Hitachi Sets High Water Mark for Consolidated Storage

Quick Note

John Webster
14 May 2007

Since the inception of networked storage in the late 1980s, the storage domain has produced amazing diversity. Storage alternatives for fundamental IT problem solving now abound, driven by combinations and permutations of hardware, software, and networking protocols. Virtualization has long existed as a way to make growth-oriented, diverse, and complex IT environments manageable. Since growth, diversity, and complexity also abound within the storage domain, storage virtualization enters as the fundamental consolidator, manager and simplifier.

IT professionals now have a spectrum of choices with regard to "virtualizing" storage. At one end of the spectrum are emerging distributed or grid-based models. At the other end are units based on centralized, massively scalable, virtualized storage controllers. In between are the network-based virtualization approaches. All have been positioned as platforms for the delivery of storage-based services to heterogeneous operational and application environments.

There can be no question as to where to place Hitachi on this spectrum. Hitachi champions the highly consolidated, massively scalable, and operationally centralized approach. Hitachi's latest embodiment of this approach is the new Universal Storage Platform V (USP V), which follows and builds upon the successful TagmaStore Universal Storage Platform (USP) introduced and delivered in September 2004. This is the platform for large-scale, heterogeneous storage consolidation and virtualization.

We have observed a general desire among IT administrators to at least cap, if not actively reduce, that number of storage and storage-related vendors housed

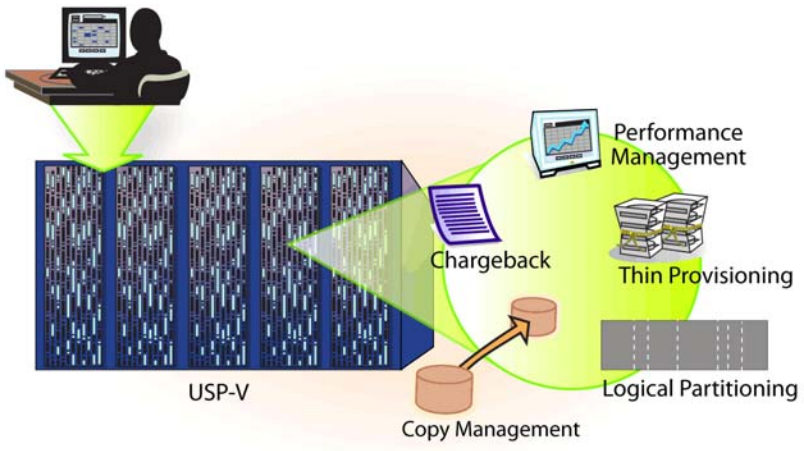

beneath one operational IT roof. For shops that adopt Hitachi as a preferred vendor, the USP V provides a platform for reducing the number of supported storage vendors—a move that also increases operational efficiency.

We also note that open systems computing and distributed computing architectures have finally run into a hard reality: that as distributed implementations grow, so too do the required energy, floor space and cooling requirements. Enterprise data centers are now running out of available electricity and cooling resources. While the most immediately realizable gains in energy and cooling efficiency are coming from server consolidation and virtualization projects, storage nevertheless should play a role as well, by doing a number of things in concert with one another. The Hitachi USP V can be seen as a platform for orchestrating storage energy efficiency as well as operational consolidation.

In this note we examine the USP V and its applicability to operational consolidation leading to increased operational efficiency, a key requirement among IT and storage administrators. With the USP V, Hitachi is now uniquely positioned to consolidate virtualized storage management operations across IBM z/OS mainframe and open systems using a common, virtualized platform.

Enabling Operational Consolidation

Much has been written about the benefits of storage consolidation. But that's often in terms of attributes such as capacity. Much less has been said about the benefits of consolidating storage-related operational functions. Yet, achieving operational efficiency is as important, if not *more* important. Reducing operational complexity yields increased management effectiveness and predictability, and reduces human-induced errors. Indeed, we believe that consolidating storage capacity without consolidating the supporting management functions can make matters worse. The USP V is interesting in this regard, as it is clearly as much about increasing operational efficiency as about scalable, consolidated, virtualized storage services.

The first USP iteration, Hitachi's TagmaStore USP, delivered a basic foundation for operational consolidation. A single USP encompasses IBM z/OS, Unix, Windows, and Linux application host environments. The USP presents SAN, NAS and DAS to these environments over ESCON, FICON, FC, iSCSI, and IP networks. Underlying storage can consist of Hitachi arrays and/or arrays from other vendors (such as EMC's Symmetrix DMX and IBM "Shark"). All of these storage environments can be managed from a single platform via the Hitachi Storage Management Software Suite.

However, the most significant aspect of the USP from the standpoint of operational consolidation is logical partitioning. IT administrators can carve the USP up into Private Virtual Storage Machines (a maximum of 32) that are managed as autonomous control units, each with its own assignment of virtualized storage capacity, cache, and I/O ports/channels. Each can be dynamically configured and reconfigured to exhibit desired performance and availability characteristics. Each can be tuned to a specific application environment or specific quality of service (QoS) levels. Each can be assigned a separate Virtual Serial Number for asset tracking and chargeback purposes. All are managed through a Virtual Partition Manager. Additionally, a My Domain feature within Hitachi's Storage Management Software Suite supports independent assignment and management of virtual storage pools by multiple qualified IT administrators. Each My Domain administrator is given a separate user interface, and can apply customized management routines to individual partitions.

With the USP, Hitachi also delivered a consolidated platform for managing heterogeneous data copy and migration functions called Universal Replicator. Universal Replicator copies and/or migrates data, both locally and over MAN/WAN links, between heterogeneous storage arrays, and between tiers in a tiered storage architecture managed by the USP.

Consolidation Moves Forward

The latest USP V builds upon all of the basic USP functions that enable operational consolidation.

USP V capabilities that contribute to increased operational efficiency include:

Consolidated capacity provisioning – Hitachi Dynamic Provisioning (HDP) is a unique application of the thin provisioning[1] concept. HDP allows administrators to provision capacity to applications based on a projection of both present and future needs of the application, but without actually having the physical storage readily available. It is common for enterprise storage administrators to allocate more physical disk to an application than actually needed, in order to avoid unplanned outages that can occur if a volume unexpectedly runs out of physical disk capacity. The result is that, on average, only 35-40% of physical disk capacity is actually used. It is also common to see storage administrators lose track of the allocated but unused capacity over time, resulting in stranded disk capacity. Allocated-but-unused and stranded disk capacity negatively impacts both operational expense and capital expense budgets in two ways:

1. It leads IT administrators to think that they have less capacity than they actually have, causing them to buy or lease more high-performance, premium disk over time than they really need, and;
2. It contributes unnecessarily to rising energy and floor-space costs.

Hitachi is also the first company to combine heterogeneous storage virtualization with thin provisioning software on an enterprise-scale. HDP allows storage administrators to continue the practice of allocating more capacity than applications need, but the capacity is carved out of pooled virtual storage managed by Hitachi's Universal Volume Manager, rather than actual physical capacity. As the physical (internal disk) capacity is consumed by an application, the storage administrator is alerted to add more when a predefined threshold is reached, an operation that can be performed dynamically, without disruption to normal operations. The USP V supports the creation of from one to 32 HDP pools composed of up to 256 Array Groups, with each pool containing a common RAID and disk type. Additional Array Groups can be dynamically added to any pool.

HDP's approach to thin provisioning is unique among other presently available thin provisioning implementations because:

It runs against a potentially huge capacity pool (up to 247 PB—yes, *petabytes*) that can consist of multiple disk types and RAID levels. While Hitachi is making thin provisioning available on internal storage during its initial release in May, 2007, we expect support to be extended to external storage a few months later. While applying HDP to external storage is technically possible now—as any controller functions in the USP V can essentially be mapped to external disk—we believe Hitachi's timing is directly tied to the company's comprehensive test and qualification process.

1. It supports multiple applications and multiple operating environments simultaneously.
2. HDP can be applied to individual partitions under Virtual Partition Manager.
3. Hitachi copy functions (Shadow Image) can be extended to HDP volumes. Hitachi Universal Replicator, Copy-on-Write, and Tiered Storage Manager following later in 2007.
4. Underlying physical storage can be composed of internal Hitachi arrays as well as external arrays from other vendors shortly after initial USP V availability.

Hitachi Global Solution Services has announced a professional service offering to assist storage administrators with planning, implementation, and ongoing internal support of HDP.

Universal Volume Manager enhancements - Here, Hitachi underlines an ability to scale its virtualized storage controller well into the petabyte range. USP V supports 247 PB of raw capacity, a greater than 600% increase over the USP, allowing IT administrators to consolidate even more internal Hitachi and external heterogeneous storage under

[1] This storage provisioning is the practice of allocating more storage to an application than is actually, physically available. Physical storage is added incrementally, non-disruptively, on an as-needed basis, without configuration changes.

the Universal Volume Manager. In addition, UVM will support allocation of second tier storage (such as SATA arrays and older external arrays) to z/OS mainframes via direct FICON attachment, whereas only FC protocol was utilized in the past.

Two additional facilities - Large Logical Storage Pools and Wide Striping are new with the USP V announcement. Large Logical Storage Pools allows storage administrators to create storage pools containing hundreds of drives and up to 1,024 volumes. Pools can also contain replicated volumes. Physical pool capacity can be expanded without configuration changes as data is striped across all drives in a pool (Wide Striping). Preformatted disks are simply added to a pool. Logical volumes can also be expanded without reformatting. The result is that file systems and databases can be expanded without physical configuration changes. In addition, performance tuning can be minimized when data is striped across all drives in a pool that contains a relatively large number of disks.

Storage-based Information Security Services – As the incidence of security breaches and identity theft are rising, Storage administrators are increasingly tasked to address information security issues. With the USP V, Hitachi integrates a number of storage-based information security services under one "roof," simplifying security operations:

1. End-to-end isolation of the I/O chain from HBA or mainframe processor channel to storage devices;
2. Controller-based data shredding
3. Write Once Read Many (WORM) software
4. Role-based administrative access;
5. Audit log for forensic activities;
6. Support for Fibre Channel Secure Protocol Authorization A;
7. Support for NeoScale and Decru encryption appliances;
8. Safe multi-tenancy with Virtual Partition Manager software

4 Gb/s Fibre Channel Switched (FSW) Backplane – Hitachi has added a unique switching capability to its back-end directors. The effects of this enhancement include:

1. Performance is improved for individual loops composed of 30 or more drives.
2. Troubleshooting required to identify a failed component is reduced.
3. Switching is inherently faster than utilizing a loop for arbitration
4. Hitachi is now "switched" all the way through the USP, from the internal architecture to the backplane connecting to the disks

Hitachi has also made the USP V easier to service and configure by reducing the physical size of board-level components

Storage Efficiency: It All Adds Up

IT has become focused on achieving energy savings through server consolidation and virtualization. No doubt, therein lies the biggest and quickest "bang for the buck." However, we believe that important results can be achieved with storage consolidation and virtualization, but via a less obvious and more incremental approach. In essence, the factors that contribute to storage-related energy efficiency are many, and often interdependent. For example, a storage tiering strategy that includes tape can lower energy costs by pushing archival storage offline, as compared to using disk as the archival media. Adding de-duplication technology to that archival strategy can also reduce both capacity and energy requirements, but de-dupe can only be applied to disk. Resulting efficiencies can be additive, but the magnitude of the beneficial result depends on how the two approaches work with each other and the applications (backup and/or archive) they support.

Therefore achieving storage energy efficiency will require a consolidated operational strategy. Factors that contribute incrementally include:

1. Thin provisioning, improving disk utilization from between 30-40% in most opens systems environments to 60% or greater

2. Load balancing to reduce or eliminate hot spots within the array enclosure, thereby reducing heat load on the ambient cooling system.
3. Data compression – compression algorithms reduce (by a factor of 2:1 or 3:1) the amount of physical resources required to store a given data object (file, volume, or dataset).
4. Data de-duplication techniques, to further reduce (up to 20:1) physical disk resources required for a given volume of data.
5. Tiered Storage that integrates slower-spinning secondary disk (for example, SATA) and tape to reduce the load placed on electrical and cooling systems by high performance disk.
6. Active archive software that moves data off high-performance disk to secondary and tertiary storage tiers.
7. Spin-down of inactive disks – Array power management can also include the ability to essentially "turn off" drives that have not been active for some pre-determined time period and then turned back on as the data they contain is needed. This process can yield substantial increases in power density (kW required for a given volume of data).

Presently, the USP V directly supports thin provisioning via HDP, tiered storage, and load balancing. VTL and archiving solutions offered by HDS and integrated with the USP V address requirements for de-duplication and active archiving. Compression and disk spin-down technologies are currently under review.

A consolidated operational platform that leverages storage virtualization can help move an energy efficiency strategy forward. Virtualization alone can offer many energy-saving services including thin provisioning and tiered storage, and data migration to devices that cost less to operate on a cost per GB basis. One must be able to analyze and manage all of these presently available tools to achieve maximum storage-related energy efficiency. It's a complex calculus that requires a consolidated and unified approach to virtualized storage; the Hitachi USP V enables that approach.

Conclusions

We are not philosophically bound to any of the storage virtualization architectures now available, including the distributed grid, network-based, and centralized controller-based incarnations. Rather, we believe that IT professionals should choose the approach or approaches that best fit their operational model. While many commonalities among IT, particularly enterprise environments, can be drawn, in the end all IT shops are unique entities. Therefore, IT professionals must deeply understand these architectural models to get the best fit of what will happily "live" within a given production environment, both now and into the future. Current USP customers will also have to query Hitachi as to how existing units can be replaced in order to get USP V benefits.

Storage virtualization is the key technology that enables a large scale consolidation of heterogeneous storage environments. In terms of scalability to the petabyte range combined with operational consolidation and depth of storage-based services deliverable to a wide range of applications and operating environments, Hitachi's USP V has no present equal.

The USP V is not "plug and play" storage. It embodies advanced technologies and delivers scalable services applicable to large enterprise IT. This isn't to say that the USP V is commensurately complex and cumbersome. Usability features have been designed into the user interface in order to reduce human error and speed administrative tasks. Nevertheless, using advanced features requires a high degree of operational understanding and sophistication on the part of IT administrators.

The Hitachi USP V can be seen as a self-contained, pre-integrated, storage services platform that scales to petabyte levels while maintaining a manageable, predictable operating environment over the lifecycle of the product. In many ways, it puts the once vaunted "storage utility" models within reach of medium-large scale users who have neither resources available nor the inclination to build their own storage utility from piece parts.