Hitachi Vantara

Solution Profile

Hitachi Virtual Storage Platform One Object

Next Gen Object Storage Supporting Data Lakehouse for AI

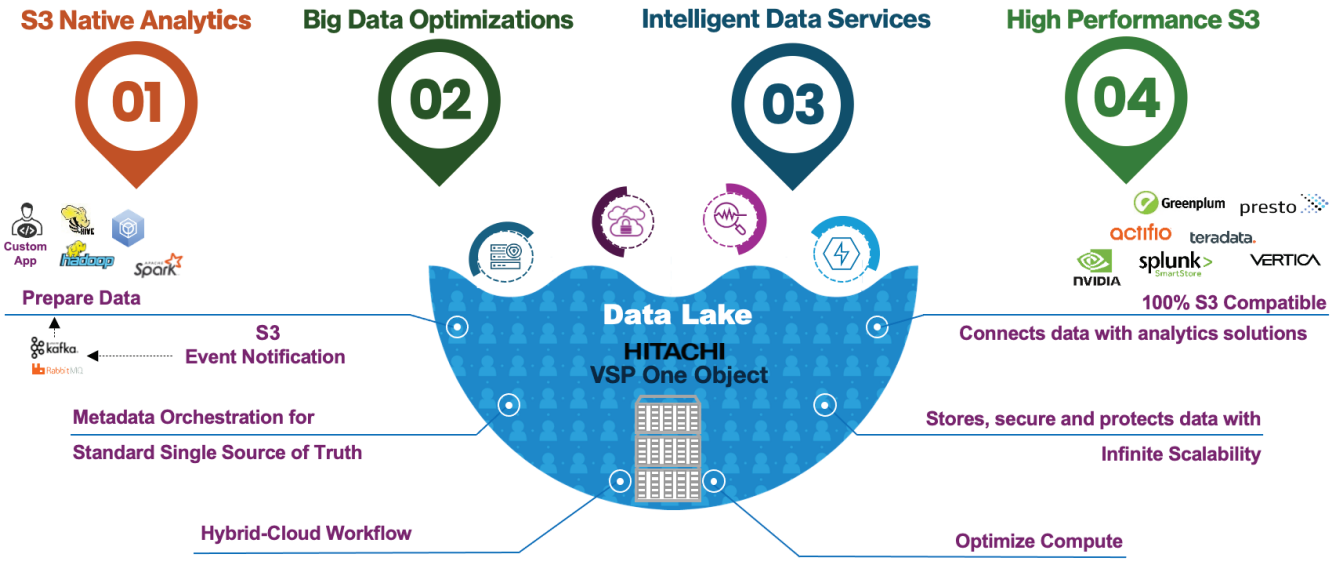

The demand for flexible and scalable storage solutions is more critical than ever, as managing immense amounts of data can be daunting and costly with outdated platforms. Customers face challenges with inefficiencies and high costs associated with siloed systems. Virtual Storage Platform One (VSP One) Object simplifies data management by providing a performant, durable, and intelligent solution. It seamlessly integrates with both cloud and on-premises systems, supports diverse workloads, and ensures data sovereignty. Unbreakable and sustainable, it delivers high throughput and low latency, empowering businesses to stay at the forefront of AI innovation.

Flexibility

VSP One Object offers massive scale, cloud capabilities, and broad protocol support. It supports all-flash configurations and policy-driven tiered storage, allowing performance and capacity to scale independently. It quickly ingests large volumes of small objects and cost-effectively stores very large objects.Customers benefit from a diverse array of configurations and deployment options. These include traditional archive and backup solutions, as well as modern cloud-native environments, AI/ML applications, and analytics workloads. The platform is optimized for a compact form factor, fitting an entire cluster into a 4U footprint. This enables real-time analytics at the edge, eliminating the need to transfer data to a central location for analysis.

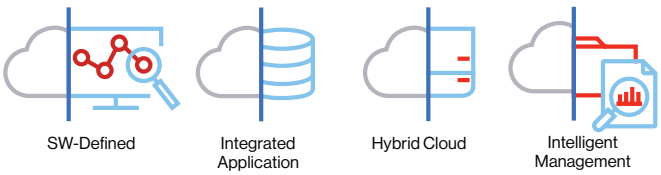

S3 Native Analytics for Lakehouse and AI/ML

Advanced Analytics from Object Data

At Hitachi, we guide you in leveraging S3 for intelligent data services. By embracing open data lake architectures, you can run analytics directly on data stored in S3, manage data preparation, and streamline ingestion—all within a unified setup that simplifies and enhances analytics workflows. Our intelligent data services include S3 event notifications that automatically trigger workflows or analytics processes, providing near real-time insights into object activities such as writes, updates, and deletions. This capability ensures that your system can quickly adapt to changes, and maintain data integrity.

For data lakes and analytics use cases, our approach to compliance and retention management is unique. These capabilities allow for efficient updates and modifications without moving the data itself. In healthcare, our solutions can extract and aggregate metadata, aiding in trend analysis and compliance reporting. Metadata management involves collecting, storing, and organizing metadata, making data easier to discover and understand. This is crucial for integrating existing repositories into a unified system, enhancing data accessibility and usability.

Event-based retention policies ensure data is governed and archived based on object lifecycle events, driving long-term data analysis. We enable comprehensive data processing and enrichment, where action-based analytics trigger operational workflows based on object activities. Automatic data enrichment ensures your data is optimized for complex analysis without manual intervention. Operational automation, such as scaling bucket quotas or restricting access based on suspicious activities, feeds into analytics dashboards, providing real-time operational intelligence and predictive analytics.

Hybrid Cloud Storage for S3-native Analytics

Hitachi helps customers better leverage their data for advanced analytics. Our hybrid cloud solution is designed to meet these demands. We recognize that access speed and performance are paramount for AI and ML use cases. VSP One Object allows access to local disks and different models within VSP One's product line without limitations. This means you can start with 7.8TB and scale to a 60TB drive or any future capacity drive that becomes available without code modifications. This flexibility ensures that your system will grow with your needs, providing a future-proof solution.

Our solution is engineered to deliver high-speed local data processing, eliminating network performance trade-offs. By unifying data management across structured, semi-structured, and unstructured data within a single platform, we simplify data handling with advanced metadata capabilities. Additionally, our solution features robust data preparation and integration functionalities. These capabilities streamline the process of preparing data for analysis, making it easier to manage and derive insights from large datasets. This approach not only optimizes the total cost of ownership but also enhances performance based on density pricing.

We are building a balanced combination of capacity and performance, architecting a high-performance environment that supports a variety of media options, including HDD, QLC SSD, and NVMe. VSP One Object is architected to support protocols such as S3, Posix, and SQL to ensure compatibility with a wide range of applications, including Elastic, Spark SQL, Starburst, Vertica, Confluent, and Red Panda. VSP One Object offers file protocols through a gateway with Hitachi Content Platform Anywhere Enterprise.

Data Protection, Backup and Cyber Resiliency

VSP One Object offers unique benefits and capabilities for lakehouse solutions. This solution caters to those who seek reliable and efficient data management to ensure business continuity. VSP One Object is able to withstand multiple disruptions while maintaining continuous availability for always-on businesses, and ensuring data is securely stored and protected. VSP One Object is compatible with a range of applications, including Commvault, Veeam, Veritas, and Rubrik. Features like immutability and end-to-end encryption enhance cyber preparedness and protect against sophisticated attacks, ensuring data security and recoverability.

Key capabilities of VSP One Object include streamlined data integration, ensuring efficient and reliable big data optimization. By defining policies, we can perform data placement, allowing for advanced metadata-type tiering based on object properties, tags, and abstracted metadata. We are also developing additional metadata-type tiering capabilities to further enhance our offerings. For example, if you have an application that requires high-performance, you can define a policy to place incoming data from this application on flash storage within the same object store. For backup purposes, where high performance is less critical, data can be sent to disk storage instead. Some workloads may require a mix of both, allowing you to send data to local disk versus an S node based on specific needs. By defining these policies according to workload requirements, you can optimize data placement and performance.

Reliability

The system is designed to handle node failures without impacting data ingestion, making it highly reliable for large-scale data operations. Built-in data resiliency keeps your data always available, protected, and secure. It enables trusted content mobility with full visibility of all the control points where data enters, exits and exists across a global IT landscape. Object storage is now a primary target for new cloud-native applications, so reliability and availability are paramount requirements.

Hardware failures are inevitable, and their impact on data can be catastrophic. VSP One Object mitigates this risk through its distributed architecture. By distributing data across multiple nodes and replicating it across multiple sites, object storage safeguards against hardware failures. In the event of a failed component, requests are automatically redirected to functional nodes, ensuring continuous access to data. This redundancy protects organizations from potential data loss or service interruptions, enhancing overall system resilience.

VSP One Object leverages patented technology and advanced capabilities to deliver unbreakable data resiliency. Customers experience continuous access to their data with the assurance of erasure coding and data replication that enables Hitachi's 100% Data Availability Guarantee extending to block, file, and object storage when using VSP storage as the data target.

VSP One Object natively has exceptional durability. Objects stored on S32 nodes have a combined durability of 15 nines. Handle the unexpected with smart load balancing that helps manage unexpected outages and activity surges. Immutable object-locking and site replication are additional features built-into VSP One Object.

Today, with the release of VSP One Object and future updates, we support flexible data protection options. Customers can set policies to tolerate up to three node failures, rather than just two. This flexibility allows you to tailor protection based on the workload and sensitivity of the application and data. For instance, you might be okay with losing availability for a backup workload but require extra protection for a high-performance, Tier 0 application.

High-performance

Object Storage for AI/ML

Put your data to work with VSP One Object's flexible architecture, which enables high performance at predictable costs. Designed for hyperscale cloud and enterprise environments, VSP One Object scales to exabyte capacities with multi-petabyte nodes, delivering high performance for AI/ML workflows and data storage in a lakehouse. Policy-driven tiering and data placement, combined with flash-accelerated storage, enhance outcomes for applications requiring hyperscale and comprehensive S3 capabilities.

Leveraging high-performance storage solutions is crucial for businesses aiming to stay ahead of the curve. Utilizing S3 for high-performance data storage and retreival offers advantages, especially when integrated with powerful analytics and data warehousing tools like Vertica, Snowflake, and Presto. Seamless integration with analytics solutions ensures rapid data retrieval and processing, essential for real-time analytics. With complete S3 compatibility, it guarantees smooth integration with existing analytics platforms, enhancing overall efficiency.

VSP One Object Platform supports disaggregated, converged, and hyperconverged storage configurations, reducing waste and promoting efficient resource use combined with a wide range of media options available in VSP One Block.

Affordable AI provides a cost-effective solution with high performance, making advanced AI capabilities accessible as part of the IQ stack. Our object storage solution integrates seamlessly with Hitachi IQ for enhanced data management and analytics.

Secure, Performant, Resilient, S3-Compatible Object Storage Software

Case Study: Sustainability for performance-intensive data lakehouse solutions, high-performance backups, and AI/ML operations

Customer: Large Enterprise

- **Challenge:** In a Data Lakehouse environment, substantial amounts of data is ingested daily. A large Hadoop cluster processes this data, which is then rapidly transferred to Hitachi object storage using S3. Currently, the system utilizes several racks to store all of this data but the customer faces space and power constraints. Therefore, efficiency improvements are necessary. The plan is to continuously add large volumes of data daily, enhance system performance with VSP One Object, and eventually handle even larger data volumes. The goal is to reduce data into less space while maintaining performance, focusing on density and sustainability to save space and power, ensuring smoother operations .
- **Solution:** Hitachi developed a new solution for the customer's existing Hitachi Gen2 S3 platform with VSP One Object and VSP One Block. This innovative architecture enabled seamless data migration between storage systems without disrupting operations. By virtualizing storage, VSP One Object optimized resource usage by dynamically allocating storage based on current needs. This reduced waste and enhanced performance. It also leveraged various media options, including high-capacity flash storage, to ensure faster data processing and improved overall system performance.
- **Results:** This solution demonstrated significant cost and power savings by reducing the number of racks by more than half. Each rack incurred substantial annual costs, so this reduction led to a 57% decrease in expenses. Additionally, it highlighted space efficiency and performance benefits. By optimizing output to VSP One Object, the system can now continue to ingest very high volumes of data daily and expand to many more petabytes, eliminating space and power constraints.

Cost Optimization

VSP One Block enables customers to integrate and utilize third-party storage assets. In the near future, VSP One Object will also be able to leverage VSP One Block. This means that with Hitachi Vantara Virtual Storage Platform One solutions, an existing block storage array can be transformed into S3-compatible object storage. Imagine transforming your existing VSP One Block storage into a powerful S3 platform without the need for new storage purchases. By virtualizing other arrays, you can activate this software and migrate storage between arrays during technology refreshes without any interruptions. This ensures your operations continue smoothly, leveraging high-capacity TLC and QLC flash storage for performance intensive tasks.

Sustainability

VSP One Object dramatically reduces rack space by an impressive 50-70% and enhances power efficiency by up to 50% compared to other leading object storage vendors. This substantial improvement not only optimizes operational efficiency but also significantly lowers data center CO2 emissions, providing a performant, durable, and intelligent solution*

* 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.

9 out of 10
Top Banks Choose Us

9 out of 10
Leading Healthcare Providers Trust Our Expertise

All top 10
Telcos & Insurers Rely on Us

Choose Hitachi Vantara: The Trusted Name in Innovation

Intelligent Data Services

Intelligent Data Services provide advanced capabilities for data management and protection. These services ensure data is securely stored and protected, integrating seamlessly with various analytics tools and platforms.

- **Unified Management** streamlines administration and boosts efficiency by integrating various tools into one interface, allowing you to focus on your core business.
- **Accelerated Insights** provide quick access to critical data, enabling real-time analytics, advanced analytics, and machine learning.
- **Compliance Monitoring** ensures your data practices align with regulations, reducing the risk of penalties.
- **Data Classification** organizes and tags data using metadata and machine learning, improving data retrieval and management.
- **Forensic Analysis** offers detailed data examination for investigations and audits, uncovering crucial information through forensic techniques.
- **Content Transformation** enhances media usability by converting it into various formats, ensuring your content is always accessible and in the right format.

These services cater to different use cases such as archiving, backup, lakehouse architecture, and AI/ML applications. For archiving, Intelligent Data Services manage long-term data storage, ensuring data is securely archived and easily retrievable. For backup, they create and manage backups, protecting data against loss and enabling restoration in case of failures or disasters. In a lakehouse architecture, they combine features of data lakes and data warehouses, providing a unified platform for both structured and unstructured data. For AI/ML, they manage and prepare data, ensuring it is accessible, clean, and ready for analysis.

Intelligent Data Services support HDD/SSD media and S3/REST protocols, and are compatible with industry leading applications.

Learn more

Explore the Hybrid Cloud Storage with Hitachi Content Platform solution brief.

About Hitachi Vantara

Hitachi Vantara is transforming the way data fuels innovation. A wholly owned subsidiary of Hitachi, Ltd., we're the data foundation the world's leading innovators rely on. Through data storage, infrastructure systems, cloud management and digital expertise, we build the foundation for sustainable business growth.

Hitachi Vantara
A Hitachi Group Company

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

HV-BTD-SP-Virtual-Storage-Platform-One-Object-28Oct24-A