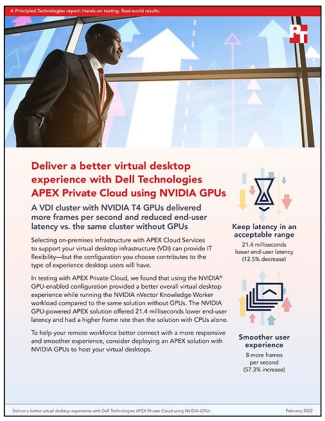

The science behind the report:

Deliver a better virtual desktop experience with Dell Technologies APEX Private Cloud using NVIDIA GPUs

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report Deliver a better virtual desktop experience with Dell Technologies APEX Private Cloud using NVIDIA GPUs.

We concluded our hands-on testing on November 23, 2021. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on November 11, 2021 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to http://facts.pt/calculating-and-highlighting-wins. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Results of our virtual desktop infrastructure (VDI) testing.

	Without GPUs	With NVIDIA® GPUs
Average end-user latency (ms) (lower is better)	170.09	148.67
Frame rate in frames per second (FPS) (higher is better)	14.19	22.33

System configuration information

Table 2: Detailed information on the systems we tested.

System configuration information	3x Dell PowerEdge™ R730	4x Dell VxRail™ V570F
Version	N/A	VxRail 7.0.200-17911444
BIOS name and version	2.10.2	2.10.2
Non-default BIOS settings	Default	APEX default
Operating system name and version/build number	VMware® ESXi™, 7.0.2, 17867351	VMware ESXi, 7.0.2, 17867351
Processor		
Number of processors	2	2
Vendor and model	Intel® Xeon® E5-2697 v3	Intel Xeon Gold 6248R
Core count (per processor)	14	24
Core frequency (GHz)	2.6	3.0
Stepping	C0 (QGEF) C1 (QGN3, SR1XF)	B1 (SRGZG)
Memory module(s)		
Total memory in system (GB)	512	768
Number of memory modules	16	12
Vendor and model	Samsung® M393A4K40BB0-CPB	SK Hynix HMAA8GR7AJR4N
Size (GB)	32	64
Type	DDR4	DDR4
Speed (MHz)	2,133	3,200
Speed running in the server (MHz)	2,133	2,933
Storage controller		
Vendor and model	Dell HBA330	Dell HBA330
Distributed storage		
Storage type	VMware vSAN™	VMware vSAN
Number of nodes	3	4
Number of disk groups	3	8
Number of disks per group	6	5
Network adapter		
Vendor and model	Intel X710	Mellanox ConnectX-4MT27710 Family
Number and type of ports	2 x 10GbE	2 x 25GbE
Driver version	1.10.9.0-1OEM.700.1.0.15525992	14.28.45.12
Power supplies		
Number of power supplies	2	2

How we tested

The APEX Private Cloud solution we tested was deployed in an offsite data center. We used four Dell VxRail V570 nodes as the cluster under test to host the VDI sessions and used three Dell PowerEdge R730 servers as infrastructure and client servers.

We performed all testing remotely after inspecting the cluster under test, server clients, the network implementation, and the storage. We had full control and unfettered access to the testbed during the duration of testing. We used the same testbed for all our tests with and without GPUs.

We conducted testing with assistance from Dell Technologies engineers, who provided custom PowerShell scripts to assist with setup and deployment of VDI infrastructure, host balancing, and test execution.

For information on how to set up and run NVIDIA nVector, please visit https://developer.nvidia.com to request access to software and documentation.

Creating pools (192 users)

1. Use remote desktop protocol to access the controller VM.
2. Open PowerShell.
3. Navigate to C:\Scripts.
4. To remove the current pool, run RemoveTestPool.ps1.
5. Check within the VMware Horizon® admin console to ensure the VMs are deleted completely.
6. Run either nVector-CreateTestPool-vGPU-t4-1b.ps1 or nVector-CreateTestPool-NoGPU.ps1, depending on whether you want a vGPU-enabled pool or not.
7. Navigate to the Horizon admin console, and log in.
8. Under Inventory, click Desktops.
9. Check the box next to Cirrus Pool ID.
10. Click Edit.
11. Click Desktop Pool Settings.
12. Scroll down to Refresh OS Disk After Logoff.
13. Set to Always, and click OK.

Preparing the environment

1. Use remote desktop protocol to access the controller VM.
2. Open PowerShell.
3. Navigate to C:\Scripts.
4. Run Logoff_users.ps1.
5. Run Restart-EndPoints.ps1.
6. If this is a non-GPU test, run the BalanceHosts.ps1 script.

Running the nVector test

1. Use remote desktop protocol to access the test VM.
2. If this is the first login, double-click Start nVector Server. This will start the nVector web server in a console window.
3. Open PowerShell.
4. Navigate to Desktop.
5. Run Ping-VMs.ps1. All of the VMs should return the FQDN. If any do not, reboot them manually from the Horizon admin console or VMware vSphere® client.
6. Navigate to C:\Cirrus.
7. Run Run-nVector-TEST-vGPU.ps1 or Run-nVector-TEST-noGPU.ps1 depending on which type of desktop pool is currently deployed. If DNS errors occur, log into the DNS server, clear cache, and scavenge stale resource records.
8. The test will start and run for several hours. Monitor the workload by running TightVNC viewer and connecting to the endpoint VMs. Do not disturb the desktop. Monitor using SnakeTail to view the GridFramework.log, which has logging output from the test.
9. Once the test is complete, run Run-nVector-Analysis.ps1.

Read the report at https://facts.pt/Of9XvuT ▶

This project was commissioned by Dell Technologies.

Principled Technologies®

Facts matter.®