VDURA Data Platform

The Industry Leading Parallel File System

The VDURA Data Platform, now in version V11, is an advanced data platform optimized for HPC, AI/ML workloads, and scientific computing. It delivers exceptional scalability, reliability, and simplified management.

VDURA Data Platform Architecture:

- Combines All-Flash and hybrid nodes for flexible, high-performance data storage, with dedicated director nodes for metadata and storage nodes for data processing.

DirectFlow Client Driver:

- Parallel, POSIX-compliant, and cache-coherent file system access, eliminating performance bottlenecks.

NFS and SMB/CIFS Gateway:

- Standard protocol support allows easy integration, with peak performance via DirectFlow.

Dynamic Data Acceleration (DDA):

- Dynamically optimizes storage performance across hybrid and flash nodes without manual tuning, ideal for mixed workloads and varying file sizes.

Security:

- Advanced data protection through software-based distributed erasure coding.
  - Secure encryption using self-encrypting drives (SEDs) and KMIP-compatible key management.

Reliability & Scalability:

- Combines self-managed, self-healing architecture with both hybrid and all-flash storage nodes, delivering enhanced reliability and unlimited scalability.

Simplicity of Management:

- Automated and centralized management, including snapshots, quota enforcement, and comprehensive reporting, reducing administrative complexity.

VDURA Data Platform Architecture

Three fundamental components comprise the VDURA Data Platform parallel file system: data storage nodes, metadata director nodes, and the DirectFlow® driver on client systems. As shown in Figure 1, the three elements work together to power the VDURA parallel file system. Director and storage nodes are commodity compute servers with high-speed networking connections dedicated to running the VDURA Data Platform software. File system workloads are divided between director nodes that process metadata and storage nodes that process data.

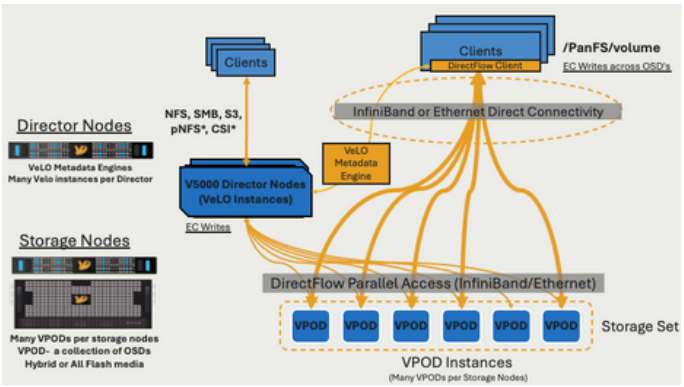

Figure 1. The VDURA Data Platform Parallel File System.

DirectFlow Client Driver

VDURA DirectFlow client driver is a loadable file system implementation that can be installed on Linux client systems and used by application programs like any other file system. DirectFlow clients work with director nodes and storage nodes to deliver fully POSIX-compliant and cache-coherent file system behavior, from a single namespace, across all the servers in the compute cluster. For each file an application accesses, the DirectFlow client reads and writes directly in parallel to all storage nodes that hold that file's data. The DirectFlow client is the key to eliminating bottlenecks, application scalability issues, and erratic performance common to scale-out NAS systems.

NFS and SMB/CIFS Gateway

The VDURA Data Platform provides high-performance NFSv3 and SMB/CIFS v3 access through director nodes that act as gateways translating NFS and SMB/CIFS operations into DirectFlow operations. This ability allows clients such as laptops and workstations to access the same namespace and data. But, as it is parallel and direct, the DirectFlow protocol will always be the highest performance path to storage.

Unlimited Scalability

For more capacity or storage performance, just add more storage nodes. For more metadata processing performance, add more director nodes. In VDURA scale-out storage systems, there simply is no maximum performance or maximum capacity.

Reliability that Increases with Scale

Any system can experience failure and, as systems grow larger, those failures can become more frequent and severe. However, VDURA's reliability actually increases with scale due to its linear scale-out reconstruction time improvement in the event of a storage node failure.

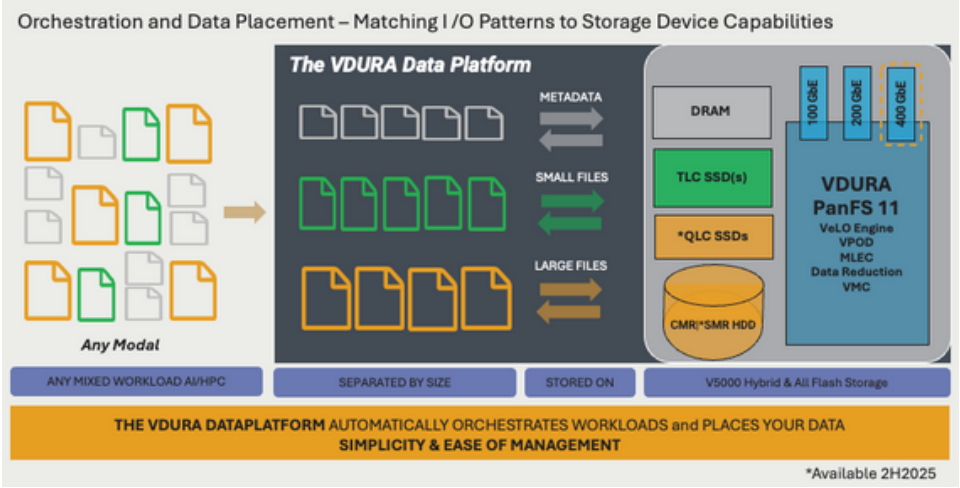

Figure 2. VDURA Data Platform Dynamic Data Acceleration.

Dynamic Data Acceleration and Mixed Workloads

VDURA's Dynamic Data Acceleration (DDA) technology takes the complexity out of tiered high-performance storage systems by maximizing the efficiency of all storage media in a seamless, all-hot system that matches I/O patterns to storage device capabilities. DDA automatically adapts to changing file sizes and mixed workloads without the need for tuning or manual intervention.

VDURA's architecture delivers exceptionally high and consistent performance for workloads that include a wide range of file sizes, access patterns, and changes in volume over time. The result is a dynamic performance environment compared to other parallel file systems that require time-consuming and laborious tuning and retuning as workloads change.

Enterprise-Grade Data Protection and Availability

The VDURA Data Platform uses software-based, network-distributed erasure codes to separately protect individual files, rather than traditional RAID groups that only offer drive-level protection. This eliminates the multiple risks, costs, and performance penalties of traditional RAID architectures.

Multiple Levels of Security

To prevent unauthorized data access while the VDURA Data Platform is online, VDURA supports file system ACLs and SELinux security labels. VDURA also offers hardware-based encryption-at-rest with zero performance impact to prevent unauthorized data access while VDURA is offline by using industry-standard self-encrypting drives (SEDs). VDURA is KMIP compliant, allowing customers to use proven KMIP security key management servers, including a tested and validated integrated solution with industry leading CipherTrust Manager from Thales.

Surprising Simplicity, Streamlined Management

As scale-up storage architectures grow in complexity, manageability often suffers. The VDURA Data Platform file system architecture simplifies management by offering a single global namespace that easily scales. VDURA automates key workflows such as new storage discovery and load balancing to streamline performance, and offers enterprise data services such as re- porting, snapshots, and user quota enforcement. In addition, automated capacity balancing and centralized management features streamline management as storage volumes grow.

More Information

For more information about the VDURA Data Platform, contact your local VDURA representative at marketing@vdura.com.

About VDURA

VDURA is a leader in AI and HPC data infrastructure, delivering high-performance, scalable storage solutions that power the world's most demanding AI applications.
For more information, visit www.vdura.com.

Worldwide Office
1-888-726-2727
Info@vdura.com

VDURA Headquarters
Milpitas, CA, USA

VDURA Research & Development
Pittsburgh, PA, USA

VDURA EMEA
Oxford, United Kingdom
emeainfo@vdura.com

VDURA APAC
Sydney, Australia
apacinfo@vdura.com

VDURA China
Beijing, China
chinainfo@vdura.com