a mid-range CPU.

Tolly found that the PowerEdge R7525 delivered better VDI performance than the predecessor system. Compared to public results from a prior test published in August 2018, the PowerEdge R7525 delivered better response time when lightly loaded as well as when running maximum load across 100, 140 and 200 users on a single system. See Figure 1. (cont. on next page)

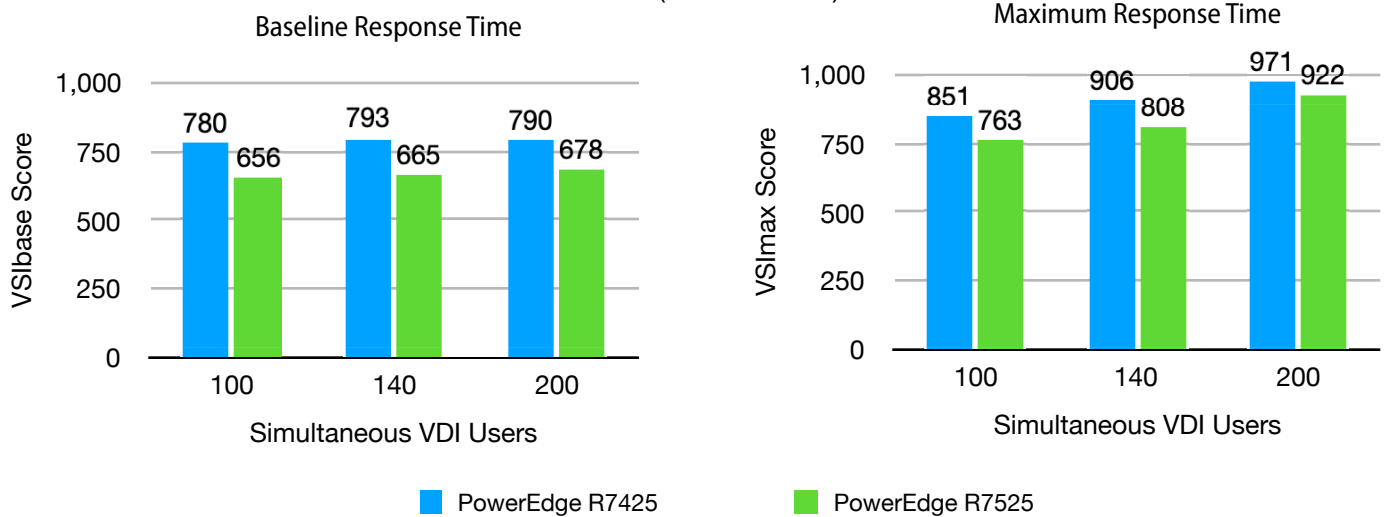
The Bottom Line

The Dell EMC PowerEdge R7525 delivers:

- 1 Improved VDI performance and scalability over the previous generation R7425
- 2 Scaling from 100 to 200 users while still providing sub-second performance
- 3 Scalability on a three-node system of 230 users per node while maintaining sub-second performance
- 4 Support for 90 simultaneous users in GPU configuration with excellent quality and low latency

Dell EMC PowerEdge R7X25 Windows 10 VDI Scaling Login VSI Knowledge Worker Profile on VMware Horizon 7

As reported by Login VSI v4.1
(Lower is better)



Notes: VSI scores are in milliseconds. Windows 10 64-bit. PowerEdge R7425 results from public report published Aug. 2018 by Demartek. Prior test used Windows 10 Pro build 1709, current test used Windows 10 Ent. build 1803. Systems did not have GPUs. R7525 BIOS optimized for VDI.

Source: Tolly, January 2019

Figure 1



The Dell EMC PowerEdge R7525 is built around 2nd Gen AMD EPYC 7002 series processors. What matters, of course, is how any new technology translates into demonstrable benefits for system users. In these tests, we built a VMware ESXi 6.7 hyper-converged infrastructure (HCI) environment that leveraged VMware Horizon 7 to deliver Windows 10 VDI desktops.

To quantify the benefits of the upgraded hardware, test scenarios were run that could be compared with prior generation hardware along with tests that demonstrated reasonable maximum user counts for both office and graphical users. Fundamental to all of these tests is that the user experience remains good as the user load increases.

Scaling vs. Prior Generation

Login VSI was used to benchmark Windows 10 VDI performance at 100, 140 and 200

simultaneous users on a single host. These users ran the Login VSI "Knowledge Worker" workload. This consists, primarily, of common office tasks such as using Microsoft Word, PowerPoint, Internet Explorer and so forth.

Login VSI computes scores for the test with VSIbase representing the lightly loaded system and VSImax representing the fully loaded system. These values represent response time in milliseconds. Thus, for example, a score of 780 means an average response time of .78 seconds.

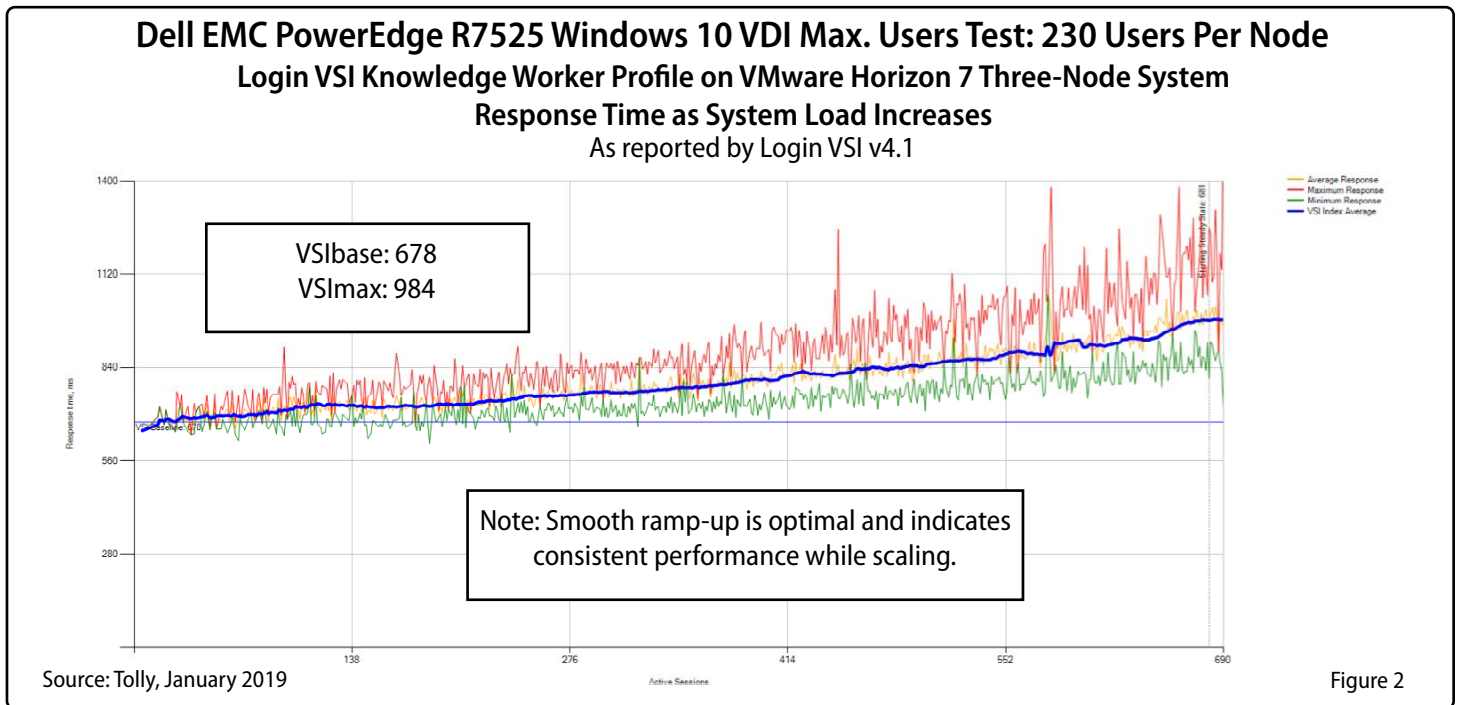
Engineers ran a series of tests designed to be comparable to tests previously published on the prior generation hardware, the Dell EMC PowerEdge R7425.¹

As shown in Figure 1, the PowerEdge R7525 provides better performance than prior generation technology across all workloads evaluated.

With 200 simultaneous users, the baseline (i.e. lightly-loaded) response time is 112 ms faster, on average, across all users. Running at full load with 200 users, the VDI response time with the PowerEdge R7525 remains, on average, under one second (.922 sec) running the later build of Windows 10 Enterprise. This result was 50ms faster than the prior test which ran an earlier build of Windows 10 Professional that was likely less taxing on the system.

Maximum Users

Engineers demonstrated further scalability by building out a three-node system with each system running 230 users. Even with the additional load, the average score at maximum load remained under 1 second (984) and the response time remained consistent as the test scaled. See Figure 2.



¹ https://demartek.principledtechnologies.com/Reports_Free/Demartek_Dell_R7425_AMD_VDI_Knowledge_Worker_2018-08.pdf



Graphical User Scaling

For certain users, graphics are vitally important. Graphics processing units (GPUs) are used to offload compute-intensive graphics rendering from the main CPU.


With VMware Horizon, physical GPUs can become virtual resources that are made available to users on virtual machines.

For this single-server test, six NVIDIA Tesla T4 GPUs (TU104GL) were installed into PCIe slots of the PowerEdge R7525.

As benchmarking graphical performance has requirements that are different than those for primarily text-based applications, a purpose-built graphics benchmarking program was used for this test.

NVIDIA designed its nVector benchmarking program to evaluate the quality of virtualized graphics delivery by measuring image quality, user-experience latency and the rate that video frames are delivered. Ideally, the quality of the image (frame) delivered would be identical to the original, the VDI user would experience minimal delay, and video would be delivered with a consistent frame rate appropriate for the content. And, for all of this to happen, one needs to avoid overloading both the virtualization host and the GPUs.

Tests showed that the PowerEdge R7525 could handle 90 graphics/GPU users providing good results and without overloading system resources. Figure 3 summarizes the results as viewed both from the client and system perspectives.





The system delivered near-perfect (0.99 out of 1.0) image quality with less than a quarter of a second lag while maintaining an average frame rate of 8.82FPS.

Furthermore, system resources were not overtaxed with the host CPU remaining under 85% and the GPUs averaging ~35%.

Dell EMC PowerEdge R7525 Windows 10 VDI Graphical Users Test: 90 Users nVector Knowledge Worker Profile on VMware Horizon 7 with 6x NVIDIA Tesla T4 GPUs System Resource Utilization & Client Performance As reported by VMware TOPS & NVIDIA nVector v1.0

Client Performance (Average)

Image Quality (1.0 is perfect)	Latency (ms)	Frame Rate (FPS)
0.9904	215	8.82



Source: Tolly, January 2019

Figure 3



Test Configuration

All testing was conducted using Dell PowerEdge R7525 servers running VMware ESXi in a three-node HCI/vSAN configuration. See Table 1.

The system BIOS profile was set to "Performance" with one modification. The Determinism Slider was set to "Performance Determinism." This sets the C-State² of the CPU to C0 (full power) and is known to improve the performance of a VDI system.

VDI & Virtual Client Configuration

VMware's Horizon 7 was used for all testing. Client configurations were identical for tests and all clients ran the current version of Windows 10 Enterprise. See Table 2.

Login VSI Benchmarking

Login VSI was used for all of the tests except for benchmarking that involved the GPUs. Login VSI measures response time across a range of applications running on the VDI client. All clients ran the "Knowledge Worker" workload profile.³ This workload consists of programs such as Microsoft Word, Microsoft Outlook, Internet Explorer, Adobe Reader, etc. See Table 3.

Each test was run for 60 minutes and was run twice. Results were very similar for each run. The better run was used for the report.

For the 100, 140 and 200-user scalability tests, tests were run on a single host of the cluster to be comparable with the prior tests. The maximum user test was run across three hosts with each host handling 230 users.

² <https://www.dell.com/support/article/us/en/04/qna41893/what-is-the-c-state?lang=en>

³ <https://www.loginvsi.com/blog-alias/login-vsi/665-simulating-vgi-users-introduction-to-login-vsi-workloads>

Factors & Metrics for Measuring Graphics UX with NVIDIA nVector

Typical VDI deployments have two conflicting goals: Achieving the best possible user experience and maximizing user density on server hardware. Problems arise as density is scaled up, however, because after a certain point it negatively impacts user experience.

NVIDIA nVector measures these trade-offs by modeling how knowledge workers use applications and what happens to performance when workloads are run at scale. There are three key metrics:

End-User Latency — Measures the responsiveness of a remote desktop or application. It measures the lag that an end user experiences when the user interacts with a remote desktop or application.

Frame Rate — Measures the perceived "smoothness" of the UX. It actually measures the rate at which frames are delivered on the screen of the endpoint device.

Image Quality — Measures the impact of "remoting" on image quality. It quantifies the impact using the Structural Similarity Index (SSIM) between a frame that is rendered on the target virtual machine and the frame that is visible to the end user.

Source: Dell Technologies

Excerpted and adapted from:

<https://www.dell.com/resources/en-us/asset/white-papers/solutions/h17917-quantify-impact-virtual-gpus-wp.pdf>

nVector Graphics Benchmarking

NVIDIA nVector was used for benchmarking performance of graphics. nVector has its own implementation of a "Knowledge Worker" workload profile that is oriented toward tasks that benefit from GPU processing. See Table 4.

For this test, six GPU cards were installed on a single host server. The test was run on a single server. The test measures quality of experience for graphical users by measuring user latency, frame rate and image quality. See sidebar.

The test also measured key system hardware resources, in particular host CPU utilization and processor utilization for each of the GPUs in the system.

Dell EMC PowerEdge R7525

The R7525 is a highly adaptable 2U 2 socket rack server that brings powerful performance and flexible configurations. Some key features include:

- Support for 2x 2nd Generation AMD EPYC 7002 processors with up to 64 cores each
- Large memory footprint supporting up to 16 DIMMs per processor (32 total)
- Flexible storage options from 8x3.5" to as many as 24x2.5" drives
- Future ready with support for PCIe Gen4 cards
- Flexible networking support including an OCP 3 slot

Source: Dell Technologies



VDI Test System Configuration Summary

Server Configuration

Make & Model	Dell EMC PowerEdge R7525
Host OS	VMware ESXi 6.7.0 Build 15160138
VDI	VMware Horizon 7.10.0. Build 14584133
BIOS	1.2.9. Profile: "Performance" except for Determinism Slider was set to "Performance Determinism."
CPU	2x AMD EPYC 7502 (32 cores, 64 threads each) @ 2.5GHz
Memory (RAM)	16x Micron 36ASF8G72PZ-3G2B2 64 GB Dual-Rank DDR4 3200 MHz (1024GB total)
AHCI Controller	Marvell BOSS-S1
Video Controller (Graphics test only)	6x NVIDIA TU104GL [Tesla T4]
Capacity Drives	4x Toshiba and Samsung 1.92TB SAS plus 2x Toshiba 800GB SAS. All non-RAID using HBA345 Front
Network Interface Cards	25GbE, dual-port Mellanox MT27800 Family [ConnectX-5] Firmware 16.25.40.62
LAN Switch	Dell EMC PowerSwitch S5248-ON with OS 10.5.0.2

Table 1

VM Configuration

Client OS	Windows 10 Enterprise Build 1803 64-bit
vCPUs	2
Memory	4 GB Ram 2048GB reserved
Hard Disk	60 GB
Adapter	VMXNET 3
VMware Horizon Client	5.3.0

Graphics Test Information

Screen Resolution	1920x1080
vMemory	6144MB

Table 2

Test Tool: Login VSI Configuration

Vendor	Login VSI
Tool	Login VSI 4.1.40.1
Workload	Knowledge Worker
Test Parameters	60 minute run time. One new session every 5.22 sec.

Table 3

Graphics Test Tool: nVector Configuration

Vendor	NVIDIA Corp.
Tool	nVector 1.0
Workload	nVector Knowledge Worker
Test Parameters	60 minutes run time. One new session every 5 seconds.

Table 4

Source: Tolly, January 2019



About Tolly

The Tolly Group companies have been delivering world-class information technology services for over 30 years. Tolly is a leading global provider of third-party validation services for vendors of information technology 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.