WHITE PAPER

June 1996

Prepared By
OS Development

Compaq Computer Corporation

CONTENTS

The Compaq SMART-2 Controller and NetWare

The speed and reliability of the storage subsystem is a key element in the overall performance and reliability of your server. A drive array subsystem, such as the SMART-2 Controller, provides high performance and hardware-based fault tolerance that can support a large number of users. In addition, the SMART-2 Controller allows:

- *Online configuration, including deleting arrays or adding spare drives to an array*
- *Online expansion of arrays without duplicating parity space*

This white paper explains how to use the SMART-2 Controller in a NetWare environment and explains some of the performance implications of expansions, rebuilding and specific fault tolerant configurations.

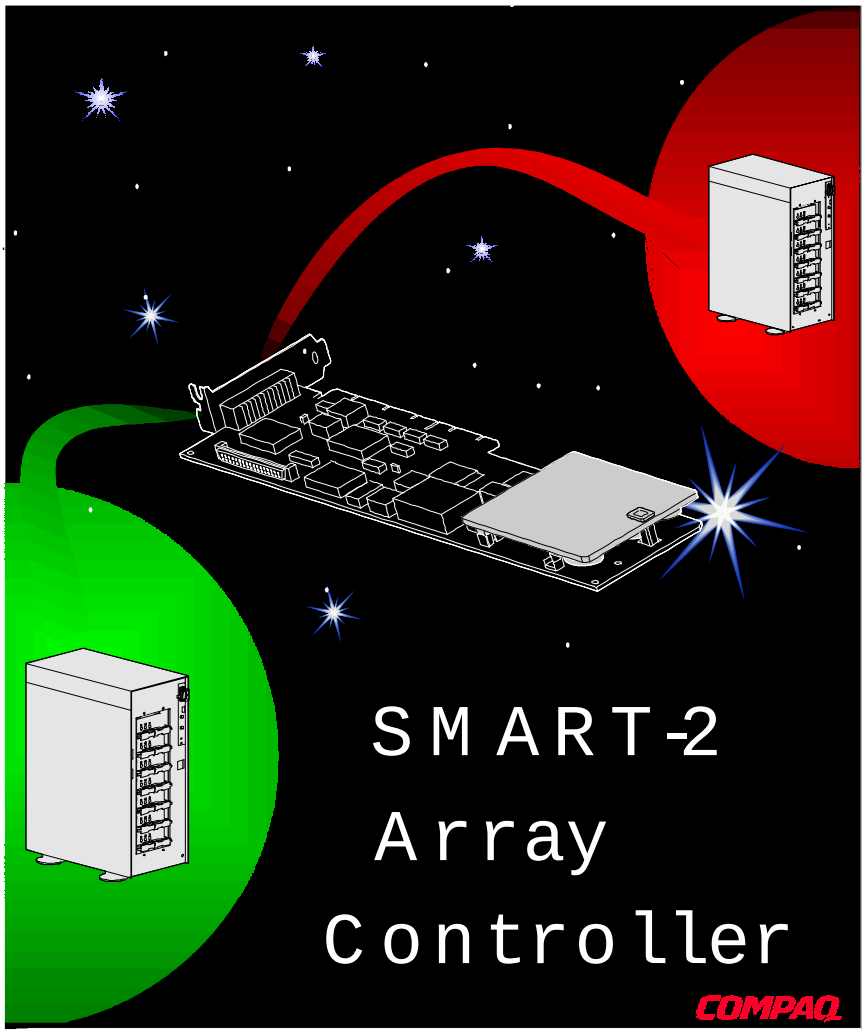

COMPAQ

NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

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 does not warrant products other than its own strictly as stated in Compaq product warranties.

The Compaq SMART-2 Controller and NetWare
First Edition (June 1996)

CONFIGURING THE SMART-2 CONTROLLER WITH NETWARE

IMPORTANT!
It is very important that NetWare be selected as the operating system. Otherwise the drive geometry may not be optimal and the system may operate slower than if the correct operating system was selected.

You should use the Compaq System Configuration Utility to configure your SMART-2 Controller when it is first installed. In the System Configuration Utility, select NetWare as your operating system. You can configure the SMART-2 Controller from the System Configuration Utility or you can do so later using CPQONLIN.NLM.

Load the CPQDA386.DSK array device driver. CPQONLIN.NLM can be used to configure the SMART-2 Controller without downing your NetWare server or dismounting any existing NetWare volumes.

CPQONLIN.NLM - INSTALLATION

In addition to configuring your SMART-2 Controller, the Array Configuration Utility for NetWare also provides basic information about the status of drives attached to SMART-2 Array Controllers indicating drive failures, expansions or waiting for expansion or rebuild (queued). Before loading CPQONLIN.NLM, you must load the array device driver, CPQDA386.DSK (v3.03 or later.)

Warning:
CPQONLIN.NLM cannot be loaded from a DOS partition on NetWare 3.12 servers. Attempting to load this NLM from DOS may cause the server to hang.

1. Copy the CPQONLIN.NLM into a directory on a NetWare volume such as SYS:SYSTEM.
2. CPQONLIN.NLM uses Novell's Media Manager. Therefore, you must have PARTAPI.NLM and VOLLIB.NLM in a NetWare directory that is in the SEARCH path. These files can be found on the NPFC diskettes in the DRV_ARRY.
3. Load CPQDA386.DSK even if the Compaq array controllers in the system have not been configured.
4. Load CPQONLIN.NLM from the system console.

CUSTOM CONFIGURATION OF THE SMART-2 ARRAY CONTROLLER

Custom Configuration allows you to create arrays and assign fault tolerance—one array at a time. When using custom configuration, you must:

- Create each array
- Create at least one logical drive per array and assign fault tolerance

CPQONLIN.NLM on-line help also provides instructions for completing these tasks.

Create Arrays and Logical Drives

You must create arrays and logical drives, and then assign fault tolerance. Keep the following points in mind when creating arrays:

- You can use different size drives in the same array; however, the smallest physical drive in the array determines how much space will be used on each drive. You will not be allowed to use the extra space at the end of the larger physical drives. For example, if you have an array consisting of a 4GB drive and a 2MB drive, only 2MB of the 4GB drive will be used. To avoid wasted space, group like-sized drives into arrays.
- You can mix Fast-SCSI-2 (narrow) drives and Fast-Wide SCSI-2 (wide) drives in arrays.
- If a wide drive fails in an array, it can be replaced by either a narrow or wide drive as long as the drive is the same size (or larger) as the failed drive. Either type drive can be used as a spare for any array.

After creating arrays, you must create logical drives. Keep the following points in mind:

- There must be at least one logical drive per array.
- The logical drives in the array can be different sizes.
- Any space not configured as a logical drive cannot be used to store data. This space will be marked unused.
- Each logical drive can be assigned its own fault tolerance level.
- Expansion is accomplished by adding physical drives and creating new logical drives within an existing array.
- A maximum of 32 logical drives per controller is allowed.

Fault Tolerance

When configuring or expanding drives attached to the SMART-2 Controller, you should assign fault tolerance. Select the appropriate method of fault tolerance based on site-specific criteria, such as cost, performance, and criticality of the data.

Careful: When assigning fault tolerance, a different level of fault tolerance can now be assigned to each logical drive in an array (with previous controllers, this was not possible because only one logical drive per array was allowed.) Be careful when assigning fault tolerance. If you inadvertently assign RAID 0 to one or more logical drives in an array and RAID 1, 4 or 5 to another logical drive in that array, the entire array is not protected. In the event of a drive failure, only the data protected by fault tolerance will be recovered—not the entire array.

Tip: *Of the fault tolerant choices, RAID 1 provides the best performance. RAID 5 uses less capacity for fault tolerance. For many customers, RAID 5 is the best cost/performance fault tolerance option.*

The SMART-2 Controller supports the following fault tolerance options:

RAID 0 - No fault tolerance. If any physical drive fails, data is lost on all RAID 0 logical drives in the array.

RAID 1 - Hardware mirroring. For each drive containing data, a mirrored drive is required. Fifty percent of storage space is used for fault tolerance. Tolerant of multiple, simultaneous drive failures as long as the failures do not include both an original data drive and its mirror. Higher performance than RAID 4 or 5. Requires an even number of drives.

RAID 4 - Tolerant of single drive failures. Stores parity on a single drive and data on all others. Requires three or more drives in an array. If you wish this type of fault tolerance, Compaq recommends that you use RAID 5 for best performance.

RAID 5 - Tolerant of single drive failures. Stores parity data across all drives. Requires three or more drives in an array. If a drive fails, the controller uses the parity data and data on remaining drives to reconstruct data requested from the failed drive and continues operation. Uses between 7 and 33% of storage space for fault tolerance depending upon the number of drives in the array.

In a NetWare environment, the method of fault tolerance can affect network server performance. The Compaq hardware fault tolerance methods described above are independent of NetWare, relieving it of the task of managing the hard drive subsystem fault tolerance.

The following chart gives one example of the performance differences seen between RAID 1, RAID 5 and NetWare mirroring (RAID 0 hardware fault tolerance). Compaq highly recommends that some type of hardware fault tolerance be used. Hardware fault tolerance protects all partitions that exist on your drive subsystem, including your DOS partition. In addition, hardware fault tolerance allows the use of online spares, interim data recovery, and automatic data recovery.

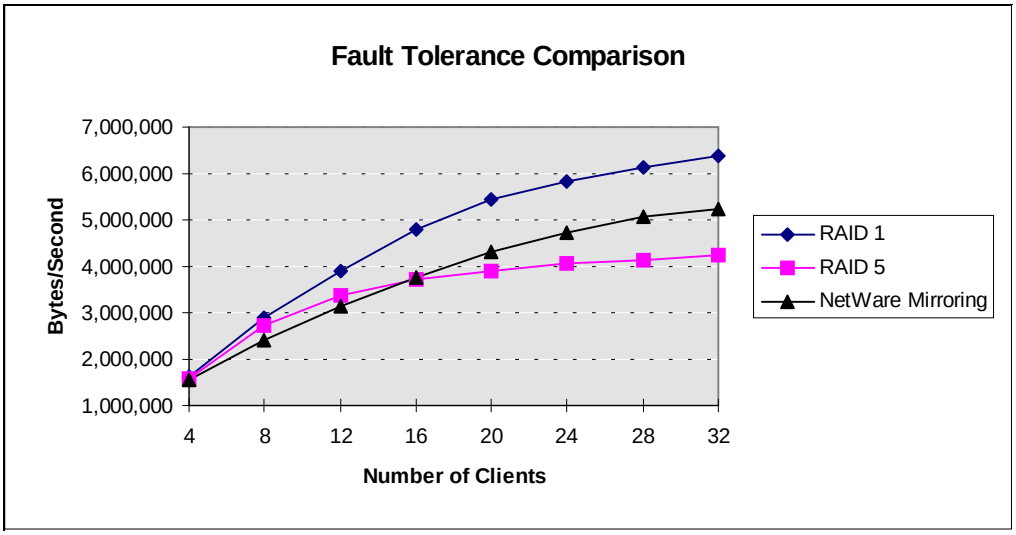

Equipment used: Compaq ProLiant 4500 with SMART-2/E Array Controller, NetFlex-3 100VG/E running NetBench v4.0 Disk Mix, Think Time=2 sec.

Spare Drives

WARNING: *An assigned spare will only recover the parts of a physical drive that were protected by fault tolerance—because fault tolerance is assigned per logical drive, to recover an entire physical drive, make sure that all logical drives on the array have fault tolerance assigned.*

A spare drive is used in conjunction with hardware fault tolerance. A drive is assigned to an array as a spare and takes the place of a single drive, should it fail. For example: RAID 5 has been selected for use by ARRAY A, and a spare has also been assigned to that array. If one of the drives in ARRAY A fails, RAID 5 rebuilds the data automatically on the spare, replacing the failed drive.

When the rebuild of the spare drive completes, the array is returned to full fault tolerance. Spares are especially useful if no one is immediately available to replace failed drives, such as at remote sites. The data is protected by fault tolerance, and the array can withstand another drive failure after the spare has taken over for the first failed drive. Requirements:

- You must have at least one unassigned physical drive in the system and be running RAID 1, RAID 4 or RAID 5.
- A spare drive can be assigned to multiple arrays.

When you select the "Assign Spare Drive" in CPQONLIN.NLM only those drives that qualify will appear. If drives do not appear that you expect to see, switch to the physical drive view (TAB key) and check the size of the drives. The spare must be as large as the smallest drive in the array.

Other Functions: Expand, Rebuild, Replacing Failed Drives

There are several other configuration options available in CPQONLIN.NLM. All of these can be performed without downing your NetWare server.

- Expand an Array
- Drive Rebuild, Expand Priority and Accelerator Ratio
- Failed Drive or Interim Recovery Mode
- Drive Failures During the Expand Process

Expand an Array

Use CPQONLIN.NLM to expand storage while on-line. The following rules apply to on-line expansion:

- Expansion uses the battery-backed cache memory on the SMART-2 Controller, thus you must have the cache board attached and the batteries must be charged. When the system is powered on, a warning message is displayed if the batteries are not sufficiently charged or if the cache board is not attached. The cache memory will NOT be used for normal cache operations during the expansion process.

Tip: *Adding physical drives not only increases capacity, it often increases performance. More drives in an array allow for execution of a greater number of concurrent requests. For example, if one drive is capable of 100 I/O requests per second, two drives would be capable of twice that number, and so on.*

- When expanding, you must add drives that are minimally as large as the smallest drive in the array. Adding larger drives will waste space because you will only be allowed to use the space that is equal to the smallest drive size.

- NetWare can continue to access data on an array that is expanding.

  Note: Any new logical drives created on an array during an expansion will not be available for use until the expansion of the existing logical drives is complete.

How it works

All logical drives that existed are re-spanned across the old and new space. Any new logical drive(s) will also span the old and new space. The following diagram demonstrates this technique.

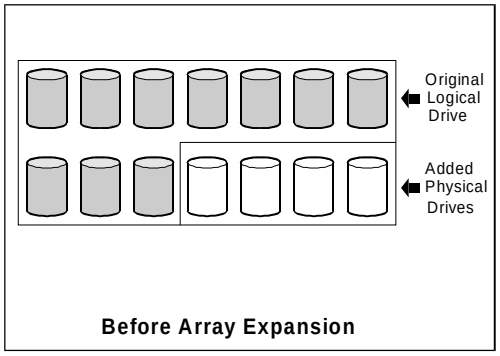

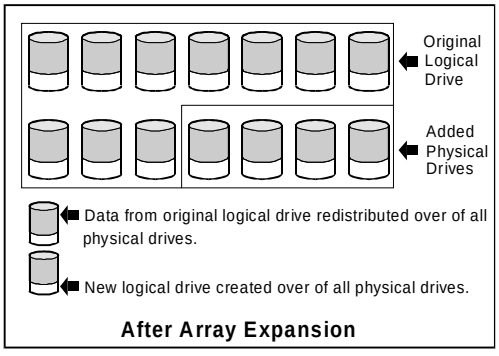

The space represented by the new logical drive can be used to either expand an existing NetWare volume or create a new volume. An existing NetWare volume can be expanded by adding the partition created on the new logical drive as a new segment to the volume. For example, if the SYS volume is not large enough, you can expand the array and then use INSTALL.NLM to create a NetWare partition and add a new segment to the SYS volume.

Changing Existing Fault Tolerance

When you expand an array and assign fault tolerance to the new logical drive, you cannot change the old fault tolerance assignment. For example, if your original system is configured for no fault tolerance and you add drives, the original logical drive will still be unprotected, regardless of whether you assign fault tolerance to the new logical drive(s). See the diagram below.

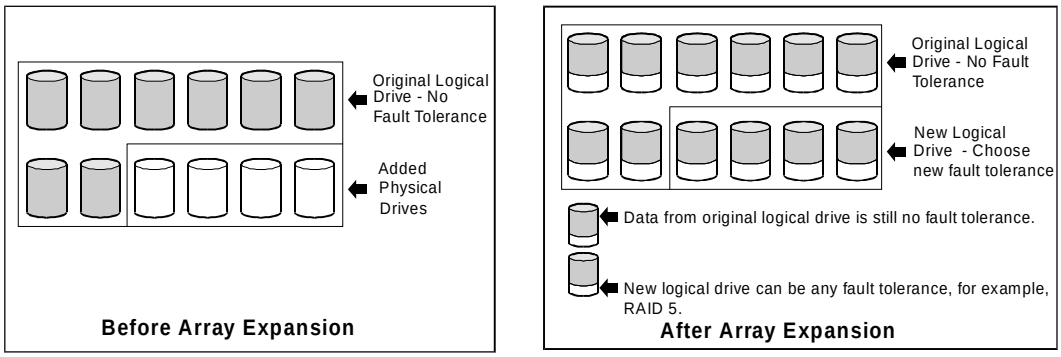

Tip: *You may want to perform the expand process during periods of low server utilization because in some configurations, the expand may affect performance. If you must expand during peak periods, the Array Configuration Utility for NetWare allows you to set the priority of the expand. Setting the priority to LOW will affect performance the least, but it will be longer before the new space is available.*

During an expand, performance may be slightly degraded. In most cases, however, any potential degradation is offset by the addition of physical drives. In the following chart, four drives were expanded to seven. Expand priority was set to high. The chart illustrates relative performance before expand, during expand and after expand.

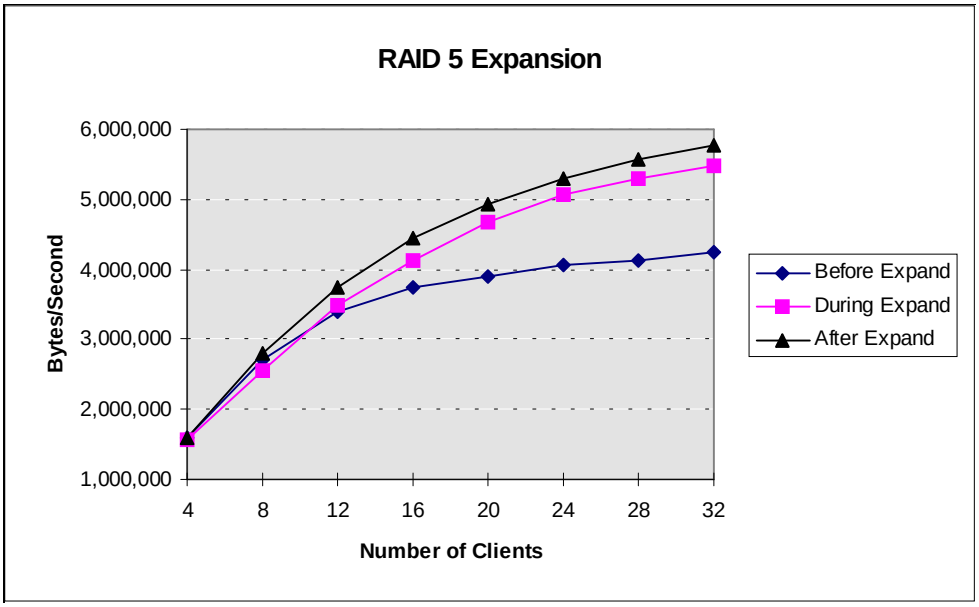

Equipment used: Compaq ProLiant 4500 with SMART-2/E Array Controller, NetFlex-3 100VG/E running NetBench v4.0 Disk Mix, Think Time=2 seconds.

Drive Rebuild, Expand Priority and Accelerator Ratio

To set the drive rebuild priority, expand priority or accelerator ratio, select the SMART-2 Controller that you wish to configure and then select the Controller Settings option.

Drive Rebuild

Tip: Expanding does not affect performance as much as rebuilding a drive. Performance during a rebuild is slowed mainly because all data that is requested from the failed drive must be regenerated from parity (RAID 4 and RAID 5), not because the rebuild is occurring.

Drive rebuilding will occur automatically if a spare drive was assigned to the array. If a spare was not assigned, rebuilding will occur after a failed physical drive is replaced. Only those logical drives on the array with the failed physical drive that were configured with fault tolerance (RAID 1, RAID 4 or RAID 5) will rebuild.

Performance during rebuild is decreased, not only because the data on the new drive is being rebuilt, but also because any data requested from the failed drive must be regenerated from parity (RAID 4 and RAID 5 configurations). In RAID 1 configurations, performance degradation after a drive failure will be negligible.

Priority Settings

When you choose low priority (the default) for drive rebuilding, rebuild takes place only when the SMART-2 controller is not busy handling normal I/O requests. Low priority has minimal effect on normal I/O operations. With high priority, drive rebuilding occurs faster at the expense of normal I/O operations. You may want to choose high priority for drive rebuilding because the array is vulnerable to additional drive failures while the drive is being rebuilt.

Accelerator Ratio

The SMART-2 Controller has an on-board cache, called an Array Accelerator, that does both write-posting and read-ahead caching. Leaving this accelerator on will improve performance in most environments. The read-ahead feature reads more data than is requested and places it in the cache in anticipation of a request for this data. The posted write feature coalesces small writes into larger ones before writing the data to the drives. The Array Accelerator was designed to protect data integrity. Batteries maintain any cached data for up to four days in case of a system power failure. When power is restored to the system, the controller writes the preserved data to the disk drives. This allows you to take full advantage of the performance improvements of the SMART-2 Controller without sacrificing reliability.

The Array Accelerator is also removable so that in case of a controller failure, it may be physically transferred to another SMART-2 Controller. Any data remaining in the cache will then be written to the disk drives when the replacement controller is powered on.

Failed Drives, Interim Recovery Mode and Other Status Indicators

If a drive fails and you have hardware fault tolerance enabled, operation will continue. You should replace a failed drive as soon as possible. The following status indicators are available from the Array Configuration Utility when F3 is pressed:

- Interim Recovery: the logical drive is operating, but a failed physical drive has not been replaced. Replace the drive as soon as possible.
- Ready for Recovery: the logical drives are queued for recovery. This status will be shown when another logical drive is already rebuilding or expanding.
- Rebuilding: the logical drive is operating. A replacement drive or an on-line spare is being rebuilt.

NOTE:
A failed status can occur on drives protected by fault tolerance if two or more physical drives fail at the same time.

- Logical Drive Failed: If you have one or more logical drives in an array that are not protected by fault tolerance, the data on these logical drives will be lost. The Array Configuration Utility will show the logical drives as FAILED. After drive replacement, any fault tolerant logical drives will rebuild. The logical drives that were not protected (FAILED) will become available for data (the devices are reactivated automatically). If you have a backup of the data, you can restore it at this time.

  If you do not replace the failed drive, the only action that you can take from the Array Configuration Utility is to delete logical drives. Do NOT delete logical drives that contain valid data. Doing so will result in lost data.

Some of these status indicators are available without pressing F3. For example, on the main menu, the FAILED status will appear next to the logical drive that has failed. EXPANDING and REBUILDING will appear next to the array in which the activity is occurring.

Drive Failures During Expand Process

IMPORTANT!
Compaq highly recommends that some type of fault tolerance be used on all logical drives.

Compaq highly recommends use of either RAID 1 or RAID 5 fault tolerance on drive arrays. If fault tolerance is used on all logical drives in an array and a drive fails while an expand is being performed, the expand process will continue until it is complete. No data will be lost. All logical drives in the array will be in interim recovery mode after the expand has completed. Interim recovery mode means that a drive has failed and you should replace it as soon as possible, however, normal data I/O continues. When the drive is replaced, the data will be rebuilt from parity (RAID 4 or 5) or from a mirrored drive (RAID 1).

The expand process will complete on the all logical drives that were protected by fault tolerance. All data that was protected by fault tolerance will not be lost. Data that was not protected by fault tolerance will be lost. When the expand is completed, you will need to replace the failed drive and restore any lost data from back-up sources.

UPGRADING THE SMART CONTROLLER TO THE SMART-2 CONTROLLER

If you wish to upgrade your SMART Controller to a SMART-2 Controller, no special steps are needed. After removing the SMART Controller and replacing it with the SMART-2 Controller, the data on the drives will be accessed by the SMART-2 Controller. More detailed installation instructions follow:

1. Down your NetWare server and remove the SMART Controller.
2. Insert the SMART-2 Controller and replace all cables.
3. When you reboot, a Slot ID mismatch will occur and you should run the System Configuration Utility. If your ports have been connected incorrectly, the following message will be displayed: "External cables attached to wrong SCSI Port Connectors. Turn system power off and swap SCSI port connectors to prevent data loss." Reconnect your cables to the correct ports before continuing—see the diagram below.

IMPORTANT!
Drives must be kept on the same port to which they attached when the SMART was installed. The external connector on SMART-2 is port 2 (visible at the back of the machine when the SMART-2 is installed). To use port 1, you need to connect the connector panel and ribbon cable to the outside of the server in which the SMART-2 is installed.

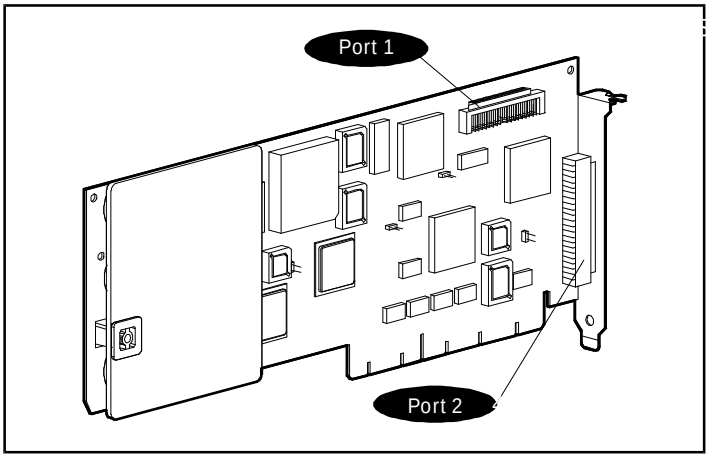

   Ensure that the drives are connected to the proper ports and reboot. The SMART-2 Controller should then continue successfully.
4. Run the Compaq System Configuration Utility. The configuration will match the old configuration except for the upgraded SMART-2 Controller. Do not change any of the SMART-2 settings as it may have an adverse affect on your data.
5. Upgrade the CPQDA386.DSK array device driver.
6. Start your NetWare server.

SUMMARY

The SMART-2 Controller has many features that make it easy to configure and operate. Using the SMART-2 Array Controller provides many important benefits, including data protection, performance enhancement, capacity changes, and data reliability. For more information on the SMART-2 Controller, see the Compaq SMART-2 Array Controller Performance Brief (document number 034B/0296) which includes performance information on competitor controllers and operating systems other than NetWare. Also, check the Compaq World Wide Web site at : WWW.COMPAQ.COM.