Josh Krischer & Associates

Enterprise Servers, Storage and Business Continuity

Hitachi Virtual Storage Platforms G1500/F1500- storage for zSystems

An Analyst Product Review

Josh Krischer, December 2014 –updated June 2016, December 2016

Josh Krischer & Associates
Enterprise Servers, Storage and Business Continuity

Contents

Josh Krischer is an expert IT advisor with more than 47 years of experience in high-end computing, storage, disaster recovery, and data center consolidation. Currently working as an independent analyst at Krischer & Associates, 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 mainframes, 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 October 24th, 2016, Hitachi Data Systems, a fully owned subsidiary of Hitachi Ltd., Japan, announced two new, enhanced models of its high-end Virtual Storage Platform; the hybrid VSP G1500, as the follow-on model for its technology leading VSPG1000 subsystems announced in 2014, and the all-flash VSP F1500. As did its predecessor, the new VSP supports all major operating systems, including IBM zSystems (or mainframe), and is designed for large data centers. The new G1500 model and delivers 20 percent more throughput and runs individual workloads up to 40% faster in comparison to the original VSP VSP G1000 at a comparable price. G1000 users can do an in-place upgrade from the HDS VSP G1000 to G1500 through a non-disruptive replacement of the virtual storage director controller cards.

The all-flash VSP F1500 contains new, high density Flash Module (FMD HD) technology with capacities of up to 14TB per module with up to 40PB of effective flash capacity.
In addition to hardware, HDS announced the next generation of its Storage Virtualization Operating System, SVOS 7.0, featuring enhanced QoS and new adaptive data reduction techniques. More details of the enhancements are down in the document.

Since the initial launch, the VSP G1000 and the G1500 were enhanced with compatible features such as the 16 Gb/sec FICON Express and several unique features such as the Mainframe Analytics Recorder (MAR) and Active Cache.

Users' Requirements and Challenges

The world of IT and communication is changing at a pace which we have never experienced before. Billions of people worldwide communicate on mobile devices and social networks, creating daily enormous amounts of data. IT organizations are changing as well with the adaption of cloud and mobile computing, and deployment of applications which benefit from "big data" analysis.

What does this mean for users? Organizations have used IT to manage and improve operational efficiencies for five decades. What is new is the volume of data, the variety of data (e.g., unstructured data), velocity and uncompromising demand for availability, business

continuity, and security. Despite to the latest buzz word of "Software Defined Storage", a robust, reliable and secure storage infrastructure is required today more than ever. New technologies such as thin provisioning and Solid State Disks (SSD) are gaining in popularity[1]. New criteria such as virtualization, cloud-readiness, energy consumption, and cost containment have emerged. The mainframe, declared dead many times, is "alive and kicking" and serves the largest and the most important organizations across all geographies.

What do users typically require from a storage subsystem?

- **Availability and Business Continuity:** These 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. In addition to complete fault-tolerance and data protection, requirements still not met by some products include service, capacity upgrades, testing, etc., which need to be performed non-disruptively, and completely transparent to users.

- **Scalability:** A storage subsystem should scale seamlessly in capacity, connectivity, and performance without impacting service levels. The capacity should support multi-tier storage media: SSDs, Flash, performance HDDs in different capacities, and high capacity "near-line" disks – in simple words, tiered storage "in-a-box".

- **Performance** This has two aspects: throughput measured in I/Os per second (IOPS), and response time measured in milliseconds. Performance should be at the SLA level regardless of capacity and workload. Erratic performance levels irritate users more than a slightly slower but constant response time.

- **Secure Multi-Tenancy:** Server virtualization creates situations where multiple tenants are using the same physical infrastructure, which requires new techniques to address security issues.

- **Virtualization:** The ability to virtualize internal and external storage, including heterogeneous storage and bring it into a common singly managed set of internal and external pooled storage.   This is the key and underpinning to software defined storage architectures.

- **Functionality and automation:** These 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. And there is a constant need to reconfigure and adapt to changing

---

[1] "The adoption of flash was remarkably quick during the year, with flash systems, both HFA and AFA, overtaking HDD-only systems and representing 51% of total value shipped in 2015 versus 40% in 2014"- said Silvia Cosso, senior analyst, European storage research, IDC.

---

requirements over time. The only way to cope with this 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, Flash 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.

Management Summary

The VSP G1000 is the industry leader in every aspect: reliability, functionality, performance and sustainability. In the Gartner's latest "Critical Capabilities for General-Purpose, High-End Storage Arrays" report, Gartner positioned the VSP G1000 as the leading product in this category. The new G1500 has all the characteristics to maintain this position.

Three vendors manufacture high-end enterprise storage systems that connect to zSystems: EMC with VMAX, Hitachi Data Systems with VSP G1500 (also OEMed by *Hewlett-Packard* Company as HPE G2 XP7 Storage), and IBM with the DS8000 series. HP's market share on zSystems is negligible. All three systems are multi-platform products; however this paper will concentrate on features support, synergy, and compatibility with zSystems.

Hitachi has a long tradition of delivering plug-compatible storage systems for IBM mainframes. The first exported devices in late 1970ies were compatible to IBM 3350 disk and OEM-ed by BASF, Olivetti, and National Advanced Systems, the predecessor of Hitachi Data Systems. Over more than 38 years Hitachi remained compatible with this critical platform, delivering storage systems, which delivered value-add features such as virtualization. Today Hitachi has a large installed base of storage systems connected to IBM S/390 and zSystems mainframes via IBM FICON channels as well as mainframe-based Linux environments via Fiber Channel networks.

IBM owns the zSystems mainframe architecture, which from time to time allows it to offer

HDS (since 1999) is the only vendor to guarantee a 100% data availability (100% availability, 100% of the time).

exclusive features for at least a limited period. EMC and Hitachi purchase the specifications for these features to remain compatible. IBM provides EMC and Hitachi with technical specifications but not the code, which the competing companies have to develop. Typically, Hitachi delivers the compatible features 6-24 months after IBM and ahead of EMC. In addition to providing support for IBM-compatible replication solutions such as Metro Mirror (PPRC), Hitachi delivers its own alternative high-performance replication solutions.

IBM successfully completed qualification testing of the Virtual Storage Platform with IBM GDPS using Hitachi support for PPRC, XRC and, in 2012, also qualified Hitachi Universal

Replicator (HUR) for use in combination with GDPS. Similar qualification testing for the Hitachi VSP G1000 was completed in 2015.This qualification testing is one more way organizations can be confident in the Hitachi commitment to providing compatibility with essential mainframe storage features. Customers gain full support of unique, innovative, Hitachi value-added replication solutions, including Hitachi TrueCopy, HUR, and Hitachi ShadowImage Replication. As mentioned above, Hitachi VSP is the only virtualized storage in the mainframe environment. It includes the ability to virtualize external multivendor storage.

Different vendors offer five or six 9s of availability, but HDS (since 1999) is the only vendor to guarantee 100% data availability (100% availability, 100% of the time).

Hitachi enterprise storage platforms are used by most Fortune 100 organizations today. These represent some of the most critical mainframe environments in the world.

VSP G1500 Design

System Design

The Virtual Storage Platform G1500 and the F1500 are enhanced Hitachi enterprise storage systems that includes the new **Hitachi Storage Virtualization Operating System 7**- **SVOS 7**. Hitachi SVOS is integrated storage system software that provides system element management and advanced storage system functions such as storage virtualization, thin provisioning, storage service level controls and performance instrumentation across multiple storage platforms.

The Virtual Storage Platforms (VSP) G1500 and F1500 unlike EMC's VMAX 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. Each VSP G and F series storage controller is in effect a clustered system with a flexible set of paired components. The 'tailored' design is one of the factors in achieving its very high performance, large scalability, and unmatched reliability.

With SVOS the VSP G1500 is a highly scalable storage system, with the following characteristics within the system (without any external in-band or out-of-band devices required to manage storage):

- A VSP G1500 solution scales **Up** for increased performance and connectivity. This applies to the hardware of a single controller system.
- A VSP G1500 solution scales **Out** for doubled internal capacity, connectivity, and performance. This applies to upgrading from a single controller to a dual controller system.
- A VSP G1500 solution scales **Deep** with virtualization of external storage systems thereby incorporating additional tier-2 and tier-3 performance and capacity

_____

Scale Up – Single Controller

Within the basic single controller system, the number of controller boards, drive chassis, and drives is highly scalable. A system can initially comprise a single rack and the minimum set of controller boards in one controller chassis along with one (optional) drive chassis. Drive chassis options include types that hold up to 96 LFF, 192 SFF, or 96 Accelerated Flash Module (FMD) drives in 8 drive trays per drive chassis. A diskless system can be configured with the minimum set of controller boards and no internal storage, only using virtualized storage attached to Front End Director (FED) Open Fiber (FC) ports.

This system can then be expanded with more controller boards (up to 20 in all per controller), 3 racks, and more drive chassis (up to 6 in all). A single controller system can support a maximum of 6 drive chassis using any combination of the three chassis types. In a pure Small Factor (SFF) drive configuration, the G1500 can support up to 1152 drives (including 192 SSD drives) in six drive chassis. If using only Large Factor (LFF) drives, the limit is 576 drives in six LFF drive chassis. There can be total of 576 Flash Module drives in six FMD drive chassis (from the total of 6 drive chassis). See the HDD and Flash options in the Capacities and Scalability section, below.

Scale Out – Dual Controllers

A VSP G1500's resources can be greatly extended by adding a second controller system with up to six racks in all. A dual controller system can start with a minimum set of controller boards and drive chassis and later scale out to a full configuration of up to 40 controller boards and 12 drive chassis in 6 racks.

A minimal (very high IOPS) dual controller storage system can be configured as a single rack with two controllers and two FMD drive chassis. Perhaps a more typical smaller configuration would include two controller chassis, two racks, and four drive chassis. This can be expanded by four more racks to add up to 8 more drive chassis (12 in all).

Each expansion rack supports two drive chassis, with each containing up to 96 LFF, 192 SFF, or 48 FMD drives. The two controller racks can be separated by up to 100 meters using longer optical connection cables. The two controller chassis are the same and are populated with the same types and numbers of controller boards. The two controller chassis are cross-attached at the cache path control (CPC) adapter board level, thus creating the global HiStar-E system.

Note that any front-end port on a dual controller storage system can access any internal (or mapped external) LUN as there is no division within the system. The only hardware division is that each controller's Back End Director (BED) boards have their own set of directly attached drive chassis. Any Virtual Storage Director (VSD) processor can manage a back-end I/O operation on any BED board in the system. Any Front End Director (FED) or BED

board can access all installed CPC boards and communicate with every VSD board. In a system configured solely with SFF drives, a dual controller chassis G1000 can contain up to 2304 SFF drives (including up to 384 SSD drives) in 12 SFF drive chassis; or if using only LFF drives, the limit is 1152 drives in 12 LFF drive chassis. It can also have up to 576 FMD drives in 6 FMD chassis (three per controller). See table 1 for capacity options in the Capacities and Scalability section, below.

Scale Deep – Virtualization

The Storage Virtualization Operating System for the VSP G series provides the Hitachi Virtualization functionality (the **Universal Volume Manager** feature) as found on the previous VSP and Universal Storage Platform V generations. Existing customer storage systems (usually repurposed upon being replaced by newer storage systems) may be directly attached to some of a VSP G1500's FED FC ports (now operated in *External Initiator Mode*). The LUNs that are supplied by these external storage systems are accessed via hosts that are attached to other FED FC or FICON ports (operated in Target Mode) on the same G1500 system. As far as any host is concerned, all virtualized LUNs passed through the VSP G1500 external ports to the host target ports simply appear to be *internal* LUNs or Volumes from the VSP G1500. The VSP G1500's FC FED ports that attach to the external storage system's front-end ports are operated in "external" or "SCSI initiator" mode (attached to storage), rather than the usual "SCSI target" mode (attached to hosts).

One advantage of virtualization is greatly simplified management of LUNs and hosts. Another advantage is the ability to dynamically and transparently relocate LUNs or volumes from Tier-1 RAID Groups down to Tier-2 or Tier-3 using other RAID levels, drive types, or even virtualized storage. These migrations are able to proceed while the original source remains online to the hosts, and the VSP G1500 will seamlessly switch over the mapping from the original LUN to the new LUN when completed. No changes on the host mount point are required.

Processors

Architecturally each VSP G1x00 Controller has tightly-tied multiple system Virtual Storage Directors (VSDs) sharing a global cache linked by a non-blocking, crossbar architecture connecting the front- and back-end processors. The major hardware modification of the VSP G/F1500 are the more powerful processors of the VSDs. The Hi-Star crossbar switch-based architecture is 8th generation Hitachi innovation and has 896 GB/sec aggregate internal bandwidth. The control unit is based on a three-level processing structure (see fig. 1 and 2):

- ❑ **VSD** processor boards (Intel Xeon E5 2.3Ghz Eight-core processor)
- ❑ **CPC** Cache Path Control Adapter board provides the internally used HiStar-E ports as well as the Cache Path Controllers. All controller boards now have four PCIe3 paths to the central HiStar-E grid (with one path per CPC board within the same controller).

- ❑ **FED** front-end director boards (16x 8Gb/sec or 16Gb/sec FC or FICON ports each)

- ❑ **BED** back-end director drive-controller boards (16x 6Gb/sec SAS ports each, utilizing a proprietary high speed Data Acceleration ASIC for parity and formatting operations),

- ❑ **BKM** battery back-up units for each **CPEX.**

All of the controller boards within a system are connected via the **Hitachi HiStar-E** network paths.

VSP G1500 and the F1500 separates the management and control processing from the front-end and back-end-processing. It is internally connected by a Hi-Star crossbar switch. This structure separates the data and control paths and delivers improved performance for the unique processing requirements of different data types (see Figures 1 and 2).

The Virtual Storage Director (VSD) is a powerful data movement engine and the heart of the Storage Subsystem. The 2-16 VSDs are installed in pairs to ensure redundancy. Each board contains an 8-core Xeon processor. These boards manage the VSP G1500's operation outside of the data path. The difference between the two and four VSDs configuration is in performance and 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.
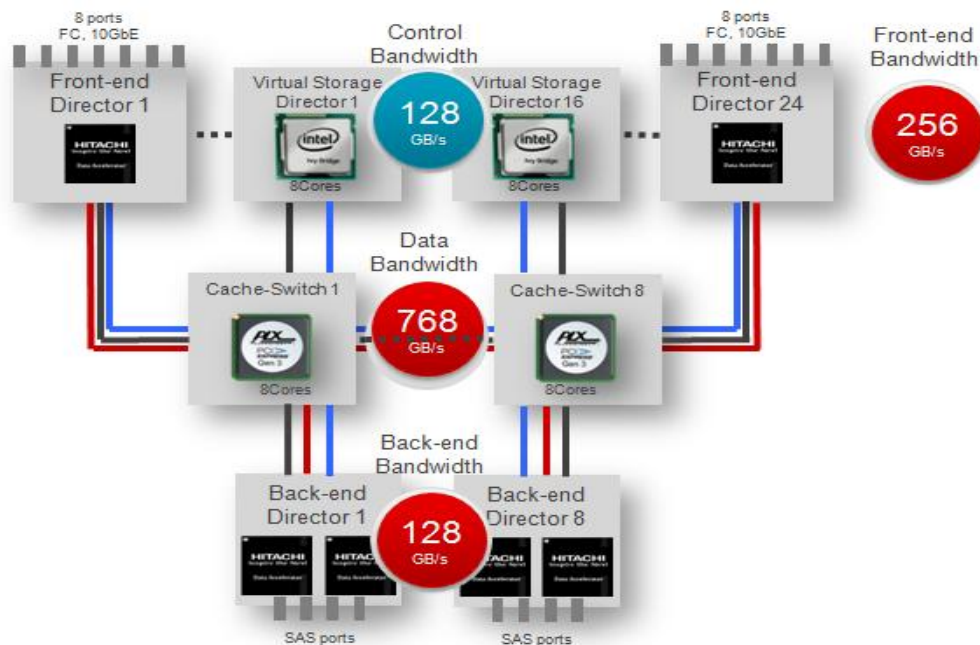

Figure 1: VSP G1500 Structure and Bandwidths (source: Hitachi Data Systems)

The operational <u>control memory</u> of VSP G1500 is physically located on the virtual storage director pairs (VSDs) which enables the main processors to access the control data memory at internal memory speeds eliminating interface latency delays. It also means that there is no internal bandwidth usage for access to control data, which results in faster processing. The maximum control bandwidth is 128 GB/sec. There can be a total of 16GB of control data memory on a single VSD (8 GB per processor) and a total of 128GB of control data memory in a single controller chassis. This control data is not mirrored so the memory capacity available for operations in a single chassis is128GB. However, this control data is backed up in cache on the first data cache adapter pair, which is mirrored, creating three copies of the control data, ensuring redundancy. The separation of the control memory and the data memory (cache) eliminates the possible contention among multiple accesses to the cache for fetching control and users' data and increasing performance.

Hitachi Accelerated Flash (HAF)

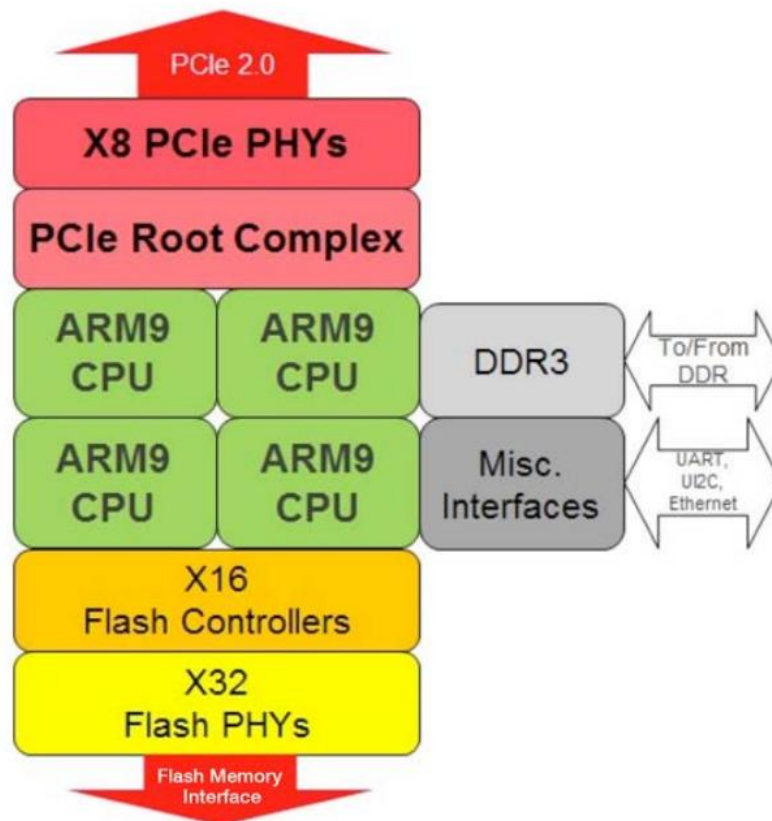**Figure 2: Hitachi custom Flash Module controller**

Additional processing power deliver the processors embedded in flash module. In addition to the industry standard enterprise SSD's, the VSP G1000 supports high-speed Flash Modules Drives (FMD) (see figure 2). Hitachi Accelerated Flash (HAF) storage is an enterprise-class, solid-state storage option that combines the unique Hitachi embedded flash controller and

_____

enterprise MLC flash memory to deliver superior performance at an acceptable price. The HAF controller is a custom-designed module featuring a quad-core ARM processor and ASICs that delivers processing power needed for housekeeping management. It manages the Error Correction Code (ECC) to ensure data integrity, wear leveling to increase endurance, in-line powerful compression and adaptive data refresh to secure data availability. The embedded processors off-load these tasks from the main processors. As of August 2015 Hitachi had registered more than 350 patents covering the FMDs and flash. As opposed to the SSDs, which are connected on the SAS bus, the FMDs are sitting directly on the PCIe interface with much lower latency.

The results are:
- Real-time inline compression (not supported for zSystems, other platforms only!):
  - Faster performance,
  - Better capacity usage, and,
  - Larger endurance.
- Advanced ECC:
  - Delivers up to 5x the write performance of SSDs,
  - Provides 4x more effective capacity for a lower TCO, and,
  - Maintain <1 ms response time.
- FMD Data Eradication Service (All user data, including data in bad blocks is physically erased),

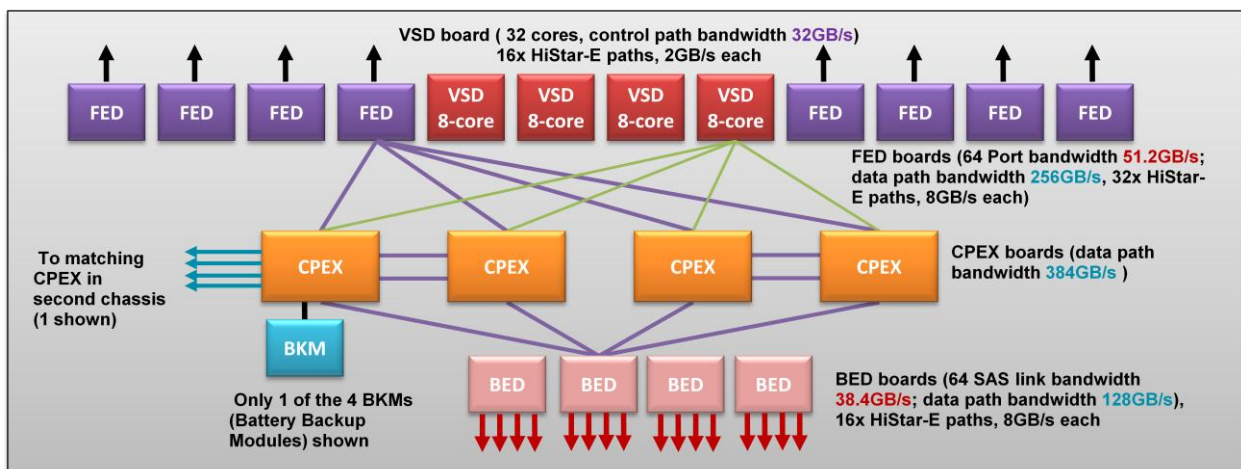The compression ratio can be as high as 5:1.

Figure 3:VSP G1500 Single Controller Components (source: Hitachi Data Systems)

Data Cache

A Cache Path Control adapter is included with each controller. The Cache Path Control adapter is the combination cache memory and 4 HiStar-E switch boards, with up to 256GB of cache each, with 16GB being the minimum and 1024GB being the maximum. Each CPC

---

board provides 16 HiStar-E paths to the FED, BED, and VSD boards, 6 paths to the other CPC if installed.

An additional cache path control adapter may be added to each controller, expanding the total data cache capacity supported to 2048TB. The data cache memory is used to cache user data blocks that are fetched from disk drives during a read and written to cache data as part of a data write operation. All write data is simultaneously mirrored to a paired CPC cache 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. Each cache board also has embedded SSDs 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 failure or shutdown. In the case of any failure, a recovery mechanism is used to back up cached user data that has not been de-staged to disk and the control memory by temporary storing this data on the embedded 128GB or 256GB SSD devices. The cache path control adapter has an onboard battery that allows the cache adapter microprocessor to de-stage the data to SSDs on system shutdown or loss of power.

VSP G1000 uses several advanced cache algorithms and software to deliver the required performance and ensure optimal cache capacity usage. For example, if data is being accessed sequentially, the VSD processors allocate additional cache slots and will pre-fetch the data in parallel so that subsequent read requests will be satisfied immediately from the cache, thus speeding up I/O.

Hi-Star crossbar switch

If the VSDs are analogous to the heart of the subsystem, then the Cache Path Control cache-switches (CPEX) represent its blood vessels. 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 with 768GB/sec bandwidth.

Front-end Directors (FED)

The 2-24 Front-End Directors (FEDs) board controls the connection between the array and servers attached to the host ports, as well as connections to external storage (see virtualization) or remote copy products (TrueCopy and Hitachi Universal Replicator). four types of FED boards are available (always installed in pairs of the same type), providing up to 176 FICON 8 Gb/sec or 16 Gb/sec[2] channels for IBM zSystems, up to 192 Fiber Channel

---

[2] In August 2015 Hitachi added **16-port 16Gb/sec IBM FICON** card hardware options: These short- or long-wave configurations enable high-speed connections to latest IBM z Systems mainframes.

8Gb/sec or 16Gb/sec host ports for other platforms, 192 Fiber Channel over Ethernet (FCoE) and 88 iSCSI10GBase T host ports.

The various types of interface options can be 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 G1000 internal volumes.

The FEDs use their own processing power and therefore don't "steal" cycles from the VSDs. The total maximum bandwidth of the FEDs is 256 GB/s.

Back-end Directors (BEDs)

The 2-8 BEDs boards act as the interface to the HDDs, SSDs or Flash Modules installed in the drive control chassis. Two or four BED boards can be installed per chassis. Each BED board has sixteen 6Gb/sec SAS links. A standard-performance single-chassis configuration uses one BED pair with 32 back-end 6Gb/sec SAS links. The high-performance single-chassis uses an additional pair for a total of 64 6Gb/sec SAS links.

In addition to data transfers, the BED boards generate RAID-5 and RAID-6 parity and perform drive rebuilds by an embedded Data Acceleration ASIC and data encryption with a dedicated Encrypting Processor. The VSP G1500 supports RAID-10 (mirrored data striping), RAID-5 (data striping with distributed parity) and RAID-6 (data striping with two distributed parity disks. Any disk can be configured for any RAID technique, though it is recommended to configure SSDs as RAID-5 and high-capacity near-line disks as RAID-6.

Hitachi VSP F1500 – All Flash Appliance

Three decades after delivering semi-conductor storage for mainframes, Hitachi delivers all-flash appliances for all platforms[3]. Powered by the SVOS 7 and based on the VSP G1500 hardware, the VSP F 1500 delivers unmatched throughput of 4.8 million IOPS with sub-millisecond response time. Due to the huge benefits of the FMDs over the SSDs, both the VSP G1500 and the F1500 support high density Flash Module (FMD HD) only. The normal SSDs are still supported on the other G series models. The F1500 supports 7 or 14 TB Flash modules, from 50TB up to 8 PB raw capacity.
By using technology with capacities of up to 14TB per device, the F1500 scales up to 40PB of effective flash capacity.

---

[3] October 1986: **National Advanced Systems** Corp. has announced a model of its 7900Semiconductor Disk Storage Subsystem (SCD). 512 MB in a cabinet.

One of the biggest advantages of using the proven hardware and the SVOS is that the VSP F series supports all the functions of the VSP series, which is rare in most of all-flash subsystems.

Hitachi Storage Virtualization Operating System

The Storage Virtualization Operating System 7(SVOS 7) of the VSP G1000 is the seventh generation of HDS' enterprise storage array operating software. The SVOS software provides software-defined storage management for both Open and Mainframe environments. Hitachi SVOS abstracts information from storage systems, virtualizes and pools available storage resources and allows automating key storage management functions. This unified environment maximizes the utilization and capabilities of the storage resources and significantly reduces operations costs and risk. Standards compatible for easy integration into IT environments, SVOS advanced storage virtualization and management capabilities enable the upmost agility and control required to build infrastructures that are continuously available, automated and agile. SVOS supports both Mainframe and Open storage environments and includes functionalities for mainframes and other platforms such as:

- ❑ **Hitachi Storage Navigator,** which obtains system configuration and status information and sends user requested commands to the VSP G1000.
- ❑ **Universal Volume Manager,** which enables virtualization of external heterogeneous storage.
- ❑ **Dynamic Provisioning,** which provides thin provisioning for simplified provisioning operations, automatic performance optimization and storage space savings.
- ❑ **Device Manager,** which provides single-point management for all Hitachi physical and virtualized storage and acts as the interface for integration with other management systems.
- ❑ **Resource Partition Manager,** which supports secure administrative partitions for multi-tenancy requirements.
- ❑ **Cache Partitioning**, which supports up to 32 cache partitions to ensure multi-tenancy.
- ❑ **Performance Monitor** feature assists with performance configuration planning, workload balancing, and analyzing and optimizing storage system performance.
- ❑ Along with other useful storage system utilities.

SVOS includes Storage Navigator for low level storage system management and is the foundation for many VSP G1000 mainframe compatibility features. SVOS provides and manages customer-defined non-disruptive data migration workflows between storage tiers and between virtualized physical systems among and across data centers for true workload mobility.

The SVOS 7.0, announced in October 2016, is engineered with an improved flash-aware IO stack that includes enhanced QoS, new adaptive data reduction and support for all-flash cloud tiering. These new features are for other Operating Systems than zOS. See appendix 2 for detailed description and benefits.

_____

Capacities and Scalability

VSP G1500 is available with a choice of standard high-performance disk drives, high-capacity near-line drives, Flash SSDs and high performance Hitachi Accelerated Flash storage. A fully populated six-rack design holds up to 2,304 drives. The different capacity options of the VSP G1500 can be seen in table1 below.

Max. Internal Raw Capacity	8106TB			
Flash Storage Options	1.6TB FMD	3.2TB FMD	7TB FMD	14TB FMD
Hard Disk Drive (HDD) SFF raw capacity 2.5" SAS	600GB 15K RPM NA	600GB 10K RPM NA	1200GB 10K RPM NA	1800GB 10K RPM NA
LFF HDD 3.5"SAS 7.2K RPM			4TB	6TB
Minimum to Max. Hard Drives	0-2,304 2.5" and/or 0-1 152 3.5" including spares			
Max. SSDs	384			
Max. Flash Modules	576			

Table 1: HDD, SSD and Flash Options for G1500

Functionality

The VSP G1500 and the F1500 support wide range of functionality for all computing platforms including zSystems.

zSystems functionality and compatibility

Hitachi and IBM share APIs to ensure compatibility between Hitachi storage and IBM mainframes. For more than 38 years Hitachi, Ltd. has provided mainframe technology to business and industry, supporting all mainframe operating systems, including MVS, z/OS, z/VM, z/VSE, z/TPF and Linux for z. During this time, license agreements have provided Hitachi with access to IBM APIs and interface specifications for mainframe storage platforms as well as IBM access to selected Hitachi APIs. These APIs enabled Hitachi storage systems to support remote mirroring such as PPRC, XRC and other functionalities such Parallel Access Volumes (PAV), HyperPAV, Multiple Allegiance (MA), High Performance FICON (zHPF), TPC-R, and Geographic Dispersed Parallel Sysplex (GDPS), and HyperSwap (a full list of supported features is in appendix 1). The license agreements have also given IBM access to Hitachi APIs for NanoCopy disaster recovery manager technology from Hitachi Data Systems and Hitachi ShadowImage Replication. This licensing agreement also enables Hitachi and Hitachi Data Systems to develop and deliver storage management software for mainframe environments. Hitachi participates in the IBM Early Ship program for z/OS, which means that Hitachi has early access to z/OS versions for testing. This access enables Hitachi to ensure compatibility with new versions of z/OS when they become generally available.

Hitachi storage systems and features are tested and qualified with IBM mainframes, operating systems, and critical features. This testing and qualification includes features such as support for GDPS/HyperSwap, FlashCopy and zHPF. IBM successfully completed qualification testing of Virtual Storage Platform with IBM GDPS using Hitachi support for PPRC, XRC and, in 2012, qualified Hitachi Universal Replicator (HUR) to be used in combination with GDPS. Similar qualification testing for the Hitachi VSP G1000 was completed in 2015.The Hitachi VSP G1x00 family, Virtual Storage Platform, Hitachi Universal Storage Platform V (USP V) and Hitachi Universal Storage Platform VM (USP VM), are qualified and supported on the IBM z196, Z9, z10, z114, zEC12, zBC12 and z13s processors. They are also supported under the latest z/OS version. Hitachi has implemented technologies that are compatible with IBM's Metro Mirror (PPRC), FlashCopy V2, and FlashCopy Space Efficient on Hitachi storage systems. Hitachi is the only storage vendor to go through this qualification. Hitachi also successfully completed the qualification tests for GDPS/XRC and a 3-data-center (3DC) multitarget replication configuration using VSP GDPS/HyperSwap support and HUR with delta resync along with Hitachi Business Continuity Manager.

The license agreements ensure that as IBM announces and delivers new features and functions for mainframe environments, Hitachi Data Systems will be able to support these new features. In addition, Hitachi has its own testing facilities to ensure that its products and features are compatible with the IBM solutions.

Hitachi Data Systems and IBM have a long-term agreement to provide assistance to each other to quickly resolve any joint customer issues related to Hitachi storage and IBM mainframes and mainframe operating systems. The focus is on determining the source of the problem, with the root cause and solution being provided by the vendor whose product or products are identified as the source of the problem.

General Functionality

In addition to the compatible features, the VSP G1500 and the F1500 systems support a wide range of storage management features which ease storage mamagement and ensure better storage economics. The most important features are:

Hitachi Dynamic Provisioning (HDP) for the mainframe and other platforms complements existing mainframe storage provisioning processes while improving performance and simplifying performance and capacity optimization. In addition to thin provisioning, HDP improves performance by automatically spreading the content of application data sets across a storage pool, containing a large number of physical disks, distributing I/O activity across available HDDs and SSDs. Striping the data among a large number of physical devices practically eliminates hot spots and allows parallel access to disks, SSD's and Flash Modules, which results in almost uniform performance.

Thin provisioning enables the option of allocating 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 and Write Same* functions return 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. In addition to improving performance Hitachi Dynamic Provisioning increases installed capacity utilization, defers upgrades, and saves storage management tasks, thus reducing CapEx and OpEx.

Complimenting Hitachi Dynamic Provisioning is **Hitachi Dynamic Tiering** (HDT). On average, 20 percent of data accounts for 80 percent of activity, which justifies tiered storage deployments. Users require features that simplify the building of tiered storage infrastructures. Hitachi Dynamic Tiering is a dynamically automated data placement feature that constantly classifies and migrates data to the most suitable tier based on actual usage. This is especially important as users move from HDD to SSD or Flash Modules for tier 1 storage. SSDs and Flash are an order-of-magnitude faster on average in reads and writes but are also high-cost, making it vital to migrate data to less-expensive storage as its use decreases.

The optimal usage scenario for Hitachi Accelerated Flash Modules or SSDs is for 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 SSD may be not economical, and partial data placement cannot be manually controlled by operators. The answer is automated data placement.

Hitachi's VSP G1500 is a tiered storage in-a-box: it supports Accelerated Flash Modules, SSDs, several performance SAS HDDs, capacity near-line disks, and various RAID techniques. HDT is an automated methodology that attempts to move the most active back-end data to drives with higher performance, while migrating less active data to slower, less expensive drives. Mainframe storage data chunklets of 38MB are dynamically placed in the appropriate tier based on host access and usage patterns, media type, and speed, RAID level and sustained I/O level requirements. This feature helps maintain peak performance under dynamic conditions without manual intervention at optimal CapEx and OpEx savings. The

> *"Hitachi's VSP G1500 is a tiered storage in-a-box: it supports Accelerated Flash Modules, SSDs, several performance SAS HDDs, capacity near-line disks, and various RAID techniques.."*

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). The system automatically optimizes the use of storage by keeping the higher tiers fully utilized.

Hitachi Virtual Storage Platforms G1500/F1500- storage for zSystems

Hitachi Dynamic Tiering for Mainframe complements existing mainframe storage provisioning processes, such as DFSMS. At the same time it offers the full benefits of Hitachi Dynamic Provisioning to improve performance and simplify performance and capacity optimization. Existing SMS storage groups and ACS routines can be aligned to differently tiered storage pools.

In addition, **Hitachi Tiered Storage Manager (HTSM)** for Mainframe is a z/OS software management product for Hitachi Dynamic Tiering for Mainframe that enables a user to control service levels based on performance and/or time to facilitate meeting mainframe service level agreements (SLAs). HTSM offers 32 levels of fine grain policy control to control storage service levels in an auto-tiered environment. A key differentiator with the competition is that HTSM for Mainframe runs on the mainframe, not in the storage control unit. HTSM for Mainframe enables management of HDT for mainframe via REXX Functions within REXX Scripts and/or an ISPF Dialog at the volume level or through DFSMS tools and provides a familiar interface for mainframe storage administrators. It allows setting tiering policies by SMS Storage Group or individual devices and provides reporting capabilities for ongoing monitoring of the dynamic tiering environment.

Existing Interesting or Unique Features

A unique feature is the **Hitachi Universal Volume Manager** introduced with the first version of the USP in 2004. The Hitachi *Universal Volume Manager* software configures, manages, and accesses external volumes as if they were VSP G1500 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, 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.

Another unique feature of VSP and USP systems is storage system partitioning by **Hitachi Resource Partition Manager** (HRPM). Enterprise storage is often connected to multiple hosts and shared by multiple divisions in a company or by multiple companies. This means that a storage system may have multiple storage administrators. As a result, isolation and compartmentalized management of the storage system are required. HRPM software addresses this need by enabling administrative partitioning of storage system resources. Creating isolating environments on a physical storage system solves security, audit, and compliance issues of shared storage assets used by different groups. Public cloud-shared storage needs to ensure stringent security for each user's data, otherwise the users require separate subsystems, which increases service costs and makes the concept less attractive. Therefore this feature is an ideal solution for External Service Providers and cloud storage providers.

_____

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. Since its introduction in early 2014, Hitachi continued the development of the VSP G1x00 mainframe storage capabilities with zSystem compatible features and unique features. The two new value added features are:

- ❑ **Active Flash:** This new Hitachi Dynamic Tiering feature accelerates performance with promotions to flash at real-time speeds, maximizing the efficiency of installed flash.
- ❑ **Mainframe Analytics Recorder** (MAR)**:** This new Storage Virtualization Operating System feature provides detailed, time-coherent VSP G1500 performance information in standard SMF record format for use by mainframe analytics tools. Maximizing mainframe systems use requires detailed performance information therefore delivering this data in native SMF format ease the performance reporting of the "whole system" mainframe analysis.

Availability and Resiliency

Hardware reliability

Availability, Business Continuity and Disaster Recovery 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 and 24/7 non-disruptive operations. The requirements from any data center are providing non-disruptive data access, movements, provisioning, and protection of the organization's data.

The VSP G1000 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 G1500 maintains HDS's position as the industry's availability leader for many storage subsystem generations already. The base configuration is non-disruptively and flexibly upgraded by adding components such as Virtual Storage Director (VSD) pairs, data cache, front-end directors, back-end directors, Cache Path Control cache-switches, and disk drives to the existing chassis. Upgrades also include adding a second controller chassis and adding cabinets with disk modules.

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. The virtualization functionality of the VSP G1000 can significantly simplify data migrations.

The HDDs are carefully screened before installation. In addition to the three RAID techniques mentioned above, Hitachi uses extensive error correction techniques for the HDDs and the SSDs. For example: Eight bytes of error correction code (ECC) are stored with each

_____

512 byte HDD data sector and are used by the 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 of a suspected drive, avoiding having to perform an entire RAID rebuild[4]. Write operations are simultaneously duplexed into VSD G1000 cluster 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.

The Hitachi multi-core embedded flash controller introduces multiple unique capabilities designed for enterprise workloads and extended endurance. Zero block write avoidance compresses, in real time, any data stream of zeroes or ones by up to 94% capability, thus speeding writes, improving performance and saving writes to the flash. It also reduces array setup time by performing flash formatting 5x faster and extends endurance by avoiding unnecessary erases. Error correction corrects up to 42 bits of errors per 1.4KB of data, ensuring data integrity better than HDDs. To manage the uneven wear characteristic, the flash controller performs a memory block analysis function. This continuous scanning and diagnostic assessment of each block effectively manages its optimal time for reclamation and ensures the expectations for enterprise data. The program decreases background task overhead and improves sustained performance while extending flash cell longevity.

The VSP G1500 supports a variety of Hitachi proprietary or IBM zSystems-compatible local and remote mirroring techniques. The management of the Hitachi replication solutions for high-end storage has remained consistent across the last three generations of devices, which eliminates the need to build new scripts and procedures and time and money investment in training.

Different vendors ensure offer five or six 9s of availability up time but HDS (since 1999) is the only vendor to offer a 100% data availability (100% availability, 100% of the time).

Mirroring

The VSP supports three techniques of local (in system) mirroring:

❑ **Hitachi Compatible Mirroring for *IBM FlashCopy*** provides support compatible with IBM FlashCopy V02 and enables application-aware, near-instant copies and restores capability.

❑ **Hitachi Compatible Software for *IBM FlashCopy SE (FCSE)*** provides an in-system solution for data protection that is IBM mainframe compatible. This solution combines

---

[4] To copy a 4 TB HDD may take 4-6 hours, to rebuild without load ca. 48 hours. A rebuild with 80% CU load can more than a two weeks.

Hitachi Virtual Storage Platforms G1500/F1500- storage for zSystems

snapshots, virtual data management and dynamic provisioning to create logical point-in-time (PiT) copies.

❑ **Hitachi proprietary *ShadowImage Replication* software** provides non-disruptive, host-independent data replication to create copies of any customer-accessible data within a single Hitachi storage system. These copies can provide immediate access and sharing of information and optimize tape backup operations. ShadowImage copies provide nearly instant recovery from logical data corruption

All these techniques may use consistency groups to ensure consistency of cross volume data when replicating or making point-in-time copies. **Multiple tiers** provide the ability with ShadowImage to make 1st- and 2nd-generation copies of volumes that are being copied by ShadowImage. The copies can be on different storage tiers, either internal or externally virtualized storage. This capability supports disk-based backup as well as additional copies for testing or backing up to tape.

Remote Replication Techniques

The remote mirroring techniques ensure continous or near-continous operation in cases of major malfunction, power loss, fire, or other disaster. Hitachi replication solutions for mainframe environments support all of the latest generations of Hitachi enterprise-class storage systems, including the VSP G1x00, VSP, USP V, and USP VM. It allows use of previous storage generations on the recovery site to control costs.

Data consistency means that, from an application's perspective, the secondary disks contain all updates to a specific point in time, and that the write update sequence is maintained. Recovery time can be significantly affected by the consistency of the remote data. Databases, catalogs, volume table of contents (VTOC) and indices are only a few examples of files that are vulnerable to "out-of-sequence writes" or "lost writes" because of rolling disasters. Database restart logic performs resolution of the "in flight and in doubt" transactions, typically in seconds or minutes. If data consistency is not maintained, then a database recovery process must be performed. This involves restoring image copies and logs, if they exist, to disk and executing forward recovery utilities to apply updates to the image copies. Database recovery is a time-consuming task and may take several days. Hitachi also provides support for extended consistency groups, providing consistency of data during replication across multiple storage systems to the storage systems at the remote location.

GDPS and consistency groups ensure data consistency and fast recovery. As mentioned above, Hitachi data replication solutions for mainframes fully support GDPS, GDPS/HyperSwap or Basic HyperSwap. 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.

Another option popular with major financial institutions is the ability to deploy three-data-center solutions with the option to synchronize the data at the alternate sites using delta

resync if the primary site should be unavailable. This can all be managed using Hitachi Business Continuity Manager software. Hitachi Business Continuity Manager enables single-point management of Hitachi remote mirroring products from mainframe environments. It provides the ability to manage complex distributed two and three data center environments easily from a single point of control. Like HTSM for Mainframe, it supports both TSO/ISPF and standard REXX scripting to customize and automate:

❑ Hitachi *TrueCopy synchronous* remote replication software, which provides a continuous, nondisruptive, host-independent, zero-data-loss solution. This technique is deployed for distances within the same metropolitan area. For mainframe environments, TrueCopy synchronous is compatible with IBM Metro Mirror (a.k.a. PPRC) and is eligible to participate in GDPS and HyperSwap deployments. This technique supports consistency across multiple storage systems (up to 12x12) and provides 3 data center configurations.

❑ Hitachi *Universal Replicator* **(HUR)**, which enables hardware-controlled asynchronous remote copying between heterogeneous storage platforms. HUR offloads the majority of processing for asynchronous replication to the remote storage system using unique Hitachi pull technology which significantly reduces the overhead on the primary storage system. HUR minimize the Recovery Point Objective (RPO) by a proprietary technique using an internal storage system journal-based method. The TrueCopy HUR supports consistency groups across volumes and across up to 12 controllers on Hitachi high-end storage systems. It participated in GDPS and with the TrueCopy in three-site configurations.

3-Data-Center (3DC) Replication

The 3DC replication solution uses both Hitachi TrueCopy synchronous replication and Hitachi Universal Replicator asynchronous replication. It provides a synchronous copy to a separate Hitachi storage system within metro distances using TrueCopy and an asynchronous copy to a second Hitachi storage system that can be located at any distance from the primary site. There are two implementations of 3DC replication: cascaded replication and multitarget replication:

❑ **Cascaded replication** provides synchronous replication from a production Hitachi high-end-class storage system (such as VSP G1500, G1000, VSP or USP V) to a secondary storage system using TrueCopy synchronous replication software. It then asynchronously replicates the data from the secondary storage system to an out-of-region Hitachi enterprise storage system using HUR.

❑ **Multitarget (Concurrent) Replication** synchronously replicates data from the production site using TrueCopy synchronous software to a local site (within metro distance) and concurrently replicates to an out-of-region remote site using HUR. With multitarget replication, protection can be resumed with no data loss between the two backup sites (hot standby site and remote site) using delta resync if the primary data center is lost. Delta resync dramatically reduces the times required for recovery from a

disaster. If the production site experiences a failure, the hot standby (local) site can take over and use the remote site for disaster recovery; it only has to copy differential data. A full copy of the data is normally not required. As a result, disaster recovery protection for the hot standby site resumes within minutes instead of hours.

Performance

On February, 16 2015 Hitachi Data systems published SPC-1 IOPS[5] results for the VSP G1000 with Hitachi Accelerated Flash. The subsystem managed to execute 2,004,941.89 IOPS in sub-millisecond response time, which positioned the VSP G1000 on the second place behind Huawei which doesn't support zSystems. The Response time versus Throughput can be seen as figure 4. EMC doesn't participate in SPC or any other public benchmark, IBM didn't publish yet the SPC-1 results for the latest DS888x. It is too early to produce SPC-1 figures for the VSP G1500, but early measurements shows ca. 20 percent throughput improvements over the VSP G1000 and lower response time. The VSP F1500 delivers unmatched throughput of 4.8 million IOPS with sub-millisecond response time and bandwidth of 48 GB/sec which most probably position it as performance leader among all storage subsystems.
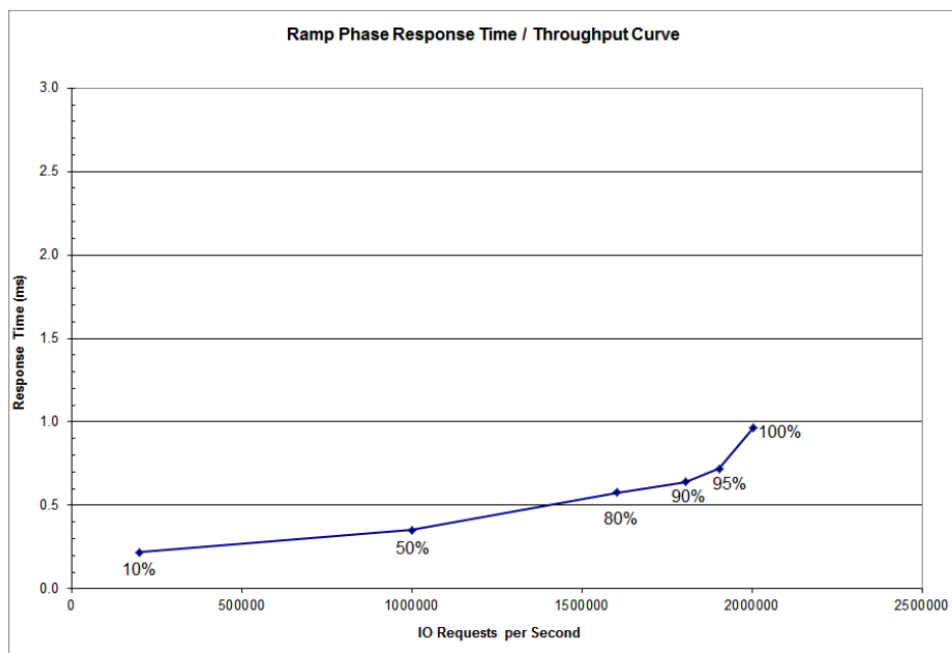

Figure 4: SPC-1 Response Time/Throughput Curve. Source SPC for G1000.

---

[5] **SPC Benchmark 1** *(SPC-1)*
SPC-1 consists of a single workload designed to demonstrate the performance of a storage subsystem while performing the typical functions of business critical applications. Those applications are characterized by predominately random I/O operations and require both queries as well as update operations. Examples of those types of applications include OLTP, database operations, and mail server implementations.

The VSP G1000 was the winner in performance and throughput in the zSystems high-end storage. The VSP's G1500 performance improvements compared to the VSP G1000 entrench the G1x00 series as the performance champion.

These high performance is contributed by:

- ❑ VSDs as a third layer of processing (in addition to FEDs and BEDs) and the largest number of processing cores of any mainframe storage platforms in the industry,
- ❑ Special Data Accelerators processors and new ASICs for the FED and BED boards,
- ❑ The enhanced Hi-Star crossbar switch-based architecture has 768 GB/sec aggregate internal bandwidth,
- ❑ Separate caches and paths for control and data
- ❑ A custom designed quad- processor and ASICs in the HAF controller offloading tasks from the VSDs.

Technology

The VSP G1500 and the F1500 are designed and manufactured by Hitachi Ltd., a huge leading international technology corporation with 335, 244 employees at the end of business year 2016 and multi- billions in revenue. 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 semiconductor technology manufactures can plan for using the latest technologies. For example, Hitachi's Telecommunication & Network division contributed knowledge in designing the HiStar-crossbar switch, while the Micro Device division supplied the tailored ASICs and other components. Another Hitachi in-house developed technology is the embedded flash controller for Hitachi Accelerated Flash (HAF) storage to deliver superior performance at acceptable price.

Integration with other Platforms and Server Hypervisors

For mixed environments such as UNIX or Windows, the VSP G1000 provides full functionality support and integration. Storage virtualization amplifies the benefits of server virtualization. For example, HDS has been cooperating with VMware since 2002, becoming a VMware Premiere Global Partner in late 2006. The alliance as resulted in comprehensive interoperability via certifications and qualifications. VSP supports also other vendors' hypervisors.

Competition

There are many storage subsystems on the market, but only few can be considered as real high-end or tier-1 and only three[6] (last generation) are currently supporting the mainframe – EMC VMAX3, Hitachi VSP G1000 and IBM DS88xx series.

There are many differences between the three subsystems: The EMC VMAX structure is loosely coupled Intel Xeon E5-2687modules connected with InfiniBand bus, The IBM DS8xxx series uses the Power 8 active/active symmetrical multiprocessor with PowerPC processors in front- and back-end ports. The VSP G1500 and the F1500 are a tightly-tied multiple system using Virtual Storage Directors (VSDs) sharing a global cache linked by a non-blocking, crossbar architecture connecting the front- and back-end processors. There are significant differences in remote Copy techniques and Flash usage. For example both Hitachi and IBM use high speed flash modules sitting on the PCIe bus in addition to normal SSDs, while EMC use SSDs only.

IBM owns the zSystems mainframe architecture, which from time to time allows it to offer exclusive features for at least a limited period. EMC and Hitachi purchase the specifications for these features to remain compatible. IBM provides EMC and Hitachi with technical specifications but not the code, which the competing companies have to develop. Typically, Hitachi delivers the compatible features 6-24 months after IBM and ahead of EMC. IBM's XIV acquisition initially impacted the DS8xxx series. However they recovered after a couple of years and are trying to get the competitive edge by delivering new functions at a faster pace than in the past.

The latest high-end EMC storage the VMAX3 was launched without FICON channels. EMC suggested that mainframe users committed to EMC may purchase the outdated VMAX series 40K. It took EMC about nine months to deliver FICON functionality. EMC is also late delivering other z Systems compatible features. Launching Tier 1 products without FICON demonstrates EMC's shortage of developers and the lack of commitment to mainframe storage. As opposed to EMC, Hitachi is continuing to show strong commitment to the zSystems customers. For example on October 20[th], 2014 Hitachi announced improvements to the Hitachi VSP G1000 mainframe storage performance with added SVOS support for z/HPF features and new releases of two Hitachi mainframe storage products - Hitachi Tiered Storage Manager for Mainframe 8.1.0 and Hitachi Business Continuity Manager 8.1.0. Among the enhancements of Hitachi SVOS for the VSP G1000 microcode V02 are SVOS support for expanded z/HPF features such as High Performance FICON (zHPF), zHPF for QSAM/BSAM, zHPF List Prefetch Optimizer, BiDi for List Prefetch Optimizer, zHPF for format writes, and mainframe support for SVOS Cache Residency Manager, including setting and releasing cache-based volumes from mainframe hosts.

---

[6] Hewlett-Packard OEM the VSP G1000 as XP9500 but its market share is negligible.

Another two new, important features delivered in 2016 were the **zHyperWrite**, which speeds the DB2 log writes by eliminating a majority of the time required for the synchronous replication of data, and the **zHPF Performance for distance**, which is very valuable for multi-site GDPS infrastructures. Some other compatibility features are replaced by comparable Hitachi features or will be delivered in 2017.

Summary and Conclusions

Hitachi and Hitachi Data Systems are committed to providing industry-leading high-end storage for IBM mainframe computing environments and other platforms. Hitachi has again enhanced its state-of-the art high-end storage system. The VSP G1500/F1500 are the industry leader in every aspect: reliability, functionality, performance and sustainability. In the Gartner's latest "Critical Capabilities for General-Purpose, High-End Storage Arrays" report, Gartner positioned the predecessor of the G1500, the VSP G1000, as the leading product in this category, and the G1500 brings even more performance and enhanced functionalities. The VSP F1500 is all-flash appliance which delivers unmatched performance, availability, and functionality.

> *"Hitachi and Hitachi Data Systems are committed to providing industry-leading high-end storage for IBM mainframe computing environments and other platforms. Hitachi has again enhanced its state-of-the art high-end storage system."*

Hitachi, Ltd., maintains the strong relationship with IBM ensuring its ability to support the latest features and functions of z/OS as they are delivered. The VSP is compatible to zSystems, fully supports GDPS, HyperSwap, other platforms, and hypervisors. Additionally, Hitachi has developed and continues to develop new technology and capabilities to enable customers to maximize the efficiency and utilization of their storage and the availability of their data. It offers impressive set of functionalities among them three unique features; the control unit partitioning, storage virtualization and Active Flash.

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 G1500 or the F1500 on any shortlist of organizations searching for excellence[7].

---

[7] "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.

Appendix 1: List of supported IBM compatible functions and features:

Parallel Access Volumes (PAV).

HyperPAV.

Multiple Allegiance.

MIDAW.

Extended Address Volume (EAV).

Dynamic Volume Expansion (DVE).

Extended-Distance FICON.

IBM Metro Mirror (PPRC).

IBM z/OS Global Mirror (XRC)[8]

FlashCopy V02.

FlashCopy Space Efficient (FlashCopy SE).

Remote Pair FlashCopy.

Geographically Dispersed Parallel Sysplex™ (GDPS) HyperSwap.

Basic HyperSwap and TPC-R.

High Performance FICON (zHPF).

zHPF multitrack.

zHPF DB2[®]

List Prefetch

 zHPF Bi-directional channel program

zHPF List prefetch optimizer

zHPF BSAM/QSAM support

zHPF Format Write

zHPF and DB2I/O

z/OS Discovery and Auto Configuration (zDAC) support.

16 Gb/sec FICON Express

DB2 Hyperwrite

zHPF Performance for distance

---

[8] Hitachi support for z/OS Global Mirror (XRC) via Hitachi Compatible Replication for IBM[®] XRC[®] is only available on Hitachi VSP G1000 series as a migration aid for organizations moving to Hitachi Universal Replicator.

Appendix 2 SVOS 7.0 enhancements (not for zSystems)

The SVOS 7.0, announced in October 2016, is engineered with an improved flash-aware IO stack that includes enhanced QoS, new adaptive data reduction and support for all-flash cloud tiering. The major benefits are:

❑ **Quality of Service (QoS):** Prevents workloads from misusing bandwidth and delivers performance information in real-time, improving overall performance and service.

❑ **Adaptive data reduction:** Smart selection process optimizes which data is deduplicated first to improve response times and provides deduplication and compression within SVOS for 2:1, 5:1 or greater capacity savings. It automatically manages data reduction techniques for optimal performance; for example If FMDs detected, FMD compression is used, saving VSP cycles for other tasks. The computing power of the quad-core ARM processors of the FMDs allow much faster process deduplication than using the main controller processing power. By offloading this task from the VSP, it improves overall subsystem performance. If encryption required, SVOS compression used.

Another enhancement of the SVOS 7 is the direct cloud connectivity, which optimizes spend and enables faster transitions to all-flash. It supports the Amazon and Microsoft Azure public clouds and the Hitachi Content Platform for private clouds. The data move is transparent to users and applications, the files migrate to cloud automatically based on policies and can be recalled if required. Moving inactive or infrequent accessed files to cloud may reduce ongoing storage and data center costs and space, power and cooling requirements of Data Centers.

Management software updates

To ease storage management efficiency SVOS 7 brings enhanced Hitachi Storage Advisor, Automation Director, and Infrastructure Analytics tools and reports.

Hitachi Storage Advisor (formerly Hitachi Infrastructure Director) is a common configuration management application for the Hitachi Virtual Storage Platform (VSP) family that simplifies storage management complex tasks. It provides administrative capabilities in suggesting configuration setting (based on best practice experiences), storage pool and volume creation, replication and monitoring, which enable end-to-end management of the IT infrastructure. The SVOS7 supports the new linked-snapshot feature to create thin copies of data sets, enhanced data reduction management, and capacity savings reporting.

The Infrastructure Analytics Advisor licensed tool defines and monitors storage service level objectives (SLOs) for improved resource performance. It assists users in identifying and and analyzing historical performance trends and helps to optimize storage system performance and planning for capacity growth. Using analytics tools, users can find which resource (storage subsystems, servers, or volume) is causing a problem. When threshold levels are set, operators are alerted to threshold violations and possible performance problems before the start getting complains from end-users. The new release brings central

analytics reporting across multiple data centers, monitoring deduplication and compression savings.

Hitachi Automation Director is a licensed tool for data center administrators, which helps to automate and simplify end-to-end storage provisioning. It is built by different blocks of pre-packaged automation templates known as Service Templates. These pre-configured templates are customized to users' specific environments and processes for creating services that automate complex tasks such as resource provisioning. The new functions are: automated SAN zoning, automated (Global-Active Device) GAD provisioning, two or three data center support, and OpenStack cloud automation