

Choose Hitachi's all-flash data reduction and realize superior VMware private cloud scaling with integrated data efficiencies and direct connect to your choice of cloud.

DATASHEET

Hitachi Storage Virtualization Operating System

Improve Your Private Cloud Return on Investment With New Efficiencies

Every business is under pressure to transform. New and existing applications must be deployed to accelerate time to market, improve customer experience and uncover new revenue streams. Organizations gain agility and flexibility by moving certain workloads to public cloud, while hosting other applications on-premises for better control and security considerations, enabling true hybrid cloud scenarios. At the same time, it is critical that your hybrid cloud delivers operational efficiency to run a business successfully. The difficulty is balancing the demands for improvement with the constant growth in data.

Challenges arise as you migrate bare-metal workloads, such as online transaction processing (OLTP) and online analytical processing (OLAP) databases to a server environment. The odds are that your current flash system is struggling to deliver on its promise. As the hypervisor creates an I/O blender effect on your sequential workload, you experience higher latency and reduced virtual machine (VM) density. And this contention moves from your storage system interface through to the controller processing.

The tipping point is often the data reduction service of all-flash arrays (AFAs). This service heavily consumes processing and memory resources to maintain data relationship and consistency. As data

and metadata levels increase, they can cause controller bottlenecks and sporadic response times.

Hitachi Storage Virtualization Operating System (SVOS) is designed to deliver superior adaptive data reduction and operational efficiency. It enables effective capacity usage by up to 5:1 or more, up to two times greater VM density and 50% lower latency than alternative AFAs (see Figure 1: Capacity savings from eliminating duplicate VM copies).

To improve return on investment and allow greater VM consolidation, SVOS adaptive data reduction intelligence is optimized for highest system throughput and response time consistency. With multithreaded capabilities and quality of service (QoS) control, SVOS adaptive intelligence has the ability to slow down or pause the data reduction processing. It takes this action once the system reaches high processor utilization level and/or elongated wait time

is experienced for data on AFA cache that is ready to be written.

SVOS covers a broad range of efficiency services, including thin provisioning, snapshots and linked clones, compression, deduplication and cloud connect.

Minimize Ongoing Costs by Reducing Database and Virtual Machine Capacity

Adaptive Data Reduction Reacts in Real Time

Enterprise IT is becoming overwhelmed by the amount and variability of data, applications and change required to support today's information-intensive businesses. Simplification is needed to reduce complexity and the operating cost of running a private cloud.

SVOS All-Flash Architecture White Paper
READ

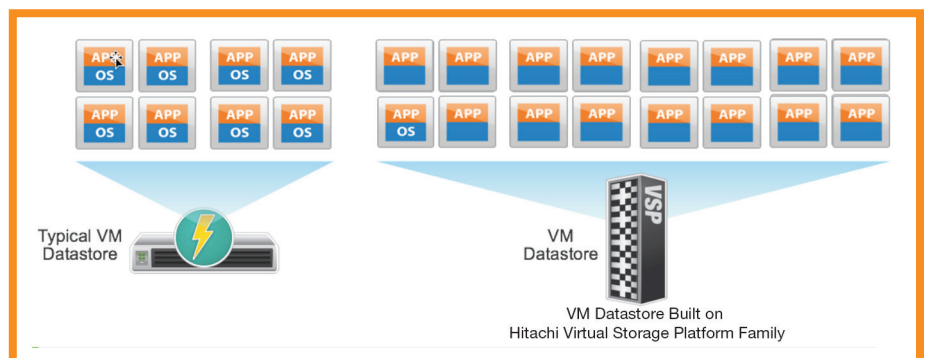


Figure 1. SVOS covers a broad range of data reduction and efficiency services for Hitachi Virtual Storage Platform family systems.

Adaptive data reduction includes selectable data deduplication and compression, intelligently adjusting the mode to minimize latency. SVOS detects whether a data pattern already exists, updates existing data inline and handles the new data post process. Data handling is optimized for performance because application I/O is prioritized over data reduction. Further, smart selection prioritizes the most stable data first, more efficiently working on the data set most likely to benefit from deduplication.

SVOS manages services used, based on configuration. If flash modules (FMD) are detected, FMD compression is used. If encryption is required, SVOS compression is used. Hitachi Data Systems guarantees that you'll get at least a twofold increase in flash capacity when you use our data reduction technologies.

Reduce Capacity Needs by Minimizing Duplicate Copies of Virtual Machines

Linked, Writeable Snapshots Using Thin Images

There simply isn't enough time to perform full backup of large databases without impacting user productivity and, ultimately, enterprise performance. Storage-based snapshot and replication technologies reduce risk by eliminating the backup window and increasing the frequency of backups.

Linked, writable snapshots can be created very rapidly and are very space efficient. They deliver performance consistent with that of source volumes, regardless of how many there are, and can be retained for long periods of time without performance impacts. Simplify VM deployments by creating new VMs consistently and in seconds. Perform global updates in seconds. Accelerate business development by copying production data to test new models. Use thin images to minimize costs and test more ideas. Thin images reduce capacity needs by up to 97%. Tiered snapshot

HITACHI STORAGE VIRTUALIZATION OPERATING SYSTEM: DATA REDUCTION*

Capability	Benefit
Pattern detection and removal	Pattern detection identifies predefined repetitive binary patterns, including zeros, prior to compressing and identifying duplicates. This process reduces the volume of data to be processed by the compression and dedupe engine.
Compression	Built on lossless extremely fast encoder/decoder; the algorithm enables the highest compression and decompression speed.
Accelerated Compression	Offload data reduction on Hitachi Accelerated Flash embedded lossless encoder/decoder; enables compression and decompression with no system or latency impact.
Deduplication	Built on a crypto hash fingerprint. This function finds and removes duplicate data without compromising its fidelity or integrity by segmenting a page into small fixed-sized chunks, identifying duplicate chunks, and maintaining a single copy of each chunk. Redundant copies of each chunk are validated with byte-to-byte comparison against original data and are replaced by a reference to the single copy.

* Expected data reduction ratios can vary and depend on individual environments and data structures.

copies offer full read/write copies or thin read-only copies. Copies can be "broken off" and used independently. Updates can then be propagated to the master.

SVOS inherently provides space-efficient virtual storage capacity. Capacity is allocated to an application without it being physically mapped until it is used. In this approach, it is possible to achieve overall higher rates of storage utilization with just-in-time provisioning. It also simplifies performance optimization by transparently spreading workloads across many physical devices, thereby reducing performance management concerns and self-optimizing performance and throughput.

Ensure Predictable Operating Expenditures Using Transparent Cloud Access

Migrate File Data Automatically, Based on Policies

With many enterprises now implementing both on-premises (private cloud) and off-premises (public cloud) services as part of their overall IT strategy, the ability to take

advantage of this hybrid model is critical. Minimize your ongoing storage and data center costs by migrating files to cloud automatically, based on policies. You can maintain choice without complicating operations via support for public cloud from Amazon and Microsoft®, as well as Hitachi Content Platform for private clouds. Where data lives is transparent to users (via stubbing).

Next Steps

SVOS is the standard operating system for the Hitachi Virtual Storage Platform family of all-flash and hybrid storage systems. SVOS delivers the foundation for digital transformation, enabling IT leaders to accelerate business insights and deliver a superior customer experience while realizing greater operational efficiency.

For more information, please visit HDS.com.



Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com | community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 31897900 or hds.marketing.apac@hds.com



Hitachi is a registered trademark of Hitachi, Ltd. VSP is a trademark or registered trademark of Hitachi Data Systems Corporation. Microsoft is a trademark or registered trademark of Microsoft Corporation. All other trademarks, service marks and company names are properties of their respective owners.

DS-432-C DG April 2017