

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
21 December 2000 (21.12.2000)

PCT

(10) International Publication Number
WO 00/77606 A2

(51) International Patent Classification⁷: G06F 3/06

151st Lane Northwest, Ramsey, MN 55303 (US). GUHA, Alok; 814 West Mulberry Street, Louisville, CO 80027 (US). HINTON, Walter H.; 9970 Newton Court, Westminster, CO 80030 (US). STAMAS, Bradley; 18178 82nd Avenue North, Maple Grove, MN 55311 (US).

(21) International Application Number: PCT/US00/15579

(22) International Filing Date: 6 June 2000 (06.06.2000)

(25) Filing Language: English

(81) Designated States (*national*): AU, CA, CN, JP.

(26) Publication Language: English

(84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

(30) Priority Data:
09/330,960 11 June 1999 (11.06.1999) US


Published:
— Without international search report and to be republished upon receipt of that report.

(71) Applicant: STORAGE TECHNOLOGY CORPORATION [US/US]; One StorageTek Drive, MS-4309, Louisville, CO 80028-4309 (US).


(72) Inventors: BAKKE, Mark, A.; 11738 Gentilly Road, Maple Grove, MN 55369 (US). FIORE, Edward, J.; 7861

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: INTELLIGENT STORAGE AREA NETWORK



(57) Abstract: An intelligent storage area network permits a plurality of host computer systems to access a plurality of storage devices connected through at least one interconnection network. Presenters connected to the network present data to host computer systems in the host computer specified data format. A data manager connected to the network receives a storage access request, examines the storage devices, determines a response to the request based on the examination, and forwards the response.



WO 00/77606 A2

INTELLIGENT STORAGE AREA NETWORK

TECHNICAL FIELD

5

The present invention relates to storage area networks interconnecting host computers and storage devices.

BACKGROUND ART

10

An increase in the demand for storage coupled with a decrease in the cost per bit of storage has greatly increased the amount of storage accessible within a computer system. Adequate performance in traditional computer systems, in which storage is directly connected to the host, is difficult to achieve with increasing storage volume. A significant amount of host resources are spent managing storage devices as well as in setting up storage access through the host. Human resources are also expended managing storage allocation, configuration, ad the like. Storage management problems may be worsened when hosts are interconnected, such as through a LAN, where each host may receive requests for storage access from clients connected through other hosts. These clients may be operating under different protocols or operating systems, such as Microsoft WINDOWS NT or UNIX, may expect data in varying formats, may expect storage with particular performance specifications, or may expect to be connected to a particular storage device type. An additional problem arises when a host fails, decoupling connected storage devices from the rest of the computer system.

25

A proposed solution to these problems is the storage area network (SAN). A SAN interconnects one or more hosts with a plurality of storage devices. Depending upon the configuration, SANs may decrease traffic on the LAN, may ease data compatibility requirements, and may provide for graceful degradation in the event of a host failure. SANs may be implemented using high speed connections such as a 100 megabyte per second Fibre Channel loop.

30

A challenge in the design of a SAN is the location and configuration of control operations needed to set up and tear down storage connections, arbitrate storage access, and manage storage devices. In one design, SAN management resides in one or more hosts. A managing host controls access to a subset of storage devices connected to the SAN. Problems with this technique include resource burdens for managing the

35

SAN and reliability in the event of a host failure. Another control design option is to place intelligence in storage devices connected to the SAN. Each storage device is responsible for arbitrating access to data stored within the storage device. Problems with this design include the need to support multiple protocols in heterogenous computer systems, the difficulty in selecting appropriate storage devices to handle initial storage requests, and the difficulty in finding the appropriate provider for desired services.

What is needed is a SAN system that can manage access requests from heterogenous hosts having different protocols to distributed heterogenous storage devices. Management tasks should not present a substantial burden to host processors or storage administrators. Data should be held in storage devices using a format best suited to each storage device and presented to each host in response to an access request in a format best suited to the host.

DISCLOSURE OF INVENTION

It is an object of the present invention to place storage management within the SAN.

It is another object of the present invention to manage distributed heterogenous storage.

It is another object of the present invention to support access requests from hosts with multiple protocols.

It is still another object of the present invention to present data to each host in an acceptable format.

It is yet another object of the present invention to reduce host resource requirements for storage access and management.

It is a further object of the present invention to examine storage devices to determine appropriate configurations for satisfying storage access requests.

It is a still further object of the present invention to provide centralized access to management services.

In carrying out the above objects and other objects and features of the present invention, a storage area network is provided. The SAN includes at least one interconnection network. Storage devices and host computer systems are connected to at least one network. Each host can access data on one or more storage devices in a host specified data format. At least one presenter is connected to at least one network. Each presenter presents data to a host computer system in the host computer specified data format. At least one manager is connected to at least one of the networks. Each manager processes a storage access request. Access request processing may include receiving a storage access request, determining a response to the request based on an examination of one or more storage devices, and forwarding the response.

In an embodiment of the present invention, the manager resolves conflicts between hosts requesting access to the same storage device. The manager may also allocate physical space on one or more storage devices in response to the received storage access request.

In another embodiment of the present invention, the storage area network includes at least one mover connected to one or more networks. Each mover reads data from one or more storage devices and writes the data into one or more storage devices. The manager receives a request to move data, examines the storage devices, and determines one or more movers to perform the data move request. In a refinement, at least one storage device is connected to at least one network through a mover.

In yet another embodiment of the present invention, the manager receives a request to modify storage allocation from a host connected through a presenter. The manager examines one or more storage devices and determines allocation of physical space on the storage devices satisfying the received request. The manager notifies the presenter through which the host is connected of the determined allocation.

In a further embodiment of the present invention, the at least one network includes one or more data networks for carrying data between hosts and storage devices and one or more communication networks transporting storage access requests to the manager. In a refinement, a user interface communicates with the manager through the communication network.

In a still further embodiment of the present invention, the manager provides a plurality of services to the hosts. One of the services is a location service permitting a host to find the remaining services.

5 A storage area network is also provided which includes at least one switch defining at least one data network. The storage area network also includes a communication network. Storage devices and hosts are interconnected through at least one of the data networks. Each host is also connected to the communication network. At least one presenter is connected to at least one data network and the communication network. A manager is connected to the communication network. The manager receives a storage access request originating in a host. The manager examines the storage devices and determines a response to the request. The response is forwarded to the originating host.

10 The above objects and other objects, features, and advantages of the present invention are readily apparent from the following detailed description of the best mode for carrying out the invention when taken in connection with the accompanying drawings.

BRIEF DESCRIPTION OF DRAWINGS

20 FIGURE 1 is a schematic diagram of a generalized intelligent SAN according to an embodiment of the present invention;

FIGURE 2 is a schematic diagram of an intelligent SAN implemented using a crossbar switch according to an embodiment of the present invention;

25 FIGURE 3 is a schematic diagram of an intelligent SAN implemented using a crossbar hub with associated processors according to an embodiment of the present invention; and

30 FIGURE 4 is a schematic diagram of an intelligent SAN implemented using a Fibre Channel switch hub according to an embodiment of the present invention.

BEST MODE FOR CARRYING OUT THE INVENTION

35 Referring to Figure 1, a schematic diagram of a generalized intelligent SAN according to an embodiment of the present invention is shown. An intelligent SAN, shown generally by 20, includes one or more interconnection networks 22. A plurality of storage devices, each of which is indicated by 24, is connected to at least

one of the interconnection networks 22. Storage devices 24 include magnetic disks, magnetic tape, optical disks, solid state memory, and the like. A plurality of host computer systems, each of which is designated as 26, are also connected to one or more of interconnection networks 22. Interconnection networks 22 permit each host 26 to
5 access data on at least one storage device 24 in a host specified data format. Hosts 26 may be interconnected by interconnection network 28, such as a LAN, which is not part of any interconnection network 22.

Intelligent SAN 20 also includes at least one presenter 30 connected to
10 one or more of interconnection networks 22. Each presenter 30 presents data to at least one host 26 in the host computer specified data format. Thus, host 26 accessing data through presenter 30 need not know or be concerned with the data structure used for each storage device 24. Host 26 may be connected to interconnection networks 22 through presenter 30 or may access presenter 30 through a direct connection to
15 interconnection networks 22. Host 26 may also include presenter 30.

Intelligent SAN 20 further includes at least one manager 32 connected to one or more interconnection networks 22. Each manager 32 receives a storage access request, examines at least one storage device 24, determines a response based on the
20 examination, and forwards the response. Typically, the access request comes from host 26 and the subsequent response is forwarded to requesting host 26. Host 26 may then directly access storage device 24. Manager 32 may be part of interconnection network 22, may be integrated into host 26, or may be separate from interconnection network 22 and host 26.

In an embodiment of the present invention, manager 32 resolves conflicts
25 between hosts 26 requesting access to the same storage device. For example, only one host 26 can typically access a tape storage device 24 at one time. If multiple hosts request the same tape storage device 24, manager 32 determines which host 26 will be granted access.
30

In an embodiment of the present invention, manager 32 allocates physical
35 space on one or more storage devices 24 in response to the received storage access request. The request may include certain parameters including amount of storage required, access time, data transfer rate, cost per megabit, relative reliability, and the like. Manager 32 attempts to find one or more storage devices 24 which best meets the specified parameters. A suitable method for allocating storage devices 24 is disclosed

in U.S. patent application Serial No. 09/232,431 entitled "Intelligent Storage Data Manager" to Mark A. Bakke *et al.*, filed January 15, 1999, which is incorporated by reference herein. If host 26 is connected through presenter 30, presenter 30 is notified of the determined allocation on storage devices 24.

5

In an embodiment of the present invention, intelligent SAN 20 includes at least one mover 34. Each mover 34 reads data from at least one storage device 24 and writes the data into at least one storage device 24. Manager 32 receives a request to move data, examines at least one storage device 24, and determines at least one
10 mover 34 to perform the data move request. Mover 34 relieves host 26 from performing data transfer requests such as tape backup, defragmentation, and the like. One or more storage devices 24 may be connected to interconnection networks 22 through mover 34. Mover 34 may also access storage device 24 which is directly connected to interconnection networks 22.

15

In a preferred embodiment of the present invention, interconnection networks 22 are split into one or more communication networks 36 for transporting storage access requests to manager 32 and one or more data networks 38 for carrying data between hosts 26 and storage devices 24. Each data network 38 is a high speed
20 network such as Fibre Channel, Small Computer System Interface (SCSI), Intelligent Peripheral Interface (IPI), High Performance Parallel Interface (HPPI), Framing Protocol, Internet Protocol (IP), ATM Adaption Layer for computer Data (AAL5), Link Encapsulation (FC_LE), Single Byte Command Code Set Mapping (SBCCSM), IEEE 802.2, and the like. Each communication network 36 typically has a lower
25 bandwidth than data networks 38 and may be, for example, Ethernet implementing TCP/IP protocol. If more than one communication network 36 is used, each network has at least one manager 32. Managers 32 on separate communication networks 36 are interconnected to permit communication between managers 32. More than one manager 32 may be placed on each communication network 36 to provide redundancy or to
30 provide different services.

In an embodiment of the present invention, user interface 40 is connected to at least one communication network 36. User interface 40 permits a SAN administrator the ability to monitor and control at least one manager 32.

35

In various embodiments of the present invention, manager 32 performs a variety of services. Manager 32 controls data network 38, establishing and breaking

connections to form subnetworks within and connections through interconnection networks 22. Manager 32 provides conflict resolution, storage space allocation, data transfer scheduling, and monitors the health of interconnection networks 22.

5 Manager 32 may provide a centralized location service by which host 26 can access other service provided by manager 32, presenter 30, mover 34, and storage device 24. Therefore, host 26 need not maintain and update a list of services and service locations.

10 Referring now to Figure 2, a schematic diagram of an intelligent SAN implemented using a crossbar switch according to an embodiment of the present invention is shown. Data network 38 is implemented using at least one crossbar switch 60. Crossbar switch 60 may be implemented using a Fibre Channel switch such as, for example, the SILKWORM SWITCH from Brocade Communication Systems, Inc. of San
15 Jose, California.

 Referring now to Figure 3, a schematic diagram of an intelligent SAN implemented using a crossbar with associated processors according to an embodiment of the present invention is shown. Crossbar 80 includes crossbar switch 82 which
20 functions in a manner similar to crossbar switch 60 to interconnect any two or more of processors 84. Each processor 84 is connected to one or more hosts 26 through a high speed channel such as Fibre Channel, SCSI, or the like, or to one or more storage devices 24 through a high speed channel such as Fibre Channel, SCSI, ESCON, or the like. Each processor 84 handles real-time forwarding, conversion, and transformation
25 of data passed between host 26 and storage device 24. Each processor 84 may be further programmed to function as presenter 30, mover 34, or both.

 In a preferred embodiment, crossbar 80 includes manager 32 connected to each processor 84 through internal communication network 86, typically
30 implemented as an Ethernet network implementing TCP/IP protocols. Manager 32 may also be connected to hosts 26 and storage devices 24 through external communication network 88 also implemented using Ethernet.

 Referring now to Figure 4, a schematic diagram of an intelligent SAN implemented using a Fibre Channel switch hub according to an embodiment of the present invention is shown. Storage devices 24, hosts 26, presenters 30, and movers 34
35 are connected to switch hub 100 through Fibre Channel ports. Manager 32 forms one or more sub-loops 102 in response to access requests from hosts 26.

While embodiments of the invention have been illustrated and described, it is not intended that these embodiments illustrate and describe all possible forms of the invention. Rather, the words used in the specification are words of description rather
5 than limitation, and it is understood that various changes may be made without departing from the spirit and scope of the invention.

WHAT IS CLAIMED IS:

1. A storage area network comprising:

at least one interconnection network;

5 at least one plurality of storage devices, each storage device connected to at least one of the at least one network;

a plurality of host computer systems, each host connected to at least one of the at least one network, each host operative to access data on at least one of the storage devices in a specified data format;

10 at least one presenter connected to at least one of the at least one network, each presenter operative to present data to at least one host computer system in the host computer specified data format; and

at least one manager connected to at least one of the at least one network, each manager operative to process a storage access request.

15 2. A storage area network as in claim 1 wherein storage access requests are processed by the at least one manager operative to:

receive a storage access request,

examine at least one of the plurality of storage devices,

20 determine a response to the request based on the examination, and

forward the response.

25 3. A storage area network as in claim 1 wherein the at least one manager is further operative to resolve conflicts between multiple hosts requesting access to the same storage device.

30 4. A storage area network as in claim 1 wherein the at least one manager is further operative to allocate physical space on at least one storage device in response to the received storage access request.

5. A storage area network as in claim 1 further comprising at least one mover connected to at least one of the at least one network, each mover operative to read data from at least one of the storage devices and write the read data into at least one of the storage devices, the at least one manager further operative to:

35 receive a request to move data,

examine at least one of the plurality of storage devices, and

determine at least one mover to perform the data move request.

6. A storage area network as in claim 5 wherein at least one storage device is connected to at least one network through at least one mover.

5 7. A storage area network as in claim 1 wherein at least one storage device is connected to at least one network through at least one presenter.

8. A storage area network as in claim 1 wherein at least one host is connected to at least one network through at least one presenter.

10

9. A storage area network as in claim 8 wherein the manager is further operative to:

receive a request to modify storage allocation from the at least one host connected through at least one presenter;

15

examine at least one of the storage devices;

determine allocation of physical space on the at least one examined storage devices satisfying the received request to modify storage allocation; and

notify the at least one presenter through which the at least one requesting host is connected of the determined allocation.

20

10. A storage area network as in claim 1 wherein the at least one network comprises:

at least one data network operative to carry data between the hosts and the storage devices; and

25

at least one communication network operative to transport the storage access request to the manager.

11. A storage area network as in claim 10 further comprising a user interface in communication with the at least one manager through the at least one communication network.

30

12. A storage area network as in claim 1 wherein the at least one manager is operative to provide a plurality of services to at least one host, one of the services being a location service permitting the at least one host to find the remaining services.

35

13. A storage area network as in claim 1 wherein the storage access request is received from one of the hosts and the response is forwarded to the host requesting access.

14. A storage area network comprising:
at least one switch defining at least one data network;
a communication network;

5 a plurality of storage devices, each storage device storing data in a storage format;

a plurality of hosts, each host connected to at least one storage device through at least one of the at least one data networks, each host connected to the communication network, each host accessing data in a host format;

10 at least one presenter connected to at least one of the at least one data network and the communication network, each presenter operative to convert data between at least one storage format and at least one host format; and

at least one manager connected to the communication network, each manager operative to process a storage access request.

15 15. A storage area network as in claim 14 wherein storage access requests are processed by the at least one manager operative to:

receive a storage access request,

examine at least one of the plurality of storage devices,

20 determine a response to the request based on the examination, and

forward the response.

25 16. A storage area network as in claim 14 wherein the at least one manager is further operative to resolve conflicts between multiple hosts requesting access to the same storage device.

30 17. A storage area network as in claim 14 wherein the at least one manager is further operative to allocate physical space on at least one storage device in response to the received storage access request.

35 18. A storage area network as in claim 14 further comprising at least one mover connected to at least one of the at least one network and to the communication network, each mover operative to read data from at least one of the storage devices and write the read data into at least one of the storage devices, the at least one manager further operative to:

receive a request from a host to move data,

examine at least one of the plurality of storage devices, and

determine at least one mover to perform the data move request.

device is connected to at least one network through at least one mover.

5 20. A storage area network as in claim 14 wherein at least one storage device is connected to at least one network through at least one presenter.

 21. A storage area network as in claim 14 wherein at least one host is connected to at least one network through at least one presenter.

10 22. A storage area network as in claim 20 wherein the manager is further operative to:

 receive a request to modify storage allocation from the at least one host connected through at least one presenter;

 examine at least one of the storage devices;

15 determine allocation of physical space on the at least one examined storage devices satisfying the received request to modify storage allocation; and

 notify the at least one presenter through which the at least one requesting host is connected of the determined allocation.

20 23. A storage area network as in claim 14 wherein the at least one manager is operative to provide a plurality of services to at least one host, one of the services being a location service permitting the at least one host to find the remaining services.

25 24. A storage area network as in claim 14 further comprising a user interface in communication with the at least one manager through the at least one communication network.

30

An intelligent storage area network permits a plurality of host computer systems to access a plurality of storage devices connected through at least one interconnection network. Presenters connected to the network present data to host computer systems in the host computer specified data format. A data manager connected to the network receives a storage access request, examines the storage devices, determines a response to the request based on the examination, and forwards the response.

5

10




Fig. 1




Fig. 2




Fig. 3




Fig. 4