

Attorney's Docket No. 02855.P001

Patent

ASSISTANT COMMISSIONER FOR PATENTS
Washington, D.C. 20231

SIR: Transmitted herewith for filing is the **nonprovisional patent application** ofInventor(s): James R. Bergsten

For: DATA STORAGE CONTROLLER FOR PROVIDING MULTIPLE HOSTS WITH ACCESS
TO MULTIPLE STORAGE SUBSYSTEMS

(Title)

Enclosed are:

- ☒ 18 sheet(s) of Drawings.
☒ An Assignment of the invention to Ark Research Corporation.
☒ Assignment Cover Sheet Form PTO-1595.
☒ A Declaration and Power of Attorney (☒ signed/ ☐ unsigned).
☒ A Verified Statement to establish Small Entity Status under 37 C.F.R. §§ 1.9 and 1.27.
☒ Information Disclosure Statement; PTO form 1449; 7 cited references

The Filing Fee has been calculated as shown below:

(Col. 1)				(Col. 2)		SMALL ENTITY		SMALL ENTITY		
For:	No. Filed				No. Extra		Rate	Fee	Rate	Fee
Basic Fee:							\$ 385		\$ 770	
Total Claims:	61	- 20	*	41		x 11	\$	x 22	\$ 902	
Indep. Claims:	12	- 3	*	9		x 40	\$	x 80	\$ 720	
	Multiple Dependent Claim(s) Presented					+ 130	\$	+ 260	\$	
* If the difference is less than zero, enter "0" in Col. 2.						TOTAL	\$	TOTAL	\$ 2392	

* If the difference is less than zero,
enter "0" in Col. 2.

- ☒ A check for \$ 2,392.00 for the filing fee is enclosed.
☒ A check for \$ 40.00 for recordation of the Assignment is enclosed.

"Express Mail" mailing label number EM564113043USDate of Deposit 3/31/97

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

Grace C. Gaoiran

(Typed or printed name of person mailing paper or fee)

(Signature of person mailing paper or fee)

X The Commissioner of Patents and Trademarks is hereby authorized to charge payment of the following fees associated with this communication, or credit any overpayment, to our Deposit Account No. 02-2666. **A duplicate copy of this sheet is enclosed.**

 X Any additional filing fees required under 37 C.F.R. § 1.16.

 Any patent application processing fees under 37 C.F.R. § 1.17.

 The Commissioner of Patents and Trademarks is hereby authorized to charge payment of the following fees during the pendency of this application, or credit any overpayment, to our Deposit Account No. 02-2666. **A duplicate copy of this sheet is enclosed.**

 Any processing fees under 37 C.F.R. § 1.17, including any extension fees.

 Any filing fees under 37 C.F.R. § 1.16 for presentation of extra claims.

 X Send all correspondence to the undersigned at BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, 12400 Wilshire Boulevard, Seventh Floor, Los Angeles, California 90025, and direct all telephone calls to the undersigned at (408) 720-8598.

Respectfully submitted,

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP

Date: 3/31/97

By Jordan M. Becker
Jordan M. Becker

12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025
(408) 720-8598

Reg. No.: 39,602



UNITED STATES PATENT APPLICATION

FOR

Data Storage Controller for Providing Multiple Hosts with Access to Multiple
Storage Subsystems

INVENTOR:

James R. Bergsten

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CALIFORNIA 90025
(408) 720-8598

Attorney's Docket No. 002855.P001


"Express Mail" mailing label number EM564113043US

Date of Deposit 3/31/97

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

Grace C. Gao

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing paper or fee)



2,392,101
68/828888

Data Storage Controller Providing Multiple Hosts with Access to Multiple Storage Arrays

FIELD OF THE INVENTION

The present invention pertains to the field of computer systems. More particularly, the present invention relates to devices which control data flow between one or more host processing systems and one or more data storage subsystems.

BACKGROUND OF THE INVENTION

Computer systems sometimes handle valuable or irreplaceable data. Data maintained by a computer system may be of vital importance, for example, in business applications such as airline reservations, bank account management, electronic funds transfers, shipping and receiving, and inventory control. Consequently, there is a need to ensure that valuable data is adequately protected against loss or damage.

It is common to store large volumes of data on non-volatile mass storage devices, such as magnetic or optical disks. Occasionally, however, a mass storage device will fail, resulting in the loss of stored data. Consequently, it is a common practice to store a back-up copy of data that is considered valuable on a separate, back-up storage device. For practical reasons, however, it is often necessary to locate the back-up storage device in the same geographic vicinity as the primary storage device. As a result, both the primary data and the back-up copy may be lost or destroyed due to theft, vandalism, or natural disasters, such as fires, floods, and earthquakes. Hence, there is a need for the capability to store multiple back-up copies of data in geographically separate locations, while still permitting quick and easy access by a host computer to any back-up copy.

SUMMARY OF THE INVENTION

The present invention includes a storage controller that includes a data interface with a host processing system and a data interface with at least one storage array. Each storage array includes at least one mass storage device.

- 5 The host interface is independent of a device type of the host processing system, and the storage interface is independent of the device type of the mass storage devices. The data storage controller further includes control circuitry coupled to the host interface and the storage interface. The control circuitry is configured to enable the host processing system to access each of the storage
- 10 arrays.

Other features of the present invention will be apparent from the accompanying drawings and from the detailed description which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

5 Figure 1 illustrates a computing system in which a number of storage controllers provide multiple host computers with access to multiple storage arrays.

10 Figure 2 illustrates a computing system according to an embodiment in which one storage controller is connected to a number of other storage controllers.

 Figure 3 is a block diagram of a storage controller.

 Figure 4 is a block diagram illustrating functional modules of a storage controller.

15 Figure 5 is a block diagram illustrating a host computer system including a storage controller according to the present invention.

 Figure 6 is a flow diagram illustrating a check routine performed in response to an input/output request.

 Figure 7 is a flow diagram illustrating a routine for mapping a host address to a logical address.

20 Figure 8 illustrates a tree hierarchy used by a storage controller to map a host address to a logical address.

 Figure 9 is a flow diagram illustrating a routine for mapping a logical address to a physical address.

25 Figure 10 is a flow diagram illustrating a routine performed by a storage controller for staging data from a mass storage device.

Figure 11 is a flow diagram illustrating a read error recovery routine.

Figure 12 is a flow diagram illustrating a routine performed by a storage controller for destaging data to a mass storage device.

Figure 13 is a flow diagram illustrating a write error recovery routine.

5 Figure 14 is a flow diagram illustrating a routine performed by a storage controller for reading data from a mass storage device in response to a request from a host.

10 Figure 15 is a flow diagram illustrating a routine performed by a storage controller for writing data to a mass storage device in response to a request from a host.

Figure 16 is a flow diagram illustrating a routine performed by a storage controller for transmitting data to a remote storage controller.

Figure 17 is a flow diagram illustrating a remote communications error recovery routine.

15 Figure 18 is a flow diagram illustrating a routine for establishing copies.

Figure 19 is a flow diagram illustrating a routine performed by a storage controller for performing automatic detection and correction of data errors.

DETAILED DESCRIPTION

A storage controller capable of providing multiple host computers system with access to multiple storage arrays is described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be evident, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to facilitate description of the present invention.

As will be described below in detail, the present invention includes a storage controller which allows multiple host computer systems at different locations to access any of multiple copies of stored data. The storage controller automatically creates and manages multiple back-up copies while the host computer systems are "on line" in a manner that is both non-disruptive of, and transparent to, the host computer systems and their users. Further, the storage controller automatically detects and correct errors in stored data and automatically replaces faulty copies. Moreover, the storage controller is not dependent upon any particular hardware or software configuration of the host computer system which it services or the mass storage devices which it accesses. The storage controller emulates a local storage array for the host computer system which it services and emulates a host computer system for the local storage array which it accesses. Host computer systems access stored data using virtual device addresses, which are mapped to real device addresses by the storage controller.

computer (PC), a mini-computer, or a mainframe. In addition, any of the host computer systems may function as a server for one or more client computer systems (not shown).

Each MSD may include non-volatile facility for storing large volumes of data, such as a magnetic disk or tape, an optical storage device such as CD-ROM (Compact Disk-ROM), CD-R (CD-recordable), DVD (Digital Versatile Disk), a magneto-optical (MO) device, or the like. The MSDs within the computing system need not be of the same device type. That is, the MSDs in any given storage array may use a different type of storage medium from those in any other storage array.

Each storage array may be located geographically distant from the other storage arrays. Multiple copies are generally maintained on different, geographically-separated storage arrays. Hence, the loss of one or more MSDs in a given storage array will not result in the complete loss of data. With respect to a given (local) storage controller, any or all of the other (remote) storage controllers, host computer systems, and storage arrays may therefore be located at distant locations to the local storage controller.

Storage controllers 3-1 through 3-M function cooperatively to provide any of host computer systems 2-1 through 2-M with access to any of storage arrays 4-2 through 4-M. Each one of storage controllers 3-1 through 3-M directly services one local host computer system and one local storage array in one embodiment. For example, in the embodiment shown in Figure 1, storage controller 3-1 services and is directly coupled to its local host computer system 2-1 via a data communication path 7. Storage controller 3-1

also is directly coupled to its local data storage array 4-1 via another communication path 8.

In one embodiment, each storage controller communicates with its local host computer system and its local storage array using standard SCSI (Small Computer Systems Interface) protocols. Consequently, operation of a storage controller of the present invention in the manner described herein is not dependent upon the particular hardware or software configuration of any host computer or storage array, as long as those devices are SCSI-compatible. Note that in other embodiments, however, the data communication paths 7 and 8 may conform to other protocols and standards, such as serial SCSI, Fiber Channel, or ESCON. Thus, because data paths 7 and 8 are conventional interfaces, a storage controller of the present invention can be used concurrently with host computers and MSDs having different configurations. For example, one host computer in the system may be a mainframe computer system while another host computer is a PC. Similarly, one storage array in the system may consist of conventional magnetic hard disk drives while another storage array consists of CD-ROM drives.

The storage controllers in the computing system cooperate to allow any of the host computer systems to access data located in any of the storage arrays. For example, host computer system 2-1 may access data located on any MSD in storage array 4-1 by communicating with storage controller 3-1. In addition, host computer system 2-1 may access data located in remote storage array 4-M by communicating with remote storage controller 3-M via storage controllers 3-1, 3-2, etc. As will be explained below, data is accessed using

virtual addressing, such that a host computer system has no knowledge of which physical storage device is being accessed.

The storage controllers operate in peer-to-peer relationships (as opposed to master-slave relationships) with each other when responding to remote access requests. Any storage controller can access data in a remote storage array serviced by any other storage controller. Communication between individual storage controllers takes place on communication links using a common protocol implemented by all storage controllers in the system. Note that various different protocols might be used for this purpose. A description of the details of such a protocol is not necessary for a full understanding of the present invention and is therefore not included herein.

The remote access capabilities of the storage controllers can be used advantageously in a variety of ways. The storage controllers can be used to enable sharing and centralized control of storage media. For example, the MSDs in any of the storage arrays may include CD-ROM devices, so that CD-ROMs in one a storage arrays may be shared by all host computers. Use of the devices can be recorded by one of the storage controllers. Access to CD-ROM data can therefore be monitored and controlled by a user, such as a system administrator, using the storage controller. The present invention also permits low-volume production of CD-Rs without specialized equipment. For example, a limited quantity of CD-R disks can be produced by writing to a single, logical CD-R drive corresponding to multiple physical CD-R devices within one or more storage arrays.

The present invention is not limited to the specific configuration shown in Figure 1. For example, the system configuration might

alternatively include only a single host computer which is coupled to multiple geographically-separated storage arrays using multiple storage controllers. In another embodiment, a storage controller may have a direct interface with more than one host or may have redundant interfaces with the same host. Each storage controller might also be coupled to more than one other storage controller, as illustrated in Figure 2, in which storage controller 3-1 is connected to a number of other storage controllers, 3-2 through 3-M.

The storage controllers provide a number of functions and services to the host computer systems. In particular, the storage controllers automatically provide multiple back-up copies of stored data in a manner that is transparent to, and non-disruptive of, all of the host computer systems. For example, when a given data file is saved by a local host computer to its local storage array, the local storage controller causes one or more back-up copies of the data file to be stored in one or more of the remote storage arrays, without any intervention from or interruption to the local host computer. The process of backing-up data can be initiated by any storage controller in the system. This allows centralized back-up of any or all of the host computers. The number of back-up copies to be created and the locations to which the back-up copies are to be stored can be configured in advance by a system administrator. For example, these parameters can be programmed into a storage controller by connecting a keyboard and a display to ports on the storage controller. Alternatively, the storage controller can be configured by using one of the host computers to transmit configuration commands and data to the storage controller. Further, a remote storage controller can be configured via a local storage controller. The storage controller of the present

invention also performs automatic block error detection and reassignment, as will be described below in detail.

The remote data access, data mirroring, and path redundancy provided by the present invention allow recovery from many possible failure modes, such as failure of communication medium, failure a host computer, or failure of a storage device.

In general, when the system of Figure 1 is first set up, or when MSDs are added to the system, all data from a "primary" MSD are automatically copied to one or more newly-added MSDs. Such copying is performed while the host computers are operating, although the copying process is transparent to the host computers 2-1 through 2-M. Copying can be throttled to reduce the impact on system performance. Once all copies are established (brought into a consistent state), only changed data are subsequently written to the secondary MSDs in response to changes dictated by a host computer.

Each of the storage controllers also provides virtualized data access and emulation, as mentioned above. A local storage controller will emulate its local storage array from the viewpoint of its local host computer system; similarly, the local storage controller will emulate its local host computer system from the viewpoint of its local storage array. Such emulation is implemented, in part, by using a common communication interface for data communication paths 7 and 8, such as SCSI. Again, in other embodiments, the data communication paths 7 and 8 may conform to other protocols and standards, such as serial SCSI, Fiber Channel, or ESCON.

A local host computer accesses data by transmitting a (virtual) host address to its local storage controller. The host address is then mapped to a

real address representing a location on one or more physical MSDs, as will be described below in detail. The mapping is completely transparent to all of the host computers. To improve performance, the storage controller may distribute a single copy among multiple MSDs, which may be located in
5 different storage arrays. Hence, in the above described mapping process, a single host address may map to multiple physical addresses, which may be distributed among multiple MSDs, and such MSDs may further be located in different storage arrays.

Figure 3 illustrates the hardware architecture of a storage controller of
10 the present invention according to one embodiment. The storage controller 3 includes a central processing unit (CPU) 10, random-access memory (RAM) 11, a non-volatile storage facility (NVSF) 12, an MSD 13, each of which is coupled to a bus 17. Bus 17 may represent multiple physical or logical buses, which may be interconnected by various adapters and/or controllers. NVSF
15 12 may be, or may include, a programmable non-volatile storage device, such as flash memory or electrically erasable programmable read-only memory (EEPROM). MSD 13 may be any conventional device that is suitable for non-volatile storage of large volumes of data, such as any of those discussed above. Also coupled to the bus 17 are a host device interface 14, a storage device
20 interface 15, a controller device interface 16, and input/output (I/O) device interfaces 18 and 19. I/O device interfaces 18 and 19 are also coupled to separate, external connectors on the storage controller 3.

The host device interface 14 connects communication path 7 to the bus
17 in order to connect a local host computer to the storage controller 3. The
25 storage device interface 15 connects communication path 8 to the bus 17 in

order to connect a local storage array to storage controller 3. In one embodiment, host device interface 14 and storage device interface 15 each comprise a SCSI adapter, such that communication between storage controller 3 and the local host computer and storage array is performed using SCSI protocols. As is well-known, the SCSI standard allows multiple peripheral devices to be connected to a host device in a daisy chain configuration. From the viewpoint of the local storage array, the storage controller 3 emulates the host device. From the viewpoint of the host device, the storage controller 3 emulates the local storage array.

The controller device interface 16 connects communication path 9 to the bus 17 in order to connect a remote storage controller to storage controller 3. Controller device interface 16 may be an Ethernet, ATM, T1, T3, or FDDI adapter, or any other suitable device, depending upon the nature of the communication link 9.

I/O device interfaces 18 and 19 may be used to connect a keyboard and a monitor to the bus 17. I/O interface 18 and 19 may therefore be used by a systems administrator to perform various functions, such as initially configuring the storage controller 3, inputting commands and control information to the storage controller 3, or obtaining status information from the storage controller 3. Further, these interfaces 18 and 19 can be used to remotely perform these same functions on a remote storage controller via (local) storage controller 3 and communication link 9.

In one embodiment, the present invention is carried out in the storage controller by the CPU 10 executing sequences of instructions that are contained in a memory. More specifically, execution of the sequences of

instructions contained in the memory causes the CPU 10 to perform steps according to the present invention which will be described below. For example, instructions may be loaded into RAM 11 for execution by the CPU 10 from a persistent store, such as the NVSF 12 or MSD 13. In other embodiments, hardwired circuitry may be used in place of, or in combination with, software instructions to implement the present invention. Thus, the present invention is not limited to any specific combination of hardware circuitry and software, nor to any particular source for the instructions executed by a computer system.

embodiments, hardwired circuitry may be used in place of, or in combination with, software instructions to implement the present invention. Thus, the present invention is not limited to any specific combination of hardware circuitry and software, nor to any particular source for the instructions executed by a computer system.

Figure 4 shows the functional modules within a storage controller 3 of the present invention. Each of these functional modules may be embodied in hardware, software, or a combination of hardware and software. In one embodiment, the storage controller 3 includes an operating system (OS) 20 coupled to emulation drivers 21, physical drivers 22, communication drivers 23, and local memory 24. Memory 24 may include any combination of RAM 11, NVSF 12, or MSD 13 (see Figure 3).

23, and local memory 24. Memory 24 may include any combination of RAM 11, NVSF 12, or MSD 13 (see Figure 3).

The OS 20 controls the overall operation of the storage controller, including emulation, data flow, caching of data in the local memory 24, control of data mirroring, and error recovery. Emulation drivers 21 are controlled by the OS 20 to provide communication of data between the storage controller 3 and the local host computer 2 via communication path 7. In particular, the emulation drivers 21 receive read and write requests from the host computer 2 and convert the read and write requests into a format recognized by the OS 20. The emulation drivers 21 further function to transfer any data received from the local storage array or a remote storage

controlled by the OS 20 to provide communication of data between the storage controller 3 and the local host computer 2 via communication path 7. In particular, the emulation drivers 21 receive read and write requests from the host computer 2 and convert the read and write requests into a format recognized by the OS 20. The emulation drivers 21 further function to transfer any data received from the local storage array or a remote storage

```
25  transfer any data received from the local storage array or a remote storage
```

array to the local host computer. The communication drivers 23 are controlled by the OS 20 to provide communication of data between the storage controller 3 and remote storage controllers via communication path 9. The physical drivers 22 are controlled by the OS 20 to transfer data between
5 the storage controller 3 and its local storage array. Such transfer includes the transformation of commands and control data from the format used by the storage controller 3 to the format recognized by the local storage array (e.g., SCSI).

In an alternative embodiment, a storage controller 3 may be an internal
10 component of a host computer, rather than a separate, external unit. Figure 5 illustrates a host computer 2-R according to one such embodiment. The host computer 2-R includes a CPU 30, a RAM 31, an NVSF 32, an MSD 33, a modem 34, a network adapter 35, and a storage controller 36, each of which is coupled to a bus 37. The local bus 37 may be any conventional local bus, such
15 as a PCI (Peripheral Component Interconnect) bus. In this embodiment, the storage controller 36 may be constructed on an add-in card inserted in the host computer 2-R. The storage controller 36 provides all of the functionality of the storage controllers described above and has essentially the same architecture. However, in contrast with the storage controllers described
20 above, the storage controller is directly connected to the local bus 37 of the host processing system. This embodiment therefore eliminates the need for the emulation drivers 21 illustrated in Figure 4. Communication paths 8 and 9 to the local storage array and the remote storage controllers, respectively, are provided through one or more external connectors on the host computer 2-R.

The modem 34 and network adapter 35 each allow the host computer 2-R to communicate other computer systems, such as clients or other hosts.

The various functions of the storage controller of the present invention will now be described in greater detail. The storage controller is configured to provide capability to gracefully recover from numerous types of failure modes, including block errors, device failures, and failure of communication paths, in a manner that is transparent to all hosts. For example, if an error occurs during an I/O operation, and the error is determined to result from a communication media failure (e.g., failure of one of the communication paths 7, 8 or 9), the storage controller automatically performs the operation using an alternate communication path, if available. If the error is determined to be a device failure (i.e., failure of an MSD), then the storage controller automatically attempts to perform the write operation to an alternate storage device. If the error is a block error, the system automatically attempts to refresh the block or, if necessary, to reassign the data to an alternate block.

In response to each I/O request (read or write) from a host, the storage controller performs a series of standard checks before carrying out the operation. Figure 6 illustrates a routine for performing these standard checks.

It is first determined, in step 601, whether the storage device identified by the host exists. As noted above, this device will be a virtual device. The method of identifying the physical MSD to be accessed is discussed below. If the device does exist, then the routine proceeds to step 602. Otherwise, the routine exits in step 610 with a "no device" error being provided to the host. If the device exists, then in step 602, it is determined whether the device to be accessed is

performance, the storage controller may map a single host address to multiple physical addresses, which may be distributed among multiple MSDs located in different storage arrays. The storage controller 3 maps a host address to one or more physical addresses using a two-part process that is transparent to all
 5 hosts. Specifically, a host address received from a host is first mapped to a logical address, and the logical address is then mapped to a physical (real) address in one or more MSDs.

Figure 7 illustrates the process of mapping a virtual address to a logical address. Because a given storage controller may have an interface with more
 10 than one host or multiple paths to the same host, each host address includes a host interface identifier (ID) that identifies the particular host or redundant path with which the access request is associated. Hence, in step 701, the storage controller receives a host interface ID and a host (memory) block number. In step 702, the storage controller determines whether the host
 15 device ID and block number map exactly to a logical device. This determination can be made by checking the status of a single bit representing whether or not the mapping is exact.

The storage controller maintains and uses a tree structure such as that illustrated in Figure 8 to map the host interface ID and block number to a
 20 logical device. The tree structure 40 of Figure 8 comprises a number of elements 41, 42 and 43. Each element includes a field designating a range of memory blocks as well as pointers to both a logical device and a logical block for that range. In addition, each element includes a pointer to another element corresponding to a different memory range. For example, the range
 25 of memory blocks represented by element 42 may be a lower range than that

represented by element 41. Accordingly, element 41 would include a pointer to element 42 which is designated as a pointer to lower memory ranges. Similarly, element 43 may correspond to a range of memory blocks that is higher than the range of element 41. Accordingly, element 41 includes a pointer to element 43 which is identified as a pointer to higher memory ranges. Each element in the tree structure 40 similarly includes pointers to elements representing lower and higher memory ranges (except for elements in the bottom-most level of the tree).

Hence, referring again to Figure 7, if it is determined in step 702 that the mapping between the host device and the logical device is exact, then in step 707 the output of the process is simply a logical device ID and a logical block number. If there is not an exact mapping, however, then in step 703 the process accesses the top element (e.g., element 41 in Figure 8) in the virtual-to-logical mapping tree structure. In step 704, the host block number is compared to the element range in the top element in the tree. If the host block number falls below that range, then the process points to the element representing the next lower range (e.g., element 42 in Figure 8) in step 706; step 704 is then repeated. If, in step 704, the host block number falls higher than the range of the top element in the tree, then in step 705 the process points to the element in the tree representing the next higher range (e.g., element 43 in Figure 8); the process then repeats from step 704. If the block number is determined to fall within the range of the top (or current) element in the tree in step 704, then the process proceeds to step 707 by outputting the logical device and logical block number in the top (or current) element.

Figure 9 illustrates the second part of the address mapping process, in which a logical address is mapped to a physical address. In step 901, the logical device ID and logical block number determined in the host-to-logical mapping process (Figure 7) is input. In step 902, the storage controller
5 determines whether the logical device maps exactly to the physical device. As in the above-described routine, this determination can be made based upon the status of a single bit. If the mapping is exact, then the process proceeds to step 907, in which the appropriate physical device ID and physical block number are output. Otherwise, the process proceeds to step 903, in which the
10 routine points to the top element of a logical-to-physical mapping tree. The logical-to-physical mapping tree is similar in structure to, and is accessed in essentially the same manner as, the host-to-logical mapping tree of Figure 8. In step 904, the logical block number is compared to the element range in the top element in the tree. If the logical block number falls below that range,
15 then the process points to the element representing the next lower range in step 906; step 904 is then repeated. If the logical block number falls higher than the range of the top element in the tree in step 904, then in step 905 the process points to the element in the tree representing the next higher range; the process then repeats from step 904. If the block number is determined to
20 fall within the range of the top (or current) element in the tree in step 904, then the process proceeds to step 907 by outputting the correct physical device and physical block number in the top (or current) element.

At various times, the storage controller must read data from one or more MSDs into its internal memory ("the cache") or transfer data from the
25 cache to one or more MSDs; these procedures are known as "staging" and

"destaging", respectively. Staging is generally performed, for example, when executing a host read request. Destaging is generally performed when executing a host write request.

Figure 10 illustrates a routine for staging data, which can be performed by a storage controller of the present invention. In step 1001, the storage controller selects the "nearest" free device housing the data of interest. In this context, the term "nearest" refers to the mass storage device associated with the smallest overall access time or earliest availability, rather than referring to physical distance. (Note, however, that physical distance may be a factor in the total access time.) In a given situation, the nearest storage device may be the local memory within the local storage controller, an MSD within the local storage array, or an MSD within a remote storage array. After selecting the nearest free device, in step 1002 the logical block address of the data is mapped to a physical block address in the manner described above. In step 1003, the block or blocks are read from the selected device using the physical block address. If the read is successful (step 1004), then in step 1005 a read error recovery routine is called if any block was bad. The read error recovery routine is discussed below in connection with Figure 11. If the read is not successful, and no additional copies of the data are available (step 1010), then a read error message is presented to the host and the routine exits in step 1011. If the read was unsuccessful and additional copies are available, then the routine repeats from step 1002 by mapping the block of the next available copy to a physical block. After a successful read (step 1004), and after calling the read error recovery routine if necessary (step 1005), then in step 1006 it is determined whether the staged data was requested by the host; if not, the

Figure 12 illustrates a routine for destaging data. It is appropriate at this point to consider certain signals that are used in conjunction with destaging. Associated with each block of data are the two single-bit binary quantities, Exclusive Lock and Shared Lock. Only a single device is permitted access to a given block of data for purposes of performing a write operation; this constraint is implemented using Exclusive Lock. A write on a block is performed by obtaining an Exclusive Lock on that block. For purposes of writing to a block, Exclusive Lock for that block can be set for only if neither Shared Lock nor Exclusive Lock is already set. When a write operation on a block is initiated, Exclusive Lock is set for that block, so that no other device may access that block during the write operation. In contrast, multiple devices are permitted to read from a given block of data at the same time. A read on a block is performed by obtaining a Shared Lock on that block. A device may obtain a Shared Lock on a data block, regardless of the current status of Shared Lock, provided Exclusive Lock for that block is not already set. When a read is initiated on a given block, Shared Lock for that block is set, so that no device may write to that block during the read operation.

The storage controller can optionally be configured to delay and "merge" multiple destaging tasks to avoid multiple updates of the same physical blocks. In other words, instead of executing each destaging task, the storage controller may simply accumulate and log several destaging tasks in internal memory without executing them, and then subsequently perform a single "merged" destaging task to implement the aggregate effect of the accumulated destaging tasks.

Ending status is then presented to the host and the storage controller disconnects from the bus in step 1213. The routine then proceeds from step 1210.

Figure 13 illustrates the write error recovery routine. In response to a write error, in step 1301 the faulty block is reassigned and written to a new block; this step is accompanied by verification of the write operation. If the write is successful (step 1302), then in step 1306 the new location of the block is logged and the routine exits. If the write was not successful, then the data is written to the internal MSD (i.e., MSD 13 in Figure 3) in step 1303. Note that the internal MSD of the storage controller is generally not accessible to a host. After writing to the internal MSD, it is determined in step 1304 whether a hot spare MSD is available for the new copy. If no hot spare is available, then in step 1307 the location of the data is logged and the routine exits. If a hot spare is available, then in step 1305 an "establish" routine is performed to establish a new copy of the data on the hot spare. The establish routine is discussed below in connection with Figure 18. The routine exits after establishing the new copy.

Figure 14 illustrates a routine for performing a standard read operation in response to a request by a host. In step 1401, a storage controller receives a data read request from a host. Next, in step 1402 the standard checks routine (see Figure 6) is performed. If any of the standards checks fail, then the routine exits in step 1410, presenting an error status message to the host. If none of the standard checks fail, then in step 1403 the storage controller identifies the appropriate logical device based on the host interface and virtual block number received from the host in step 1403. This is done in the

manner discussed above in connection with Figure 8. Next, if the requested block is cached (stored in the storage controller's internal memory 24), then it is determined in step 1405 whether Exclusive Lock is set for the requested block. If Exclusive Lock is set or if the block was not cached, then in step 1411 the storage controller disconnects from the host, and in step 1412 a staging task is queued. The routine exists after queuing a staging task. If Exclusive Lock was not set in step 1405, then in step 1406 Shared Lock is set and the data is transferred to the host from the cache. After transferring the data to the host, Shared Lock is cleared in step 1407. Then, it is determined in step 1408 whether additional blocks have been requested. If so, the routine repeats from step 1403. If no additional blocks have been requested, then in step 1409 an ending status message is presented to the host, and the storage controller disconnects from the bus 8. The routine then exits.

As noted above, the storage controller may maintain multiple copies of a given set of data on different MSDs (which fact may be transparent to the hosts). Consequently, to improve performance, the storage controller can optionally be configured to concurrently read different subsets (i.e, blocks) of the set of data from different copies that are stored on different devices. This technique, which may be referred to as "spreading" reads, may reduce the overall time required to read the complete set of data in some cases.

Figure 15 illustrates a routine for performing a standard write operation in response to a request from a host. In step 1501, the storage controller receives a write request from a host. Next, in step 1502 the standard checks (Figure 6) are performed. If any of the standard checks fail, then the routine exits in step 1513, presenting an error status message to the host.

Assuming none of the standard checks fail, then in steps 1503 and 1504, respectively, it is determined whether the device or blocks to be written to are designated as read-only. If the device or any of the blocks are read-only, then the routine exits, presenting an "illegal request" message to the host (steps

5 1514 and 1515). If neither the device nor any block is read-only, then in step 1505 the storage controller maps the host interface to the appropriate logical device using the technique described above (Figure 7). After identifying the appropriate logical device, it is determined in step 1506 whether the block is already cached in the storage controller's internal memory 24. If the block is
10 already cached, then in step 1507 it is determined whether it possible to obtain Exclusive Lock on the requested block. As noted above, in order to obtain Exclusive Lock, neither Shared Lock nor Exclusive Lock for that block can be set. If Exclusive Lock can be obtained, then in step 1508 the data is transferred from the host to the appropriate block in the internal cache 24. If Exclusive
15 Lock is not possible, then in step 1516 the storage controller disconnects from the host and waits for Exclusive Lock in step 1517. Once Exclusive Lock is obtained, then in step 1518 the storage controller reconnects to the host and, in step 1508, transfers the data from the host to the internal cache 24, as described above.

20 Referring again to step 1506, if the block to be written was not already cached, then in step 1519 it is determined whether the write operation is an update write (i.e., the writing of a partial block). With SCSI protocol, partial blocks cannot be written; only complete blocks can be written. However, other protocols which might be used in conjunction with the present
25 invention might permit an update write. Accordingly, if the requested write

specific job. In such cases, the storage controller simply writes the data to its internal memory 24 and then, if required, presents ending status to the host.

As mentioned above, a local storage controller of the present invention can communicate with one or more remote storage controllers for purposes of accessing a remote storage array. Figure 16 illustrates a routine which can be performed by a local storage controller for transferring data to a remote storage controller. If there is data to be sent (step 1601) to a remote storage controller, then in step 1602 the local storage controller transmits the data onto the link 9 in packet form and waits for an acknowledgment from the remote storage controller. If a wait time-out occurs (step 1603) before an acknowledgment is received, then in step 1612 a line error recovery routine is called. The line error recovery routine is discussed below in connection with Figure 17. If no wait time-out occurs before the acknowledgment is received, then if a request complete message is required by the host (step 1604) the storage controller waits in step 1605 for a "request complete" message from the remote storage controller. If no "request complete" message is received before a wait time-out occurs (step 1606), then the storage controller calls the line error recovery routine in step 1607. If no "request complete" message is required, or if no time-out occurs before the "request complete" message arrives, the routine repeats from step 1601.

Figure 17 illustrates the line error recovery routine. In step 1701, a communications data packet is sent to the remote storage controller. If the transmission is successful (step 1702), then if alternate paths to that storage controller exist (step 1703), the routine repeats from step 1701 using an alternate path. Otherwise, the routine exits in step 1704 with the notification

of an error condition. If the transmission was successful in step 1702, then the routine exits.

Figure 18 illustrates the establish routine for establishing (ensuring consistency between) copies. Assume now that at least one copy of a given file has already been created from an "original" (which may be a previously-created copy). In step 1801, the storage controller identifies one or more contiguous blocks of data in the file. In step 1802, it is determined whether the corresponding blocks are identical between the original and the copy. If so, then in step 1803 one or more other contiguous blocks of data are examined, starting from the next block copy boundary. Next, in step 1804, if the extent of the storage device has not been reached, then the routine repeats from step 1801. If the extent of the device has been reached, then the copy has been established and the routine exits. If the blocks are determined not to be identical in step 1802, then in step 1805 the storage controller obtains (sets) Exclusive Lock for each of those blocks and reads the blocks from the source file. If the read is successful (step 1806), then Exclusive Lock is exchanged for Shared Lock and the blocks are written to the destination device in step 1808. If the read was not successful, then in step 1807 the read error recovery routine is called. After exchanging Exclusive Lock for Shared Lock and writing to the destination (step 1808), in step 1809 it is determined whether the write was successful. If not, the write error recovery routine is called in step 1811. If the write was successful, then in step 1810 the corresponding blocks of the two different copies are marked as identical and the routine proceeds from step 1803 (described above). Note that during the establish routine

double buffering may occur; that is, one or more subsequent reads may overlap in time with previous writes in order to improve performance.

A storage controller of the present invention also implements automatic block error detection and reassignment. When the storage controller is idle, it automatically examines data stored in any of the MSDs for errors and corrects detected errors, when possible. Figure 19 illustrates a routine performed by the storage controller for implementing these functions. In step 1901, the storage controller attempts to read a block of data from a given location. Note that this given location may be within the local storage array or it may be within one of the remote storage arrays. This attempted read (step 1901) operation is not performed in response to any request from a host computer; it is performed on the storage controller's own initiative in response to the system becoming idle. In step 1902, it is determined whether the block was successfully read. If the block was successfully read, then in step 1903 it is determined whether all blocks have been examined; if so, the routine ends. If all blocks have not been examined, then the storage controller attempts to read another block of data in step 1904, and the process repeats from step 1902. If the block was not read successfully, then in step 1905 the storage controller calls the read error recovery routine (Figure 11).

The storage controller of the present invention can also be used to migrate data from one storage array to another. The storage controller can be configured to perform data migration automatically in a manner that is transparent to, and non-disruptive of, any host computer. Further, in contrast with prior systems, which maintain only a single back-up copy of

given file, the storage controller of the present invention allows two or more back-up copies of a file to be created and maintained in addition to a primary copy. Any copy is then accessible to any host computer connected in the system.

- 5 A storage controller of the present invention further allows data blocks to be write protected, so that a block cannot be modified from any host computer. Write protection may be desirable for purposes such as virus protection or implementation of security firewalls. Write protection can be achieved by configuring the storage controller appropriately at set-up time or
- 10 by inputting a write protect command to the storage controller from a host computer.

- Thus, a storage controller capable of providing multiple host computers system with access to multiple storage arrays has been described. Although the present invention has been described with reference to specific
- 15 exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention as set forth in the claims. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

- 1 1. A storage controller, comprising:
 - 2 a host interface providing a data interface with a host processing system
 - 3 independent of a device type of the host processing system;
 - 4 a storage interface providing a data interface with at least one storage
 - 5 array, each of said at least one storage arrays including at least one mass
 - 6 storage device, the storage interface further providing the data interface with
 - 7 said at least one storage array independent of a device type of the mass storage
 - 8 devices; and
 - 9 control circuitry coupled to the host interface and the data interface, the
 - 10 control circuitry configured to enable the host processing system to access each
 - 11 of said at least one storage array.
- 1 2. A storage controller according to claim 1, wherein the host processing
 - 2 system has a hardware configuration and a software configuration, wherein
 - 3 the storage controller is configured to enable the host processing system to
 - 4 access said at least one storage array regardless of the hardware configuration
 - 5 or the software configuration of the host processing system.
- 1 3. A storage controller according to claim 1, wherein the storage controller is
 - 2 configured to enable the host processing system to access a mass storage
 - 3 device having any of a plurality of different device types.

1 4. A storage controller according to claim 1, wherein the storage controller is
 2 configured to:
 3 receive from the host processing system a request specifying a virtual
 4 storage location of data; and
 5 determine a physical storage location of the data based on the request.

1 5. A storage controller according to claim 4, wherein the virtual storage
 2 location corresponds to a plurality of physical storage locations distributed
 3 among a plurality of mass storage devices.

1 6. A storage controller according to claim 1, wherein:
 2 the storage controller is coupled to a local storage array;
 3 the storage controller is further coupled to a remote storage controller;
 4 the remote storage controller is coupled to a remote host processing
 5 system; and
 6 the storage controller is further configured to provide the remote host
 7 processing system with access to the local storage array.

1 7. A storage controller according to claim 6, wherein the remote storage
 2 controller is coupled to a remote storage array, wherein the storage controller
 3 is further configured to provide the local host processing system with access
 4 to data stored in the remote storage array.

1 8. A storage controller according to claim 1, wherein each of a plurality of
 2 copies of a set of data is stored on a different one of the mass storage devices,

3 each copy of the set of data having a plurality of subsets, and wherein the
4 storage controller is configured to access a different subset of the set of data in
5 each of the plurality of copies to access the complete set of data.

1 9. A storage controller according to claim 1, wherein a first one of the storage
2 arrays may comprise mass storage devices having a different storage medium
3 than a second one of the storage arrays.

1 10. A storage controller according to claim 1, wherein:
2 the plurality of processing systems comprises at least three host
3 processing systems; and
4 the plurality of storage arrays comprises at least three storage arrays.

1 11. A storage controller according to claim 1, wherein the control circuitry is
2 configured to communicate with a plurality of other storage controllers.

1 12. A data storage controller for providing a plurality of host processing
2 systems with access to a plurality of storage arrays, each of the storage arrays
3 including a plurality of mass storage devices, the data storage controller
4 comprising:
5 a processor; and
6 a memory coupled to the processor, the memory having instructions
7 stored therein which, when executed by the processor, cause the data storage
8 controller to:

9 provide a data interface with a local host processing system
 10 independent of a device type of the local host processing system;
 11 emulate a local storage array to the local host processing system;
 12 provide a data interface with the local storage array independent
 13 of the type of mass storage devices in the local storage array;
 14 emulate the host processing system to the mass storage devices
 15 of the local storage array;
 16 provide an interface to a remote data storage controller, the
 17 remote data storage controller for coupling to a remote storage array; and
 18 control communication of data such that any of the plurality of
 19 host processing systems can access data in any of the plurality of storage
 20 arrays.

1 13. A data storage controller according to claim 12, wherein the instructions
 2 further comprise instructions which, when executed by the processor, cause
 3 the data storage controller to receive from one of the host processing systems
 4 a request specifying a virtual storage location of a data block and to determine
 5 a real storage location of the data block based on the virtual storage location.

1 14. A data storage controller according to claim 12, wherein the instructions
 2 further comprise instructions which, when executed by the processor, allow
 3 the data storage controller to provide any of the host processing systems with
 4 access to any of the storage arrays using virtual addresses, such that actual
 5 locations of stored data in the storage arrays are transparent to the host
 6 processing systems.

1 15. A data storage controller according to claim 14, wherein one of the virtual
2 addresses may be mapped to a plurality of real storage locations distributed
3 among a plurality of the mass storage devices.

1 16. A data storage controller according to claim 12, wherein the instructions
2 further comprise instructions which, when executed by the processor, cause
3 the data storage controller to provide a remote host processing system of the
4 plurality of host processing systems with access to the local storage array, the
5 remote host processing system coupled to the remote data storage controller.

1 17. A data storage controller according to claim 16, wherein the remote
2 storage controller is coupled to a remote storage array, and wherein the
3 instructions further comprise instructions which, when executed by the
4 processor, cause the data storage controller to provide the local host
5 processing system with access to data stored in the remote storage array.

1 18. A data storage controller according to claim 12, wherein the data storage
2 controller is coupled to the host processing system via a SCSI interface.

1 19. A data storage controller according to claim 18, wherein the data storage
2 controller is coupled to the local storage array via a SCSI interface.

1 20. A data storage controller according to claim 12, further comprising means
2 for connecting a mass storage device having any of a plurality of different
3 device types to the host processing system.

1 21. A data storage controller according to claim 20, wherein a first one of the
2 storage arrays may comprise mass storage devices having a different storage
3 medium than a second one of the storage arrays.

1 22. A data storage controller according to claim 12, wherein the local host
2 processing system has a hardware configuration and a software configuration,
3 the data storage controller further comprising means for providing the data
4 interface to the host processing system, such that the local host processing
5 system can access the local storage array independent of both the hardware
6 configuration and the software configuration of the local host processing
7 system.

1 23. A data storage controller according to claim 12, wherein:
2 the plurality of processing systems includes at least three host
3 processing systems; and
4 the plurality of storage arrays includes at least three storage arrays.

1 24. A data storage controller according to claim 12, wherein the instructions
2 which cause the data storage controller to control communication of data
3 comprise instructions which, when executed by the processor, cause the data

4 storage controller to communicate with a plurality of other data storage
5 controllers.

25. In a storage controller coupled to a plurality of host processing systems and a plurality of storage arrays, a method of providing the plurality of host processing systems with access to the plurality of storage arrays, the plurality of host processing systems including a local host processing system, the plurality of storage arrays including a local storage array, each of the storage arrays including a plurality of mass storage devices, the method comprising:

- providing a data interface with the local host processing system;
- emulating the local storage array;
- providing a data interface with the local storage array, wherein the local storage array may include any of a plurality of different storage mediums;
- emulating the local host processing system;
- providing an interface to a remote data storage controller coupled to a remote storage array of the plurality of storage arrays; and
- controlling communication of data such that any of the plurality of host processing systems can access data in any of the plurality of storage arrays.

1 26. A method according to claim 25, wherein the local storage array includes
2 mass storage devices of a first storage medium, and the remote storage array
3 includes mass storage devices of a second storage medium different from the
4 first storage medium.

1 27. A method according to claim 25, further comprising the step of providing
2 the host processing system with access to the storage arrays using virtual
3 addressing, wherein physical locations of data stored in the storage arrays are
4 transparent to the host processing systems.

1 28. A method according to claim 25, wherein the plurality of processing
2 systems comprises at least three host processing systems.

1 29. A method according to claim 28, wherein the plurality of storage arrays
2 comprises at least three storage arrays.

1 30. A method according to claim 25, wherein the step of controlling
2 communication of data comprises the step of communicating with a plurality
3 of other storage controllers to cause a plurality of copies of a set of data to be
4 stored distributed among the plurality of storage arrays.

1 31. A system comprising:
2 a plurality of host processing systems;
3 a plurality of storage arrays, each of the storage arrays including at least
4 one mass storage device; and
5 a plurality of storage controllers, each of the storage controllers coupled
6 to one of the host processing systems, coupled to one of the storage arrays, and
7 coupled to another one of the storage controllers, each of the storage
8 controllers having a first interface with one of the host processing systems,

10 including means for supporting any one of a plurality of types of mass storage
11 media in a corresponding one of the storage arrays.

1 36. A system according to claim 35, wherein each of the storage controllers
2 further comprises means for supporting any one of a plurality of
3 configurations of a corresponding host processing system.

1 37. A system according to claim 36, each of the storage controllers comprising
2 an interface with a host processing system for emulating a mass storage
3 device and an interface with a local storage array for emulating a host
4 processing system.

1 38. A system according to claim 37, each of the storage devices further
2 comprising translation means for representing the storage arrays as at least
3 one virtual device to the host processing systems.

1 39. A system according to claim 35, wherein the storage controllers function
2 as peers with respect to each other.

1 40. A system according to claim 35, wherein the plurality of processing
2 systems comprises at least three host processing systems.

1 41. In a storage controller for providing at least one host processing system
2 with access to a plurality of storage arrays, a method of providing back-up
3 copies of a set of data, the method comprising the steps of:

1 46. A method according to claim 41, further comprising the steps of, at
 2 predetermined time intervals, automatically creating a new copy of the set of
 3 data and discarding a previously-created copy of the set of data.

1 47. A method according to claim 41, further comprising the step of
 2 automatically migrating one of the copies of the data from a first one of the
 3 storage arrays to a second one of the storage arrays in a manner that is
 4 transparent to and non-disruptive to each of the at least one host processing
 5 systems.

1 48. A method according to claim 41, further comprising the step of write-
 2 protecting one of the copies of the data to prevent said one of the copies from
 3 being modified from any of the at least one host processing systems.

1 49. A method according to claim 41, further comprising the steps of:
 2 receiving an access request from one of said at least one host processing
 3 systems, the access request including a virtual address of one of the copies;
 4 and
 5 mapping the virtual address to at least one physical address in at least
 6 one of the plurality of storage arrays.

1 50. A method according to claim 49, wherein the mapping step comprises the
 2 step of mapping the virtual address to a plurality of physical addresses
 3 distributed among a plurality of the storage arrays.

8 if the failure resulted from a failure of one of the mass storage devices
 9 automatically identifying an alternate mass storage device and attempting to
 10 write the data to the alternate mass storage device; and
 11 if the failure resulted from a block failure, automatically attempting to
 12 write the data to an alternate block.

1 54. In a storage controller coupled to a plurality of mass storage devices, a
 2 method of accessing a stored set of data, wherein each of a plurality of copies
 3 of the set of data is stored on a different one of the mass storage devices, the
 4 set of data having a plurality of subsets, the method comprising the step of
 5 accessing a different subset of the set of data in each of the plurality of copies
 6 to access the complete set of data.

1 55. In a storage controller coupled to a host and to a plurality of mass storage
 2 devices, a method of accessing stored data, the method comprising the steps
 3 of:
 4 receiving a single virtual address from the host; and
 5 mapping the single virtual address to a plurality of physical addresses
 6 corresponding to a plurality of storage locations distributed among the
 7 plurality of mass storage devices.

1 56. In a storage controller coupled to a plurality of mass storage devices, a
 2 method of destaging data to at least one of the mass storage devices, the
 3 method comprising the steps of:
 4 logging a plurality of destaging tasks; and

5 executing a single destaging task representative of the plurality of
6 destaging tasks.

1 57. A storage controller for providing a host processing system with access to
2 a plurality of mass storage devices, the storage controller comprising:
3 a processor;
4 a memory coupled to the processor; and
5 an input interface coupled to the processor, the input interface for
6 receiving first information responsive to user inputs for configuring the
7 processor to control functions of the storage controller.

1 58. The storage controller of claim 57, wherein the input interface comprises
2 an interface for connecting an input device directly to the storage controller,
3 the user inputs originating from the input device.

1 59. The storage controller of claim 57, wherein the input interface comprises
2 a network interface for connecting the storage controller to a remote
3 processing system, wherein the first information is received by the storage
4 controller from the remote processing system.


1 60. The storage controller of claim 57, wherein the input interface comprises
2 an interface to the host processing system, the first information originating
3 from the host processing system.

1 61. The storage controller of claim 57, wherein the input interface is further
2 for receiving second information responsive to user inputs requesting status
3 information, wherein the processor is configured to cause the storage
4 controller to respond to the second information by outputting the requested
5 status information.

46 "EE0" 28882880

08/23/88

FIG. 1



08/23/88

FIG. 2





FIG. 3



08/884688

FIG. 4





FIG. 5



03/03/2008 15:00:00

Standard Checks




FIG. 6

08/12/88




FIG. 7

40




FIG. 8




FIG. 9

Staging

00/00000000





FIG. 10

Read Error Recovery

02/07/88



00000000.033197

FIG. 11

Destaging




FIG. 12

08/828888




FIG. 13

Host Data Read




FIG. 14

Host Data Write

08/228888




FIG. 15

08/828888





FIG. 16



08/828888

Copy Establish




FIG.18




FIG. 19