Hitachi's USP VM: Services Oriented Storage for SME

Quick Note

John Webster
10 September 2007

As expected, Hitachi Data Systems (HDS) has introduced a smaller version of its large enterprise-class, virtualized disk storage subsystem (USP V). This smaller USP VM can be ordered with disk (the standard configuration) or as a diskless, rack-mounted virtual controller. Either way, it's a scaled-down version of Hitachi's USP V virtualized, storage services platform originally introduced for large enterprises.

With the announcement of the USP VM, Hitachi continues to execute on its strategy of scaling-down enterprise-class storage technology and making it available in smaller, modular, more affordable form factors, thereby expanding its addressable markets. The NSC55, introduced in mid-2005, was also a scaled-down version of the original TagmaStore USP; a diskless version was introduced in mid-2006 for those customers who needed virtualization services and no additional storage. With the USP VM, HDS replaces the Fibre Channel Arbitrated Loop (FCAL) connection between the controller and disk arrays with a switched FC network, and significantly enhances performance and internal bandwidth over the NSC55 by approximately 75%.

The USP VM puts Hitachi's unique, controller-based storage virtualization engine within the financial reach of small-to-medium enterprises (SMEs). SME IT administrators can now buy the Hitachi USP V controller, along with many of Hitachi's information management services, in a smaller, more affordable package.[1] This will be of particular interest to SME users who now plan to migrate from a networked server environment to a virtualized server environment (VMware, XenSource, etc.). Small to medium-sized mainframe users who are running open systems-based applications alongside z/OS-based applications could also see a way to consolidate and upgrade both application sets to virtualized storage. Indeed, with the USP VM, one could uniquely support both a virtualized mainframe and a virtualized open systems server environment from the same consolidated storage subsystem.

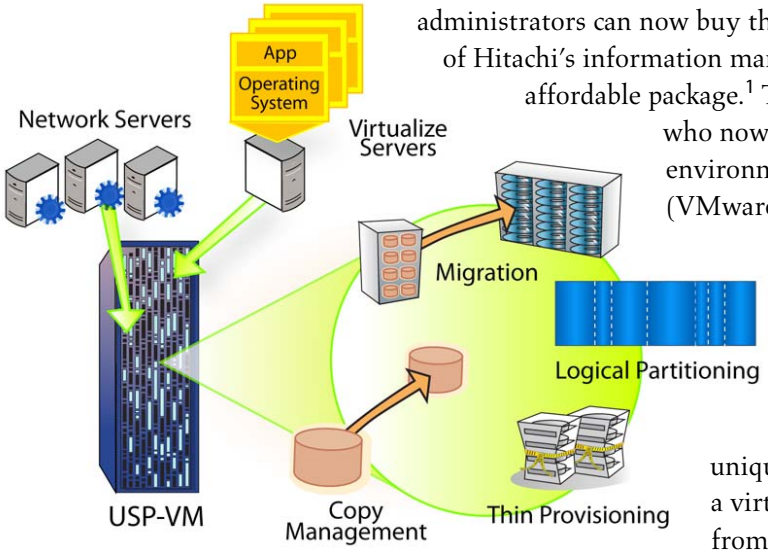

[1] For a more detailed discussion of the Hitachi USP V and its platform-based services, see our *Hitachi Sets High Water Mark for Consolidated Storage*

Going For Affordability

HDS has done two things with the USP VM that make it more affordable than its flagship USP V. First, pricing for the USP VM software that provides functions like logical partitioning, thin provisioning, and non-disruptive data migration between unlike disk arrays has been greatly simplified. Users can buy basic bundles and later graduate to bundles that add more functionality and support larger operating environments. Of course, they can also simply go with the more functional bundles right from the start.

Second, HDS offers both a standard edition that includes Hitachi disk arrays and a "diskless" version. The standard edition includes a minimum of 584 GB of disk storage (146GB 15k FC HDDs), expandable to 72 TB internal, plus external support for up to a total of 96 PB. It also has up to 48 FC and FICON ports; and up to 64 GB of cache.

With the diskless edition, users can source the underlying disk arrays themselves. These arrays can be provided new from Hitachi, or bought new or used from another vendor (as long as it's supported for external attachment to the USP VM). Of special note, users can also simply attach existing, supported arrays that are already installed and in service at their sites.

USP VM's diskless edition comes with a minimum of 8 GB of controller memory, 4 GB cache, and 8 Fibre Channel ports. Everything else can be user-added into the rack or attached to the USP VM controller, giving IT administrators significant flexibility with regard to the use of existing storage infrastructure. Storage mounted inside the USP VM rack must be sourced from HDS, but SAN switches mounted inside the rack can be sourced directly by the user. External arrays from other vendors may be attached to the USP VM controller for up to 96 PB of total managed disk capacity. The USP VM rack can be expanded with additional shared memory, cache memory, service processor, and another battery chassis for supporting larger stores of shared memory and cache.

The diskless edition supports many of Hitachi's virtualized controller-based services including thin provisioning, heterogeneous data migration, and VMware support. However, for mixed open systems and z/OS workloads, virtual tape library (VTL), and support for Hitachi's Content Archive Platform (HCAP), users must order the standard version.

USP VM Software Bundles

For the USP VM, HDS has done away with capacity-based pricing for core subsystem software components. Rather, users can now purchase graduated software bundles, each under a single license key as follows:

USP VM Operation – For USP VM operation and virtualization of underlying storage assets, users will need Hitachi Basic Operating System (BOS) software. BOS includes the subsystem OS, HiCommand Device Manager and all Resource Manager components, Server Priority Manager, and Virtual Partition Manager with support for up to 4 partitions. If support for attachment of external storage arrays is required, an additional Hitachi BOS Virtualization (BOS V) bundle is required that includes Universal Volume Manager. BOS V also includes support for up to eight logical partitions.

Data Copy and Migration – For managing the replication of data within a single disk array, or between disk arrays attached to a single USP VM, users will require the In-System Replication software bundle that includes both cloning and snapshot copy software. Both of these capabilities add significant value in VMware environments (more on this later). Hitachi In-System Replication software includes ShadowImage, which creates fully functional, RAID-protected clone copies of data objects (files, volumes, virtual machines, etc.) for open systems and z/OS, and Copy-on-Write Snapshot software that provides logical, pointer-based copies of data objects. All copy functions are supported between heterogeneous arrays attached to the USP VM with both BOS and BOS-V bundles installed. An additional bundle is offered for z/OS

users requiring IBM FlashCopy-compatible software.

Business Continuance and Disaster Recovery – The Hitachi Disaster Recovery software bundle is needed for replicating data between two USP VM images at local, metro-cluster, or geographically remote sites. DR software includes TrueCopy for synchronous replication and Hitachi Universal Replicator for asynchronous data replication. DR software supports both open systems and z/OS environments. For users requiring more robust data replication and recovery capabilities requirements, users will need Hitachi Disaster Recovery Extended software instead. The extended version supports three-data-center (3DC) disaster recovery configurations for open systems and z/OS and (4x4) for z/OS only.

The Services Oriented Storage USP VM

We have defined Services Oriented Storage as operational services built around virtualized storage—services that can be deployed between heterogeneous server and OS environments above, and heterogeneous storage capacity below. Services in this context are defined as outcomes, with the real work performed by the virtualization engine.[2]

For VMware users in particular, understanding virtualized storage controllers like the USP VM is highly useful. While VMware is essentially an operating system, VMware now also provides systems management tools and extended facilities ("layered services") along with the base ESX hypervisor. In fact, it is increasingly add-ons like VMotion and DRS that set VMware apart from its competitors. Indeed, VMware reports that sales of the hypervisor alone now account for only 20% of annual revenue.

As VMware grows in popularity, we will see more add-on tools and management applications from both VMware itself and other vendors wishing to add value to VMware environments. Users therefore have a choice: standardize on the VMware platform including the add-ons, or supplement VMware installations with value-add toolsets and applications from third-party vendors.

With the USP VM, HDS offers Dynamic "Thin" Provisioning, non-disruptive data replication and migration tools, and volume managers. All of these have analogs to some degree in VMware. In some cases the decision whether or not to go with the HDS storage and data management facilities integrated in the USP VM will be fairly obvious and straightforward. In other cases, the trade-offs will need to be considered. For those implementing USP VM in a VMware environment, the question isn't binary, but when to use which facility, and how. We offer two examples:

Snapshots and Clones

In VMware ESX 3.0, an administrator takes a snapshot or a clone copy of a virtual machine, which can be used to restore the VM, either on the same server or a different one. This process requires the creation of a stored map of changes (called a "redo log file") that occur up until the next snapshot is taken. Cloning creates a usable copy of a virtual machine. These processes are heterogeneous with regard to the underlying infrastructure. However, they all occur on the application server. Cloning also requires server-storage processing and I/O traffic.

Hitachi SOSS Model

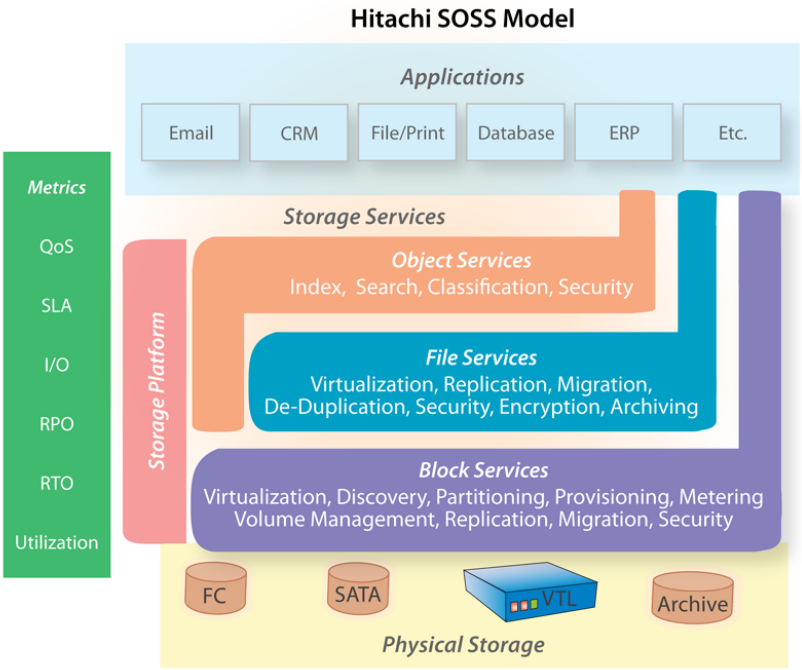

[2] See our *Services Oriented Storage*

In contrast, the processing of HDS' ShadowImage and Copy-on-Write Snapshot functions takes place entirely within the USP VM virtual controller, requiring no host processing cycles or server-storage I/O. Copy functions therefore run faster, as do the applications running on the VM server. The same argument holds for using storage-resident TrueCopy and Hitachi Universal Replicator vs. using the VMware functions in a disaster recover scenario. Here, performance is a critical consideration. As a result, many IT administrators have traditionally opted to run DR-related remote copy functions between storage controllers rather than application hosts.

While neither sets of copy functions are mutually exclusive and could, in fact, be used in conjunction with one another, the VMware copy functions result in less desirable outcomes than those of HDS' clone and snapshot functions. Using both together, while technically feasible, adds management complexity.

There is at least one trade-off to consider. When choosing the HDS copy function over one provided in VMware ESX, one must be prepared to either manage this process from a separate USP VM console, or write scripts to integrate the management of these processes from VMware's VirtualCenter.

Thin Provisioning

Using the VMware `vmkfstools` command, one can create individual VMDK files that are "thinner" than VMware's normal "zeroed thick" VMDK files. While useful, this capability is not thin provisioning as the concept is now commonly understood in the context of virtualized storage. The VMware concept refers to individual files that have been reduced in size rather than how a physical storage pool is made to appear, through virtualization of the underlying physical storage, as though it were larger than it actually is. In outcome, using the `vmkfstools` command is more comparable to data compression than to thin provisioning.

In contrast, HDS offers Dynamic Provisioning Software on the USP VM. Dynamic Provisioning allows administrators to apportion storage to users from a logical pool of storage capacity that appears to be larger than the actual physical pool of storage. Doing so greatly facilitates the job of provisioning storage on an as-needed basis. It also tends to keep storage utilization rates pleasingly high, because storage administrators are not forced to overprovision physical storage in order to keep up with continual demand—a situation that often results in "stranded" unused disk capacity as time goes on. Thin provisioning lets storage administrators add physical storage to the pool on a "just-in-time" basis, saving administrative time and money.

Here, the two functions are complementary. An administrator can have "thin" VMDK files stored within the USP VM running Dynamic Provisioning. However, when implementing Dynamic Provisioning software, storage administrators need some way of knowing when physical capacity is in danger of becoming used-up, which Hitachi provides with its thin provisioning services.

Conclusion

All announced USP VM functionality will be generally available in early October. We expect HP and Sun to announce OEM versions shortly after HDS' September 10 announcement. Upon shipment, the USP VM will support attachment to virtualized server environments running VMware ESX 3.0.2. The USP VM, its Universal Volume Manager (for management of heterogeneous external storage volumes), and array management software will also have been certified by VMware.[3] As such, HDS will be the first vendor to obtain external storage virtualization certification directly from VMware.

The USP VM will find applicability in a wide range of mid-tier IT environments from mainframe to open systems, and in shops that use both in

[3] For more information regarding USP and NSC VMware certification, see tinyurl.com/ysn4bb

combination. We have highlighted server virtualization in general, and VMware in particular, as an interesting application scenario because an IT administrator can see a virtualized storage controller like the USP VM as a bridge between networked servers on one side, and virtualized servers on the other. Most virtualized server installations begin in a test bed. Once users are familiar with managing the new environment, they begin to move server virtualization into mainstream production by migrating a few applications at a time. The USP VM can be attached simultaneously to the legacy server network and the virtualized server farm, transferring data non-disruptively from the old to the new platform as IT administrators migrate applications. As such, it acts as a bridge from old to new.

The USP VM can act as a bridge in other ways. It can be used, for example, to migrate data from a legacy array that is coming off lease, for example, to its replacement array, without disruption. Or, by using the USP VM to create a storage tier above the existing storage infrastructure that does not disrupt normal operations, it can be used to retire expensive copy functions that are bound to like arrays from the same vendor and replace them with copy functions that operate between unlike arrays from different vendors.

The rapidly growing acceptance of server virtualization is having a knock-on effect with the acceptance of storage virtualization. Now becoming more common in large IT organizations, storage virtualization is catching on in smaller enterprises. We advocate a common-sense evolution into storage virtualization delivered not as hunks of sheet metal or numerous vendor APIs, but as flexible, specially-built services mated with a highly functional platform. The USP VM demonstrates that Hitachi is clearly thinking much along the same lines. It is now placing its Services Oriented Storage solutions within the grasp of SME users.