



## Industry Support for NVIDIA vComputeServer

“Enterprises today increasingly need powerful, GPU-accelerated virtual machines to support their compute-intensive workloads. NVIDIA vComputeServer running on Cisco HyperFlex and UCS will enable our customers to get the performance and flexibility needed for their deep learning and data analytics applications.”

— Kaustubh Das, vice president of Product Management and Strategy, Computing Systems Product Group, Cisco

“Dell EMC PowerEdge servers with NVIDIA vComputeServer greatly simplify GPU virtualization for analytics, AI and HPC workloads. This combination can reduce IT expenses while boosting efficiency and agility, and it’s available with the ease of Dell Technologies’ world-class infrastructure solutions, expertise and financing.”

— Rajesh Pohani, vice president of Product Management, Server & Infrastructure Systems, Dell EMC

“Our customers are increasingly looking for ways to accelerate their AI development. Deploying NVIDIA vComputeServer on Lenovo ThinkSystem servers capitalizes on the power of scale out to drive AI deployments faster and more economically. We are excited about its potential to fast-track AI for our customers and speed their efforts in solving humanity’s greatest challenges.”

— Scott Tease, executive director for HPC and AI, Lenovo Data Center Group

“Increasingly, enterprises are turning to AI-augmented work and decision science to stay competitive and to bolster ROI. However, traditional CPU and networking technologies are no longer sufficient to support the size and complexity of machine learning workloads. vSphere’s Advanced Paravirtualization over Mellanox RDMA (PVRDMA) Networking Technology and NVIDIA vComputeServer powered by the T4 has been proven to significantly accelerate modern applications in virtualized environments, enabling the most efficient VDI deployments.”

— Motti Beck, senior director of Enterprise Market Development, Mellanox Technologies

“Organizations see a rising need for GPU acceleration to support a variety of workloads from data science to rendering. With the launch of NetApp HCI with NVIDIA T4 GPUs, NetApp’s support for vComputeServer will help companies manage their AI workflows, improve infrastructure utilization and achieve better time to insights.”

— Russell Fishman, senior director of Product Management, HCI Solutions, NetApp

“As AI and data-driven computing are added as integral components of modern IT environments, these workloads greatly benefit from accelerated computing power from GPUs. As part of our commitment to customer choice in hybrid cloud deployments, our customers are able to deploy Red Hat OpenShift on virtual machines running on vSphere and, with vComputeServer and NGC containers, our customers will be able to add GPUs to these virtual environments to run even more complex workflows.”

— Julio Tapia, director of Cloud Platforms Partner Ecosystem, Red Hat

“Increasingly, businesses are relying on GPUs to accelerate HPC and deep learning workloads. NVIDIA’s new vComputeServer software empowers data scientists to better streamline access for their GPUs. Leveraging Supermicro GPU systems and vComputeServer, data center administrators can deliver powerful performance users expect to run almost any workflow on a virtual machine.”

— Jim Sangster, senior director of Product Marketing, Supermicro

“Today’s enterprises are running compute-intensive AI and data analytics workloads, and the need for hardware acceleration is becoming more prominent. With our NVIDIA partnership and our recent acquisition of Bitfusion, we are delivering the next generation of efficient and performant AI infrastructure. NVIDIA vComputeServer and NGC support for vSphere allows our customers to easily handle AI and HPC apps with the best performance possible on a virtual machine and we extend this to remote servers for GPU pooling with Bitfusion.”

— Krish Prasad, senior vice president and general manager of Cloud Platform Business, VMware