Hitachi Virtual Storage Platform 5000

The Hitachi Virtual Storage Platform 5000 continues the high end enterprise system with Hitachi's virtualized storage architecture with an updated PCIe switching interconnect, newly designed controller cards, and NVMe attached devices. The VSP 5000 is a purpose-built, Intel-based hardware with the same embedded Storage Virtualized Operating System (SVOS) for the prior high-end and midrange systems.. The VSP 5200H and 5600H are hybrid systems and the VSP 5200 and 5600 are all flash systems that build on the high-end, highly reliable systems before it. Advances in the technology have increased performance with the new interconnect and NVMe devices.

The seamless continuation of the Hitachi VSP will allow customers familiar with the earlier systems to expand on their value derived from the system with the new capabilities and performance. Environments where the Hitachi VSP high end systems are used can be expanded with the new capabilities and simplified operations.

This document provides an overview and analysis of the Hitachi VSP 5000 that was announced in October 2019. HPE resells the VSP 5000, along with software options, as the HPE XP8. HPE's software options have different names and may also differ in terms of functionality in some areas. Competitive systems that support mainframe and open systems attachment alternatives include the Dell EMC PowerMax and IBM DS8900.

Contents

Highlights

The Hitachi 5000 is one of the few systems supporting mainframe and open systems from the same platform. With each generation of systems, Hitachi has continued to add capabilities that bring more value. Previous generations, such as VSP 5100 and VSP5500, utilized custom FMD devices and SSDs. The current models no longer support FMDs and instead soly utilize SSDs and SCM devices for solid state technology. Both SAS and NVMe SSDs are supported.

Highlights of the VSP 5000 include:

- Purpose-built controllers using standard Intel components
- Embedded software SVOS – Storage Virtualized Operating System
- Command Suite software – system management & SRM
- Virtualizes more than 200 external storage systems
- Up to 287 PB external storage support
- Thin provisioning for external volumes, fat-to-thin conversion
- Real-time tiering for optimizing cachine of data in DRAM
- Custom designed Flash Module Devices (previous models only) with inline hardware compression and optional hardware encryption with internal or externally managed keys
- Software compression and deduplication
- Automated Dynamic Tiering (Hybrid models)
  FMDs, SSDs, SAS, and NL-SAS devices

Overview

With the VSP 5200H and 5600H hybrid systems and the VSP 5200 and 5600 all flash systems, Hitachi offers a storage system that can operate at the very high-end of enterprise storage demands with a rich set of features and a history of very high reliability.

The VSP has Virtual Storage Machines (VSMs) which abstract the storage resources from the applications, enabling more advanced usage of the system. Each VSM can be managed as a separate storage system. The value the customer can derive from an VSM include:

- Maximize storage resource utilization
- Multi-tenancy for storage – isolation of access
- Advanced features
  - Non-disruptive migration
    - Data and remote copies
  - Global Active Device: active-active stretched clusters
  - High Availability operations
- Isolates application – no restarts, reboots on changes

Basic Characteristics of the VSP (all data per Hitachi)

	VSP 5200	VSP 5200H	VSP 5600	VSP 5600H
Maximum devices	96 NVMe 33 NVMe SCM 768 SSD	96 NVMe 33 NVMe SCM 768 SSD, SFF HDD 384 LFF HDD	288 NVMe 99 NVMe SCM 2,304 SSD	288 NVM 576 FMD 2,304 SSD, SFF HDD 1,152 LFF HDD
Device types	375 GB NVMe SCM 1.9, 3.8, 7.6, 15, 30 TB NVMe SSD 30, 15, 7.6, 3.8, .96 TB SAS SSD	375 GB NVMe SCM 1.9, 3.8, 7.6, 15, 30 TB NVMe SSD 30, 15, 7.6, 3.8, .96 TB SAS SSD 2.4TB 10K SAS HDD 14TB 7.2K SAS HDD	375 GB NVMe SCM 1.9, 3.8, 7.6, 15, 30 TB NVMe SSD 30, 15, 7.6, 3.8, .96 TB SAS SSD	375 GB NVMe SCM 1.9, 3.8, 7.6, 15, 30 TB NVMe SSD 30, 15, 7.6, 3.8, .96 TB SAS SSD 2.4TB 10K SAS HDD 14TB 7.2K SAS HDD
Hitachi measured IOPs	5.1M	5.1M	33M	33M
Controllers	2	2	4, 8, 12	4, 8, 12
Host ports	32 Fibre Channel 32 FICON 16 iSCSI	32 Fibre Channel 32 FICON 16 iSCSI	192 Fibre Channel 192 FICON 96 iSCSI	192 Fibre Channel 192 FICON 96 iSCSI
Cache size options	1 TB	1TB	2, 4, 6 TB	2, 4, 6 TB
Local copy pairs	32K	32K	32K	32K
Remote copy pairs	65,280	65,280	65,280	65,280
LUNs/LDEVs	65,280	65,280	65,280	65,280

Table 1: VSP 5000 Characteristics

For a complete listing of technical and physical specifications for the VSP 5000, the Evaluator Group Comparison Matrices should be referenced. .

The VSP 5000 is positioned as the high-end enterprise storage systems for Hitachi. One or two controller chassis are included with up to six 19" racks to hold a maximum of 2,304 SFF (2.5") devices.

Figure 1: Graphic from Hitachi for the VSP

Supported Devices

The VSP 5000 features use of 2.5" and 3.5" (for HDD in hybrid models) devices in a high-end virtualized storage array. The VSP supports 32 SAS or NVMe paths per Back End Director (BED) and up to 32 NVMe or SAS BED links at 12 Gb/s per controller chassis.

The latest VSP 5000 models support solid-state technology in the form of SSDs and SCM. Previously, VSP 5000 models utilized Flash Module Devices that are custom designed devices that utilize Flash technology. Wear leveling, error recovery, and drive sparing are supported features. Both SSDs and FMDs are supported under Hitachi's Dynamic Provisioning (Hitachi version of Thin Provisioning) and Dynamic (automated) Tiering.

Cooling

VSP 5000 supports installation in data centers that are configured for alternate hot and cold aisles. Airflow is from the front to the rear of a rack. Fan thermal control has three levels to reduce noise.

External Storage

Another major feature available with the VSP 5000 and its predecessor models is the ability to connect to external storage so that existing LUNs in a heterogeneous multivendor environment can be centrally managed. The Hitachi website should be consulted for the latest information on supported systems.

External storage connectivity is via an External LUN (ELUN). An ELUN operates in a manner similar to the Private Virtual Storage Machine, but it maps the LDEV through an external port of the VSP to a physical or target port on an external storage system. In effect, the external storage is treated by the VSP as a Virtual Storage Machine and maps LDEVs through it on a pass-through basis.

The current interoperability matrix for the Hitachi series in the support section of the Hitachi website should be consulted for the latest information on the external storage systems supported. External storage will be connected to the VSP via the front-end ports using Fibre Channel and Ethernet adapters using iSCSI.

Hardware Architecture

VSP high-end systems are composed of controller chassis, disk chassis, FMD chassis, and standard 42U x 19” racks.

Maximum VSP 5000 System Configurations

- Up to six racks per system
- One rack may contain a control chassis and up to two drive chassis, or up to three drive chassis
- One system contains up to two control chassis and up to 16 drive chassis
- A control chassis supports up to 96 I/O ports. The maximum port count is doubled when a second control chassis is added.
- In total, the drive chassis can contain up to 2,304 2.5” devices or 1,152 3.5” devices.

VSP Controller Architecture

The VSP controller architecture is comprised of controller blocks, nodes, and the Hitachi Accelerated Fabric. The Fabric is an internal PCIe switch implemented using an FPGA. There are basic components that comprise the controllers:

- Multiprocessor-based nodes called Storage Directors which contain the main system control multi-core processors and control data memory.
- Cache memory modules (CMs) which present a single global cache image to all of the different directors (see below for more detail)
- Front End Directors (FEDs) which manage the front end port connections to servers and to external storage. Front end directors are HBAs connected via the PCIe interconnect.
- Back End Directors (BEDs) which manage connections to internal storage drives. The back end directors are SAS HBAs or NVMe protocol connected via the PCIe interconnect. The back end directors also perform encryption when FMDs with the encryption feature (designated FMD-HDE) are not utilized. The encryption is AES-256 with internal or externally managed keys.

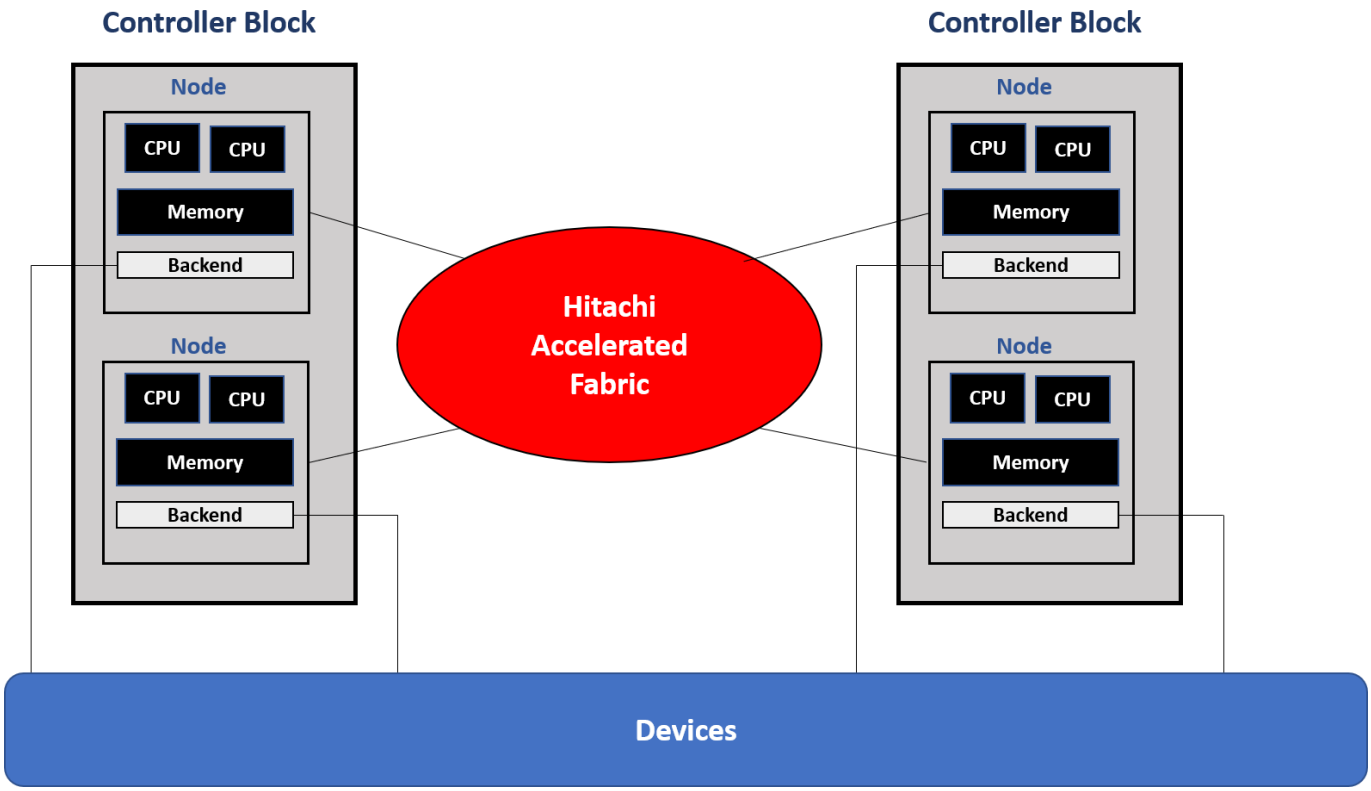

Figure 2: VSP 5000 Controller Architecture

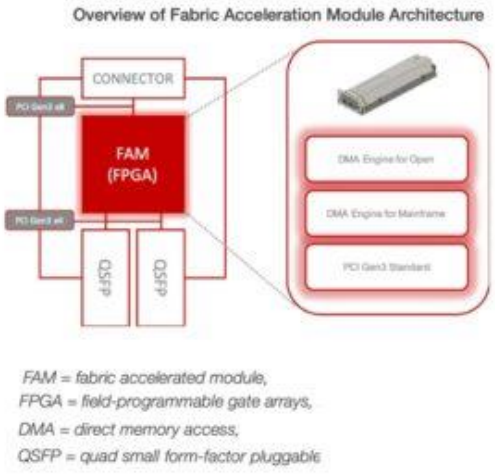

Figure 3: Hitachi Accelerated Fabric Source: Hitachi

Control Chassis Configuration

The minimum controller configuration can start with 2 x VSD, 2 x FED, 2 x CM and no BED for a diskless entry version of the Virtual Storage Platform for use with existing FC storage systems. These boards slot into a 4u high x 19-inch rack module. If/when internal disks are needed; customers can add 2-4 BED boards, which will provide 8 to 16 SAS links.

Maximum Controller Configurations:

- Up to 16 VSDs
- Up to 24 FEDs for a maximum 192 port connections
- Up to 2,048 GB of cache.
- Up to 16 BEDs for a maximum of 64 SAS links to SAS disk

Upgrades can be done without disruption to applications.

Virtual Storage Director

The VSD contains the main VSP system multi-core processors and control data memory. It stores and manages internal metadata and state data for the internal elements.

- Array groups
- Internal LDEVs and external LDEVs residing on virtualized third party storage arrays
- Runtime tables
- Mapping data for other Hitachi add-on software facilities

The VSD is where overall state-of-the-system data is stored, referenced and executed. When required, I/O processing is distributed to the appropriate I/O offload processors on the front and back-end directors. This is also true during microcode exchanges that are performed on both the VDS cores independently and the VSP cores vs. the I/O directors.

The VSD also handles processing for data copy functions. The use of multi-core processors allows copy functions to be isolated to and processed by cores dedicated to these functions, significantly improving copy performance vs. the previous generation USP design.

Distributed Control Memory

What was once centralized control memory in the earlier USP models is now local memory; resident on each of the VSDs. Data Cache is separate from control memory as can be seen in the above graphic. The VSP distributes control memory in 4 GB segments across all VSDs reducing the latency between system processors and control memory compared to earlier USP implementation.

Evaluator Group Comment: Hitachi has traditionally separated control memory from data cache so that usage of one does not impact the performance of the other. Other competing systems use data cache for their control memory such that control data can contend with production data in cache resulting in a performance impact. If partitioned, the amount of memory available to the system as cache is reduced.

Global Cache

Global cache contains application data blocks from drives via the BEDs during read operations, and application data blocks from FEDs during write operations. Global cache retains only one copy of read data whereas all write data is mirrored in cache. If a power or system failure occurs, data and system metadata in cache is protected by backing up cache to onboard flash drives. Two mirrored data copies are kept in the backup flash drives.

Flash Module Devices - Previous Models Only

The current models, VSP 5200 and 5600, no longer support FMDs.

FMDs are custom designed flash device that with embedded processor and logic for managing flash storage devices. Inline hardware data compression and optional data encryption is supported in the FMDs. FMDs with the encryption feature are designated FMD-HDE. The encryption support is enabled with a software license key. Encryption keys may be internally managed or exported for management with a KMIP. With the second generation of FMDs, 7TB and 14TB capacities are available. Twelve FMDs may be installed in an FMD chassis.

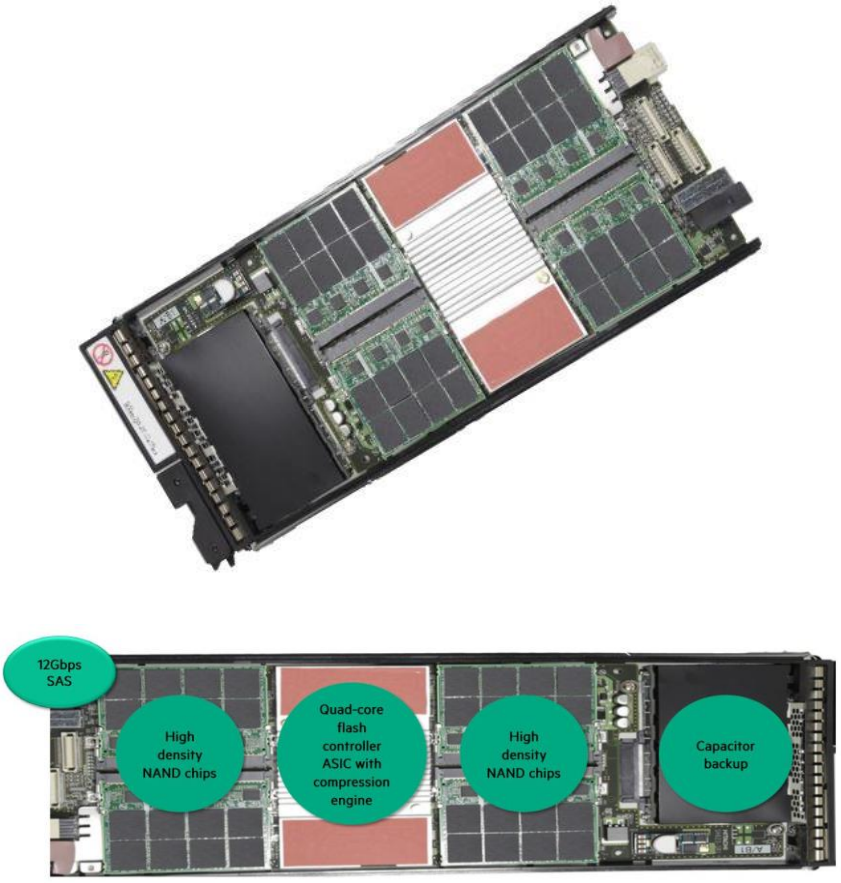

Figure 4: Flash Module Device Source: Hitachi

The Flash controller in the FMDs supports inline compression of data, which is implemented with a Lempel Ziv algorithm to compress data stored in the FMD. The inline compression reduces the capacity required and reduces the number of writes to the Flash technology. Optional encryption is done after compression.

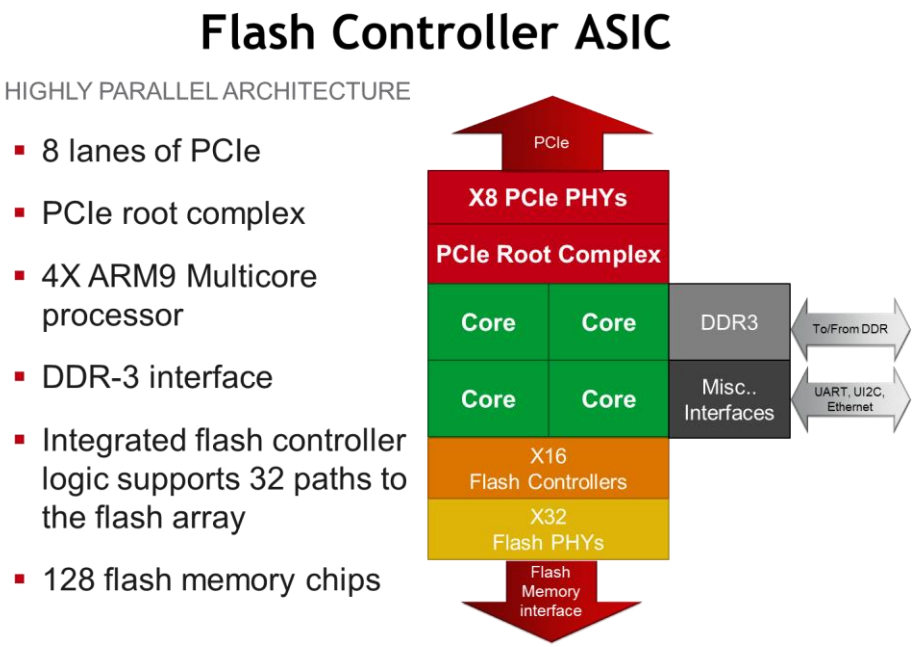

Figure 5: Flash Controller Design in the FMD

The FMDs support RAID-1, 5, and 6 for protection of device failure (managed by the VSP controller) and have added additional error correction for data integrity. A Secure Erase feature is implemented with the FMDs allows the PDEV erase function of SVOS to erase Flash memory blocks in the FMD for selected devices.

The FMDs are connected to the SAS backend directors of the VSP using an adapter card that converts the PCIe interface to 12 Gb/s SAS.

Adaptive Data Reduction

Adaptive Data Reduction is the name for software compression and deduplication. Each is individually selectable at a LUN level. The compression is inline and the deduplication has two types of operation. The first is a pattern match to determine if a block has been seen previously. If so, it is deduplicated. If not, then it is put on an ordered list for post storage (post process) deduplication. If the block is still highly active, it is pushed down on the list for post process deduplication, avoiding the overhead of inflating highly active data. FMDs still have inline hardware compression and it is recommended to not use software compression and deduplication on FMDs which would not yield the maximum performance. The software based compression and deduplication works on all devices (SSDs, HDDs, and FMDs) and does allow the data at rest to be encrypted either in the FMD or by the backend director.

Virtual Storage Ports

The Virtual Storage Port concept is a unique Hitachi feature that was first implemented in the 9900 V. Each physical Fibre Channel port may have multiple Virtual Storage Ports, with each virtual port supporting a different open systems server. Each server is assigned a Host Storage Domain (HSD), with each HSD having its own set of LUNs (including LUN0). Hosts are matched to their assigned HSD on a unique World Wide Name identifier.

Evaluator Group Comment: This feature is unique to this family of products from Hitachi. It can be used to reduce the number of front-end directors required to support multiple servers in those environments without significant throughput requirements.

Virtual Storage Machines

A significant feature in the VSP is the ability to be partitioned into up to thirty-two (32) private Virtual Storage Machines. The partitioning can be done at the capacity (internal and/or external storage), cache and port levels of granularity. The user allocates the resources via the Virtual Partition Manager software. The partitions can also be changed dynamically. VSM's enable an abstraction of the internal resources which provides these capabilities:

- Maximize storage resource utilization
- Multi-tenancy for storage – isolation of access by allocation of underlying physical resources to the Virtual Storage Machine on a percentage basis
- Advanced features
    - Non-disruptive migration – data and remote copies
    - Global Active Device: active-active stretched clusters
    - High Availability operations
- Isolates application – no restarts, reboots on changes

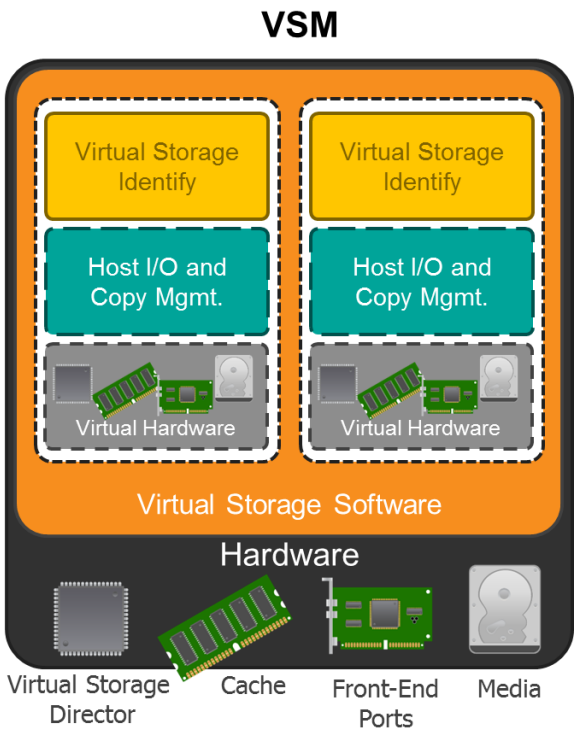

Figure 6: Virtual Storage Machines

Array Groups

An Array Group, which is also referred to as a Parity Group, is the basic unit of storage capacity for the VSP. Each Array Group can be configured as RAID-1, RAID-5 or RAID-6.

RAID-1 Array Group

Two types of RAID-1 Array Groups are supported: 2D+2D and 4D+4D. A 2D+2D group consists of a pair of disk drives in a mirrored configuration. The 4D+4D group combines two 2D+2D groups.

Data is striped to two drives and mirrored to the other two drives, with each stripe consisting of two data chunks. Each data chunk consists of either eight logical tracks for mainframe or 768 logical blocks for open systems. Hitachi states that the RAID-1 option is ideal for workloads with low cache-hit ratios.

RAID-5 Array Group

Also supported are two types of RAID-5 Array Groups: 3D+1P and 7D+1P. The data is written across the drives in a stripe that has three or seven data chunks and one parity chunk. As with RAID-1, each chunk in a mainframe environment contains eight logical tracks, while a chunk in open systems consists of 768 logical blocks.

RAID-6 Array Group

With RAID-6, blocks of data are striped across an array like RAID-5, except that it calculates two redundant sets of parity information for each stripe of data with Array Groups of 4D+2P, 6D+2P, 12D+2P, and 14D+2P. Hitachi states that the goal of RAID-6 is solely to improve fault tolerance since it can recover from the failure of two drives. It also states that RAID-6 is good for applications characterized by sequential reads on largest devices.

Device Emulation Intermix

A variety of device emulations for mainframe and open system environments is available. It supports intermix of all device emulations on the same back-end director pair, with the restriction that the devices in each Array Group have the same type of track geometry or format.

The Virtual LVI/LUN function enables different logical volume types to coexist. When it is not being used, an Array Group can be configured with only one device type (e.g. 3390-3 or 3390-9). When it is being used the user can intermix 3390 device types, and intermix OPEN-x devices types, but cannot intermix 3390 and OPEN device types.

The standard logical device types for open systems supported are listed in the following table.

LU Type	OPEN-3	OPEN-9	OPEN-E	OPEN-L	OPEN-V
Capacity (GB)	2.461	7.384	14.568	36.450	48.1MB to 64.422GB

Table 2: Standard LU Types for Open Systems

The VSP allows users to configure custom-size LUs that are smaller than standard LUs, as well as size-expanded LUs that are larger than standard ones. LUN Expansion (LUSE) volumes can range in size from two to thirty-six times larger than the standard size LUs. Sizes will vary from 4.92GB (OPEN-3 * 2) to a maximum of 1.312TB (OPEN-L * 36). The OPEN-V LU can vary in size from 48.1MB to 64.422GB.

Software Architecture

Storage Virtualization Operating System (SVOS)

The Storage Virtualization Operating System or SVOS is the main software element for the embedded system software that controls the VSP system. Included with SVOS are the integrated management tools required to manage the system. As basic management capabilities, these are not optional elements but can be augmented with additional advanced features. The integrated management features include:

Dynamic Provisioning

Dynamic Provisioning is the Hitachi implementation of thin provisioning for Hitachi storage. Capacity is only committed when space is actually being used. Dynamic Provisioning is used with internal storage and, with the VSP series; it can be used with external storage. A Dynamic Provisioning Pool can be created where volumes are assigned that are thinly provisioned whereby capacity is only committed as it is required. The parity groups are assigned to that pool and cannot be specifically selected when a dynamic provisioning volume is created.

Evaluator Group Comment: Thinly provisioned volumes eventually become "fat" over time. One way to retain "thinness" is to reclaim space that has been deleted from file systems or virtual machine volumes. This requires software or APIs with system software, typically filesystems, to identify the space to be reclaimed and issue a SCSI UNMAP command to the system. As part of volume migration, the VSP series will detect zero pages and assume those are available space that can be returned to the free space pool. Called Zero Page Reclaim, the function is part of the administrative action of either volume migration or administrators determining that space is not efficiently being used.

Dynamic Tiering

Dynamic Tiering is storage system-based software in Hitachi VSP G series systems to manage movement between different tiers of storage internal to the storage system. Storage may be high-capacity disks, high performance disks, or solid-state technology in the form of SSDs and FMDs. Tiering is not required with the all-flash F series.

Virtual Partition Manager

The Virtual Partition Manager enables the partitioning into managed Virtual Storage Machine of ports, cache, and disk parity groups for up to four partitions.

Performance Monitor

The Performance Monitor provides a graphical representation of the storage system performance for use with capacity and configuration planning, workload balancing, analyzing, and optimization.

LUN Manager / LUN Expansion

LUNs may be defined, configured, added, deleted, and reassigned to specific paths without the requirement for a reboot. It also supports the definition of paths, the reconfiguration of LUN-to-port assignments, and viewing of the remote service information messages. Logical unit size expansion is also supported with the Hitachi LUN Expansion (LUSE). This allows up to thirty-six physical LUNs to be presented to the operating system as a single, large LUN.

Two levels of password protection are available: user and administrator. Users can view only configuration information, while administrators can access all LUN configuration information and functions. Privileges for individual users can also be customized.

Virtual LVI / LUN Manager

Administrators can create custom size logical volumes when not using Dynamic Provisioning. This allows for improvement in utilization and remote copy performance when non-thin provisioned volumes are utilized. VLVI Manager has a mainframe-based utility that allows the customer to configure multiple logical volume images (LVIs). Each LVI can be as small as one cylinder or as large as a full 3390-9 volume, and definable in cylinder increments. Different types of LVIs can be defined within an array group.

Volume Port Security and Volume Port Security Option / SANitinel

As a security measure, Volume Port Security can be used to deny access to unauthorized users at the port level. The Hitachi Volume Port Security or SANtinel product controls host access to LUNs in open systems, multi-platform, and SAN environments. It enables users to restrict server access to authorized LUNs only.

In open system environments it allows users to define multiple Host Storage Domains, with up to 128 Host Storage Domains per physical Fibre Channel port on the system.

Connections from host servers arriving at the physical Fibre Channel port are routed to the logical Fibre Channel port within the appropriate Host Storage Domain based upon the World Wide Name of the host. Volume Security can also be used to authorize multiple hosts to access a particular Host Storage Domain. It can also be used to permit or deny a particular host to access individual LUNs within the Host Storage Domain.

For S/390 environments, this option is referred to as the Volume Security Port Option. It allows users to restrict host access to the logical devices and prevent other hosts from seeing the secured logical device or from accessing the data.

Server Priority Manager

Hitachi Server Priority Manager enables prioritization of a specific server's port access. This allows a single storage system to limit the I/O activity of non-critical servers and prioritize other servers that access the same port. Service levels can be established using the Server Priority manager software. The granularity of prioritization can be limiting the I/O rate (IOPS) or throughput (bandwidth). The prioritization can be established by port or by worldwide name.

Cache Residency Manager and Cache Management Host Agent

The Cache Residency Manager allows the user to ensure the highest possible performance for critical data. By "locking" the data in the cache it is assured a 100% hit ratio for all read and write operations. When using this feature, the performance is equivalent to that of a solid state disk. Hitachi refers to these portions of cache as cache extents, and they are defined by logical volume or LUN image. Cache extents are dynamic and can be added, deleted or changed at any time. This can be accomplished by the user via the Cache Manager or by Hitachi service representatives via the Service Processor. Each cache extent is defined as a range of logical blocks within a logical device or as an entire logical volume. Logical blocks are defined by starting and ending Logical Block Addresses (LBAs).

The user can select one of two modes of operation for Cache Residency Manager: Priority or Bind. The mode of operation is selectable for each cache extent.

Priority Mode

Hitachi states that Priority will be the normal mode for most users, and that it provides the most effective use of cache resources. In Priority mode the cache extents are not staged until the data is referenced. Therefore the first reference will be a miss, and a single copy of the data will be loaded into cache. All write data is de-staged to disk via standard algorithms, and it is duplexed to maintain full data integrity. Once referenced, the data will remain in cache.

Bind Mode

Bind mode provides the highest performance, but requires more cache resources. As with the Priority mode, data will not be staged until it has been referenced. The major difference is that write data, which is duplexed in cache, is not destaged to disk unless the cache extent is deleted or the subsystem is powered off.

Cache Requirements

The amount of cache available for Cache Residency Manager will vary based upon the installed cache size and the mode of operation. All installed cache greater than 512MB is accessible to Cache Residency Manager and Hitachi recommends that additional cache capacity be installed to support this feature. The cache requirements for Cache Residency Manager will depend upon the mode of operation and the type of logical volume image. The following table lists the cache requirements for each track of data. A cache segment is 16KB in size.

Mode of operation	Priority	Bind
Segments per mode in 3390, Open-9, -8, -3	4 (64KB)	16 (256KB)
Segments per mode in 3380, Open-K	3 (48KB)	12 (192KB)

Table 3: Cache Residency Manager Cache Requirements

Data Retention Utility for Open Systems and IBM z/OS

Volumes can be set into a WORM (Write Once Read Many) mode to prevent erasure or alteration of data. The setting is done through the management software and administered through the software.

Volume Shredder

Volume Shredder enables the shredding engine in the storage system to delete and overwrite data that has be designated as capable of being deleted after a retention period has expired.

Volume Shadow Copy Service (VSS)

The Microsoft VSS for applications is supported to utilize storage system-based point-in-time copies. With the support for VSS, the copy can be done in the storage system while the VSS capable application is temporarily is quiesced.

SNMP Agent

The SNMP agent enables SNMP reporting for Hitachi storage using Hitachi specific MIBs.

Database Validator

Hitachi implements the Oracle HARD technology to check the data transferred to ensure that the data is valid before it is written to the disk and is checked upon retrieval for additional integrity.

Virtual Partition Manager Extension

The extension takes the four-partition capability of the Virtual Storage manager and expands it to 32 partitions.

Universal Volume Manager

The Universal Volume Manager virtualized heterogeneous storage systems into a common storage pool. It also provides the common storage management across the storage systems and data migration and control capabilities.

Data Protection

Thin Image Snapshot

The Thin Image Snapshot capability is software that creates a pointer-based point-in-time snapshot copy with space only consumed when a write occurs by copying the original data to a new location. The space efficient snapshot can be created quickly and provide a significant capacity saving over a full copy.

In-System Replication

The In-System Replication feature is a bundle of replication functions meant for creation of copies of data within the storage system. The bundle includes Hitachi ShadowImage In-System Replication (clone replication) and Copy-on-Write Snapshot. All copies of data are RAID-protected and can be split from host applications.

Replication Manager

The Hitachi Replication Manager centralizes the management of the replication software for Hitachi storage systems in both open systems and mainframe environments. This includes both the In-Systems Replication as well as the remote replication solutions. Rather than have individual software management, the centralized management of Replication Manager combines the reporting, monitoring, configuration, and automation.

Role-based access controls are implemented with Replication Manager. Split mirror and consistency group management is also included.

TrueCopy Remote Replication

Hitachi TrueCopy Remote Replication consists of a suite of remote copy products that are unique to the VSP and USP series. TrueCopy Basic operates in synchronous mode and can duplicate from storage device to storage device any number and types of logical volume images from a primary source to a secondary source at distances up to 43 kilometers. Both source and target arrays must be Hitachi storage platforms.

Hitachi offers two versions of asynchronous remote copy: TrueCopy Asynchronous and Hitachi XRC. With TrueCopy Asynchronous, all records are time stamped to allow for the reassembling of data at the target site for data that is sequence-sensitive. It also supports Consistency Groups.

Hitachi XRC Replication

The Hitachi XRC Replication is a software based feature that is compatible with the IBM Extended Remote Copy (XRC) feature with the IBM Data Mover host software component. It is supported on the VSP and USP series.

XRC Replication supports primary and secondary systems at distances greater than 43 kilometers. The primary and secondary systems may be from different vendors provided they honor the IBM compatible software commands for remote copy implementation. It supports the IBM Data Mover host software component for primary, secondary and intermediate locations.

ShadowImage

The Hitachi ShadowImage product provides a means of non-disruptively producing "point-in-time" copies at the logical volume, or LUN, level on the same system for a variety of purposes, including application development and testing, and off-line backup. Through a technology exchange with IBM, the ShadowImage Flashcopy extension software solution responds to all IBM FlashCopy commands from TSO, DFSMSdss and SIBBATCH. This capability is called **Hitachi Compatible Mirroring software for IBM FlashCopy**.

As illustrated below, up to three target copies can be created in an open system or System z environment. Open system platform support also includes the ability to create two additional "cascade" copies per target volume, providing a total of up to 1,024 copies of a source volume.

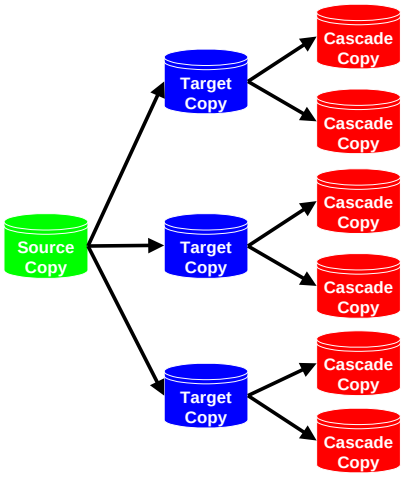

Figure 7: ShadowImage Copies

ShadowImage pairs can be created, split, resynchronized, suspended and deleted using the Remote Console. In a split state the source volume is not coping data to the target. While the pair is split, changed data to both the source and target volumes is kept track of in differential bit maps. When the pair is resynchronized the bit maps are compared and the changed data is sent to reconstruct the desired volume.

ShadowImage In-System Replication Manager

The Hitachi ShadowImage In-System Replication Manager feature creates a clone copy of a volume within a storage system. The copy will be RAID protected and is immediately available after the start of the copy operation. The full volume mirroring solution is non-disruptive and the mirrors may be split for independent usage. Consistency groups are supported for seamless dependent mirroring and splitting.

Universal Replicator

The Hitachi Universal Replicator software provides asynchronous replication for internal and external storage attached to the VSP and between VSP and USP systems. The feature support three data center replication configurations. The implementation uses journaling of data rather than caching in order to minimize system resource tied up during the asynchronous replication. Both mainframe and open systems are supported with Hitachi Universal Replicator.

	ShadowImage Clones	Thin Image Snapshot	TrueCopy	Universal Replicator
Copy Type	Full Clone	Bit Map	Synchronous Mirror	Asynchronous Delta
Max Copies	1,024 / LUN	1,024 / LUN 32K / System	64K	64K
Use for DR	Local only	Yes	Yes	Yes
Overhead	100%	Delta only	100%	100%
Available for Restore	After copy completes	Immediate	After Mirror split	After Mirror split
Access to Copy	After initialized	Immediate	After Mirror split	After Mirror split
Writeable	Source and Destination	Source and Destination	After Mirror split	After Mirror split
Re-sync	Optional	No	N/A	N/A
Performance Impact	During copy	Copy on write	I/O delay for ack	Limited

Table 4: Data Protection Features

Integration with Systems and Software

VMware Integration

Hitachi makes available a number of VMware vSphere "plug-ins" to integrate Hitachi storage and data management functionality with VMware environments. These include:

- vSphere Multi-pathing Plug-in for Hitachi Dynamic Link Manager
- Hitachi Tuning Manager VMware Plug-in
- Hitachi Command Director VMware Plug-in
- Hitachi Device Manager VMware Plug-in

Integration with VMware environments is also supported for Hitachi Dynamic Tiering. Hitachi Universal Replicator and True Copy Remote Replication support integration with VMware Site Recovery Manager (SRM).

Support for VMware vMotion is provided by Hitachi High Availability Manager, Hitachi Dynamic Link Manager on ESXi, and with Hitachi Global Active Device (active-active stretched clusters). vMotion of workloads across vCenter instances within the data center, across campus and metro distances, and across long distances is supported.

VMware VAAI Support

The VMware vStorage APIs for Array Integration (VAAI) are supported on the VSP. VAAI is a set of mechanisms that allow processing for certain data-related services—copying data when creating a new VM, for example—to be offloaded from the ESXi host to a storage array. The intent of these APIs is to streamline the functioning of the ESXi server and speed-up delivery of storage-supported services.

Full copy — Enables the storage system to make full copies of data within the storage system without having the ESX host read and write the data.

Block zeroing — Enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines.

Hardware-assisted locking — Provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large ESX host farms sharing a datastore.

Beginning with vSphere 5.0, enhancements for thin provisioning were added that are supported.

Dead Space Reclamation — Reclaims blocks from VMFS deleted files. The SCSI Unmap command is used to free physical space.

Out of Space Condition Support — A VM is paused if disk space is exhausted as reported by the storage system. This allows administrators to mitigate the situation rather than causing VM to fail.

Volume Shadow Copy Service (VSS)

The Microsoft VSS for applications is supported to utilize storage system-based point-in-time copies. With the support for VSS, the copy can be done in the storage system while the VSS capable application is temporarily is quiesced.

Microsoft Hyper-V Support

The VSP G series supports ODX – offload data transfer for Microsoft Hyper-V environments.

SMI-S Support

The VSP series supports SMI-S version 1.6. This level is required for Microsoft System Center Virtual Machine Manager to be able to manage the VSP system.

Supported Operating Systems

VSP Supported Operating Systems		
Open Systems	HP	▪ HP-UX ▪ Tru64 Unix ▪ Open VMS
	IBM	▪ AIX 5L
	Microsoft	▪ Windows Server (2003, 2008, 2012), Hyper-V
	Novell	▪ NetWare ▪ SUSE Linux Enterprise Server
	Red Hat	▪ Red Hat Enterprise Linux
	Oracle	▪ Solaris
	VMware	▪ VMware ESXi Server
Mainframe	IBM	▪ OS/390 ▪ z/OS ▪ z/OS e VSE/ESA ▪ z/VSE VM/ESA ▪ z/VM Linux for S/390 ▪ zSeries MVS/XA, MVS/ESA

Table 5: Supported Host Operating Systems

Advanced Features

Business Continuity Manager

The Business Continuity Manager provides replication management for IBM z/OS mainframe environments. The manager provides a single interface with similarity to ISPF panels to automate Hitachi replication products including Universal Replication, ShadowImage, In-System Replication, and TrueCopy Remote Replication. Business Continuity Manager can also access system performance data for monitoring and management. The auto-discovery feature for replication objects can be used in configuring the replication tasks.

The Business Continuity Manager allows administrators to define copy groups that contain multiple replication objects that can be managed with a single comment. The real-time status of replication objects can be viewed and the replication metrics monitored for performance.

Cross-OS File Exchange

Cross-OS File Exchange allows the one way transfer of data from a mainframe to an open systems platform. It consists of the following components.

- Embedded SVOS code
- Hitachi supplied open system File Access Library (FAL)
- Hitachi supplied open system File Conversion Utility (FCU)
- Minimum of two shared logical volumes designated as 3390-3A or 3390-3B

The shared logical volumes are accessed by S/390 via normal access methods through the FICON interfaces, and by open systems with the Hitachi supplied FAL API through the fibre channel interfaces. They must be defined as either a 3390-3A or 3390-3B logical volume. A 3390-3A can be used for bi-directional operations (System z to open and open to System z). A 3390-B can only be used for uni-directional operations (System z to open). Although these volumes can be used by non-HDME tasks, they are typically dedicated to Cross-OS File Exchange since they cannot be managed by SMS.

Any data that can be unloaded, created, or copied as a sequential dataset in z/OS can be read directly by a UNIX server. It uses an application program similar to DB2 extractor to create a sequential flat file on a specific device on the system. Cross-OS File Exchange also includes a Hitachi supplied File Access Library (HFAL), containing subroutines to handle the CKD flat file. These routines read the mainframe passed data and transfer the data to the requesting open systems platform application, or create an open platform data file.

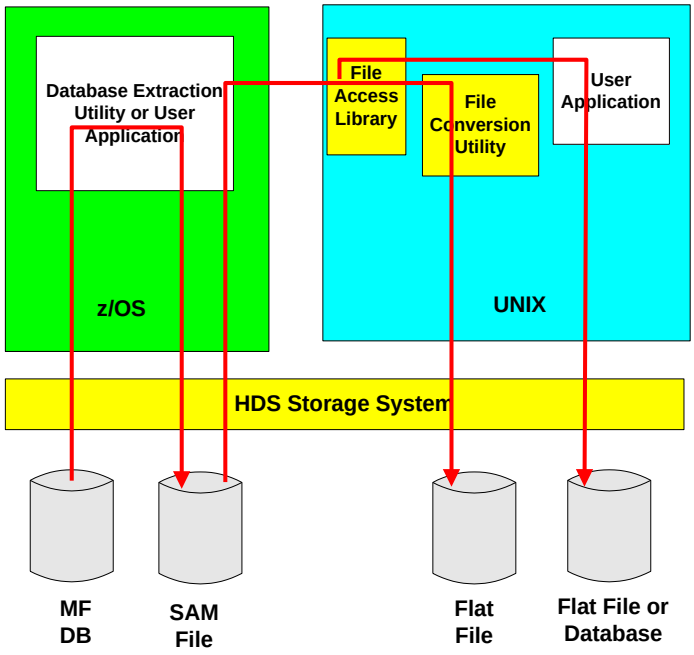

Figure 8: Cross-OS File Exchange

Command Control Interface

The Hitachi Command Control Interface (CCI) feature allows open systems users to create batch streams to manage Hitachi TrueCopy and ShadowImage from programs running on the host. The CCI software interfaces with the system and high-availability software on the UNIX/PC server host, as well as the TrueCopy and ShadowImage software on the VSP G series and/or previous generations. It provides failover and other functions such as backup commands.

The CCI functions can also be integrated with host-based failover software such as HP MC/Service Guard, FirstWatch and IBM HACMP.

The CCI interface is not used for SMI-S support. The Hitachi Device Manager interface is the route for SMI-S support within Hitachi Storage Products.

Four Data Center Multi-target and Cascade Support for VSP

Hitachi TrueCopy and Universal Replicator also support four data center disaster recovery configurations with SVOS-based systems in a multi datacenter implementation. Synchronous copy between two short distance sites is accomplished using TrueCopy while two long distance sites are supported using Universal Replicator asynchronous.

Encryption Support

FIPS 140-2 compliant Encryption is available with the VSP. There are two forms of data encryption with the VSP. One is with the encryption-enabled FMDs (FMD-HDE) where the encryption is offloaded to the processing in the FMD. The second is where encryption is performed on data in the backend director hardware. In both cases, keys may be managed internally or be exported for management by a KMIP. Enablement for the FMD encryption is through a software license.

Hitachi claims no performance degradation when enabling the controller-based or FMD encryption. The SAS HBA to connect internal storage enclosures performs the encryption. Encrypting HBA's are an optional feature. Encryption can be applied to any media type being used as internal storage within the array.

Administrative Security

External authentication of privileged users (i.e., storage administrators) allows for the VSP to be integrated with existing authentication data. Users can provision and de-provision storage management access with existing infrastructure. Active Directory and the RADIUS protocol are supported for authenticating against a wide variety of authentication sources.

Multiplatform Backup

Multiplatform Backup supports volume level backup and restore operations between S/390 and open systems platforms. Any open systems LUN on the system can be backed up to MVS or VSE system using standard utilities such as DFDSS, FDR, or VSE FASTWRITE.

The backup process is initiated by making the LUN unavailable to the open systems platform. The mainframe device is then varied on-line and a copy is created. When the copy operation is completed, the device is varied off-line to the mainframe and the LUN is made available to the open platform. To perform a restore, the mainframe software recreates the required data while the device is unavailable to the open system.

Compatible PAV for IBM z/OS

The Hitachi Compatible PAV (Parallel Access Volumes) for IBM z/OS feature allows multiple applications running on an IBM z/OS system to access the same information simultaneously on the storage system. Multiple Allegiance extends the simultaneous access to multiple z/OS systems and applications. Allowing simultaneous access reduces queuing or wait time on the storage resource, which can be very significance in high performance transaction environments.

Compatible HyperPAV for IBM z/OS

The Hitachi Compatible HyperPAV for IBM z/OS feature is an extension of the Compatible PAV feature that provides the management of multiple I/O requests after every I/O simultaneously enabling faster reaction to workload changes.

Reliability, Availability, and Serviceability Features

System Monitoring Network

The VSP is designed with a system monitoring network. The monitoring cables are connected to all of the cabinets in a system, and report a variety of system conditions such as component failure, fan speed, voltages, and temperature. This network provides predictive maintenance as well as failure alert capabilities. The information is passed to the Control Chassis where it can be viewed either by the Service Processor or remotely. System status and alerts are also sent to the Hitachi Customer Support Centers via the Hi-Track "call-home" service/remote maintenance tool.

Service Processor

A Service Processor (SVP) is integrated into the controller chassis. It can only be used by authorized service personnel to configure, maintain and upgrade the system. It also collects performance data for all of the key system components for diagnostic testing and analysis. A single service processor is standard and a redundant service processor is an option.

Redundant Power Supplies

Each storage cluster is powered by its own set of redundant power supplies, with each power supply capable of providing power for the entire system if required.

Global Cache Data Protection

Global Cache caches application data blocks from drives via the BEDs during read operations, and application data blocks from FEDs during write operations. Global cache retains only one copy of read data whereas all write data is mirrored in cache. If a power or system failure occurs, data and system metadata in cache is protected by backing up cache to onboard flash drives. Two mirrored data copies are kept in the backup flash drives.

Backup Batteries

Each frame contains internal nickel-hydride batteries that provide backup power for the cache memory, shared memory, front-end and back-end directors, and disk drives.

System Management

Hitachi Storage Advisor

The Hitachi Storage Advisor (formerly the Hitachi Infrastructure Director) is a configuration management application for the VSP G series midrange systems. A new GUI simplifies the administration and provides recommended configuration practices. The recommendations are Hitachi's best practices for the VSP systems. Key to the simplifying the management is the provisioning operations guidance through the GUI. Storage resources are automatically grouped into tiered storage profiles for service level selection when provisioning volumes for applications. Many of the storage tasks in managing a high-function system like VSP are automated including pool creation, port assignment, zone creation, and parity group initialization.

The Hitachi Storage Advisor is the element manager for management operations across the midrange VSP series. A REST-based API and CLI is available for management programmability. A detailed system log file of all administrative action is maintained for change management controls and problem determination.

Hitachi Infrastructure Analytics Advisor

The Hitachi Infrastructure Analytics Advisor HIAA is a SaaS-based offering that analyzes telemetry data to detect anomalies for preventive actions, predict issues when planning for changes with use of historical data, provide correlation to infrastructure changes and performance, and monitor elements for optimizing performance. HIAA allows analysis and resolution with the collected data. Reports are generated from the multi-vendor devices monitored.

HIAA monitors the topology across the storage environment. The monitoring can be done from an application, network, or server level to storage. It can report on service level management for storage volumes. Real-time analytics are used for problem isolation.

Hitachi Automation Director

Hitachi Automation Director (HAD) expands beyond the specific Hitachi devices to automate many of the tasks for the infrastructure. Virtual machine (VMware, Hyper-V, and KVM) resources may be automatically configured and assigned. Switch zoning for Brocade and Cisco can be set with resiliency and performance parameters. Specific applications can have resources assigned and QoS set at the LUN level. Data protection policies for the Hitachi Data Instance Director can be set based on best practices.

Performance

When properly configured, the VSP 5000 is an extreme performance system. Currently, no SPC numbers are available and only Hitachi reported numbers for cache hit reads have been released.

Evaluator Group believes that the best source currently available for providing comparative performance data is the Storage Performance Council (SPC). Hitachi is a member of this organization so it should be consulted to obtain the latest information.

Evaluator Group comments

Strengths:

The VSP 5000 is a logical hardware update to the earlier hybrid G1500 and the all flash F1500 systems. The hardware updates solve some internal limitations and repackages the controllers with new processors and logic. The new fabric is an implementation of a custom PCIe switch in an FPGA. The same SVOS embedded software executes on the new controllers.

With this update, the VSP again sets the bar for high-end shared storage systems and are in a narrow group that can also allow mainframe connectivity. There are additional capabilities that expand the value of the software but there is fundamental shift in the direction of Hitachi with their storage systems:

- ***Separation of hardware and embedded software into unique elements:***
    - ***The embedded software called SVOS is single code stream for Hitachi going forward. It is used on both the high-end VSP systems and the midrange VSP models. This means that new developments with the software can be leveraged to add value across the entire product line. It also brings focus to support and allows the resources to be applied more efficiently.***
    - ***With the hardware independent of the software, the hardware designs will have a defined operating environment, which will lead to greater commonality. This will allow greater flexibility in delivering new hardware platforms with advances in the processor and interface technology.***

Other high value capabilities of the high-end VSP include:

- ***Global Active Device, which is an active-active stretched cluster implementation enabling critical functionality for business continuity and workload mobility.***
- ***Advances in system management and storage resource management.***
- ***Non-disruptive migration to allow for seamless, transparent scaling. This feature will be exploited with technology upgrades that enable longer lifespans for the hardware systems.***

The Hitachi VSP continues to be one of the most reliable, full features, highest performing systems available for top tier enterprises. Advances made will continue to draw customers with those needs.

Perceived Challenges:

There are not many challenges for the VSP 5000 being successful from a product standpoint. The issues may be focused on the awareness of the capabilities, especially the more recent one. This requires crisp and focused messages about the products and usage environments.

The removal of FMDs from the latest generation of VSP 5000 systems will help to clarify some of the confusion regarding the support for both FMDs and SSDs in previous models.

More detailed information is available at http://evaluatorgroup.com