

DELLTechnologies

Dell Technologies AI Solutions

 **nVIDIA**

Table of contents

Drive innovation with AI	3
AI is already making an impact	4
The Dell Technologies point of view.....	5
Overcoming the challenges of AI adoption	6
Easy for IT.....	7
Validated design for AI in virtualized environments.....	8
NVIDIA technologies are built in	9
Power your AI journey with PowerScale.....	11
Eliminate I/O bottlenecks with PowerScale	11
Accelerate GPU applications with RDMA data access on PowerScale	12
PowerScale deployment options for AI	13
Learn more.....	14

Drive innovation with AI

In the drive to digital transformation and the rise of the data-driven organization, artificial intelligence (AI) has emerged as one of the most important technologies to help organizations innovate. AI is helping modern businesses save time and money while reducing risk by quickly identifying trends and patterns, improving IT management, and eliminating manual tasks.

Business leaders are recognizing the game changing potential, and competitive advantage, of AI. Organizations that fail to leverage AI today will fall behind their competitors tomorrow.

92% of businesses report that the pace of AI investment is accelerating¹

49% of businesses are driving innovation with data¹

77% of Fortune 500 CEO's plan to start or already employ AI initiatives²

The need to modernize, coupled with the massive amounts of data available, has fueled the adoption of AI as a business game changer.



AI is already making an impact

AI offers tremendous potential for specific industries. If implemented in the right way — it can help you optimize operations, improve overall productivity, and utilize manpower for more important tasks. Discover the possibility, power, and performance of AI computing from concept to reality in your market.

<p>Healthcare and life sciences</p>	<p>Healthcare and life sciences demands new computing paradigms to meet the need for personalized medicine, next-generation clinics, enhanced quality of care, and breakthroughs in biomedical research to treat disease. Now, healthcare institutions can harness the power of artificial intelligence and high-performance computing (HPC) to define the future of medicine.</p>
<p>Manufacturing</p>	<p>Harness AI and technology breakthroughs in engineering, simulation, production, remote collaboration, and visualization to revolutionize design, transform engineering, and power the factory of the future. From CAD to structural analysis to CFD, our software can fundamentally change the way you work.</p>
<p>Retail</p>	<p>From supply chains to stores—online and physical—AI helps retail run more efficiently and keeps employees and customers happier. Discover how retailers are using our solutions to power autonomous shopping, improve forecasting, automate warehouse logistics, and deliver better, more personalized customer experiences.</p>
<p>Public Sector</p>	<p>In one of the largest technology transformations ever undertaken, a revolution is happening right now amongst the government agencies and municipalities. AI is impacting new developments in cybersecurity, healthcare, life sciences, platform sustainment, humanitarian response, and disaster relief. We are at the forefront of this democratization of AI, building the technology that will make communities safer and more connected everywhere.</p>
<p>Financial</p>	<p>Intelligent technology can address critical challenges within the modern financial services industry. With our solutions—including deep learning, machine learning, and natural language processing—institutions can boost risk management, improve data-backed decisions and security, and enhance customer experiences.</p>



The Dell Technologies point of view

While AI used to be regarded as the future of technological innovation, Dell Technologies now recognizes AI as the present. Organizations that wish to remain relevant and competitive must fully adopt and embrace the benefits of AI.

Recent Gartner research found that many organizations are essentially in an experimental phase with AI.³ They know that AI holds great potential and that AI adoption is fundamental to the long-term success of the organization, but they struggle to integrate AI into many of their daily operations. Another recent study found that cultural challenges and a lack of skills are impeding the adoption of AI within many businesses.⁴

Businesses are looking for ways to handle the complexity of AI, even though they are challenged by the lack of in-house expertise. For these organizations, there is a solution developed in partnership between three companies that are each dominant in their respective fields:

- › **Dell Technologies**, industry leader in simplifying IT environments
- › **VMware**, industry leader in software server virtualization
- › **NVIDIA**, a full-stack computing company, developing CPUs, DPUs, GPUs, and AI software that fuel data center-scale computing solutions

78% of AI/machine learning projects stall before deployment — one-third of them at the proof-of-concept stage.⁵

3. "Gartner Forecasts Worldwide Artificial Intelligence Software Market to Reach \$62 Billion in 2022," Gartner, November 22, 2021

4. "As Company Culture Warms to AI Adoption, Lack of Skills and Difficulty Hiring Still Present Significant Barriers to Entry, According to New O'Reilly Research," Business Wire, April 19, 2021

5. "Almost 80% of AI and ML Projects Have Stalled, Survey Says," Robotics Business Review, May 2019.

Overcoming the challenges of AI adoption

Several key challenges regarding the implementation of an AI solution can cause a business's digital transformation to lose momentum. Identification and remediation are key to a successful AI strategy.

Together, Dell Technologies, VMware, and NVIDIA deliver simplified and high performance AI-ready solutions backed by partners with deep AI expertise.



Complexity

The implementation of an AI solution can take weeks to months to architect and implement depending on the sophistication, scope and technical depth of the solution. And, once assembled, ongoing maintenance of the environment can be a complex and time-consuming task.

Simplified AI

Jointly engineered and validated solutions make it quick and easy to deploy a hardware and software stack optimized to accelerate AI initiatives.

Dell Technologies Validated Designs for AI enable **18-20% faster configuration and integration**, saving 12 employee hours a week with automated reconciliation feeds.



Speed

Increased data acquisition to support business insights requires the adoption of new AI solutions. As these solutions evolve, organizations are challenged to keep pace and ensure alignment with business strategy.

High Performance

High performance improves AI model training accuracy with fast access to larger data sets, enabling live inferencing at scale to drive real-time, actionable responses.

Dell Technologies Validated Designs for AI enable **18X faster AI models**—two to three days instead of one to two months.



Lack of expertise

78% of AI/ML projects stall before deployment — one-third of them at the PoC stage due to lack of effective resources or dedicated technical experts to help facilitate an AI implementation.⁶

AI expertise

Deploying an engineering-tested validated AI solution backed by world-class services and support helps you confidently plan, implement and optimize AI solutions anywhere you are in your journey.

Customers can **deploy with confidence** knowing that mission-critical AI projects will stay on track with world class support.

6. Alegion White paper: What Data Scientists Tell Us About AI Model Training Today

7. Robotics Business Review, Almost 80% of AI and ML Projects Have Stalled, Survey Says, May 2019.

Easy for IT

Run AI workloads alongside existing applications.

Many organizations already use VMware vSphere for server virtualization and NVIDIA GPUs for workload acceleration. However, as data grows, data centers are starting to hit performance and scale limitations — especially when it comes to AI.

Building separate systems for AI workloads — with different hardware, software and toolsets — leads to management complexity, poor utilization rates and lengthy provisioning times. Data scientists might have to wait weeks or even months for IT to provision the resources required for AI development.

Dell Validated Designs for AI are jointly engineered and validated to power AI workloads alongside existing applications in VMware-virtualized data centers with near-bare-metal performance.

Disaggregating resources at a more granular level optimizes hardware investments — including server GPUs — increasing utilization and creating an incredibly flexible compute foundation for AI, all on infrastructure that is easy for IT to deploy and manage.

These solutions offer:

AI simplified

Jointly engineered and validated solutions make it quick and easy to deploy optimized hardware-software stacks.

Faster AI insights

Automated ML and streamlined operations delivered with the AI tools and frameworks in an optimized infrastructure enable faster time to production for development and for IT teams.

Proven AI expertise

Confidently deploy an engineering-tested AI solution backed by world-class Dell Technologies services and support. Select ProSupport Plus for a single point of contact for software and hardware support.

Validated design for AI in virtualized environments

Configurations are based on Dell Technologies NVIDIA-Certified Systems™ with NVIDIA AI Enterprise Suite, VMware vSphere virtualization with Tanzu containers.



Compute

4x Dell VxRail HCI V670
or
4x Dell PowerEdge R750xa
with
NVIDIA A100 or A30 GPUs



Networking

Dell PowerSwitch
25GbE S5248F-ON
or
an out-of-band
PowerSwitch
S4148T-ON



Storage

PowerScale F600
or
H600



Software

H2O.ai Driverless AI
cnvrg.io MLOps
NVIDIA AI Enterprise
VMware vSphere with
Tanzu

[Learn more about the Dell Validated Design for AI in Virtualized Environments >>](#)

NVIDIA technologies are built in

The VxRail hyperconverged infrastructure (HCI) or PowerEdge servers at the heart of your solution come with integrated NVIDIA technologies that help speed ML workloads — and results.

vGPUs

NVIDIA virtual GPUs (vGPUs)

NVIDIA vGPU software enables sharing GPU resources across multiple virtual machines to make them accessible to any device, anywhere.

MIGs

NVIDIA multi-instance GPUs (MIGs)

NVIDIA MIGs expand the performance and value of GPUs by partitioning them into as many as seven instances to support every workload and extend accelerated resources to more users.

A100

NVIDIA A100 GPUs

Accelerate AI workloads with up to 20X higher performance over the prior generation.⁸ The A100 supports NVLink bridge, the world's first high-speed GPU interconnect offering a significantly faster alternative for multi-GPU systems than traditional PCIe-based solutions.⁹

NVIDIA technologies are built in

The VxRail hyperconverged infrastructure (HCI) or PowerEdge servers at the heart of your solution come with integrated NVIDIA technologies that help speed ML workloads — and results.

AI

NVIDIA AI Enterprise

NVIDIA AI Enterprise is an end-to-end, cloud-native suite of AI and data science tools and frameworks optimized and exclusively certified by NVIDIA to run on VMware vSphere with NVIDIA-Certified Systems. It includes key technologies for the rapid deployment, management and scaling of AI workloads.

 NVIDIA
CERTIFIED

NVIDIA-Certified Systems

As NVIDIA-Certified Systems, VxRail HCI and PowerEdge bring together NVIDIA GPUs, NVIDIA ConnectX® smart network interface cards (SmartNICs) and NVIDIA BlueField DPUs in optimized configurations. These are validated for performance, manageability, security and scalability and are backed by enterprise-grade support from NVIDIA and Dell Technologies.



Power your AI journey with PowerScale

Dell PowerScale delivers high-performance machine learning and deep learning



All-Flash Performance

Eliminate the I/O bottleneck up to 18x more bandwidth; up to 15GB/s per chassis, up to 540 GB/s per cluster



Flexibility

Bring analytics to the data with in-place analytics leveraging multi-protocol support and direct access to the Snowflake platform



Extreme Scale-Out

Scaling to 21x* more capacity than the competition with industry leading storage efficiency up to 80%



Enterprise Features

Advanced data management tools including ransomware protection

Eliminate I/O bottlenecks with Powerscale



Accelerate Innovation

Up to 2.7x* faster performance to accelerate model training cycles



Increase Model Accuracy

Deeper higher resolution data sets with access to 10s TBs up to 10s PBs per cluster



Improve Data Science Productivity

Flexible in-place analytics and pre-validated solutions for faster, low-risk deployments




Flexible Deployments

Start small and independently scale-out compute and storage for large scale deployments

*Compared to competition


Accelerate GPU applications with RDMA data access on PowerScale



New NFS over RDMA for faster, secure and low latency data access

Lower latency and lower CPU utilization on both client and storage

Compatible with NVIDIA GPUDirect Storage protocol



Linear performance scaling up to 252 scale-out NAS nodes on PowerScale

PowerScale deployment options for AI

From edge to core to cloud



AI SANDBOX

Storage starter bundle for Data science workstation



SCALE OUT

Validated designs and reference architectures of appliances for large scale workloads



AI ON CLOUD

Support for distributed multi-cloud AI/ML workflows including direct access to the Snowflake platform.

Learn more

Visit: Dell.com/AI

Contact: ai.assist@dell.com

Get a feel for enterprise AI on Dell

NVIDIA LaunchPad is a free program that provides short-term access to a large catalog of hands-on labs so you can experience end-to-end solution workflows in the areas of AI, data science, 3D design collaboration and simulation, and more.

[Try It Now](#)

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell and other trademarks are trademarks of Dell Inc. or its subsidiaries. NVIDIA®, the NVIDIA® logo, AI Enterprise Suite™, NVIDIA-Certified Systems™, Spectrum®, BlueField® and ConnectX® are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. VMware® is a registered trademark or trademark of VMware, Inc. in the United States and other jurisdictions. H2O® is a trademark of H2O.ai. Other trademarks may be the property of their respective owners. Published in the USA 07/22 eBook virtualized-ai-design-EB-101

Dell Technologies believes the information in this document is accurate as of its publication date. The information is subject to change without notice.