

Solution Profile

Modernize Data Lakehouse Environments With VSP One Object Storage

Future-ready data lakehouse storage — Amazon S3 Tables enabled



Simplified Data Lakehouse Architecture

Consolidate storage and analytics with on-premise native S3 Table support — no external services needed — simplifying architecture and operations

Faster, Cost-Efficient Analytics

Query data in place using open table formats like Iceberg — no need to move or duplicate data into external engines or warehouses

Compatibility With Modern Analytic Tools

S3 Table and Iceberg Catalog support ensures seamless compatibility with modern analytics tools, enabling faster insights directly from your object storage

The complexities of managing a data lakehouse — fragmented storage, slow queries and costly data movement, which hinder performance, scalability and integration — require a modern storage solution.

Today's data-driven organizations are under immense pressure to extract value from ever-growing volumes of unstructured and semi-structured data. Yet many struggle to keep their data lakehouse environments efficient, scalable and analytics-ready.

The challenges start with fragmented infrastructure. Data is often siloed across multiple systems, requiring complex pipelines and costly data movement just to make it accessible to analytics tools. These inefficiencies slow time to insight and increase operational overhead. Without native support for modern table formats, customers often rely on external metadata services and custom integration layers, creating brittle architectures that are hard to maintain and scale.

Performance is another key concern. As datasets grow and analytic workloads become more frequent and more complex, traditional object storage solutions fall short — unable to support fast, reliable, consistent and accurate processing of queries or interactive analytics across large-scale data.

Security and governance also suffer. With disconnected systems and inconsistent policies, enforcing data lifecycle management, access controls, or compliance becomes cumbersome and error-prone.

Customers are left asking: How can we modernize our data lakehouse without reinventing our architecture?

The answer lies in rethinking the foundation — starting with storage. A modern object storage system that natively supports S3-compatible tables like Apache Iceberg or Delta Lake addresses these challenges head-on. It enables direct, high-performance access to structured data without requiring data movement or external services, while seamlessly integrating with existing analytics tools and pipelines.

The result? Simpler operations, faster insights, and a future-ready data infrastructure.

Address These Challenges With a Modern Object Storage Solution

A modern object storage solution with native Amazon S3 Table support is purpose-built to meet the demands of data lakehouse workloads. By supporting open table formats like Apache Iceberg directly within the storage layer, it eliminates the need for external metastores and complex data movement. This enables high performance, in-place querying of structured data, which is crucial for fast, scalable analytics.

ACID-compliant operations* ensure reliable, concurrent reads and writes, which are essential for maintaining consistency across evolving datasets.

Seamless integration with popular analytics engines like Apache Spark, Presto and Trino streamlines data workflows and accelerates time to insight. As lakehouse environments grow to support AI, ML and real-time business intelligence (BI), native S3 Table support ensures the underlying storage keeps pace — delivering the scale, simplicity and performance modern data teams require. This next-generation object storage solution becomes the foundation for a more efficient, cost-effective and agile lakehouse architecture.

* Atomicity, Consistency, Isolation, Durability (ACID) compliance ensures that operations like reads, writes, updates and deletes happen reliably.

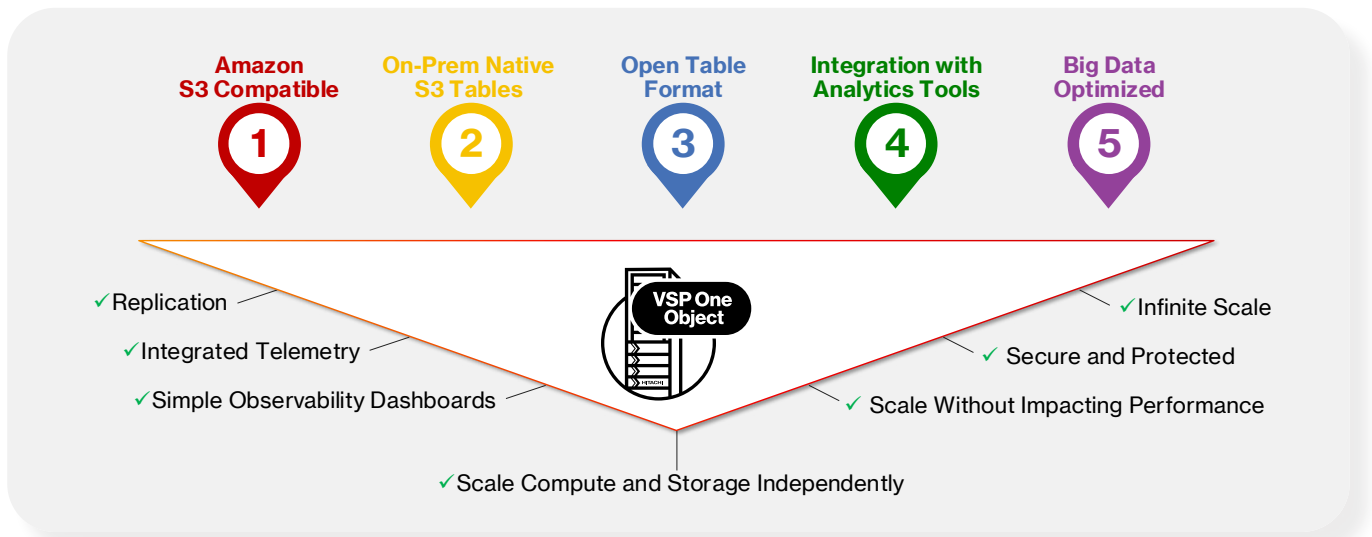


Figure 1: Modern data lakehouse infrastructure unifies storage and analytics by combining S3 compatibility, native S3 Table support, open formats and intelligent services for scalable, tool-integrated and big data performance.

Run Your Data Lakehouse on VSP One Object

Hitachi Vantara Virtual Storage Platform One Object (VSP One Object) is a modern, enterprise-grade solution designed to power data lakehouse environments with unmatched scalability, performance and simplicity. Fully S3 compatible, VSP One Object supports on-premise native S3 Tables and integrates directly with Apache Iceberg catalogs, enabling fast, ACID-compliant, in-place analytics on structured data — without external metastores or complex data movement.

Organizations can seamlessly scale from terabytes to exabytes, supporting long-term data growth and AI/ML workloads with ease. VSP One Object allows independent scaling of compute and storage resources, optimizing performance and cost efficiency. Its native integration with leading analytics tools like Apache Spark, Trino and Presto accelerates time to insight, while built-in support for data

replication, security controls and object-level access policies ensures enterprise-grade reliability and governance.

Built with intelligence at the core, VSP One Object automatically detects sensitive data, such as personally identifiable information (PII), helps enable compliance and streamlines data classification. Optimized for big data, it delivers low-latency access and supports mixed workloads — structured, semi-structured and unstructured data — making it ideal for modern lakehouse architectures.

VSP One Object brings high durability, robust metadata management and intelligent data services to deliver operational efficiency and deep insights from your data. Whether you're managing petabytes of IoT telemetry or building real-time AI pipelines, VSP One Object enables a future-ready foundation for your data lakehouse strategy.

Features

On-Prem Native S3 Table Support

Enables in-place, ACID-compliant operations on structured data using open table formats like Apache Iceberg. No need for external metastores or duplicated data.

Run Fast, Reliable Queries Directly on Object Storage Without Moving Data

Seamless Integration With Apache Iceberg Clients

Works out of the box with Spark, Trino, Presto and other tools, accelerating time to insight and simplifying lakehouse data workflows.

Accelerate Insights Using Familiar Tools Like Spark, Trino and Presto

S3-Compatible API Support

Fully supports S3 APIs, enabling easy integration with existing tools, cloud-native applications and data pipelines — reducing friction and accelerating adoption across hybrid or multicloud environments.

Easily Integrate With Existing Tools and Cloud-Native Applications

Scalable Architecture, From Terabytes to Exabytes

Effortlessly supports growing data volumes, from pilot projects to enterprise-scale lakehouses, ensuring consistent performance and capacity without disruptive migrations or rearchitecting.

Grow Your Lakehouse Confidently, Without Performance or Capacity Restrictions

Independent Compute and Storage Scaling

Allows organizations to grow analytics capacity and storage independently, optimizing performance and costs for dynamic, high-volume data workloads.

Optimize Performance and Cost by Scaling Resources as Needed

Apache Iceberg Catalog Integration

Supports seamless use of Iceberg tables across query engines, enabling schema evolution, time travel and optimized metadata operations at scale.

Simplifies Structured Data Management With Scalable, Open Table Support

Enterprise-Grade Replication and Data Protection

Delivers high availability and data protection through flexible replication, advanced storage services, strong consistency and secure object-level access controls.

Ensure High Availability With Built-in Replication

QLC and TLC Support

With QLC/TLC NVMe-based VSP One Block 20 array support.

Get the price/performance you deserve

Summary

VSP One Object is purpose-built to address the challenges of modern data lakehouse environments. With S3 compatibility, native S3 Table support and built-in Apache Iceberg catalog integration, it enables fast, in-place analytics without complex pipelines or external services.

VSP One Object simplifies lakehouse architectures by unifying storage and structured data access, allowing seamless integration with tools like Spark and Trino. Its ability to scale from terabytes to exabytes allows you to independently scale compute and storage resources and intelligently detect PII data, ensuring performance, governance and cost-efficiency.

VSP One Object provides the scalable, intelligent foundation organizations need to modernize and future-proof their data lakehouse strategy.

Hitachi Proven: 20 Years of Object Storage Innovation

With two decades of proven innovation and customer trust, Hitachi stands as a global leader in enterprise-class object storage. Our VSP One Object is a storage platform that is built for today's data lakehouse environments — delivering the performance, scalability and flexibility needed to power modern analytics. Backed by a robust ecosystem of expert support and technology partners, Hitachi is ready to help organizations unlock the full potential of their data with confidence and ease.

Ready to modernize your data lakehouse environment with VSP One Object?

Get in touch with us. →

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 2026. 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-SP-VSP-One-Object-Data-Lakehouse-13Apr26-C