

Datasheet

Hitachi Vantara Virtual Storage Platform One Object (VSP One Object)

Transforming object storage

VSP One Object is transforming object storage from a passive archive into an active foundation for your business innovation.

VSP One Object transforms object storage into a dynamic engine for innovation — powering AI, analytics and lakehouse workloads with scalable, intelligent data services. On-premises native Amazon S3® table support and Apache® Iceberg catalog integration, enable fast, secure, in-place querying.

With support for VSP One Block QLC (quad-level cell) arrays, you can maximize infrastructure investments while scaling compute and storage independently. The new PII service feature automatically identifies PII data and takes actions to promote good data governance. VSP One Object simplifies operations, reduces costs and unlocks deeper insights from your data.

Unify Data. Unleash Innovation.

Modern enterprises face explosively expanding volumes of unstructured data, fragmented architectures and slow, costly analytics pipelines. Legacy object storage systems can't keep up with the scale, speed or flexibility required for AI, lakehouse, and compliance workloads. Siloed infrastructure, brittle integrations, and manual governance slow innovation and increase risk. VSP One Object solves these challenges with scalable, intelligent object storage — enabling in-place analytics, automated PII discovery and seamless integration with VSP One Block arrays to maximize performance and cost-efficiency.

Reimagine What Object Storage Can Do

VSP One Object redefines object storage as an intelligent foundation for modern data-driven innovation. It enables on-premises analytics with native S3 table support, and Apache Iceberg integration — eliminating the need for complex data movement or external metastores. Built for scale, it supports structured, semi-structured and unstructured data across AI, ML and lakehouse workloads. Automated PII discovery and enterprise-grade governance simplify compliance and data lifecycle management.

Why VSP One Object Delivers More

- Seamlessly scales from petabytes to exabytes without performance loss.
- Supports diverse workloads, formats and deployment models.
- Automates data management with built-in intelligence and PII discovery.
- Independently scales compute and storage.
- Enables fast, in-place analytics with on-premises native S3 table support.
- Reduces footprint and emissions with sustainability-optimized architecture.
- Maintains VSP One Block guarantees when connecting to VSP One Block arrays.

With support for QLC- or TLC- (triple-level cell) based VSP One Block arrays, organizations can maximize existing infrastructure investments while independently scaling compute and storage to meet performance and budget needs. VSP One Object delivers agility, insight and resilience — empowering businesses to simplify operations, reduce costs and unlock the full value of your data.

Core Capabilities

Query Data Where It Lives

VSP One Object enables fast, in-place analytics with on-premises native S3 tables support and Apache Iceberg integration — eliminating the need for external metastores or data movement. VSP One Object allows direct querying of object data using familiar BI tools, accelerating insights while simplifying architecture. This capability empowers data teams to run scalable, ACID*-compliant analytics directly on object storage, reducing complexity and cost while boosting performance across lakehouse, AI and ML workloads.

* ACID = atomicity, consistency, isolation and durability

Protect Sensitive Data Automatically

VSP One Object includes a built-in PII service feature that automatically identifies and classifies sensitive PII data such as names, addresses and phone numbers. It supports redaction, masking and policy-driven actions to streamline compliance with regulations like GDPR and CCPA. This automation reduces manual overhead, enhances data governance and ensures privacy across massive datasets — making it easier for organizations to stay ahead of evolving regulatory demands while maintaining trust and security.

Scale Smart With QLC or TLC-Based Arrays

VSP One Object integrates seamlessly with QLC or TLC-based VSP One Block arrays, allowing organizations to leverage existing infrastructure for high-density, cost-efficient storage. This support also enables independent scaling of compute and storage, optimizing performance and budget for dynamic workloads. By combining object and block storage, businesses gain flexibility to support AI, analytics and long-term retention use cases — without rearchitecting or sacrificing speed, durability or sustainability.

Built To Power What's Next

- 1st in the industry to offer on-premises native S3 table support for fast, in-place analytics.
- Direct SQL querying of object data without extra tools.
- Apache Iceberg integration for scalable, open-format analytics.
- PII services for automated discovery to simplified compliance and governance.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA
hitachivantara.com | community.hitachivantara.com

Contact Information
USA: 1-800-446-0744
Global: 1-858-547-4526
hitachivantara.com/contact

© Hitachi Vantara LLC 2025. All Rights Reserved. Amazon and Amazon S3 are registered trademarks of Amazon Technologies, Inc. Apache is a registered trademark of Apache Software Corporation. All other trademarks, service marks and company names are properties of their respective owners.

HV-BTD-DS-Virtual-Storage-One-Object-18Sept25-A

VSP One Object

- On-premises native S3 tables and Apache Iceberg support.
- External storage via S-Node or VSP One Block midrange.
- Up to 1.25B objects per node
- Independent scaling of compute and storage resources.
- 25GbE and 100GbE network options.

- Seamless integration with VSP One Block midrange QLC arrays.
- Independent scaling of compute and storage for cost efficiency.
- High-density design reduces rack space and power consumption.

Modernize Your Data Strategy

Modern enterprises are overwhelmed by fragmented data, siloed infrastructure and legacy storage that can't keep pace with AI, analytics and lakehouse demands. VSP One Object solves these challenges by transforming object storage into an intelligent, scalable foundation for innovation.

With on-premises native S3 table support and Apache Iceberg integration, VSP One Block enables fast, in-place analytics — eliminating the need for complex data movement. Its automated PII discovery streamlines compliance and governance, while support for VSP One Block QLC arrays allows organizations to scale compute and storage independently, maximizing performance and cost-efficiency.

Built on a microservices architecture, VSP One Object delivers high availability, flexible deployment and enterprise-grade security. Whether you're powering AI/ML pipelines, modernizing your lakehouse, or protecting long-term data, VSP One Object helps you simplify operations, accelerate insights and future-proof your data strategy.