LAB VALIDATION REPORT

HDS Adaptable Modular Storage 2000 Family

Intelligent, Enterprise Class, Modular Storage

By Tony Palmer
With Brian Garrett

March, 2009

Table of Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by HDS.

Introduction

The Hitachi Data Systems Adaptable Modular Storage (HDS AMS) 2000 family is a series of next-generation midrange storage systems which use enterprise-class technologies to improve operational efficiencies and significantly reduce a company's administrative burden. This ESG Lab Report documents the results of hands-on testing of the AMS 2000 with a focus on ease of implementation and management, availability, and agility as well as performance and capacity scalability.

Background

Today's microeconomic climate is driving users to leverage advanced infrastructure and management technologies, particularly in the storage and server areas, to better support top-priority business requirements. At the same time, these technologies help customers improve resource utilization by enabling IT teams to minimize time spent on routine operations and support activities. Medium-sized organizations—defined as having between 100 and 999 employees—are under increased pressure to more efficiently support their businesses while improving information management and keeping costs down. Storage technology providers, in turn, have traditionally responded with storage systems that include scaled-down and compromised feature sets. But medium-sized businesses are increasingly in search of "enterprise-class" information management and infrastructure technologies—and these organizations are making significant investments in networked storage as a key enabler for improved information management.

As shown in Figure 1, recent ESG research shows that the top IT infrastructure initiative influencing server and storage spending in medium-sized businesses during the next 24 months will be improvement and optimization of IT service levels.[1] Other initiatives with major impact include application and database deployment/upgrades and business continuity and disaster recovery programs. In each case, addressing an initiative's priorities requires comprehensive IT services to support new and existing applications and processes while optimizing day-to-day workflows and business processes.

FIGURE 1. SERVER AND STORAGE INFRASTRUCTURE SPENDING

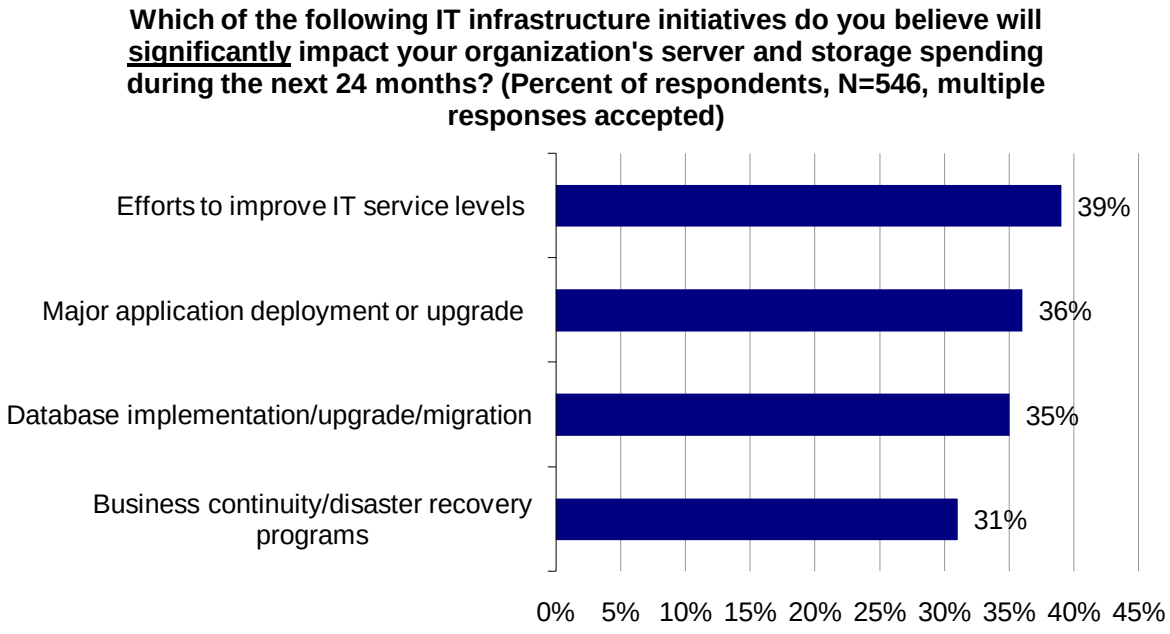

[1] Source: ESG Research Report, *Medium-Size Business Server & Storage Priorities,* June 2008.

HDS Adaptable Modular Storage

HDS has focused on these customer trends and priorities with its AMS 2000 family of next-generation midrange storage systems. These flexible, modular storage systems are designed to address the most crucial challenges faced by storage and IT managers. An AMS 2000 storage system can be configured with as few as four and as many as 480 drives in a single array as shown in Figure 2. Performance of the new AMS 2000 family has increased dramatically and all systems support Fibre Channel or iSCSI host connections.

FIGURE 2. THE HDS AMS 2000 FAMILY

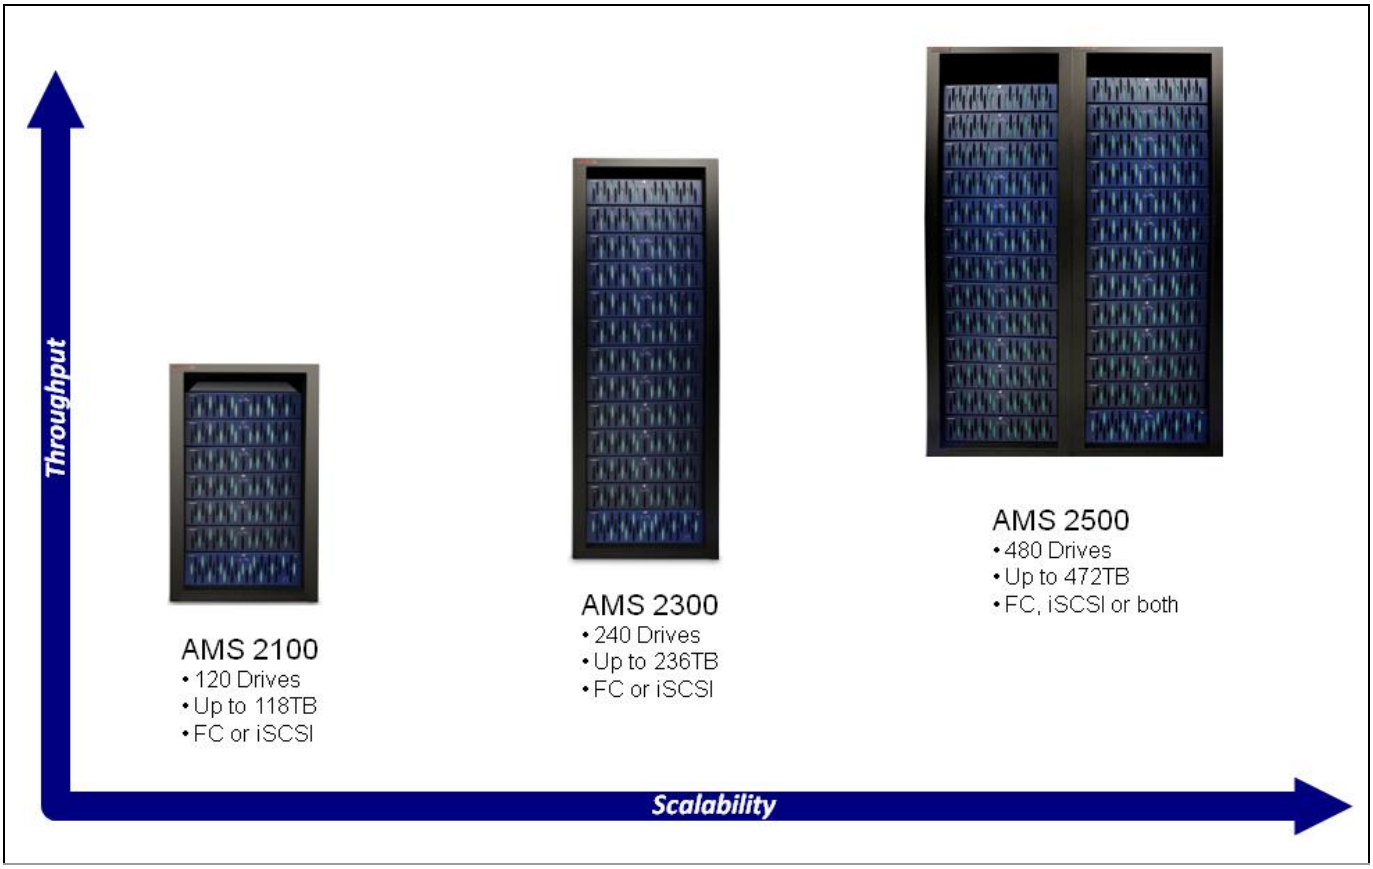

The HDS AMS 2000 family's features include:

- **Enterprise-class Technology** - Symmetric, active/active controllers with load balancing.
- **Green IT** - Automatic power savings with SATA-II drives and programmable spin down for all drives.
- **SAS Architecture** - Up to 32 point-to-point links to the back-end with fully intermixable SAS & SATA-II.
- **Enhanced Mega LUNs** – Up to 60 TB LUNs for simplified performance and scalability.
- **Cache Partitioning** – Cache can be allocated to applications according to their requirements.
- **Modular Volume Migration –** Dynamic, online, non-disruptive volume migration.
- **Online RAID Group Expansion** - Drives can be added to an existing RAID group while the system remains online with non-disruptive data re-striping.
- **LUN grow/LUN shrink** - The capability to increase or reduce the size of any selected LUN

This report documents ESG Lab hands-on testing of the HDS AMS 2000 storage family with a focus on its ability to leverage enterprise functionality to increase performance, availability, and investment protection as it reduces cost, complexity, and energy requirements.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of the HDS AMS 2000 family at HDS' corporate headquarters in Santa Clara, California. Testing began with a look at the Web-based Storage Navigator Modular 2 (SNM2), the management tool for HDS entry level and midrange storage systems. The SNM2 server component runs on either a Windows or Solaris server in the user's environment; both Internet Explorer and Mozilla-based browsers are supported as clients.[2]

Ease of Implementation and Management

ESG Lab testing was conducted on pre-wired, rack mounted AMS 2300 storage arrays. The ESG Lab test bed, as presented in Figure 3, consisted of two Dell PowerEdge R905 servers SAN attached to an AMS 2300 storage array configured with 105 15K RPM SAS disk drives and sixty 7200 RPM SATA drives.[3] One server was running VMware ESX server with Windows 2008 installed as a Guest OS and the other ran Windows 2008 natively. A thirty-two port 4 Gb/sec Brocade switch was used for Fibre Channel SAN connectivity.

FIGURE 3. THE ESG LAB TEST BED

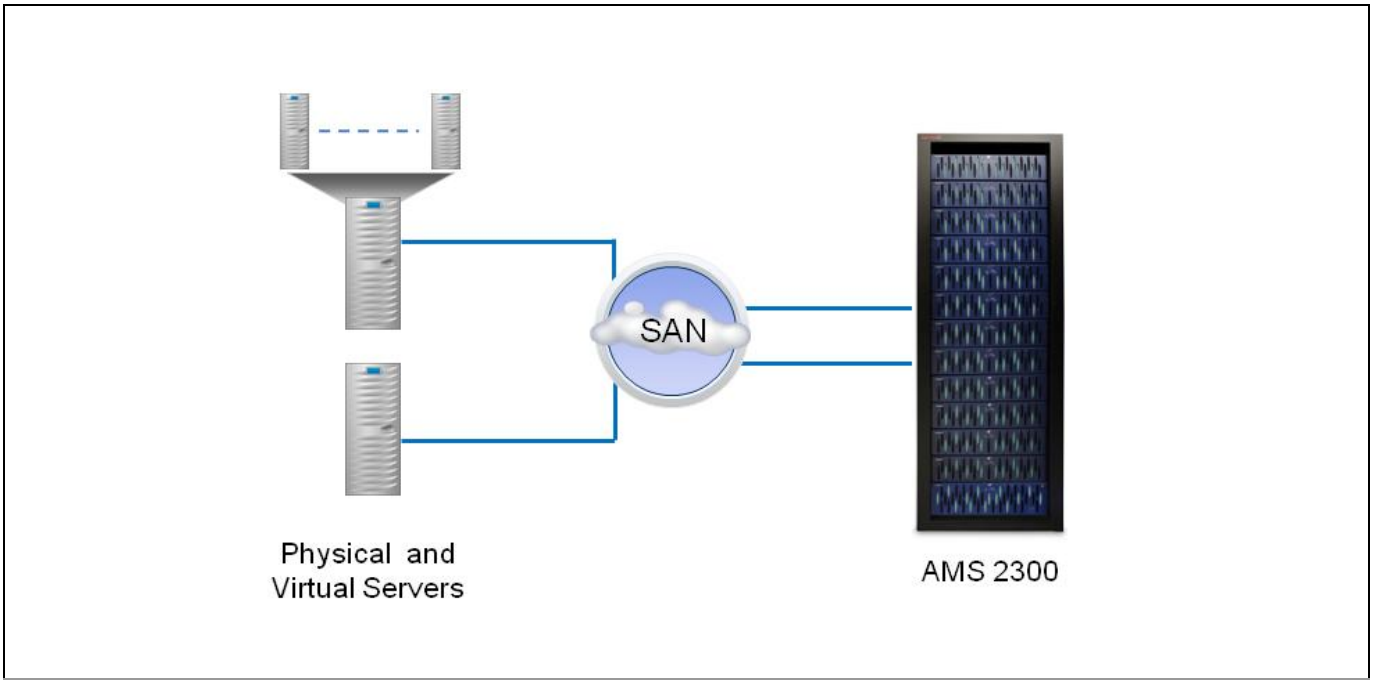

ESG Lab Testing

ESG Lab began by pointing a web browser at the IP address of the SNM2 server. The Add Array Wizard prompted for the IP addresses of the arrays to be managed. IP addresses on AMS arrays are set by HDS at the time of installation.

As seen in Figure 4, the Add Array Wizard can discover systems to be managed using either the specific IP addresses of controllers or a range of addresses. ESG Lab entered a range of four addresses and the wizard discovered two arrays.

[2] SNM2 requirements can be found at: http://www.hds.com/products/storage-software/system-requirements/storage-navigator-modular2-system-requirements-and-support-matrix.html
[3] The test bed configuration is listed in detail in the Appendix.

FIGURE 4. ADDING AN ARRAY

Figure 5 depicts the Hitachi SNM2 GUI. The layout was straightforward and easy to navigate. The navigation pane on the left offers a familiar tree-based menu system, allowing the user to access any function in detail. Context sensitive, common task wizards are shown with friendly icons and clear descriptions on the main page. ESG Lab clicked "Create Logical Unit" to launch the Create & Map Volume Wizard.

FIGURE 5. HITACHI STORAGE NAVIGATOR

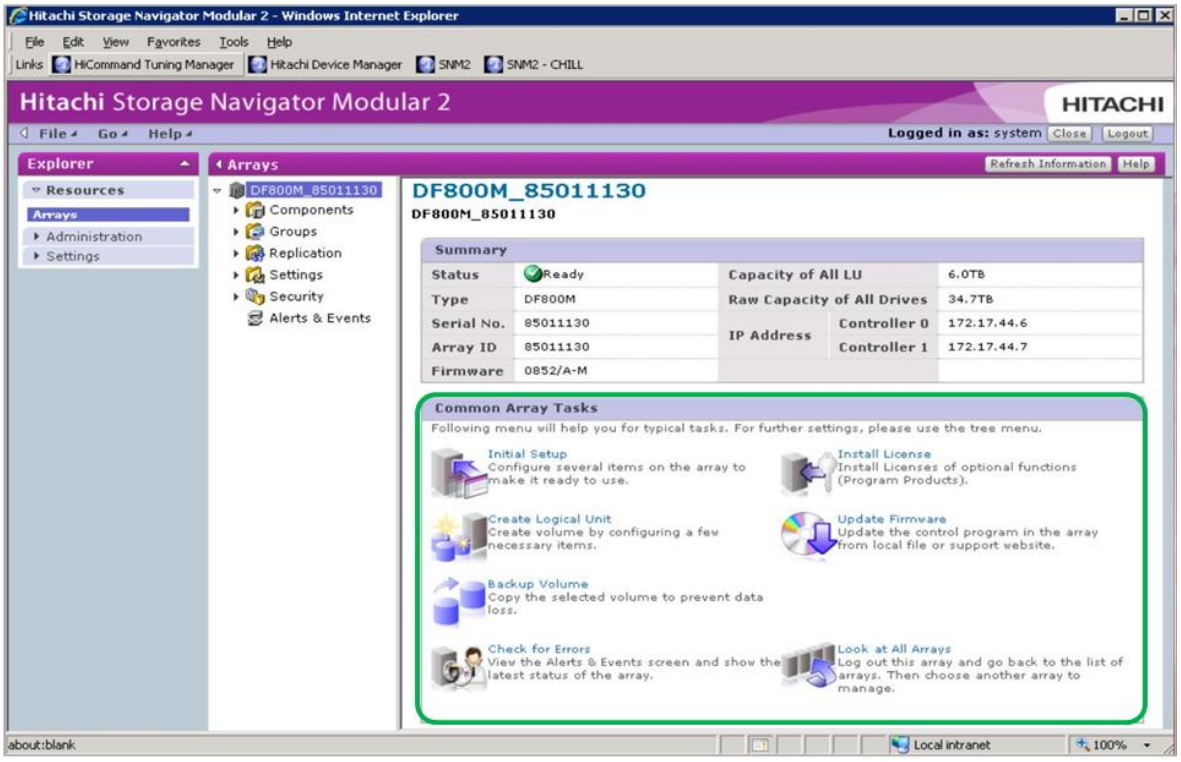

SNM2 uses a single wizard to create new RAID groups and volumes. In this test, the wizard was used to create new volumes from raw disk and assign them to an attached host. As seen in Figure 6, ESG Lab selected the type of disk, the desired RAID protection, and the size of the RAID group. In this example, a RAID 5 group was created using fifteen SATA drives.

FIGURE 6. CREATING AND MAPPING A LUN: CREATE A RAID GROUP

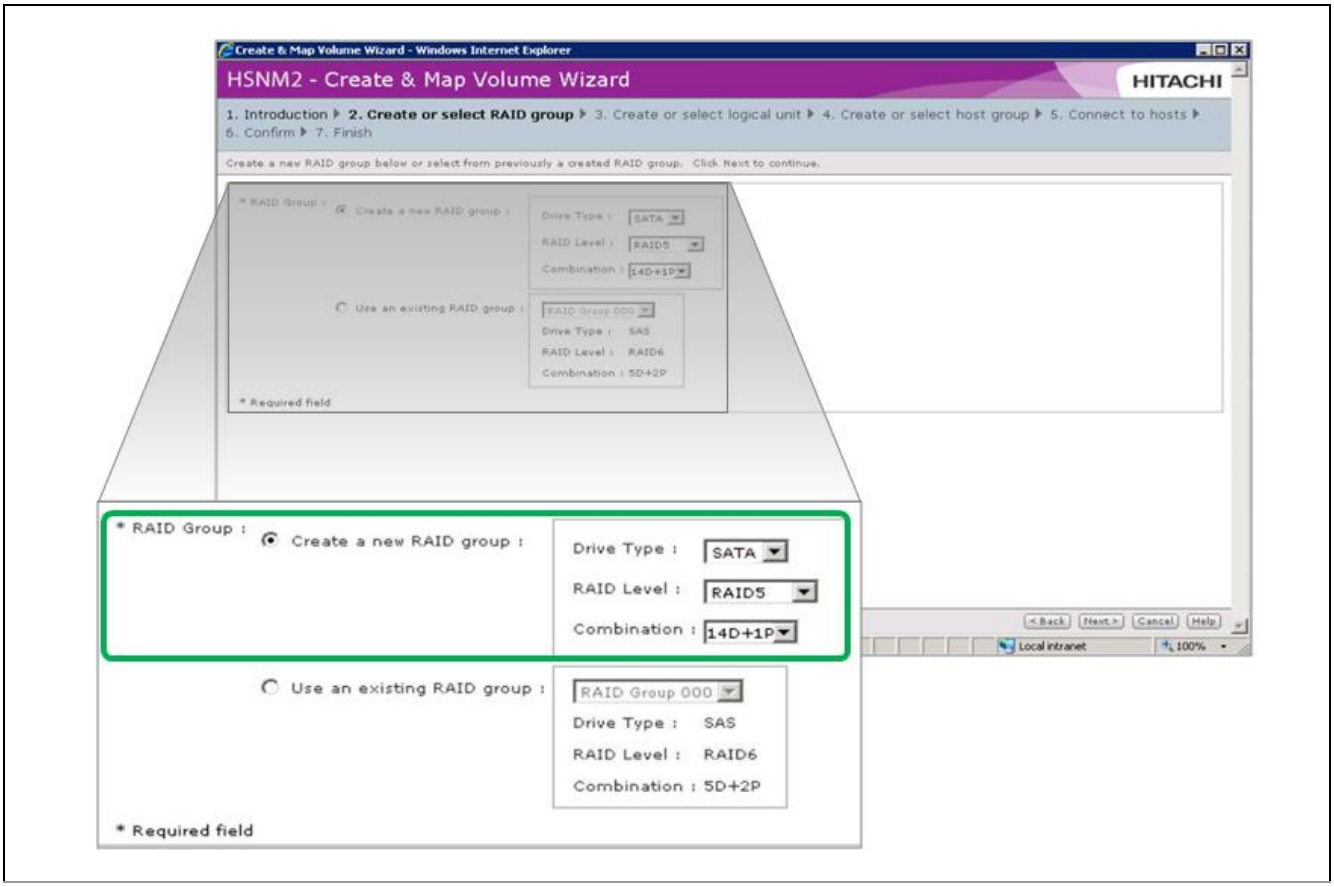

Next, five 25 GB LUNs were created and assigned to the ESX server in the test bed.

The five volumes created can be seen in the confirmation screen, shown in Figure 7.

FIGURE 7. CREATING AND MAPPING A LUN: CONFIRMATION SCREEN

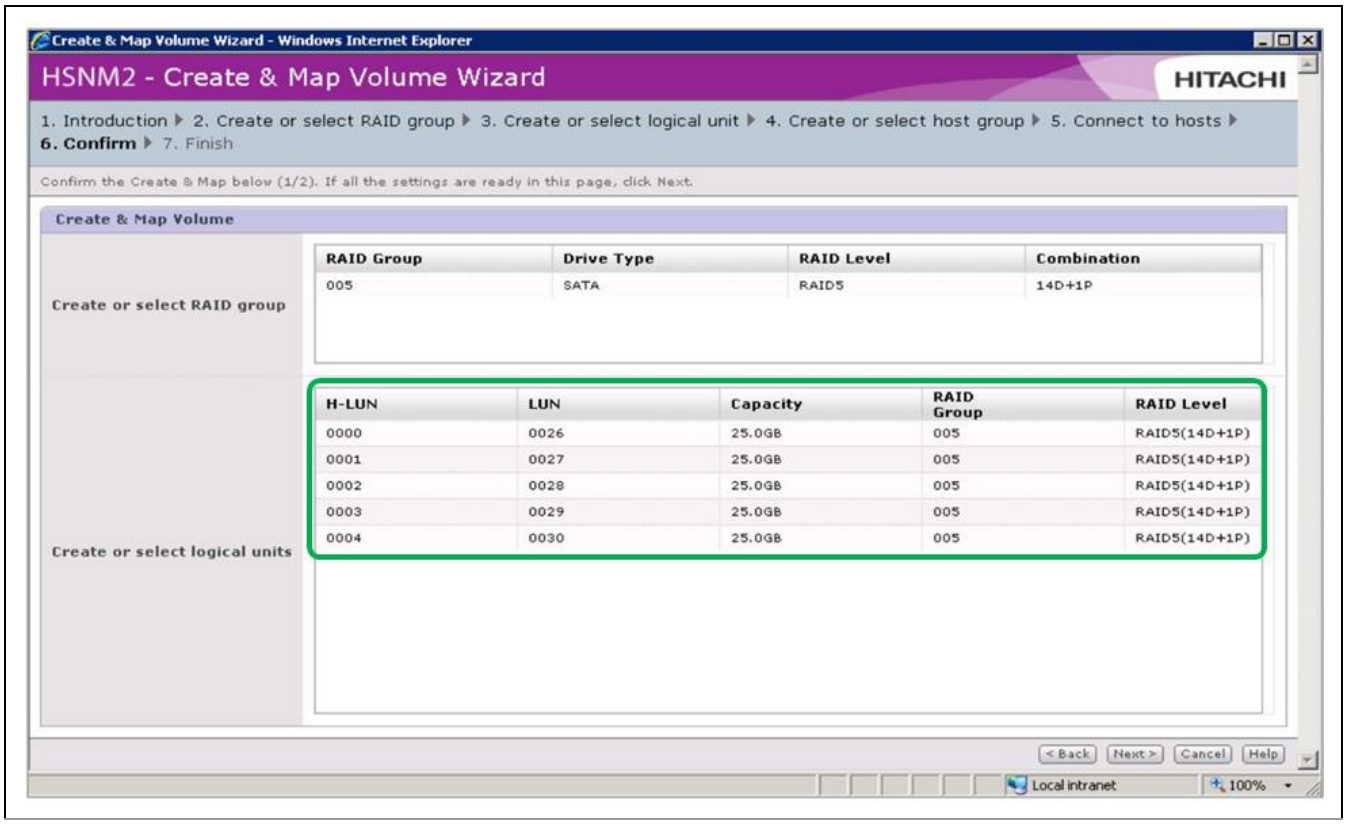

When ESG Lab clicked "Finish," the AMS 2300 proceeded to create the RAID group and initialize the LUNs. The entire elapsed time of the process, from the first click of the wizard to the moment the LUNs were visible to the host, was less than five minutes.

ESG Lab also examined the process for creating LUNs on a previous generation AMS 1000 in order to get a feel for the differences between the two systems. Creating a LUN on an AMS 1000 was very similar to the procedure used for the AMS 2300, but with one notable exception: after the new volume was created, ESG Lab had to edit the properties of the volume to assign it to the correct controller as shown in Figure 8. This step is not required with the AMS 2300 due to an Active-Active architecture, which allows hosts to access any volume through either front-end controller.

FIGURE 8. ASSIGNING A LUN TO A CONTROLLER IN THE AMS 1000

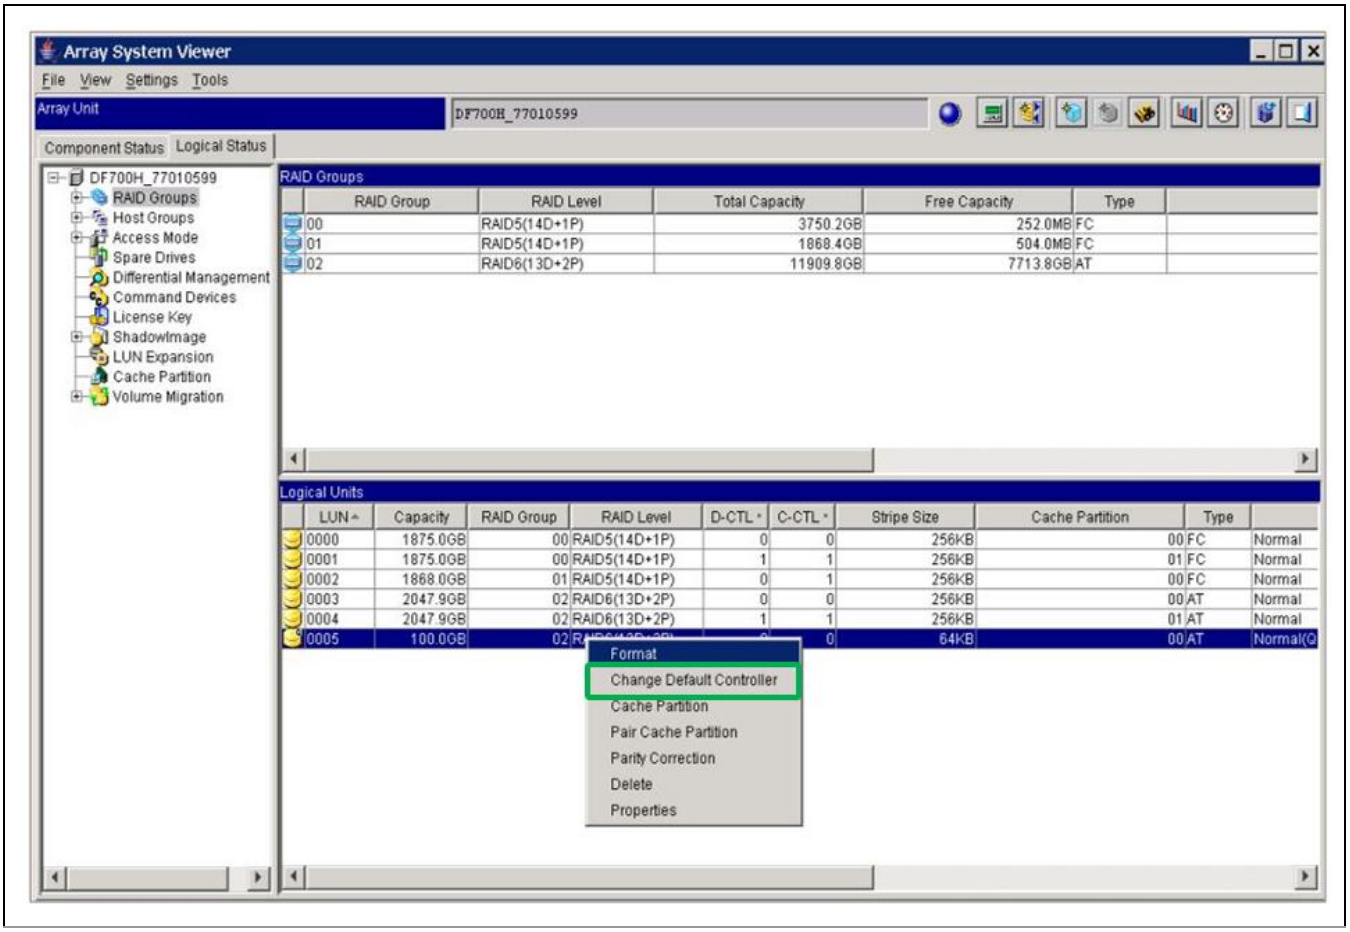

Next, ESG Lab explored the AMS family's ability to expand and reduce the size of existing volumes online. Both functions are executed using the same wizard, the Change Logical Unit Capacity Wizard. First, ESG Lab selected a 100 GB LUN for reduction and clicked 'Change Capacity.' The Change Logical Unit Capacity Wizard, seen in Figure 9, was launched. The wizard displayed the capacity of the LUN, the free capacity in the array, and allowed input of the volume's new capacity. ESG lab chose to reduce the capacity of this 100 GB LUN by half, to 50 GB.

FIGURE 9. SHRINKING A LUN

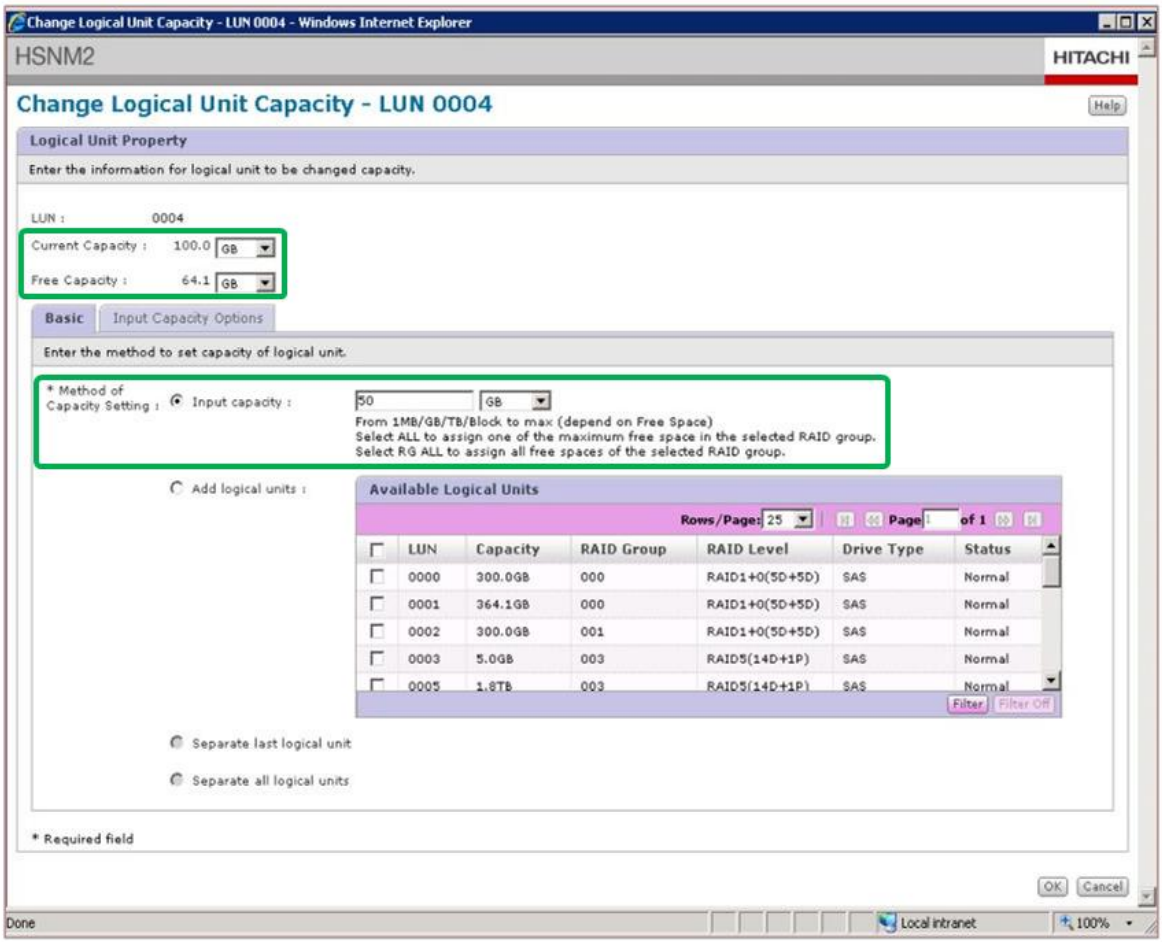

Clicking 'OK' brought up a warning screen that informs the user that shrinking a LUN is a destructive operation; the administrator must click a checkbox acknowledging this before the procedure can finish. Once the LUN was re-sized, ESG Lab performed a re-scan of the LUN from the Windows server that owned it, followed by a re-format.

Figure 10 shows detail from the Microsoft Disk Administrator before and after the shrink operation.

FIGURE 10. SHRINKING A LUN: BEFORE AND AFTER

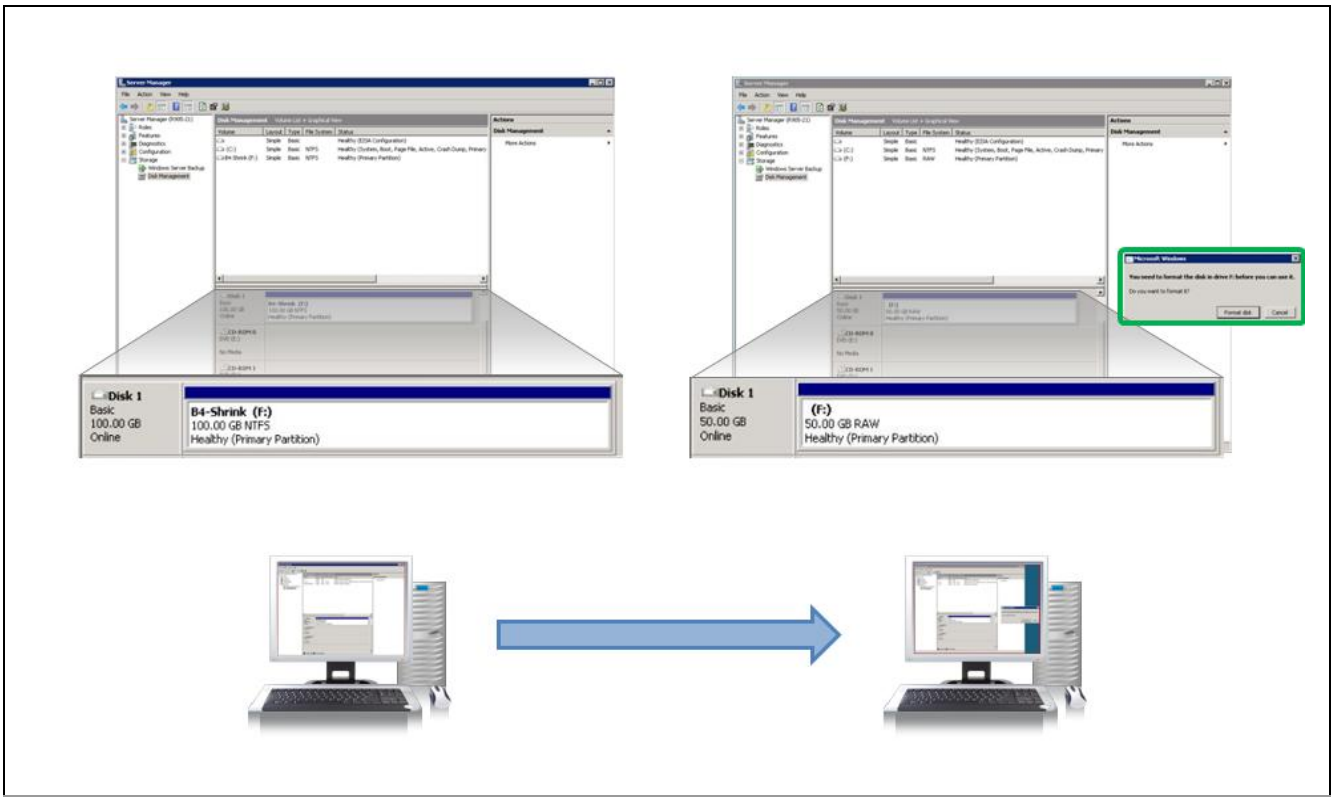

After re-formatting the volume, new data was added and the LUN was expanded from 50 GB to 200 GB. Using the SNM2 console, the process was exactly the same as the volume shrink. When the volume was re-scanned using Windows Disk Management in the MMC Computer Management console, the additional capacity appeared as unused space at the end of the existing volume. ESG Lab was able to take advantage of the additional capacity by expanding the volume using Windows Disk Management as seen in Figure 11.

.

FIGURE 11. EXPANDING A VOLUME

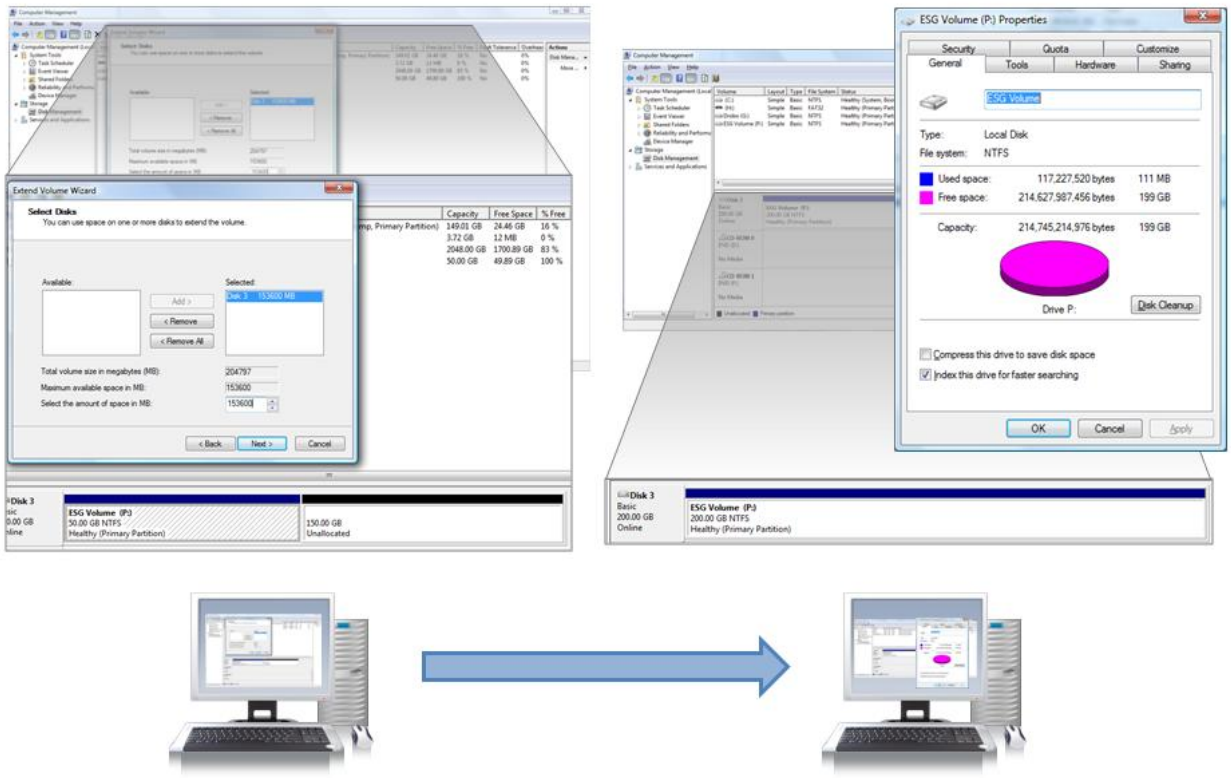

ESG Lab confirmed that all files copied into the volume before expansions were still accessible and intact.

Why This Matters

Modular storage deployments are growing in capacity and complexity within organizations of all sizes. IT managers are increasingly being asked to manage more storage capacity with stagnant, or shrinking, budgets and staffing.

HDS has significantly improved the ease of deployment and manageability of the AMS product line. During this latest round of testing of the AMS 2000 family, ESG Lab confirmed that HDS has loaded the Storage Navigator Modular 2 with usability improvements and made nearly all administrative functionality available through easy to use wizards. This enables customers to take advantage of the advanced architecture of the AMS 2000 family to provide higher service levels to their applications and users while making it easier to manage the system over its lifetime. These usability improvements help end-users to further reduce costs and do more with less.

Simple, Scalable Performance

HDS has made significant changes to the architecture and design of the AMS 2000 series with a goal of improving performance while reducing cost and complexity. The AMS 2000 is the first midrange storage system to offer a high performance SAS back-end and active-active load balancing front-end controllers. SAS (Serial Attached SCSI) is a point-to-point data transfer technology and enterprise SAS drives provide the same reliability and performance as Fibre Channel drives at significantly lower cost. The SAS backplane in the AMS 2000 family can accept both SAS and SATA drives in the same enclosures.

Performance testing began with evaluation of IOPS and throughput characterization tests—often referred to as "the corners" in storage benchmark testing. Corners testing, which provides a good first impression of the theoretical limits of a storage system, was followed by tests designed to measure the performance capabilities of the AMS 2000 family when running a real-world application, Microsoft Exchange. First, ESG Lab audited tests run in HDS performance labs using the industry standard Iometer benchmarking utility.[4] Figure 12 shows the relative performance of the maximum configurations of the previous and current generation of AMS storage systems.

FIGURE 12. AMS FAMILY PERFORMANCE COMPARISON

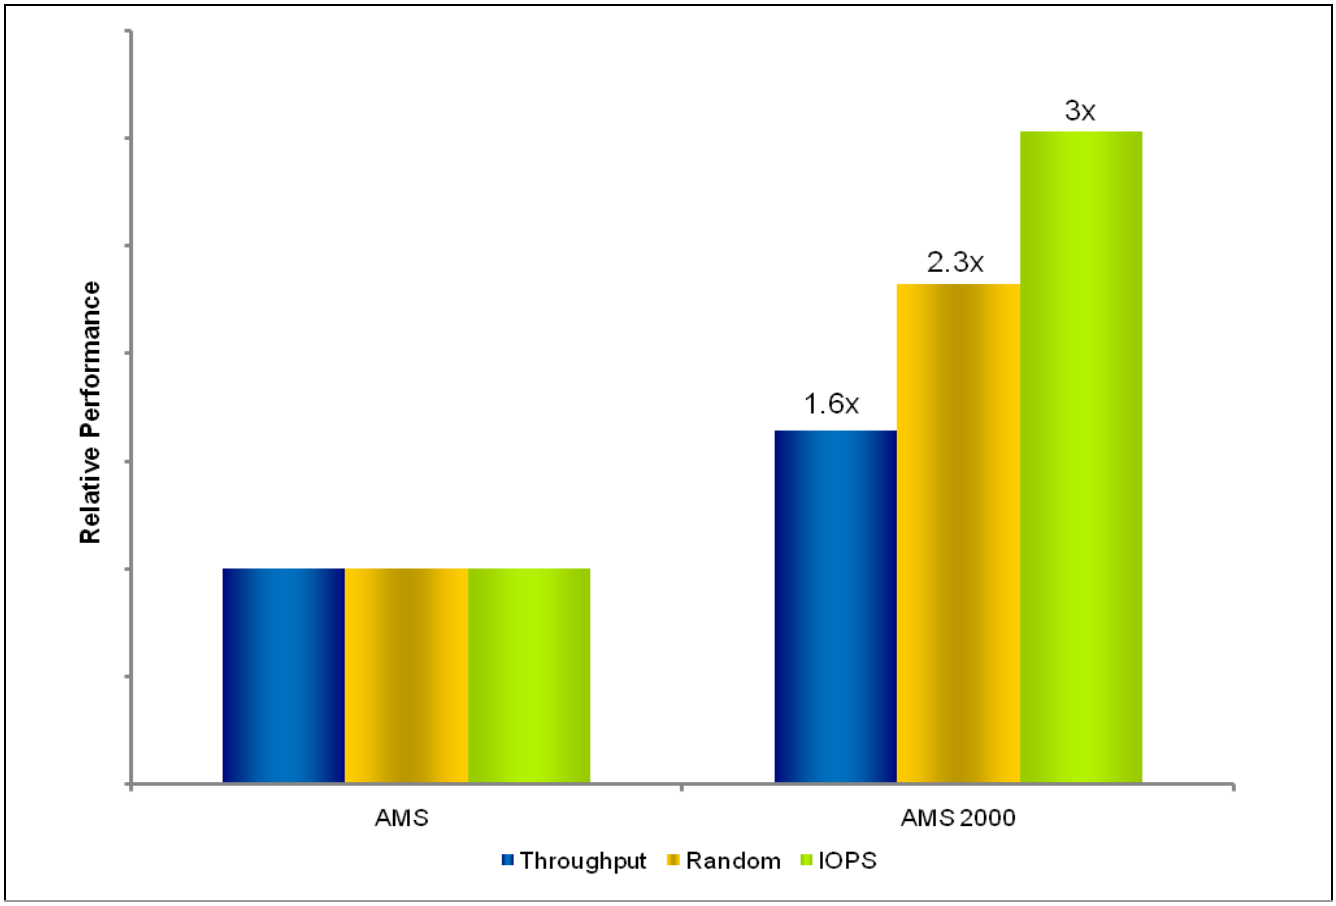

The "corners" tests depicted here measured three industry standard metrics: IOPS, throughput, and random I/O. IOPS (I/Os per second) are commonly used to demonstrate the processing power and front-end efficiency of the disk array, while throughput measures how much data the AMS could move from disk in a given amount of time. The random I/O workload is designed to simulate the I/O pattern that makes up the majority of traffic generated

[4] Iometer is an open-source I/O measurement and characterization utility, available for download at: http://www.iometer.org/

by interactive, response-time sensitive business applications (e.g. databases and e-mail). The chart shows that HDS has significantly raised the bar, with the AMS 2000 family delivering from 1.6 to 3 times the performance headroom of the previous generation.

ESG lab next analyzed AMS family performance in an Exchange environment via an audit of Microsoft ESRP 2.0 results and hands-on testing using Iometer workloads designed to simulate a Microsoft Exchange 2007 environment. The AMS 1000 results were extracted from published Microsoft Exchange Solution Review Program (ESRP) test results, while ESG Lab performed a series of tests against the AMS 2300 to simulate a Microsoft Exchange 2007 workload using Iometer.

ESRP is a Microsoft program designed to facilitate third party storage testing and solution publishing for Exchange Server. The program combines a storage testing harness (JetStress) with publishing guidelines for Microsoft Gold Certified and Storage OEM Partners. ESRP employs the Jetstress utility to create real exchange traffic that runs against real exchange databases—with logging and file attachments—exactly as in the real world. The testing is designed to measure both the performance and reliability of a given solution. The performance test runs for two hours while the reliability test runs for 24 hours. Both tests must run without exceeding a prescribed disk latency threshold (20 milliseconds) and a reliability test is performed to check for database and log corruption at the end of the run. Manufacturers use the ESRP framework to test storage solutions and then submit results to Microsoft for review. Approved solution results are posted on the Microsoft Exchange ESRP website.[5]

FIGURE 13. AMS FAMILY EXCHANGE PERFORMANCE

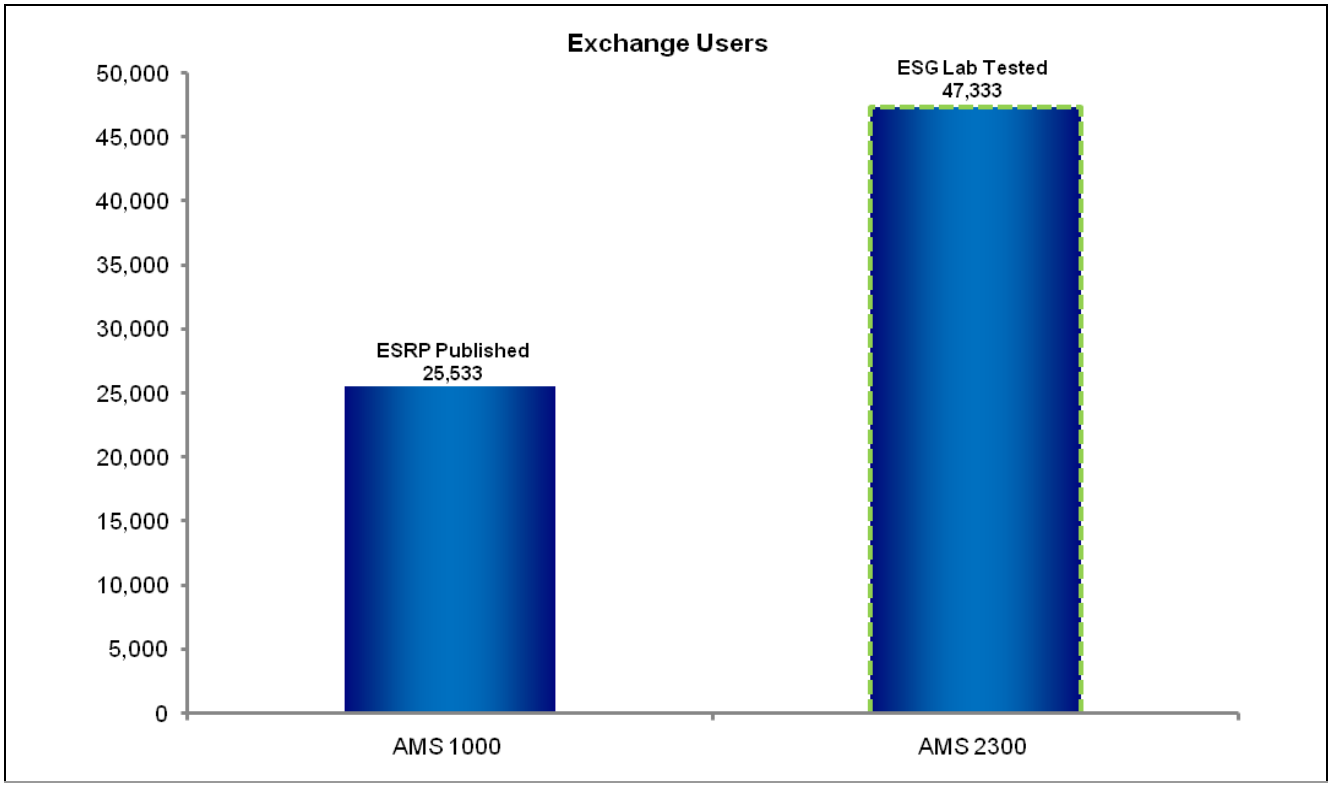

ESG Lab Testing was performed against an AMS 2300 using the same quantity and configuration of disks employed for the ESRP test. Iometer was used to generate a workload similar to the IO generated by JetStress to provide a quick estimate of the number of Exchange 2007 users that the system could support as compared to the previous generation architecture.

[5] http://technet.microsoft.com/en-us/exchange/bb412164.aspx

As seen in Figure 13, the AMS 2300 was able to sustain enough IO to support nearly twice the number of Exchange users than the previous generation using the same quantity and configuration of disks. It's important to note that the ESRP published numbers are the result of a rigorous testing program audited by Microsoft to evaluate many facets of a solution in an exchange environment. The ESG Lab tested results represent a methodology that ESG Lab has developed using our experience in the field to quickly estimate performance in an Exchange environment. ESG Lab is confident that when HDS completes ESRP 2.0 testing for the AMS 2000 family their results will be similar.

ESG Lab constructed the Iometer test bed for the AMS 2300 with the same number of disks, raid protection, and layout as were used to obtain the ESRP results.[6] The number of supported users was derived using a 'Very Heavy' user profile, defined by Microsoft as .48 IOPS per user plus 20% for overhead, or .576 total IOPS per user. Table 1 lists the raw results.

TABLE 1. ESG LAB AND ESRP EXCHANGE PERFORMANCE

Platform	Exchange DB IOPS	Response Time (Read Latency)	User Profile	Number of Users
AMS 1000	14,707	16ms	Very Heavy	25,533
AMS 2300	27,264	16ms	Very Heavy	47,333

What the Numbers Mean

- The AMS 2300 was able to sustain IO sufficient to support 47,333 Exchange users, as projected by ESG Lab.
- Latency was also excellent, remaining below 16 ms throughout the test. This translates to a positive user experience with fewer storage-induced application delays.
- The performance observed by ESG Lab proved that the improved architecture of the AMS2300, including the switch from a FC to SAS back end, improves performance for real-world applications, In this example, ESG has confirmed that the AMS2300 can support up to 85% more Exchange users than the previous generation AMS1000.
- The AMS 2000's active-active architecture made volume configuration and layout for the MS Exchange tests easier and faster than with the previous generation arrays.
- The combination of additional headroom plus excellent application performance shows that HDS has substantially improved the operational efficiency of the system.

Why This Matters

ESG research indicates that performance is a key concern when deploying applications in a highly consolidated environment. With multiple applications relying on a shared infrastructure, there is a concern that performance requirements can't be met affordably. As a matter of fact, 51% of ESG survey respondents that have already deployed virtual servers connected to networked storage report that performance is their top concern.[7]

With architectural enhancements at the storage system level and enterprise class load balancing technology, ESG Lab has verified that the AMS 2000 family can be deployed to cost-effectively provide easy-to-configure storage for mission-critical applications with excellent performance. As customers continue to search for ways to drive down costs, the AMS 2000 family offers a simple, scalable, modular platform with enterprise class functionality designed to enhance performance without increasing complexity.

[6] Storage configuration and layout details can be found in the Appendix.
[7] Source: ESG Research Report, *The Impact of Server Virtualization on Storage,* December 2007.

Symmetric Active–Active controllers

Medium-sized businesses are increasingly being asked to do more with less. Hosting multiple operating systems and applications on a single storage array reduces capital costs for new equipment as well as the operational costs associated with power, cooling, and data center space. A major challenge with midrange storage arrays has been the administrative effort required to manage performance. With traditional midrange storage systems, administrators must carefully map servers to resources to avoid bottlenecks and balance the load evenly across controllers, busses, and disks. This functionality is often automated in enterprise class arrays.

ESG believes that virtualized, networked storage is essential if organizations are going to unlock the full potential of their IT infrastructure. The ability to map physical or virtual machines to storage resources without the need for manual load balancing can significantly reduce the managerial effort required for deployments, implementations, and ongoing administration.

The AMS 2000 family is the only dual controller midrange storage system with symmetric active-active controllers. An AMS 2000 system behaves more like an enterprise class array, providing integrated and automated front-to-back I/O load balancing. Figure 14 compares traditional active-passive controller architecture with the AMS 2000 family. In the traditional architecture, each controller 'owns' a set of LUNs and each server accesses individual LUNs one path at a time. Load balancing is a manual, administrative task where LUNs are distributed between front-end controllers. If I/O is unbalanced, the administrator must manually change ownership of LUNs to the opposite controller until the I/O load is balanced between the two, potentially forcing reconfiguration of the server and SAN to maintain access.

FIGURE 14. HDS AMS ACTIVE-ACTIVE LOAD BALANCING CONTROLLERS

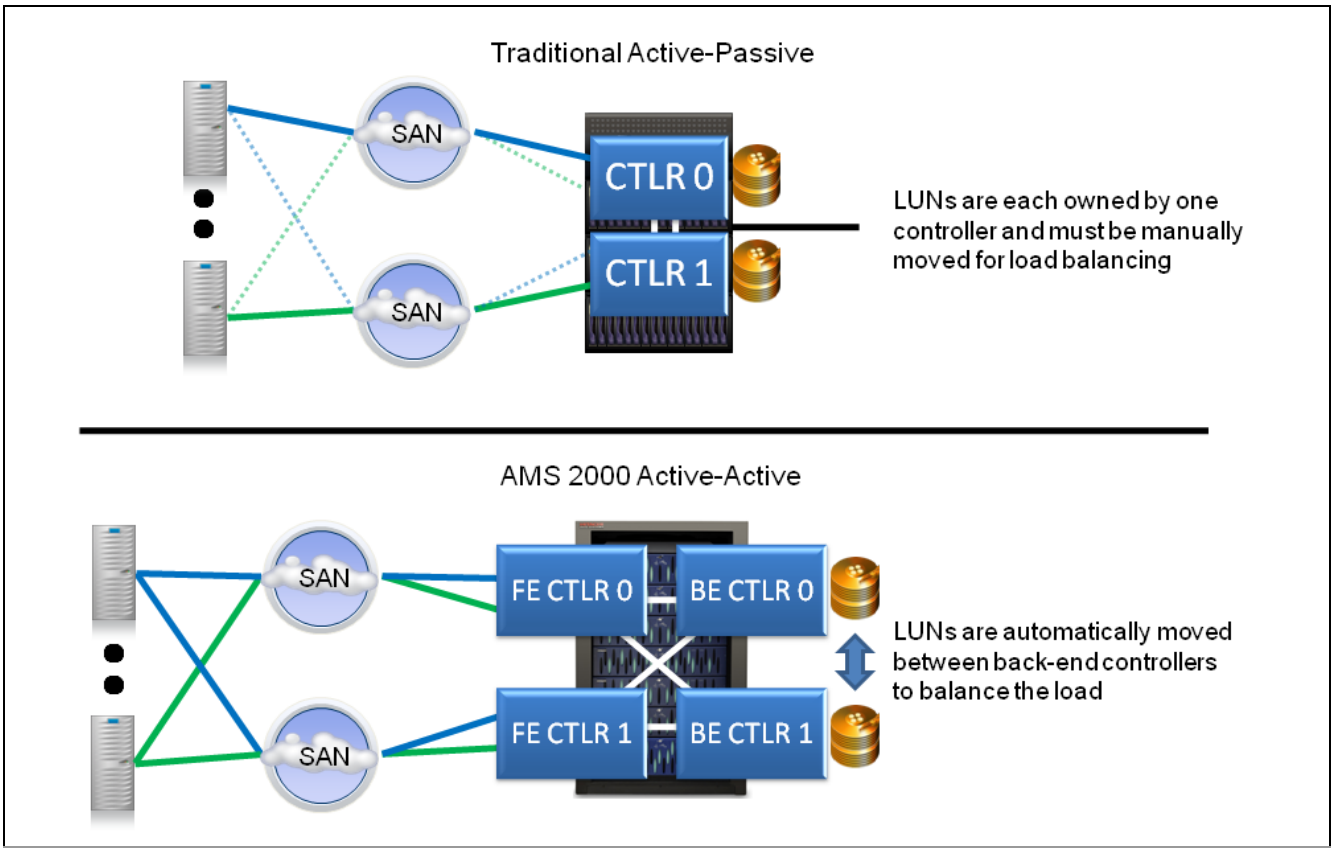

In the AMS 2000 architecture, LUNS are not owned by the front-end controllers and servers can access their LUNs through either controller. This has two immediate benefits: administrators no longer have to assign a LUN

to a particular controller and the server can be attached to the array through both controllers without having to keep track of the 'active' controller for each LUN. Server side load balancing can then use all available paths to each device while the array automatically performs back-end load balancing—internally without disruption.

ESG Lab Testing

ESG Lab evaluated the active-active, load balanced front-end architecture of the AMS 2300 with a single Windows 2008 server SAN attached to the AMS system using two Fibre Channel HBAs. Iometer was used to generate traffic to a single LUN on the AMS 2300 array. A bandwidth intensive 512 KB sequential read workload, similar to a nightly backup job or video streaming application, was run for five minutes. No other I/O was running against the array during this test. As shown in Figure 15, Hitachi Dynamic LUN manager was managing the paths to LUN 0004 on the host. The server was accessing the volume through ports 0D (on controller 0) and 1A (on controller 1).

FIGURE 15. ACTIVE-ACTIVE LOAD BALANCING CONTROLLERS IN ACTION

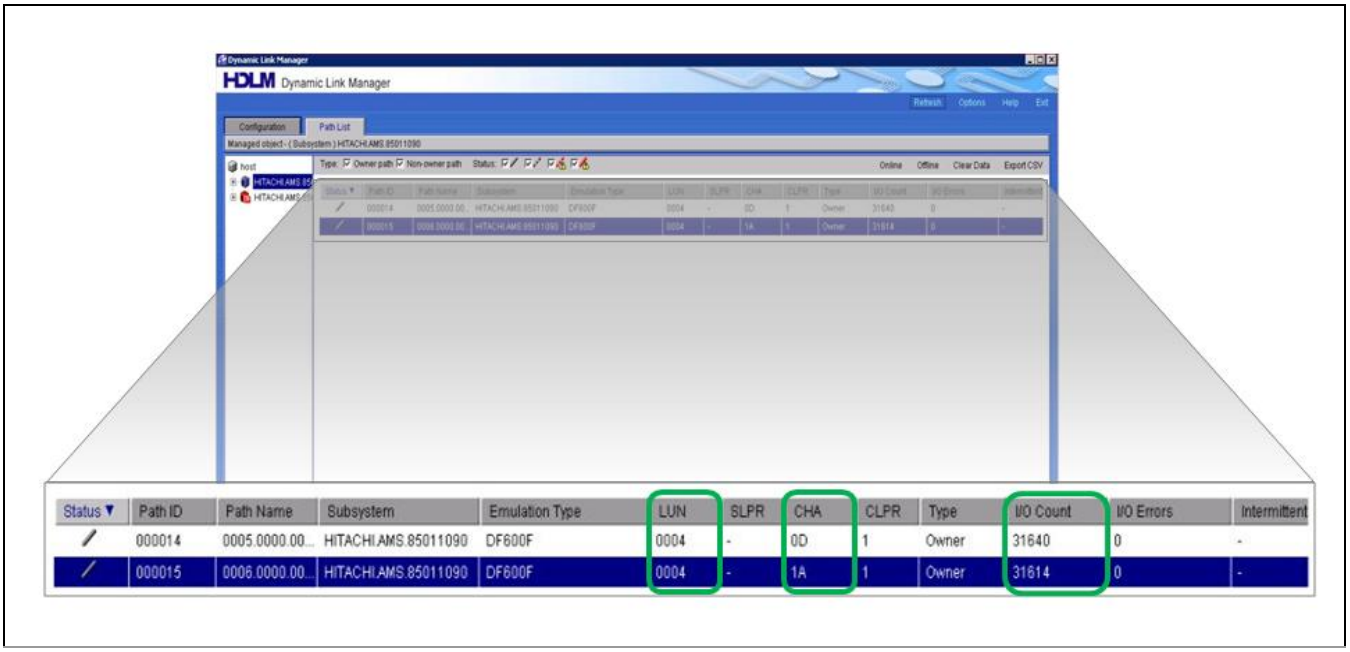

Iometer reported an average of 676 IOPS (338MB/sec) overall for the test run. Figure 16 shows the traffic as seen from the SNM2 Performance Monitor. The two charts show IOPS on Controller 0 and Controller 1. Each controller is reporting between 300 and 350 IOPS over the five minutes of the test run, perfectly in sync.

FIGURE 16. MONITORING LOAD BALANCING FROM THE SNM2 CONSOLE

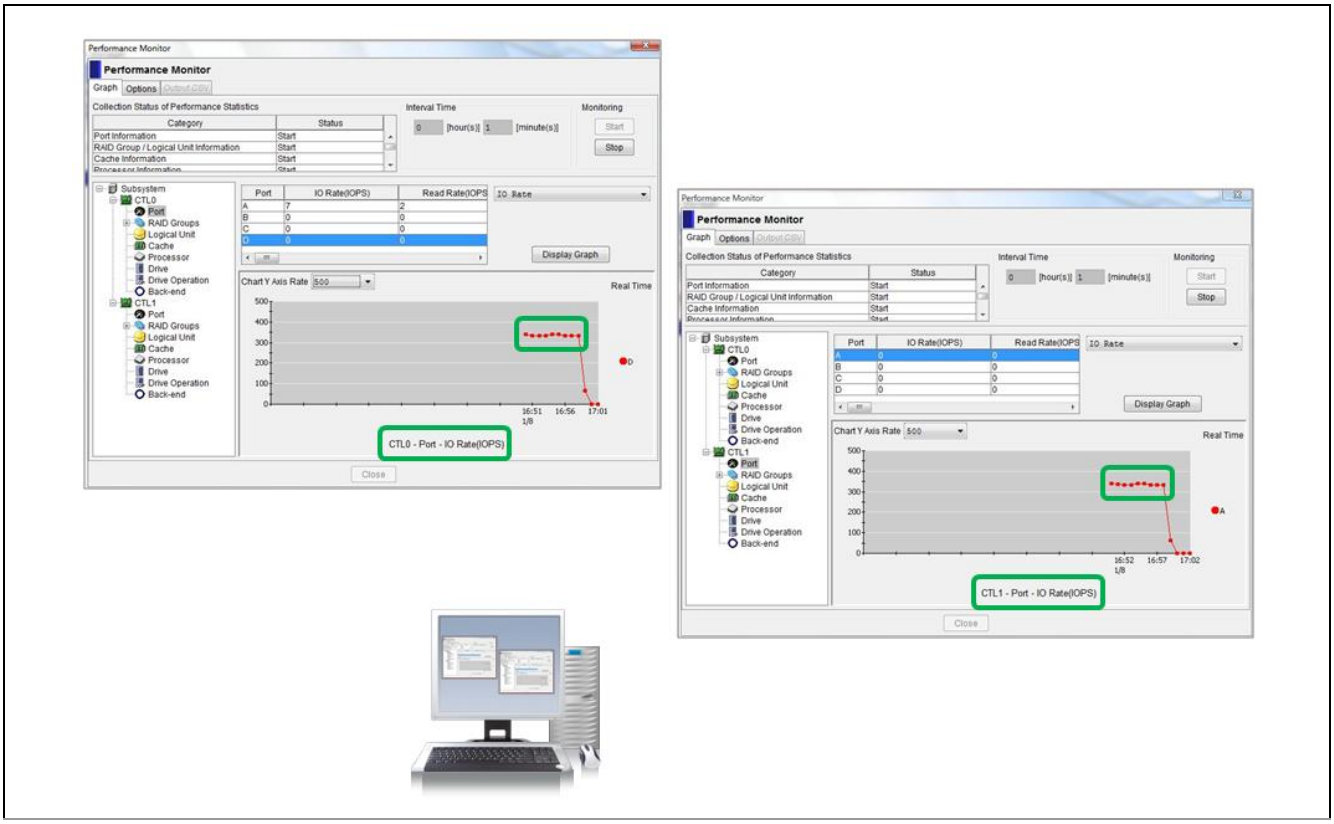

ESG Lab examined the SNM2 Performance Monitor periodically throughout testing. Without exception, traffic was evenly balanced between the front-end controllers during every test.

Why This Matters

ESG research found that more than half of IT professionals surveyed cited performance concerns as among their largest challenges with networked storage when implementing server virtualization and consolidation.[8] A major challenge with midrange storage systems in this space has been the administrative effort required to manage performance. With traditional midrange storage systems, administrators must carefully map servers to resources to avoid bottlenecks and balance the load evenly. The ability to host multiple operating systems and applications on a single storage array can reduce capital costs for new equipment as well as operational costs associated with power, cooling, and data center space.

The ability to map physical or virtual machines to storage resources without the need for manual load balancing can significantly reduce management effort required for deployments, implementations, and ongoing administration. ESG Lab validated that the HDS AMS 2000 family's active-active controller architecture with automatic load balancing provided consistent, stable performance while greatly reducing the effort required for array configuration and LUN mapping. The AMS 2300 automatically load balanced and provided access to LUNS through both controllers simultaneously, a feature normally found in enterprise class arrays.

[8] Source: ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007.

Availability – Remote Replication

The HDS AMS 2000 family provides a broad range of simple and fast volume protection services that enable organizations to provide a highly available storage service to their users across multiple generations of storage systems. Hitachi Copy-on-Write software provides capacity-efficient snapshots, Hitachi ShadowImage Replication software creates point-in-time local clones, and Hitachi TrueCopy software offers synchronous and asynchronous replication.

ESG Lab Testing

ESG Lab tested Hitachi TrueCopy software in synchronous mode between two AMS family storage systems attached to separate SAN domains to simulate an intra-campus disaster recovery strategy. As shown in Figure 17, data in production servers is protected by synchronously mirroring the volumes to a second on-campus data center.

FIGURE 17. LOCAL AND REMOTE REPLICATION

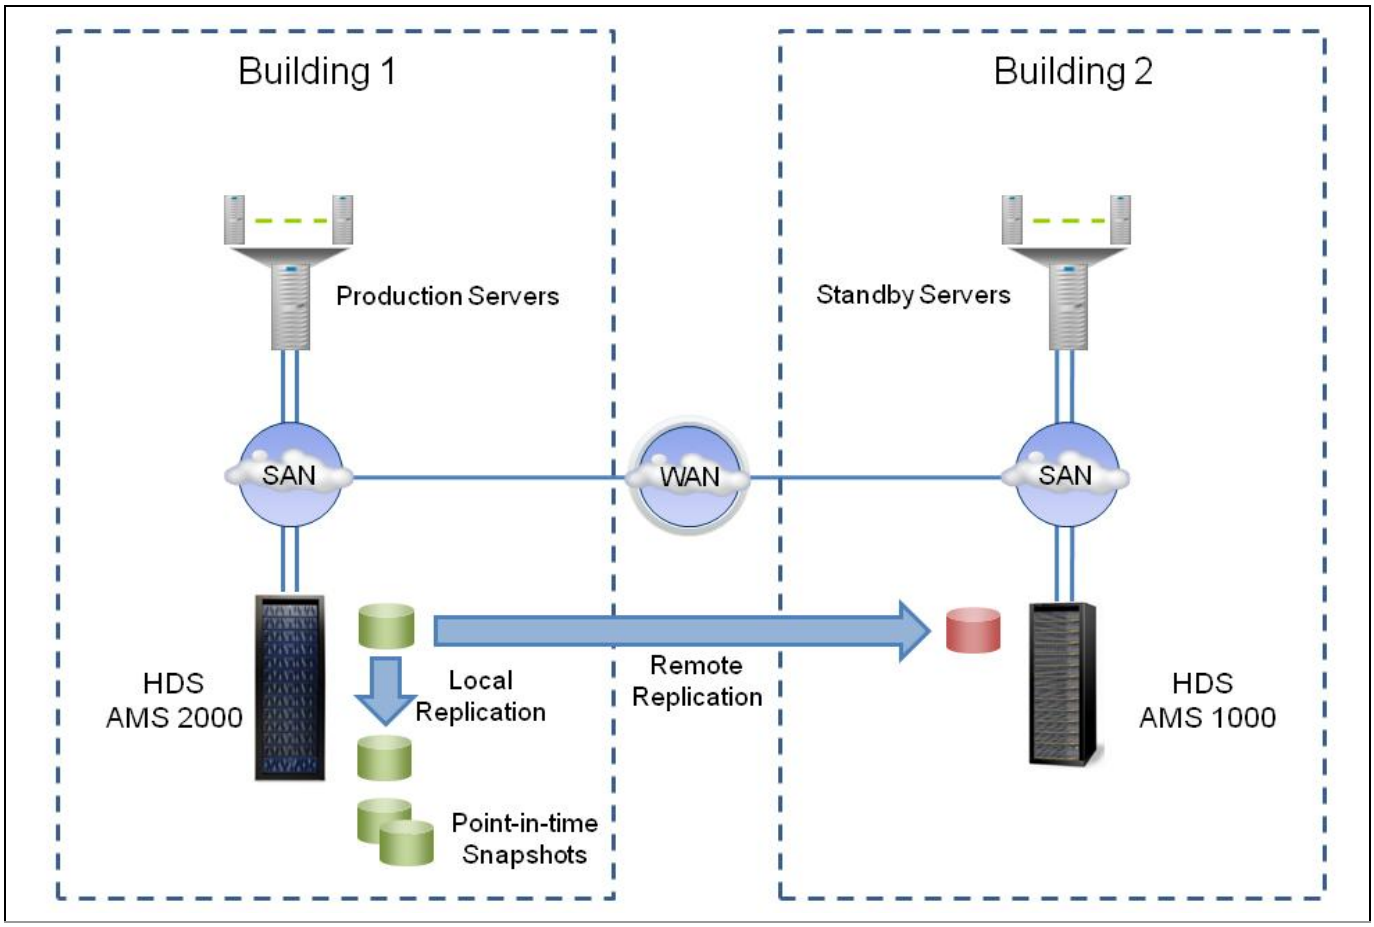

ESG Lab provisioned the 'Production' server, a Windows 2003 virtual machine running under VMware ESX, with two 25GB volumes provisioned on the 'local' AMS 2300 array. Two 25 GB target volumes were assigned on a 'remote' AMS 1000 array that was connected to the 'local' AMS2300 via 4 Gb/sec Fibre Channel. Next, ESG Lab configured synchronous remote mirroring between the two arrays using the SNM2 Create Remote Path Wizard, seen in Figure 18.

FIGURE 18. REMOTE REPLICATION SETUP

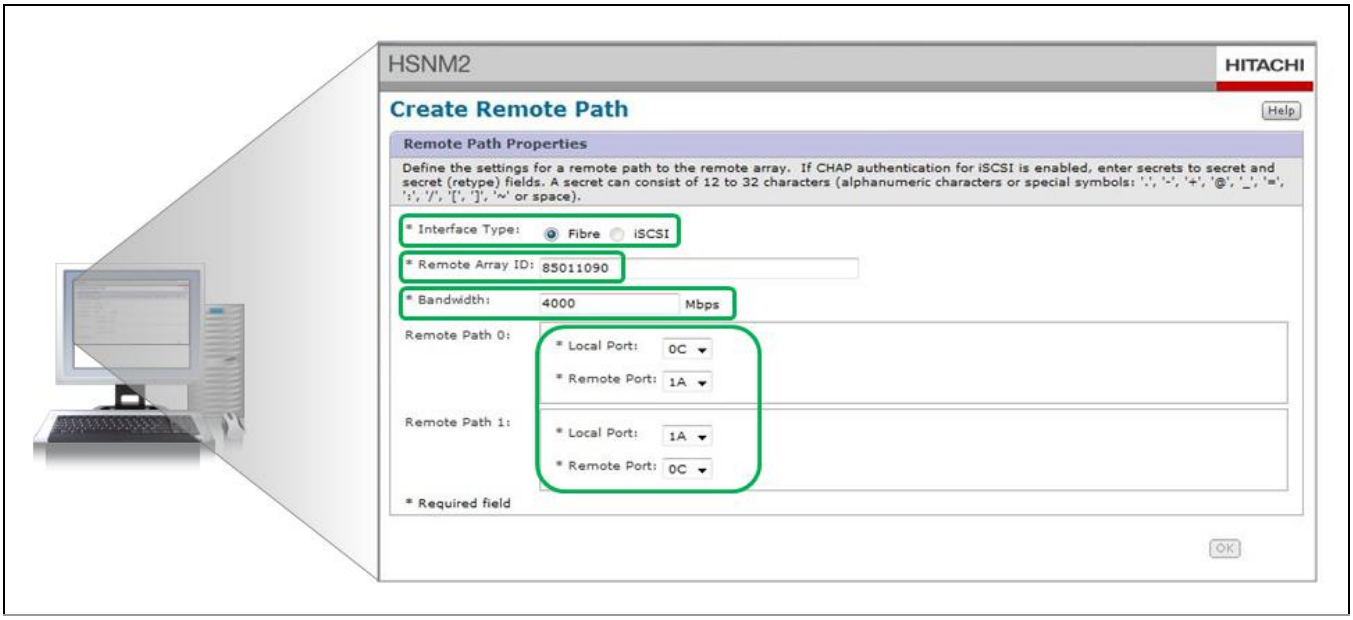

ESG Lab generated I/O on a virtual machine in the local data center by running a continuous 4 KB random write workload using Iometer and by playing a video file stored on the mirrored volume.

Figure 19 illustrates the steady state, with applications running on the source side and volumes being mirrored between arrays. In a full disaster recovery configuration, servers would use clustering technology to fail applications over to the target data center.

ESG Lab verified I/O was being read from and written to the correct volumes on both arrays. The total time, from first click to volume synchronization, was less than two minutes.

The next step was to simulate a major failure. This was done by physically disconnecting the Fibre Channel ports between the two arrays. ESG lab then assigned the volumes to a server attached to the AMS 1000 and verified that the volumes were accessible by playing back the video and copying additional files into the volume.

FIGURE 19. REMOTE REPLICATION IN ACTION

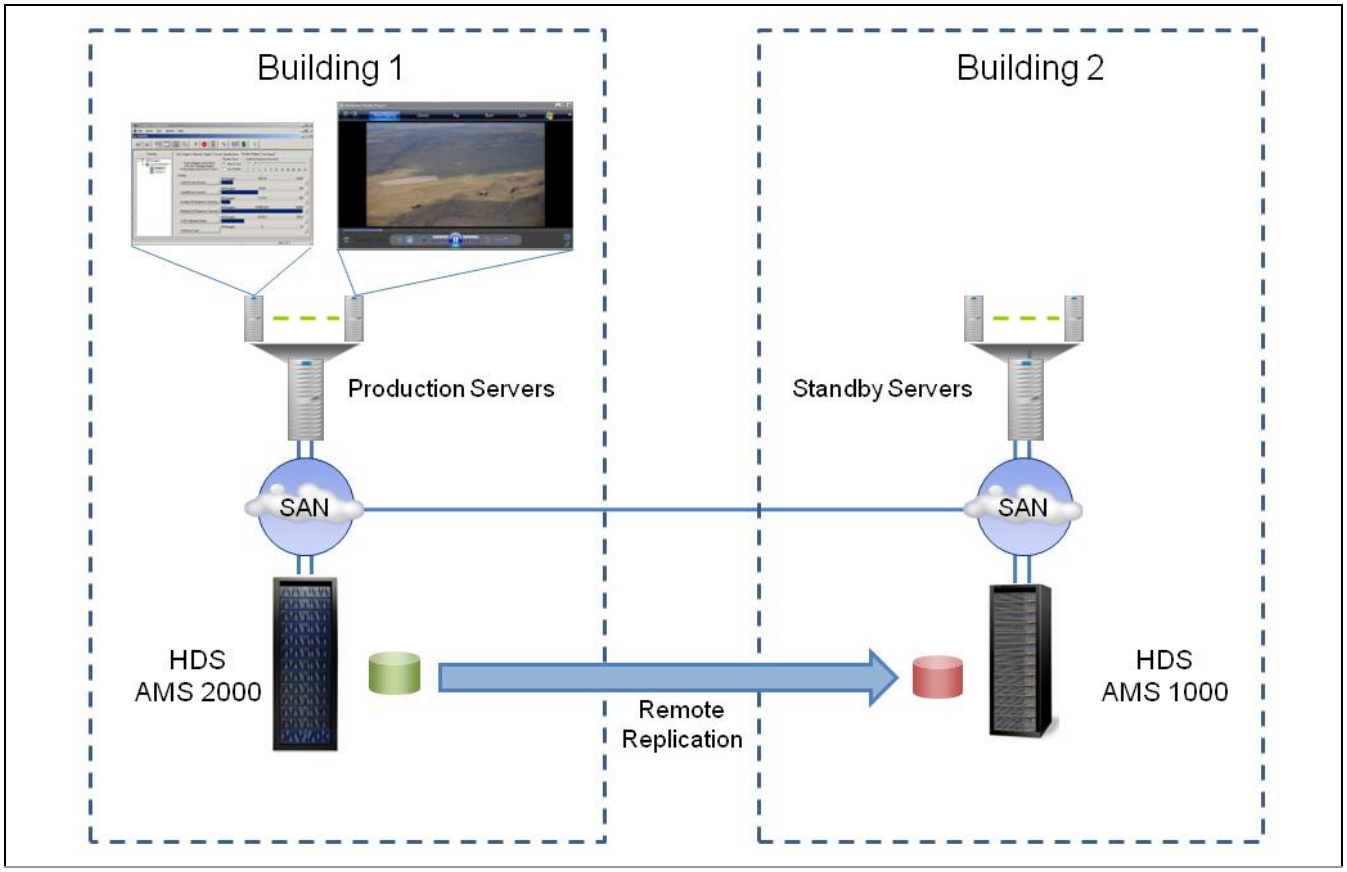

A failback was then performed by reconnecting the FC links between the systems, removing the volumes from the remote server and re-establishing the path between the two arrays. Once connected, the new data was synced back to the production side and ESG Lab verified that all data written during the 'disaster' was propagated back to the production system.

Why This Matters

Organizations replicate business critical applications to a secondary site because they cannot tolerate interrupted access to those applications before, during, or after a major failure in a production data center. In enterprise class environments, storage arrays' remote mirroring capabilities have been used to recover quickly from such disasters.

ESG lab validated that Hitachi's AMS 2000 family's implementation of Hitachi TrueCopy provides storage mirroring services which can be used to implement disaster tolerant applications running on multi-generational storage platforms. Customers could conceivably deploy a smaller, more affordable, or older storage system at a remote site to reduce the total cost of disaster avoidance.

ESG Lab Validation Highlights

- ☑ The Hitachi SNM2 GUI was intuitive and made the AMS200 very easy to configure and manage.
- ☑ The AMS2000 family performed very well, showing up to a 3x performance improvement over previous generations in both engineering-style benchmarks AND real world applications.
- ☑ ESG Lab found the AMS 2000's active-active controller technology with automatic load balancing provided consistent, stable performance while greatly reducing the efforts of array configuration and LUN mapping.
- ☑ In less than two minutes, ESG Lab was able to configure and start Hitachi TrueCopy remote mirroring between two arrays, then simulate a disaster and failover/failback.

Issues to Consider

- ☑ Thin provisioning, a feature offered on Hitachi's higher-end storage systems, is the ability to provision volumes that only consume capacity based on the amount of data actually written to disk. Thin provisioning can be used to effectively eliminate allocated, but unused, capacity and stranded storage. ESG believes supporting thin provisioning in the AMS family would provide more high-end enterprise capabilities for end-users with midrange storage budgets.
- ☑ While Hitachi has scaled up the AMS 2000 family with impressive performance gains over the previous generation, the system supports roughly the same storage capacity as the previous generation ' AMS 1000 family. As more and more organizations turn to modular storage systems for increased consolidation and savings, ESG recommends that HDS consider increasing the total amount of capacity supported per AMS system.
- ☑ The AMS 2000 family supports a range of connectivity options including FC, iSCSI or a mixture of both. FC connectivity is currently supported over 4Gb/sec host connections and iSCSI is supported over 1Gb Ethernet. As organizations continue to increase the number of physical and virtual servers sharing a common pool of modular storage, ESG recommends that HDS consider supporting 8 Gbps FC and 10 Gbps Ethernet host connectivity.

ESG Lab's View

Medium-sized organizations are increasingly being asked to improve information management to more efficiently support their businesses while keeping costs down. Storage technology providers, in turn, have traditionally responded with 'midrange' storage systems, frequently defined as scaled-down and compromised versions of enterprise class arrays. Medium-sized businesses, however, are in need of "enterprise-class" information management and infrastructure technologies—without the enterprise price tag.

HDS is once again delivering innovation—the level of capability and scalability in the Hitachi AMS 2000 family is compelling and extremely valuable to end-users.

The AMS 2000 is a highly reliable, flexible, and scalable family of storage systems designed to enable medium-sized businesses and small enterprises to deliver highly available storage services to their users and customers. The AMS 2000 family provides a series of capabilities that individually provide great value; but when these capabilities are combined, the stakes are raised; offering enterprise class functionality and availability in a modular platform—provided by a leading, world-class storage vendor.

ESG Lab found the AMS 2000 to be easy to set up and manage, providing impressive performance for critical applications like Microsoft Exchange. The AMS 2000's SAS back-end was able to sustain up to 3 times the performance of previous generation models using more expensive Fibre Channel disks. Enterprise class SAS drives, as offered in the AMS 2000 family provide the same reliability and performance as Fibre Channel drives at significantly lower cost. The SAS backplane in the AMS 2300 tested by ESG lab intermixed SAS and SATA drives in the same drive enclosures.

The AMS 2000 also demonstrated enterprise class functionality automatically and transparently; active-active controllers with automatic load balancing allowed for greatly simplified host SAN and LUN provisioning. Availability options are easy to use as well: ESG Lab used Hitachi TrueCopy remote replication technology to construct a synchronously mirrored configuration in minutes.

Hitachi aimed high with the improved scalability and feature-rich redesign of the AMS family of modular storage arrays. The AMS2000 series' combination of enterprise class features with easy to manage midrange usability and reduced operational costs is powerful and worthy of serious consideration by any IT organization being asked—once again—to do 'more with less' In their data center.

Appendix

TABLE 2. TEST ENVIRONMENT AND CONFIGURATION

HDS AMS 2300 – 105x 146 GB SAS drives 60x 1 TB SATA drives 8x 4 GB FC connections	Firmware 0852/A-M
2x Brocade 5000 32 port 4 Gb FC switches	Firmware 5.2.1
VMware ESXi v. 3.5 Servers – 2x Dell PowerEdge R905 4x Quad core Opteron CPU 64 GB RAM 2x QLogic 2462 4 Gb/sec FC HBA	Guest Operating Systems: RedHat Enterprise Linux 5 (64 bit) MS Windows Server 2008

IOMETER		
Version		2006.07.27
Disks		20x 146GB SAS RAID 1+0
Access Specification	**Random/Sequential Distribution**	**Read/Write Distribution**
512 byte Read	100% Sequential	100% Read
4K byte Read	100% Random	100% Read
512 KB Read	100% Sequential	100% Read
Exchange 2007 EDB	90% Random	80% Read
Exchange LOG	100% Sequential	100% Write

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com