

Datasheet

# Virtual Storage Platform One Object

**Resilient, supercharged, sustainable, and cost-optimized object storage built for the demands of modern data-driven innovation**

**A high-performance and flexible solution with built-in data smarts for intelligent management, AI and analytics.**

In an era where data-driven success is crucial, fueling AI and business growth is paramount for businesses. Organizations need scalable, secure solutions to manage and harness value from exponential unstructured data growth. Hitachi's Virtual Storage Platform One (VSP One) Object is meticulously designed to manage the overwhelming volume and growth of an ever increasing variety of data and process it at increasing speeds due to new cloud-native, analytics, and AI applications. It ensures quality and security while representing data in a way that is easy to understand, visualize and interpret, or to feed it into other applications. Ensuring data validity is critical for adherence to legal and compliance regulations, accurate AI and analytics, and safeguarding against data breaches.

**Unbreakable for Always-On Businesses**

Our next generation object storage solutions offer Hitachi's trusted and robust cybersecurity measures, while providing flexible scaling by supporting converged and disaggregated storage configurations with NVMe, HDD, and SSD options. The need for 24/7 business continuity and always-on operations is more important than ever. VSP One Object offers data resilience and 100% data availability, guaranteeing uninterrupted service and protecting critical data against loss. It ensures security with durability, cyber resilience, and built-in data compliance and governance controls. It includes FIPS-certified encryption and integration with LDAP and Active Directory for seamless authentication. Asynchronous replication supports long distance replication for enhanced data availability, disaster recovery, and business continuity. Furthermore, it ensures consistent and accurate metadata management across big data systems for reliable data access.

*VSP One Object gives you the flexibility to choose the right storage option, offering access to S-nodes, local disks, and VSP One Block to adapt to diverse workload demands and data value. This is a solution built to match varying requirements, empowering organizations to store and manage data based on its unique workload requirements and importance. The platform ensures high availability through redundancy across multiple nodes, maintaining reliability even for deep archival needs. This approach not only optimizes performance but also enables smarter, value-driven data storage strategies.*

## Virtual Storage Platform One Object

### Robust Data Protection

VSP One Object delivers the latest, differentiated features to ensure data is securely stored and protected, with the ability to scale as needed. These include:

- **High-Density and Efficient:** 4U appliance supporting large object storage with 750TB NVMe and 2.5 billion objects per node.
- **Scalability:** 5-node cluster ensuring the system can grow with your data needs.
- **Resilience:** Tolerates 2-node failures, provides robust data protection and system reliability. The system supports a configurable resiliency factor with an out-of-the-box 3+2 EC schema.
- **Intelligent Data Services:** Multiple data protection options, including async-replication for data availability, disaster recovery, and business continuity.
- **S3 Object Lock:** Prevents object deletion for a specified retention period, ensuring compliance with regulatory requirements.
- **Legal Hold:** Preserves data for legal purposes, ensuring that important information is retained and protected from deletion.
- **External KMS Support:** Integrates with external key management systems for enhanced security.

### Performant, Future-ready Unified and Multicloud Solutions

Customers need solutions that leverage AI to enhance infrastructure and accelerate time to value for AI innovation and end-user outcomes. VSP One Object is engineered to deliver rapid data retrieval with high IOPS and low latency, making it perfect for data lakehouse environments. It provides modern, high-performance data storage with resiliency, scalability, and SLAs, ensuring your data is always accessible and ready to maximize value.

### TIER 1 S3-Native and AI/ML Workloads

Hitachi Vantara's common data platform seamlessly combines the power of on-premises infrastructure with cloud experience and scale. As automation advances digital transformation, massive volumes of unstructured data are generated continuously, especially with new data sources. VSP One Object is inherently scalable, accommodating immense amounts of data without needing re-architecture or migration. Its ability to handle petabytes of data makes it perfect for developing customer-facing applications and deployment.

### Unified Data Management Architecture

Strategic initiatives often require a mix of structured, semi-structured, and unstructured data, which traditional storage architectures may not efficiently support. VSP One Object is designed to handle unstructured data and can accommodate various data formats (e.g., images, videos, logs), making it well-suited for environments that need to unify disparate data types.

### Empowering Your Business with Data

Hitachi understands that organizations are at various stages of embracing new technology advancements and modernizing their information management practices. This includes cloud, big data, mobile, software-defined, and more. Organizations are also dealing with a growing number of technologies, devices, access points, business requirements, and types of data to be stored.

### Transforming Data Management

Hitachi aims to address enterprise data challenges through advancements in data management features for object storage. These features, offered by VSP One Object, include creating custom

data sets, categorizing data by specific attributes, discovering Personally Identifiable Information (PII) in objects, and synchronizing data across multiple platforms. VSP One Object enables robust data management for corporate data, ensuring it remains accessible and secure throughout its lifecycle.

As data continues to grow, we must access and manage new endpoints and formats, such as machine-to-machine communication and the Internet of Things (IoT). A piecemeal approach to data endpoints won't suffice. Instead, we need a unified pool of storage resources that spans all endpoints and common repositories. This pool must be virtualized to integrate new technologies and manage itself automatically based on policies or business events.

As VSP One Object evolves with new use cases, managing Personally Identifiable Information (PII) becomes more crucial than ever. The PII Discovery feature is designed to seamlessly integrate into your data governance strategy, providing automated identification and categorization of PII such as names, addresses, and phone numbers. This feature not only ensures compliance with data privacy regulations like GDPR and CCPA but also offers robust redaction and masking capabilities to protect sensitive information. By leveraging advanced algorithms, PII Discovery can handle large volumes of data, making it scalable for enterprises with extensive data assets. This empowers your business to maintain data integrity and privacy, ensuring that your data governance strategy is always a step ahead.

Regardless of the path you choose for the future, it's wise to avoid point storage solutions that create data silos. Instead, focus on embracing data mobility. Future challenges from unstructured data require technologies that go beyond traditional storage solutions.

## Virtual Storage Platform One Object

### Your Landscape, Our Expertise: Solving Problems Together

VSP One Object is a cutting-edge object storage solution, designed to meet the diverse needs of modern enterprises. Whether you're managing a data lake/lakehouse, running AI/ML and analytics workloads, or ensuring robust backup and data protection, our solution offers unparalleled performance, scalability, and flexibility.

#### Data Lake and Lakehouse

- **Flexible Solution:** Efficiently handles both small and large objects, optimizing storage costs.
- **High Object Count:** Supports over 2.5 billion objects per node with scalable metadata drives.
- **Data Lakehouse Workloads:** Supports unified data storage, ETL processes, data governance, security, and multi-engine support.
- **Advanced Metadata-Type Tiering:** Manages data efficiently based on object properties, tags, and abstracted metadata.

#### AI/ML and Analytics

- **High Volume of Objects:** Utilizes NVMe drives and a 1000 GbE network for efficient data management and high-throughput AI/ML workloads.
- **High-Capacity SSDs:** Enhances performance with large storage capacities for AI/ML data.
- **Data Ingestion Tools:** Uses Kafka and Rabbit MQ to prepare and feed data into S3 for analytics.
- **AI Workloads:** Supports training, inference, data preprocessing, model storage, and versioning.
- **Big Data Analytics Workloads:** Integrates with Hive, Spark, Kafka, and Presto for data ingestion, processing, querying, and archiving.
- **High Throughput and Low Latency:** Delivers exceptional performance.
- **S3 Event Notification:** Enables real-time analytics workflows.

#### Backup and Data Protection

- **Massive Network Bandwidth:** 200 GbE per worker node, totaling 1000 GbE across 5 nodes.
- **Flexible Network Configuration:** Supports both 25 Gb and 100 Gb NICs, configurable up to 250 Gb per node.
- **Extensive Connectivity:** Capable of attaching to other IP storage through the front-end network.
- **High-Performance S3:** Converged and disaggregated storage options allow independent scaling of storage and compute.
- **Storage Class Configuration:** Configure local, S3 external, and block external storage for varied performance needs.



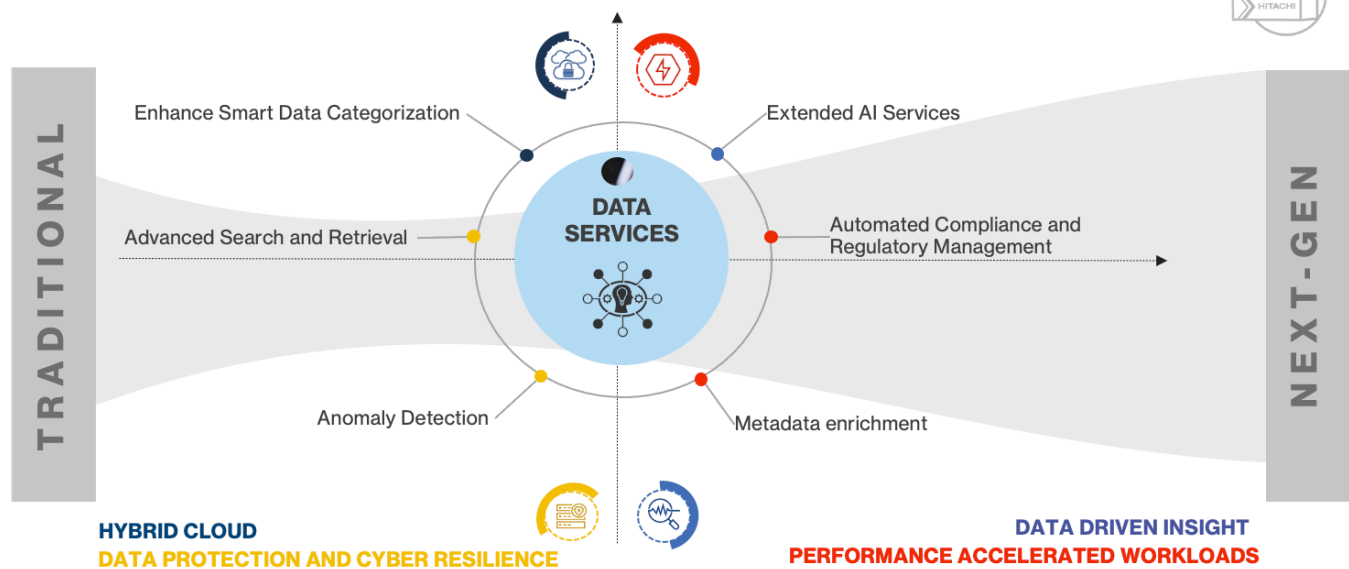
# Virtual Storage Platform One Object

## Data Services: The Key to Innovation

BY 2027, at least 40% of organizations will deploy data storage management solutions for classification, categorization, insights and optimization, up from 15% in early 2023<sup>1</sup> VSP One Object is built with an emphasis on data services which enables analytics and seamless management.

When Hitachi Vantara discusses object storage data services, it means true intelligence applied to data, such as classification, categorization, metadata enrichment, content awareness, and data placement. This goes beyond basic Kubernetes management or backup. VSP One Object supports a variety of data services, including high-performance

ingestion and analysis, intelligent storage services, and AI-driven insights, which enable smart data categorization, advanced search and retrieval, anomaly detection, and automated compliance and regulatory management. Customers can expect improved data organization, faster data retrieval, and strengthened security compliance.



## Cost Efficiency, Flexibility, and Sustainability

VSP One Object integrates tightly across the Virtual Storage Platform One data plane, unifying access to block, file, and object storage. It leverages microservices-based architecture for flexible scaling and a best-in-class container management system, Kubernetes, providing flexible and automated deployment mechanisms.

The solution supports both converged and disaggregated hardware configurations, allowing users to deploy a hyperconverged setup or mix and match components to implement a tiered

approach tailored to specific workload requirements. VSP One Object uses self-encrypting drives (SEDs) and a variety of drive types to meet different performance targets, ensuring optimal efficiency and security. By incorporating all-flash storage, Hitachi Vantara offers significant environmental benefits, making it an excellent choice for large customers seeking both high performance and cost-efficiency.

The high-density design, featuring S-node spinning drives and mainline SAS drives, delivers industry-leading efficiency. This setup reduces rack space by 50-70% and improves power efficiency by up to 50%

compared to other leading object storage vendors, significantly lowering data center CO2 emissions. This combination of density, operational efficiency, and environmental sustainability makes VSP One Object a standout in the market.<sup>2</sup>

<sup>1</sup> Top Trends in Enterprise Data Storage 2023

<sup>2</sup> This estimation is based on an internal analysis of VSP One Object with 2 system units and 24TB NLSAS drives, compared to other leading vendors in a comparable 10PB configuration. Results will vary depending on each customer's specific configuration and load.

# Virtual Storage Platform One Object

**Table 1: Packaged Offerings**

Feature	Details
Nodes	8 Nodes with 6 NVMe SED slots each
Node Roles	5 Worker Nodes, 3 Management Nodes
Memory	512 GB RAM
Metadata Storage	Min. 1 NVMe slot per worker node reserved for Metadata
Local NVMe Storage Options	7.68TB, 15TB, 30TB
External Storage Options	S-Node, VSP One Block
Network Options	25 GbE, 25 GbE + 100 GbE
Node CPU Options	32 Core, 32/64 Core, 64 Core
Sustainability	Pre-packaged configuration to reduce risk and simplify deployment
Expansion Options	Cluster expansion in 8 worker node increments (4U)

**Table 2: Capacity Specifications**

Specification	VSP One Object 44
Min. Capacity per node (internal)	7.68TB NVMe
Max. Capacity per node (internal, with one drive reserved for meta data)	150TB (30TB NVMe)
Max. Capacity 5 worker cluster (internal, one drive for metadata)	750TB (30TB NVMe)
Disk Sizes	7.68TB, 15TB, 30TB NVMe
Max. External Capacity (VSP One Object S32)	Unlimited
Max. External Capacity (VSP One Block)	Unlimited
Networking Options	6 x 25GbE SFP28, 2x 10Base-T (Management) 2x 25GbE SFP28 and 2x 100GbE QSFP28, 2x 10Base-T (Management)
Scale-out Architecture	Hyperconverged and disaggregated storage
Objects per node	1.2M per 7.68TB metadata drive, up to 6 metadata drives per node
Licensing	Capacity based Subscription Term License

# Virtual Storage Platform One Object

**Table 3: System Specifications**

Specification	VSP One Object 44
Cores	32 Cores for Management Nodes 64 Cores for Worker Nodes
Minimum Cluster	8 nodes
Maximum Cluster	Unlimited
Data Protocols	Amazon Simple Storage Service (S3)
Height/ Width Depth (Controller)	448 x 177 x 737 mm (17.63" x 6.96" x 29")
Max. Weight (Controller excluding media)	112.71 KG / 248.49 lbs.
Form Factor	4U Rackmount
Power Consumption	4470W @ 127 VAC 4377 @ 240 VAC
Total BTU/hour (max)	15256 BTU/hour

**Table 4: Software Specifications**

Specification	VSP One Object 44
Management Features	<ul style="list-style-type: none"> <li>● Identity Provider integration: LDAP and Active Directory integration for Role-Based Access Control (RBAC).</li> <li>● Local User accounts, allowing for user management directly within the system.</li> <li>● User Groups &amp; Roles Management enables the creation and management of user groups and roles to streamline access control and permissions.</li> <li>● Storage Component (S-Node) Configuration to optimize storage performance and capacity</li> <li>● Service Configuration &amp; Scaling Supports configuration and scaling of services to meet varying workload demands and ensure optimal performance.</li> </ul>
Software Enhancements	<ul style="list-style-type: none"> <li>● Distributed Database Enhancements: Improvements aimed at increasing overall system scalability and stability.</li> <li>● Rearchitected Object Storage Software: Based on Kubernetes, this rearchitecture enhances the efficiency and flexibility of object storage.</li> <li>● Native K8s &amp; Rancher Management: Provides ease of deployment, configurability, flexibility, and scalability to meet workload demands.</li> </ul>
Data Services	<ul style="list-style-type: none"> <li>● S3 Fidelity: Ensures compatibility with S3 APIs.</li> <li>● S3 Object Lock (Compliance mode): Prevents object deletion for a specified retention period.</li> <li>● S3 Select API, S3 SSE-S3 APIs, S3 Expiration Lifecycle Policies (metadata maintenance, multiple functions, independent services, and event-driven)</li> </ul>

**Table 4: Software Specifications (continued)**

Specification	VSP One Object 44
Data Protection	<p><b>EC (Erasure Coding) Storage Classes</b></p> <p><b>Supported Classes:</b></p> <ul style="list-style-type: none"> <li>● 1+0: Single data block with no redundancy.</li> <li>● 1+1: Single data block with one redundant block for fault tolerance.</li> <li>● 3+2: Three data blocks with two redundant blocks for enhanced fault tolerance.</li> <li>● Flexible EC configuration is also supported on top of the 1+0, 1+1, and 3+2 EC storage classes.</li> </ul> <p><b>Disaster Recovery:</b></p> <ul style="list-style-type: none"> <li>● Async-Replication for Disaster Recovery: This feature supports disaster recovery by asynchronously replicating data to geographically distant sites, ensuring data availability and integrity without impacting performance.</li> <li>● Async-Replication (BC): Business continuity support through asynchronous replication</li> </ul> <p><b>Backup and Recovery:</b></p> <ul style="list-style-type: none"> <li>● Fast backup ingest/restore, disaster recovery, ransomware protection, tape alternative, cloud backup</li> </ul>
Supportability Features	<ul style="list-style-type: none"> <li>● Node Management: Add/remove node support</li> <li>● Notifications: Event notifications via syslog &amp; SMTP</li> <li>● Monitoring: Prometheus metrics, Grafana dashboards</li> </ul>
Security Features and Certification	<p><b>DARE Licensing (s-node):</b></p> <ul style="list-style-type: none"> <li>● Data-at-Rest Encryption licensing for secure nodes.</li> <li>● Authentication and authorization protocols: SAML, OIDC, MFA, etc.</li> <li>● External KMS support: Interoperable with any external key management server that is compatible with KMIP V1.3 or later</li> <li>● Client &amp; server certificate management: Manages certificates for secure communication.</li> <li>● Compliance and Security Certifications: Meets SEC-17A compliance and other certifications.</li> <li>● Legal Hold: Ensures data cannot be deleted until the hold is removed.</li> <li>● Versioning</li> <li>● All data transmissions are protected with mTLS encryption, v1.3</li> <li>● FIPS 140-2 certified library (s-node): Uses FIPS certified cryptographic library for secure nodes.</li> </ul>

*The reliability of Hitachi Vantara storage is legendary. That's why 90% of the world's top banks and 85% of Fortune 100 companies use VSP One storage solutions to support their open systems and mainframe environments.*

**Click here for more information about VSP One Object.**

**Learn more**



**Corporate Headquarters**  
 2535 Augustine Drive  
 Santa Clara, CA 95054 USA  
[hitachivantara.com](http://hitachivantara.com) | [community.hitachivantara.com](http://community.hitachivantara.com)

**Contact Information**  
 USA: 1-800-446-0744  
 Global: 1-858-547-4526  
[hitachivantara.com/contact](http://hitachivantara.com/contact)

© Hitachi Vantara LLC 2024. All Rights Reserved. HITACHI and Pentaho are trademarks or registered trademarks of Hitachi, Ltd. All other trademarks, service marks and company names are properties of their respective owners.

HV-BTD-DS-Vsp-One-Object-4Dec24-A