

SOLUTION BRIEF

Meet AIRI//S AI Infrastructure For Modern Innovators

Next-gen AI infrastructure that accelerates end-to-end GPU workflows.

In recent years, artificial intelligence (AI) has grown dramatically, becoming one of the fastest and largest areas of enterprise technology investment and innovation. There are so many practical applications for this technology, it's no surprise that AI is now supporting mainstream use cases, ranging from healthcare and life sciences to semiconductor and chip manufacturing, automotive, financial services, and beyond.

“Do-It-Yourself” Causes Roadblocks

The IT infrastructure historically used to support legacy data science efforts can't provide data access fast enough to unleash the advantage of GPU-powered AI-based data science. Organizations are looking for solutions to enable AI simply, quickly, and cost effectively.

When implementing home-grown AI solutions, organizations contend with a variety of roadblocks, including:

- **Unexpected hardware expenses** due to the high-performance, dedicated, and optimized hardware required to successfully process each stage of the AI pipeline and store massive amounts of data. Legacy servers are not up to the task.
- **Complex AI software** requiring careful, precise, and frequent configuration by highly specialized developers.
- **Endless maintenance cycles** demanding constant software patching, tuning, and compiling, along with ongoing tweaking of the AI algorithms and frequent hardware upgrades to keep up with the unending deluge of new data.



AIRI//S Technology Stack:

- Pure Storage FlashBlade//S
- NVIDIA DGX Systems
- NVIDIA Spectrum and Quantum¹ Networking
- Software includes: DGX software stack, Purity//FB OE, Pure1[®] Management, Pure RapidFile Toolkit.



AIRI//S Simplifies AI-at-scale

- Reduces racks of complexity into a complete solution
- Data scientists can focus on algorithms, not infrastructure



AI-at-scale Is an Advantage

- More compute = faster training
- More data = higher accuracy
- AIRI makes it simpler to run multi-node training



Uncomplicate Your Data

Fortunately, there is a much simpler path to AI enlightenment. AIRI//S™ is the latest evolution of the industry's first complete AI-ready infrastructure. Architected by Pure Storage® and NVIDIA, and powered by the latest NVIDIA DGX systems and the new Pure Storage FlashBlade//S™, AIRI//S is a cutting-edge, disaggregated, and modular scale-out storage platform for unstructured data. The AI solution is out-of-the-box and simple-to-use, enabling you to reap the benefits of AI right away and delivering a simple, fast, and future-proof infrastructure to meet your AI demands at any scale.

Delivering the Next Generation AI Platform



Full Stack Pre-integration for AI

The goal of any AI factory is to maximize training performance. To achieve this, you need storage performance that meets the needs of the compute layer. Once this has been achieved, the next major focus should be on efficiency and density. Only AIRI//S achieves this sweet spot of performance, density, and efficiency that maximizes your GPU investment and optimizes TCO.

Ease of Use

AIRI//S can be set up, deployed, and managed quickly as an end-to-end AI pipeline solution. It ensures that your data teams remain productive—spending more time delivering insights and less on worrying about their infrastructure.

Sustainability


AIRI//S provides significant gains in data center efficiency. FlashBlade//S delivers industry leading 1.3 watts/TB which is 48% more efficient than another popular NAS solution. The DGX A100 MiG is seven times more performant than its predecessor the V100 GPU—in the same footprint.



Simplifies AI Infrastructure



Maximizes Scalability and Efficiency



Performance at Any Scale

Start your AI journey at any scale and seamlessly grow and adapt as needs evolve—all while delivering predictable performance to keep your GPU fed and data teams productive. FlashBlade//S scales to 1.92PB per single chassis and is greater than two times more performant than the previous generation FlashBlade®.

Future Ready

AIRI//S is built for future upgrades to grow with your expanding AI needs. FlashBlade//S software and hardware upgrades are always included and non-disruptive. Future support includes Bright Cluster Management software, DGX H100, and more.



AIRI//S Technology Stack



AIRI represents a combination of groundbreaking solutions from two industry leaders. All-flash storage from Pure Storage is paired with NVIDIA's latest AI hardware and software. AIRI//S is optimized out-of-the-box and can seamlessly scale storage and AI compute performance, allowing organizations to begin their AI journey at any scale and grow and adapt as their needs evolve.

Pure Storage: Enabling AI Success with NVIDIA

Pure Storage has worked with NVIDIA to help organizations realize their AI initiatives across a variety of industries, including the federal and public sector, healthcare, manufacturing, financial services, and telecommunications. The following are just a few of the areas Pure and NVIDIA have enabled AI success:

- Advancing the national security mission of the United States and other nations' intelligence agencies (private references).
- Accelerated genetic research and medical imaging 7x at the Center for AI in Medicine at [Chang Gung Memorial Hospital](#).
- Accelerated medical image processing which increased customer satisfaction and market competitiveness at [Olympus](#), a leading MedTech provider.
- Accelerated training models to advance the capabilities and safety of self-driving vehicles for [Zenuity](#).
- Advanced AI analysis of large data sets at [NavInfo](#) to produce real-time maps and navigation models used by car manufacturers around the world.
- Delivered significant expansion in the volume data science research projects for the [Gwangju Institute of Science and Technology \(GIST\)](#).
- Accelerated EDA to support the pace of chip development where for each nanometer a die is reduced in size, the volume of data required to be processed in the development almost doubles from the prior size (private references).

Learn More

AIRI//S from Pure Storage powered by NVIDIA accelerates and simplifies the process of deploying and running complex AI infrastructure, so our customers' data teams can focus on their core mission and discover critical insights and outcomes faster. [Explore the AIRI//S solution.](#)

¹ Quantum networking supported for larger DGX clusters

purestorage.com

800.379.PURE



PURESTORAGE®
Uncomplicate Data Storage, Forever