Hitachi USP VM - the "Little Big Man" Disk Storage Subsystem

Josh Krischer is an expert IT advisor with 37 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.

Hitachi Data Systems' new USP VM intelligent virtual storage controller reminds me of the 1970 movie ***Little Big Man,*** which features award-winning performances by Dustin Hoffman and Chief Dan George. The story's main character is Jack Crabb (played by Hoffman), a dying centenarian who recalls several facets of his long and unusual life. The movie's central theme is Crabb's adoption by the Cheyenne band of American Indians who give him the name "Little Big Man", because although he is short, he is very brave.

On September 10th 2007, Hitachi Data Systems announced the USP VM, a rack mounted high-end disk storage subsystem for System z and other platforms. Its impressive specifications and functionality for a subsystem of such a small size inspired me to characterize it as the "little big man" storage subsystem.
Two models are being introduced; a capacity model and a "disk-less", model which is designed to deliver the most advanced storage functionality to tier-one or 3rd party storage subsystems which are missing these functions.

Basic Design

The USP VM is a smaller,"" but "very brave" version of Hitachi's industry-leading Universal Storage Platform V, (announced on May 14th, 2007and technically equivalent to HP's XP20000 and the Sun StorageTek 9985V) It replaces the NSC55 (HP XP10000, Sun StorageTek 9985) launched in 2005. Similar to its bigger sibling, its design is based on a massively parallel crossbar switch architecture (called the Hitachi Universal Star Network V), mirrored data cache, mirrored control cache, channel host Front-end Directors (FeDs) and Back-end Directors (BeDs). The central point of this design are the Application Specific Integrated Circuits (ASICs) of the non-blocking crossbar switch architecture technology, which have embedded logic for checking, routing and managing data and have been designed specifically for the USP V series. They were developed collaboratively by engineers from Hitachi's supercomputer, semiconductor, networking, and data storage research divisions. As opposed to some other vendors, Hitachi has access to research and intellectual property from multiple IT disciplines.

Major Specifications

	USP VM
Raw Maximum Supported Internal Capacity	**72TB**
Maximum Supported Total Capacity (Connected Externally)	**96 PB**
Internal Bandwidth	**13.3GBps**
Maximum Cache	**64GB**
Maximum Control Memory	**16GB**
Back-end Switched Fabric	**4 GBps**
Front End Fibre Channel Ports	**48 x 4 GBps**
Maximum Internal Disks	**240 x 4 GBps**
Maximum LDEVs	**64k MF & OS**
External Storage Connection	**FC**
Maximum ESCON PORTS	**24**
Maximum FICON PORTS	**24**
Size of Single LDEV	**2TB (to be expanded later)**
NAS and iSCSI Support	**High Performance NAS Platform iSCSI – Possible Future Capability**

Hitachi's USP VM supports a cache of up to a maximum capacity of 64 GBytes for data. The cache meta-data is held in a separate, dedicated cache. Only the "write portion" of the cache is mirrored with a threshold that is automatically adjusted depending on activity. Hence, the effective cache size is reduced by ca.20% for a typical workload. The dynamic cache structure and the separate control cache allow for dynamic configuration changes in the data cache by changing bits in the control store through a service processor. Hitachi's cache design is the most advanced in the industry and provides any-to-any connectivity between any host port and disk array. Access to storage can also be load-balanced across multiple host ports since they can all view the same cache image. This provides additional resilience since the failure of any one or two components would not be noticed by the end-user. The maximum available bandwidth is 13.3 GBytes/s which puts it more in the high-end scale than the mid-range.

Scalability

The Hitachi USP VM supports up to 72 TBytes of internal capacity and 92 Petabytes of externally virtualized storage. See the Universal Virtualization Layer in the next section.

Availability

The USP VM ensures non-disruptive upgrades and maintenance, “hot swap” components and online microcode changes. The cache is mirrored for writes, and control data is stored in separate mirrored “control” memory. Its channel adapters include spare microprocessors, which allows microprogram updates even if alternate path and path-balancing software is not installed. All data stored in cache is protected by batteries, which allows it to be transferred to disks in the event of a main power failure.

The subsystem supports three RAID techniques: RAID-1 (2D+2D), RAID-5 (3D+1P) or (7D+1P) and RAID-6 (6D+2P). The RAID-6 method, which allows recovery from two simultaneous disk failures without data loss, offers rebuild times that are significantly faster than RAID-5 rebuild times.

Remote Copy Techniques

The USP VM supports all the remote-copy services of the USP V series, including Hitachi’s TrueCopy Synchronous, Asynchronous, and the Universal Replicator. All these techniques include advanced data consistency and integrity-keeping mechanisms. The subsystem also supports Hitachi ShadowImage (point-in-time copy) for all computing platforms and Hitachi Compatible FlashCopy (snapshot type) Version 2 for z/OS.

Functions and features

Most Hitachi-developed features are based on control unit virtualization, starting in 2002, when Hitachi launched the 9980V with a new feature called Virtual Storage Ports. In the USP VM, this virtualization layer provides up to 1,024 virtual ports for each of the 48 physical ports, including LUN0 for booting. A mode set is specified at the sub-system to set the appropriate server platform and provides separate storage pools for each host. With separate LUN addressing, QoS (access priorities), and LUN security, each storage domain appears as a separate virtual array despite using the same physical port. This ensures safe multi-tenancy as there is no danger of overwriting another server’s data. Multiple hosts can safely share a common physical storage system, since each host can be assigned its own virtual private storage. This embedded virtualization layer is particularly useful for supporting heterogeneous clusters and server virtualization such as VMware ESX. In fact, Hitachi claims that it is the first vendor to obtain external storage virtualization certification from VMware and the USP VM has ‘Day 1’ support for VMware ESX Server 3.0 to facilitate the integration of server and storage virtualization infrastructures..

This port virtualization is similar to idea of the NPIV (N-Port ID Virtualization) Fibre Channel standard, which 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 as if each image had its own unique physical port. The adoption of the NPIV standard is planned for 2008.

In addition to the Virtual Ports mentioned above, the USP V supports the following unique features:

- **Hitachi Dynamic Provisioning or "Thin provisioning"** enables allocation of virtual storage as needed without the need to dedicate physical disk storage up-front. Additional capacity can be allocated without any disruption to mission-critical applications from existing or newly-installed capacity. This feature, in addition to saving investment and running costs (less energy consumption, smaller floor space), also improves the performance by striping the data across all the disks in the array. Striping the data among a large number of physical devices practically eliminates "hot spots", which results in almost uniform performance. Currently, this feature is supported on internal storage only; however starting in 4Q07, it will be extended to externally connected storage as well.

- **Universal Virtualization Layer** (introduced with the first version of the USP in 2004). The virtualization layer is embedded in the processors of the USP VM channel adapter cards. These cards function as a normal port for volumes which reside internally or as a host bus adapter for accessing external storage, which may be Hitachi's or from a third party. Hitachi Data Systems' Universal Volume Manager software configures, manages and accesses external volumes in a similar way as if they were USP V internal volumes. Externally connected storage may use the same functionality as internal storage, which means that data replication software and other applications 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 building fluid tiered storage environments without compromising on functionality. It gives customers the ability to store non-critical data or to archive mainframe data on low-cost SATA systems, for example.

- **Virtual Partition Manager** is sub-system partitioning (introduced in 2004 on the original USP) that allows resources (internal and externally attached) such as capacity, cache and ports to be dynamically partitioned into "virtual machines," each with its own virtual serial number (for asset tracking and charge-back purposes). Up to 16 of these virtual machines can be created, each separately managed and password-protected, to provide better resource allocation and enhanced protection by isolation between the various partitions. This capability enables users to build different internal service levels, to separate test from production systems, and to reduce the costs for users that previously, for data security reasons, may have required separate storage sub-systems.

- **Storage Security Services,** which includes several functions, some of them introduced as early as with the 7700 subsystem in the mid '90s. These functions include:

- Controller-based data shredding.
- Write Once Read Many (WORM) software for tamper-proof data protection (required by most compliance regulations)
- An Audit Log file storing a history of all user access operations performed on the system to allow the tracing of unauthorized access to data.
- LUN Security for LUN-level access control to world wide names (WWN)
- Fibre Channel Secure Protocol Authorization A (FC-SP Auth A) for the authentication of Fibre Channel entities, as well as support for encryption appliances such as Decru and NeoScale

- **Hitachi Universal Replicator,** which is asynchronous, storage-agnostic data replication software for internal and externally attached storage. This technique, instead of using the cache, uses a disk to temporarily log the data before transferring it to the remote site(s), and thus, significantly reduces cache utilization and bandwidth requirements. This technique is particularly useful for users deploying disaster recovery with heterogeneous storage subsystems and external service providers.

Product Positioning

It is not easy to position the USP VM, and in fact, it deserves to be in a class of its own. There is a huge gap between the USP VM and the many mid-range storage subsystems such as EMC's CLARiiON, HP's EVA, IBM's DS 4x00 series, NetApp's FAS6000/9000 and the 3PAR InServ line. The USP VM leads the industry from every aspect: functionality, scalability, connectivity, serviceability, ease-of-use, and availability.
It also supports the System z mainframe, a platform which is not supported by mid-range subsystems. The two closest competitors are the EMC Symmetrix DMX-4 950 and IBM's DS6800; however, these subsystems cannot match the impressive specifications of the USP VM either. Clearly, Hitachi's innovation engine is moving the company further ahead of competitors in entry level enterprise-class storage technologies.

Possible Solutions with the USP VM

The advanced storage virtualization services of the USP VM enables clients to build technically advanced infrastructures at very affordable costs. The "disk-less" USP VM can be deployed as the industry's most advanced and functionality-rich virtualization platform. The two obvious solutions are virtualization of heterogeneous storage subsystems, or the ability to connect mainframes to tier-two or tier-three storage systems. Data migration can often pose a challenge; therefore, another solution is to use the USP VM as a non-disruptive data migration solution. The sub-system partitioning enables the consolidation of several clients on the same subsystem while providing robust security services. This may save operational costs, particularly for external service providers. TrueCopy Synchronous and Hitachi Universal Replicator can be used by external disaster recovery service providers to establish remote data replication for heterogeneous storage platforms.

Pricing

Hitachi will price the USP VM above conventional mid-range storage offerings but below its USP V high-end series. Its specifications relative to typical mid-range subsystems justify a premium price above the mid-range price level. The ability to support third party storage subsystems allows extending the life of these subsystems and ensures investment protection. The software features are priced “per frame” and not per capacity, which allows a lower price and simplifies procurement and upgrades. Taking into consideration that it delivers more advanced functionality than high-end storage subsystems from Hitachi’s competitors, the USP VM provides the “best value for your money” in the storage market.

Summary

As the new “Little Big Man” of the storage industry, the USP VM is a unique product with no real competition. It delivers the most advanced functionality in a compact, rack-mount infrastructure at an affordable price. It enables the deployment of common storage services, original solutions, solves technical constraints and saves capital and operational expenses. It should be evaluated by large enterprises but also small and medium organizations seeking excellence for their IT infrastructure.