

SOLUTION BRIEF

Evergreen//One for AI

Solving the unpredictability of AI with SLA-driven Storage-as-a-Service

Nearly every organization recognizes the transformative power of AI to solve real world problems and deliver competitive advantages for their organization. But the exponential growth of AI applications and use cases, coupled with the proliferation of GPU-accelerated computing, makes sizing the storage environment to support these data-hungry workloads difficult to achieve. Deploying too little storage creates performance bottlenecks for your expensive compute, data science, and AI development investments, while over-provisioning leads to cost inefficiencies. Compounding these challenges, many AI projects start small and grow over time. This makes storage sizing even more difficult to forecast and guessing at the beginning of a 3-5 year refresh cycle isn't going to cut it. Leveraging the cloud is an option, but as AI projects grow, costs become prohibitive. Additionally, much of the critical data needed to train differentiated AI models often resides in on-premises storage environments where data privacy, security, compliance and performance requirements necessitate keeping it under your control.

Introducing Evergreen//One for AI

Skip the sizing guesswork with Pure Storage® Evergreen//One™ for AI, a unique service from Pure Storage designed specifically for AI infrastructure deployments. Evergreen//One for AI is the perfect fit for the unpredictability of AI. Driven by its experimental and iterative development processes, AI's data requirements, performance needs, and scalability demands can vary dramatically and evolve rapidly. By delivering performance guarantees based on the maximum bandwidth requirements of your GPUs and accelerators, you never have to worry about storage becoming a bottleneck. Evergreen//One for AI can meet the most extreme performance requirements of Nvidia OVX servers, DGX BasePOD, and even DGX SuperPOD¹ clusters. That's a promise we'll keep with SLA-driven storage-as-a-service (StaaS).

With the Uncertainty of AI, Predictability Is Paramount

Various AI workloads have different storage requirements. Data variability, model complexity, data quality, training dynamics, evolving performance needs, and additional external factors add to the difficulty of sizing the right storage footprint to support these workloads. Evergreen//One for AI delivers an SLA that guarantees



Certified

Proven for AI with Nvidia OVX server, DGX BasePOD, and DGX SuperPOD* certification.



Energy Efficient

Dedicate energy expenditures to GPUs and accelerators. Pure Storage products reduce energy use and carbon emissions by up to 85%.



Performant

No performance ceiling. Evergreen//One for AI guarantees performance for any AI workload, up to multiple TB/s and beyond.

consistent and predictable storage performance, regardless of workload intensity or the number of GPUs, ensuring valuable AI investments are not underutilized. Storage costs are known upfront, no matter the workflow, with a simple rate based on the maximum bandwidth required. Incremental capacity is billed at competitive, cold storage-like rates. There's no need to delete precious artifacts, checkpoints, or logs. With predictable spend per month and the certainty of an SLA, you get the storage capacity you need today and can invest the savings into more strategic AI expenditures tomorrow.




FIGURE 1 Evergreen//One service catalog

Continuous innovation Delivered Non-disruptively

Evergreen//One for AI is true to the Pure Storage Evergreen® architectural principles by delivering continuous innovation non-disruptively. As AI workloads change and grow, Evergreen//One for AI ensures that any upgrades to your storage infrastructure are effortless on your part. The Evergreen architecture delivers the latest technology advancements, capacity expansions, performance enhancements, and feature updates, without the need for costly forklift upgrades or data migrations, ensuring AI workflows are never interrupted.

Included with your Evergreen//One for AI subscription is Pure1®, an AI-driven data-services platform. With Pure1 Meta, an artificial intelligence engine, you can see into your future and accurately forecast your application and infrastructure or service needs, with continuous monitoring and proactive resolution of issues. Pure1 is the single place for you to subscribe, monitor, and extend your Evergreen//One subscription with predictive intelligence, full-stack analytics, global end-to-end visibility, and assessments for sustainability and data protection.



Focus on AI Innovation, Not Managing Your Storage Footprint

By offering comprehensive management and support throughout the lifecycle of your storage infrastructure, Evergreen//One for AI eases the burden of administration, allowing your organization to focus resources and expertise on advancing your AI initiatives. Pure handles every aspect of management, from initial deployment and configuration to ongoing monitoring, maintenance, and optimization. Whether for training, tuning, inferencing, deep learning or deploying RAG workflows, by subscribing to a service, Evergreen//One for AI can support any predictive or generative AI workload with guaranteed performance levels that maximize your investments.

Additional Resources

- Learn more about [Evergreen//One](#).
- Read the Evergreen//One [ESG Economic Validation](#).
- Try the Evergreen//One [pricing calculator](#).
- Learn more about Pure [Artificial Intelligence solutions](#).
- Learn more about [AIRI AI Infrastructure](#).

1 Nvidia DGX SuperPOD certification expected Q3 CY2024. While Pure Storage is committed to pursuing these certifications, it should be understood that any forward-looking statements about certifications are based on current expectations and are not promises or guarantees.

purestorage.com

800.379.PURE

