

HPC for Artificial Intelligence and Data Analytics

Run AI, data analytics and advanced computing workloads on one system

Customer results

3X

computational power at TACC¹

70 petaFLOPS

of sustainable computing at ENI²

1K datapoints

streaming per second at McLaren³

The advantage in today's marketplace goes to the data-driven enterprise. For many organizations, advanced computing workloads like artificial intelligence (AI), data analytics, modeling and simulation are important sources of competitive advantage. As high performance computing (HPC) systems become smaller, simpler and less costly, enterprise IT teams are adopting HPC to provide the throughput and capacity needed for advanced computing and other enterprise workloads.

Converging HPC with AI and data analytics on a single system gives you the horsepower to run high-performance data analytics (HPDA), boost high-frequency trading, enhance risk analysis, improve fraud detection, and speed motion picture special effects cycles — just to name a few options.

However, to optimize budgets along with performance, it's critical to match advanced computing resources to business requirements. System configuration can be a complex task, requiring a balance between workload requirements, performance targets, data center constraints and pricing. Many enterprises don't have the time to research, optimize and deploy advanced computing systems to deliver the required return on investment.

With the HPC Validated Design for AI and Data Analytics, Dell Technologies engineers have done the heavy lifting, so you can quickly deploy a high performance computing solution that matches the needs of the business. Dell Technologies can help you optimize investments based on your budget, with the ability to tune solutions for specific workloads and scale as needed with modular building blocks.

Essentially, the HPC Validated Design for AI and Data Analytics can deliver the throughput and capacity to manage rapid data growth and increased workload demands, so business users can dream bigger and act faster than ever before.

One System for AI, Data Analytics and HPC

Dell Technologies Validated Designs are workload-optimized rack-level systems with servers, software, networking, storage and services to scale faster with the confidence of an engineering-tested solution while saving valuable time and resources. The HPC Validated Design for AI and Data Analytics has been optimized, tested and tuned for a variety of applications on the Kubernetes® stack, with ongoing testing and validation to expand the list of validated options. This solution runs HPC, AI and data analytics workloads on the same system, simplifying deployment and management, and keeping costs low.

¹ Compared to Lonestar5. Source: [TACC Lonestar6](#), October 2021.

² Dell Technologies case study, [Accelerating the path to sustainable energy](#), August 2020.

³ Dell Technologies case study, [Data-driven innovation starts at racing's edge to improve race car aerodynamics — and speed](#), April 2021.

Resources

- See the [Validated Design](#)
- See the [Validated Design with DKube](#)
- Performance testing and engineering documentation at hpcatdell.com

Learn more

- delltechnologies.com/hpc
- delltechnologies.com/ai
- delltechnologies.com/innovationlab
- delltechnologies.com/coe
- dellhpc.org

Flexible workload management enables dynamic movement of jobs between Slurm® and Kubernetes based on user demand, with a scalable shared filesystem to support both. Bright Cluster Manager® provides a single-pane-of-glass management experience for Dell EMC hardware, Slurm and Kubernetes.

The HPC Validated Design for AI and Data Analytics comes with a best practices guide and toolkit to help you take systems from factory-installed operating system to full Kubernetes cluster with a repository on GitHub®. With Bright Cluster Manager software, system administrators can quickly get clusters up and running and keep them running reliably throughout their lifecycle.

Components

Management node	Compute nodes	Acceleration nodes	Accelerators	Networking	Storage
PowerEdge R750 Servers	PowerEdge R750xa, R7525, C6520, C6525 Servers	<ul style="list-style-type: none"> • PowerEdge XE8545 or C4140 • R750xa or R7525 and/or • DSS8440 	<ul style="list-style-type: none"> • NVIDIA A100, A10 or T4 • AMD® Instinct™ Mi100 • Intel PAC FPGAs 	Choice of: PowerSwitch 25/100Gb Ethernet or NVIDIA® Mellanox® HDR100 InfiniBand®	<ul style="list-style-type: none"> • PowerScale • HPC NFS or BeeGFS® Storage

Because the optimum solution configuration will depend on the specific mix of applications and types of data analytics, artificial intelligence and advanced computing workloads, recommended and engineering-validated options are provided, along with relevant criteria to consider when making these selections. As always, Dell Technologies HPC and AI experts are available to assist you with designing a solution for your specific needs. And Dell Technologies [Services](#) — ranging from consulting and education to deployment and support — are available when and where you need them.

