D≪LL Technologies

NVIDIA.

Simplifying AI systems

University flexes to accommodate new users and workloads with VMware and NVIDIA AI Enterprise software on Dell EMC infrastructure.

"

"The biggest benefit of virtualized GPUs is flexibility, in the sense that you can design and adapt your enterprise infrastructure to AI workloads."

Maurizio Davini
Chief Technology Officer,
University of Pisa

Business needs

The University of Pisa needs to support massive increases in new and emerging high performance computing workloads, including those for artificial intelligence.

Business results

- Simplifying the deployment and management of AI systems

- Increasing system utilization and reducing costs

- Building flexibility into IT processes and services

- Extending the power of NVIDIA GPUs to more users

Solutions at a glance

- Dell EMC VxRail hyperconverged infrastructure

- Dell EMC PowerEdge servers with NVIDIA GPUs

- Dell EMC PowerScale storage

- NVIDIA AI Enterprise software suite

Supporting AI and other modern workloads

In today's academic and enterprise environments, IT organizations need to support growing numbers of artificial intelligence applications, along with many other new and emerging high performance computing workloads. This mission creates challenges that aren't easily solved with conventional IT technologies and approaches. In a typical scenario, IT administrators set up separate systems to accommodate AI workloads, and then manage many aspects of those systems with time-intensive manual processes.

Today, there is a better way forward — thanks to the integration of technologies from VMware® and NVIDIA®. With NVIDIA now working directly with VMware code, IT organizations can run traditional and AI workloads in the same environment and on the same server and storage systems. This makes life much easier for the people tasked with delivering the systems to support modern workloads. AI applications can now be easily managed with the same VMware flexibility as with other applications. As a result, it's now easier than ever to develop, deploy and manage diverse AI workloads.

In addition, with the tight integration of VMware and NVIDIA software, organizations can now virtualize multiple technologies inside the systems. For example, they can virtualize and share the GPUs inside the servers, so multiple data scientists can simultaneously accelerate their deep learning workloads. This increases utilization rates and saves money that would have otherwise been spent on the procurement and management of additional, more isolated hardware.

This integration also saves time and administrative steps. And perhaps best of all, IT organizations can manage everything centrally with VMware vCenter and easily allocate resources as needed.

These are the kinds of benefits that the University of Pisa is realizing with its use of a groundbreaking solution that leverages the combined capabilities of Dell EMC VxRail, Dell EMC PowerScale and NVIDIA® AI Enterprise software suite.

Solution components

Dell EMC VxRail

VxRail is the only hyperconverged infrastructure (HCI) jointly engineered with VMware, continuously integrating the latest in software and hardware advancements for a seamless, curated and optimized experience.

A simple, cost effective HCI, VxRail solves a wide range of IT challenges and supports almost any use case, including tier-one applications and mixed workloads. VxRail enables faster, better and simpler delivery of VMware-virtualized applications.

This deep integration across the VMware ecosystem, combined with the simplicity of VxRail HCI system software, provides an ideal platform across core, edge and cloud.

Dell EMC PowerScale

Dell EMC PowerScale storage is designed to serve as the foundation for data from proof of concept to production. These all-flash scale-out network-attached storage (NAS) systems deliver the performance and extreme concurrency at scale to consistently feed the most data-hungry deep learning algorithms.

With the PowerScale OneFS operating system, Dell EMC PowerScale storage has the data storage governance, management, security, compliance and protectio to help ensure AI solutions conform to regulatory and enterprise security policy requirements.

NVIDIA AI Enterprise

NVIDIA AI Enterprise is a software suite of enterprise-grade AI tools and frameworks that is optimized, certified and supported by NVIDIA with VMware vSphere. With this software, IT professionals at the hundreds of thousands of enterprises that use vSphere for virtualization can now support AI with the same tools they use to manage large-scale data centers and hybrid cloud environments. NVIDIA AI Enterprise provides scale-out, multi-node, AI application performance on vSphere that is indistinguishable from bare-metal servers.

And forget about AI silos. With NVIDIA AI Enterprise running on vSphere, organizations can avoid silos of AI-specific systems that create management headaches and security risks.

With this combination of technologies, organizations can overcome the challenges that stem from deploying individual AI applications, as well as the potential failures that can result from having to manually provision and manage different applications and infrastructure software.

A center of excellence

The University of Pisa is both a Dell Technologies and a VMware Center of Excellence. As part of this designation, the University's IT team regularly evaluates and tests new technologies. That's the case today with the combination of Dell EMC VxRail with VMware, Dell EMC PowerScale for AI, and NVIDIA AI Enterprise.

"We are running AI workloads on top of VMware, and we are using Dell EMC PowerStore for the storage in this virtualized environment," says Maurizio Davini, chief technology officer for the University of Pisa. "And we have a Dell EMC PowerScale all-flash environment for AI and HPC as a sort of traditional scale-out in our fast systems."

Davini notes that his organization has both bare-metal and virtualization implementations of PowerEdge servers with NVIDIA GPUs.

"We have traditional bare-metal GPUs, which are typically used for research like language processing, image processing, deep learning, deep neural network research and so on," he says. "So we are still increasing our bare-metal capabilities on GPUs. And we now also have clusters of GPUs inside our VMware production environment."

The University currently has several VMware vCenter implementations and several clusters running databases or authentication or administrative applications together in parallel with the HPC and AI workloads.

"The biggest benefit of virtualized GPUs is flexibility, in the sense that you can design and adapt your enterprise infrastructure to your AI workloads," Davini says. "You can distribute the workload without having to reconfigure everything."

In a complementary innovation, the University is using VMware vSphere with containers inside virtual machines, Davini says. And looking ahead, the team has plans to add support for Docker containers with integrated VMware Tanzu. VMware Tanzu automates the delivery of containerized applications.

"The main benefit of this is good performance, and that

The University of Pisa is both a Dell Technologies and a VMware Center of Excellence. As part of this designation, the university's IT team regularly evaluates and tests new technologies.

you have a packet application that you can run preloaded inside your environment," Davini says. "We are moving to Kubernetes to manage these Docker containers, and then as soon after we will move to VMware Tanzu, and then the Kubernetes part will be on the VMware side."

The benefits of virtualization

Virtualization of the infrastructure is one of the keys to gaining the greatest value from the University's investments, Davini notes. With vSAN, for example, the University can gain the benefits of hyperconverged storage stretched across several data centers. That enables the easy migration of workloads from one data center to another, while simplifying management and disaster recovery.

"That was our first approach," he says. "And then we introduced VMware Horizon for VDI."

VMware Horizon is a solution designed to simplify the management and delivery of virtual desktops and apps to end users via a single platform. It helps the University's IT administrators control, manage and protect all of the Windows resources end users want, at the speed they expect.

Davini and his colleagues initially used Horizon to meet the computing needs of laboratories, administrative offices and other University units. They then expanded the focus to the use of virtual machines for HPC.

"VMware gives us the possibility to be flexible and to use the infrastructure for a lot of things — enterprise workloads, VDI, remote workstations, support for smart working, scientific computing, HPC — all in the same infrastructure in a very flexible way," Davini says. "And this is the problem that VMware and Dell have helped us to solve."

Delivering Developer- and AI-Ready Infrastructure to Advance Digital Business

- NVIDIA exclusively certified the new VMware vSphere release for NVIDIA AI Enterprise suite, a cloud-native collection of optimized AI applications and frameworks, for an end-to-end AI solution.

- vSphere introduced support for the NVIDIA A100 Tensor Core GPUs in NVIDIA-Certified Systems™. This, with NVIDIA AI Enterprise, enables customers to fold-in an end-to-end AI solution with confidence on their existing enterprise virtualization platform, instead of running AI projects in separate unmanageable IT silos; and,

- Customers are also able to incorporate the latest generation of NVIDIA GPUs into their virtual environment and take advantage of features like Multi-Instance GPU (MIG) allowing GPU cycles to be shared across multiple users; vSphere vMotion to provide live migration for non-disruptive operations; and vSphere Distributed Resource Scheduler (DRS) for automatic initial workload placement to avoid performance bottlenecks.

"

"NVIDIA AI Enterprise is a software suite optimized, certified and supported on VMware vSphere that enables customers to rapidly deploy, manage and scale AI in production with confidence," said Justin Boitano, Vice President and General Manager of Enterprise and Edge Computing, NVIDIA. "Through NVIDIA's collaboration with VMware, IT professionals can now support business teams with the industry's most trusted AI tools across their hybrid cloud infrastructure."

VMware news release, "VMware Evolves Developer and AI-Ready Infrastructure to Advance Digital Business," March 9, 2021.

Learn more about Dell EMC advanced computing

Unlock the value of data with artificial intelligence

DELLTechnologies

Connect on Social