Everpure

Ready your data systems for AI scale

Move from experimentation to production with a unified data foundation.

Enterprise organizations are rapidly shifting focus from AI experimentation to operational monetization. However, executive teams often face a critical structural bottleneck: traditional data architectures, designed for legacy workloads, struggle to keep pace with the real-time velocity and scale modern AI demands. Gartner[1] predicts that over 60% of enterprise AI projects will be abandoned through 2026 without AI-ready data. While a stalled initiative may look like a model failure from the outside, the actual crack is almost always underneath in the data layer.

Beyond data silos, technical leaders have stretched existing systems to support decades of digital innovation. However, enterprise AI demands a fundamentally different architectural model. Adding more GPUs doesn't fix a throughput ceiling. It just leaves expensive compute sitting idle while capital spend climbs and output doesn't. Combined with fragmented governance, traditional point-solution expansions introduce manual operational drag and unpredictable timelines that ultimately stall AI velocity.

The AI scale imperative and path forward

Technical leaders are now being charged with helping their businesses create competitive advantage through AI. Bridging the gap between business ambitions and operational execution requires navigating structural hurdles across resources, data readiness, and governance.



Massive scale

Supports 10,000+ GPUs to power exabyte-scale workloads



Unmatched speed

Eliminates data movement delays and integration complexity while outperforming closest competitors in MLPerf Storage 2.0 and SPECstorage testing.



Decreased TCO

58% lower TCO, including 25% lower cost of power and cooling and 75% lower cost of administration, compared to competing solutions

To bridge this execution gap and transform unstructured data into a governed, high-velocity asset, organizations can modernize their enterprise data strategy by transitioning from fragmented storage to a unified platform built on three core pillars.

1 Unlocking AI-ready data for faster time to value

Dark data, fragmented silos, and regional residency rules keep generative AI models from accessing live operational context. Further, AI teams are uncertain if their models are relying on current, trustworthy information. Accelerating time to value requires infrastructure that helps teams effortlessly find, contextualize, and validate the right data at the source. By delivering real-time readiness and automated context directly at the data layer, organizations can improve response accuracy, eliminate redundant processing, and significantly reduce large language model (LLM) token costs.

2 Maximizing AI performance for scale and ROI

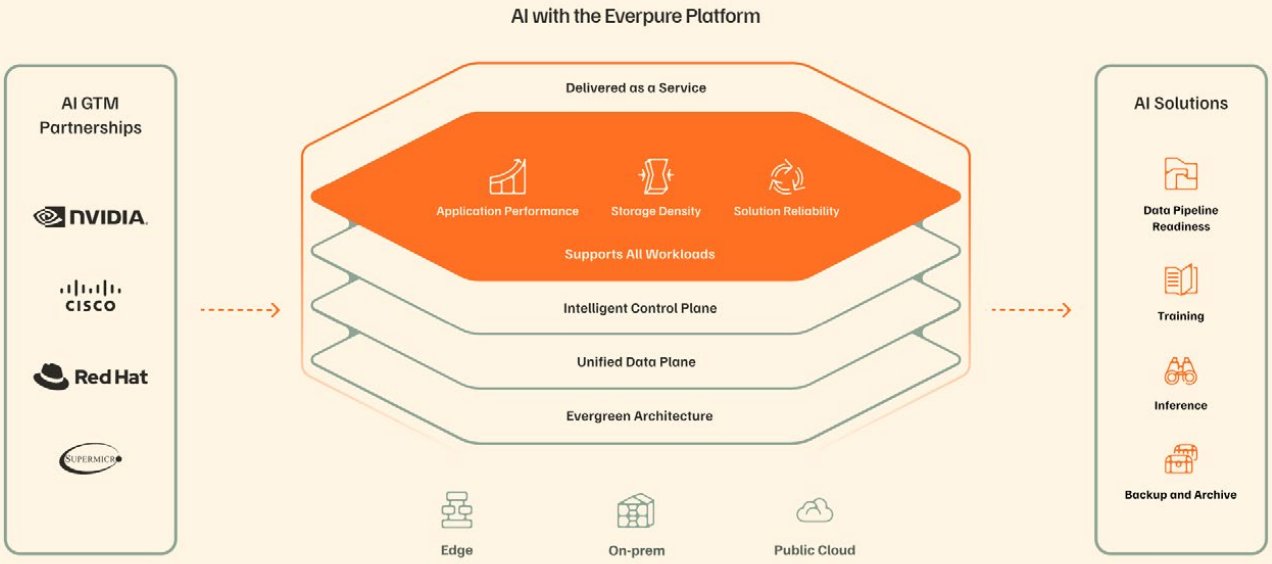
Millions of dollars in high-performance GPU compute sit idle waiting for storage I/O, eroding AI investment returns. To shorten the time from AI pilot to production, storage infrastructure must evolve to match the massive processing power and throughput of modern GPU clusters. Aligning storage performance directly with GPU capacity enables organizations to keep compute clusters continuously saturated at scale, maximizing training efficiency and model throughput.

3 Automating to increase operational efficiency

Complex, manual processes create operational bottlenecks that slow AI deployment and increase compliance risks. Streamlining enterprise AI execution requires policy-driven presets that automate continuous governance, compliance enforcement, and performance tuning at the infrastructure layer. By embedding intelligence directly into operational workflows, organizations can accelerate deployment speed, ensure consistent regulatory adherence, and eliminate friction across multicloud and on-premises environments.

Everpure for AI scale

To bridge the gap between legacy infrastructure limits and production readiness, Everpure™ replaces app-centric storage silos with a unified Enterprise Data Cloud (EDC) architecture built on the Everpure Platform. By embedding intelligence, resiliency, and performance into a single platform, Everpure directly addresses the three core challenges facing technical leaders.



1 Delivering AI-ready data through universal data intelligence

AI models are only as good as the data feeding them. Right now, dark data, siloed systems, and residency rules keep live operational context out of reach, and teams often can't tell if a model is working from data that's current or trustworthy. Everpure Data Intelligence, part of the Everpure Platform, automatically finds data and attaches security, privacy, and residency policies at the source so teams can verify what's feeding a model before they trust its output. Everpure Data Stream™ then converts raw data into AI-ready vectors directly in the data path, feeding retrieval-augmented generation (RAG) pipelines and inference engines without a separate extract, transform, and load (ETL) step. That cuts processing time and lowers token costs.

2 Matching parallel performance to GPU scale

Millions of dollars in GPU compute often sit idle waiting on storage. FlashBlade//S™ closes that gap. It moves data fast enough, delivering up to 7.2 million metadata IOPS to keep GPU clusters busy instead of stalled. For massive training datasets, FlashBlade//EXA™ adds exabyte-scale, all-flash capacity that hit 6,300 concurrent AI jobs in SPECstorage benchmark testing while cutting power, cooling, and rack footprint demands by up to 85%.

3 Eliminating operational drag with policy-driven automation

Manual governance slows AI deployment and raises compliance risk every time a policy change requires infrastructure to be updated by hand. The Everpure Platform runs that work on policy instead: presets handle governance, compliance, and performance tuning at the infrastructure layer, and one control plane applies those policies consistently across multicloud and on-premises environments. Teams stop rebuilding for every new location and start deploying faster, with regulatory adherence built in instead of checked after the fact.

Strategic outcomes for technical leaders

Transitioning to an Everpure AI data foundation directly fulfills the following core strategic mandates:

- **Keep GPU and compute clusters running**: 99.9999% availability with nondisruptive updates means clusters stay saturated, not idle.
- **Strengthen governance and compliance**: Automated semantic governance and resilient sourcing protect data, respect regional sovereign boundaries, and keep project timelines on track.
- **Cut manual work**: Automated policy execution removes 70–80% of manual operational tasks, so teams manage data instead of babysitting it.
- **Keep the architecture flexible long-term**: One software-defined foundation spans on-premises, edge, and cloud environments, so scaling doesn't mean a forklift upgrade.
- **Improve ROI**: Embedding compliance directly into the data path cuts the need for third-party audit tools and advisory consulting, freeing that budget for AI work instead.

> “Building a scalable AI factory requires an infrastructure foundation that combines raw performance, total reliability, and energy efficiency. Everpure allows us to unlock secure AI at scale without increasing operational complexity or environmental impact.”
>
> **Wojciech Stramski, CEO, Beyond.pl.**

Read the Full Story

Summary and next steps

An AI-ready foundation isn’t more hardware or another point solution bolted onto what you already have. A unified platform gets technical leaders past experimentation and into production, with the performance and security to back it up.

Everpure treats storage as active infrastructure, not passive capacity: unified, governed, and understood wherever your data lives. That’s what turns audit spend into budget you can put back into building instead of defending what you already have.

See What That Looks Like with Enterprise Data Cloud

1 | https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk

Visit Our Website **800.379.PURE**

©2026 Everpure, the Everpure P Logo, Pure Storage, Evergreen, Everpure Data Stream, FlashBlade//EXA, FlashBlade//S, and the marks in the Everpure Trademark List are trademarks or registered trademarks of Everpure, Inc. or its licensed subsidiaries in the U.S. and/or other countries. The Trademark List can be found at everpuredata.com/trademarks. Other names may be trademarks of their respective owners.

ID55705-01-en-09/26