TECHBYTE
Insights Into Research Computing

ON DEMAND systems

FUTURE-PROOFING DATA STORAGE: THE ROLE OF ON DEMAND ARCHIVE IN HPC & AI WORKLOADS

Authored by:

Rakesh Sabharwal
Founder & CEO
On Demand Systems Pte Ltd

www.odspte.com

The Data Storage Bottleneck in HPC & AI

Data-intensive workloads in High-Performance Computing (HPC) and Artificial Intelligence (AI) are growing exponentially. With petabyte-scale datasets becoming the norm, organizations need a well-designed storage strategy that balances performance, scalability, and cost efficiency.

Traditional storage architectures, whether all-flash or disk-based systems struggle with scalability, access latency, and cost overheads. This is where **On Demand Archive** comes in, an intelligent, high-performance archiving solution designed to optimize HPC and AI data workflows and storage.

In association with **On Demand PFS** and other third-party partner solutions, the solution can support advanced features of Intelligent Data Tiering and Policy Based archiving and retrieval.

ON DEMAND systems

Understanding HPC and AI Storage Demands

High-Performance Computing (HPC) Storage Needs:

HPC workloads involve parallel computing clusters running large-scale simulations, modeling, and data analytics. These workloads demand:

- High-throughput I/O to process terabytes per second.
- Low-latency access for real-time computations.
- Scalable capacity to store datasets generated over time.

Artificial Intelligence (AI) Storage Challenges:

AI pipelines involve:

- Training datasets requiring rapid access for iterative model updates.
- Inference workloads that need quick retrieval of archived models.
- Hybrid storage workflows, where AI dynamically moves data between hot and cold tiers.

Traditional NAS, SAN, Object Storage and backup solutions cannot efficiently handle the performance-cost tradeoff, leading to excessive CapEx and OpEx.

How On Demand PFS & On Demand Archive Solves These Challenges

On Demand PFS and **On Demand Archive** provide an intelligent, scalable, and high-performance parallel file system integrated with archiving system tailored for HPC and AI workloads. Both solutions can also be deployed independently.

1. Intelligent Tiering & Policy-Based Archiving

- Automated data movement between high-performance storage (NVMe, SSD) and cost-efficient archive layers (HDD, tape, object storage).
- Policy-driven archival, data is classified based on usage frequency, enabling seamless tiering.
- Hierarchical Storage Management (HSM) integration for tape, object, and cloud archival.

2. High-Performance Metadata Indexing

- Uses distributed metadata management to accelerate file lookup.
- Metadata caching with SSD/NVMe layers to enable sub-millisecond search times.
- Eliminates cold storage delays through predictive prefetching algorithms.

How On Demand PFS & On Demand Archive Solves These Challenges

3. Parallel Data Access for HPC & AI

- Supports high-speed parallel file systems.
- Optimized for InfiniBand fabric reducing network latency.
- Provides S3-compatible APIs for seamless integration with AI data lakes.

4. Scalability & Cost Optimization

- Linear scalability across petabyte to exabyte-scale deployments.
- Object-based deduplication & compression reduce storage footprint.
- Pay-as-you-go model for cloud-based archival, reducing TCO.

Note:
These capabilities are included in the software stack using BeeGFS and Ceph and other third-party solutions which are supported as part of the overall solutions offering.

ON DEMAND systems

Why On Demand PFS and On Demand Archive is Future-Proof for HPC & AI

- **Accelerates Data-Intensive Workflows:**
Speeds up HPC simulations and AI training.

- **Optimizes Storage TCO:**
Automates intelligent tiering to reduce cost.

- **Integrates with HPC & AI Ecosystems**:
Supports PFS, S3, InfiniBand.

- **Scales with Demand:**
Handles exponential data growth without infrastructure overhauls.

- **Reliable long-term storage:**
On Demand Archive ensures data integrity, which guarantees that all data is stored correctly on the underlying media.

Conclusion: The Future of HPC & AI Storage is On Demand PFS & On Demand Archive

As data volumes grow exponentially, the need for intelligent, scalable, and cost-efficient storage becomes critical.

On Demand PFS & **On Demand Archive** empower organizations to take full control of their data lifecycle management, from high-performance processing to long-term archival, with automation, efficiency, and speed at the core.

By ensuring AI models and HPC applications always have timely access to the right datasets, these solutions eliminate bottlenecks, reduce storage overhead while ensuring seamless scalability.
With **On Demand Systems**, you're not just storing data, you're building a foundation for innovation.

Smarter. Faster. Future-ready.
That's the On Demand Systems way.

Ready to Future-Proof Your HPC & AI Storage?

Whether you're managing petabytes today or planning for exabyte-scale tomorrow,
On Demand Archive & **On Demand PFS** are built to evolve with your data.

Let’s explore how we can:

- Streamline data workflows
- Scale effortlessly
- Minimize storage overhead without compromising performance or cost-efficiency

Contact us to learn more, request a consultation or discuss your unique storage challenges.

Our Contact Details:
Email: mail@odspte.com
Website: www.odspte.com