

TECHNOLOGY BRIEF

July 1997

Compaq Computer
Corporation

CONTENTS

Introduction	3
Intelligent I/O	
Architecture	3
Optimized Hardware	4
Embedded Intelligence.....	4
Intelligent Array Engines.....	4
Read-Ahead Caching.....	5
Data Availability	5
Fault Prevention.....	5
Fault Tolerance and Recovery	6
I/O Throughput and Data Bandwidth	8
Improved Storage Management	11
Compaq Array Configuration Utility	11
Compaq Insight Manager.....	13
Conclusion	13

EXECUTIVE SUMMARY

Since 1989, Compaq has been a leader in developing array controller technology. Compaq's fifth-generation product, the Compaq SMART-2 Array Controller, provides high availability, high performance, storage fault tolerance, and significant expansion capabilities for centralized computing of large corporations. Now Compaq has developed new array controller technology that provides dependable hardware RAID protection for mission-critical data to businesses and work sites where it has previously been cost prohibitive.

This new technology has been implemented in the Compaq SMART-2SL Array Controller, developed for use by small and mid-sized companies and by remote sites of large companies. This technology brief describes the intelligent I/O architecture of this controller and how it provides fault prevention, fault tolerance, and fault recovery. It explains the level of performance the controller achieves and demonstrates that it outperforms software RAID. It also explains how the controller interacts with two powerful Compaq management tools to improve storage manageability.

Please direct comments regarding this communication to the ECG Technology Communications Group at this Internet address: tech_com@bangate.compaq.com

COMPAQ

NOTICE

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Compaq and Compaq Insight Manager are registered with the United States Patent and Trademark Office.

NetWare and Novell are registered trademarks of Novell, Inc.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1997 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Compaq SMART-2SL Array Controller Technology

First Edition (July 1997)

Document Number 667A/0697

INTRODUCTION

Compaq is a leader in developing array controller technology, and since 1989 has been providing state-of-the-art drive array controllers to meet the storage needs of corporate customers. In early 1996 the company introduced its fifth-generation product, the Compaq SMART-2 Array Controller, to provide high availability, high performance, storage fault tolerance, and significant expansion capabilities for centralized computing of large corporations.

With today's trend for decentralized computing, branch offices and remote sites now access and produce business-critical data that needs equal storage protection. Moreover, the range of companies with large quantities of data to store and back up has grown to include small and mid-size companies. They, too, require high availability and storage fault tolerance to protect swelling volumes of data. Meeting this need for remote sites and small to mid-size businesses is a special challenge, however, because they are highly cost-sensitive operations that require a solution tuned to their price/performance/capacity needs.

Compaq has developed new array controller technology to provide RAID (Redundant Array of Inexpensive Disks) protection for mission-critical data to businesses and work sites where cost sensitivity previously kept it out of reach. This new technology, based on industry-leading Compaq SMART-2 technology, has been implemented in the Compaq SMART-2SL Array Controller product. This new controller makes proven, dependable, hardware RAID technology available to operations that previously could not afford it.

INTELLIGENT I/O ARCHITECTURE

The Compaq SMART-2SL Array Controller is based on proven Compaq SMART-2 technology. It has an advanced intelligent I/O architecture designed specifically to provide economical, high-performance hardware RAID protection for small- to medium-sized companies and for remote sites of large companies. This architecture consists of optimized hardware and firmware components working together efficiently, with functions and embedded intelligence distributed across appropriate points in the design. See Figure 1.

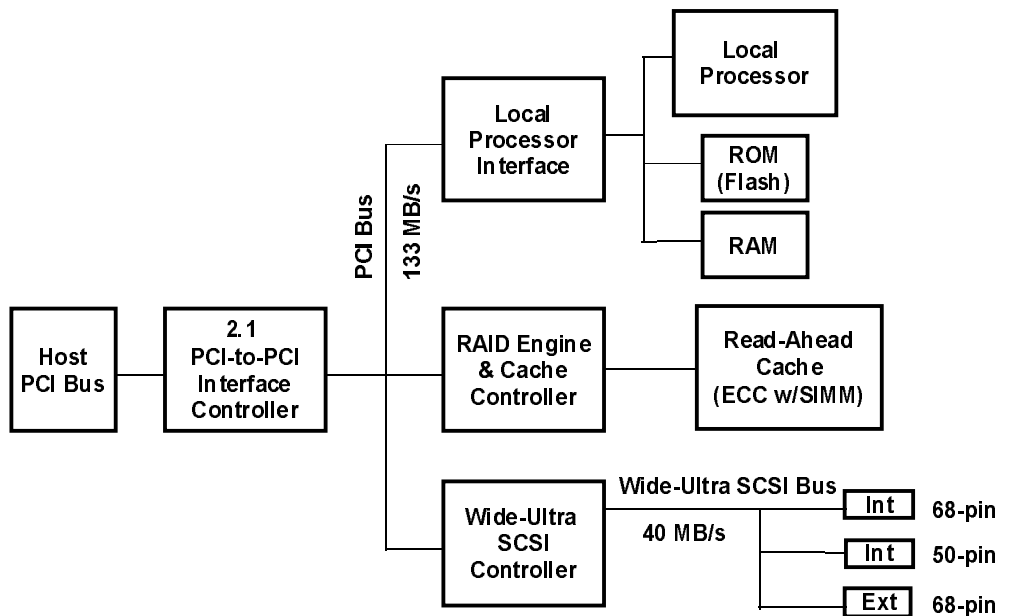


Figure 1. Block diagram of Compaq SMART-2SL Array Controller architecture.

Optimized Hardware

The host interface of the SMART-2SL Controller is a low-latency, 32-bit PCI bus master. It transfers data directly to system memory at rates of up to 133 MB/s (maximum burst rate). Because the industry is moving to PCI and because the bandwidth of EISA controllers is limited to 33 MB/s, Compaq is producing only PCI models of this controller.

The SMART-2SL Controller has a single Wide-Ultra SCSI-3 channel supported by a Symbios 53C875 RISC-based SCSI-3 protocol chip capable of running 32-bit Wide-Ultra SCSI-3 at 40MB/s. The SCSI protocol chip itself contains an embedded RISC processor that operates independently from the local processor. The protocol chip is capable of executing SCSI commands from start to finish with only a single interrupt to the local processor. It is also capable of posting interrupts "on the fly," allowing the local processor to service multiple completions per interrupt and greatly reducing overhead on the SCSI bus.

The SMART-2SL Controller supports up to 32 logical drives. Its Wide-Ultra SCSI-3 channel supports up to seven 9.1-GB drives for a maximum storage capacity of 63 GB. Attaining the peak transfer rate of 40 MB/s on the Wide-Ultra SCSI-3 channel requires use of Wide-Ultra SCSI-3 hard disk drives, drivers, servers, and storage subsystems. For backward compatibility, the SMART-2SL Controller supports 16-bit Fast-Wide SCSI-2 disk drives or 8-bit Fast SCSI-2 disk drives. Where Wide-Ultra SCSI-3 is not supported for external storage systems, the SMART-2SL Controller will default to Fast-Wide SCSI-2.

Embedded Intelligence

The SMART-2SL Controller is effectively a computer system on a board. The local processor is a 32-bit RISC processor from Advanced Micro Devices (Model AM29045-50), an industry-standard part that has been optimized for high-performance, embedded applications. It provides approximately twice the processing power of the local processor on the original Compaq SMART Controller. This increased power improves I/O throughput. The local PCI bus also provides high bandwidth (133 MB/s burst rate) and low latency.

The intelligence of the SMART-2SL Controller comes from ROM-resident firmware that is executed by the local processor. At power-on time, the firmware code is copied from flash ROM into high-speed RAM, from which instructions are fetched and executed by the processor. A small and efficient real-time OS provides basic task and control functions for the embedded environment. Algorithms have been tuned for optimal I/O performance. The host interface protocol has also been optimized. OS device drivers running on the host can submit hundreds of commands at once to the SMART-2SL Controller. Similarly, the controller allows the host driver to retrieve multiple command completions under a single status interrupt. The firmware is flashable (that is, software upgradable) for convenient field upgrades.

Intelligent Array Engines

On the SMART-2SL Controller, processing tasks are divided between two intelligent array engines. These engines are Compaq-developed ASICs (application-specific integrated circuits). The two engines work simultaneously, one generating fault tolerance information and managing data flow, while the other prepares and sorts array storage commands.

The Compaq-developed RAID engine optimizes RAID 5 performance. This is a hardware-based solution that employs an automatic read-XOR-write cycle to generate parity information as data is being written to the drives. It allows parallel execution of multiple XOR operations for different I/O threads, which eliminates resource contention for a single parity generator. It also allows XOR operations to occur as soon as data is written from the host to the controller. This eliminates data staging and greatly reduces the amount of memory required for parity generation. The RAID engine also allows automatic data recovery and automatic drive rebuild operations to execute very quickly. Because these

operations are performed in hardware, no local processor intervention is required. The processor is free to perform other vital controller functions.

Read-Ahead Caching

The SMART-2SL Controller architecture includes a 6-MB read-ahead cache. The controller uses an intelligent, read-ahead algorithm that can anticipate data needs and reduce wait time. It can detect sequential read activity on single or multiple I/O threads and predict that sequential read requests will follow. It can read ahead, or pre-read data, from the disk drives before the data is actually requested. When the read request does occur, it can then be serviced out of high-speed cache memory at microsecond speeds rather than from the disk drive at millisecond speeds.

This adaptive read-ahead scheme provides excellent performance for sequential small block read requests. At the same time it creates no penalty for random read patterns, because read-ahead is disabled when non-sequential read activity is detected. Thus, the SMART-2SL Controller overcomes the problem of some array controllers in the industry today that use fixed read-ahead schemes that increase sequential read performance but degrade random read performance.

DATA AVAILABILITY

The SMART-2SL Controller provides the highest levels of data availability for demanding business requirements. The intelligence of the SMART-2 Controller allows it to monitor I/O activity, track key parameters, predict potential problems, take corrective actions, provide automatic fault recovery, and deliver full fault management.

Fault Prevention

The SMART-2SL Controller can detect potential hard disk drive failures before they occur so that preventive action can be taken. Three features of the controller provide this capability: drive parameter tracking, dynamic sector repair, and environment tracking.

Drive Parameter Tracking with Pre-Failure Warranty

The SMART-2SL Controller includes a drive parameter tracking feature that monitors more than 15 parameters of drive operation as a background task. If a drive fails to meet certain factory-preset criteria, the controller notifies Compaq Insight Manager,¹ which then alerts the system administrator to the potential problem. Compaq disk drives that fail to meet the expected criteria may be eligible for replacement under the Compaq Pre-Failure Warranty.

Dynamic Sector Repair

Under normal operating conditions, even disk drive media that was initially defect free can develop defects caused by variances in the high-precision drive mechanisms over time. Development of such defects is a common phenomenon in the industry. To protect data from media defects, Compaq built a dynamic sector repair feature into the SMART-2SL Controller.

During normal operation, the SMART-2SL Controller performs a background surface analysis (Auto Reliability Monitoring, or ARM) by continually scanning all drives for media defects. If a defect is found during ARM or during normal operation, the controller automatically remaps the bad sector: It moves data from the defective sector or its regenerated equivalent to a new storage location, marks the defective sector as bad, and no longer writes data there.

¹ Compaq Insight Manager is discussed later in this document.

Environment Tracking for ProLiant Storage System

The SMART-2SL Controller also monitors fan and cabinet temperature of the ProLiant Storage System or ProLiant Storage System/F in which it resides. If either temperature exceeds a preset value, the controller notifies Insight Manager, which issues an alert to the system administrator.

Fault Tolerance and Recovery

The SMART-2SL Controller architecture provides hardware RAID protection of data integrity and availability. Prior to any possible drive failure, the controller proactively generates RAID parity data so that it can keep all data available and the server running during replacement of any failed drive. Moreover, the controller supports quick, reliable fault recovery through two hardware features: on-line spares and hot-pluggable drives.

RAID Support with Automatic Data Recovery

Redundant Arrays of Independent Disks (RAID) was an idea originally conceived at the University of California at Berkeley to improve disk subsystem reliability and performance. It is a method of distributing data redundantly across multiple disk drives such that user data can be recovered or regenerated in the event of a single disk drive failure. Compaq was the first company to introduce RAID subsystems into the network server marketplace in 1989. Today RAID is an industry standard, and most on-line network data storage is protected with some level of RAID.

The value of RAID comes from the ability of the SMART-2SL Controller to perform automatic data recovery in the event of a disk drive failure. In any fault tolerant RAID configuration, if a drive fails and a read request is received for the missing data, the controller will automatically regenerate missing data from parity information.

The SMART-2SL Controller supports² the standard RAID levels 0, 1, and 5 illustrated in Figure 2.

RAID 0 provides no data redundancy; therefore, it provides no fault protection. Data is simply striped across all drives in the array. This configuration provides high performance at low cost but also incurs high risk.

RAID 1 (referred to as drive mirroring) simply creates two copies of all the data, each copy residing on a separate disk drive. This method has performance equal to or better than RAID 0 but doubles the disk drive cost, since it requires twice as much space to store the same amount of data. RAID 1 can protect against data loss during multiple drive failures as long as the failing drives are not mirrored to each other.

RAID 5 (referred to as distributed data guarding) is the most popular RAID configuration. It provides cost-effective fault tolerance with high performance because parity information is distributed across all drives in the array. The amount of disk space used for redundancy is equivalent to the capacity of a single drive.

² The Compaq Array Configuration Utility will allow customers to configure a drive array for RAID 4, which provides parity checking but with slightly lower performance than RAID 5. However, Compaq does not support RAID 4 with the SMART-2SL Controller.

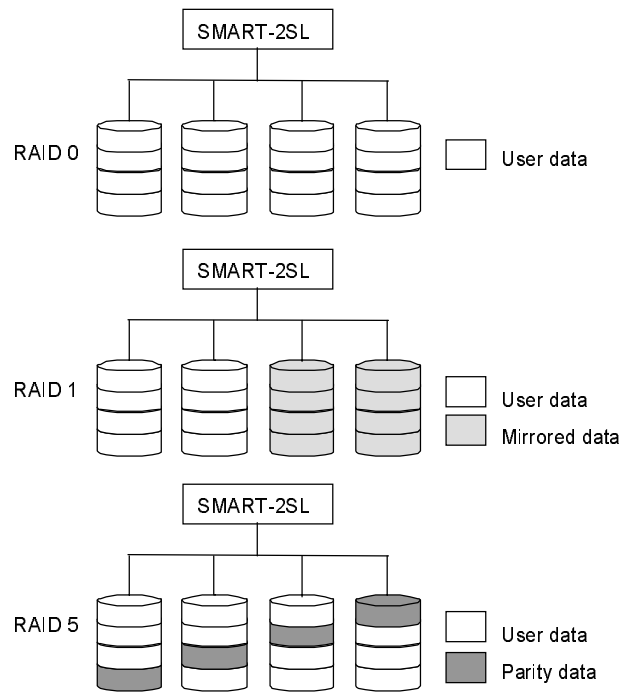


Figure 2. RAID levels supported by Compaq SMART-2SL Array Controller

Hot Plug Drive Support with Automatic Rebuild

The SMART-2SL Controller supports SCSI hot plug drives. If the storage subsystem contains hot plug drives, users can insert or remove drives from fault tolerant configurations while the system is still up and running.

The SMART-2SL Controller detects when a failed drive is removed and replaced. Then from parity information, it automatically rebuilds the data from the failed drive onto the replacement drive. After the rebuild operation is complete, data can again be read directly from the drive and no longer needs to be regenerated.

In general, the time required for a rebuild operation is approximately 15 minutes per gigabyte. That is half the time required by the original Compaq SMART Controller. However, actual rebuild time is dependent upon the amount of I/O activity occurring during the rebuild operation, the number of drives, and the disk drive speed. For example, in an idle system with three 9-GB drives, a RAID 5 rebuild requires approximately 10 minutes per gigabyte.

During configuration the system administrator can set the priority of rebuild operations to high, medium, or low:

- High priority means that the rebuild operation is given top priority. Service of I/O requests continues but at a lower priority.
- Medium priority means that the rebuild operation and I/O requests are given equivalent priority.
- Low priority means that the rebuild operation is given lowest priority. Rebuild occurs, but I/O requests are granted top priority.

During intervals of no I/O activity, the controller is fully dedicated to the rebuild operation, regardless of priority setting.

The priority set during configuration can later be changed on line by means of the Compaq Array Configuration Utility, which will be described later in this document.

On-Line Spares

The SMART-2SL Controller enables users to designate up to four disk drives in an array as on-line spares. During system operation, these spare drives remain up and running but not active; that is, no I/O operations are performed to them during normal array operation. Spare drives are held in reserve in case one or more of the active drives should fail.

An on-line spare (Figure 3) differs from a parity drive. Parity drives are active or involved with all I/O operations. On-line spares power up and run, but they remain in standby mode until needed. If an active drive fails during system operation, then the SMART-2SL Controller automatically and immediately begins a data rebuild operation onto a spare. The administrator does not have to intervene manually. Once the rebuild operation is complete, the system is again fully fault tolerant. The failed drive can be replaced at a convenient time. Once the administrator installs a replacement drive, the controller will restore data to that replacement drive. The original on-line spare will return to standby mode and again be available as an on-line spare.




Figure 3. RAID-5 with an on-line spare drive

I/O THROUGHPUT AND DATA BANDWIDTH

Real user applications depend on both I/O throughput and data bandwidth for optimizing performance. I/O throughput is the rate at which I/O commands are processed. It is usually measured in I/Os per second. Data bandwidth is the rate at which data is transferred from source to destination (for example, from disk to system memory during read operations). It is usually measured in megabytes per second.

In typical on-line transaction processing environments, I/O throughput is the more important performance factor because it determines the rate at which random data base requests can be serviced. In typical image processing environments, data bandwidth is the critical factor because it determines the rate at which large sequential blocks of data can be transferred. In both these examples and in most other situations, end user performance is fully optimized, however, if *both* I/O throughput and data bandwidth are maximized.

The SMART-2SL Controller optimizes I/O throughput without penalizing data bandwidth, and vice versa. This means users can achieve excellent performance in all environments—applications with random I/O or sequential I/O, applications with small block size transfers or large block size transfers, and applications with any combination of I/O activity.

Because of its streamlined architecture and Wide-Ultra SCSI-3 channel, the SMART-2SL Controller provides excellent data bandwidth for small- to medium-sized organizations. Figure 4 compares performance capabilities of the SMART-2SL Controller and the Compaq SMART-2DH Array Controller in terms of data bandwidth measured in a benchmark test.

In this NetBench 5.0 test, the two controllers were installed in identical Compaq ProLiant 800 Servers. Each system contained the maximum number of drives that its array controller could support: 7 Wide-Ultra SCSI-3 drives for the SMART-2SL Controller and 14 Fast-Wide SCSI-2 drives for the SMART-2DH Controller. As Figure 4 indicates, the SMART-2SL provides excellent data bandwidth for its target market of remote sites and small to mid-sized networks. Because it was outperformed by the SMART-2DH Controller by an average of 40 percent, however, the SMART-2SL Controller is not the preferred solution for large networks.

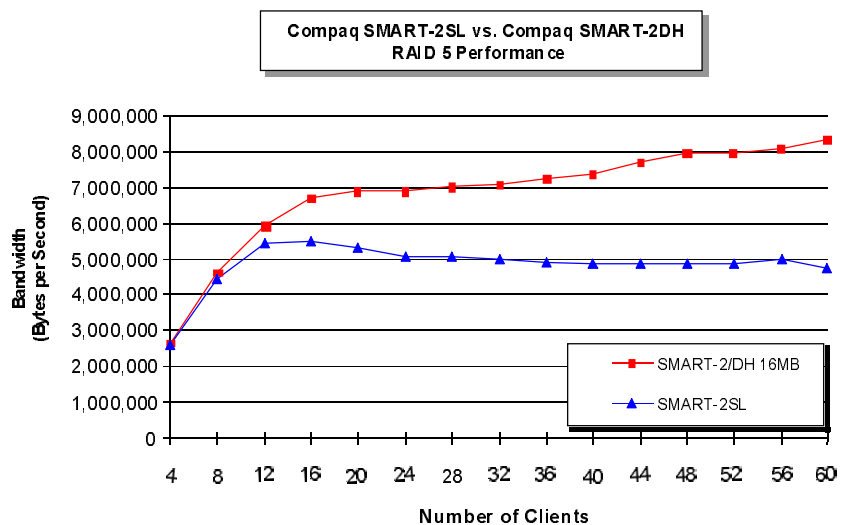


Figure 4. Results of comparative performance test of RAID 5 configurations run with NetBench 5.0 under the Novell NetWare 4.1 operating system

The SMART-2DH Controller is Compaq's mainstream server RAID product. It provides high availability, capacity, and expandability for network servers in data centers and other environments that are not highly cost sensitive. Table 1 compares key features of the SMART-2SL and SMART-2DH Controllers. For more information about the SMART-2DH Controller, refer to the technology brief titled *Compaq SMART-2 Array Controller Technology*.

Figure 5 compares performance of SMART-2SL hardware RAID against that of software RAID provided by Novell NetWare. With software RAID, the RAID functions are performed through software; and SCSI I/O requests are serviced by the standard embedded C875 SCSI Controller. In this NetWare NetBench RAID 1 performance test, the SMART-2SL Controller outperformed the NetWare Software RAID by an average of 36 percent.

TABLE 1: COMPARISON OF KEY FEATURES IN COMPAQ SMART-2SL AND SMART-2DH ARRAY CONTROLLERS

Feature	SMART-2SL	SMART-2DH
Target market	Small- and mid-sized businesses or remote sites of large businesses	Mainstream network servers in data centers or other environments that are not highly cost sensitive
Host interface	32-bit PCI bus master	32-bit PCI bus master
I/O bus(es)	Single Wide-Ultra SCSI-3 channel	2 Fast-Wide SCSI-2 channels
I/O bus speed	40 MB/s	20 MB/s per channel
No. physical drives supported	7 maximum per controller	14 maximum per controller
No. logical drives supported	Up to 32 per controller	Up to 32 per controller
Hot plug drive support	Yes	Yes
No. on-line spares supported	Up to 4 per array	Up to 4 per array
On-line capacity expansion	No	Yes
Enhanced Array Accelerator	No	Yes
Cache	6 MB read-ahead	16 MB write-back and read-ahead
Dynamic Sector Repair	Yes	Yes

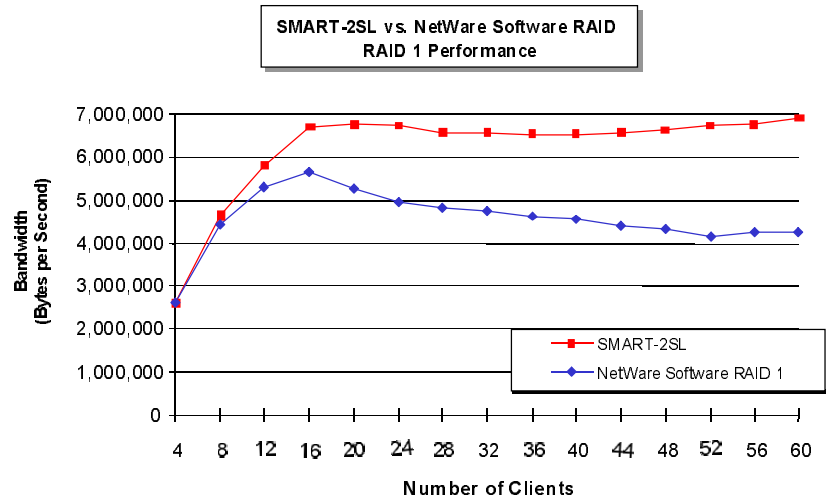


Figure 5. Results of comparative performance test of hardware and software RAID 5 configurations run with NetBench 5.0 under the Novell NetWare 4.1 operating system

IMPROVED STORAGE MANAGEMENT

Improving manageability means making it easier for IS organizations to manage their network data storage. Teaming the SMART-2SL Controller with two proven Compaq management programs significantly improves storage manageability.

Compaq Array Configuration Utility

The Compaq Array Configuration Utility is a powerful yet flexible storage management tool. This easy-to-use, GUI-based³ utility simplifies array configuration. It contains intuitive on-line configuration wizards that provide guided paths for complete array configuration. It also contains customized paths for administrators needing complete control over configuration options. This utility provides an interface to intelligent features of the SMART-2SL Controller that make it easy for administrators to configure, modify, expand, manage, and monitor their storage.

The concept of an array has a special meaning in the Compaq Array Configuration Utility. An *array* is defined as a set of physical disk drives that may be subdivided into logical drives striped across all disks in the set. There can be from 1 to 32 logical drives per controller.

Capacity expansion is done at the array level, not at the logical drive level. In most cases, all disk drives attached to a controller should be grouped together into a single array. This provides the most efficient use of RAID redundancy.

By means of the Compaq Array Configuration Utility, physical drives can be assigned to an array (Figure 6); and up to four drives can be assigned to become on-line spares.

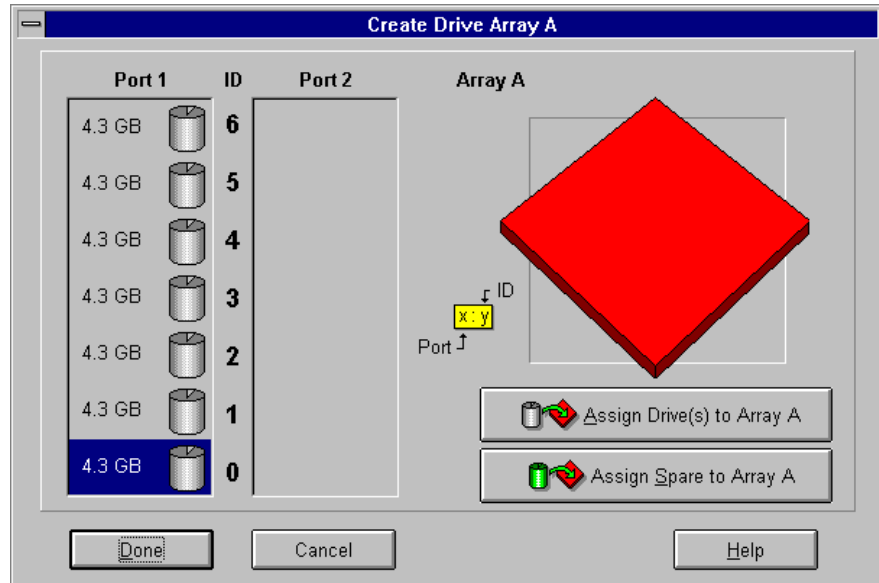


Figure 6. Creating a drive array using the Compaq Array Configuration Utility

The Compaq Array Configuration Utility provides an intuitive interface for creating logical drives within the arrays (Figure 7). Fault tolerance levels can be selected on a logical drive basis.

³ GUI is an acronym for Graphic User Interface.

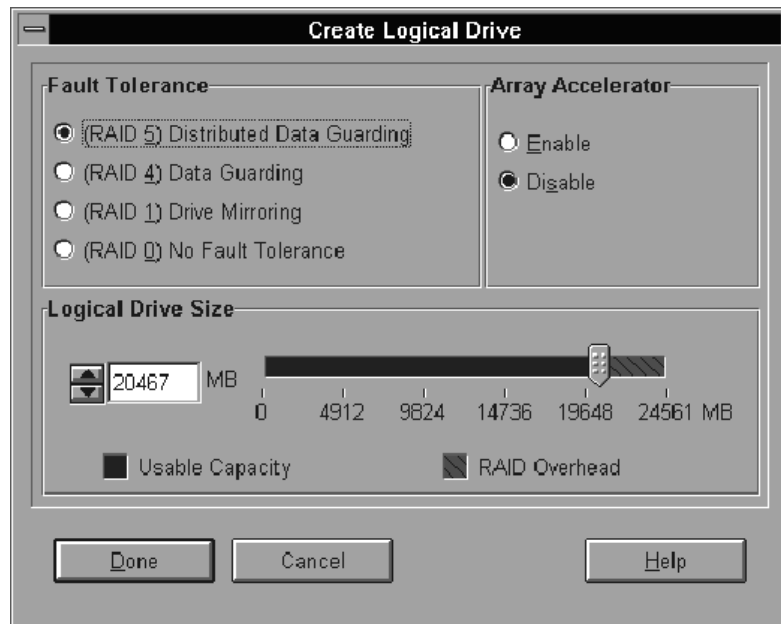


Figure 7. Creating a logical drive using the Compaq Array Configuration Utility

Graphical views of the array (Figure 8) and the logical drives make it easy to visualize how the storage is configured. The Compaq Array Configuration Utility steps through each logical drive configuration until all space is allocated.

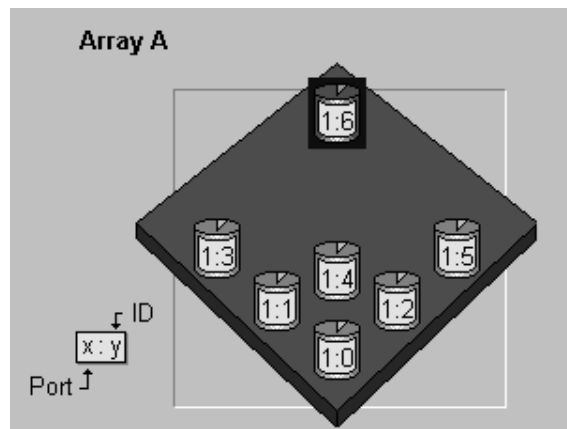


Figure 8. The Compaq Array Configuration Utility provides graphical views of arrays being built.

Custom configuration paths allow the administrator to examine the details of both logical and physical views and make specific configuration changes if necessary (Figure 9).

For more information about the Compaq Array Configuration Utility, refer to the *Compaq SMART-2SL Reference Guide*, P/N 295196-001.

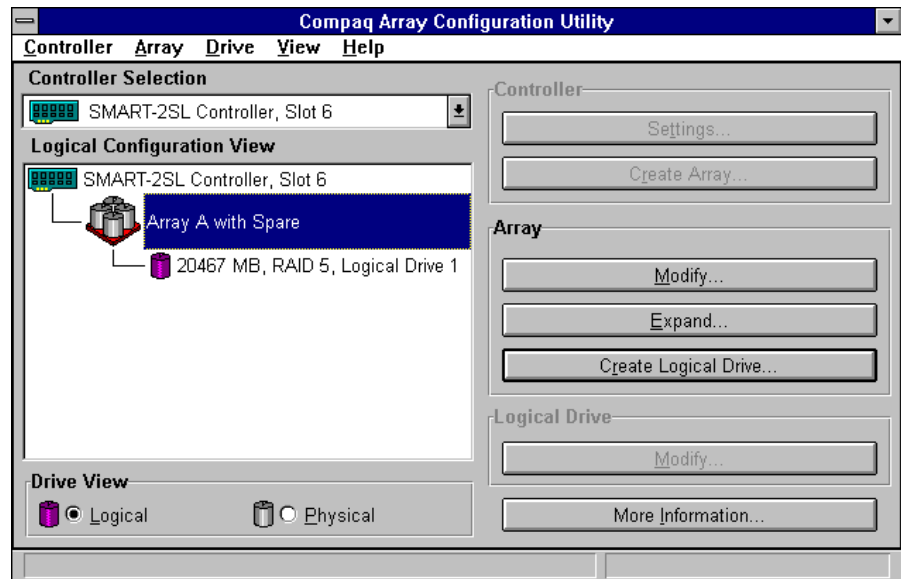


Figure 9. Using the Compaq Array Configuration Utility, system administrators can view array configuration details and make changes as needed.

Compaq Insight Manager

Compaq Insight Manager, a client/server SNMP-based tool for integrated management of the server environment, provides the capability for monitoring performance and other operational characteristics of SMART-2SL Controller-based storage. It displays configuration information, OS device driver version numbers, controller firmware version numbers, pre-failure warranty information, and operating statistics.

Insight Manager gives the system administrator a window to look at low-level performance characteristics of the SMART-2SL Controller in the administrator's environment. It monitors three basic SMART-2SL Controller performance parameters for proactive management of the storage subsystem.

- I/O command per second
- average command latency
- local processor utilization

Analysis of these key parameters can assist administrators in fine tuning configurations for performance. Insight Manager also provides a means for graphically charting performance over time for each of these parameters.

Compaq Insight Manager can also warn administrators of potential drive problems before they occur. A background task on the controller monitors several key drive parameters and notifies Insight Manager if a drive fails to meet certain factory-preset criteria. Insight Manager, in turn, alerts the administrator to the potential problem.

CONCLUSION

As a technology leader, Compaq continually develops new and improved products and tools to meet the expanding computing needs of all our customers. The new Compaq SMART-2SL Array Controller brings proven, dependable hardware RAID protection for mission-critical data to businesses and work sites where it had previously been cost prohibitive.