Strategic Options for Virtualized Storage Services

An Analysis of Hitachi's New Universal Storage Platform VM

Executive Summary

Over the past decade, storage network infrastructures have dramatically grown through merger and acquisition activity as well as a series of new initiatives that often favored completing a project over developing a storage infrastructure strategy. As a result, a sea of incompatible storage hardware and software has emerged in which servers and applications cannot gain access to desired data. Seemingly overnight, operating expenses have begun trending upward and storage islands are now constricting the flexibility of organizations that want to be more responsive to change.

In addition to a proliferation of largely incompatible storage area network (SAN) infrastructures, there are two related factors gripping organizations today:

- Expensive tier-one hardware, software and maintenance have forced customers to seek alternative strategies to accommodate hyper growth of storage cost-effectively.

- Costs related to moving data and making changes to storage infrastructure are substantial. Migration and provisioning complexities often require shutting down applications to move data or provision new storage capacity, often costing in excess of $50,000 for each migrated array.

The latter fact is noteworthy, as case studies show that, on average, over the life of an array, for every $1 spent on storage hardware and software another $.50 is expended on provisioning and migration costs. This additional outlay is being reduced dramatically to well under $.10 for each $1 spent by implementing virtualization, thin provisioning and automated tiered storage management.

This white paper examines solutions to these challenges. In addition to reviewing and comparing the various approaches, a main objective of this paper is to provide guidance for readers as to how best practices are shaping the way customers approach automated tiered storage and storage virtualization. The paper will examine the newly announced Hitachi Universal Storage Platform VM and provide guidance as to where the product fits in the market and how storage administration can best utilize its capabilities.

Introduction

Infrastructure assets in data centers around the world are being virtualized at a rapid rate. While mainframes were virtualized years ago, all open systems platforms—UNIX, Windows and Linux—are now following suit. ITCentrix has heard overwhelmingly from its clients that data centers are implementing or actively planning storage and server virtualization. A 2007 survey of nearly 300 IT professionals by InformationWeek indicates that more than 75% of installations are either implementing virtualization or planning to implement virtualization within 12 months.

The main business driver for virtualization is to improve IT responsiveness. The ability to react to change, to apply more computing resources to the applications and initiatives that need it and take it from those that don't, is becoming a fundamental competitive prerequisite for companies. This change requirement is now measured in hours, as opposed to the days or months that IT departments could traditionally deliver.

The historical challenge for data centers is that *physical* hardware of systems (servers, storage, networking, memory and systems software) had to be connected together, tested and brought online. It takes significant time to assemble, connect, test and bring into production these systems. Even though resources were available on equipment, it was too difficult to move those resources within the data center. The only way data centers could react was to add more equipment, even though there were resources aplenty within the data center.

The key technology for rapid provisioning (and, with it, better use of resources) is virtualization, which separates the logical systems from the physical resources. Virtualization means standard (virtual) server configurations that can be enabled in hours. It is the foundation that reduces the time required to commission and decommission storage from months to hours, and enables the implementation of storage services such as thin provisioning and tiered storage. Virtualization is the key enabler for a computing-on-demand model.

Virtualization for file-based systems (e.g., CIFS or NFS protocols against a NAS subsystem) is much simpler to achieve, as the protocol is already at a high level. ITCentrix has written about the technologies that have been introduced in this space, such as the Hitachi High-performance NAS system powered by BlueArc™. The challenge for most installations is the mission-critical systems running block-based I/O, and this white paper will focus on the strategic choices for implementing storage virtualization for block-based storage volumes.

Where Hitachi's USP VM Fits

The Hitachi USP VM comes in two main flavors: USP VM Diskless, in which it acts as a "stand-alone virtual controller," and USP VM, including internal disk.

As a stand-alone virtual controller, it provides full virtualization services to other storage arrays in the data center. It does this within a very small 10U rack-mounted form factor. This provides the basis for much faster commissioning and decommissioning of arrays and the foundation for tiered storage and thin provisioning. It is upgradeable to a full-function USP VM with internal disk.

As a full-function storage controller with internal disk, the USP VM provides virtualization services for both internal and external storage. It also provides services-oriented storage solutions (SOSSs), such as host credential authorization protocol (HCAP), for virtualization and volume migration, and virtual tape library (VTL).

Both flavors of the USP VM provide:

- the same storage management software suite as the rest of the high-end USP V Series;
- simplified software pricing that for most configurations uses frame-based pricing rather than capacity pricing;
- support for VMware ESX (Hitachi is the first storage vendor to obtain external storage virtualization certification for VMware Infrastructure 3); and
- the ability to operate on a standard 220V power supply, which reduces the power consumption required.

Overall, the USP VM is the industry best-of-breed in its performance class for environmentals (space, power and cooling).

Storage Virtualization: The Key to Flexibility

Storage virtualization separates the physical resources of specific LUNS from the logical resources the application sees, and maintains a map connecting those resources. Because the application continues to see its environment unchanged, it enables the seamless moving of data. For example, swapping out storage arrays at the end of a lease for a new storage array traditionally requires the interruption of applications to move the data. Scheduling those interruptions and migrations traditionally means that bringing on a new array can take months of meticulous planning, testing and implementation. Virtualization is a key enabler that allows this to be completed instead in days or hours.

For example, one ITCentrix client traditionally took about five months, on average, to provision new arrays. For this organization, lack of flexibility represented 43% in additional costs over the cost of the array, made up of the following components:

- Loss on buying earlier than required ~13% cost of array
- Loss on additional lease payments ~13% cost of array
- Cost of array migration (in-house or outsourced) ~17% cost of array
- Total cost of commissioning/decommissioning storage ~43% cost of array

Virtualization is planned to bring down the time from several months to five days and the cost of commissioning/decommissioning storage from 43% to less than 5%.

In addition to the advantages above, virtualization of storage enables other storage services. Two important ones are:

- Tiered storage: This is the ability to create storage pools with different performance characteristics, storage management capabilities and cost structures, and to move data seamlessly between these pools as the requirements of applications dictate. This allows significant savings in storage costs.
- Thin provisioning: The traditional method of provisioning storage is to physically allocate the necessary planned space; removing excess storage was deemed impractical. Thin provisioning allows the application to see a virtual pool of storage, but the physical storage is allocated only when required. This results in significant savings, as storage is allocated on demand from a storage pool.

There are three major strategic choices to be made when looking at storage virtualization. These are as follows:

1. Storage virtualization via an external appliance, such as IBM's SAN Volume Controller (SVC) or EMC's Invista. Such products provide the benefits of virtualizing heterogeneous storage arrays.
2. Virtualized arrays that provide storage virtualization (and sometimes thin provisioning) within the array itself, such as 3PAR® InServ arrays**.** These products offer good in-box tiering of homogenous capacity.
3. A combined tiered-storage approach through controller-based virtualization, which offers internal tier-one storage and external virtualization of heterogeneous arrays. The advantage here is that storage tiers can be managed with the same set of software tools without making performance concessions.

The Hitachi USP VM is the only enterprise-class example of a solution that can provide both appliance-like heterogeneous connectivity (via a diskless version) and tier-one storage support via internal disk.

Storage virtualization via an external appliance

Storage virtualization within an appliance provides some significant advantages in smaller environments that are less performance-critical. These advantages are:

- the ability to connect to and manage a heterogeneous set of storage arrays;
- the low initial cost of implementation;
- the implementation of a tiered-storage strategy with seamless migration of storage between arrays; and
- the lower cost of using storage management software in the appliance rather than in the storage arrays.

The connection of the appliance can be either a simple in-band solution (e.g., IBM SVC) or a more expensive but higher performing out-of-band solution (e.g., EMC Invista). Hitachi has introduced the best of both worlds via the USP VM Diskless, which allows connection either through shared ports linked to the SAN or by direct attachment of external storage to the USP VM.

The disadvantages of the traditional appliance approach are as follows:

- No thin provisioning to date. Traditional appliances such as the IBM SVC and the EMC Invista do not offer thin provisioning for either homogeneous or heterogeneous storage arrays.
- Performance latency overhead. The appliance introduces significant additional latency for very high I/O-intensive applications. These arrays have to be put in a separate storage pool without being virtualized.
- Virtualization pool size limited by performance of the appliance. If a larger storage virtualization pool is required, a second appliance completely separate from the first has be installed, with a separate set of software licenses.
- Additional stress and management complexity introduced to the SAN network.

Hitachi has addressed all these disadvantages in the newly announced USP VM. Diskless, it supports heterogeneous storage arrays and provides thin provisioning (called “Dynamic Provisioning” by Hitachi). In addition, this version can be upgraded to the USP VM with internal storage for higher performance tier-one applications.

Storage virtualization within the storage array itself

Some suppliers of storage arrays have provided virtualization within the array. This provides several significant advantages:

- Tiered storage can be implemented within the storage array.
- Thin provisioning is sometimes provided within the storage array.
- High-performance tier-one applications can be managed up to the capability of the storage array controllers with a common set of software products.

The disadvantages of this approach are also significant:

- The size of the storage pool is dictated by the largest array that can be implemented.

- Virtualization cannot help migrate storage to or from other arrays. As shown in the example above and in the case study, this is a significant cost disadvantage.
- Thin provisioning, if supported, is restricted to the storage array only (homogeneous thin provisioning) and cannot be extended to other storage arrays.

Storage virtualization in a controller (combined approach)

Hitachi is the only storage vendor that has introduced this type of approach, in 2004. The UPS VM is the latest model in the series. The storage controller allows direct attachment of external heterogeneous storage arrays, which can then be virtualized and managed as if they were internal storage.

The advantages of this approach are as follows:

- Tiered storage can be extended to both internal and external storage, using the same management software.
- Thin provisioning can be extended to both internal and external storage, using the same constructs (e.g., virtual storage pools) and the same management software.
- Performance latency issues can be overcome by migrating external storage to internal storage.
- Implementation of a tiered-storage strategy with seamless migration of storage between arrays is possible.

The disadvantage of this approach is a higher capital expense of controller-based solutions for very small configurations.

The case study below investigates the application of these strategies.

Case Study

The management of this installation was satisfied by the performance and availability of 200 terabytes of tier-one platform storage arrays. Some of the applications supported by the tier-one storage were very I/O-intensive. A significant amount of the storage had been commissioned within the last 12 months.

There were also approximately 100 terabytes of tier-two modular storage (mainly Windows-based applications) showing rapid growth, and increased demand for storage to support unstructured content was creating constraints on the business. Overall there were 550 open system servers.

The management was concerned about the length of time required to provision new servers and storage and very concerned about the cost and time required to migrate to new arrays. Most of the storage arrays were leased, so every year certain equipment needed to be migrated. It typically took between five and six months to complete the commissioning and decommissioning of an array.

The company decided to adopt an aggressive virtualization strategy for both servers and storage to improve the flexibility and cost effectiveness of its infrastructure. The server virtualization strategy was based on proprietary hardware and OS solutions for its IBM and Sun servers, and mainly VMware for its Windows and Linux Intel servers.

The storage management team was tasked with recommending an approach to virtualize storage after consider the three strategic approaches outlined above: 1) appliance-based; 2) array-based; and 3) in-controller combination.

Appliance-based: The team was very attracted to the appliance route. There were a large number of successful installations, and the potential cost of the storage management software was lower. The team liked the incremental way that storage virtualization could be implemented and the low-risk step-at-a-time approach.

The major concern with the appliance approach was in regards to performance, particularly of the I/O-intensive applications on tier-one storage. The team also noticed that there were a number of tier-two applications that had significant and critical I/O activity at certain times of the month. Team members were also concerned about the lack of thin provisioning support, which strategically they were hoping to embrace.

Array-based: The team looked at storage virtualization within the storage array itself and quickly rejected that option. A large amount of tier-one storage had just been installed and the team would have to either take a significant loss by getting rid of the storage or wait a long time before the leases were up for renewal and new storage could be brought in and virtualized.

In-controller combination: The team was introduced to Hitachi's USP VM Diskless product and found that it was a very good fit for the strategic requirements. The advantages of the approach were as follows:

- The team was not forced to buy additional storage as part of the virtualization strategy.
- The team could take a step-by-step approach to storage virtualization and virtualize the most important storage arrays first.

- The team could try the I/O-intensive applications first, behind the USP VM configuration with direct attachment, and see if the performance was "good enough."
- If the performance of any particular application was constrained, the team could upgrade the USP VM and add internal storage.
- The team was able to include thin provisioning as a part of the project, which it believed would release at least 10% of additional storage within three years.
- The USP VM fully supported VMware's ESX offerings and would assist an integrated approach to server and storage virtualization.
- The initial cost of the USP VM and software was competitive with the other appliance offerings.

The only disadvantage of this approach was that the USP VM was a new product. This concern was lessened by the fact that Hitachi had delivered over 6,000 virtualization-capable controllers over the last three years.

Virtualization Strategy	Advantages	Drawbacks	Overall Assessment
Appliance	-Good references -Good services -Low initial cost	-Doubts about availability and performance of tier-one behind the appliance -Problems scaling the appliance with future growth -No thin provisioning	***
Virtualized Storage Array	-Good references -Good thin provisioning	-Cost of removing existing tier-one storage -One type of storage array for tiered storage -No heterogeneous tiered storage or thin provisioning	*
Hitachi USP VM Diskless	-Thin provisioning of internal and external storage resources -More options for tier-one I/O-intensive storage -Low initial cost -Strong USP enterprise references -Performance -VMware certification	-New product	*****

Table 1 – Summary of strategic positioning study of different approaches

Recommendations and Conclusions

Hitachi has had a clear leadership role in providing virtualization storage services for high-end users with its USP and USP V series. The availability of the USP VM and the USP VM Diskless fills out its virtualization offerings to meet the needs of all parts of the market and all virtualization approaches. At the low end, Hitachi provides a stand-alone virtual controller approach with the superior functionality of thin provisioning and superior growth options for performance-critical tier-one storage. For modular storage, it provides the same virtualization, thin provisioning and tiered-storage functionality as the high end, but at a much lower entry point and with a much smaller footprint.

Customers facing sprawling SAN infrastructure, high migration costs, expensive tier-one storage and maintenance costs, and a desire to reduce overall storage TCO should act as follows:

- Gain agreement with lines of business, establish clear guidelines for tier-one storage and communicate these to the organization.
- Allow storage administration to apply these policies, default all other storage to tier-two and archive accordingly based on retention policies.
- Virtualize front-end server and back-end storage resources.
- Strongly consider a combined approach such as the USP VM; start diskless, virtualize storage assets (including tier-one) and, over time, add internal capacity for more demanding tier-one applications.

By introducing the USP VM Diskless as a direct competitor to virtualization appliances, Hitachi offers a broad variety of configurations for and approaches to heterogeneous storage virtualization. Behind this flexibility is a consistent set of architectures and storage management products. ITCentrix believes that the USP V and USP VM offer the most advanced and flexible storage platform in the industry by far, with heterogeneous virtualization, cross-array heterogeneous tiered storage and heterogeneous tiered storage. No competitor comes close.

ITCentrix believes that there is significant synergy in addressing virtualization of both storage and servers within a data center. Even though Hitachi has to expand the scope and size of its data center services group, ITCentrix would strongly recommend Hitachi as the storage partner of choice.

ITCentrix
4 Mount Royal
Marlborough, MA 01752
Tel: 774-463-3400
www.itcentrix.com

About ITCentrix

ITCentrix is a consultancy that primarily serves the needs of CIOs and technology professionals. Its main emphasis is on using tools and analytic modeling techniques to advise clients on increasing company performance through improved resource allocation and better infrastructure management.

The company's products and services have been used at several hundred organizations in North America, Europe, the Asia-Pacific region and emerging countries to focus investments on returning optimal business value.

In 2007, the company launched the Wikibon project (http://www.wikibon.org), a community of practitioners, consultants and analysts dedicated to improving technology adoption through the open sharing of business and advisory knowledge.