Tyrone

VELOX

TYRONE PARALLELSTOR VELOX

PERFORMANCE AND SIMPLICITY FOR THE HYBRID CLOUD

SOLUTION OVERVIEW

Tyrone ParallelStor Velox is a software-defined, high-performance data storage solution that helps enterprises manage explosive data growth with ease. It provides a parallel, scalable file system that consolidates unstructured data across on-premises and cloud environments into a single, unified platform. Designed to meet the needs of modern big data and AI applications as well as traditional workloads, ParallelStor Velox delivers extreme performance, reliability, and security on a scale.

Organizations can eliminate data silos and support both new-era analytics and legacy applications on one platform. The result is faster insights and streamlined operations – those who leverage their data more quickly gain a competitive edge in their industry. Backed by IBM's decades of expertise, ParallelStor Velox is a proven enterprise-grade solution used in mission-critical environments worldwide, offering the robustness and support that IT decision-makers demand.

KEY FEATURES AND CAPABILITIES

Unified Global Namespace:

Consolidate storage silos across datacenters and clouds into one global namespace, simplifying management and enabling seamless data access everywhere.

Global Collaboration and Multi-site Access

Built-in Active File Management (AFM) technology lets distributed teams access and share data globally with local speed reads and writes. It creates a single synchronized dataset across geographies, masking WAN latency and outages.

Policy-driven Tiering & Archiving:

Intelligent lifecycle management automatically moves data to the optimal storage tier – hot data to flash, cold data to object storage or tape – cutting data retention costs by up to 90% through automation.

Multi-protocol Data Access

ParallelStor Velox supports POSIX files, NFS/SMB, and object (S3/Swift) interfaces simultaneously on the same data. Applications can access a common data lake using the protocol they need without creating duplicate copies.

Scale out Performance Architecture:

A truly parallel, scale-out design with distributed metadata and no single controller means performance grows linearly as you add resources.

Enterprise Security & Compliance

The solution has built-in encryption, authentication, and role-based access control to protect sensitive data. It supports compliance requirements (e.g. HIPAA, GDPR) with immutable file options and audit logging.

Container and Cloud Integration:

Deployable in VMs, containers, or on bare metal, ParallelStor Velox is fully containerized for modern cloud-native environments. It can orchestrate data for Kubernetes and Red Hat OpenShift, making it ideal for AI/ML and microservices use cases.

Advanced Data Services:

Comprehensive features like snapshots, backup/restore, synchronous and asynchronous replication, and geo-distributed erasure coding are included. An intuitive GUI and REST APIs simplify administration, while integration with IBM Cloud Pak® and other ecosystem tools enables automated data workflows and AI-driven insights.

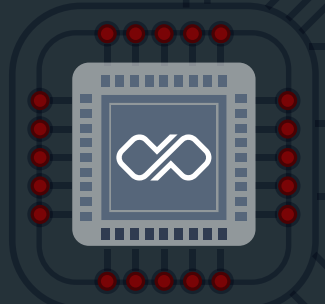

PERFORMANCE AND SCALABILITY HIGHLIGHTS

Scales to Global Dimension:
A ParallelStor Velox cluster can scale to thousands of nodes (16,000+ nodes in a unified cluster) and manage virtually unlimited data under one filesystem. It efficiently handles billions of files and many petabytes – even into the yottabyte range – without breaking a sweat.

Blazing Throughput and IOPS:
In modern all-flash deployments, ParallelStor Velox has demonstrated industry-leading speed – for example, a single 4U ParallelStor Velox-based storage node can sustain 330 GB/s read and 155 GB/s write throughput, and over 13 million IOPS when leveraging NVMe fabrics.

24*7 Consistency at Scale:
Even at multi-petabyte scale, ParallelStor Velox maintains reliability and consistency. Its distributed recovery and replication mechanisms can handle node or disk failures transparently, so performance remains steady under failure conditions – a critical factor in sustaining performance SLAs.

Note: *In summary, ParallelStor Velox combines extreme throughput, IOPS, and horizontal scalability to support the most data-intensive workloads on the planet. It provides performance headroom for future growth, ensuring that your storage will not be the bottleneck in your IT environment.*

WHY BUSINESSES CHOOSE PARALLELSTOR VELOX?

Proven at Enterprise Scale:

With thousands of deployments worldwide and over two decades of pedigree, ParallelStor Velox is a field-tested platform trusted for critical workloads. Its IBM backing ensures world-class support and continuous innovation.

Unmatched Performance and Scalability:

Businesses can deliver insights faster and meet growth demands without redesigning storage architectures. ParallelStor Velox handles petabytes of data and billions of files with ease, so IT can scale infrastructure to support business expansion or new projects.

Operational Efficiency & Cost-Savings:

By consolidating data and automating tiering, organizations achieve greater efficiency and lower TCO. ParallelStor Velox helps streamline data workflows and reduce storage costs (e.g. saving on expensive primary storage by auto-archiving cold data).

Improved Service and Collaboration:

ParallelStor Velox improves application uptime and responsiveness, leading to better service delivery. Global teams gain secure, fast access to shared data from anywhere, boosting productivity and collaboration.

Robust Resiliency and Security:

Businesses need to manage risk and protect data – ParallelStor Velox's no-single-point-of-failure design and integrated data protection mean continuous operations even in face of node or site failures.

Future-Ready Flexibility:

ParallelStor Velox supports modern and legacy workloads equally well. It can host AI training data lakes, Hadoop analytics, virtualization storage, backup archives – all on one platform.

COMPETITIVE DIFFERENTIATORS

True Scale Out Architecture (No Single Bottleneck):

Unlike traditional NAS systems (NFS, CIFS, etc.) that funnel data through a single controller or server, ParallelStor Velox is fully distributed. There is no single choke point – multiple servers concurrently serve data and share the load.

Integrated All-in-one Solution:

ParallelStor Velox provides a complete storage software stack in one product – it includes its own cluster management, volume management, and advanced features out-of-the-box. Competing cluster file systems often require multiple add-on software packages (external cluster managers, separate file systems, etc.), adding complexity and cost.

Multi-protocol and Heterogenous Support:

Many alternatives are siloed – for example, one system for files, another for objects, others limited to Linux only. ParallelStor Velox is one platform for all data, supporting POSIX, Windows, object, and Hadoop workloads together. It runs on heterogeneous infrastructure (x86, Linux, Windows), offering greater flexibility.

Extreme Performance + Reliability:

Competing storage offerings often force a trade-off between performance and resiliency on a scale. ParallelStor Velox delivers both: its parallel I/O with client-side caching accelerates throughput for demanding applications, while built-in features like failover clustering, replication, and encryption ensure enterprise-grade reliability and data integrity.

Advanced Data Management & Analytics Integration:

Tyrone ParallelStor Velox can perform in-place analytics on stored data without needing costly ETL duplication. This “analytics in place” capability – along with IBM’s ecosystem is a key advantage over basic storage systems that merely store and retrieve data.

Enterprise Support & Ecosystem:

We deliver a system with rigorously engineered and validated key integrations. By leveraging the existing NVIDIA GPU Direct Storage certification from the IBM core, our system is immediately production-ready for the most demanding AI/ML, HPC, and data-intensive workflows.

TYPICAL USE CASES ACROSS INDUSTRY

Financial Services:
High-frequency trading, risk analytics, and fraud detection systems rely on ParallelStor Velox for ultra-fast, concurrent access to large datasets (market data, risk models). Banks also use it to build secure data lakes for compliance and regulatory archives.

Telecommunications & Media:
ParallelStor Velox provides the throughput needed for real-time analytics on data deluge in telco. In media and entertainment, studios use it to store and stream 4K/8K video content and to support render farms, thanks to its ability to serve huge files to many editors or rendering nodes in parallel without lag.

Healthcare & Lifesciences:
Hospitals and research institutions use ParallelStor Velox to manage big data from genome sequencing, medical imaging (e.g. MRI, CT scans), and electronic health records. It can securely store petabytes of sensitive patient data with encryption and replicate it across sites for disaster recovery.

Public Sector and Research HPC:
Government labs and universities running simulations for weather forecasting, climate research, defense, or astrophysics need a storage backbone that can keep up with their supercomputers. ParallelStor Velox's parallel file system is ideal for HPC clusters – it has been a foundation in many Top500 supercomputer sites.

AI & Big Data Analytics:
For AI/ML model training and big data pipelines, ParallelStor Velox provides a unified data hub. It can feed GPUs and big data frameworks (Spark, Hadoop) with high-throughput access to training datasets, eliminating IO bottlenecks that slow down AI jobs. Companies like autonomous vehicle firms and internet giants use ParallelStor Velox to store exabytes of data (e.g. self-driving car sensor data, clickstreams) and tier cold data to cheaper storage transparently.

Hybrid Cloud and Multi-Cloud Deployments:
Modern enterprises often need to span on-prem and cloud environments. ParallelStor Velox is used to create a hybrid cloud data fabric – for example, extending on-prem data into public cloud for burst processing, or consolidating data from edge sites to a central cloud repository. Its support for cloud-object storage and containerized deployment makes it a backbone for multi-cloud strategies.

Note: *These are just a few examples – ParallelStor Velox's flexibility means it can support everything from manufacturing IoT platforms and automotive design simulations to media archives and academic supercomputers. Across all industries, any scenario requiring fast, scalable, and reliable access to large volumes of data is a natural fit.*

VELOX

Ready to transform your enterprise storage?

Experience the power of Tyrone ParallelStor Velox for yourself.

Contact us today for a consultation or demo and see how ParallelStor Velox can unlock the full potential of your data – from accelerating AI initiatives to simplifying your hybrid cloud strategy. Let's scale your business to new heights, together.

Tyrone

VELOX

E-mail: info@netwebindia.com
Call: 1800-419-0688

facebook.com/tyronesystems
x.com/tyronesystems
linkedin.com/company/tyrone-systems
www.tyronesystems.com