



The 3D Scaling Advantage

How 3D Scaling Drives Ultimate Storage Efficiency

By Hitachi Data Systems

September 2010

Table of Contents

Executive Summary	3
Introduction	4
The Rise of Virtual Machines	4
Forecast: Data and Server Growth May Outpace Virtualization Capabilities	5
Introducing the Only 3D Scaling Architecture	5
3D Management: Manage Up, Manage Out and Manage Deep	7
Dynamic Mobility of Data across Virtual Storage Tiers	8
3D Scaling Architecture Benefits	9
Summary	10

Executive Summary

With the creation of data increasing exponentially, the demand on corporate IT resources is growing dramatically. At the same time the pressure to reduce staff, be more environmentally friendly and reduce costs is also increasing. While these seem to create a "no win" situation, advancements in hardware design, virtualization and better tools for the IT staff are making it possible to address these demands.

Server virtualization has provided a foundation for significantly increasing the utilization of server resources while reducing the number of real physical servers. Blade servers have reduced the footprint of the servers, conserving floor space and cooling in the data center. This consolidation of virtual servers is driving an increasing need to consolidate and scale up storage systems. Virtual machines in a virtual server fire off I/Os independent of each other. But the only storage solutions for these random I/Os are very expensive flash drives or wide striping of the I/O load across many disk spindles. Server virtualization has seen major adoption because it can consolidate servers and reduce cost. Server virtualization is entering the next phase where its value is seen in adding business agility for mainline applications. This new phase places a heavier load on storage systems and requires tighter integration with storage virtualization.

Hitachi introduced storage virtualization over six years ago with the Hitachi Universal Storage Platform® to help reduce the cost of migrating and tiering storage and to extend the life of existing storage assets. In 2008, Hitachi introduced the Hitachi Universal Storage Platform V with Hitachi Dynamic Provisioning software, which has dramatically reduced the time to provision storage, eliminated the waste of allocated but unused space and increased disk performance with automatic wide striping. While this has been the leading storage virtualization in the industry, it is time to step up the game to meet the demands of server virtualization.

With the introduction of the Hitachi Virtual Storage Platform, Hitachi Data Systems is addressing the need for better performance and scalability. The Virtual Storage Platform employs an enhanced switch architecture and page-level tiering, and it meets reduced floor space and cooling requirements through innovative new technologies and packaging. It also enables tighter integration with server virtualization and improved storage administrator productivity through the Hitachi Command Suite of software. This all adds up to a reduction in the total cost of servers as well as storage and an ability to transform the data center.

The Virtual Storage Platform is able to scale in three dimensions. It can scale up to provide improved performance and throughput as its host servers consolidate. It can scale out to deliver capacity, performance and bandwidth from a shared pool of resources as more and more host servers are connected to it. And, it can scale deep by connecting external storage; it allows existing multivendor storage or lower cost modular storage to add their resources to the scale up and scale out demands of host servers. This extends the life of existing storage assets and increases the return on assets.

This platform is based on a new design that builds on years of Hitachi enterprise storage experience. Building on industry-leading Hitachi storage virtualization, the Virtual Storage Platform incorporates high performance computing technology to provide even better performance and reliability than its predecessor, the Universal Storage Platform V. While the Universal Storage Platform V was designed to be the best product for storage virtualization, the Virtual Storage Platform in conjunction with Hitachi Command Suite is designed to integrate with server virtualization and provide the best virtual storage for server virtualization with increased scalability, usability and ease of management.

Introduction

Data is at the core of the data center, and any effort to transform the data center must involve the movement, provisioning, access and protection of data, which is provided by storage systems. This transformation becomes increasingly difficult as more data is generated and more applications are dependent on access to that data storage. Storage could be the biggest inhibitor to data center transformation, unless it is addressed through the type of storage virtualization that can meet the requirements for transformation.

Storage virtualization should be as easy to implement as virtual servers in a hypervisor. Just as simple as it is to virtualize external storage, it should be as simple to de-virtualize, so there is no vendor lock-in. The storage virtualization controller should be able to enhance whatever storage it virtualizes. An appliance with limited connectivity, limited cache and limited processing power cannot enhance the storage it virtualizes, except to make some limited copies and moves between external storage systems.

The storage virtualization controller must provide the abilities to scale up, scale out and scale deep. It must:

- **Scale up** dynamically to meet the consolidation demands of increasingly powerful virtual server clusters
- **Scale out** dynamically across a pool of shared virtual storage resources instead of a cluster of standalone storage silos
- **Scale deep** to extend scale up and scale out capabilities to externally attached storage

Finally, storage virtualization must include transparency. This is the ability for applications to see into the virtual infrastructure to observe the health of the underlying physical components that support their virtual storage, to monitor their service level objectives and to track their usage trends. This requires an integrated set of software tools that can gather data from the infrastructure, correlate the data to an application or server, and present it through an easy-to-understand dashboard, with "drill down" capabilities and report generation.

The Rise of Virtual Machines

Virtual servers have solved the problem of server proliferation, which in the past saw power hungry servers configured for peak demand, but idling at 10% utilization for most of the day. It also improved business agility by enabling data centers to spin up new servers on demand, load balance workloads and do site recovery across a pool of server resources. Today, there are more virtual servers deployed than there are physical servers.

According to Gartner, Inc., hypervisor-based workloads will grow from 19% in 2009 to as much as 48% by 2012¹. Multiple virtual machines on a single physical server can address the problem of the physical server proliferation and the accompanying need for more floor space, power and cooling.

¹ "How to Implement High-Availability Storage for Server Virtualized Environments," Valdis Filks, Robert E. Passmore, © 2010 Gartner, Inc., and/or its Affiliates

However, these virtualized environments introduce new management problems related to performance tuning, storage management and availability. In a virtualized environment each physical server hosts multiple virtual machines. The physical server is oversubscribed in order to maximize its utilization, resulting in tens or even hundreds of virtual machines running on virtual servers. This consolidation of virtual machines may multiply the I/O workload by tens or even hundreds of times.

Forecast: Data and Server Growth May Outpace Virtualization Capabilities

The dramatic growth in data across all businesses, according to InfoPro², is exceeding 30% for both Fibre Channel SAN and network attached storage. This data growth as well as the introduction of new applications to process the data are driving significant growth in the demand for servers and storage. At the same time, businesses are looking to lower the total cost of ownership to help address budget constraints. This demand may quickly outpace the capability of virtualization technology. Additionally, most companies are being asked to meet this demand with smaller staffs and budgets. According to Gartner, Inc.³, only about 18% of the enterprise data center workloads that could be virtualized had been virtualized at the end of 2009; however, by end of 2012 it is expected to reach more than 50%. In a recent survey, IDC found that 75% of new server implementations default to virtualization or virtualization is strongly recommended.⁴ Today, in most companies, less than 50% of the servers are virtualized. This means they must add physical servers at the same time they are moving to a virtualized environment in order to leverage their existing assets and optimize the use of floor space and the associated electrical and cooling demands.

Introducing the Only 3D Scaling Architecture

With the Hitachi Virtual Storage Platform, Hitachi Data Systems is delivering the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, out and deep in a single storage system, the new Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. The Virtual Storage Platform provides virtual storage that is designed to meet the growing demands of server virtualization.

The trend in server virtualization with powerful multicore processors is to consolidate the I/O workload of many servers onto a single storage system. As more virtual machines are consolidated through a virtual server, the storage system must be able to dynamically add more storage resources to keep up with the I/O demand. Performance, capacity and connectivity are increased by adding cache, processors, connections and disks to the base system. A virtual server that accesses the storage system can utilize all these resources, which act as one system managed as a common pool of resources. The ability to scale up the storage will result in more virtual server consolidation, better utilization of resources and lower cost.

² "The InfoPro™ Storage Study," Industry Profile Report, Wave 13, Q4 2009

³ Gartner, Inc., press release, "Gartner Says 60 Percent of Virtualized Servers Will Be Less Secure than the Physical Servers They Replace through 2012," March 15, 2010


⁴ "IDC Server Virtualization 2009," IDC, September 2009

Scaling out is required when multiple servers, virtual servers or standalone servers are spun up to meet business demands. With the Virtual Storage Platform, additional storage resources can be added dynamically to grow the base system for greater capacity, performance and connectivity. Because it is a tightly coupled system, host servers can access the resources that they need out of the common pool of storage resources. The Virtual Storage Platform also has the ability to dynamically partition the storage resources between the servers for safe multitenancy. This prevents any data leakage or escalation of management privileges between servers while they share the same pool of resources. The result is the highest performance available in the highest density footprint, which consumes the least power for the capacity stored of any enterprise storage platform.

Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage. By dynamically virtualizing new and existing storage systems, those systems become part of the Virtual Storage Platform pool of storage resources, which can scale up and scale out. Once virtualized, external data can then be migrated, tiered, replicated and managed by the Virtual Storage Platform. In this manner, older data storage systems can gain a longer useful life. Distance replication for business continuity can be extended to lower cost, lower function storage systems through virtualization behind a Virtual Storage Platform.

The architecture that makes it possible is a switch matrix architecture, which connects the basic components, front end directors, back end directors, global cache modules and virtual storage directors. This enables redundant pairs of directors and cache modules to be added as required, without disruption to connected host servers (see Figure 1). All these resources are tightly coupled through a global cache, which creates a common pool of storage resources. These resources can include external storage, which is connected through front-end director initiator ports.

Figure 1. Fully Configured Hitachi Virtual Storage Platform



3D Management: Manage Up, Manage Out and Manage Deep

Hitachi Command Suite is the management software that provides 3D management for Hitachi Virtual Storage Platform and the earlier generation Hitachi Universal Storage Platform V. Hitachi Command Suite simplifies operations and lowers costs in three distinct dimensions. It works with these storage systems to:

- **Manage up** to scale with the largest infrastructure deployments: over five million objects
- **Manage out** with breadth to manage storage, servers and the IT infrastructure
- **Manage deep** with the integration of storage and server virtualization and application transparency for today's complex data centers

The Hitachi Command Suite provides application-centric management of the entire virtual data infrastructure composed of virtual machines, host servers, networks and storage (see Figure 2). It manages the infrastructure from a variety of perspectives. It can be administered not only from a storage view, but also from business application views across all tiers of data.

With Hitachi Command Suite, externally virtualized storage can be grouped into tiers based on availability, recoverability, performance and cost, as if it were internal storage. The Hitachi Command Suite actively tunes and balances the workloads within this common storage pool, and it lets administrators configure, manage and move data across the entire storage pool, including third party storage systems. Hitachi Command Suite offers a 3D management software framework to lower costs and efficiently manage all data types. With end-to-end reporting and simplified management of the virtual infrastructure, you can maximize use of IT assets and improve application data access, availability and performance. The Hitachi Command Suite capabilities help you properly align the management practices needed for your infrastructure today and plan appropriately for future growth.

Hitachi Command Suite with the Virtual Storage Platform provides support for VMware vStorage API for Array Integration (VAAI). VAAIs offload repetitive writes to the storage, the use of storage-based copies to support VMware copies, and the replacement of the software SCSI reserve lock, which locks the entire VMFS volume with an atomic test. They also offload set command, which limits the lock to the extents of a particular virtual machine. The net result of this integration is increased scalability for VMware servers. Hitachi Command Suite also provides a bottom up view of the storage infrastructure behind the VMware Virtual Machine File System (VMFS) and the VMware Disk (VMDK) or virtual disks that belong to virtual machines. It also provides the capability to diagnose and fix storage problems.

Figure 2. Hitachi Command Suite provides end-to-end visibility for traditional and virtual environments



Dynamic Mobility of Data across Virtual Storage Tiers


Most data is rarely accessed after it is created. As a result, it should not be stored on your most expensive tier of storage, but rather moved to a lower, less expensive storage tier. Managing the movement and placement of data is the promise and premise of data lifecycle management methodologies. However, defining where and for how long data should reside at any point in its lifecycle is complex and problematic. With the introduction of the Hitachi Virtual Storage Platform, Hitachi has introduced Hitachi Dynamic Tiering, a revolutionary new software solution. Dynamic Tiering eliminates the time consuming manual processes of data classification and movement to different storage tiers, optimizing tiered storage usage while actually improving performance.

Dynamic Tiering takes the automation of tiered storage to a new level. It enables the management of multiple storage tiers as a single entity. It presents a new kind of virtual volume with embedded smart tiering that monitors access and moves data at the fine grain 42MB page level. It breaks the volume into pages and automatically moves pages that are being referenced infrequently to lower cost tiers of storage, avoiding the time and storage space to move the entire dataset or file. It allows the right data to be in the right place at the right time by self-optimizing for high performance and space efficiency.

(Note: Data comes in at the highest tier then migrates down based on activity. If a page on a lower tier begins to see more activity it is promoted up.)

Dynamic Tiering is very good for database applications. It can improve performance for databases stored in a file system by keeping the highly referenced data, such as indices on tier 1 storage, while moving the less frequently referenced pages to a lower tier of storage. Dynamic Tiering is also excellent for file and content data, which requires high performance up front but becomes inactive over time. It is also ideal for Microsoft Exchange 2010, which provides very large mailboxes that are replicated for fast recovery. The mailboxes can be allocated to Dynamic Tiering pools where recent mail would reside on tier 1 and inactive mail would automatically migrate down to lower tiers. The self-adjusting system can optimize performance based on demand. It enables content service providers to manage their storage assets more effectively and allows the use of less expensive, higher capacity disk drives, such as SATA drives.

Figure 3. Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering Software



3D Scaling Architecture Benefits

With the advantages of 3D scaling, the Hitachi Virtual Storage Platform creates a more agile storage infrastructure, which can increase the productivity of IT operators, reduce storage cost and increase return on storage assets. Its unique virtualization scaling capabilities expand its management efficiencies to other multivendor storage and enable the creation of a more efficient and agile data center.

Storage virtualization de-references the application and server view of data from the physical storage infrastructure so the physical storage infrastructure can be changed and transformed without disruption to the application. The first thing it can do is transform your legacy storage infrastructure without the need to "rip and replace." Once Fibre Channel storage is attached to the storage virtualization controllers and applications are redirected to the virtual ports, applications will be able

to access their existing volumes through the high performance global cache of this virtualization engine. They will be able to leverage new capabilities like high speed distance replication for business continuity, dynamic tiering for lifecycle management, and wide striping and the latest high speed media for improved random performance.

Most virtualized storage systems will see an increase in performance just by sitting behind the large global cache, but if more performance is needed, you can wide stripe your volumes or move them onto the tier 1 storage and do both. If you still have several years of useful life in your existing storage but the warranty is about to expire, use that storage as tier 3 where the expense of tier 1 maintenance is not required and convert to time and materials. If it makes more sense to replace the older storage with greener, more cost-effective storage, you can do the migration without stopping the application.

If you are converting servers or applications as part of this transformation, you can create nondisruptive clones of the data for conversion, extraction, translation and loading, as well as development testing on lower cost tiers of storage. Or you can dynamically spin up new allocations of virtual storage to support virtual servers. Not only can storage virtualization protect your applications and servers from changes in the physical storage infrastructure, but it also can enable your applications and servers to change and grow dynamically.

Capabilities like storage virtualization, dynamic tiering and dynamic provisioning were also available in the previous generation Universal Storage Platform V. The difference in the Virtual Storage Platform is that its storage virtualization capabilities are designed to meet the increasing demands of server virtualization. The first requirement for support of server virtualization is the ability to scale, which has been greatly enhanced. The second is to provide integration with virtual server APIs to make the virtualization of servers more efficient. The third is to simplify the management and provide application and server visibility or transparency into the storage virtualization infrastructure for openness, communication and accountability.

The Hitachi Data Systems approach to data center transformation provides the integration of server and storage virtualization, which allows organizations to consolidate resources, technologies and applications, and reduce the complexity of years of "bolting on" physical servers and storage. If the requirements discussed above are met, then by design you are incorporating future flexibility. This also greatly helps the bottom line.

Summary

The Hitachi Virtual Storage Platform is the only 3D scaling storage platform designed for the integration of server and storage virtualization. It is the only storage architecture that flexibly adapts data storage infrastructures for performance, capacity and multivendor storage to optimize return on storage and server assets. 3D management increases automation and efficiency for storage, computing and virtual infrastructure. Data mobility reduces impact on the business when adapting to change. Improved operations increase data center efficiency and flexibility. The Hitachi Virtual Storage Platform combined with the unique Hitachi Command Suite management platform transforms the data center so IT is more agile.

 **Hitachi Data Systems Corporation**

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627 USA
www.hds.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

© Hitachi Data Systems Corporation 2010. All Rights Reserved. WP-382-A DG September 2010