Pure//Launch Blog March Edition

At Pure Storage, we're constantly evolving to meet the ever-changing needs of our customers. We strive to provide solutions that break through barriers and deliver unmatched performance. Here's a look at our latest advancements designed to propel your business forward with unparalleled efficiency and security.

Simplified Storage Management with Pure Fusion for FlashBlade

Managing disparate storage arrays creates unnecessary complexity in modern data centers. The latest **Purity//FlashBlade 4.5.5** release extends Pure Fusion™ capabilities to support all protocols including file and object on FlashBlade®,

transforming how enterprises manage their unstructured data at scale.

This powerful upgrade enables:

- **Unified management** of entire FlashBlade fleets as a single entity
- **Remote provisioning** and resource allocation while minimizing management overhead
- **Unifying** the management capabilities of Pure Fusion with FlashBlade to deliver the option of managing FlashBlade and FlashArray™ resources together as a single fleet

By streamlining storage management processes and improving cross-site data access, organizations can reduce operational overhead while maintaining data integrity across geographically distributed environments.

Learn more about Pure Fusion

Accelerate AI Training and Inference

We know that data access efficiency is often the bottleneck in AI workflows. Our collaboration with NVIDIA to integrate **S3 over RDMA into FlashBlade** tackles this challenge by fundamentally reimagining how AI applications interact with object storage.

FlashBlade Object Store is a native first-class offering that provides high-performance, highly resilient AWS S3-based object storage for next-generation applications.

Remote Direct Memory Access (RDMA) is a technology that significantly improves data transfer efficiency, as well as read/write latency for AI/ML environments. S3 over RDMA will enable RDMA capability for S3 access.

This integration delivers:

- **Better performance** through improved throughput with data transferred

via direct memory access, speeding up the data transfer process

- **Reduced system overhead** via better CPU utilization with data transfer happening directly from GPU memory, bypassing the CPU bounce buffers
- **Faster training and inference cycles** due to the fact that RDMA transfers are handled directly at the network card level, bypassing the kernel and network stack
- **More energy-efficient GPU farms**, enabled by RDMA improving client-side CPU utilization, which allows more data to be processed without the need for additional compute

By optimizing this critical data path, organizations can achieve faster time to results for both training and inference workloads, ultimately accelerating their AI initiatives while maximizing infrastructure investments.

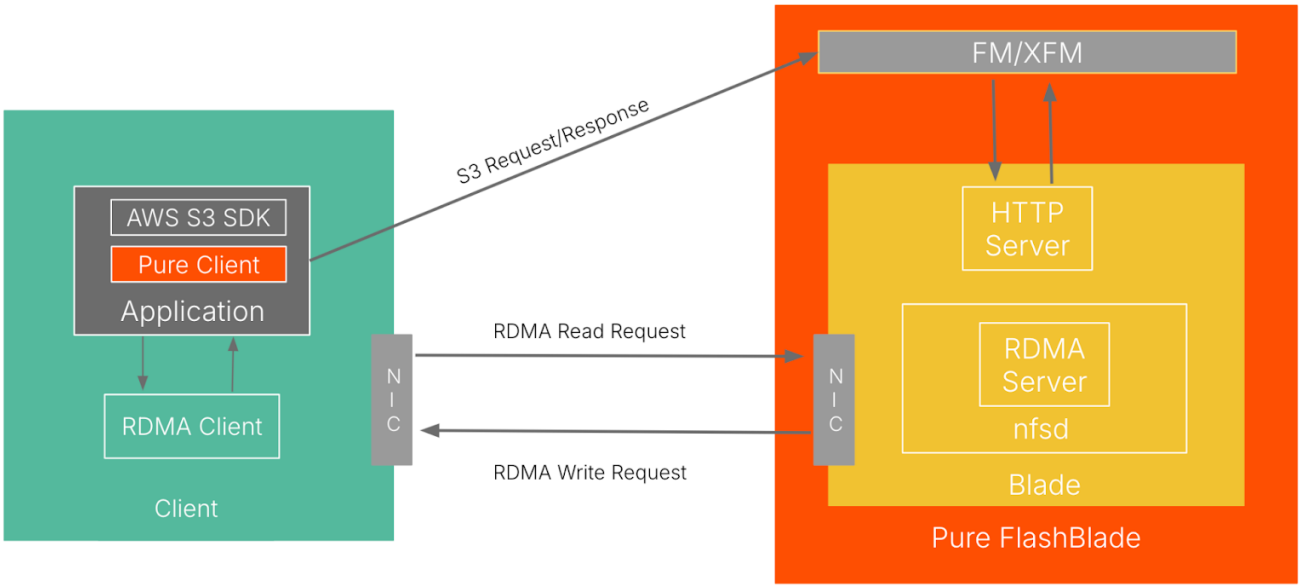

Figure 1. How S3 over RDMA works for FlashBlade.

Contact your Pure Storage representative to learn more!

Industry-leading Performance to Power Extreme-scale AI and HPC Workloads: Introducing FlashBlade//EXA

AI and HPC workloads are pushing infrastructure to new extremes, demanding a data storage platform built for performance, scale, simplicity, and adaptability. **FlashBlade//EXA™** is the industry's most powerful data storage platform, designed to power the world's largest AI factories and accelerate AI innovation.

FlashBlade//EXA enables large-scale AI and HPC workloads with:

- **Massively parallel data access:** Delivers unmatched multi-dimensional performance, enabling rapid access to large multimodal data sets, optimizing AI training and inference, and maximizing GPU utilization.
- **Metadata-optimized architecture:** Overcomes bottlenecks associated with legacy parallel file systems with a proven metadata core, powered by Purity//FB, deployed across thousands of AI and HPC environments.
- **Disaggregated design for future-ready scalability:** Independently scale metadata and data to adapt to evolving AI models, eliminating rigid infrastructure constraints.
- **Simplified AI infrastructure management:** Streamline storage management, reducing complexity compared to legacy parallel file systems and first-generation disaggregated models.

With FlashBlade//EXA, AI innovators can maximize GPU utilization, eliminate bottlenecks in the AI pipeline, and accelerate time to insight—all with the simplicity, agility, and efficiency that Pure Storage is known for.

Discover FlashBlade//EXA

Learn more about how Pure Storage powers AI workloads

High-performance Storage (HPS) Certification for NVIDIA Cloud Partners

The complexity of deploying AI data infrastructure at scale presents significant challenges for cloud providers. Pure Storage is now officially certified as a high-performance storage (HPS) platform for NVIDIA Cloud Partners (NCP). This certification acknowledges Pure Storage as a trusted storage partner and validates the ability of Pure Storage to support NVIDIA Cloud Partners with a proven reference architecture that meets all performance and technical requirements for building state-of-the-art GPU clouds.

For AI cloud providers, this translates directly to reduced time to market and lower implementation costs. It also enables providers to deliver exceptional performance to their customers without the traditional integration challenges and deployment risks.

Learn more

NVIDIA’s AI Data Platform for Enterprise

The gap between raw data and actionable insights remains a challenge for many enterprises. Our collaboration with NVIDIA includes support for the NVIDIA AI Data Platform for Enterprises, a next-generation AI infrastructure that integrates enterprise storage with NVIDIA-accelerated computing to power AI agents with near real-time business insights. Pure Storage enhances this collaboration by delivering a high-performance, future-proof AI infrastructure that simplifies data management, reduces costs, and accelerates time to insight.

This platform combines the enterprise-grade storage capabilities of Pure Storage with NVIDIA’s accelerated computing to enable:

- Automated data ingestion and embedding at scale
- Sophisticated reasoning and task execution
- Contextual semantic search across enterprise data

- Generative capabilities that transform raw information into business insights

The result is a transformative infrastructure that empowers AI agents to deliver near real-time intelligence, turning data from a static asset into a dynamic business advantage.

Get more information

NVIDIA-Certified Storage for Foundation and Enterprise-level Reference Configurations

Building effective AI infrastructure requires careful coordination between storage, compute, and networking. Our **NVIDIA-Certified Storage** solutions address this complexity by providing pre-validated configurations that ensure optimal performance for accelerated analytics, classical AI, and generative AI workloads.

These reference architectures deliver:

- Guaranteed performance across diverse AI use cases

- Enterprise-grade security and data protection
- Simplified deployment and management
- Seamless scalability as demands increase

By partnering with systems integrators and enterprise customers, we're enabling the construction of high-performance AI data centers that serve as the foundation for manufacturing intelligence and efficient AI factories—eliminating guesswork and accelerating time to value.

Discover more about Pure Storage and NVIDIA

Protecting Large-scale Deployments of Unstructured Data

Unstructured data growth creates significant data protection challenges. **Rubrik NAS Cloud Direct** for Pure Storage addresses these complexities with an intelligent, policy-driven approach to data protection.

This comprehensive solution delivers:

- Automated protection for files, databases, and other unstructured data types
- Simplified management through intuitive policy configuration
- Reduced risk exposure through consistent, reliable backups
- Minimized human error through automation

By combining efficient data management from Pure Storage with Rubrik's protection capabilities, organizations can ensure their most valuable unstructured data remains secure and recoverable, regardless of scale.

Go deeper

Security and Simplicity Powered by Pure1

Administrative overhead and access management complexity can drain IT resources. The **Pure1® Identity Access Manager 2.1** update streamlines security administration while enhancing protection.

This update introduces:

- Simplified authentication that eliminates multiple login requirements
- A new “Operator” role for efficient multi-organization management
- Streamlined access workflows that reduce administrative burden

Additionally, **Pure1 Log Center** centralizes critical operational data across your Pure Storage environment, delivering:

- Comprehensive user behavior insights
- Complete audit trails for compliance and security
- Activity correlation for faster troubleshooting
- Seamless integration with monitoring tools like Datadog

These enhancements work together to simplify operations while strengthening security posture—a winning combination for today’s resource-constrained IT teams.

Try out Pure1 Identity Access Manager (Pure1 login required)

Try out Pure1 Log Center (Pure1 login required)

We’re committed to continuously improving our offerings and ensuring our customers have access to the best technology available. Stay tuned for more updates and innovations from Pure Storage, and reach out to learn how these solutions can transform your data infrastructure.

Contact us today to schedule a demo and see these innovations in action.