STEELDOME

HyperSERV

The Future of Data Infrastructure

StratiSTOR and StratiSERV unite to form HyperSERV

HyperSERV provides a comprehensive data platform which seamlessly integrates all the essential storage, network, and compute capabilities needed to run workloads at any scale. What sets it apart is its unmatched flexibility and total control over any hardware asset from any vendor, regardless of the deployment model chosen.

Hardware-Agnostic Flexibility

HyperSERV is compatible with any server or storage hardware, empowering customers to leverage existing investments and integrate new technologies without costly overhauls. As a hardware-agnostic HCI solution, it enables complete control over both performance and costs.

High-Performance

By integrating compute, storage, and networking into a unified system, HyperSERV ensures that applications can directly access the data they need without unnecessary latency or delays. This results in faster processing and improved overall performance, especially for demanding workloads such as AI, databases, and high-transaction environments.

AI-Ready

HyperSERV supports all the latest NVIDIA hardware for AI, allowing you to scale across any combination of hardware with a software-defined, hardware-agnostic solution. We will always support the latest and greatest hardware available today. Our platform includes all the necessary frameworks to deploy AI and get it up and running quickly.

Simplified Management and Reduced Costs

Control both the virtualization layer and our storage cluster through one interface, with support for both virtual machines and containers. HyperSERV automates key tasks like data protection, scaling, and self-healing, which reduces the complexity of managing storage at scale. This automation cuts operational costs.

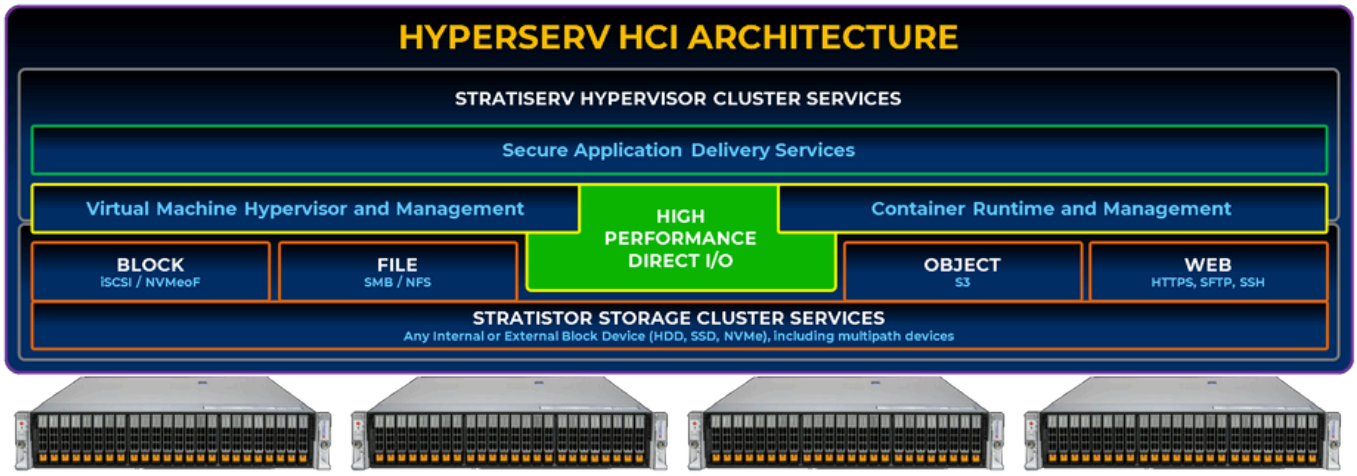

HYPERSERV HCI architecture integrating STRATISERV hpervisor cluster services with STRATISTOR storage cluster services. Combines high-performance virtual machine and container management, secure application delivery, and direct I/O with unified access to block (iSCSI/NVMeoF), file (SMB/NFS), object (S3), and web protocols (HTTPS, SFTP, SSH). Optimized for scalable, high-throughput enterprise environments.

Distributed, Scale-Out Architecture

HyperSERV eliminates single points of failure with a fault-tolerant design that scales from a few to hundreds of nodes—supporting both vertical and horizontal growth

Metadata/Data Separation

Decouples metadata from data paths to maintain predictable, low-latency performance at scale—critical for real-time and high-throughput applications.

Next-Gen I/O Acceleration

SDCache transforms any disk into a high-performance device, reducing latency and boosting throughput. HyperSERV's parallel design enables concurrent access for demanding workloads like virtualization, analytics, and ML.

Advanced Data Protection

Flexible erasure coding and replication ensure resilience and high availability, with built-in integrity checks for fault tolerance.

Multi-Protocol & High-Speed Networking

Supports iSCSI, NVMeoF, SMB, NFS, and S3, with seamless integration into RoCE, InfiniBand, and NVIDIA BlueField environments—ideal for GPU-accelerated and AI workloads.

Self-Healing & Automated Recovery

Continuously monitors and repairs faults without manual intervention—ensuring maximum uptime in mission-critical environments.

Deployments

Single-Node

Ideal for organizations needing resilient, high-density block, file, or object storage—without node-level high availability. Perfect for development environments, ROBO, or edge computing, this model enables rapid deployment and easy management with minimal infrastructure, making it a cost-effective choice for new projects.

Multi-Node Cluster

Best suited for mid-sized businesses or departmental data centers needing high performance and availability. This model distributes data across multiple nodes for load balancing and failover, supporting diverse workloads like virtualization, large-scale file sharing, and data-intensive applications. It scales seamlessly as demand grows.

Large-Scale Clusters

Built for large-scale enterprise operations that demand ultra-high performance and scalability. Ideal for cloud providers, research institutions, and data-intensive industries like media, finance, and healthcare. This configuration supports complex workloads across extensive storage and compute networks with consistent performance and reliability.

Regional Federated Super Clusters

Ideal for large enterprises unifying IT infrastructure across multiple sites. A federated HyperSERV super-cluster spans data centers, regions, or continents—creating a single, manageable storage ecosystem. It enables data replication, cross-site redundancy, and resource sharing, future-proofing your infrastructure for global operations and next-gen workloads.