

VSP Architecture Overview

HDS GLOBAL ACCOUNT TEAM

INTEGRATED SCALABLE PORTFOLIO

HITACHI
Inspire the Next

Hitachi Unified
Storage 100

Midrange



HUS-VM
Unified Virtual Controller
Entry Level Enterprise



Virtual Storage
Platform


Enterprise



Fully populated Dual Chassis VSP has 6 racks

2 DKC racks, each with a DKC box and 2 DKU boxes

4 DKU racks, each with 3 DKU boxes




	4 VSDs	8 VSDs
HDD (SFF)	1,024	2,048
FED ports	64 (80/96 ^{*1})	128 (160 ^{*2})
Cache	256GB (512GB) ^{*3}	512GB (1,024GB) ^{*3}

^{*1} 80 ports with 1 BED pair
96 ports in a diskless (all FED) configuration


^{*2} 160 ports with 1 BED pair per DKC module
(Diskless not supported on 2 module configurations.)

^{*3} Enhanced(V02)


- **FED – Front End Directors**
 - The Front End Director boards provide the interface to host servers and also to any external storage that may be attached to the VSP. FEDs can be either 16 port 8Gb/s FICON or 16 port Fiber Channel
- **VSD – Virtual Storage Directors**
 - A Virtual Storage Director is the central processor and data movement engine. Either two or four VSDs are installed in each chassis, four are used for added performance.
- **Cache**
 - The Cache boards or DCAs (Data Cache Adapter) hold the system memory. This contains transient user IO activity and also configuration details like RAID setup, dynamic tiering status and remote copy operations
- **BED – Back End Director**
 - The Back end Director boards interface to the disk or SSD devices. Each chassis can hold two or four BED boards and BED has eight 6Gb/s SAS links




VSP Single Chassis Architecture w/ Bandwidths



VSP Single Chassis Grid Overview



VSP Second Chassis - Uniform Expansion



- VSP is a completely new, highly scalable enterprise array
- VSP is the first “3D Array”
 - **Scales up** within a single chassis by adding logic boards (I/O processors, cache, host ports, disk controllers), disk containers and disks (to 1024 disks)
 - **Scales out** by adding a second fully integrated chassis to double the cache, disk capacity and host connectivity of a single chassis (to 2048 disks)
 - **Scales deep** with external storage
- VSP continues support of **Hitachi Dynamic Provisioning** and **Universal Volume Manager** (virtualized storage), as well as most other Hitachi Program Products available on the USP V
- VSP has a new feature within HDP named **Hitachi Dynamic Tiering** to migrate data among different storage tiers (SSD, SAS, SATA) located within a single HDP Pool based on historical usage patterns
- **VSP provides up to 40%** better power efficiency than USP V and a much smaller footprint

- The VSP shares no hardware with the USP V
- The VSP architecture is 100% changed from the USP V
- VSP does reuse much of the USP V software, such as HDP and other Program Products
- Major changes from the USP V include:
 - The previous *Universal Star Network* switch layer (PCI-X, 1064MB/s paths) has been upgraded to a new *HiStar-E* grid (PCI-e, 2048MB/s paths)
 - The *MP FED/BED processors* have been replaced with Intel Xeon quad-core CPUs located on a new *Virtual Storage Director* I/O processor board
 - The discrete *Shared Memory* system has been replaced by a *Control Memory (CM)* system. This uses processor board **local memory** plus a master copy in a region of **cache** that is updated by the individual VSDs
 - Each VSD board *manages* a discrete group of LDEVs that may be accessed from any port, and has a reserved partition in cache to use for these LDEVs
 - Individual *processes* on each VSD Xeon core dynamically execute tasks for the different modes: Target, External, BED (disk), HUR Initiator, HUR Target, various mainframe modes, and various internal housekeeping modes

- A **single chassis array** can include up to:
 - 3 racks and one logic box
 - 4 VSD boards
 - 64 8Gbps FC or FICON ports (no ESCON) – 8 FED boards*
 - 256GB cache – 8 DCA boards (using 4GB DIMMs)
 - 1024 2.5" disks (or 640 3.5" disks) – 64 HDUs
 - 32 6Gbps back-end SAS links – 4 BED boards
 - 65,280 Logical Devices
- A **dual chassis array** can have up to:
 - 6 racks and two logic boxes
 - 8 VSD boards
 - 128 8Gbps FC or FICON ports – 16 FED boards*
 - 512GB of cache – 16 DCA boards (using 4GB DIMMs)
 - 2048 2.5" drives (or 1280 3.5" drives) – 128 HDUs
 - 64 6Gbps back-end SAS links – 8 BED boards
 - 65, 280 Logical Devices

* More if some DKAs are deleted

- 2.5" SFF Disks (SFF DKU):
 - 200 GB SSD (3 Gbps**)
 - 146 GB 15K RPM SAS (6 Gbps)
 - 300 GB 10K RPM SAS (6 Gbps)
 - 600 GB 10K RPM SAS (6 Gbps)

- 3.5" LFF Disks (LFF DKU):
 - 400 GB SSD (3 Gbps**)
 - (~20% slower on writes than the 200 GB SSD)
 - 2 TB 7.2K SATA (3 Gbps)

** In the future, the SSDs will have the 6 Gbps interface.

- Disks of different interface speeds may be intermixed in the DKUs as the BEDs drive each "conversation" at the speed of the individual drive over the switched SAS back-end.

- Each **FED** board has a **Data Accelerator chip** (“DA”, or “LR” for *local router*) instead of 4 MPs. The DA routes host I/O jobs to the VSD board that owns that LDEV **and performs DMA transfers of all data blocks to/from cache**.
- Each **BED** board has 2 Data Accelerators instead of 4 MPs. They route disk I/O jobs to the owning VSD board and move data to/from cache. Each BED board has 2 **SAS SPC Controller chips** that drive 8 SAS 6Gbps switched links (over four 2-Wide cable ports).
- Most MP functions have been moved from the FED and BED boards to new multi-purpose **VSD** boards. **No user data passes through the VSD boards!** Each VSD has a 4-core Intel Xeon CPU and local memory. Each VSD manages a private partition within global cache.
- Unlike the previous Hitachi Enterprise array designs, the FED board does not decode and execute I/O commands. In the simplest terms, a VSP FED accepts and responds to host requests by directing the host I/O requests to the VSD managing the LDEV in question. The **VSD processes the commands**, manages the metadata in Control Memory, and creates jobs for the Data Accelerator processors in FEDs and BEDs. These then transfer data between the host and cache, virtualized arrays and cache, disks and cache, or HUR operations and cache. The VSD that owns an LDEV tells the FED where to read or write the data in cache.

- In VSP, VSDs manage unique sets of LDEVs, and their data is contained within that VSD's cache partition. Requests are routed to the VSDs by the Data Accelerator chips on the FED and BED boards using their local LDEV routing tables.
- LDEV ownership can be looked at in Storage Navigator, and may be manually changed to another VSD board.
- When creating new LDEVs, they are round-robin assigned to the installed VSDs in that array.
- If additional VSDs are installed, groups of LDEVs will be automatically reassigned to the new VSDs. There will be a roughly even distribution across the VSDs at this point. This is a fast process.
- LDEV ownership by VSD means that VSP arrays don't have an LDEV coherency protocol overhead. There is only one VSD board that can manage all I/O jobs for any given LDEV, but any core on that Xeon CPU may execute those processes.

- Note that a VSD controls all I/O operation for an LDEV, whether it is processing a host I/O, a disk I/O, an external I/O, or a Copy Product operation.
- Copy products **PVOLs and SVOLs must be on the same VSD, as the data has to available from the same cache partition.**


- Basically, on the USP V, we know that
 - Small block I/O is limited by **MP busy rate** (FED-MP or BED-MP busy)
 - Large block I/O is limited by **path saturation** paths (port MB/s or cache switch path MB/s, etc.)
- On VSP, the “MPs” are separated from the ports.
 - Where there are multiple LDEVs on a port, these can be owned by different VSD boards.
 - Where there are multiple LDEVs on a port that are owned by a single VSD board, the 4 cores in the VSD board can be processing I/Os for multiple LDEVs in parallel.
- VSP can achieve very high port cache hit IOPS rates. Tests using 100% 8KB random read, 32 15K disks, RAID-10 (2+2), we saw:
 - USP V: 1 port, about 16,000 IOPS (2 ports-2MPs, 31,500 IOPS)
 - VSP: 1 port, about 67,000 IOPS (2 ports, 123,000 IOPS)

VSP Architecture Overview

Fully populated Dual Chassis VSP has 6 racks

2 DKC racks, each with a DKC box and 2 DKU boxes

4 DKU racks, each with 3 DKU boxes



	4 VSDs	8 VSDs
HDD (SFF)	1,024	2,048
FED ports	64 (80/96 ^{*1})	128 (160 ^{*2})
Cache	256GB (512GB) ^{*3}	512GB (1,024GB) ^{*3}

^{*1} 80 ports with 1 BED pair
96 ports in a diskless (all FED) configuration

^{*2} 160 ports with 1 BED pair per DKC module
(Diskless not supported on 2 module configurations.)

^{*3} Enhanced(V02)

- The VSP can be configured as a single or dual chassis array. It is still a single homogeneous array.
- A VSP might be set up as a dual chassis array from the beginning, with a distribution of boards across the two chassis.
- A single chassis VSP can be later expanded (Scale Out) with a second chassis.
- The second chassis may be populated with boards in any of these scenarios:
 - Adding 2 or 4 Grid Switches and 4-8 Cache boards to provide larger amounts of cache
 - Adding 2 or 4 Grid Switches and 2-4 VSDs to add I/O processing power (for random I/O)
 - Adding 2 or 4 Grid Switches and 2-8 FEDs to add host, HUR, or external ports
 - Adding 2 or 4 Grid Switches and 1-2 BEDs to add disks and SAS paths
 - Any combinations of the above

VSP Second Chassis - Uniform Expansion




Table of Limits	VSP		USP V
	Single Chassis	Dual Chassis	Maximum
Data Cache (GB)	32-256 (512)	32-512 (1024)	512
Raw Cache Bandwidth	64GB/s	128GB/s	68 GB/sec
Control Memory (GB)	8-48	8-48	24
Cache Directories (GB)	2 or 4	6 or 8	-
SSD Drives	128	256	128
2.5" Disks (SAS and SSD)	1024	2048	-
3.5" Disks (SATA, SSD)	640	1280	-
3.5" Disks (FC, SATA)	-	-	1152
Logical Volumes	65,280	65,280	65,280
Logical Volumes per VSD	16,320	16,320	-
Max Internal Volume Size	2.99TB	2.99TB	2.99TB
Max CoW Volume Size	4TB	4TB	4TB
Max External Volume Size	4TB	4TB	4TB
IO Request Limit per Port	2048	2048	2048
Nominal Queue Depth per LUN	32	32	32
HDP Pools	128	128	128
Max Pool Capacity	1.1PB	1.1PB	1.1PB
Max Capacity of All Pools	1.1PB	1.1PB	1.1PB
LDEVs per Pool (pool volumes)	1024	1024	1024
Max Pool Volume size (internal/external)	2.99/4TB	2.99/4TB	2.99/4TB
DP Volumes per Pool	~62k	~62k	8192
DP Volume Size Range (No SI/TC/UR)	46MB-60TB	46MB-60TB	46MB-4TB
DP Volume Size Range (with SI/TC/UR)	46MB - 4TB	46MB - 4TB	46MB - 4TB

Q and A