



# NetApp E-Series High-Performance Storage Solutions for Artificial Intelligence

## BACKGROUND

Artificial intelligence (AI), machine learning (ML), and deep learning (DL) capabilities are rapidly becoming a crucial component for achieving business and scientific breakthroughs. From keeping manufacturing lines up and running by predicting repairs, to saving lives with more accurate ways to diagnose cancer, to developing autonomous vehicles, organizations across all industries are changing the world with the help of AI, ML, and DL technologies.

## CHALLENGE

To get the most value from AI, ML, and DL solutions, organizations must be able to quickly process and analyze massive amounts of data. The more data, the better the results. However, with petabytes of data to process and analyze, data management can be a major challenge.

To deliver faster, more accurate insights that help reinvent the way we work, live, and play, organizations need a cost-effective storage solution that enables fast, continuous feeding of data to the AI application. And the solution must scale seamlessly to accommodate data coming in from the Internet of Things (IOT) as well as data generated from machine learning and deep learning training.

## SOLUTION

The NetApp® E-Series high-performance computing (HPC) solution features a complete line of high-performance, highly reliable storage systems. A modular architecture with industry-leading price/performance offers a true pay-as-you-grow solution to support storage requirements for the massive amounts of data needed for AI, ML, and DL operations. Storage solutions are deployed with the leading parallel file systems, including BeeGFS, IBM Spectrum Scale, Lustre, and others to handle the performance and reliability requirements of the most demanding AI workloads.

## THE SPEED YOU NEED

To keep your AI, ML, and DL operations running smoothly, your storage must be able to keep pace with your compute power.

The high-performance NetApp E-Series storage solution delivers top performance in public SPC-1 and SPC-2 benchmarks. It can process up to 1 million random-read IOPS at less than 250 microseconds and 14GBps sustained (maximum burst) write bandwidth per scalable building block. Optimized for both flash and spinning media, the solution includes built-in technology that monitors workloads and dynamically adjusts system parameters to maximize performance.

## AROUND-THE-CLOCK RELIABILITY

Because organizations have become so dependent on their AI environments, downtime of any sort is unacceptable. Processing and analyzing data for AI, ML, and DL uses can take months to complete. And downtime during that run can mean starting over, greatly delaying time to discovery and possibly compromising the integrity of the data. In addition to negatively impacting the business, downtime can put the safety—in some cases even the lives—of consumers at risk.

The NetApp E-Series HPC solution offers nonstop reliability with a fault-tolerant design that delivers greater than 99.9999% availability. With industry-leading redundancy, the solution provides outstanding resilience, proven by more than 1 million systems deployed. Built-in data assurance features protect data accuracy by avoiding drops, corruption, and missed bits, from host to storage media and back.

Our extensive partner ecosystem helps us validate configurability, stability, interoperability, and reliability. You can be certain that your NetApp E-Series HPC solution will deliver the 24/7 availability your AI, ML, and DL operations require.

## EASY DEPLOYMENT AND SUPPORT

The NetApp E-Series HPC solution makes enterprise storage easy. As systems are deployed, the solution is easy to install and manage as a single unit or as hundreds of units. The modular design enables your IT staff to nondisruptively add performance and capacity without complex deployments or migrations. Scripting allows dynamic replication, so you can spontaneously configure new systems for faster deployment and automate common tasks for easier management.

The NetApp enterprise architecture also features proactive monitoring and support to automate issue resolution and reduce management overhead. Worldwide support and 4-hour parts delivery help keep operations flowing. NetApp also offers secure support options for sensitive government sites.

## LIMITLESS SCALABILITY

The modular design of the NetApp E-Series HPC solution offers a granular, building-block approach to growth. You can scale seamlessly from terabytes to petabytes by adding capacity in any increment—one drive or multiple drives at a time.

## AI, ML, AND DL USE CASES

### Laboratories

- Discover cures for diseases like cancer and diabetes
- Forecast weather and natural disasters more accurately
- Improve global ocean and fisheries management

### Manufacturing

- Improve product quality with consistent checks
- Use robots to increase efficiency with production and distribution
- Enhance customer service and improve product support

### Healthcare

- Improve the speed and accuracy of patient diagnosis
- Reduce surgical complications and long hospital stays with robotic-assisted surgery
- Save time and money with 24/7 virtual nursing assistants

### Energy

- Forecast supply and demand
- Optimize yield from oil and gas reserves
- Identify new reserves or sources of energy

## AI, ML, AND DL USE CASES

### Financial Services

- Predict market trends
- Execute trades at high speeds and high volumes
- Detect fraud

### Media and Entertainment

- Create and deliver more personalized content
- Put relevant content at the fingertips of video production staff and TV and radio hosts
- Deliver innovative experiences through virtual reality

### Government

- Increase cybersecurity by forecasting, preventing, and mitigating risks
- Improve public safety with more accurate security screening at airports
- Reduce case-load backlogs by filtering requests and applications

### Retail

- Predict purchasing patterns
- Enable more personalized marketing
- Increase sales with personalized coupons and other discounts

### LOWER TCO

One constant in AI, ML, and DL computing is that datasets continue to grow. Without forethought, storage costs can easily spiral out of control. With price/performance-optimized building blocks from NetApp, configurations of all sizes, from small to large, are cost efficient. In addition, NetApp HDDs and SSDs are required to meet strict quality, performance, and interoperability requirements before acceptance. This extra scrutiny results in failure rates that are four times lower than those of commodity HDD and SSD devices. Also, NetApp employs the industry's best performance density per rack unit, resulting in lower power, cooling, and support costs. Finally, you enjoy the flexibility to support Remote Direct Memory Access (RDMA) through 100Gb iSCSI Extension over RDMA (iSER), 100Gb SCSI RDMA Protocol (SRP), 100Gb nonvolatile memory express over InfiniBand (NVMe/IB), 100 Gb NVMe over RDMA over Converged Ethernet (NVMe/RoCE), 32Gb Fibre Channel (FC), 12Gb Serial Attached SCSI (SAS), and 25 Gb iSCSI connectivity. All of these capabilities add up to significantly lower TCO.

### MORE INFORMATION

To learn more about NetApp high-performance computing solutions for AI, ML, and DL, visit [netapp.com/hpc](https://netapp.com/hpc).

### ABOUT NETAPP

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit [www.netapp.com](https://www.netapp.com). #DataDriven