




AI Moves Fast Your Data Infrastructure Should Too





The Pace of AI Innovation is Unprecedented





Preparing & Architecting Enterprise AI with Everpure

Your AI Strategy depends on the right Data Strategy

Data Governance

- Accessibility
- Security - Audit
- Backup/Compliance

Data Readiness

- Data Prep and Cleaning
- Data Logging
- Data Refresh

AI Readiness

- Validated AI solutions
- Scalability
- Observability



Everpure is the #1 Enterprise AI Data Platform

Enterprise Leadership

#1

IDC Market Share Changes
Everpure has the fastest growing market share

11X

Gartner® Magic Quadrant™
Positioned as a Leader for eleven consecutive years

81

Net Promoter Score
One of the highest in the industry

AI Leadership



NVIDIA Certifications, Validated AI Solutions, and Services across all AI use cases



The Only Platform with **Non-Disruptive Upgrades and 99.9999% Uptime****



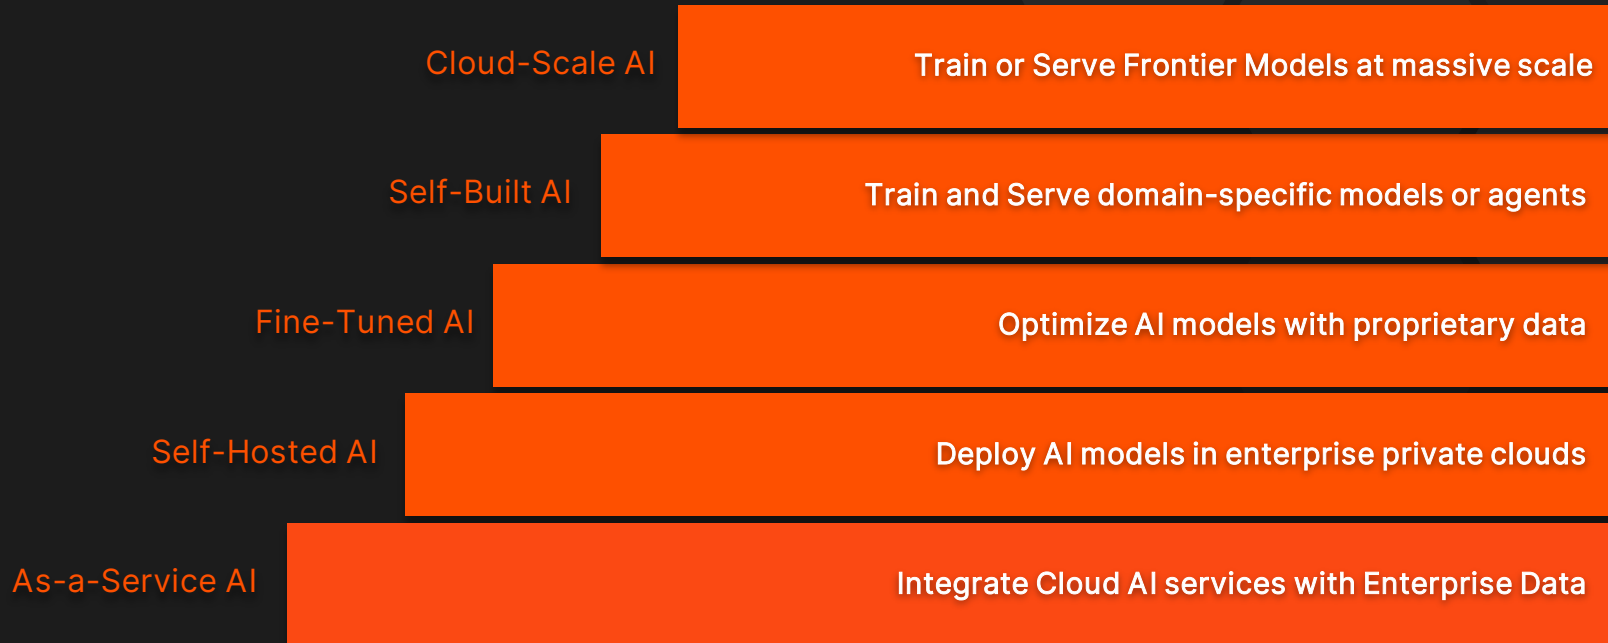
Exabytes of Storage Serving Cloud-Scale AI Training and Inference



Installed in **13,000+** Mission Critical IT Environments Across Every Vertical

*Source: IDC, Market Share Changes From 2014 Through 2023
**99.9999% SLA with Evergreen / One

Enterprise AI Capability Index



Everpure AI Wins Across Enterprises



Cloud-Scale AI



Self-Built AI



Fine-Tuned AI



Self-Hosted AI



As-a-Service AI



DEPAUL UNIVERSITY



SoftBank



paylocity
Forward Together.

side includes CONFIDENTIAL wins- NDA required



As-a-Service AI

Goals	Pain Points	Solved with Everpure
Unified Data	Data lives in enterprise silos	Everpure Platform, Pure Fusion Enterprise Data Cloud File, Object, Block with MCP APIs
Data Readiness	Unstructured & inconsistent data	GenAI Pod, Starburst, Dremio GenAI Solutions, Lakehouse Architectures and services
Data Governance	Lack of audit/control for AI data flows	Platform Governance, Varonis, Rubrik Fine-grained Access Control, Auditing, Threat Detection, Cyber Resilience
	<i>Data Access complexity causes many AI project to stall and fail after initial prototype</i>	<i>Unified, governed, and connected data foundation ensures AI readiness and unlocks value from data</i>



Self-Hosted AI

Goals	Pain Points	Solved with Everpure
Availability	Keeping infrastructure current without disrupting service level or uptime	Evergreen Non-Disruptive Upgrades
Automation	Lack of automation for DevOps/MLOps	Portworx Simplified DevOps & MLOps Cloud Container Portability
De-risk Deployment	Software-Hardware Imbalance	Nvidia Run ai, MLflow, Milvus Full-stack AI Solutions & Certifications
	<i>Manual operations and bespoke upgrades using point integrations create fragility</i>	<i>Non-disruptive upgrades, fleet operations, and validated stacks reduce operations risk</i>



Fine-Tuned AI

Goals	Pain Points	Solved with Everpure
Model Accountability	Data versioning & lookup	Everpure Data Platform Application Aware Snapshot Fast Metadata lookup
Validated Design	Difficulty choosing the right infrastructure	AIRI, Flashstack NVIDIA BasePOD, SuperPOD NVIDIA Certified & Validated AI Solutions
Consistent Performance	Supporting varying SLA requirements	Evergreen One Guaranteed SLAs, Storage-as-a-Service
	<i>Complex data pipelines, infrastructure uncertainty, and SLA pressures slow reliable model fine-tuning</i>	<i>Integrated snapshots, validated stacks, and SLA-backed services streamline enterprise AI fine-tuning</i>



Self-Built AI

Goals	Pain Points	Solved with Everpure
TCO Optimization	Growing AI workloads strain both GPUs and storage budgets	Everpure Platform capabilities NFS + S3/RDMA, GDS (GPU Direct Storage)
Elasticity	Unable to meet unexpected AI workload spikes	Portworx, Fusion, Everpure Platform Container/GPU burst-to-Cloud Storage optimizer, Fairness policy
Financial Flexibility	High upfront costs and unpredictable scaling expenses	Evergreen One Flexible Consumption Models Opex and Capex
	<i>AI growth drives up costs, limits elasticity, and locks spend into inflexible infrastructure</i>	<i>Optimized GPU/storage, elastic scale, and flexible consumption reduce cost and risk</i>



Cloud-Scale AI

Goals	Pain Points	Solved with Everpure
Metadata Performance	Metadata bottlenecks break performance at massive scale	Purity Metadata Operations at Scale Trillions of Files & Objects, Zero Lag
Extreme Capacity	AI at scale stalls under capacity, performance, and namespace limits	FlashBlade//Exa Unlimited Scalability to Exabytes/TBps
Custom Architecture	Requirement of unique customizations for latency, reliability or performance	Direct Flash Custom integrated ultra-low latency, Reliability-optimized flash module
	<i>Metadata strain and scaling limits stall AI; extreme scale demands custom infrastructure design</i>	<i>Pure unlocks cloud-scale AI with limitless scale, ultra-low latency, and the ideal foundation for custom environments</i>





Everpure AI Value Across Enterprises



Cloud-Scale AI

Metadata Ops at Scale | Unlimited Scalability | Foundation for Custom Architecture

Self-Built AI

Optimize GPU & Storage | Hybrid-Cloud Ready | Flexible Consumption & Elastic Scale

Fine-Tuned AI

Model Accountability point-in-time copies | Validated Designs | SLA-backed Availability

Self-Hosted AI

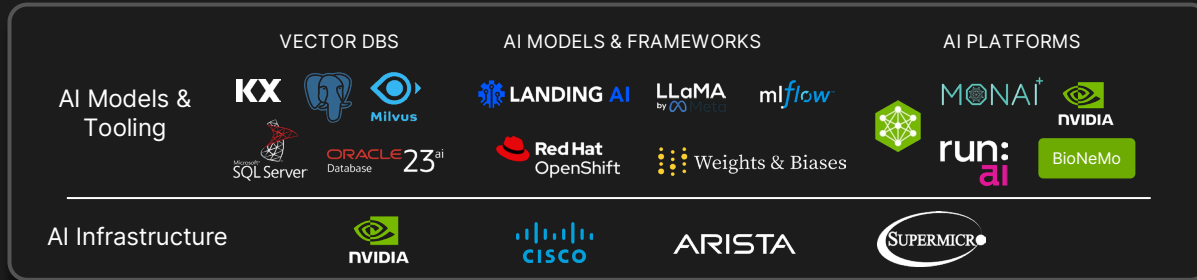
Non Disruptive Upgrades | Cloud Portability & Storage Automation | De-risk Operations

As-a-Service AI

Enterprise Data Cloud | Unified File, Object & Block | MCP (Model Context Protocol) APIs

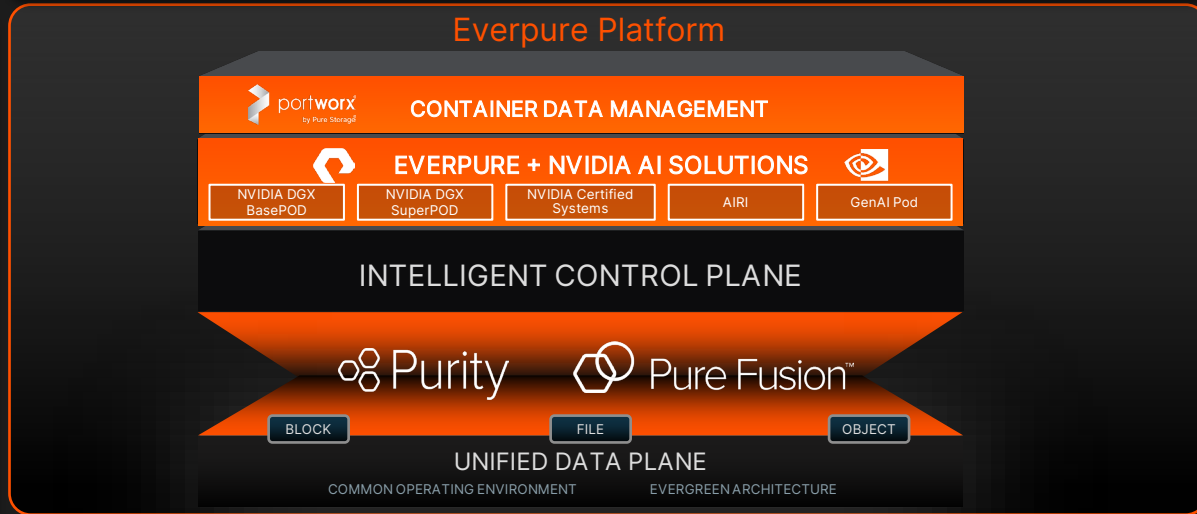


One Data Platform for All Enterprise AI Capabilities



Data Engineering

- dremio
- kafka
- Redpanda
- APACHE Spark
- Starburst



Cyber Resilience

- Commvault
- CROWDSTRIKE
- veeam
- rubrik
- superna
- VARONIS



Next Steps

Get Started!

Attend
Everpure AI
Workshops &
Assessments

Evaluate Validated Designs

Everpure Certified
AI Solutions &
Integrations

Collaborate with Experts

Work with Everpure
AI Experts for
advanced initiatives







Appendix

Everpure AI Solutions



Cloud-Scale AI



Self-Built AI



Fine-Tuned AI



Self-Hosted AI



As-a-Service AI



AI Messaging for 2H'2026

Positioning statement	<i>Accelerate AI Inference & Training with Simplicity, Performance, Value and Reliability.</i>			
Customer Value	Simplicity of Fully Validated AI Solutions	High Performance at Massive Scale	Industry-leading AI Economics	Always-ON AI Inferencing & Training
Benefits:	<ul style="list-style-type: none"> • Faster Time to Deployments with Fully Validated Solutions, RAs and Certified Systems with NVIDIA • Storage Automation & Cloud portability with Enterprise Data Cloud • Ecosystem integration such as MCP to unlock LLM based storage workflows at massive scale 	<ul style="list-style-type: none"> • 20X Faster Inference with the First KV Cache for S3 and NFS • Exabytes of Storage Serving Cloud-Scale AI Training and Inference • Improve Metadata performance with high parallelization 	<ul style="list-style-type: none"> • Drive cost savings and efficiency for better GPU utilization through Storage optimization • Address Data Readiness and Data Governance for AI deployments • Accelerate your AI Factory to get to Mission Critical AI Outcomes Faster 	<ul style="list-style-type: none"> • De-risk Operations with Non Disruptive Upgrades of Evergreen • Embrace peace of mind with our SLA guarantee of 99.9999% uptime and zero planned downtime guarantee.
Customer Validation:	<p style="text-align: center;">NTT</p> <p>Our partnership with Everpure has given us an AI-ready infrastructure that can keep up with the high-throughput demands for big and fast data to power our clients' AI initiatives</p>		<p style="text-align: center;">NASA</p> <p>Aims to accelerate development of new ideas at NASA by 25%</p>	<p style="text-align: center;">AC Milan</p> <p>AC Milan Innovates On and Off the Pitch with Everpure</p>

