

# Acceleration-Optimized servers and accelerator portfolio



**R**edefine data visualization and insights with AI





# Accelerate insight and innovation

For the digital enterprise, success hinges on leveraging big, fast data. But as data sets grow, traditional data centers are starting to hit performance and scale limitations — especially when it comes to ingesting and querying real-time data sources.

While some have long taken advantage of accelerators for speeding visualization, modeling and simulation, today, more mainstream applications than ever before can leverage accelerators to boost insight and innovation with generative AI models customized to deliver precise results with enterprise business data. Accelerators such as graphics processing units (GPUs), complement and accelerate CPUs, using parallel processing to crunch large volumes of data faster. Accelerated data centers can also deliver better economics, providing breakthrough performance with fewer servers, resulting in faster insights and lower costs.

Organizations in multiple industries are adopting server accelerators to outpace the competition — honing product and service offerings with data-gleaned insights, enhancing productivity with better application performance, optimizing operations with fast and powerful analytics, and shortening time to market by doing it all faster than ever before.


Dell Technologies offers a choice of server accelerators in Dell PowerEdge servers, so you can turbo-charge your applications.

over

# 77%

of organizations are either exploring potential use cases or investing significantly in GenAI technologies<sup>1</sup>


<sup>1</sup> IDC Future Enterprise Resiliency and Spending Survey, Wave 6, July 2023





# Emerging and traditional use cases for AI

- **Generative AI and large language models** – Accelerators are powering generative AI transformers and language processing technologies which can enable more intelligent systems with a richer understanding of language than ever before. These tools can now combine natural language processing, computer vision, and audio analysis to accept complex queries and deliver multimodal results.
- **Large-scale recommendation engines** – Accelerators excel in powering deep learning models to continuously improve advertising and search recommendations, both on relevance and timeliness, from advertisers to reach their audience and affect ad ranking models, for example.
- **Natural Language Processing (NLP)** – accelerators help boost, via machine learning, the programming of systems to process and analyze language data from spoken to written. A model can then accurately extract information and insights as well as learn new natural language tasks including language modeling, parsing, summarizing and other syntactic/semantic analysis methods, across global languages.
- **Digital twins** – these are virtual representations of objects, systems, processes, updated from real-time data and using simulation, machine learning and reasoning to drive decision-making. Digital twins are synchronized to real-world systems and data to help organization simulate, optimize products, people, equipment, and processes in real-time before ever going to production.
- **Machine and deep learning** – Accelerators have taken AI from theory to mainstream by enabling the parallel processing power required to speed both training and inferencing workloads.
- **Accelerated databases** – Accelerators can help speed aggregations, sorts and grouping operations to solve complex analytics operations that overload traditional databases.
- **Streaming data** – The Internet of Things (IoT) has created a firehose of data. Accelerators enable simultaneous ingestion, exploration and visualization of streaming data for real-time analysis.
- **Visualization** – Accelerators enhance performance for 3D visualization applications such as computer-aided design, enabling software to draw models in real time as the user moves them.
- **Modeling and simulation** – Accelerators can provide modeling and simulation for early evaluation, fast testing of design modifications enabling more iterations.
- **Financial modeling** – Accelerated HPC and artificial intelligence (AI) solutions are revolutionizing analytics tools, enabling the industry to leverage massive data sets to better understand risk and return.
- **Seismic processing** – Oil & Gas companies are finding new and better ways to extract information from massive seismic data stores, leveraging accelerators to speed time to results and shave costs.
- **Signal processing** – Accelerators enable providers to model and analyze signal data streams coming in from computers, radios, videos and cell phones in real-time.



# Leveraging Innovation and accelerated architectures

As the prior use cases suggest, the continued adoption of AI, ML, HPC workloads and VDI is adding complexity to data center and business operations, as workforce grows globally and remotely, as well as demanding use cases becoming more mainstream. For example, Artificial Intelligence has generated a wide range of new and hyper-tailored solutions for customers. Companies now leverage AI to automate many business processes, shifting human resources from one business unit to other areas for value creation.

Choosing GPUs and other accelerated architectures and products is a key decision IT teams have in their hands. And once that decision is made, for the appropriate workloads, then infrastructure strategy and product choices are addressed.

## Accelerated Insights – the leading edge of innovation from PowerEdge Servers

To design an infrastructure to deliver the capabilities which can make organizations successful with AI and other demanding workloads, requires a modern architecture approach where one of the biggest innovations is improved performance with the addition of dense acceleration, at scale. Improved performance is not only about implementing complete solution and infrastructure strategy, but also starts with innovations in the building blocks to also help provide other benefits, including improved costs, security, and thermal/power design.

There are a number of innovations within the PowerEdge server family which enable drastic performance improvements. From architectures specifically designed to support acceleration to thermally optimized designs, today's workloads demand higher quality components and subsystems to flawlessly drive workload operations.

The PowerEdge Adaptive Compute approach enables servers engineered to optimize the latest technology advances for predictable profitable outcomes. Here are a few of the improvements in the PowerEdge portfolio:

- **Focus on Acceleration** – Support for the most complete portfolio of GPUs, delivering maximum performance for HPC modeling & simulation, generative AI/ML/DL training and inferencing, analytics and rich-collaboration application suites and workloads
- **Thoughtful Thermal Design** – New thermal solutions and designs to address dense heat-producing components, and in some cases, front-to-back air-cooled designs
- **Dell Multi Vector Cooling** – Streamlined, advanced thermal design for airflow pathways within the server
- **Dell Direct Liquid Cooling** – Extending liquid cooling support across more PowerEdge servers and their CPUs for exceptional heat removal capability

Dell PowerEdge XE9680  
delivers the industry's best  
AI performance<sup>2</sup>

<sup>2</sup> Based on Dell analysis of publicly available performance results and specifications of comparable OEM Servers as of 17 May 2023.

## Accelerated AI Insights

Engineered to optimize the latest technology advances for predictable profitable outcomes



# PowerEdge servers for accelerated workloads



## No-compromise accelerated AI

XE9680\* is designed to drive business insights in the most demanding Deep Learning and modeling applications, from large natural language processing models and recommendation engines to complex research and academia problems.

- Highest performance for HPC and Enterprise
- 8x AMD Instinct MI300X GPUs with Infinity Fabric or 8x NVIDIA H100 Tensor core GPUs with NVLink
- Air-cooled operation

**Ideal workloads:** Generative AI, Large Language Models, Natural Language Processing, large recommendation engine training, molecular dynamics, genomic sequencing modeling and simulation.

### Applicable GPUs:

AMD MI300X OAM or NVIDIA H100 SXM



## Dense acceleration

XE9640\* boosts insights from your growing data sets with AI acceleration technology designed for optimal performance, fastest time-to-value, in a liquid-cooled environment.


- Mainstream 2U form factor enables highest GPU density per rack AI operations
- 4x Intel Data Center GPU Max 1550 with XeLink or 4x NVIDIA H100 Tensor core GPUs with NVLink
- Liquid-cooled CPU and GPU operation

**Ideal workloads:** Natural Language Processing, large recommendation engine training, modeling & simulation, Artificial Intelligence and ML/DL training for object recognition

**Applicable GPUs:** Intel GPU Max 1550 OAM or NVIDIA H100 SXM GPUs



# PowerEdge servers for accelerated workloads



## Purpose-built performance

XE8640\* helps businesses unlock insights with purpose-built performance in a dense air-cooled server for AI, removing traditional computational boundaries of real-time insights.

- Optimized balance of performance for diverse applications
- 4x NVIDIA H100 Tensor core GPUs with NVLink
- Air-cooled operation with liquid-assisted GPU cooling radiator

**Ideal workloads:** Medium data set language Models, Natural Language Processing, modeling & simulation, Artificial Intelligence, ML/DL training and inferencing, image recognition

**Applicable GPUs:** NVIDIA H100 SXM



## Purpose-built scale up server for GPU applications

R760xa\* maximizes results from AI to Modeling & Simulation applications with maximum flexibility and the latest 4<sup>th</sup> or 5<sup>th</sup> Generation Intel® Xeon® Scalable Processors.

R760xa is optimized to tackle GPU workloads and deliver outstanding performance for demanding and emerging applications.

- Maximize performance with Dell's broadest selection of PCIe GPU configurations
- Front-to-back air-cooled design
- R760xa supports up to 12 Single-wide GPUs or 4 Double-wide GPUs, up to 350W  
Supports all GPU cards

**Ideal workloads:** AI & ML training and inferencing, data analytics, HPC, VDI & Performance graphics

**Applicable PCIe GPUs:**

AMD MI210

Intel GPU Max 1100, Flex 140

NVIDIA H100, L40S, L40, L4, A40, A30, A16, A2, L4

# Accelerated GPU servers, at-a-glance

Model	Workloads	Memory	Processor	Storage	Accelerators	Details
PowerEdge XE9680	AI ML DL Training, HPC, CRISP, Healthcare, CSP/HPCaaS, Finance, Academia	32 (4TB)	Two 4 <sup>th</sup> or 5 <sup>th</sup> Generation Intel® Xeon® Scalable processors	8 x 2.5" or 16 x E3.S	8 x 700W SXM or 8 x 750W OAM	<a href="#">Family page</a> <a href="#">Product Video</a> <a href="#">Specification Sheet</a> <a href="#">Technical Guide</a>
PowerEdge XE9640	AI ML DL Training, HPC, Modeling & Simulation, Healthcare, Life Sciences, Finance	32 (4TB)	Two 4 <sup>th</sup> or 5 <sup>th</sup> Generation Intel® Xeon® Scalable processors	4 x 2.5"	4 x 700W SXM or 4 x 600W OAM	<a href="#">Family page</a> <a href="#">Product Video</a> <a href="#">Specification Sheet</a> <a href="#">Technical Guide</a>
PowerEdge XE8640	AI ML DL Training, HPC, Oil & Gas, Healthcare, Life Sciences, Finance	32 (4TB)	Two 4 <sup>th</sup> or 5 <sup>th</sup> Generation Intel® Xeon® Scalable processors	8 x 2.5"	4 x 700W SXM	<a href="#">Family page</a> <a href="#">Product Video</a> <a href="#">Specification Sheet</a> <a href="#">Technical Guide</a>
PowerEdge R760xa	AI-ML/DL training and inferencing, HPC, render farms and virtualization	32 (4TB)	Two 4 <sup>th</sup> or 5 <sup>th</sup> Generation Intel® Xeon® Scalable processors	8 x 2.5" or 6 x 2.5" NVMe 6 x E3.S	4 x 350W DW or 12 x 75W SW	<a href="#">Family Page</a> <a href="#">Product Video</a> <a href="#">Specification Sheet</a> <a href="#">Technical Guide</a>

## PCIe GPUs, DPUs for Dell PowerEdge servers

Turbo-charge your applications with performance accelerators available in select Dell PowerEdge tower and rack servers. The number and type of accelerators that fit in PowerEdge servers is based on the physical dimensions of the PCIe cards.

Double-wide (DW) accelerators take up two slots and include: AMD MI210 and Intel Flex 140, GPU Max 1100, and NVIDIA H100, L40S, L40, A40, A30 and A16 GPUs and, Single-wide (SW) accelerators, including the NVIDIA L4 and A2, take up one PCIe slot. Dell PowerEdge engineering qualifies accelerators with servers based on demand. Dell Technologies also works with a wide range of partners to create and sell specific combinations for particular vertical market applications.

GPUs vary in number of cores, amount of memory, and power and cooling requirements. For example, the NVIDIA Hopper® H100 has up to 80GB memory, and uses up to 700 watts.

### GPUs

Graphics processing units (GPUs) are co-processors designed to accelerate compute performance. A GPU typically has thousands of cores designed for efficient execution of mathematical functions. Portions of a workload are offloaded from the CPU to the GPU, while the remainder of the code runs on the CPU, improving overall application performance.

Dell offers a range of GPUs as PCIe cards that fit into server PCIe slots. The Dell PowerEdge XE product family offers support for 4x or 8x GPU assemblies on Open Compute Project Accelerator Module (OAM) or NVIDIA SXM modules mounted on the server motherboard.

### DPUs

A Data Processing Unit (DPU) combines computing, networking, and programmability to offload CPUs and deliver software-defined, hardware-accelerated solutions for the most demanding workloads.

**Parallel processing**  
Parallel processing is a method of simultaneously breaking up and running program tasks on multiple microprocessors, reducing processing time.

**Optimize the code**  
To take full advantage of server accelerators, optimize the software code. For many applications, four lines of code can provide a boost.

## NVIDIA Hopper and Ampere and Tensor Core GPUs



NVIDIA Hopper and Ampere Core GPUs deliver the horsepower needed to run deep learning training, high performance data analytics, visualization and other workloads faster than ever before. Plus, NVIDIA GPUs deliver high performance and user density for virtual desktop infrastructure (VDI). Deliver mainstream AI on with NVIDIA AI Enterprise.

- [Hopper core GPU](#)
- [Ampere core GPU](#)
- [NVLink™ Fabric interconnect](#)
- [GPU CLOUD™](#) containers
- [Software application catalog](#) and [developer resources](#)
- [NVIDIA AI Enterprise](#)

Model	Workloads	Memory	Graphic Bus/ System interface	Slot width	Max Power Consumption	Server support
H100	HPC/AI/Database Analytics	80 GB HBM2e	PCIe Gen5 x16/ NVLink bridge <sup>3</sup>	Double-wide	300-350W	R750xa, R750, R7525
H100	HPC/AI/ML/DL Training	80 GB HBM3	NVLink bridge	N/A	700W	XE9680 (8xH100), XE8640 (4xH100)
L40	Performance graphics/VDI	48 GB GDDR6	PCIe Gen4 x16	Double-wide	300W	R750, R750xa, R7525
A40	Performance graphics/VDI	48 GB GDDR6	PCIe Gen4 x16/ NVLink bridge <sup>3</sup>	Double-wide	300W	R750xa, R750, R7525, XR12, DSS8440, R740, R740xd, T550
A30	mainstream AI	24 GB HBM2	PCIe Gen4 x16/ NVLink bridge <sup>3</sup>	Double-wide	165W	R750xa, R750, R7525, R7515, R740, R740xd, XR12, XE2420, T550
A16	VDI	64 GB GDDR6	PCIe Gen4 x16	Double-wide	250W	R750xa, R750, R7525, R7515, R740, R740xd
A10	mainstream graphics/VDI	24 GB GDDR6	PCIe Gen4 x16	Single-wide	150W	R750xa, R750, R7525, R740, R740xd, XE2420
L4	Inferencing/Edge/ VDI	24 GB GDDR6	PCIe Gen4 x16	Single-wide	72W	R750, R7525, R650
A2	Inferencing/Edge/ VDI	16 GB GDDR6	PCIe Gen4 x8	Single-wide	60W	R750xa, R750, R7525, R7515, R650, C6520, R6525, R6515, C6525, XR12, XR11, R740, R740xd, R640, T550
T4	Inferencing/Edge/ VDI	16 GB GDDR6	PCIe Gen3 x16	Single-wide	70W	R750xa, R750, R7525, R7515, R650, C6520, R6525, R6515, C6525, XR12, XR11, DSS8440, R740, R740xd, R640, XR2, XE2420, XE7100

<sup>3</sup> 8A100 w/Nvlink bridge is supported on R750XA and DSS8440; A40 w/Nvlink bridge is supported on R750XA, DSS8440 and T550; A30 w/NVLink bridge is supported on R750XA, DSS8440 and T550; ; MI210 w/Infinity Fabric Link bridge is supported on R750XA; H100 and A800 w/Nvlink bridge will be supported on R750XA; Max1100 w/ XeLink bridge is supported on R760XA

DW - Double Wide, SW - Single Wide, FH- Full Height, FL - Full Length, HH - Half Height, HL - Half Length





NVIDIA-Certified Dell Systems brings together NVIDIA GPUs and NVIDIA networking in servers and hyperconverged infrastructure from Dell Technologies in optimized configurations.


These systems are validated for performance, manageability, security, and scalability and are backed by enterprise-grade support from NVIDIA and Dell Technologies.

- Deliver infrastructure to drive a diverse range of accelerated workloads for the enterprise
- Excellent performance
- Reduce time to deployment
- Secured, no-compromise operations and workflows
- Designed for single to multi-node configs, optimal Scale-out and clusters

Learn more about Dell PowerEdge servers with NVIDIA-Certified solutions [here](#).

Consult our [matrix of supported PowerEdge servers and partner accelerators](#) to deliver the optimal configuration for your applications and workloads.

## Intel GPUs



The Intel® Data Center GPU Max Series is designed to take on the most challenging high-performance computing (HPC) and AI workloads.


Available on Dell XE9640 servers.

Unleash the Power of Intel Data Center GPU Max Series through software: For the data center GPU, Intel oneAPI and AI tools help you realize maximum performance from the innovative hardware's advanced capabilities like Intel® Xe Matrix Extensions (Intel® XMV), vector engine, Intel® Xe Link, data type flexibility, and more.

- [Intel® Data Center GPU Max Series Overview](#)
- [Explore the Max 1550](#)
- [Developer tools for Intel Data Center GPU Max Series](#)

Model	Workloads	Memory	Graphic Bus/ System interface	Slot width	Max Power Consumption	Server support
Max 1550	HPC, AI/DL training	128 GB	Intel Xe Link	N/A	600W	XE9640
Flex 140	Inferencing/Edge	12 GB GDDR6	PCIe Gen4 x8	SW	300W	R760xa, R760, R660
Max 1100	AI / HPC	48 GB HBM2e	PCIe Gen5x16/ XeLink bridge <sup>4</sup>	DW	75W	R760xa, R760, XR7620

## AMD GPUs



Built on CDNA architecture, AMD MI210 delivers industry best single-precision (FP32) performance. The AMD Instinct GPU family accelerates HPC, AI workloads, and reduces the overall cost of ownership. Now available on Dell PowerEdge R750xa and PowerEdge R7525 servers.

- [Explore MI210 Accelerator](#)
- [Watch Dell PowerEdge XE9680 with AMD Instinct MI300X product video](#)
- [Read top 5 reasons to choose Dell PowerEdge XE9680 with AMD Instinct MI300X Infographic](#)
- [Read Dell PowerEdge XE9680 with AMD Instinct MI300X and porting LLMs to AMD ROCm Infographic](#)
- [Read Dell PowerEdge XE9680 with AMD Instinct MI300X and ROCm eBook](#)
- [Learn how the ROCm™ open software platform enables HPC GPU computing](#)

Model	Workloads	Memory	Graphic Bus/ System interface	Slot width	Max Power Consumption	Server support
MI300X	AI / HPC	192 GB HBM3	AMD Infinity Fabric Links	N/A	750W	R760xa, R760, XR7620

# Dell AI Solutions

Save time with Dell Technologies and partner solutions with accelerators inside.



## Dell PowerEdge Server – Accelerator Combinations

The number and type of accelerators that fit in [PowerEdge servers](#) is based on the number and type of PCIe slots in the server chassis and the accelerator form factor (FF), or the physical dimensions of the PCIe cards.

Solution	Description	Resources
AI for Virtualized Environments	VMware-enabled AI with NVIDIA AI Enterprise on Dell infrastructure.	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>
AI MLOps with cnvrg.io	Standardize machine learning pipelines with cnvrg.io to minimize friction for data science and engineering teams from research to production.	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>
Automatic Machine Learning	Automate algorithm selection, feature generation, hyperparameter tuning, and model assessment to ease and speed time to AI.	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>
Conversational AI	Deliver extraordinary, effective, and efficient AI-enabled customer and employee experiences on voice and digital channels (including chatbots and virtual assistants)	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>
Generative AI with NVIDIA for Inferencing	Quickly get up and running with a pre-trained model and start producing outputs and value with a joint architecture from Dell Technologies and NVIDIA.	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>
Generative AI with NVIDIA for Model Customization and Tuning	Learn how to re-train an existing GenAI model for your own use cases, with examples of standard customization techniques such as transfer learning and prompt tuning.	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>
Red Hat OpenShift AI on APEX Cloud Platform for Red Hat OpenShift	Implement a digital assistant by leveraging a Large Language Model (LLM) and the Retrieval Augmented Generation (RAG) framework.	<ul style="list-style-type: none"><li>• <a href="#">Solution Brief</a></li><li>• <a href="#">Validated Design</a></li></ul>

## Rest easier from day one with our comprehensive services

Utilize Dell Technologies Services to maximize the life and value of your PowerEdge Servers on a global scale, across 170 locations and benefit from the deep expertise of our 60K+ employees and partners.

- [ProDeploy Factory Configuration](#) – Factory-based services deliver PowerEdge servers configured to your specifications, ready to install
- [ProDeploy Rack Integration](#) – Receive PowerEdge fully configured and racked direct from our facility with optional onsite final configuration
- [ProDeploy or ProDeploy Plus](#) – ProDeploy experts are here to help, with 24/7 field-based deployment services, from planning through implementation and beyond. Choose from guided remote to fully onsite hardware and software implementation
- [Data Migration Services](#) – Efficiently move data from where it is to where it will drive innovation
- [ProSupport Plus for Enterprise](#) – Leverage predictive issue detection and proactively improve the performance of your critical systems, while taking advantage of an assigned Service Account Manager

Availability and terms of services vary by region. For more information and details on our entire range of offerings, please contact your Dell Technologies representative or visit us online at [Dell.com/services](https://Dell.com/services).

## Solutions available with Dell Technologies partners

[Amulet Hotkey® virtual desktop solutions](#) combine enterprise-class servers with virtual GPU accelerators to deliver high-density, data center–optimized solutions to simplify the transition to Windows® 10. In addition, virtual GPUs help address the growing demand for graphics-accelerated virtualization of everyday programs like Windows 10, Microsoft® Office 365®, YouTube® and more for an exceptional virtual desktop experience. [Read about Amulet Hotkey customer successes](#).

[Kinetica®](#) is an insight engine that includes a GPU-accelerated database, visual discovery and machine learning capabilities, and accelerated parallel computing. Running on Dell PowerEdge servers with NVIDIA GPUs, Kinetica helps organizations meet the challenges that come with huge quantities of complex, unpredictable data. Read the article: [Explaining GPUs to Your CEO: The Power of Productization](#).

[Tracewell Systems®](#) deliver powerful, off-the-shelf computing technology for businesses, government agencies and OEMs in places where environmental factors create unique computing challenges, such as in the air, at sea or on the ground, in fixed and mobile installations, or in situations where integration with specialty hardware or software is required.

[Get data sheets, videos and resources](#).

---

## Dell Technologies Acceleration Software partners



[VMware® BitFusion®](#) software disaggregates GPUs, FPGAs and/or ASICs and dynamically attaches them anywhere in the data center.



[NVIDIA GRID™](#) Virtual Apps improve virtual desktops and accelerate server applications, with proven performance built on NVIDIA® GPUs.



[AMD ROCm™](#) delivers an open-source exascale-class platform for accelerated computing in HPC and cluster deployments.



[Kinetica®](#) software dramatically speeds up traditional online analytics processing (OLAP) workloads using GPUs for parallel computing.



[SQream Technologies®](#) GPU-accelerated data warehouse is capable of scaling from terabytes to petabytes, adapting to any scale and workload.



[FASTDATA.io](#) PlasmaENGINE® GPU-native software enables real-time processing of infinite data in motion, over multiple nodes, with multiple GPUs.



[RAPIDS](#) is a suite of data science libraries built on NVIDIA CUDA-X for executing end-to-end data science training pipelines in NVIDIA GPUs.



## Become a Dell Technologies Partner

When you join the Dell Technologies Partner Program, you are joining a partner ecosystem that together is making digital, IT, workforce, and security transformation real to organizations across the globe - every single day. Underpinning the industry's most robust portfolio from the edge to the core to the cloud is the Dell Technologies Partner Program, designed to be Simple. Predictable. Profitable.

## Resources

Ready your data center to handle any workload with PowerEdge Servers. PowerEdge tower servers are designed to grow with your organization, at your pace. PowerEdge rack servers combine a highly scalable architecture and optimum balance of compute and memory to maximize performance across the widest range of applications. Shop Dell PowerEdge servers at [dell.com/poweredge](https://dell.com/poweredge).

Server advanced engineering provides guidance at [Support for Servers Solution Resources](#). White papers are also available at [delltechnologies.com/accelerators](https://delltechnologies.com/accelerators) > [resources](#) > [white papers](#). For reference architectures, visit [delltechnologies.com/referencearchitectures](https://delltechnologies.com/referencearchitectures).

### See performance results

Get benchmarking data by workload, reference architectures and blogs from HPC/AI engineering at [hpcatdell.com](https://hpcatdell.com) and download from [GitHub](#).

### Access Education Services

Get the skills, training and certifications you need at [education.emc.com](https://education.emc.com). [Learn how to solve problems with deep learning](#) at the Deep Learning Institute by Dell Technologies.

### Community resources

Join the Dell Technologies HPC/AI Community at [dellhpc.org](https://dellhpc.org). Connect with the AI Builders Community at [builders.intel.com/ai](https://builders.intel.com/ai).

### Visit a Dell Technologies Customer Solution Center

Experience our solutions and products with a customized engagement designed to help you address your business challenges or innovate for success. Work with our subject matter experts in our dedicated labs – stacked with the latest and greatest products and solution showcases. Remote connectivity enables you to include global team members, or work with us from your own location. Learn more at [delltechnologies.com/csc](https://delltechnologies.com/csc).



[Learn more](#) about  
Dell solutions



[Contact](#) a Dell  
Technologies Expert



[View more](#) resources



[Join the conversation with  
#HashTag](#)