The Virtual Storage Platform (VSP) from Hitachi Data Systems–Setting New Levels of Excellence

Josh Krischer
24-September-2010

The Virtual Storage Platform (VSP) from Hitachi Data Systems–Setting New Levels of Excellence

Contents

Josh Krischer is an expert IT advisor with more than 40 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Krischer & Associates GmbH, he was formerly a Research Vice President at Gartner, covering enterprise servers and storage from 1998 until 2007. During his career at Gartner he covered high-end storage-subsystems and spoke on this topic at a multitude of worldwide IT events, including Gartner conferences and symposia, industry and educational conferences, and major vendor events.

Announcement

On September 27th, 2010, Hitachi Data Systems, a full subsidiary of Hitachi Ltd., Japan, announced its new high-end Virtual Storage Platform (VSP) as the follow-on model for its technology leading USP-V and VM subsystems. The new VSP is market leading in scalability, availability, performance, and storage economics. It is the most automated storage subsystem and most suitable for large-scale server-consolidation and deployments infrastructures for public clouds.

Users' Requirements and Challenges

Despite the tough economical situation of the last couple of years, the demand for storage capacity has remained almost unchanged. What has changed is the set of different procurement processes, new purchasing criteria, and higher acceptance of tiered storage infrastructures. New technologies such as thin provisioning or Solid State Disks (SSD) re-emerged and changed storage landscapes. New criteria such as virtualization, cloud-readiness, energy consumption, and cost containment have emerged. What do users typically require from a storage subsystem?

Availability and Business Continuity remain on top of the list of requirements. The non-stop global economy, the fierce competition, and new levels of service requirements raise the requirements for business continuity. Disaster recovery continues to be a major issue. Europe started to deploy disaster recovery in the late 80ies, whereas it would not play a major role In the US until much later.

For many American companies, awareness of the need for business continuity/disaster recovery plans was partially triggered by the tragic events of September 11th, Hurricane Katrina, and the north-east blackout of August 2003 (a massive power outage that occurred throughout parts of the north-eastern and mid-western US, and Ontario, Canada). The 2004 Hurricane season in Florida presented situations that had not been experienced before anywhere in North America, further emphasizing the need for disaster recovery. Another trigger was the introduction of various compliance regulations, such as Sarbanes-Oxley or SEC 17a(4). Deploying disaster recovery infrastructures for organizations with heterogonous, multi-vendor, or multi-platform storage-subsystems still poses a number of challenges today, such as incompatible remote copy techniques, for example. Additional requirements still not met by some products are that service, capacity upgrades, testing, etc., should be performed non-disruptively, and completely transparent to users.

Scalability: A storage subsystem should be able to scale seamlessly in capacity, connectivity, or performance without lowering service levels. The capacity should support multi-tier storage media: SSDs, performance HDDs in different capacity (FC or SAS), and capacity SATA disks – in simple words, tiered storage "in-a-box".

Secure Multi-Tenancy: Server virtualization creates situations where multiple tenants are using the same physical infrastructure, which requires new techniques to address security issues. Public cloud-shared storage needs to ensure stringent security for each user's data, or otherwise the users require separate subsystems, which increases service costs and makes the concept less attractive.

Performance has two aspects: the first is the throughput measured in number of I/O per second (IOS), and the second is the response time measured in milliseconds. Performance should be at the SLAs level regardless of the used capacity and the workload. Erratic performance levels irritate users more than a slightly slower but constant response time.

Functionality and automation are required to cope with storage management. The average organization's storage capacity grows by approximately 50-60 percent per year while the size of the storage management staff generally remains the same. Datacenters which supported terabytes at the end of the previous decade today support petabytes. The only way to cope with this capacity explosion is sophisticated functionality, advanced automation, and user-friendly management tools and interfaces.

Storage Efficiency in usage and energy consumption is equally important. Storage efficiency can be achieved by using thin provisioning, tiered storage, automated data placement,, compression, small-factor HDDs, SSDs and virtualization. Efficient subsystems allow for better storage utilization, which translates to lower capacity and lower CapEx and OpEx, reduced floor space requirements, lower energy consumption and effective usage of people.

VSP Control Unit Architecture

The Virtual Storage Platform (VSP), unlike other high-end storage subsystems, is a purpose-built storage array. The architecture and many of the components are specially designed and produced by different divisions of Hitachi Ltd. The 'tailored' design is one of the factors in achieving its very high performance, large scalability, and unmatched reliability. In contrast, the other high-end storage subsystem use storage control programs which run on stripped-down operating systems such as AIX and Linux, and commercial server hardware.

"The 'tailored' design is one of the factors in achieving its very high performance, large scalability, and unmatched reliability."

As opposed to previous Hitachi high-end products, there is only one model which scales out by combining resources in up to two control chassis. Each control chassis is housed in a standard 19" rack and each control up to two disk expansion racks. The scalability is achieved by adding boards (in pairs) to the chassis and HDDs or SSDs in storage units or cabinets. See single- control chassis block diagram in Figure 1.

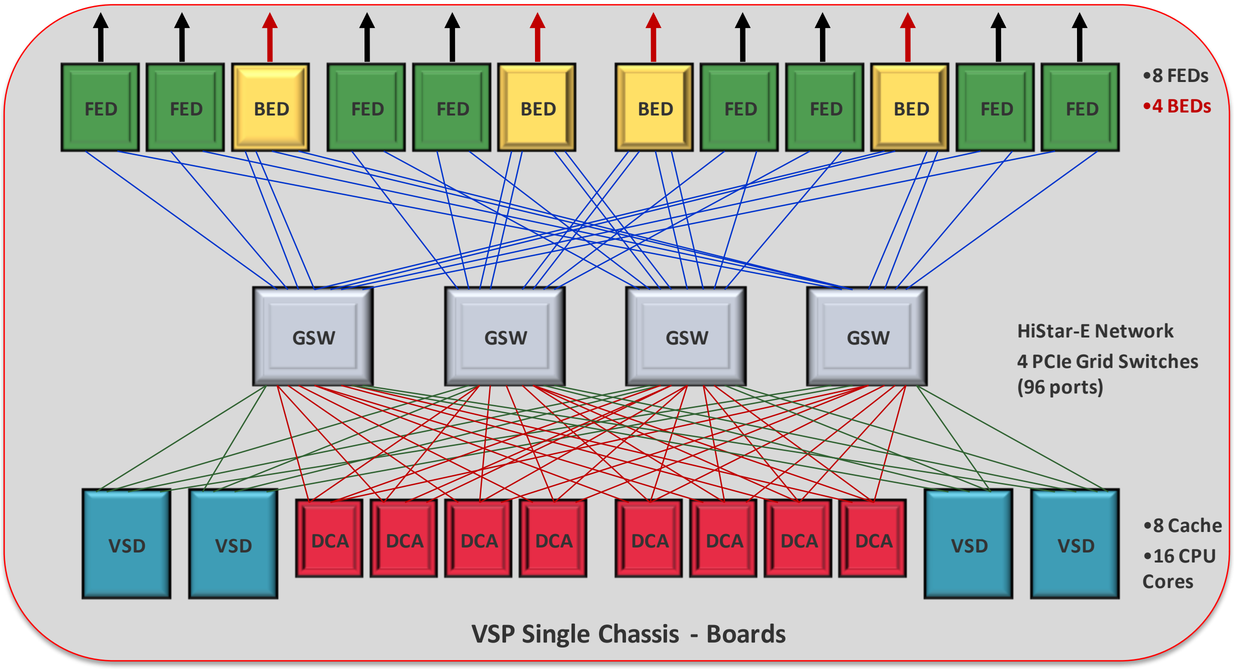

Figure 1 : VSP Single Chassis block diagram/Boards, source: Hitachi Data Systems

FED: Front-end Director - (host interface)
BED: Back-end Director - (disk interface)
VSD: Virtual Storage Directors – (central processor boards)
DCA: Cache Memory Adapter
GSW: Grid Switch

The control chassis hosts the CU hardware containing all the control boards in a single rack and one or two drive chassis. The expansion racks can hold up to three drive chassis each. Each drive chassis supports up to 80 3½ inch HDDs or SSDs, or 128 small-factor 2½ inch HDDs or SSDs. Two control units are cross-connected at the grid switch level across racks and operate as a single integrated array as opposed to a cluster.

The Virtual Storage Director (VSD) is a powerful data movement engine and the heart of the VSD. There are two or four VSDs installed per control chassis. Each board contains a quad-core Xeon CPU with 4GB of local RAM. These boards manage the VSP's operation outside of the data path. The difference between the two and four VSDs configuration is in performance, and it is unrelated to the number of installed FEDs or BEDs. In case of a failure, the other VSD (out of the pair) reads the control table information and meta data from

a master control-memory copy maintained in the system cache, and takes over the functionality of the failing module. This procedure is reversed after a VSD swap.
If the VSD is the heart of the subsystem, then the **Grid Switches** (**GSWs**) represent its blood vessels. The GSWs span the *HiStar-E Network*, a PCI-Express-based switch system. The HiStar-E Network is a fully fault-tolerant, high-performance, truly non-blocking switched architecture. The HiStar-E network provides the high speed cross-connections among all controller boards. There can be two or four GSWs installed per chassis. Each board has 24 high speed ports. Each port supports a 1024MB/s sending rate and a concurrent 1024MB/s receiving rate. Eight of the 24 ports are used to connect to the FED and BED ports. These ports transfer both data and metadata traffic inter-mixed. Four more ports are used to connect to the VSD boards. These ports transfer job request traffic or system metadata (such as the in-cache address for a target slot). Eight other GSW ports are attached to DCA boards, for data transfer (user read and writes) and control-memory updates. The last four ports of the 24 GSW board ports are used to cross-connect that GSW to the matching board in the second chassis of the dual chassis configuration.

The **Data Cache Adapter** (**DCA**) boards are the memory boards (cache) for read/write data and the control-memory (metadata) master copy. There are up to 8 DCAs installed per chassis, with 8GB to 32GB of cache per board, which yields cache sizes from 64 GB up to 256 GB per chassis and 512 GB for a fully-configured subsystem. Each DCA board also has one or two on-board SSD drives (32GB each) for use in backups of the control-memory regions, the cache directories, and all write-pending buffers, which are important in the event of an array shutdown. Write operations are simultaneously duplexed into VSD segments in both cache-A and cache-B by the FED data accelerator processor handling the request. The cache is not fully mirrored, only the individual write blocks are mirrored, which means that users have more cache for their applications. The DCA boards contain control-memory that contains all information about the array configuration details, the RAID groups and LUNs, all tables used to manage external I/O and remote-copy operations, and all dynamic-provision and dynamic-tiering (see explanations in later sections) control information.

The **Front-end Director (FED)** board controls the connection between the array and servers attached to the host ports, as well as connections to external storage or remote copy products (TrueCopy and Universal Replicator). There are four types of FED boards available (always installed in pairs of the same type): 16-port FC (8Gbps[1]), and 16-port FICON. The various types of interface options can be by mixed on a pair basis. In addition to the host and remote copy connections, the FEDs also support Hitachi or third-party storage-subsystems with their embedded virtualization functionality. This feature provides the ability to configure, manage, and access external volumes as if they were VSP internal volumes.

The **Back-end Director** (**BED**) boards act as the interface to the HDDs or SSDs installed in the DKUs. Two or four BED boards can be installed per chassis. Each BED board has eight

[1] Each port can auto-negotiate to one of three host rates: 2Gbps, 4Gbps, or 8Gbps.

6Gbps SAS links. A standard-performance single-chassis configuration uses one BED pair with 16 back-end 6Gbps SAS links. The high-performance single-chassis uses an additional pair for a total of 32 6Gbps SAS links.

In addition to data transfers, the BED boards generate RAID-5 and RAID-6 parity and perform drive rebuilds by an embedded XOR processor. The VSP supports RAID-10 (mirrored data striping), RAID-5 (data striping with distributed parity) and RAID-6 (data striping with two distributed parity disks), with RAID-10 and RAID-5 supporting 4 or 8 disks in a RAID group; a RAID-6 group contains 8 disks. Any disk can be configured for any RAID technique, though it is recommended to configure SSDs as RAID-5 and SATA disks as RAID-6.

Scalability–Hitachi's Unique 3-Dimensional Scalability

There are three ways to increase performance, connectivity, and storage capacity in the VSP:

- Scale **Up** by adding hardware boards within a chassis and HDDs or SSDs for increased performance, connectivity and capacity.
- Scale **Out** (dual-chassis arrays) for doubled internal capacity, connectivity, and better performance
- Scale **Deep** by adding external arrays (Hitachi or third party) controlled by the embedded virtualization layer

Scale Up

The VSP subsystem can be built with Lego stones-like modularity. Users may scale up or tailor their subsystem to best match their requirements. Scaling capacity, adding host connections, or increasing performance can be performed independently. The entry-level storage configuration may start with 10 logic boards (2 VSPs, 2 DCAs, 2 FEDs, 2 BEDs, and 2 GSWs) and one disk container[2], then more boards (up to 28) and disk containers (up to eight) can be added. There are two types of disk containers (DKUs): a 128 disk Small Form Factor (2.5" SFF) and an 80-disk Large Form Factor (3.5" LFF) model. There are four models of small-factor performance media: a 200GB SSD, a 146GB 15K, a 300GB 10k, and a 600GB 10K, all using a 6Gbps SAS interface. The large-factor units may contain 400GB SAS SSDs or 2TB SATA disks. Since the SATA disks only support 3Gbps, a speed matching logic is used[3]. In a pure SFF configuration, a single control chassis array may contain up to 1024 disks while for LFF drives, the limit is 640 disks.

[2] Hitachi sells also a disk-less , BED-less unit as a "virtualization engine". Additional FEDs can be installed (instead of BEDs) to increase the connectivity.

[3] All SAS drives are dual interface attached with one drive port to each switch. SATA drives have a special disk canister that has a logic board to expand the single SATA disk port to a dual SAS port.

Scale Out

The VSP's scalability can be doubled by the use of a dual-chassis array, which is able to deliver close to double the throughput, double host connectivity and supports double the capacity of a single control chassis. The tight connection between the sides (it is not a cluster) enables any host port direct access to any back-end RAID group.

Scale Deep or Out-of-the- Box Scale

The VSP supports the same embedded virtualization mechanisms as the USP V/VM. External storage subsystems connected to a VSP are controlled by Hitachi's Universal Volume Manager software. This feature provides the ability to configure, manage, and access external volumes as if they were VSP internal volumes and provides transparent access to internal and external volumes. Storage management tools, data replication software, and other applications can be used in the same way, regardless of whether the data resides internally on the VSP or on an external third-party storage subsystem's volumes.
The use of externally-connected storage allows the user to build tiered storage with outdated subsystems while making use of all the benefits of the VSP's functionality and virtualization (e.g., easy data migration), as well as unified disaster recovery with Hitachi Universal Remote copy.

Functionality

In addition to new functionality, the VSP supports all the functionality of the previous Hitachi Universal Storage Platform (USP), which eases migration to the VSP and eliminates storage management training time and costs. Some of these important functions are explained later in this section.

New Functionality

The most important new functionality expanding upon the Hitachi Dynamic Provisioning feature is ***Hitachi Dynamic Tiering***. In average 20 percent of data accounts for 80 percent of activity which justify tiered storage deployments. However, one ability that users require of storage is to simplify the building of tiered storage infrastructures. Using tiered storage gains higher importance in times when IT budgets remain unchanged or get slashed. One of the tasks when deploying a tiered infrastructure is to determine the appropriate storage tier for the data. An SSD is made up of flash memory that emulates an HDD. There is no dispute that the performance of enterprise solid-state disks is higher than that of traditional magnetic disks (no mechanical movements, no disk rotation and seek latency delays), however the higher-cost factor (x10-20) prevents their use as general-purpose disks. The optimal usage scenario for SSDs is data that is cache-unfriendly to read (random access, non-sequential). Typical candidates include hot database tables, database temporary areas, metadata, indices, control areas, etc. The question is how to identify such data? Placing a whole volume or LUN on an SSD may be not economical, and partial data placement cannot be manually controlled by operators.

Hitachi's VSP is a tiered storage in-a-box: it supports SSDs, several performance SAS HDDs, capacity SATA disks, and different RAID techniques. Hitachi Dynamic Tiering is a dynamically automated data placement feature that, based on actual usage, constantly classifies and migrates data to the most suitable tier. Data chunklets of 42MB are dynamically placed in the appropriate tier based on host access and usage patterns. This feature helps maintain peak performance under dynamic conditions without manual intervention at optimal CapEx and OpEx savings. The time of the data placement and the sampling frequency can be user-set, and, in addition, users retain the option to move data manually to a selected platform if required (e.g., high-performance financial application running once a month).

"Hitachi Dynamic Tiering is a dynamically automated data placement feature that, based on actual usage, constantly classifies and migrates data to the most suitable tier."

Existing Interesting or Unique Features

Hitachi Dynamic Provisioning (HDP) was introduced in May 2007 for the USP V/VM arrays and enhanced in the VSP. In addition to thin provisioning, it also enables wide striping across all member LUNS in its pool for the creation of high-performance volumes. Striping the data among a large number of physical devices practically eliminates hot spots and allows parallel access to disks, which results in almost uniform performance.

Thin provisioning enables the allocation of virtual storage as needed without having to dedicate the full physical disk storage capacity up-front. Additional capacity can be allocated without any disruption to mission-critical applications from existing or newly-installed capacity. The *Zero Page Reclaim* function returns unused storage blocks back to the storage pool and reclaims storage space, while *Automatic Dynamic Rebalancing* automatically re-stripes existing virtual volumes when physical volumes are added to the pool for workload rebalancing. Hitachi is the only high-end storage vendor that offers automatic rebalancing of the virtual volume pages through active re-striping in order to take advantage of new disks when the pool is expanded. Dynamic Provisioning capabilities can also be used with external third-party attached storage systems. Hitachi Dynamic Provisioning increases installed capacity utilization, defers upgrades, and saves storage management tasks, thus reducing CapEx and OpEx.

More than 50 percent of USP V/VM users have deployed the Dynamic Provisioning feature, saving 30-40 percent of real capacity on average.

A unique feature is the **Hitachi Universal Virtualization Layer** (introduced with the first version of the USP in 2004). The Hitachi *Universal Volume Manager* software configures, manages, and accesses external volumes in a similar way as if they were VSP or USP-V internal volumes. Externally-connected storage may use the same functionality as internal storage, which means that data replication software, transparent data movements, dynamic provisioning, and other high-end subsystems features can be used in the same way, regardless of whether the data resides on internal or external volumes. The virtualization of heterogeneous storage systems simplifies storage management, enables easier migrations[4], reduces the complexity of disaster recovery schemes, and allows the building of tiered storage without compromising on functionality. Users may extend the life of outdated storage subsystems by giving them high-end storage functionality.

A disk-less VSP can be used as an enterprise storage-virtualization appliance. Such a subsystem does not need back-end BED boards and may use the available slots for additional FEDs to increase the number of host ports.

Two other important security features are **Virtual Ports** and **Data Encryption**. Server virtualization creates situations where multiple virtual partitions share the same physical port to access data, so virtual ports increase security by enabling proper zoning. Virtual ports have been unique feature for the last two generations of Hitachi high-end and mid-range storage subsystems. Similar to the Fibre Channel NPIV feature, it allows a single Fibre Channel port to appear as multiple distinct ports providing separate port identification and security zoning within the fabric for each operating system image. It appears as if each image had its own unique physical port, which helps administrators manage virtual servers on a FC SAN. Virtual ports allow each server to have a unique WWN (just like a real server).

A hardware-based (storage controller) encryption option encrypts the internal drives using strong encryption (AES-256) without impacting throughput or latency. Hitachi's data-at-rest encryption feature also includes integrated user-friendly key management functionality.

Availability

Another unique feature introduced in 2009 is the **Hitachi High Availability Manager**–a local and remote storage array clustering option that allows non-disruptive failover across VSP and USP V storage subsystems. This option enables instant data access recovery at the primary site or at a remote site in case the primary subsystem malfunctions or is unavailable. Unlike the comparable mainframe *HyperSwap* software, this feature is implemented in hardware and targeted at non-mainframe SAN infrastructures. Hitachi High Availability Manager will also, in the near future, provide simplified, non disruptive, easy migration from USP-V, VM to VSP subsystems.

The USP series provides full redundancy, non-disruptive upgrades, and maintenance, hot-swappable components, and online microcode changes. Hitachi storage subsystems were the first devices to support non-stop operation, and the VSP maintains HDS's position as the industry’s availability leader for many storage subsystem generations already.

The HDDs are carefully screened before installation. In addition to the three RAID techniques mentioned above, Hitachi uses extensive error correction on each HDD sector. 8 bytes of error correction code (ECC) are stored with each 512 Bytes data sector and are used by the

[4] Data migration, in general, is one of the most complicated storage-related tasks. It can cause data loss, generally takes a long time, and is usually very expensive. It may require retaining obsolete subsystems for longer than planned and additional costs through leasing rates and maintenance.

drive when trying to reconstruct lost bits of data during a read operation. Pre-emptive techniques detect drive malfunctions as early as possible, which allows copying a suspected drive before having to perform an entire RAID rebuild[5]. Small-factor HDDs with smaller capacities allow for a fast disk copy or rebuild as well.

Write operations are simultaneously duplexed into VSD segments in both Cache-A and Cache-B by the FED Data Accelerator processor processing the request. In case of a main power loss or malfunction, the on-card batteries maintain power to enable flushing the outstanding write blocks to the internal SSDs.

Like its predecessor, the VSP supports Hitachi's *TrueCopy Synchronous* remote mirroring technique and *Universal Replicator*, which enables hardware-controlled asynchronous remote copying between heterogeneous storage platforms. The VSP fully supports Geographically Dispersed Parallel Sysplex (GDPS), which is a mainframe multi-site application availability solution with fast recovery time and highly-automated control. Many of the world's largest organizations in financial services and other key industries have deployed GDPS to protect mission-critical applications. It manages application availability in and across sites for both planned maintenance and unplanned situations such as site failures or full-blown disasters. GDPS was initially designed for mainframe z/OS systems but, with continuous development, was later enhanced to support select open systems platforms, as well. Many of the existing GDPS installation use Hitachi disks.

Performance

Early measurements show significant performance improvements in comparison to the USP-V, which currently holds the SPC records. Storage Performance Council (SPC) benchmarks[6] are modeled on real-world applications, and therefore, help provide customers with meaningful performance results.

SPC-1[7] simulates transactional operation and SPC-2 [8]simulates sequential access. The performance figures achieved by the USP V are 200,245.73 IOPS in the SPC-1 benchmark and an aggregated average of 8,724.67 MB/sec in SPC-2.

[5] To copy a 2 TB HDD may take 2-3 hours, to rebuild without load ca. 48 hours. A rebuild with 80% CU load can more than a week.

[6] The SPC is a non-profit corporation founded to define, standardize, and promote storage system benchmarks, as well as to disseminate objective, verifiable performance data to the computer industry and its customers. SPC membership is open to companies, academic institutions, and individuals. The SPC created the first industry-standard performance benchmark in 2001 targeted at the needs and concerns of the storage industry, and its goal is to serve as a catalyst for performance improvement in storage.

[7] SPC-1 is designed to demonstrate the performance of a storage component product while performing the typical functions of business-critical applications. Such applications are characterized by predominately random I/O operations and require both queries as well as update operations. Examples of these types of applications include OLTP, database operations, and mail server implementations.

[8] The SPC-2 benchmark consists of three distinct workloads designed to demonstrate the performance of a storage subsystem during the execution of business critical applications that require the large-scale, sequential movement of data.

The VSP's performance improvements compared to the USP are due to the new control unit structure and the use of more powerful technology which contribute to increased bandwidth and lower latency throughout the system

. For example:

- VSDs as a third layer of processing (in addition to FEDs and BEDs)
- The newly designed HiStar-E PCI Express Switched Grid replaced the Universal Star Network crossbar switch.
- Special Data Accelerator processors and new ASICs for the FED and BED boards
- The 800MHz NEC RISC CPUs and support chips (on USP V FEDs and BEDs) have been replaced by 2.33GHz Intel Xeon Core Duo quad-core CPUs
- BEDs' dual SAS controllers, with eight 6Gbps SAS links per BED board, instead of four 4Gbps FC loops as on the USP V.

Technology

The USP-V is designed and manufactured by Hitachi Ltd. A huge leading international technology corporation with 360.000 employees in business year 2009 (ending 31 March 2010) and consolidated revenues of ca. US$ 96.4B. It manufactures a wide range of products from medical equipment, to fast trains, and turbines for electric plants. The corporation exploits synergies between the different divisions to produce state-of-art products. It takes several years to design a new storage subsystem, which is why only a company with close relations with disk and semiconductor technology manufactures can plan for using the latest technologies. For example, Hitachi's Telecommunication & Network division contributed knowledge in designing the HiStar-E PCI Express Switched Grid, Hitachi Global storage Technologies, the world-wide third-largest HDD manufacturer and leader in 2½ inch HDDs cooperated on media design, while the Micro Device division supplied the tailored ASICs and other components.

Hitachi Storage Command Suite v7

Hitachi's Storage Command Suite combines advanced data management and protection capabilities in a simple, common (GUI and CLI) user-friendly centralized interface to manage the Hitachi dynamic storage infrastructure, whether it is comprised of virtualized heterogeneous storage or non-virtualized homogeneous storage. All other large storage vendors use different storage management applications for each platform as opposed to this single interface which manages the complete Hitachi Data Systems storage portfolio, from the high-end VSP to the HCAP archiving platform. Users are able to see all their heterogeneous storage assets as a unified pool as opposed to managing disparate islands of storage, which dramatically improves productivity, reduces risks of losing data, and lowers storage management costs. Version 7 of the Command Suite contains many enhancements that simplify usage and enhance functionality.

"single interface which manages the complete Hitachi Data Systems storage portfolio, from the high-end VSP to the HCAP archiving platform"

The Storage Command Suite includes:

Hitachi Device Manager is a core component of the new Storage Command Suite that provides a common framework for configuring and provisioning storage on Hitachi storage systems and virtualized external storage.

Hitachi Replication Manager is a software tool for management, consolidation and simplification of all replication operations.

Hitachi Storage Capacity Reporter reports both current and historical capacity views, forecasting and predictive analysis, and custom reporting capabilities.

Hitachi Tuning Manager is an automated, intelligent, and path-aware storage resource management tool that maps, monitors, analyzes, and reviews storage network resources from the application to the storage device. It helps storage administrators troubleshoot performance problems and preemptively identifies and alerts to potential problems.

Hitachi Dynamic Link Manager Advanced provides load-balancing and recovery capabilities protecting critical data against various types of failure,\ such as disk drive malfunctions and other hardware and logical failures.

Hitachi Tiered Storage Manager seamlessly migrates data (with a single click) across tiers and subsystems, at file and at object levels. HDS is the only vendor supporting such migration across discrete subsystems, other vendors use host-based software to move data across platforms

Integration with Server Virtualization

Storage virtualization amplifies the benefits of server virtualization.
In fact, users can quickly create hundreds of virtual volumes to support any virtual server platform and optimize the placement of virtual applications without wasting time, money or disk space. Tight integration of storage with server virtualization allows organizations to deploy robust, agile, and effective IT infrastructures

VMware

HDS has been cooperating with VMware since 2002, becoming a VMware Premiere Global Partner in late 2006. Their alliance as resulted in comprehensive interoperability via certifications and qualifications. The cooperation was initially focused on integration of HDS storage subsystems with VMware Site Recovery Manager and later expanded to create tight integration in initiatives such as vStorage and vCloud. This integration leverages storage subsystem functions such as thin provisioning and dynamic resource management in

VMware environments. The *Command Suite Tuning Manager* eases capacity planning, for example: The *Correlation Wizard* analyses and displays performance metrics of storage resources correlated to a VM resource and the *Historical Report* shows the trend of *datastore* utilization correlated to a VM's resources. VMware's *vStorage APIs for Array Integration (VAAI)* eases VMware deployments, improves security by hardware-assisted locking of VMFS cluster file system metadata, and allows effective use of VSP functions such as Hitachi Dynamic Provisioning and its replication techniques.

Microsoft Hyper-V

The VSP fully supports Microsoft's virtualization and business continuity technologies. Hitachi's Storage Cluster solution integrates the 2008 R2 Hyper-V™ Live Migration capabilities of Hyper-V with Hitachi replication technologies. This integration enables simple and reliable remote replication with automated failover and automatic resynchronization of disk resources for simplified failback in Hyper-V virtual environments. It fully supports Hitachi TrueCopy Synchronous and Hitachi Universal Replicator mentioned above.

Microsoft SharePoint

SharePoint is gaining momentum, particularly for dispersed workplace organizations. SharePoint Server enables quick and easy development of new applications by individuals, small work groups, and departments within these organizations. Over the past few years, organizations that use SharePoint Server and SQL Server as their main data repositories have seen an unprecedented increase in the volume of unstructured content.

Hitachi Data Discovery for Microsoft SharePoint offers significant improvements in scalability, performance, and backup of SharePoint environments. It enables customers to do a federated search, including, in addition to NAS and HCAP, SharePoint environments. Moreover, the advanced functionality provides the ability for archiving content, assigning retention policies, and protecting content from modifications according to compliance regulations. Another improvement is the synergetic search between discovery of SharePoint and Hitachi file services. For example, consider SharePoint files, which are stored in block format on an SQL Server database. When these files are migrated to an HCAP, a *stub* pointing to the HCAP will be automatically created in the database (transparent to users). This provides dramatic improvements in scalability, performance and backup/recovery time for Microsoft SharePoint.

Sustainability

Many vendors talk about "Green IT" or "Green Storage", but such things do not exist. Each IT product consumes energy and takes up floor space, so the real differentiation is in how effective it is in comparison to other products. VSP uses small form-factor HDDs and SSDs that consume less energy per capacity than large form-factor devices. The smaller 2½ inch HDDs use 40 percent of the power required for traditional 3½ inch drives. An industry-first ultra dense drive packaging holds 80 drives of 3½ inch size or 128 of the smaller 2½ inch

drives in a 13U high, 19 inch wide rack-mounted chassis. A fully configured VSP uses approximately six standard floor tiles to hold 2,048 disk drives–less space than required by any competitive subsystems.
An intelligent cooling design allows automated adjustments of the blower speed. The result is the best energy efficiency without compromising reliability, availability, or performance, which translates to over 40% reduction in power, cooling, and floor space costs.

Hitachi Data Systems Storage Portfolio

In addition to high-end storage subsystems such as the VSP or the predecessors USP V and USP VM, Hitachi Data Systems offers an impressive range of storage solutions:

- **Hitachi Adaptive Modular Storage AMS 2000 family** is one of the leading midrange subsystems on the market. HDS is the only storage company able to manage its high-end and the mid-range series using one common storage management software (Hitachi Storage Command Suite) using the same GUI/CLI.
- **Hitachi High-Performance NAS Platform (HNAS 3000 Series)** is based on the BlueArc Titan family and offers the most scalable, highest-performing clustered network storage system available in the market today. HNAS series scales up to eight nodes and more than four PB of storage while delivering up to 200,000 IOs per second and up to 1,600 MB/second per node for sequential workloads.
- **Hitachi Content Archive Platform (HCAP)** is a content-addressed storage (CAS) active-archive solution for long-term retention of fixed content. HCAP is available as a stand-alone appliance with integrated storage or as a disk-less appliance solution supporting all Hitachi block and file storage platforms. HCAP supports up to 40 PB in an 80 node archive system. It supports options for digital signatures to ensure authenticity, data compression for efficient use of network bandwidth, in-flight data encryption, compression, and data de-duplication.
- **Hitachi Data Discovery Suite (HDDS)** is a key enabler for Hitachi's file and content services vision. HDDS is a content-services software-solution designed to search across the boundaries of all Hitachi storage subsystems or NetApp NAS devices (running OnTap 7.3 and above). It can use content search to trigger the movement of data among the storage tiers as if it were a single entity.

Hitachi storage subsystems are trusted by the global business community: 44 of the top-50 *Fortune Global 500* companies, the top-15 global commercial banks, the top 16 global property and casualty insurers, the top-10 global telecommunications companies, and the top-10 global aerospace and defense companies use Hitachi storage.

Competition

There are many storage subsystems on the market but only few can be considered as real hi-end or tier-1. EMC announced its VMAX clustered subsystem based on clustered servers in April 2009, but it is often still offering its predecessor, the DMX, as the first choice. It is rare in the industry that a company does not start ramping the new product approximately half a year after launching. Why EMC is not pushing its newer product more than a year after its announcement remains a mystery. The VMAX scales up by adding disks (if there are empty slots in a module) or adding modules, the latter of which are significantly more expensive. EMC does not participate in SPC benchmarks.

IBM has introduced many enhancements to its high-end DS8x00 series since October 2007, but has since shifted the attention to the XIV, which is missing some high-end functionality.

Summary and Conclusions

Hitachi has again created a state-of-the art storage subsystem. The VSP is the industry leader in every aspect: reliability, functionality, performance, scalability and sustainability. For more than twenty years, Hitachi, having the fastest turn-around times in the industry, has brought to market new control unit designs approximately every four to five years, with a mid-life kicker two years after launching the original product to answer users' changing requirements.

The *Hitachi Storage Command Suite v7* storage management software delivers new, user friendly functionality for the complete portfolio of Hitachi Data Systems storage solutions. It integrates with other Hitachi Data systems products, delivering a comprehensive solution to ensure that data is kept on the most cost-effective media for daily, medium-term and long-term data retention requirements. The VSP's virtualization layer offers the only storage architecture that flexibly adapts in terms of performance and capacity, and extends to multi-vendor storage. The Universal Replicator simplifies deployment of disaster recovery for heterogeneous storage infrastructures. Built-in intelligence and automation constantly adapt to changes ensuring that the data is available to satisfy changing business conditions. The net result is superior data center efficiency, manageability and cost savings, which should put the VSP on any shortlist of organizations searching for excellence[9].

[9] "*Excellence is an art won by training and habituation. We do not act rightly because we have virtue or excellence, but we rather have those because we have acted rightly. We are what we repeatedly do. Excellence, then, is not an act but a habit*". Aristotle (384 BC-322 BC) Greek philosopher.