

SCI StorInt™ Dispatch

HDS announces Hitachi VSP

HDS recently announced their latest version of enterprise storage, the Hitachi Virtual Storage Platform (VSP). The Hitachi VSP is a scale up and out update of their popular USP-V storage subsystem, which supports storage virtualization, and 3D scaling. The VSP is a significant technological advance over prior generation Hitachi storage.

Hitachi VSP

Hitachi VSP provides 3D scaling, i.e., scale out, up and deep in enterprise class storage. Scale out is accomplished by dynamically combining multiple units into a single logical system with shared resources, scale up is by dynamically adding processors, connectivity and capacity in a single unit and scale deep is by virtualizing external storage.

VSP hardware

The VSP incorporates a controller blade cabinet design and now uses 19" racks. A single controller rack contains one controller chassis and includes up to two drive chassis. One controller chassis can support from 2 to 4 Virtual Storage Director (VSD) blades, from 2 to 8 Cache Adaptors (CA), from 2 to 12 Front-end Director (FED) blades, and from 2 to 4 Back-end Director (BED) blades. Two controller chassis can be connected together to form one storage subsystem and would support up to 8 VSD blades, up to 24 FED blades with 192-8Gbs FC/FICON ports, up to 16 CAs with 1TB of total cache, and up to 8 BED blades with 64-6Gbs SAS storage links.

A minimum internal storage configuration would include 2-VSDs, 2-FEDs, 2-CAs and 0-BEDs supporting a diskless storage subsystem with from 64GB to 1TB of cache storage and no BEDs. Although, a diskless system makes sense as a storage virtualization subsystem used with external storage it could also be used for migration.

Up to 4 additional storage only racks can be interconnected to the storage complex for a total of 16 drive chassis across 6 racks. The VSP also includes a 5th generation Hitachi cross bar switch for internal connections and uses PCI-E as an interconnect between the two control chassis. Together the total system acts as one storage subsystem complex with any VSD processor core accessing any individual FE, BE, or cache resource to support I/O and other storage processing activity.

A single VSD contains a quad-core Intel processor, which speeds up I/O processing and storage administration. Also the shared or control memory is now distributed in L2 cache across all VSDs rather than centralized as in the USP-V. IO and other processing activity are parceled out across all available VSD processor cores in the cluster to balance subsystem performance.

VSP cache is used to hold all read and write data, is global across the cluster. Write data is always mirrored in cache and backed up to onboard flash drive(s). Each control chassis supports up to 512GB of cache.

BED and FED blades have onboard data accelerator ASICs to assist in data path processing, IO routing, parity calculations and encryption activities. Any IO coming into a Front-end port can be directly routed to whichever core handles that particular LDEV.

Drive storage is attached to control chassis via redundant, 6 Gbs SAS links. Disk storage comes in 2.5" and 3.5" drive configurations, with each 14U storage bay containing up to 128-2.5" or 80-3.5" drives per chassis. As expected, any drive can be replaced without powering down the subsystem. Also, SAS or SATA drives can be intermixed. In addition, the VSP also supports 2X the flash drives or SSDs in USP-V, using highly reliable SLC NAND technology. SSDs support up to 70X random read performance or up to 14X random write performance more than disk drives.

Finally, cooling airflow for the VSP is from front to back and power is now single phase vs. three-phase for USP-V. The VSP uses about 48% less power than the USP-V.

VSP functionality

Hitachi's new 3D management is to manage UP, OUT and DEEP. Manage UP refers to their new automated data placement/dynamic tiering, dynamic provisioning, the ability to manage up to 5M logical objects, 255PB of virtualized capacity and the ability to scale up virtual server deployments.

- Dynamic Tiering extends to external storage and applies on a 42MB page (sub-LDEV) basis to a storage pool. Three storage tiers can be supported. Dynamic Tiering can be done at the page level or at the object level (LDEV or file). Data is written first to highest performing tier and then later migrated to other tiers based on reference activity.
- Dynamic (thin) provisioning can now be extended to external storage. Fixed provisioned volumes can be converted to thin volumes by migrating them into a Dynamically Provisioned storage pool. Also, Dynamic Provisioning automatically stripes data across all drives in the pool; increasing IO performance considerably over fixed provisioned LDEVs.

Manage OUT refers to their new unified management scheme under the Hitachi Device Manager (HDvM) for block, file and content storage across all Hitachi storage systems. This includes end-to-end visibility of applications, virtual machines/servers and logical storage devices for traditional and VMware environments.

Manage DEEP refers to the tight integration between Hitachi Device Manager and Hitachi Command Suite to supply consistent implementation of management activities, service level management for capacity and performance, and common GUI, data repository and workflows across the two software suites.

In addition, the VSP shares advanced storage functionality of the USP-V subsystem with the addition of dynamic tiering at the page level and dynamic provisioning for external storage. For example, Hitachi VSP supports multi-tenancy for guaranteed QOS, storage virtualization for externally attached storage, automated data mobility of objects or blocks to different storage tiers, VMware VAAI support, storage encryption, Hitachi

Universal Replicator, High Availability Manager, and Virtualized storage replication to name just a few.

At initial release there are some restrictions with storage replication, e.g. no support for Extended consistency groups and no delta resynch support. However the VSP can replicate to a USP-V.

VSP performance

Preliminary performance numbers were provided that shows the substantial performance advantage of the VSP. Specifically, random read is 1.9X and sequential read throughput is 1.4X better than the current USP-V. Also, random writes are 2.7X better and sequential write throughput is 1.2X better than the USP-V subsystem. These results represent full VSP and USP-V subsystems, the VSP per front-end port performance numbers are even better. Also these performance numbers come from hard disk drive subsystems, which don't include Flash or SSDs in the VSP or USP-V.

Announcement significance

Obviously, Hitachi has taken another major step to upgrade their enterprise class storage subsystems with their VSP offering. The VSP offers a number of advantages over their current USP-V system including better performance, lower power consumption, better storage density, to name just a few.

VSP compares well to EMC's VMAX Symmetrix subsystem. While VMAX is also scaleable, EMC uses separate nodes, which contain processors as well as cache, front-end, and back-end control and uses a switching fabric for inter-node communications. VMAX was a radical and total revision of their previous Symmetrix architecture. At announcement, VMAX was described as a cloud enabler because of its scaleability.

However, Hitachi's move to a blade system is not as radical a departure from their USP-V storage architecture. Yes, adding more and faster cores, L2 control cache, 2.5" disks, SSDs, SSD backed up cache, and using PCI-E for control chassis inter-connects all advance Hitachi storage technology to the next level. But one can see in all this a continuation and expansion of the USP-V architecture rather than a complete restructuring like EMC's VMAX. It will be interesting to see how VSP is received in the market place

Perhaps just as important is the continuing advance of storage functionality including Dynamic Provisioning for external storage and Dynamic Tiering at sub-LDEV level for both internal and external storage. Adding this new functionality to Hitachi's extensive storage virtualization capabilities, one can have all the latest functionality applied to older storage subsystems.

Silverton Consulting, Inc. is a Storage, Strategy & Systems consulting services company, based in the USA offering products and services to the data storage community.