WHITE PAPER

Virtual Storage Delivers a Three-Dimensional Approach to Storage Scaling

Sponsored by: Hitachi Data Systems

Benjamin Woo

October 2010

IDC OPINION

The new Hitachi Virtual Storage Platform (VSP) updates the Hitachi Data Systems (HDS) flagship platform. The "3D scaling" enables users to independently scale performance, scale capacity, and scale wide, integrating legacy assets and allowing the VSP to ultimately scale to multiple data centers.

IDC believes that the new VSP has taken into consideration not only the future needs of users but also how to evolve users from their current infrastructures, by protecting their existing investments, while allowing users to leverage the VSP's capacity and power-efficient features.

The next-generation data center of the future will need a highly virtualized and dynamic storage infrastructure. It needs to address the block, file, and object storage requirements of today and tomorrow. IDC believes that the Hitachi VSP will help customers do this.

SITUATION OVERVIEW

The New Business Initiative Is Efficiency

In the current challenging economic conditions, businesses large and small are looking to boost efficiency in IT. To achieve the desired efficiencies, businesses need to:

- Consolidate their capital investments
- Improve utilization of IT resources
- Protect investments by reusing existing resources
- Seamlessly migrate from one platform to another

A more efficient data center will enable IT departments to be more responsive to business unit needs through:

- Dynamic provisioning of IT resources
- Storing and protecting the digital assets of the business

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

Finally, businesses are also seeking to change the capitalization equation for IT spending, preferring where possible to operationalize their investments with the desired result of being able to deliver standardized predictable and repeatable results that are more closely tied to business trends.

Achieving IT Efficiency

IDC forecasts that storage capacities shipped will continue to grow at a compound annual growth rate (CAGR) of nearly 50% between 2009 and 2014. Much of this growth is a result of the increasingly digitized and regulated business environment in which we operate.

The combination of explosive data growth and the desire and need to achieve much higher efficiencies in IT overall is driving much of the evolution from traditional data centers. Currently, data centers tend to be silos of disparate IT resources (such as servers, applications, storage, networking, etc.). Conversely, the next-generation data center will be distinguished by highly virtualized resources that can be dynamically provisioned and reprovisioned on an as-needed basis.

Additionally, the increasing pressure created by the growing number of client devices — including desktops, laptops, netbooks, smartphones, and tablet clients — that generate data is driving the need for data mobility to new heights.

Despite regulatory and legislative requirements, storing, protecting, and leveraging data must be made more efficient. Otherwise, the increasing cost of supporting the data storage element of the data center will become cost prohibitive, which in turn can result in missing internal and external service levels.

Virtualized server and desktop requirements are also creating new challenges for storage platforms. Traditional data centers have a higher probability of a 1:1 relationship between servers and storage. In highly virtualized, dynamic data centers, a single server will host multiple virtual machines and thus the I/Os of many mixed workloads. Higher-performance interconnects will be required to address this.

Recognizing the increasing needs to provide quality of service (QoS) for individual virtual machines (or desktops), hypervisor vendors such as VMware have released APIs that allow storage systems vendors to better integrate with the hypervisor and guests to enable better server and data mobility and migration.

Finally, the next-generation data center cannot be a rip-and-replace strategy. Instead, organizations need to evolve into the next-generation data center.

However, simply optimizing the infrastructure in the data center is insufficient. IT organizations need to evolve and align themselves with business units as opposed to being IT focused.

The Changing IT Organization

IDC believes the opportunity is imminent to review the functions of internal IT organizations to better align the IT organization with the next-generation data center. IDC suggests that the IT organization of the future will be organized into three groups:

- **Information facilities.** This team will be the most similar to the current IT organization. Members of this team will be tasked with the operational aspects of the data center — including servers, storage, and networks. The team will also be responsible for optimizing efficiency and costs (both capital and operational).

- **Information services.** This team will be responsible for designing and enhancing business applications and decision support tools. Members of this team will be leaders in transforming the data center from being IT centric to business centric.

- **Information access.** This will be the part of the IT organization that ensures that the appropriate people can securely access information services anytime, anywhere, and on any desired device.

FUTURE OUTLOOK

Moving to a Business-Centric Data Center

As businesses look at different options to become more business centric, rather than IT centric, they are presented with one of two solutions:

- Outsource to a managed services or software-as-a-service (SaaS) provider

- Improve IT operations to meet the changing way IT supports business

Both of these options are built on the next-generation highly virtualized and dynamically provisioned data center. While the consumption of either of these options differs, the underlying infrastructure design is not dissimilar.

They are both built on a highly virtualized compute, network, and storage platform. Virtualization achieves the objective of improved utilization of IT resources while simultaneously minimizing the capital outlay required to sustain or improve service levels to the business.

IDC predicts that while worldwide server shipments will continue to grow at a CAGR of 6.5% through 2013, the worldwide server installed base will remain relatively constant (with a CAGR of 0.6% through 2013). At the same time, the worldwide virtual server installed base will increase at a CAGR of 28.1% for the same period (see Figure 1).

FIGURE 1

Server Installed Base, 2003–2013

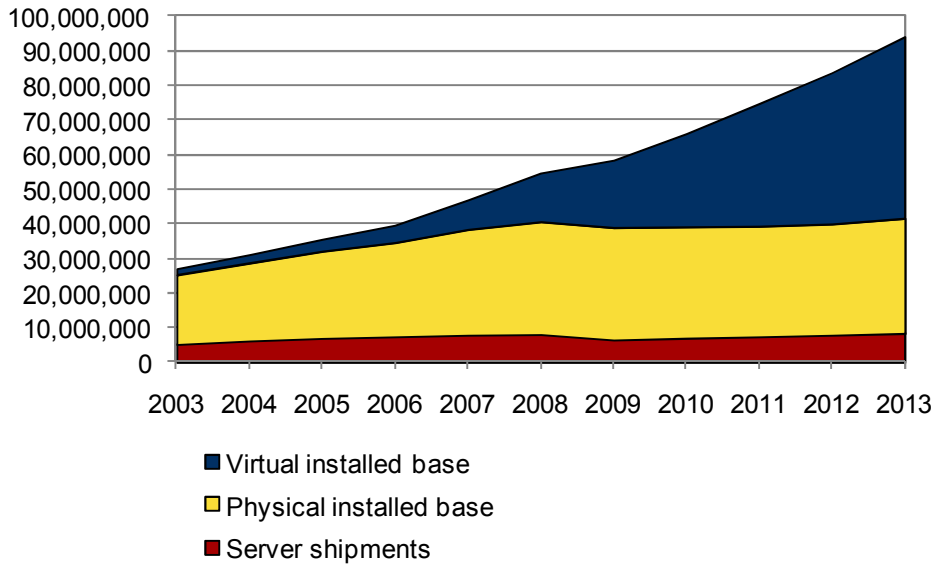

Source: IDC, 2010

Making Virtualized Compute, Network, and Storage Work in Harmony

IDC believes that this is a three-step process:

1. Expand the use of networked storage
2. Virtualize networked storage
3. Take advantage of converged IT infrastructures

The research that IDC has conducted through a survey of users suggests that roughly two-thirds of users are expected to increase the attach rate of storage to virtualized environments by at least 20% by the end of the first quarter of 2011 (see Figure 2).

FIGURE 2

Storage Attach Rates to Virtualized Environments

Q. How much do you expect the storage attached to virtualized servers to increase in the next 12 months?

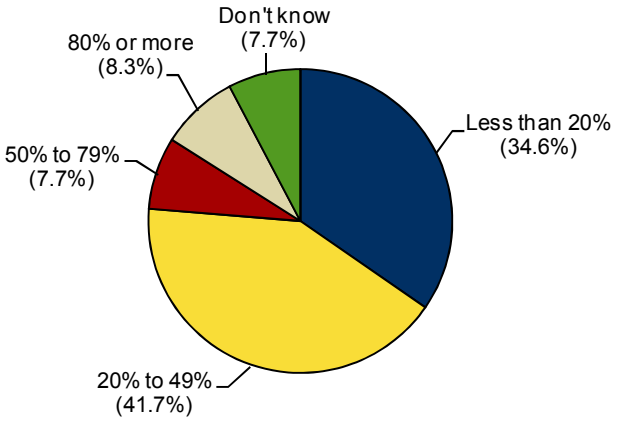

n = 156

Source: IDC's Storage and Virtualized Environments QuickPoll, 2010

With this in mind, expanding the use of network storage is critical. Virtualized environments are optimized when attached to centralized network storage solutions.

Networked storage solutions enable the efficient provision of storage capacity to virtual machines and desktops. They also enable virtual machine mobility (through replication) of both applications and data.

Virtualizing networked storage will also enable many of the functions that will create a more efficient storage architecture. Flexible allocation through the use of dynamic and thinly provisioned volumes will boost utilization rates; can often provide nondisruptive data migration in support of tiering; and will minimize the cost of future storage systems purchases.

Converged IT infrastructure and management can reduce the time needed to provision IT resources and minimize human error. Convergence also provides the ability to integrate security, business continuity, disaster recovery, and other business processes in a coordinated and automated fashion.

Managing Structured and Unstructured Data

IDC research shows that unstructured data that is stored in the context of a content depot or in the public cloud will grow at a CAGR of over 75% through 2014. Conversely, traditional structured data (such as database data) will see a CAGR of only 23.6% for the same period.

This means that by 2014, over 78% of the storage capacity in a datacenter will be either general unstructured data or data that is stored in a content depot or an archive.

In virtualized environments, almost across the board, over half of the users surveyed by IDC indicated that they either are using or plan to use the following advanced storage technologies to help deal with this data and data accessibility explosion by the end of the first quarter of 2011:

- Deduplication
- Solid state disks (SSDs)
- Replication
- Compression
- Encryption
- Snapshots
- Thin provisioning

This makes these technologies essential parts of the storage infrastructure as it relates to the next-generation data center.

The Need to Scale to the Third Dimension

For many years, storage solutions have had the ability to scale on two dimensions: improved and expanded performance and increasing capacity.

Scale Up

Traditionally, this has been achieved by migrating from one generation of product to the next. Newer generations typically offer better performance and capacity capabilities. IDC refers to this as scaling up. The architecture is the same, but the performance/capacity ratio improves as the components (disks and controllers) evolve.

With the need to protect and preserve data over longer periods of time, scaling up presents the challenge that it often involves the risk of data migration.

Scale Out

More recently, the industry has embraced the concept of scaling out. Scale-out technologies have the benefit of leveraging industry-standard components and the linking of these components in grids, thereby improving the cost efficiency of the solution. They also offer the benefits of being able to scale performance and capacity independently.

If an environment requires more performance, storage compute nodes could be introduced into the scale-out environment to provide this performance. However, capacity may not change. Conversely, if the performance is sufficient, but the data capacity needs grow, capacity can be added independent of the storage compute components of the architecture.

It is important to note that this does not replace the need for scaling up. As new generations of products are introduced by vendors, scale out can take advantage of the new generation of "scaled up" products.

Scale Wide: The Third Dimension

Scaling wide is the third dimension of the new storage equation. It enables the ability to integrate legacy assets and extend to multiple data centers.

This third dimension is important to support not only legacy assets but also legacy applications that may not be ready to take advantage of scale-out architectures or as a way to protect existing investments.

It also enables businesses to take a universal approach to the management, distribution, and allocation of storage assets across multiple data centers.

The Hitachi Virtual Storage Platform

The Hitachi VSP is the latest generation of storage systems from HDS.

This platform succeeds the Hitachi Universal Storage Platform (USP) with an emphasis on addressing the next-generation highly virtualized and dynamic data center with a focus on efficiency, data mobility, and scaling.

At the core of its design is what HDS calls the "3D architecture." It scales up, scales out, and, in HDS nomenclature, scales deep.

HDS believes that the new VSP will address the following business challenges:

- Application integration and business-level view of infrastructure services
- Page-level tiering for simplification and optimization of tier costs and performance
- Virtualization and dynamic provisioning that enhance existing storage assets, facilitate infrastructure changes, and dynamically provision storage
- Sustainability with reduced footprint, power, and cooling as well as increased efficiency
- Simplicity — simple to manage, simple to implement, simple to migrate

The VSP is managed through Hitachi Command Suite. It brings a singular management platform for managing not only the VSP but also instrumentation that will enable businesses to optimize their environments.

Addressing File-Based Access and Content Depots

The next-generation data center not only must be virtualized but also must be dynamic and flexible in its ability to provision. One of the ways to address this is to provide dynamism in the way data on the Hitachi VSP is accessed.

HDS has addressed this issue through its Hitachi NAS platform and its Hitachi Content Platform.

Hitachi NAS provides the ability to access the VSP through standard Windows CIFS and NFS interfaces. IDC research shows a fast acceleration in the use of NFS as a preferred interconnect for virtual environments.

Additionally, with the need to extract greater value from archived data in particular, storing the data in an object-based environment improves the ability to put content in the right context, thereby increasing its value.

Sustainability

In designing its new flagship platform, HDS has focused on many sustainability aspects of VSP. In particular, HDS addressed the following aspects in its hardware design:

- **Density.** The VSP supports twice the number of disk drives compared with the previous USP platform.
- **Media choices.** By leveraging SSDs, HDS has been able to sustain performance while simultaneously reducing the power footprint. With so many companies faced with the prospect of not being able to draw any more power than they already have, augmented with the rising cost of electricity, this will appeal to many users.
- **Dynamic and automated tiering.** Providing choice in storage media can actually increase the complexity of managing a storage platform. HDS is addressing this by introducing dynamic and automated tiering functions within the VSP and automating this often time-consuming and error-prone task.
- **Air flow.** HDS has taken care to reduce "hot" spots associated with the VSP to maximize air flow and cooling efficiency. Many corporations have already made public commitments to reduce their carbon footprint and emissions. Focusing on efficient power and cooling approaches, this will gain the favor of many prospective users of VSP.

CHALLENGES/OPPORTUNITIES

The VSP platform, with its ability to scale up, scale out, and scale deep (HDS' term for scaling wide), certainly addresses many of the requirements for the next-generation data center.

Hitachi's decision to leverage a specialized hardware design as opposed to simply leveraging industry-standard servers will allow it to bring optimization to all levels of the storage system but may make HDS less price competitive than its competitors that have chosen to leverage commodity off-the-shelf designs.

Nonetheless, the ability to integrate non-Hitachi, third-party storage systems will continue to be a differentiator for HDS and appeal to those customers that have a requirement to extend the useful life of their current storage assets.

HDS will also need to address quickly its partnership with application providers and the integration of Hitachi Command Suite with system management software.

Increasingly, users are looking to infrastructure vendors to provide the necessary hooks and, in some cases, client software that will provision not only the storage but also the provisioning of the virtual machine and the associated application deployment templates. HDS should focus on the leading software products, including (but not exclusively) Microsoft Exchange, SQL, and SharePoint as well as the Oracle database.

CONCLUSION

The foundation for the effective and efficient transformation of the traditional siloed data center to the next-generation highly virtualized dynamic data center requires not only the virtualization of the compute and server platform.

A storage infrastructure that can be as flexible, capacity efficient, and power efficient as the server infrastructure is mandatory.

With the VSP, HDS has delivered a contemporary architecture that addresses the needs of the new data center for the future. The VSP scales performance and capacity independently and also adds a third dimension of scaling wide. This third dimension enables HDS to help data centers evolve cost-effectively from their current state to the future desired state.