



# Hitachi Dynamic Virtual Controller Technology

Optimize Dynamic Workloads

*By Mark Adams*

April 2012

## Table of Contents

|   |           |
|---|-----------|
| <b>Executive Summary</b>  | <b>3</b>  |
| <b>Introduction</b>   | <b>4</b>  |
| <b>Storage Landscape: Complexity and Management Challenges with<br/>No Easy Solution</b>                      | <b>4</b>  |
| <b>Controller Technology: How It Works</b>  | <b>5</b>  |
| <b>SAN Administrator Challenge: Optimize Dynamic Workloads</b>  | <b>6</b>  |
| <b>Introducing the Dynamic Virtual Controller</b>   | <b>7</b>  |
| <b>Hitachi Unified Storage 100 Family Systems Support Optimized<br/>Dynamic Workloads for Top Performance</b> | <b>9</b>  |
| <b>File Sharing and Content Archiving Environments Benefit As Well</b>  | <b>9</b>  |
| <b>Conclusion: Greater Efficiency for Today and Tomorrow</b>  | <b>10</b> |

---

## Executive Summary

Storage administrators, already confronted with unprecedented data growth, now encounter a new imperative: to discover and exploit efficiencies at every level, from storage administration to data flow, in order to keep to an absolute minimum the amount of new staff, hardware and software required to manage ever-increasing volumes of data. The Hitachi Unified Storage 100 family addresses the needs of growing businesses by supporting operational efficiency at a high level and providing throughput and scalability at every stage of growth.

At the heart of the Unified Storage 100 family, innovative Hitachi dynamic virtual controller technology plays a pivotal role. It automatically and dynamically manages storage workloads, reduces the risks of misconfigured systems, and dramatically simplifies storage administration, while providing the superior levels of performance that expanding businesses require. Its dynamic virtual controller technology offers something unique: by decoupling the storage controllers from specific LUN assignments, it allows I/O processes to be automatically routed along the most efficient path, for unprecedented flexibility and optimized SAN performance.

This paper examines the complexity and management challenges presented by the current storage landscape, the explosive growth of data storage requirements and what it all means to businesses. In addition, the paper reviews various controller technologies as well as issues faced by administrators working to optimize dynamic workloads. Finally, it provides a detailed discussion of Hitachi dynamic virtual controller technology. It looks at the controller's ability to simplify storage management, reduce risk, improve efficiency and performance, and provide the capabilities required for future growth.

## Introduction

Today, organizations of all sizes must manage, store and protect exponentially increasing amounts of data. Regulatory requirements for preservation grow ever-more stringent as data continues to be generated at this unprecedented rate. This explosion of data has left businesses across the globe struggling with ways to economically store and protect data, while staying within budget and maintaining acceptable levels of service.

The continued expansion of data storage requirements is critical to the success of many organizations. To meet current and future needs, IT departments must look to new approaches for storing, protecting and managing corporate, customer and partner data. But first they must overcome the complexity and management issues associated with the current storage landscape.

## Storage Landscape: Complexity and Management Challenges with No Easy Solution

Growing data requirements have, understandably, brought certain pressures on IT management to improve productivity and maintain performance levels. Other segments of the business regard IT as a strategic asset as well as an operational necessity. As a result, other parties, namely CEOs and senior management, are looking to IT departments for both more strategic value and more ideas about where and how to improve business operational efficiencies. In many cases, they have not yet found what they are looking for. According to a recent ESG survey, 42% of businesses report that cost-reduction initiatives and 33% report that business process improvement initiatives will have the greatest impact on their IT spending decisions over the next 12 to 18 months.<sup>1</sup>

Organizations face not just strategic but also technical challenges. They must keep their data searchable, available and accessible while enforcing data protection requirements. They have to find ways to efficiently store and archive data, back up applications and recover quickly from disasters. They must meet increasingly stringent government regulations in the areas of storage, reporting and archiving. And they must find solutions to all these problems that not only meet today's requirements, but also have the flexibility and scalability to meet tomorrow's. To truly build operational efficiencies into the infrastructure, organizations must acquire a platform or portfolio of technologies that is futureproof and leverages the latest innovations.

And, because budgets are never unlimited, every business faces cost considerations. IT departments must accommodate data growth by adding storage capacity while minimizing storage costs. They must strive to reduce backup windows in order to avoid downtime, which could prove costly to the business. They must minimize storage management tasks so that they can apply a minimal amount of personnel resources and keep power consumption as low as possible in the face of rising energy costs. These cost control measures often come into conflict with 2 other imperatives: protecting the organization's storage investment through regular maintenance and upgrades, as well as ensuring that the storage environment can grow comfortably along with the business.

---

<sup>1</sup> Source: ESG Research Report, 2011 IT Spending Intentions Survey, January 2011.

IT departments are left trying to bridge the gap between what their business needs and what it can afford, with an assortment of less-than-desirable solutions. They use old tools to manage rapidly evolving, rapidly growing data sets. Their overworked staff has to work even harder to keep up with management tasks, even as increased workloads lead to more errors. Nobody on staff can work on new projects, or projects that could create more value for the business, because they are too busy treading water.

The old fixes just won't work anymore. What's required is an entirely new data storage solution that addresses these challenges at every level: strategic, technical and budgetary. This storage solution must be easy to administer and able to adapt readily to changes in workflows. And, in order to do that, it must provide a high level of operational efficiency. This level of efficiency is achieved through virtualizing the storage infrastructure, which then leads to dynamic load balancing of performance levels.

Concerns about server utilization, including physical space, reliability, disaster recovery, costs and the growth of data, have driven many businesses to adopt virtualization solutions. Many organizations have come to realize that just as data growth adds complexity to storage issues, virtualization also dramatically increases the level of complexity. For example, an organization with fewer than 25 servers will grow over time to as many as 10 virtual machines (VM) per CPU with adoption of virtualization using current VMware sizing guidelines. This transforms fewer than 25 physical servers into as many as 500 virtual servers. This level of complexity has a profound impact on storage. It increases pressure on other network operations, such as customer environments, and makes it impossible for SAN administrators to manually optimize a SAN that includes virtualization adoption.

## Controller Technology: How It Works

As data storage requirements escalate and IT departments begin to reach the limits of virtualization's ability to control the necessary hardware, they must examine new approaches to storing and managing data. Controller technology offers a proven solution to data storage, but not all controller technology is created equal.

A storage controller manages the flow of information between the server and the data, assigning 2 paths, in case one of the paths fails or is overloaded. "Active" and "passive" refer to the relative ability of a controller to assume an application workload: whether it sits idle until it is called upon to balance a load, or it is actively available at all times. An asymmetric controller architecture, also referred to as asymmetric logical unit access (ALUA), requires the administrator to define a preferred path and offers faster performance for an I/O following the preferred path. A symmetric controller does not require preferred path definitions and offers the same level of performance regardless of the path.

Despite many advances in storage technologies, no host or network balancing algorithm is perfect. For example, the Microsoft Windows implementation Multi Path IO (MPIO) is based on a round-robin algorithm. Round-robin is a scheduling algorithm that assigns time slices to each process in equal portions and in order. It handles all processes without priority. However, the round-robin approach does not balance server performance, and if scheduled jobs vary in size the round-robin algorithm favors large jobs over others.

For the best levels of performance and availability, every layer of technology must be balanced: at the host, network and storage levels. Host-level balancing generally requires installing and updating

host software as well as determining configuration and compatibility dependencies based on applications, operating system type, cluster configuration and virtualization requirements. Similarly, a balanced network eliminates bottlenecks and reroutes I/Os in order to keep data moving efficiently.

Imbalances at the host and network levels must be absorbed and mitigated at the storage system level. As a result of these many dependencies and complexities, it is far from easy to develop a well-balanced, highly available, high-performance system. For example, if a network has one large block of data that is requested as an I/O request, it will be limited to the bandwidth of a single link, creating a load imbalance. Even if link aggregation is used on both the server and the storage system, any traffic flow between a single server and the storage system could not exceed the bandwidth of a single link in the aggregate. Because none of these balancing technologies are perfect, administrators must rely on the solid foundation of the SAN to compensate. These limitations drive up the administration cost or at least make it impossible to get efficient use of a SAN without SAN administrators.

## **SAN Administrator Challenge: Optimize Dynamic Workloads**

In addition to the host of problems associated with developing a well-balanced, high-performance storage system, SAN administrators face a number of challenges both in storage provisioning planning and in managing workloads that are, by their nature, dynamic.

In planning for storage provisioning, SAN administrators have historically had to perform tedious, manual processes for each server, application and volume. They had to determine the required capacity and performance for each in a series of tedious and error-prone iterations. They had to decide on the RAID type and number of disks, as well.

The weak link in the storage provisioning process, an area of great ambiguity and increased likelihood of error, is in the characterization of applications. Online applications, email and database profiles tend to be dynamic. This is because they are so closely linked to individual users; they involve processes that are time sensitive or they are tied closely to the volume of business an entity performs. As a result, IT storage administrators are constantly playing "catch up" as they try to understand what is already required of these applications and what they will likely require in the future.

Applications do not constitute the only constantly shifting benchmark for storage provisioning; however, workloads themselves are dynamic. They vary throughout the day, week, month or quarter, depending on when the business day starts or finishes, the frequency of internal deadlines and the incidence of external events that can be neither controlled nor predicted. An email system with a large number of mailboxes will obviously experience greater-than-normal stresses during working hours when traffic is heavy, for example.

Workloads vary not only by the day but also by the time at which different processes run, such as the loading of databases, regular reporting and the running of backups. All of these can take needed storage offline, leaving the remaining available storage under pressure to accommodate a greater load.

The dynamic workload makes it a constantly shifting target; even if a design is solid and administrators have planned diligently and to the best of their ability, it still can take a number of permutations before an optimized configuration is attainable. Therefore, the old architectures are failing due to their inability to provide the flexibility required to efficiently optimize dynamic workloads.

## Introducing the Dynamic Virtual Controller

Most midrange storage systems today are dual controllers with an asymmetric active-active architecture, in which both controller nodes are available to process I/Os for their assigned set of LUNs and provide standby capability for the other, non-assigned LUNs. This arrangement requires the administrator to make manual LUN-to-controller assignments based on forecasted workloads. Rebalancing of performance across the controllers as workloads change is also a manual task. Since the performance to a particular LUN is not equal through both controllers, the system is vulnerable to major I/O bottlenecks without an automated means of balancing the load. Moreover, the asymmetric relationship to host server dependencies creates serious maintenance issues. If an organization has 25 servers today, that could easily grow to as many as 500 VMs in the near future, plus host agents. Because this architecture requires host software on the servers, and even virtual servers require management, the host solution is highly time consuming and difficult to maintain.

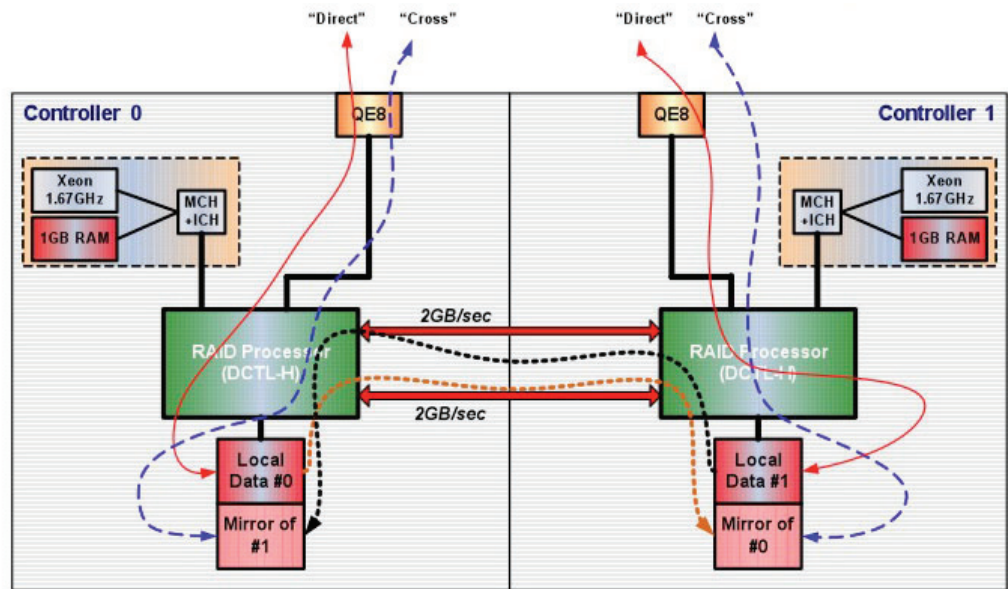
Asymmetric active-active architecture can also significantly decrease productivity and increase setup times when matched to VMware environments. For example, the time to configure path failover can take up to 1 hour per physical server. If an organization supports 20 servers, each running 25 VMs, the configuration process can take 2-1/2 days. In a symmetric active-active environment, where the controller removes the need for LUN ownership, the configuration takes less than 5 minutes per server.

In addition, asymmetric active-active architecture has inherent weaknesses and dependencies, including the following:

- Performance may be adversely affected if I/O is routed to the non-owning storage controller of a LUN. This can be the result of a misconfiguration of the SAN paths.
- Failover software must be configured so that it only load balances across the active-optimal path for a given LUN.
- SAN changes, such as component failures or replacements, may cause I/O paths to change; the administrator must ensure that hosts are still using optimal paths to LUNs.
- Firmware upgrades to the storage system are time consuming as the primary paths must be failed over before a controller is upgraded.

Another option is presented by the Hitachi dynamic virtual controller architecture (see Figure 1). Instead of providing unequal access to LUNs, which must be manually managed, they virtualize the relationship between controllers and LUNs. The dynamic virtual controllers work simultaneously to automatically assign the access paths from the back of the controllers to the LUN and even dynamically re-assign these paths as workloads shift. All LUNs are accessible, regardless of the physical port or the server from which the access is requested. Utilization rates of each controller are monitored and can be automatically brought into balance whenever necessary. This dramatically simplifies management and offers a number of administrative benefits.

Figure 1. Dynamic Virtual Controller Architecture



In a VMware environment, the dynamic virtual controllers' automatic load balancing reduces setup costs and eliminates performance bottlenecks. While VMware offers both fixed and MRU (most recently used) multipathing policies, MRU policies can result in nodes using non-owner controllers with no mechanism to fail back to the owner path. Fixed multipathing, on the other hand, uses the configured path to a LUN in preference to any other. In the event of a path error, VMware will fail to the next available path. However, VMware continues to test the failed path; when the preferred path recovers, VMware will route all activity back through the preferred path for that LUN. In this way, using fixed multipathing with active-active symmetric load balancing allows a configuration to dynamically repair itself.

The dynamic virtual controllers do not rely on proprietary or 3rd-party path management software for failover and load balancing; instead, those functions are native to the storage operating system. The removal of path management software from the equation makes it easier to optimize SAN performance and perform maintenance. For example, storage controller microcode updates can be installed quickly and easily, without any interruption to host I/O requests, because only the individual microprocessor is rebooted instead of the entire controller. This allows the application to continue running smoothly, even if the server has only 1 access or path (host bus adapter) to the storage controller.

## Hitachi Unified Storage 100 Family Systems Support Optimized Dynamic Workloads for Top Performance

The Hitachi Unified Storage 100 family systems also enable dynamic load balancing across the 2 controllers to improve total cost of storage, performance and flexibility. Overall system performance improves as bottlenecks are eliminated at the controller; these improvements in turn lead to greater ability to meet application service level agreements (SLAs). Installation is faster and easier because administrators are not required to set paths or assign controller ownership of LUNs. Ongoing maintenance is also easier because administrators won't need to monitor systems and adjust LUN ownership to correct any load balances that lead to bottlenecks.

When workload is automatically rebalanced across the controllers, the storage administrator is relieved of the burden of resetting the primary and failover paths to the storage. This is the advantage that the dynamic virtual controller design provides. With asymmetric active-active controllers, whenever LUN ownership is re-assigned to balance workload, the primary and failover paths must also be redone. The ability of the Hitachi Unified Storage 100 family to automatically self-tune keeps the storage operating at optimal levels of performance.

Compared to traditional controller architecture still used by many organizations, the dynamic virtual controller architecture offered by Hitachi represents a revolutionary advance in data storage. With 2 active controllers virtualizing LUN ownership without manual intervention, management is radically simplified, microcode updates can be effortlessly installed and system performance is made more predictable. For businesses that rely on efficient data storage, the Hitachi Unified Storage 100 family offers the capabilities required for continued growth.

## File Sharing and Content Archiving Environments Benefit As Well

Unstructured data growth continues to outpace well laid out plans of IT administrators. With its file module, the Hitachi Unified Storage 100 family provides the ability to simultaneously support network attached storage (NAS) file sharing protocols, such as NFS for UNIX and Linux operating system environments, and CIFS for Microsoft Windows environments. File sharing on midrange storage often requires the servicing of millions of small random I/O requests from thousands of clients, who expect consistent and dependable response times. The dynamic virtual controller architecture makes this happen. Without it, file systems would also have to be mapped to specific LUNs and require periodic manual adjustment as file systems and file shares hit capacity limits. This would lead to a painful migration exercise that most storage administrators can do without. Additionally, load balancing by the controllers in the block module also ensures that file reads and writes always have predictable access latency times and performance; this, in turn, enables the achievement of service level objectives (SLOs).

A key benefit of the dynamic virtual controller architecture is in the ability to create and provision large dynamic storage pools since manual LUN mapping is no longer required. This common storage pool enables the storage of multiple data types, such as block, file and object. For example,

Hitachi Content Platform, with its large object store capability, can use the same storage pool to store content either for active archiving purposes, or for cloud infrastructures in the data center. This capability is unprecedented in the industry and provides organizations flexibility and long-term cost savings.

## Conclusion: Greater Efficiency for Today and Tomorrow

Efficiency comes in many forms. It may be reducing the amount of time a storage administrator spends assigning LUNs, reducing the risk of costly downtime or finding a more direct path for individual I/O requests. It may be determining which storage system most closely aligns with an organization's needs for security, reliability and performance. It's all part of how organizations must decide what brings maximum benefit to the business for a minimum of cost and effort.

The Hitachi Unified Storage 100 family meets all these definitions of efficiency. As it automates, streamlines and accelerates the increasingly complicated task of managing dynamically changing workloads, it helps growing businesses meet today's storage needs and stands ready to address their future needs.

Hitachi Data Systems has a long and proven track record of providing customers with storage solutions that maintain required service levels for essential business functions and processes. We ensure data availability and resilience, and aid in the prevention and mitigation of interruptions caused by planned or unplanned events. The availability of the Hitachi Unified Storage 100 family allows companies to increase data storage productivity and efficiency while reducing costs associated with managing and maintaining SANs. In an environment where data storage requirements continually grow, organizations that take advantage of this revolutionary Hitachi storage solution gain a competitive differentiator and are well positioned for future growth.

 **Hitachi Data Systems**

---

**Corporate Headquarters**

750 Central Expressway  
Santa Clara, California 95050-2627 USA  
[www.HDS.com](http://www.HDS.com)

**Regional Contact Information**

**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto: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 2012. All Rights Reserved. WP-309-C DG April 2012