

# Dell EMC PowerScale and NVIDIA DGX Systems

Artificial Intelligence (AI) and Deep Learning (DL) techniques have evolved tremendously over the last few years, allowing organizations to tap into the potential of their data and transforming entire industries overnight. However, success in AI initiatives isn't guaranteed. Many projects fail due to improper planning and lack of adequate IT infrastructure needed to effectively run data science processes. The development, deployment and training of data hungry algorithms is only one piece of the puzzle.

Dell Technologies and NVIDIA have worked together to deliver a high-performance architecture for DL by combining NVIDIA DGX A100 systems with NVIDIA A100 Tensor Core GPUs, NVIDIA Mellanox SN3700V and QM8700 switches, and Dell EMC Isilon F800 all-flash nodes from the industry leading Dell EMC PowerScale scale-out NAS storage portfolio. This new reference architecture extends the commitment that Dell Technologies and NVIDIA have to make AI simple and accessible to every organization with our unmatched set of joint offerings. Together, we offer our customers informed choice and flexibility in how they deploy high-performance DL at scale.

## Dell EMC PowerScale

Isilon F800



## NVIDIA DGX A100



### Solution Highlights

- Leverage modular AI architecture building blocks offering simplicity and scale from proof of concept to production
- Plan for AI innovation with an end-to-end strategy that accelerates time to insights
- Get the blueprint for independently scaling compute, storage and networking components
- Experience seamless tiering of storage between All Flash, Hybrid and Archive nodes with PowerScale SmartPools
- Grow storage capacity, throughput, IOPs, cache and CPU via simple node additions to clusters
- Run various AI processes on data in place with multi-protocol access

## Plan for data growth with PowerScale

PowerScale is made for storing the large, expanding datasets required for, and created by, AI/DL. Organizations of all types and sizes are adopting PowerScale for their burgeoning AI initiatives. Successful AI initiatives start with planning for massive data growth, considering both the data needed for neural networks training and the data created through training and inference. With PowerScale storage, your data stays in place, from proof-of-concept to production, avoiding the need for costly and time-consuming data migration.

PowerScale eliminates I/O bottlenecks allowing for:

- Faster training and validation of AI models
- Higher model accuracy
- Improved data science productivity
- Maximized ROI for compute and GPU investments

The PowerScale OneFS operating system provides the intelligence behind all Isilon scale-out storage systems. With enterprise data management and governance capabilities to eliminate data silos, you can centralize enterprise storage processes including data management, performance management, data protection and data security. This consolidation also improves the economics of data storage, lowering ownership costs and reducing risk.

## NVIDIA DGX A100 compute and GPU

The NVIDIA DGX A100 is a universal system for all types of AI workloads. It integrates eight NVIDIA A100 Tensor Core GPUs, delivering performance up to 5 petaFLOPS. It allows data science practitioners to accelerate complete workflows, from data preparation and analytics to training and inference. DGX A100 systems are optimized for AI demands.

NVIDIA DGX Systems are:

- Powered by the NVIDIA DGX Software Stack for end-to-end AI development
- Designed for fast start and effortless productivity from proof of concept to production
- Flexible AI systems that adapt to evolving needs
- Complete hardware/software platforms backed by NVIDIA DGXperts



To learn more about our offerings visit our website

[DellTechnologies.com/StorageforAI](https://DellTechnologies.com/StorageforAI)