VDURA DATA PLATFORM V11

Built to Power the AI Pipeline

WHITE PAPER

May 2025

Contents

A Legacy of Innovation, Reimagined for the AI Era 1
The VDURA Advantage: Core Capabilities 2
Built for Every Workload That Demands Performance and Trust 3
The AI Workload Storage Problem 4
The AI Data Pipeline: What Happens at Each Stage 4
Common Challenges in the AI Pipeline 4
Modern Storage for the AI Era 6
The VDURA Data Platform 7
- Software-Defined Architecture Engineered for Performance and Scalability 7
- Direct, Parallel Data Access for the AI Pipeline 8
- Linear Scalability, Seamless Expansion 9
- Director Nodes: The Brain Orchestrating Control in a Parallel World 10
- Storage Nodes: AI Pipeline Performance from Every Layer 10
- Optimized for Every Phase of the AI Pipeline 11
- VPOD Architecture: Virtualization for Resilience and Efficiency 11
- Dynamic Data Acceleration: The Smart Storage Fabric 12
- Resilient Data Reconstruction and Integrity 12
- AI-Aware Placement and Automation 13
- VDURA V11 Data Reduction 13

VDURA Data Platform V5000 System Overview 13
- Key Features 14
- V5000 Details 14
  - Director Nodes 15
  - All-Flash Storage Nodes 15
  - All-Flash with HDD Capacity Expansion Storage Nodes 15
  - Connectivity 15
- Expansion Options 16
- Physical Connectivity 16
- Network Configuration Options 16
  - Active/Active Link Aggregation Mode 16
  - Single Link Mode 17
  - Network Failover Mode 17
- Storage Configuration Options 17
  - Storage Sets 17
  - Volumes 17
- RAID and Erasure Coding 18
  - Reliability That Increases with Scale 18
  - An Architecture of High Reliability 19
  - Per-File RAID 20

Comprehensive Data Services 20
A Competitive Edge 20
Where Velocity Meets Durability 21

Abstract

The VDURA® Data Platform™ represents a revolutionary approach to data infrastructure for AI and high-performance computing workloads. This white paper details how VDURA combines the performance of a true parallel file system with the resilience and cost-efficiency of object storage in a unified software-defined architecture. Built on a legacy of innovation in high-performance computing, VDURA delivers unprecedented parallel throughput, ultra-low latency metadata operations, and superior data protection across the entire AI pipeline from ingest and training to checkpointing and inference.

With linear scalability to thousands of nodes, integrated value-tier storage, and comprehensive data services, VDURA eliminates the traditional compromises between performance, durability, and cost that plague AI infrastructure today. This document outlines the architectural components, performance capabilities, and deployment options of the VDURA Data Platform and V5000™ system, showcasing why leading organizations trust VDURA to power their most demanding workloads.

Revisions

Date	Description
May 2025	Initial release.

A Legacy of Innovation, Reimagined for the AI Era

The VDURA Data Platform is built for the AI era, combining true parallel performance, hyperscale durability, and effortless scalability in a single, software-defined infrastructure.

Before the rise of AI, cloud-native workloads, and modern data infrastructure, there was Panasas®, a company that reshaped the high-performance computing (HPC) landscape with the industry's first true parallel file system and defined the pNFS standard. For more than 20 years, Panasas PanFS® set the bar for scalable performance, mixed workload efficiency, and enterprise-grade reliability in environments where data is everything.

Today, that legacy powers VDURA.

Built on the core Panasas architecture focused on data protection and resiliency, VDURA is the modern evolution of a platform trusted by the world's most data-intensive environments. VDURA has reimagined that foundation into a microservices-based, software-defined solution that's simple to deploy, operate, and grow, while preserving the unmatched performance and resilience the industry has come to expect.

We're not new to this space; we helped define it.

VDURA was born from a deep understanding of what it takes to serve data at massive scale with zero compromise. Our architecture is the culmination of decades of real-world deployments in the most demanding environments across thousands of the world's leading research labs, government institutions and Fortune 500 companies. Now, we've evolved that foundation to meet the AI data challenges of tomorrow.

VDURA is where velocity meets durability.

The name says it all; lightning-fast NVMe flash throughput meets industry-leading durability in a platform that scales linearly to thousands of nodes. VDURA combines the scalable and linear high performance of a true parallel file system with a cost-efficient, resilient object store, unifying active and bulk storage under one architecture. The VeLO™ metadata engine powers intelligent data flow and fast namespace operations, delivering a software-defined platform built for AI and HPC that's simple to deploy and effortless to scale.

The VDURA Advantage: Core Capabilities

Through its sophisticated software, the VDURA Data Platform can transfer terabytes of data per second to and from your compute cluster. VDURA manages this orchestration without manual intervention, continuously balancing the load across those systems, automatically ensuring resilience, scrubbing the stored data for the highest levels of data protection, and encrypting the stored data to safeguard it from unwanted exposure.

- **True Parallel Performance**
  VDURA bypasses bottlenecks with direct, parallel data transfers from NVMe flash Storage Nodes to the client. Unlike NFS or "sort-of-parallel" systems, VDURA's shared-nothing architecture and separate metadata plane eliminate contention, delivering maximum throughput, lowest latency, and the consistent performance AI workloads demand at scale.
- **Blazing Metadata Performance**
  VDURA's VeLO metadata engine is a distributed, flash-optimized key-value store that delivers ultra-low latency for billions of file operations. Built for AI, it accelerates metadata-heavy tasks like model staging, small-file access, and checkpointing, minimizing bottlenecks and enabling consistent performance across large-scale, mixed I/O workloads.
- **Integrated Value Tier**
  VDURA natively integrates high-capacity extensions as a value tier, combining NVMe flash and HDD storage within a single platform. This eliminates siloed object stores and delivers cost-efficient, long-term storage under the same namespace, making VDURA ideal for AI data lakes, model checkpoints, and archival workflows.
- **Hardware Freedom**
  VDURA runs on commodity-agnostic storage with AI-grade speed and cost efficiency.
- **End-to-End Simplicity**
  VDURA delivers effortless scalability and management through a single-vendor, software-defined solution. The dynamic architecture enables seamless NVMe flash and HDD capacity expansion, while automatically balancing workloads and self-healing from failures. With a unified global namespace and intuitive management interface, VDURA minimizes operational overhead and eliminates manual tuning, letting customers focus on innovation rather than complex infrastructure.
- **Advanced Data Protection**
  Reliability improves with scale through client-side, file-level erasure coding that protects each file individually, eliminating the need for legacy RAID or costly HA hardware. VDURA's patented Multi-Level Erasure Coding™ (MLEC) provides superior data protection, delivering unmatched resilience for critical workloads.
- **Built-In, Enterprise-Grade Security**
  The VDURA Data Platform secures data from edge to core with a multi-layered "defense in depth" architecture. SELinux enforcement, native ACL support, and encryption-at-rest via self-encrypting drives (SEDs) and KMIP-based key management ensure robust protection without compromising performance. All layers are software-defined, scalable, and seamlessly integrated across the platform.

Built for Every Workload That Demands Performance and Trust

VDURA Data Platform V11 is the product of decades of engineering, customer experience, and innovation. It marks the beginning of a new chapter, one that merges HPC reliability with the velocity of AI and scales to wherever data needs to go next.

True Parallel File System	Blazing NVMe Performance	Integrated Value Tier	Advanced Data Protection	Enterprise Grade Security	Hardware Freedom

From AI frontier models to fluid dynamics and national research initiatives, VDURA is deployed in environments where failure is not an option. With every release, we've eliminated complexity, strengthened durability, and accelerated performance, delivering a data platform that works smarter, not harder.

- Flexible deployment across on-premises, hybrid, and private cloud environments.
- True parallel performance at scale with separate control and data planes under a single, global namespace.
- Native support for industry-standard protocols including NFS, SMB, S3, and CSI.
- VDURA DirectFlow® Client is a fully POSIX-compliant parallel, loadable Linux filesystem module that delivers cache coherence and parallel I/O via RDMA or TCP/IP networks.
- Linear growth in performance and capacity through scalable architecture of up to 1,500 nodes.
- Automated, workload-aware tiering from NVMe flash to HDD capacity expansion via Dynamic Data Acceleration (DDA).
- NVMe-optimized metadata engine (VeLO) with distributed key-value store for ultra-fast namespace operations.
- Built-in software-defined and scalable security via SELinux, ACLs, encryption-at-rest (SEDs), and KMIP integration.
- Multi-level erasure coding (client plus node level) for up to 12 nines of durability with three to five nines availability guaranteed.
- Self-healing and auto-balancing infrastructure with non-disruptive upgrades and seamless expansion.
- Integrated lifecycle management through quotas, data orchestration, and multi-tenancy controls.
- Optimized for AI training, checkpointing, inference, HPC simulations, and large-scale scientific workloads.
- Easy management through GUI, CLI, or API.
- High-speed connectivity across 400/200/100GbE, InfiniBand™ NDR/HDR, RoCE v2.

This isn't just a new file system. It's the most powerful, efficient, and resilient infrastructure available for the next era of data.

The AI Workload Storage Problem

AI workloads have redefined what modern data infrastructure must deliver: speed, scale, and precision under constant I/O pressure. But most storage architectures weren't built for this.

Each area of the AI pipeline is unique with different storage requirements to keep the factory running efficiently and better than your competitors. The current approach has been to pick different systems for each area or to default to data infrastructure that may meet some requirements but not be ideal for all.

The AI Data Pipeline: What Happens at Each Stage

- **Data Ingest**: Raw data (images, videos, text, or sensor streams) flows in, ready to be converted, stored, and staged for AI processing.
- **Model Load**: Massive pretrained model files (often hundreds of GBs to multiple TBs) are loaded into GPU memory before training or inference can begin.
- **Training**: Large datasets are read in randomized batches while updated model states, gradients, logs, and metrics are continuously written out.
- **Checkpointing**: At regular intervals, the entire model state is saved. This ensures progress is not lost and allows recovery or restarts without starting over.
- **Fine-tuning**: A previously trained model is refined using a smaller, domain-specific dataset, adapting it for new tasks or environments. Smaller checkpoints occur than during training.
- **Inference**: The model is deployed in production to generate predictions at high scale, often requiring ultra-fast reads of small files or objects.

Common Challenges in the AI Pipeline

- **GPU idle time**: GPU stalls caused by poor I/O throughput and lengthy checkpointing waste compute and slow AI innovation.
- **Metadata overload**: Traditional systems collapse under billions of small file operations, especially during inference and model versioning.
- **Flash waste**: Overprovisioning high-cost NVMe for rarely-accessed data drains budgets and inflates TCO.
- **Manual data tuning**: Static policies and manual tiering cannot keep up with constantly shifting data access patterns across stages.

Figure 1 below displays the storage workload at each stage. Note that every stage has different requirements to keep the factory running efficiently.

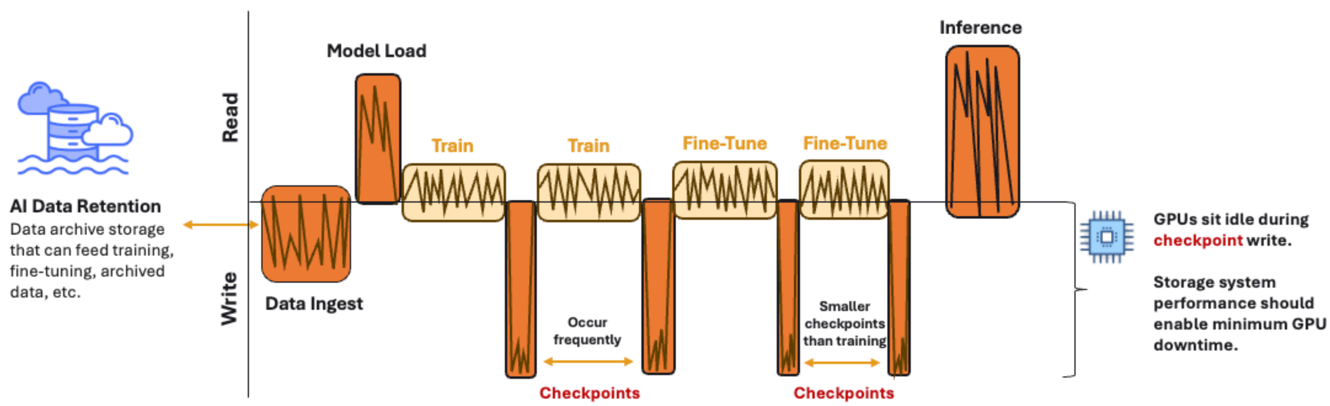

Figure 1. AI data pipeline storage workloads.

The following table displays the complexities of AI data infrastructure. Each stage in the AI pipeline has different read, write, throughput, capacity, and IOPs requirements that must be optimized.

Stage	Read	Write	Data Size	AI Workload Insights
Data Ingest	Low	**High**	TBs to PBs	Bulk writes require fast speeds. Data retention requires high capacity.
Model Load	**High**	-	GBs to TBs	High throughput required. Any delay in loading model holds back entire pipeline.
Training	Low	Low	TBs to PBs	Fast I/O crucial to saturate GPUs.
Checkpoint (Train)	-	**Very High**	GBs to TBs	GPUs idle during checkpointing. Checkpointing must be fast to prevent burning GPU dollars.
Fine-Tune	Low	Low	GBs	Smaller data sets than training. Typically lighter on reads/writes.
Checkpoint (Fine-Tune)	-	**Very High**	GBs	Workloads smaller than training checkpoints, but with high-speed write requirements.
Inference	**High**	Low	GBs	Fast reads crucial for real-time inference. Writes low for logs, caching, and vector outputs.
AI Archive / Data Retention	-	-	PBs	Acts as long-term, cost-efficient storage for raw or processed datasets that feeds pipeline.

Designing an AI infrastructure requires more than just performance at a single stage. It demands data storage infrastructure that can handle the full pipeline.

From high-throughput ingest to latency-sensitive inference, from massive training datasets to frequent checkpointing, each stage places distinct and evolving demands on the system. The right storage architecture must seamlessly balance read/write performance, scalability, data protection, and cost-efficiency.

Modern Storage for the AI Era

AI and HPC pipelines demand precision for fast writes during ingest, training, and fine-tuning checkpoints; high throughput reads during model loading and inference; and scalable, cost-effective storage for AI data retention and reuse. Most vendors force tradeoffs. Shared-everything architectures rely on centralized head nodes to handle all I/O, introducing performance chokepoints. Writes slow dramatically during cache flushes due to compression and deduplication. Bolt-on, third-party object stores for data lake functionality add latency, break the namespace, and shift complexity to the user.

These disjointed approaches cannot keep pace with modern AI workloads.

VDURA eliminates these limitations with a true parallel file system and software-defined architecture that separates the control plane from the data plane. This shared-nothing design enables scalable, high-performance throughput with no single-node bottlenecks. AI training data flows directly from NVMe flash to clients. AI archive and retention data lives cost-efficiently in high-capacity hybrid nodes all under a single global namespace.

Every stage of the AI pipeline is covered.

- Data ingest, fine-tuning, inference, and AI data retention leverage high-performance NVMe and cost-effective, high-capacity HDD storage.
- Model load, training, and checkpointing run at full speed on all-NVMe flash, delivering up to 1.4 TB/s throughput and 45 M IOPs per rack.
- Metadata-heavy tasks like small file access and checkpoint orchestration are accelerated by the VeLO distributed key-value engine, designed specifically to address the AI pipeline and built for scale.

Intelligent orchestration automates tiering, eliminating the need for manual tuning, extra software layers, or external storage systems.

VDURA is the software-defined data infrastructure platform that is purpose-built to power every stage of the AI pipeline. We combine the scalable, linear performance of a true parallel file system with the resilience and cost efficiency of object storage.

One data plane, one control plane, one namespace. Simple to deploy, operate, and grow.

The VDURA Data Platform

Software-Defined Architecture Engineered for Performance and Scalability

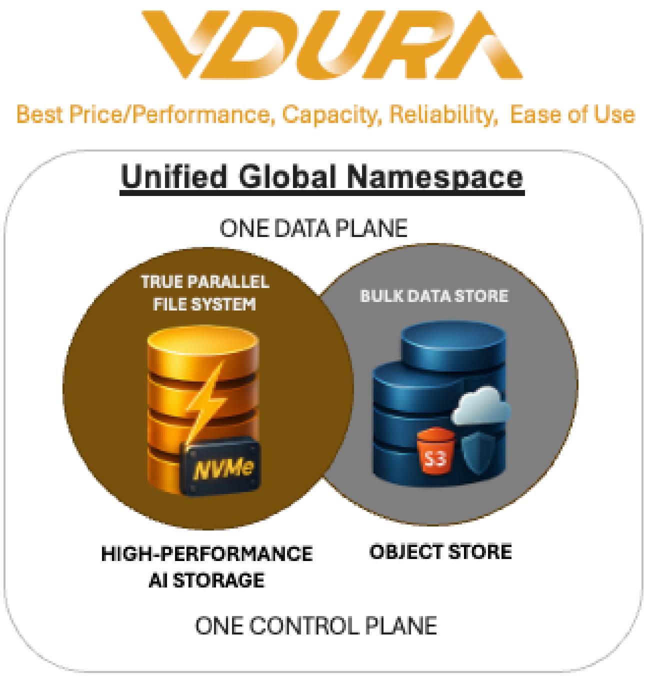

Figure 2. VDURA combines the best features of parallel file systems and object storage to power AI.

The VDURA Data Platform V11 is built on a fully software-defined, microservices architecture that combines the speed and efficiency of a true parallel file system with the durability and cost-effectiveness of resilient object storage.

This unified design ensures high performance and simplicity for active and bulk data storage and is designed specifically to address the complexities and requirements of the AI pipeline.

The VDURA Data Platform explicitly separates the control plane handling metadata operations from the data plane, which is dedicated exclusively to user data storage.

With relentless focus on performance, simplicity, and scalability, VDURA empowers organizations to push the boundaries of AI. The platform delivers data at the speed of innovation, turning infrastructure into a competitive advantage for those shaping the future.

Three key components work together to power the VDURA Data Platform:

- Director Nodes are the core of the control plane.
- Storage Nodes are the foundation of the data plane.
- The DirectFlow Client is our high-performance parallel file system driver.

Director Nodes are the core of the control plane. They orchestrate and manage all metadata operations, coordinate the actions of Storage Nodes and DirectFlow Client drivers for file access, maintain the health and membership status within the storage cluster, and oversee all recovery and reliability functions. These nodes are simple, powerful compute servers featuring high-speed networking, substantial DRAM, and NVMe flash optimized for metadata transaction logs.

The proprietary **VDURA VeLO metadata engine** runs on each Director Node. VeLO is a distributed, flash-optimized key-value store designed specifically for high-speed, parallel metadata operations. This integration ensures ultra-low latency, efficient handling of billions of file operations, and consistent metadata performance at scale.

Storage Nodes form the foundation of the data plane, dedicated exclusively to storing and managing user data. Available in configurations of either all-NVMe flash for peak performance or

NVMe flash with HDD capacity expansion for high-performance, economical bulk storage, Storage Nodes deliver versatile and optimized infrastructure.

Each node hosts multiple **Virtualized Protected Object Device (VPOD™)** instances, enabling granular, scalable data management and enhanced reliability through Multi-Level Erasure Coding. VPOD architecture ensures linear scalability and consistent parallel performance, accommodating thousands of nodes seamlessly within a single cluster. This flexible yet robust structure guarantees data integrity, high throughput, and balanced resource utilization, aligning performance and economics precisely with workload requirements.

The VDURA DirectFlow Client is a high-performance parallel file system driver specifically engineered for Linux-based compute environments. Deployed directly on compute servers, DirectFlow seamlessly integrates with existing Linux applications, presenting itself like any conventional file system. It provides fully POSIX-compliant, cache-coherent file operations across a unified global namespace, tightly collaborating with Director and Storage Nodes. By enabling direct, parallel I/O paths from compute servers to Storage Nodes, DirectFlow eliminates traditional bottlenecks and intermediary processing overhead found in NFS or legacy storage solutions. Its support spans all major Linux distributions and versions, ensuring frictionless adoption and integration within existing infrastructure.

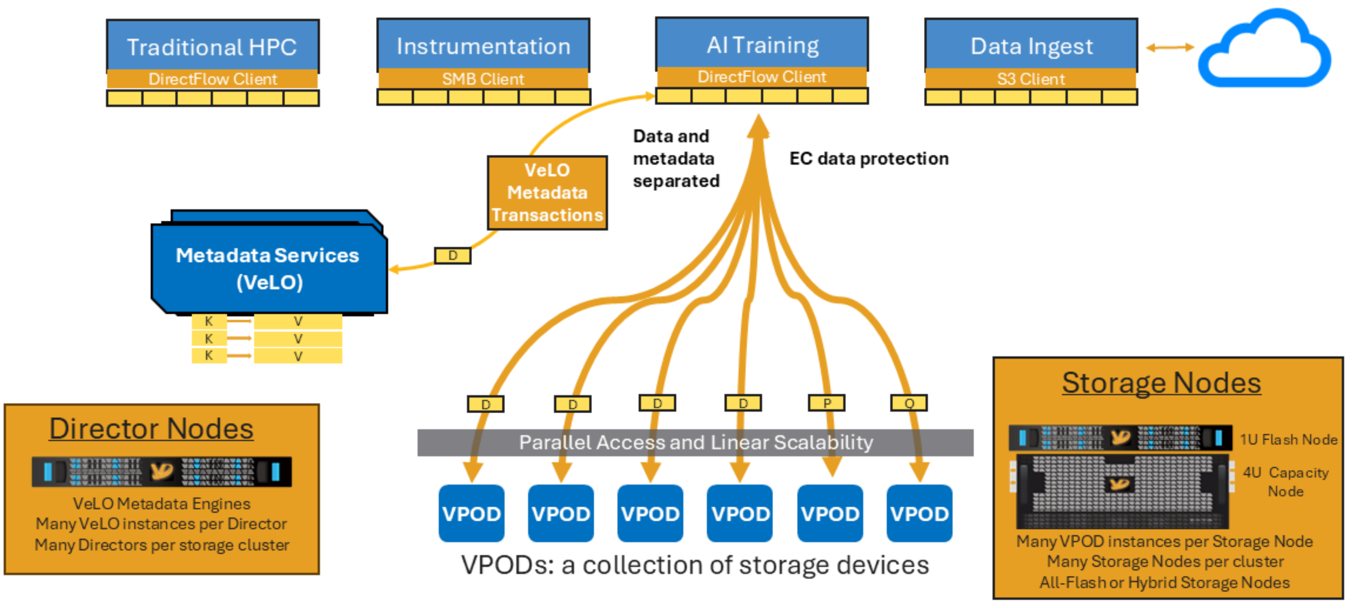

Figure 3. The VDURA Data Platform architecture.

Direct, Parallel Data Access for the AI Pipeline

The VDURA Data Platform is built as a true parallel file system, engineered to handle the intense I/O demands of modern AI and HPC workloads. Each file stored by the VDURA Data Platform is individually striped across many Storage Nodes, allowing each component piece of a file to be read and written in parallel, increasing the performance of accessing every file.

VDURA's parallel architecture dramatically accelerates data access, significantly boosting performance and throughput.

Unlike other enterprise systems which route data through limited head nodes, causing potential bottlenecks and requiring additional backend network infrastructure, VDURA's DirectFlow Client communicates directly with all relevant Storage Nodes. Each compute server directly accesses the nodes holding the data, bypassing intermediary bottlenecks. Director Nodes manage metadata and coordinate system activity out-of-band, ensuring efficient data flow without interference or congestion.

This direct and parallel design eliminates traditional NAS hotspots, ensures predictable and scalable performance, and simplifies infrastructure by removing the need for a separate, costly backend network. VDURA architecture delivers seamless scalability, consistently high performance, and exceptional efficiency across every stage of the AI pipeline, from ingest and training to inference and long-term data retention.

Linear Scalability, Seamless Expansion

The VDURA Data Platform delivers true linear scalability across both metadata and data services without compromise or complexity. AI workloads evolve fast, from early experimentation to scaled production across global clusters. Add Director Nodes to boost throughput for metadata-heavy tasks like model versioning and checkpoint tracking. Add Storage Nodes to scale bandwidth and capacity to support more training data, inference logs, or multi-tenant pipelines.

VDURA enables linear scalability, and growth is seamless and predictable. A 50 percent increase in Storage Nodes delivers 50 percent more throughput and capacity—no bottlenecks, no architectural redesigns.

This flexibility is powered by VDURA's fully virtualized, distributed software stack.

- **VeLO Metadata Engine** instances run in-memory across Director Nodes, scaling to billions of parallel metadata operations, perfect for high-frequency file creation, access pattern analysis, and rapid AI job cycles.
- **VPODs** manage user data in independently scalable units, each with its own erasure-coded stripe and logic, ideal for bursty checkpoint writes, long-term data lake retention, or active model training sets.

Together, these software-defined services form an intelligent, self-balancing system that grows effortlessly with the pipeline. Whether scaling GPU clusters, expanding training data, or retaining more checkpoints, VDURA adapts instantly.

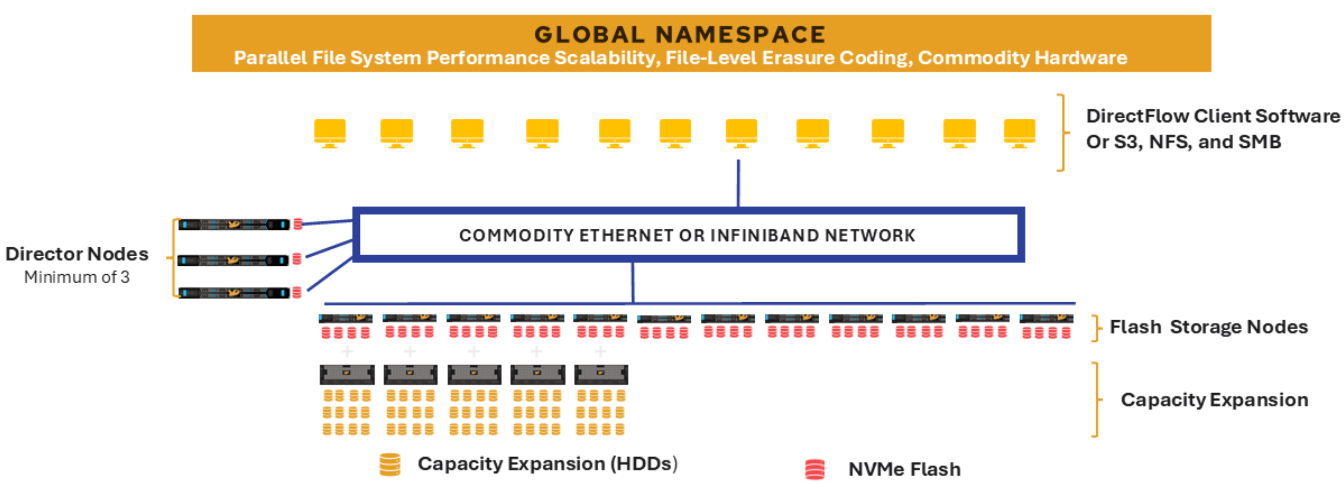

Figure 4. Linear scalability of Director and Storage Nodes.

Director Nodes: The Brain Orchestrating Control in a Parallel World

VDURA separates the **control plane**, which handles metadata, orchestration, and policy, from the **data plane**, which handles user I/O.

Director Nodes serve as the brain in the VDURA architecture. As the control plane's core, they command every stage of the AI pipeline, from ingestion and training to checkpointing and inference. Director Nodes continuously adapt to workload changes, ensuring optimal throughput and seamless orchestration across the system.

Each Director runs **VeLO**, a distributed, flash-optimized key-value metadata engine built to handle billions of operations per second. For modern AI, where performance is dictated as much by metadata velocity as data throughput, VeLO is essential. Tiny files, checkpoint indices, model versions—VeLO accelerates them all.

Director Nodes form the authoritative layer of VDURA's control structure. Every deployment requires a minimum of three. Administrators configure either three or five of the total Director Nodes as a *repset*, a voting quorum that maintains a synchronized, fully replicated configuration database. One node from the repset is elected *realm president* and is tasked with managing configuration, status monitoring, and leading failure recovery. If the current president fails, a new one is elected instantly and automatically.

Beyond coordination, Director Nodes also perform essential tasks at the president's request. These include managing volumes, serving as protocol gateways (NFS, SMB, S3), performing background data scrubbing, recovering failed Storage Nodes, and executing Active Capacity Balancing across VPODs. All changes are non-disruptive to clients; gateways and volumes can migrate transparently across nodes when necessary.

From petabyte-scale model staging to tiering outputs across flash and hybrid storage, Director Nodes orchestrate every aspect of the control plane. This intelligent coordination empowers the data plane to operate with high efficiency across every stage of the AI pipeline.

Storage Nodes: AI Pipeline Performance from Every Layer

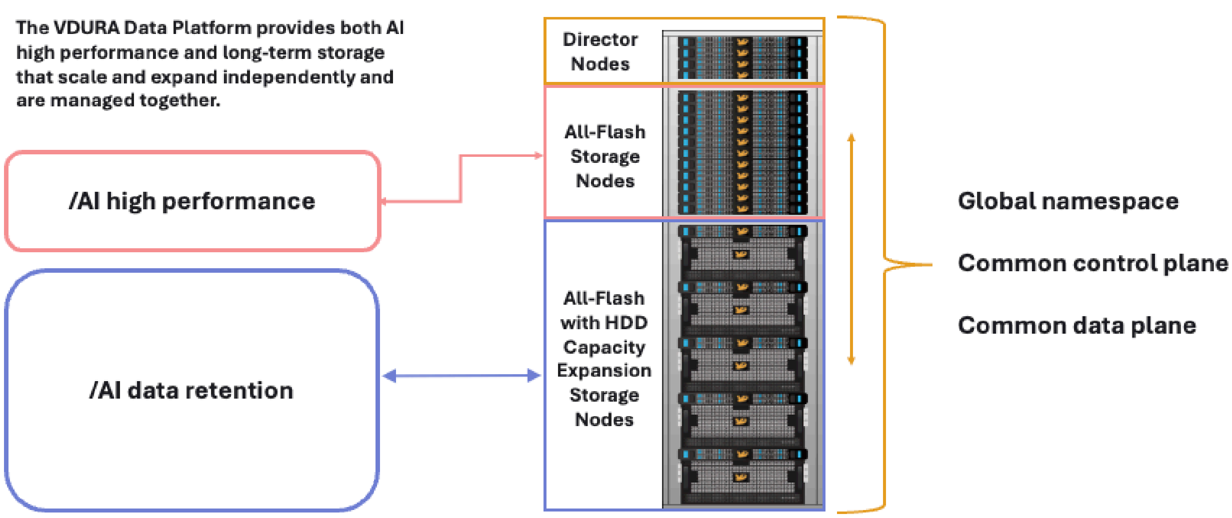

Figure 5. One control plane, one data plane, one single global namespace.

Storage Nodes are the backbone of VDURA's data plane, enabling seamless scale and sustained performance throughout every stage of the AI pipeline. Designed with flexibility and resilience, these nodes combine the best of both all-NVMe flash and flash with HDD capacity expansion storage, orchestrated under a unified control plane and single global namespace.

Optimized for Every Phase of the AI Pipeline

From high-frequency ingest and bursty checkpointing to real-time inference and long-term retraining, each phase of AI benefits from storage tiers purpose-built for performance and durability:

- **All-NVMe flash nodes** deliver ultra-low latency and high IOPs for AI high-performance data and latency-sensitive phases like model loading, active training, and checkpointing.
- **NVMe flash nodes with HDD capacity expansion** combine flash for metadata and active datasets with high-capacity HDDs for scalable, cost-efficient AI data retention. This is ideal for archived model weights, retraining inputs, inference logs, and data lakes that need fast access but are less frequently touched.

All Storage Nodes operate under a common data plane and global namespace, allowing applications to seamlessly span performance and capacity tiers without complexity or disruption.

VPOD Architecture: Virtualization for Resilience and Efficiency

Each VDURA Storage Node hosts multiple Virtualized Object Storage Devices, or **VPODs**, rather than treating the entire server as a single failure domain. This architecture introduces a finer unit of failure isolation:

- **The unit of failure is now one VPOD**, not the physical server.
- More VPODs per node increases operational granularity and cluster flexibility.
- **Device-level failures are isolated to a single VPOD**, eliminating the risk of full node failure.
- Storage reconstruction only affects the failed VPOD's component objects, not the entire node.

Files are striped across component objects in multiple VPODs using **N+2 erasure coding**, ensuring high fault tolerance with efficient space utilization. Large POSIX files benefit from this distributed protection model, while small POSIX files are triple replicated across VPODs, delivering optimal performance and storage efficiency.

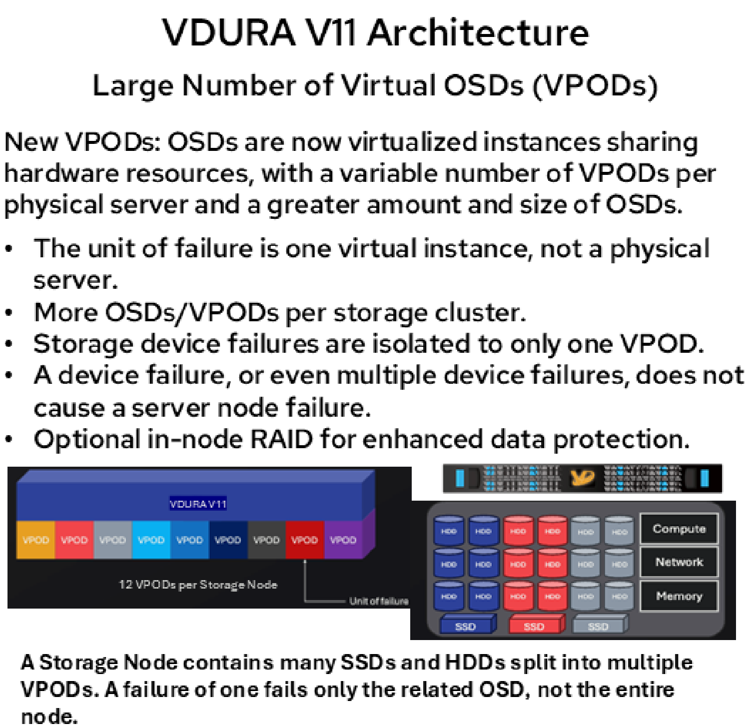

Figure 6. VDURA V11 architecture – Virtual Protected Object Devices (VPODs).

Dynamic Data Acceleration: The Smart Storage Fabric

VDURA **Dynamic Data Acceleration (DDA)** intelligently aligns I/O patterns with the most suitable media layer in real time:

- **Intent-log protection** powered by SSDs replaces legacy NVDIMMs for inflight data integrity.
- **Low-latency NVMe SSDs** store metadata databases for rapid namespace access.
- **High-IOPs SSDs** handle small file workloads.
- **High-bandwidth HDDs** manage large file sequential reads/writes.
- **System DRAM** provides caching for unmodified data and metadata.

Together, these layers form a high-performance, self-optimizing data fabric that minimizes latency and maximizes cost-efficiency.

Resilient Data Reconstruction and Integrity

In the event of a Storage Node failure, the VDURA Data Platform reconstructs only the affected component objects, not the full node's data. Files are rebuilt by pulling erasure-coded data fragments from other nodes. Continuous background scrubbing verifies data consistency across the system by validating erasure codes against stored data.

AI-Aware Placement and Automation

VDURA's intelligent orchestration engine continuously analyzes file size, access pattern, and data temperature to automate data placement across flash and hybrid tiers. Key features include:

- **Flash prioritization** for small and recently accessed files.
- **HDD capacity expansions** for large, sequential, or infrequently accessed data.
- **Continuous Active Capacity Balancing** to eliminate hotspots and evenly distribute load.
- **Real-time system adaptation** to shifting model training cycles, inference loads, and checkpoint bursts, requiring **zero manual tuning**.

The result is a storage system that evolves with the volatility of the AI pipeline, scaling performance and capacity without trade-offs.

VDURA V11 Data Reduction

VDURA V11 performs data reduction at the Storage Node level, ensuring zero impact on client-side CPU or memory resources. Unlike architectures that shift compression or deduplication tasks to the client, consuming valuable compute and memory, VDURA handles all reduction operations within the storage layer itself. This design keeps GPU and application nodes fully dedicated to AI and HPC workloads, maximizing performance and system efficiency.

The data reduction feature can be toggled on or off at any time via the graphical user interface (GUI) or command-line interface (CLI).

VDURA Data Platform V5000 System Overview

The VDURA Data Platform V5000 is an all-NVMe flash appliance engineered for AI and HPC pipelines that demand relentless GPU feed rates. Built on industry-standard servers, V5000 pairs flash performance with optional HDD capacity expansion, giving organizations a cost-balanced path from pilot to petabytes.

V5000 runs VDURA V11, VDURA’s flash-tuned parallel file system, streaming multiple terabytes per second from a single global namespace. Working with the DirectFlow Client, VDURA offers parallel redundant data paths that scale linearly, safeguard data with enterprise-class durability, and keep day-to-day management simple.

Each system begins with a minimum of **three Director Nodes** and **three Storage Nodes**, which can be either **all-flash** or **flash with HDD capacity expansion.** Additional nodes can be added seamlessly to expand performance, capacity, or metadata throughput independently.

V5000 has ultimate configurability from 100 percent NVMe Flash to 98 percent HDD.

Add all-flash Storage Nodes to boost throughput and IOPs or attach capacity expansion enclosures for cost-efficient bulk storage.

Figure 7. V5000 configuration options, left to right: all flash, 50 percent flash, 98 percent HDD.

Key Features

- 1U chassis for Director Nodes and Storage Nodes.
- 4U chassis for HDD capacity expansion.
- Configurable as either **all-flash NVMe** or **all-flash NVMe with HDD capacity expansion**.
- Up to **1.4 TB/s throughput per rack** in all-flash configurations.
- Up to **280 GB/s throughput per rack** in all-flash with HDD capacity expansion configurations.
- Up to **1.2 M IOPs per Storage Node**, supporting low-latency, high-frequency workloads. Up to **45 M IOPs per rack with all-flash configuration**.
- **Multi-Level Erasure Coding** with up to **12 nines durability (all-flash)** or **11 nines durability (all-flash with HDD capacity expansion)**.
- Supports InfiniBand™ (NDR/NDR200) and Ethernet (400/200/100 GbE) networking.
- Hardware-agnostic deployment with full support for commodity components.

V5000 Details

The VDURA V5000 system represents the culmination of decades of engineering expertise in parallel file systems and distributed storage technology. Built for AI/ML and HPC workloads, the V5000 combines enterprise-grade reliability with maximum throughput and flexibility. Its modular architecture allows organizations to independently scale performance, capacity, and metadata operations to create the ideal balance for their specific workload requirements without overprovisioning or underutilization. Each component, from Director Nodes to Storage Nodes, is engineered for maximum efficiency and resilience, with built-in redundancy and intelligent self-healing capabilities that ensure continuous operation even during component failures.

Director Nodes

- Host the **VeLO Metadata Engine**, a distributed key-value store optimized for flash.
- **Contains 12 x NVMe SSDs slots**.
- Up to 333 M inodes and up to 225 K creates/deletes per Director.
- Requires a **minimum of three nodes** for metadata triplication and fault tolerance.
- Scales out to hundreds of instances for greater metadata throughput.

All-Flash Storage Nodes

- Configurable with **12 x NVMe SSDs** in 7.68 TB, 15.36 TB, or 30.72 TB capacities.
- Delivers **>40 GB/s per node and up to 1.4 TB/s throughput per rack**.
- Supports up to **1.2 M IOPs per node**.
- Ideal for AI training pipelines, inference, and checkpointing workloads.
- Fully compatible with Multi-Level Erasure Coding and data reduction features.

All-Flash with HDD Capacity Expansion Storage Nodes

- Combined SSD tier up to 12 3.84 TB, 7.68 TB, 15.36 TB, or 30.72 TB.
- Supports JBOD expansion using **78 or 108 HDDs per node** with 16 TB, 24 TB, 30 TB drive options.
- Up to **280 GB/s** flash accelerated throughput.
- Scalable to **26 PBe per rack** (effective) with inline compression.
- Supports up to **1.2 M IOPs per node**.
- Optimized for high-capacity/data lake storage with flash-accelerated performance.
- Supports dynamic tiering and intelligent placement via the **VDURA Orchestration Engine**.

Connectivity

- **Supported Network Options**
  - InfiniBand: NDR/NDR200
  - Ethernet: 400/200/100 GbE ports
- **Connectivity per Node**
  - Up to 2x NDR200 InfiniBand
  - Up to 4x 100 GbE Ethernet

Expansion Options

Each VDURA V5000 cluster can expand incrementally and non-disruptively.

- Add **Director Nodes** to increase metadata and protocol performance.
- Add **All-Flash Nodes** for higher throughput and IOPs.
- Add **All-Flash with HDD Capacity Expansion Storage Nodes** for cost-efficient capacity growth.
- Mix and match node types within the same realm with **no architectural redesign**.
- Tiered performance and capacity levels managed via **Storage Sets**, ensuring isolation and quality-of-service (QoS).

Physical Connectivity

The VDURA V5000 Director and Storage Nodes support 400/200/100 Gigabit Ethernet (GbE) networks via two network ports in the rear of each node. The default configuration upon initial installation is link aggregation across two ports—a 2 x 200/100 GbE configuration using two 200/100 GbE SFP28 cables, with one attached to each port. The VDURA V5000 nodes support Link Aggregation Control Protocol (LACP) by default; static Link Aggregation Group (LAG), single link, and failover modes are also available.

VDURA V5000 Director and Storage Nodes contain two 25/10 GbE ports for corporate network connectivity.

All nodes also contain a single 1 GbE port that may be used as a general administrative network port or for troubleshooting.

Network Configuration Options

There are four network configuration options:

- Dynamic LACP
- Static LAG
- Single link
- Failover network

The default network configuration for V5000 nodes is LACP across the dual 100 GbE ports.

Generally, protocols other than LACP and static LAG operate in active/passive mode.

Active/Active Link Aggregation Mode

When load balancing is required to optimize performance, V5000 systems can be configured to use either dynamic LACP or static LAG. LACP is preferred, as it is significantly more robust than static LAG.

In static LAG mode, the physical ports are bonded with the IEEE 802.3ad static LAG link-layer protocol. This provides both load balancing and fault tolerance if a port loses its physical carrier status. Static LAG may fail to detect when the port stops functioning properly, but its carrier state will remain active.

In LACP mode, the physical ports are bonded with the IEEE 802.3ad LACP link-layer protocol. This provides load balancing, better fault tolerance, and protection against misconfiguration than static LAG.

Single Link Mode

While single link mode is supported on V5000 systems, it is not optimal since it is a single point of failure and suffers reduced bandwidth. Thus, single link mode should be used with caution, as loss of the one single link will make the node inaccessible.

Network Failover Mode

Network failover is used on V5000 systems when active/passive redundancy is required.

Storage Configuration Options

VDURA provides two mechanisms to manage namespace and capacity: Storage Sets and volumes.

Storage Sets

The **Storage Set** is a physical mechanism that groups Storage Nodes into a uniform storage pool. It is a collection of three or more Storage Nodes grouped together to store data. You can grow a Storage Set by adding more hardware, and you can move data within a Storage Set.

Volumes

A **volume** is a logical mechanism, a sub-tree of the overall system directory structure. A read-only top-level root volume ("/"), under which all other volumes are mounted, and a /home volume are created during setup. All other volumes are created by the user on a particular Storage Set, with up to 1,200 per realm.

A volume does not have a fixed space allocation but instead has an optional quota which can set a maximum size for the volume. Effectively, the volume is a flexible container for files and directories, and the Storage Set is a fixed size container for multiple volumes.

When planning volume configuration, keep the following points in mind:

- Volumes can be used to manage capacity.
- Volumes can be created to complement backup strategies.
- Performance can be enhanced by assigning volumes to different Directors.

RAID and Erasure Coding

Storage nodes in the VDURA Data Platform are highly sophisticated Virtualized Protected Object Devices (VPODs), and we gain the same scale-out and shared-nothing architectural benefits from our VPODs as any object store would.

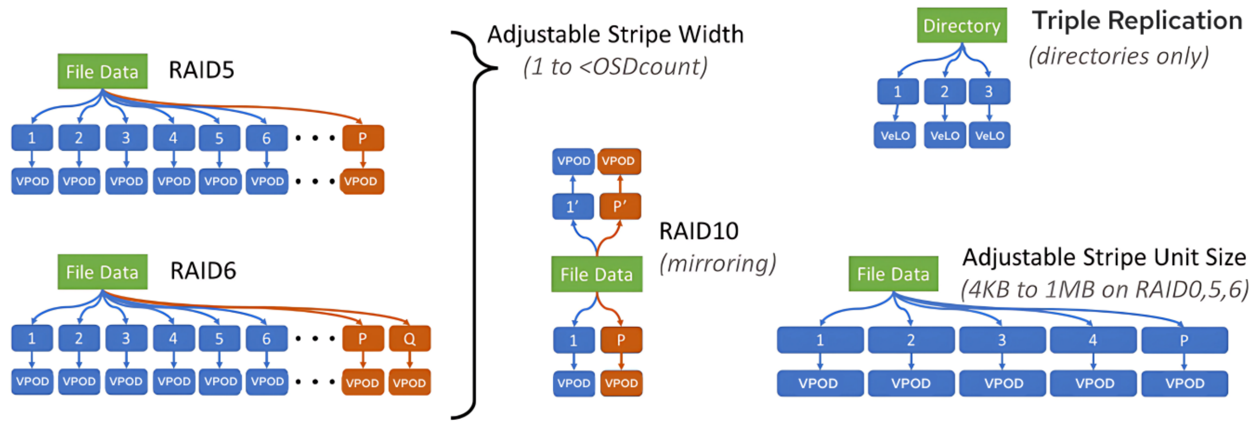

Figure 8. Per-file erasure coding layouts.

VDURA defines Objects used in our VPODs per the Small Computer System Interface (SCSI) standard definition of Objects rather than the Amazon S3 Object definition. The VDURA Data Platform uses SCSI Objects to store POSIX files, but it does so differently than how S3 Objects are typically used to store files. Instead of storing each file in an Object like S3 does, VDURA stripes a large POSIX file across a set of VPODs and adds additional VPODs into that stripe that store the P and Q data protection values of an N+2 erasure coding scheme. Using multiple VPODs per POSIX file enables the striping that is one of the sources of a parallel file system's performance.

A RAID array reconstructs the contents of drives, while VDURA reconstructs the contents of files.

While large POSIX files are stored using erasure coding across multiple VPODs, small POSIX files use triple-replication across three VPODs. This approach delivers higher performance than can be achieved by using erasure coding on such small files, while being more space efficient. Unless the first write to a file is a large one, it will start as a small file. If a small file grows into a large file, the Director Node will transparently transition the file to the erasure coded format at the point that the erasure coded format becomes more efficient.

When a file is created, and as it grows into a large file, the Director Node that manages those operations will randomly assign each of the individual VPODs that make up that file to different Storage Nodes. No two VPODs for any file will be in the same failure domain.

Reliability That Increases with Scale

Any system can experience failures, and as systems grow larger, their increasing complexity typically leads to lower overall reliability. For example, in an old-school RAID system, since the odds of any given HDD failing are roughly the same during the current hour as they were during the prior hour, more time in degraded mode equals higher odds of another HDD failing while the RAID system is still degraded. If enough HDDs were to be in a failed state at the same time, there would be data loss, so recovering back to full data protection levels as quickly as possible becomes the key aspect of any resiliency plan.

If a VDURA Storage Node fails, the system must reconstruct only those VPODs that were on the failed Storage Node, not the entire raw capacity of the Storage Node like a RAID array would. The system would read the VPODs for each affected file from all the other Storage Nodes and use each file's erasure code to reconstruct the VPODs that were on the failed node.

The VDURA Data Platform has linear scale-out reconstruction performance that dramatically reduces recovery time in the event of a storage node failure, so reliability increases with scale.

When a Storage Set is first set up, it sets aside a configurable amount of spare space on all the Storage Nodes in that Storage Set to hold the output from file reconstructions. When the system reconstructs a missing VPOD, it writes it to the spare space on a randomly chosen storage node in the same Storage Set. As a result, during a reconstruction, the system uses the combined write bandwidth of all the storage nodes in that Storage Set. The increased reconstruction bandwidth results in reducing the total time to reconstruct affected files, which reduces the odds of an additional failure during that time and increases the overall reliability of the realm.

VDURA also continuously scrubs the data integrity of the system in the background by slowly reading through all files in the system, validating that the erasure codes for each file match the data in that file. Data scrubbing is a hallmark of enterprise-class storage systems and is only found in one HPC-class storage system, the VDURA Data Platform.

An Architecture of High Reliability

Based on system configuration, the N+2 erasure coding that VDURA implements protects against either one or two simultaneous failures within any given Storage Set without any data loss. The realm can automatically and transparently recover from more than two failures, as long as there are no more than two failed Storage Nodes at any one time in a Storage Set.

If, in extreme circumstances, three Storage Nodes in a single Storage Set were to fail at the same time, VDURA has one additional line of defense that would limit the effects of that failure. All directories are independently stored triplicated—three complete copies of each directory, with no two copies on the same Director Node.

If a third Storage Node were to fail in a Storage Set while two others were being reconstructed, that Storage Set would immediately transition to read-only state. Only the files in the Storage Set that had VPODs on all three of the failed storage nodes would have lost data. All other files in the Storage Set would be unaffected or recoverable using their erasure coding. The number of affected files in these situations becomes smaller as the size of the Storage Set increases.

Since the system will have independent metadata storage that can survive against two simultaneous failures, it can identify the full pathnames of precisely which files need to be restored from a backup or reacquired from their original source and can therefore also recognize which files were either unaffected or recovered using their erasure coding.

VDURA is unique in the way it provides clear knowledge of the impact of a given event, as opposed to other architectures which leave you with significant uncertainty about the extent of the data loss.

Per-File RAID

Instead of relying on hardware RAID controllers that protect data at a drive level and computing RAID on the disks themselves, VDURA architecture uses per-file distributed RAID in software using erasure codes, i.e., per-file erasure coding rather than hardware RAID. Files in the same Storage Set, volume, and even directory can have different RAID/erasure code levels. A file can be seen as a single virtual object that is sliced into multiple component objects.

Users have three RAID levels available: RAID 6, RAID 5, and RAID 10.

- RAID 6 is the default RAID level for all volumes, as it provides a balance between performance, capacity overhead, and RAID protection.
- RAID 5 volumes are optimized for performance and capacity overhead.
- RAID 10 volumes combine striping and mirroring to provide high performance for applications that perform small random writes, while at the same time providing resiliency against a single disk failure.

You can mix RAID levels and any volume layout together in the same Storage Set, with each volume evaluated independently for availability status.

Comprehensive Data Services

VDURA's commitment to data management extends beyond performance and scalability. The platform offers a full suite of data services designed to ensure business continuity, protect data, and streamline management operations.

Snapshots, quotas and data migration are all standard features of VDURA, providing enterprises with the tools they need to manage their data effectively.

Tiered storage is another key feature of VDURA, with the platform's Dynamic Data Acceleration (DDA) technology supporting up to four different performance tiers. This ensures that data is placed on the most appropriate storage media based on its characteristics, optimizing both cost and performance.

Security is woven into every aspect of VDURA, with industry-leading AES-256 encryption for data in flight and at rest, comprehensive access controls, and detailed logging and auditing capabilities. End-to-end encryption provides stronger data confidentiality, integrity, and compliance alignment, surpassing traditional TLS + SED approaches.

A Competitive Edge

In the competitive landscape of high-performance data infrastructure, VDURA distinguishes itself through a combination of high performance, durability, ease of use, and superior scalability.

VDURA is often described as the "easy button" for high-performance storage. Its intuitive user interface and robust management capabilities make it accessible to organizations that may not have extensive in-house storage expertise. Ease of use, combined with VDURA's advanced features, provides a compelling value proposition for enterprises looking to simplify storage operations without sacrificing performance.

VDURA's commitment to continuous innovation ensures that it remains at the cutting edge of the industry. The platform's transition from a hardware-centric to a software-focused model has accelerated innovation, enabling VDURA to deliver new features and enhancements at a rapid pace.

Where Velocity Meets Durability

VDURA's expertise stems from decades of innovation, building on the industry's first parallel file system. This heritage informs a platform that simplifies management, reduces complexity, and transforms storage into a driver of AI success. Scalable from terabytes to exabytes, VDURA supports every stage of the AI pipeline.

With a relentless focus on performance, simplicity, and scalability, VDURA empowers organizations to push the boundaries of AI. The platform delivers data at the speed of innovation, turning infrastructure into a competitive advantage for those shaping the future.