# Symmetric Active/Active Controller

## Optimizing Dynamic Workloads

Application Brief

*By Patrick Allaire and Mark Adams*

**October 2008**

# 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 Adaptable Modular Storage 2000 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.

Within the Adaptable Modular Storage 2000 family, the innovative architecture of the Hitachi Dynamic Load Balancing Controller plays a pivotal role by automatically and dynamically managing storage loads and dramatically simplifying storage administration, while providing the superior levels of performance that growing businesses require. Its symmetric active/active controller technology offers something unique: by decoupling the host system from specific LUN assignments, I/O processes can 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 the Hitachi Dynamic Load Balancing Controller — a symmetric active/active implementation — and its ability to simplify storage management, improve efficiency and performance, and provide the capabilities required for future growth.

**For best results use Acrobat Reader 8.0.**

# Contents

**Introduction** ........ 1

**Storage Landscape: Complexity and Management Challenges with No Easy Solution** ........ 1

**Forecast: Data Growth May Outpace Virtualization Capabilities** ........ 3

**Controller Technology: How It Works** ........ 4

**SAN Administrator Challenge: Optimizing Dynamic Workloads** ........ 5

**Introducing the Symmetric Active/Active Controller** ........ 6

**The Hitachi Adaptable Modular Storage 2000 Family with the Hitachi Dynamic Load Balancing Controller Improves TCO, Performance and Flexibility — and Delivers 99.999 Percent Availability** ........ 8

**Conclusion: Greater Efficiency for Today and Tomorrow** ........ 9

# Symmetric Active/Active Controller

## Optimizing Dynamic Workloads

Application Brief

*By Patrick Allaire and Mark Adams*

## Introduction

Today, organizations of all sizes must manage, store and protect exponentially increasing amounts of data. Preservation requirements grow ever more stringent as data continues to be generated at this unprecedented rate, and 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. The fact that other segments of the business now regard IT as a strategic asset as well as an operational necessity means that 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 one survey, only 45 percent of businesses are satisfied with their current level of storage administration staff and skills. About one-tenth reported a “problematic shortage” in staff or skills.[1] In other words, this is a strategic problem for the business, and one that current arrangements all too often fail to address.

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 future proof 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, 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 two other imperatives: protecting the organization’s storage investment through regular

[1] “Medium-Size Business Server and Storage Priorities,” Mary Johnston Turner and John McKnight. Enterprise Strategy Group, June 2008, p. 8.

maintenance and upgrades, as well as ensuring that the storage environment can grow comfortably along with the business.

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 any more. 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 dynamic load balancing, which optimizes performance rapidly, automatically and with a minimum of administrative overhead. Because midrange storage system architecture includes dual controllers, true active/active symmetrical design means that both controllers are operational and actively working with no preferred path and no performance penalty for any path; a situation in which both controllers are up, but the secondary controller does nothing, is more accurately described as "active/passive." Vendor nomenclatures can be confusing in this regard, but the functional difference between these two technologies is quite clear.

Concerns about server utilization, including physical space, reliability, disaster recovery, costs and the growth of data, are driving many businesses to adopt virtualization solutions. SearchDataCenter.com's 2008 Purchasing Intentions Survey reported that 61 percent of respondents run less than 10 virtual machines (VMs) per server, with 33 percent running 10 to 25, and only 5 percent running more than 25 VMs on a server. However, many organizations don't 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 VMs per CPU with adoption of virtualization using current VMware sizing guidelines, transforming less than 25 physical servers into as many as 500 virtual servers. This level of complexity has a profound impact on servers, 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.

# Forecast: Data Growth May Outpace Virtualization Capabilities

In a typical customer environment, storage administrators must provide increasing levels of storage capacity and improve performance while data continues to grow daily. To cope with this data growth, IT departments add servers and rely increasingly on virtualization to manage capacity without adding more physical hardware.

Figure 1 shows the annual growth rate of data storage. Businesses from small to large are adding production servers at a rapid pace. While large organizations are the most likely to be growing their storage at the fastest rate, over one-third of businesses with under 100 production servers are still increasing the number of servers under management by 11 to 20 percent annually.[2] Thus, as data continues to grow, storage grows as well, and virtualization can only go so far in controlling the amount of hardware required to cope with the onslaught.

***Figure 1. Approximate Annual Growth Rate of Total Amount of Data Storage (by number of production servers/percent of respondents)***

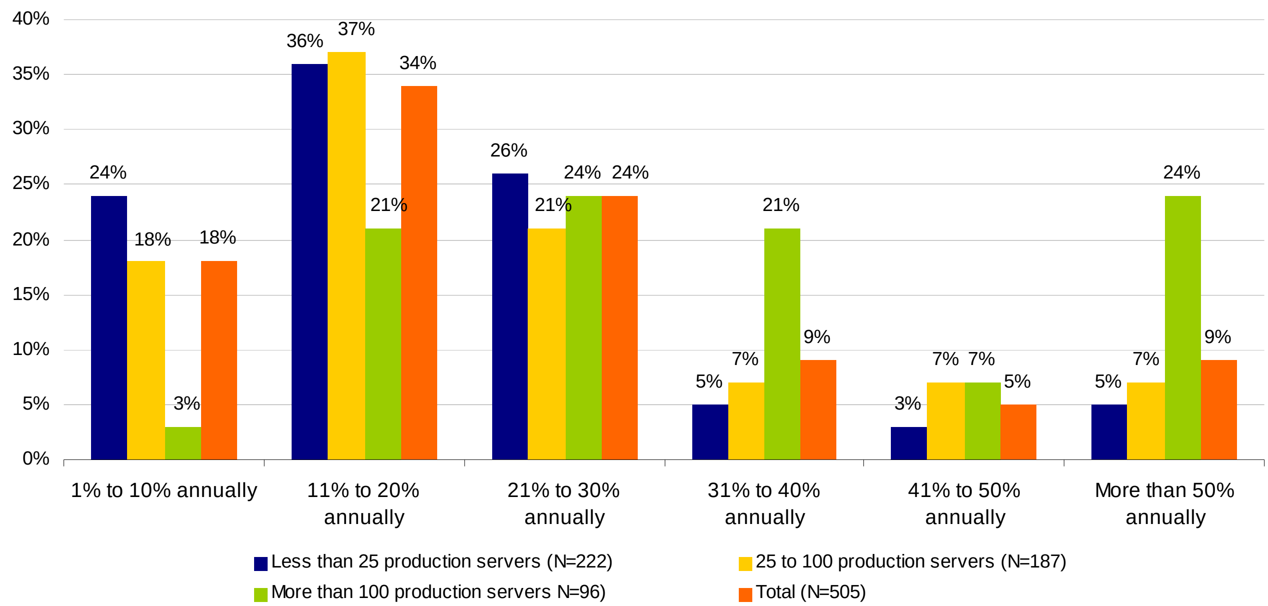


*Source: Medium-Size Business Server & Storage Priorities, Enterprise Strategy Group Research Report, June 2008, By Mary Johnston Turner and John McKnight, with Jennifer Gahm*

[2] Ibid, p. 21.

# 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 revolutionary, automated 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 two 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 requires the administrator to define a preferred path and offers faster performance for an I/O following the preferred path, while 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 — handling 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, OS 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 administrator(s).

# SAN Administrator Challenge: Optimizing 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, e-mail and database profiles tend to be dynamic 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 e-mail 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, which means that — 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 Symmetric Active/Active Controller

In a traditional controller architecture — referred to here as an active/passive architecture — one controller is actively available to process I/O requests while the other is idle in standby mode, ready to take over I/O activity if the active primary controller fails or is taken offline. The shortcoming of this mode has mainly to do with the fact that all the user data residing in LUNs must be assigned manually to a primary controller; this process is tedious and time consuming, and involves human effort, which in turn involves the possibility of human error. In the event of hot spots or utilization imbalance issues (such as a Microsoft® Exchange change of I/O pattern or VMware Virtual Servers moving by effect of VMotion), the administrator must manually reassign the workload to the secondary controller. If this does not happen, the application will experience worst performance, which will have a direct impact on the business. Sometimes this process involves a system being rebooted, which is a disruptive operation.

There are also inherent weaknesses in the traditional controller architecture when used in a VMware environment. For example, virtual machine I/O may be delayed for almost 60 seconds while failover takes place, particularly on an active/passive system.[3] In the case of an active/passive system with path policy fixed, path thrashing might be a problem. Path thrashing occurs when two hosts access the LUN through different storage controllers and, as a result, the LUN is never actually available. This may result in a virtual machine's failure if all paths to the storage device where the virtual machine disks are stored become unavailable.[4]

A slightly better arrangement involves the asymmetric active/active architecture, in which both controller nodes are available to process I/Os for a given LUN and provide a standby capability for the other. This arrangement eliminates the path/LUN thrashing but does little for storage system productivity or SAN controller balancing. And 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 one hour per physical server. If an organization supports 20 servers, the configuration process can take two and one-half days. In a symmetric active/active environment, where the controller removes the need for LUN ownership, the configuration takes less than five minutes per server.

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

- Performance may be adversely affected when I/O is routed to the non-owning storage controller of a LUN.
- 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.

[3] "Fibre Channel SAN Configuration Guide," VMware, Inc., 2006, p. 43.

[4] Ibid.

A third option presented by the Hitachi Dynamic Load Balancing Controller offers symmetric active/active architecture (see Figure 2). Instead of having an active and a passive controller, both controllers are active, able to work simultaneously, and able to dynamically and automatically assign the access paths from the back of the controller to the LUN. 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 2. Active/Active Symmetric Features***

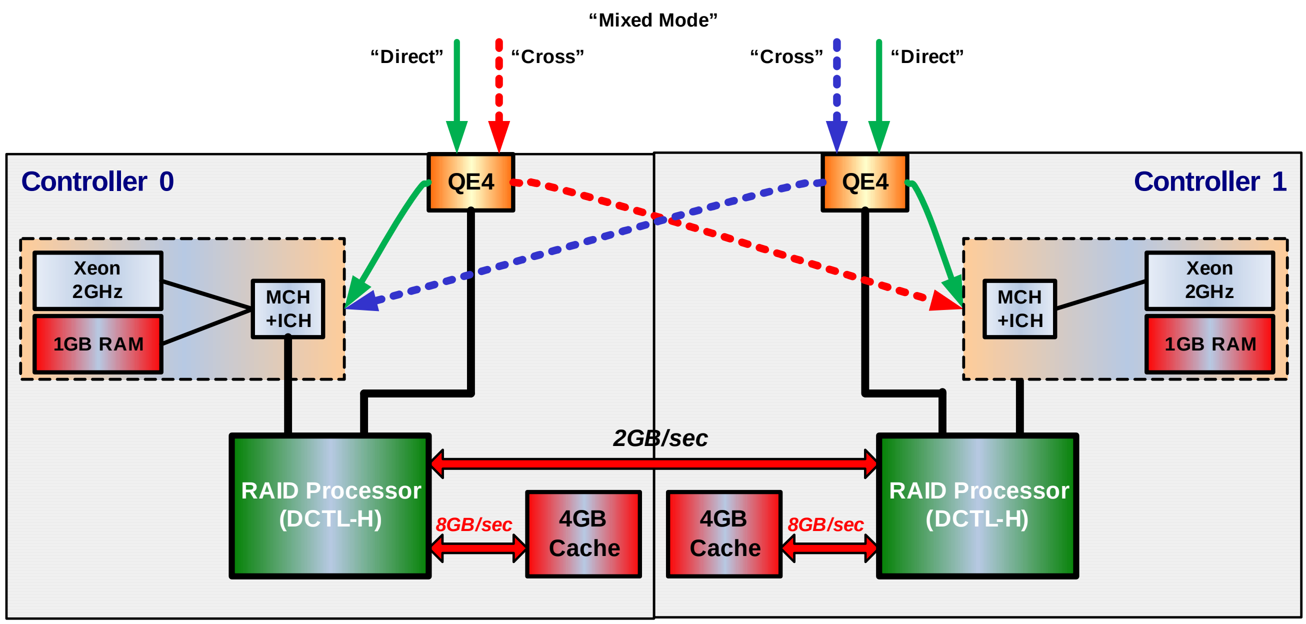


In a VMware environment, the active/active symmetric controller's automatic load balancing reduces setup costs and eliminates performance bottlenecks. While VMware offers both fixed and MRU (Most Recently Used) multipathing policies, MRU policies can result in nodes using non-owner controllers with no mechanism to fail back to the owner path.[5] 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, and 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.

An active/active symmetric controller does not rely on third-party host software for failover and load balancing; instead, those functions are native to the operating system. The removal of third-party host 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 with little impact to operations 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 one access or path (host bus adapter) to the storage controller.

[5] "Recommendations for VMware Multipathing with Hitachi Modular Storage Systems," Hitachi Data Systems, March 2008.

# The Hitachi Adaptable Modular Storage 2000 Family with the Hitachi Dynamic Load Balancing Controller Improves TCO, Performance and Flexibility — and Delivers 99.999 Percent Availability

Even better, the storage administrator no longer carries a number of burdens: installing the host core touching software; maintaining third-party software that they may or may not be familiar with; resolving conflicts between the vendor's software and other installed system software; resolving conflicts between the software vendor's failover processes and the storage vendor's failover processes; and discovering which products (if any) the vendors' third-party failover supports. Freed of these obligations, the organization no longer has to waste money on software that duplicates operating system capabilities while spending untold amounts of time in all the above mentioned activities.

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. The intelligence that is built into the Hitachi Adaptable Modular Storage 2000 family with the Hitachi Dynamic Load Balancing Controller identifies unbalanced workloads and dynamically reallocates controller to drive I/O activity so that neither controller is overloaded. The controller coordinates well with path management software, offering balanced performance over multiple data paths when used with the front end load balancing function. And its back end load balancing capability implemented in the Adaptable Modular Storage 2000 family is unique in the market of midrange modular storage systems.

Compared to traditional active/passive controller architecture still used by many organizations, the symmetric active/active architecture offered by the Hitachi Dynamic Load Balancing Controller represents a revolutionary advance in data storage. With two active controllers dynamically assigning LUN ownership without manual intervention, management is radically simplified, microcode updates can be effortlessly installed and system performance is significantly improved. For businesses that rely on efficient data storage, the Hitachi Dynamic Load Balancing Controller offers the capabilities required for continued growth.

# Conclusion: Greater Efficiency for Today and Tomorrow

Efficiency comes in many forms: whether it's reducing the amount of time a storage administrator spends assigning LUNs, reducing the risk of costly downtime, finding a more direct path for individual I/O requests, or 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 Dynamic Load Balancing Controller meets all these definitions of efficiency — and 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, ensure data availability and resilience, and aid in the prevention and mitigation of interruptions caused by planned or unplanned events. For example, Hitachi Services Oriented Storage solutions apply service oriented architecture (SOA) concepts to storage to deliver a platform that can be readily reconfigured and optimized to changing business requirements. The availability of the Hitachi Adaptable Modular Storage 2000 family with the Hitachi Dynamic Load Balancing Controller 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.

The Hitachi Adaptable Modular Storage 2000 family with the Hitachi Dynamic Load Balancing Controller offers the only midrange storage solution that provides this enterprise capability. This unique feature makes Adaptable Modular Storage 2000 family systems the most suitable solution for virtualized environments.

Hitachi Data Systems Corporation

**Corporate Headquarters** 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

**Asia Pacific and Americas** 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

**Europe Headquarters** Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.emea@hds.com



WHP-309-00 October 2008