



(/t/blogs-70,^

Home (/) > Integrated Products (/integrated-products/) > vSAN Ready Nodes (/vsan-ready-nodes/) > Blogs (/t/blogs-70/) > Dell EMC vSAN Ready Nodes: Taking VDI and AI Beyond “Good Enough”

in

f

🐦

Dell EMC vSAN Ready Nodes: Taking VDI and AI Beyond “Good Enough”

October 18th, 2021 | Read Time: 8 minutes

TF

Tony Foster

Some people have speculated that 2020 was “the year of VDI” while others say that it will never be the “year of VDI.” However, there is one certainty. In 2020 and part of 2021, organizations worldwide consumed a large amount of virtual desktop infrastructure (VDI). Some of these deployments went extremely well while other deployments were just “good enough.”

If you are a VDI enthusiast like me, there was much to learn from all that happened over the last 24 months. An interesting observation is that test VDI environments turned into production environments overnight. Also, people discovered that the capacity of clouds is not limitless. My favorite observation is the discovery by many

IT professionals that GPUs can change the VDI experience from “good enough” to enjoyable, especially when coupled with an outstanding environment powered by Dell Technologies with VMware vSphere and VMware Horizon.

In this blog, I will tell you about how exceptional VDI (and AI/ML) is when paired with powerful technology.

This blog does not address cloud workloads as it is a substantial topic. It would be difficult for me to provide the proper level of attention in this blog, so I will address only on premises deployments.

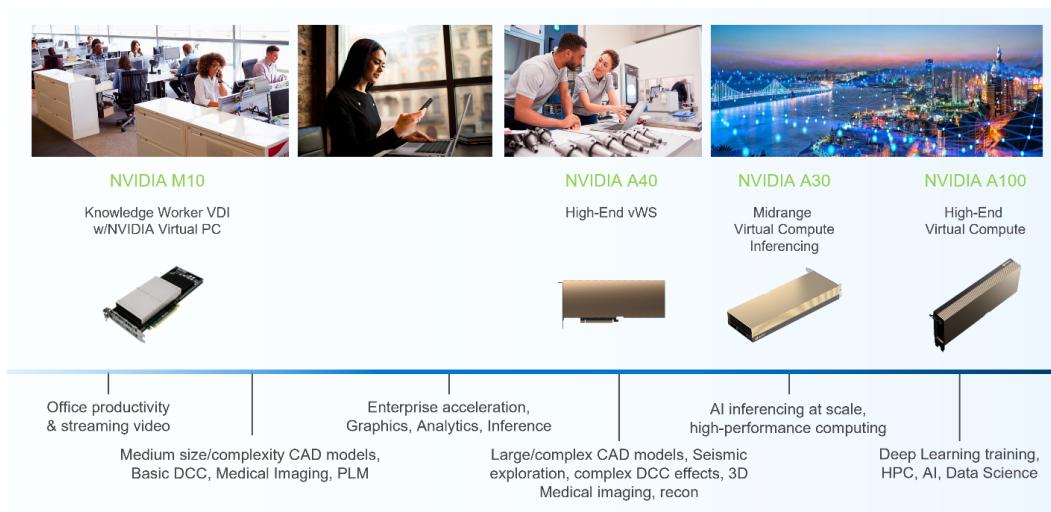
Many end users adopt hyperconverged infrastructure (HCI) in their data centers because it is easy to consume. One of the most popular HCIs is Dell EMC VxRail Hyperconverged Infrastructure (<https://www.delltechnologies.com/en-us/converged-infrastructure/vxrail/index.htm>). You can purchase nodes to match your needs. These needs range from the traditional data center workloads, to Tanzu clusters, to VDI with GPUs, and to AI. VxRail enables you to deliver whatever your end users need. Your end users might be developers working from home on a containers-based AI project and they need a development environment, VxRail can provide it with relative ease.

Some IT teams might want an HCI experience that is more customer managed but they still want a system that is straightforward to deploy, validate, and is easy to maintain. This scenario is where Dell EMC vSAN Ready Nodes (<https://www.delltechnologies.com/en-us/converged-infrastructure/hyper-converged-infrastructure/vsan-ready-nodes.htm>) come into play.



Dell EMC vSAN Ready Nodes provide comprehensive, flexible, and efficient solutions optimized for your workforce’s business goals with a large choice of options (more than 250 as of the September 29, 2021 vSAN Compatibility Guide (https://www.vmware.com/resources/compatibility/pdf/vi_vsan_rn_guide.pdf)) from

tower to rack mount to blades. A surprising option is that you can purchase Dell EMC vSAN Ready Nodes with GPUs, making them a great platform for VDI and virtualized AI/ML workloads.



Dell EMC vSAN Ready Nodes supports many NVIDIA GPUs

(<https://www.delltechnologies.com/en-us/servers/server-accelerators.htm>) used for VDI and AI workloads, notably the NVIDIA M10 and A40 GPUs for VDI workloads and the NVIDIA A30 and A100 GPUs for AI workloads. There are other available GPUs depending on workload requirements, however, this blog focuses on the more common use cases.

For some time, the NVIDIA M10 GPU (<https://www.nvidia.com/en-us/data-center/tesla-m10/>) has been the GPU of choice for VDI-based knowledge workers who typically use applications such as Microsoft PowerPoint and YouTube. The M10 GPU provides a high density of users per card and can support multiple virtual GPU (vGPU) profiles per card. The multiple profiles result from having four GPU chips per PCI board. Each chip can run a unique vGPU profile, which means that you can have four vGPU profiles. That is, there are twice as many profiles than are provided by other NVIDIA GPUs. This scenario is well suited for organizations with a larger set of desktop profiles.

Combining this profile capacity with Dell EMC vSAN Ready Nodes, organizations can deliver various desktop options yet be based on a standardized platform. Organizations can let end users choose the system that suites them best and can optimize IT resources by aligning them to an end user’s needs.

Typically, power users need or want more graphics capabilities than knowledge workers. For example, power users working in CAD applications need larger vGPU profiles and other capabilities like NVIDIA’s Ray Tracing (<https://developer.nvidia.com/discover/ray-tracing>) technology to render drawings.

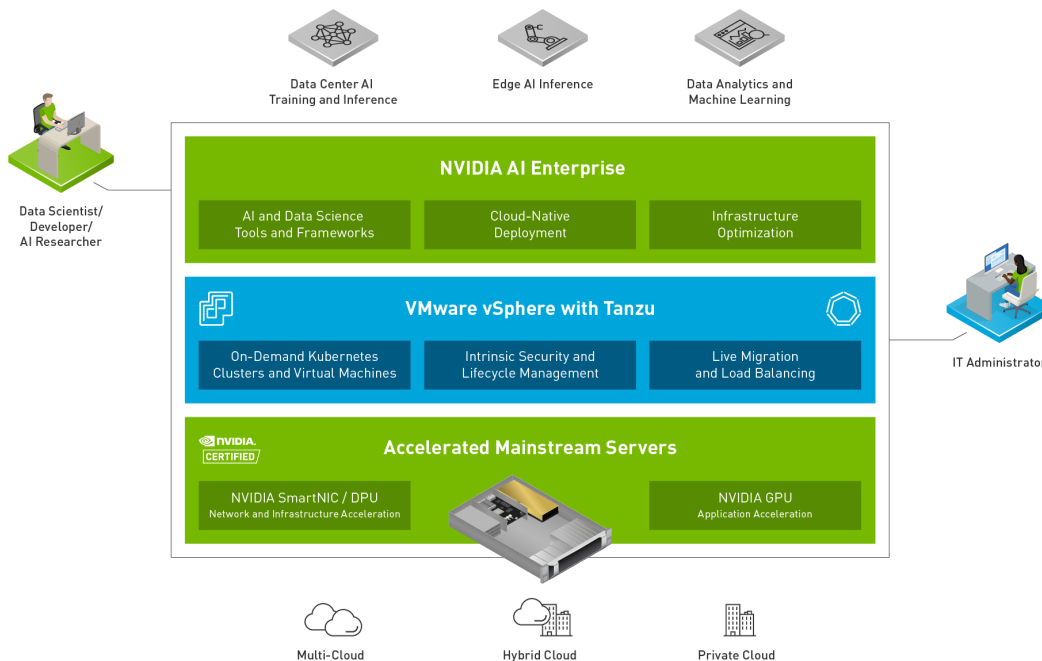
These power users’ VDI instances tend to be more suited to the NVIDIA A40 GPU (<https://www.nvidia.com/en-us/data-center/a40/>) and associated vGPU profiles. It allows power users who do more than create Microsoft PowerPoint presentations and watch YouTube videos to have the desktop experience they need to work

effectively.

The ideal Dell EMC vSAN Ready Nodes platform for the A40 GPU is based on the Dell EMC PowerEdge R750 server. The PowerEdge R750 server provides the power and capacity for demanding workloads like healthcare imaging and natural resource exploration. These workloads also tend to take full advantage of other features built into NVIDIA GPUs like CUDA (<https://developer.nvidia.com/cuda-faq>). CUDA is a parallel computing platform and programming model that uses GPUs. It is used in many high-end applications. Typically, CUDA is not used with traditional graphics workloads.

In this scenario, we start to see the blend between graphics and AI/ML workloads. Some VDI users not only render complex graphics sets, but also use the GPU for other computational outcomes, much like AI and ML do.

I really like that I can run AI/ML workloads in a virtual environment. It does not matter if you are an IT administrator or an AI/ML administrator. You can run AI and ML workloads in a virtual environment (<https://www.vmware.com/products/vsphere/ai-ml.html>).



Many organizations have realized that the same benefits virtualization has brought to IT can also be realized in the AI/ML space. There are additional advantages, but those are best kept for another time.

For some organizations, IT is now responsible for AI/ML environments, whether delivering test/dev environments for programmers or delivering a complete AI training environment. For other IT groups, this responsibility falls to highly paid data scientists. And for some IT groups, the responsibility is a mix.

In this scenario, virtualization shines. IT administrators can do what they do best: deliver a powerful Dell EMC vSAN Ready Node infrastructure. Then, data scientists can spend their time building systems in a virtual environment consuming IT resources instead of racking and cabling a server.

Dell EMC vSAN Ready nodes are great for many AI/ML applications. They are easy to consume as a single unit of infrastructure. Both the NVIDIA A30 GPU and the A100 GPU are available so that organizations can quickly and easily assemble the ideal architecture for AI/ML workloads.

This ease of consumption is important for both IT and data scientists. It is unacceptable when IT consumers like data scientists must wait for the infrastructure they need to do their job. Time is money. Data scientists need environments quickly, which Dell EMC vSAN Ready Nodes can help provide. Dell EMC vSAN Ready Nodes deploy 130 percent faster with Dell EMC OpenManage Integration for VMware vCenter (OMIVV) (Based on Dell EMC internal competitive testing of PowerEdge and OMIVV compared to Cisco UCS manual operating system deployment.)

This speed extends beyond day 0 (deployment) to day 1+ operations. When using the vLCM and OMIVV, complete hypervisor and firmware updates to an eight-node PowerEdge cluster took under four minutes compared to a manual process, which took 3.5 hours. (Principle Technologies report commissioned by Dell Technologies, New VMware vSphere 7.0 features reduced the time and complexity of routine update and hardware compliance tasks (<https://www.principledtechnologies.com/Dell/vSphere-7.0-update-compatibility-check-0820-v3.pdf>), July 2020.)

Dell EMC vSAN Ready Nodes ensures that you do not have to be an expert in hardware compatibility. With over 250 Dell EMC vSAN Ready Nodes available (as of the September 29, 2021 vSAN Compatibility Guide (https://www.vmware.com/resources/compatibility/pdf/vi_vsan_rn_guide.pdf)), you do not need to guess which drives will work or if a network adapter is compatible. You can then focus more on data and the results and less on building infrastructure.

These time-to-value considerations, especially for AI/ML workloads, are important. Being able to deliver workloads such as AI/ML or VDI quickly can have a significant impact on organizations, as has been evident in many organizations over the last two years. It has been amazing to see how fast organizations have adopted or expanded their VDI environments to accommodate everyone from knowledge workers to high-end power users wherever they need to consume IT resources.

Beyond “just expanding VDI” to more users, organizations have discovered that GPUs can improve the end-user experience and, in some cases, not only help but were required. For many, the NVIDIA M10 GPU helped users gain the wanted remote experience and move beyond “good enough.” For others who needed a more graphics-rich experience, the NVIDIA A40 GPU continues to be an ideal choice.

When GPUs are brought together as part of a Dell EMC vSAN Ready Node, organizations have the opportunity to deliver an expanded VDI and AI/ML experience to their users. To find out more about Dell EMC vSAN Ready Nodes, see Dell EMC vSAN Ready Nodes (<https://www.delltechnologies.com/en-us/converged-infrastructure/hyper-converged-infrastructure/vsan-ready-nodes.htm>).

Author: Tony Foster **Twitter:** @wonder_nerd **LinkedIn:**

<https://linkedin.com/in/wondernerd> (<https://linkedin.com/in/wondernerd>)

Tags: AI PowerEdge VxRail VDI machine learning
vSAN GPU

Related Blog Posts



(/p/new-frontiers-dell-emc-powerededge-r750xa-server-with-nvidia-a100-gpus/)

NVIDIA AI PowerEdge machine learning HPC GPU

New Frontiers—Dell EMC PowerEdge R750xa Server with NVIDIA A100 GPUs

(/p/new-frontiers-dell-emc-powerededge-r750xa-server-with-nvidia-a100-gpus/)

DC FF SP Deepthi Cherlopalle, Frank Han, Savith Pareek

June 1st, 2021 | Read Time: 7 minutes

(/p/new-frontiers-dell-emc-powerededge-r750xa-server-with-nvidia-a100-gpus/)

Dell Technologies has released the new PowerEdge R750xa server, a GPU workload-based platform that is designed to support artificial intelligence, machine learning, and high-performance computing solutions. The dual socket/2U platform supports 3rd Gen Intel Xeon processors (code



(/p/comparison-of-mlperf-tm-inference-v1-1-results-of-dell-emc-powerededge-r7525-servers-with-nvidia-gpus/)

NVIDIA PowerEdge machine learning GPU MLPerf

Comparison of MLPerf™ Inference v1.1 Results of Dell EMC PowerEdge R7525 Servers with NVIDIA GPUs

(/p/comparison-of-mlperf-tm-inference-v1-1-results-of-dell-emc-powerededge-r7525-servers-with-nvidia-gpus/)

M FF RV Manpreet Sokhi, Frank Han, Rakshith Vasudev

October 15th, 2021 | Read Time: 11 minutes

(/p/comparison-of-mlperf-tm-inference-v1-1-results-of-dell-emc-powerededge-r7525-servers-with-nvidia-gpus/)

This blog showcases the MLPerf Inference v1.1 performance results of Dell EMC PowerEdge

© 2021 Dell Inc. (<https://www.dell.com/learn/us/en/uscorp1/site-terms-of-use-copyright>)

Privacy (<https://www.dell.com/learn/us/en/uscorp1/policies-privacy>)

Terms Of Use (<https://www.dell.com/learn/us/en/uscorp1/site-terms-of-use>)

Legal (<https://www.dellemc.com/en-us/customer-services/product-warranty-and-service-descriptions.htm>)

Anti-Slavery and Human Trafficking (<https://i.dell.com/sites/doccontent/corporate/corp-comm/en/Documents/dell-california-trafficking.pdf>)