(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
27 May 2004 (27.05.2004)

PCT

(10) International Publication Number
WO 2004/044753 A1

(51) International Patent Classification⁷: **G06F 12/06**

(21) International Application Number: PCT/US2002/040199

(22) International Filing Date: 16 December 2002 (16.12.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/425,867 12 November 2002 (12.11.2002) US

(71) Applicant *(for all designated States except US)*: **ZETERA CORPORATION** [US/US]; University Research Park, Suite 160, 5251 California Avenue, Irvine, CA 92612 (US).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **FRANK, Charles** [US/US]; Zetera Corporation, 4800 Camino Costado, San Clemente, CA 92673 (US). **LUDWIG, Thomas** [US/US]; Zetera Corporation, 4800 Camino Costado, San Clemente, CA 92673 (US). **HANAN, Thomas** [US/US]; Zetera Corporation, 27022 Manscal, Mission Viejo, CA 92691 (US). **BABBITT, William** [US/US]; Zetera Corporation, 22350 Gavilan Road, Perris, CA 92570 (US).

(74) Agents: RUTAN & TUCKER LLP et al.; Suite 1400, 611 Anton Blvd., Costa Mesa, CA 92626 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ (utility model), CZ, DE (utility model), DE, DK (utility model), DK, DM, DZ, EC, EE (utility model), EE, ES, FI (utility model), FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK (utility model), SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: DATA STORAGE DEVICES HAVING IP CAPABLE PARTITIONS

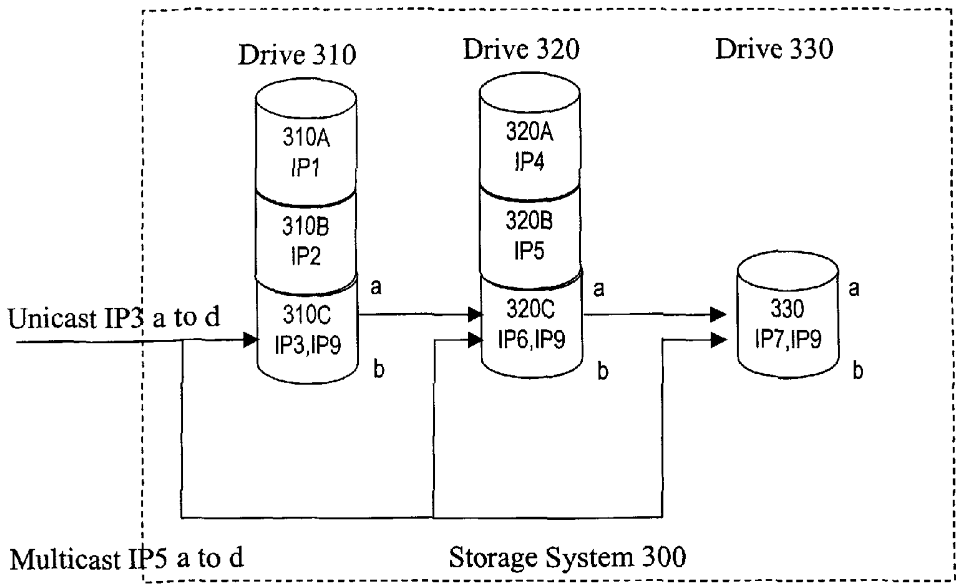

(57) Abstract: A storage device (300) has partitions (310A, 310B, 310C, 320A, 320B, 320C, 330) that are separately addressed by distinct IP addresses (IP1, IP2, IP3, IP4, IP5, IP7, IP9). This allows direct access of the partitions (310A, 310B, 310C, 320A, 320B, 320C, 330), on a peer-to-peer basis, by any other device that can communicate using IP. Preferred storage devices support spanning between or among partitions of the same device, as well as between or among different storage devices. Both multicast and proxy spanning are contemplated. Combinations of the inventive storage devices with each other, and with prior art storage devices are contemplated, in all manner of mirroring and other arrangements. In still other aspects of the invention, a given storage device can comprise one or more types of media, including any combination of rotating and non-rotating media, magnetic and optical, and so forth.

WO 2004/044753 A1

European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*
— *with amended claims*

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.

DATA STORAGE DEVICES HAVING IP CAPABLE PARTITIONS

This application claims the benefit of U.S. provisional application number 60/425867 incorporated herein by reference in its entirety.

Field of The Invention

The field of the invention is data storage devices.

Background of The Invention

There is a trend within the field of electronics to physically (i.e. geographically) disaggregate functionality, and to rely instead on networked resources. Of special interest are resources available over a packet communications network such as the Internet. In addition to the data being transferred, packets include header information such as type of data contained in the packet, i.e. HTML, voice, ASCII, etc., and origination and destination node information. The header information permits error checking, and routing across packet switched networks such as the Internet between devices that may be widely spaced apart. The header information also allows extremely disparate devices to communicate with each other - such as a clock radio to communicate with a computer. Recently published US patent application no. 20020031086, (Welin, March 14, 2002) refers to linking "computers, IP phones, talking toys and home appliances such as refrigerators, microwave ovens, bread machines, blenders, coffee makers, laundry machines, dryers, sweepers, thermostat assemblies, light switches, lamps, fans, drape and window shade motor controls, surveillance equipment, traffic monitoring, clocks, radios, network cameras, televisions, digital telephone answering devices, air conditioners, furnaces and central air conditioning apparatus."

Communications with storage devices has not kept pace with the trend to disaggregate resources. Disk access has always been under the control of a disk operating system such as DOS, or Microsoft® Windows®. Unfortunately, putting the operating system at the conceptual center of all computing devices has resulted in a dependence on such operating systems, and has tended to produce ever larger and more complicated operating systems. Now that many electronic devices, from personal digital assistants to telephones, digital cameras, and game consoles, are becoming smaller and ever more portable, the dependence on large operating systems has become a liability. One solution is to provide a stripped-down operating system that requires much less overhead. Microsoft® CE® is an example. That solution, however, sacrifices considerable functionality present in the larger systems.

What is needed is a storage device that can be directly accessed by multiple other devices, without the need to go through an operating system.

Summary of the Invention

In the present invention a storage device has partitions that are separately addressed by distinct IP addresses. This allows direct access of the partitions, on a peer-to-peer basis, by any other device that can communicate using IP. Many limitations on access to the storage device can thereby be eliminated, including geographical limitations, and the need for a given storage partition to be under the central control of a single operating system.

Preferred storage devices support spanning between or among partitions of the same device, as well as between or among different storage devices. Both multicast and proxy spanning are contemplated.

Combinations of the inventive storage devices with each other, and with prior art storage devices are contemplated, in all manner of mirroring and other arrangements.

In still other aspects of the invention, a given storage device can comprise one or more types of media, including any combination of rotating and non-rotating media, magnetic and optical, and so forth.

Various objects, features, aspects and advantages of the inventive subject matter will become more apparent from the following detailed description of preferred embodiments, along with the accompanying drawing figures.

Brief Description of The Drawing

Fig. 1 is a schematic of a prior art disk drive split into multiple partitions, but where the entire memory is accessed using a single IP address.

Fig. 2 is a schematic of a prior art storage system in which three disk drives are addressed in their entireties using three different IP addresses.

Fig. 3 is a schematic of a storage device having multiple partitions that are separately addressed by different IP addresses.

Fig. 4 is a schematic of a storage device having multiple partitions that are separately addressed by different IP addresses, and some of the partitions are addressed using multiple IP addresses.

Fig. 5 is a schematic of a storage device having multiple partitions comprising different storage media.

Fig. 6 is a schematic of a storage device having multiple partitions, two of which are spanned using multicast spanning.

Fig. 7 is a schematic of a storage device having multiple partitions, two of which are spanned using proxy spanning.

Fig. 8 is a schematic of a storage system in which three storage devices are logically coupled using multicast spanning.

Fig. 9 is a schematic of a storage system in which three storage devices are logically coupled using proxy spanning.

Fig. 10 is a schematic of a storage system in which partitions of a first storage device are mirrored on partitions of one or more additional storage device using multicast mirroring.

Detailed Description

Prior art **Figure 1** generally depicts a disk drive 10 that is split into multiple partitions 10_A, 10_B, 10_C ... 10_N. The entire storage area is addressed using a single address IP_1, with individual blocks of data being addressed by a combination of IP_1 and some other information such as partition and offset, or Logical Block Address (LBA). The data is thus always accessed under the control of a disk operating system that provides the additional information. For that reason drive 10 is usually located very close to the processor that runs the operating system, and is usually connected to a hard bus of a computer, RAID or other system.

It is known to format the various partitions 10A ... 10N differently from one another, under control of different operating systems. However, the entire memory space comprises a single media type, namely rotating magnetic memory, even though there may be some sort of RAM buffer (not shown).

It should be appreciated that the term "IP" is used herein in a broad sense, to include any networking protocol. Thus, an IP address is used as a euphemism for a network address.

Prior art **Figure** 2 generally depicts a storage system 20 in which three disk drives 21, 22, 23 are addressed using three different IP addresses, IP1, IP2, and IP3. The drives can have multiple partitions (drive 21 has three partitions 21_A, 21_B, 21_C (not shown), and drive 23 has two partitions 23_A and 23_B (not shown)), but here again individual blocks of data are addressed using a combination of the IP address, some other information such as partition and offset, or LBA. Drives 21, 22, 23 can be spanned and/or mirrored, but the data on each drive is always accessed using that drive's particular IP address.

In **Figure** 3 is a storage device 30 according to the present invention has three partitions 21_A, 21_B, 21_C, which are separately addressed by different IP addresses IP_1, IP_2, IP_3, respectively. Those skilled in the art will appreciate that showing a small plurality of partitions is merely a matter of convenience, in this and other figures, and that storage device 30 could have any practical number of partitions. Similarly, it should be appreciated that depicting storage devices without partitions indicates that such devices have no partitions.

Utilizing IP addresses to route packets directly to and from partitions facilitates the use of very light communication protocols. In particular, the partitions may be directly addressed at the IP level of TCP/IP or UDP/IP stack. It should be appreciated, however, that in order make use of the IP addresses, the storage device 30 (and indeed the various partitions) would need to have sufficient functionality to communicate using IP. That functionality could be designed into the devices (or partitions), or it could be added onto storage devices using an IP adapter 32 (not shown). Indeed, the adapter in such circumstances would essentially be a simple block-to-packet and packet-to-block translator.

Storage device 30 can be connected to any suitable bus by any suitable means. Thus, the operative principles herein can operate across a wide variety of physical buses and protocols, including ATA, ATAPI, SCSI, Fiber CH, PCMCIA, CardBus, and USB. Storage device 30 can also alternatively or additionally operate across a network acting as a virtual IP bus, with the term "IP" being used herein generically with reference to any internetworking protocol that handles packets. It is contemplated, for example, that a user may have a stand-alone storage device that communicates wirelessly with a Local Area Network (LAN), which in turn may be connected to a WAN or to the Internet. Other devices that are also connected to the network (whether in the home, office, or elsewhere) could directly access one or more partitions of the storage device. For example, an IP capable television (not shown) could display images or movies stored on one partition, while a digital camera (not shown) could store/retrieve images on another partition. Still another partition might hold an operating system and office software for use with a laptop, or even an IP capable display and IP capable keyboard and mouse. Printing from any of the partitions might occur on an IP capable printer that is also connected wirelessly, or by hardwire, to the network.

An interesting corollary is that the partitions or other elements can all communicate as peers on a peer-to-peer network. As used herein, the term "element" refers to a hardware unit that is a functional portion of a device, and traditionally communicates with other units of the same device across a bus, without having its own IP address. This can completely eliminate dependence on any particular operating system, and can eliminate operating systems altogether. In addition, many of the elements attached to the network will be dependent on other elements attached to the network to perform tasks that are not within their individual capacities, and will be able to discover, reserve, and release the resources of other peers needed to perform such tasks. Peers will preferably be able to discover the other elements attached to the network, the characteristics of the other elements attached to the network, and possibly the contents of at least some of the elements attached to the network. Such discovery is accomplished without the assistance of a master device, and will preferably involve direct communication between the peer elements.

Preferred networks will be masterless in that all elements have equal access to the network and the other elements attached to the network. The peer elements of the network will preferably communicate with each other utilizing low-level protocols such as those that would equate to those of the transport and lower layers of the OSI model. Preferred embodiments will utilize TCP and UDP IP protocols for communication between elements.

Storage device 30 is preferably able to dynamically create partitions upon receipt of requests from network elements. For example, when a network element requests use of device 30, the network element may provide a unique identifier, possibly a name, to storage device 30, which in turn associates the identifier with any newly created partition. In some instances the network element may also request a particular storage size to be allocated, including all of the remaining storage available on the storage device 30.

In preferred embodiments, the IP addresses for such partitions are obtained from an address server such as a DHCP server upon request from the storage device 30. It is important to note, however, that address allocation devices such as DHCP servers are not masters, since they don't control the network, elements coupled to the network, or the sharing of resources between elements. Assignment of IP addresses to partitions may additionally or alternatively occur during initialization of the device, such as when it is first turned on.

Since storage device 30 may be associated with only a single network interface card (NIC), it is preferred that storage elements be able to obtain multiple IP addresses despite having a single NIC and a single media access control (MAC) address. This can be accomplished by providing a unique partition identifier to an address server when trying to obtain a IP address from the address server. It is contemplated that associating a name provided by an element with any partition created for that element makes it possible to identify each of the partitions of a storage element, despite the fact that IP address associated with each partition may have changed since the partition was created.

Additional details can be found in concurrently filed PCT application no. ________, entitled "Communication Protocols, Systems and Methods" and PCT application no. ________, entitled "Electrical Devices With Improved Communication", the disclosures of which are incorporated herein by reference.

In **Figure 4**, storage device 40 is similar to storage device 30 in that it has multiple partitions 41_A, 41_B, 41_C, 41_D that are separately addressed by different IP addresses IP_1, IP_2, IP_3, IP_4, respectively. But here some of the partitions are addressed using multiple IP addresses. In particular, partition 41_A is addressed with IP_1 and IP_5. Partition 41_D is addressed with IP_4, IP_6 and IP_7.

In Figure 5 a storage device 50 has multiple partitions comprising different storage media. In this particular example there are 2 partitions of rotating media 50_A, 50_B, one partition of flash memory 50_C. All other practical combinations of these and other media are also contemplated. As in Figure 3, the various partitions are separately addressed by different IP addresses IP_1, IP_2, IP_3, respectively.

In **Figure** 6 a storage device 60 has multiple partitions 60_A, 60_B, 60_C, 60_D, addressed by IP addresses IP_1, IP_2, IP_3, IP_4, and IP_5 (multicast) respectively. Two of these partitions, 60_A and 60_C, are spanned in that partition 60_A extends from logical address a to logical address b, while partition 60_C continues from logical address b+1 to logical address c. The spanned set is thus logical address a to logical address c. The spanning here is multicast spanning, because the partitions share multicast IP5 which is used to address both partitions 60_A and 60_C.

In **Figure** 7 a storage device 70 has multiple partitions 70_A, 70_B, 70_C, 70_D, addressed by IP addresses IP_1, IP_2, IP_3, IP_8, respectively. (The use of IP_8 here rather than IP_4 is intended to illustrate that the IP addresses need not be consecutive in any manner.) Here again two of the partitions are spanned, 70_A and 70_C, in that partition 70_A extends from logical address a to logical address b, while partition 70_C continues from logical address b+1 to logical address c. The spanned set is thus once again logical address a to logical address c. Here, however, we are dealing with proxy spanning as opposed to multicast spanning. IP_1 is used to address partition 70_A, while the second part of the spanned data, in partition 70_C, is addressed by the IP1 proxy using IP_3. Of course, it is possible to combine multicast spanning and proxy spanning within the same storage device .

In Figure 8 a storage system 100 has three storage devices 110, 120, and 130 coupled to depict multicast spanning. Device 110 has three partitions 110_A, 110_B and 110_C, which are separately addressed using IP addresses IP_1, IP_2, and IP_3, respectively. Device 120 has four partitions 120_A, 120_B, 120_C, and 120_D, which are separately addressed using IP addresses IP_4, IP_5, IP_6, and IP_7, respectively. Device 130 is not partitioned, which for our purposes is the same as saying that it only has one partition. The entirely of the storage area of device 130 is addressed using IP address IP_8. The spanning in this case is among all three drives. Partition 110C extends from logical address a to logical address b; partition 120D continues from logical address b+1 to logical address c, and the data space of device 130 extends from logical address c+1 to logical address d. The data set extends from logical address a to logical address d.

Figure 9 is similar to Figure 8, in that spanning occurs across three drives, and the data set extends from logical address a to logical address d. The main conceptual difference is that the storage devices are logically coupled using proxy spanning rather than multicast spanning. Here, device 210 has three partitions 210_A, 210_B and 210_C, which are separately addressed using IP addresses IP_1, IP_2, and IP_3, respectively. Device 230 is not partitioned. The entirely of the storage area of device 230 is addressed using IP address IP_4. Device 220 has three partitions, 220_A, 220_B and 220_C, which are separately addressed using IP addresses IP_4, IP_5, and IP_6, respectively. Partition 210_C extends from logical address a to logical address b; the data space of partition 220_C continues from logical address b+1 to logical address c, and partition 230 extends from logical address c+1 to logical address d.

As elsewhere in this specification, the specific embodiments shown with respect to Figure 9 are merely examples of possible configurations. A greater or lesser number of storage devices could be utilized, and indeed spanning may be protean, in that devices and/or partitions may be added to or dropped from the spanning over time. There can also be any combination of multicast and proxy spanning across and/or within storage devices, which may have the same or different media. Moreover, the use of IP addresses facilitates physically locating the various storage devices virtually anywhere an IP network can reach, regardless of the relative physical locations among the devices.

In **Figure 10** a storage system 300 provides mirroring of partitions between three different physical storage devices 310, 320 and 330. This could be done by proxy, in a manner analogous to that described above for proxy spanning, or in higher performance systems using multicasting. Thus, partitions in multiple storage devices are addressed using the same IP address. In this particular embodiment, storage device 310 has partitions 310_A, 310_B, and 310_c, addressed using IP addresses IP_1, IP_2, IP_3 and IP9. Storage device 320 has partitions 320_A, 320_B, and 320_C, addressed using IP addresses IP_4, IP_5, IP6 and IP_9. Write requests to IP_3 or IP_9 will result in partition 310_C 320_C and 330_C storing the same data. Read requests to IP_1 address will result in 310_C, 320_C and 330_C responding with the same information, with presumably the requester using whichever data arrives first. In the Multicast form it may be preferred that device 310,320 and 330 listen for the first data returned by any member of the mirrored set, and then remove that request from their request que if another device completes the request before they complete the request.

Communications

In preferred embodiments, communications between a storage element and a non-storage element, will utilize a datagram protocol in which data blocks are atomically mapped to a target device. A datagram sent between elements will preferably comprise command (CMD), logical block address (LBA), data, and token fields, and no more than X additional bytes where X is one of 1, 2, 7, 10, 17, and 30. The data field of such a datagram is preferably sized to be the same as the block size (if applicable) of the element to which the datagram is addressed. As such, an element sending a quantity of data to a storage element where the quantity of data is larger than the block size of the storage element will typically divide the quantity of data into blocks having the same size as the blocks of the storage element, assign LBAs to the blocks, and send each block and LBA pair to the storage element in a datagram.

It is preferred that the datagrams be communicated between elements encapsulating them within addressed packets such as IP packets, and the IP address of the encapsulating packet be used to identify both the element a packet is intended to be sent to, and the partition within the element that the datagram pertains to.

It is preferred that datagram recipients handle datagrams on a first come, first served basis, without reordering packets, and without assembling the contents of the data fields of datagrams into a larger unit of data prior to executing a command identified in the CMD field. As an example, an storage element may receive a datagram containing a block of data, an LBA, and a write command. The storage element, without having to wait for any additional packets, utilizes the IP address of the packet enclosing the datagram to identify the partition to be used, and utilizes the LBA to identify the location within the partition at which the data in the data field is to be written.

Handling the data in individual datagrams as they arrive rather than reassembling the data permits the use of an implied ACK for each command. Using an implied rather than an explicit ACK results in a substantial increase in performance.

Marketing of Storage Devices and Adapters

It is contemplated that once persons in the industry recognize the benefits of having storage devices having partitions that are accessed using their own IP addresses, companies will start producing and/or marketing such devices It is also contemplated that companies will start producing and/or marketing adapters that includes a functionality (hardware or software, or come combination of the two) to permit traditional disk drives, flash memories, and other storage devices to operate in that manner.

Thus, methods falling within the inventive subject matter include manufacturing or selling a disk drive or other storage device in which the partitions can utilize their own IP addresses to execute packet communication with other network elements. Other inventive methods include manufacturing or selling adapters that enable prior art type storage devices to do the same. Indeed it is contemplated that companies will recognize that such adapters are available, and will continue to manufacture or sell prior art type storage devices, knowing (or even advertising) that users can employ such adapters to enable the prior art type storage devices to use in an infringing manner.

Thus, specific embodiments and applications of the inventive storage devices have been disclosed. It should be apparent, however, to those skilled in the art that many more modifications besides those already described are possible without departing from the inventive concepts herein. The inventive subject matter, therefore, is not to be restricted except in the spirit of the appended claims. Moreover, in interpreting both the specification and the claims, all terms should be interpreted in the broadest possible manner consistent with the context. In particular, the terms "comprises" and "comprising" should be interpreted as referring to elements, components, or steps in a non-exclusive manner, indicating that the referenced elements, components, or steps may be present, or utilized, or combined with other elements, components, or steps that are not expressly referenced.

CLAIMS

What is claimed is:

1. A storage device having a first type of media logically split into a first partition directly addressed by a first IP address and a second partition directly addressed by a second IP address that is distinct from the first IP address.

2. The storage device of claim 1, wherein the first partition is also directly addressed by a third IP address distinct from both the first and second IP addresses.

3. The storage device of claim 1, further comprising a third partition that is directly addressed by a third IP address.

4. The storage device of claim 3, wherein the third partition comprises a second type of media that is different from the first type of media.

7. The storage device of claim 3, wherein the third IP address is distinct from the first and second IP addresses.

6. The storage device of claim 3, wherein the third IP address is the same as the first IP address.

7. The storage device of claim 6, wherein the first IP address directly addresses both the first and third partitions using multicast spanning.

8. The storage device of claim 6, wherein the first IP address directly addresses both the first and third partitions using proxy spanning.

9. The storage device of claim 1, wherein the first type of media comprises a rotating media.

10. The storage device of claim 1, wherein the first type of media comprises a non-rotating memory.

11. The storage device of claim 1, wherein the first type of media comprises an optical memory.

12. The storage device of claim 1, wherein the first partition is formatted differently from the second partition.

13. The storage device of claim 1, further comprising a connector that connects to a hardware bus.

14. The storage device of claim 1, further comprising a wireless IP connection.

17. A storage system comprising the storage device of claim 1, and another storage device that is directly addressed by the first IP address using multicast spanning.

16. A storage system comprising the storage device of claim 1, and another storage device that is directly addressed by the first IP address using proxy spanning.

17. A storage system comprising the storage device of claim 1, and another storage device having a partition that is directly addressed by the first IP address.

18. A storage system comprising the storage device of claim 1, and another storage device, substantially the entire storage area of which is directly addressed by the first IP address.

19. A storage system comprising the storage device of claim 1, and a second storage device having additional partitions that mirror data stored in the first and second partitions.

20. A storage system comprising the storage device of claim 19, wherein the additional partitions of the second storage device are directly addressed by IP addresses that are distinct from the first and second IP addresses.

21. A storage system comprising the storage device of claim 20, wherein the storage devices participating in a mirrored set discontinue processing requests successfully completed by another member of the mirrored set.

AMENDED CLAIMS
[received by the International Bureau on 02 July 2003 (02.07.2003); original claims 1-21 replaced by new claims 1-30 (4 pages).

1. A storage device having a first type of media logically split into a first partition directly addressed by a first IP address and a second partition directly addressed by a second IP address that is distinct from the first IP address, wherein:

   the storage device is adapted to receive and process datagrams comprising a command field, a logical block address field, a data block field, a token field, and no more than 30 additional bytes encapsulated within the data portion of an IP packet, and the data block field of such datagrams is sized to correspond to a set block size of the storage device;

   the storage device is adapted to associate a first name provided to it by an external device with the first partition and to associate a second name provided to it by an external device with the second partition, wherein the names are partition identifiers that are not network addresses;

   the device is adapted to obtain the first and second IP addresses from an address server upon initialization of the storage device, and is adapted to allow the obtained IP addresses to change whenever the storage device is initialized while maintaining the associations between provided names and partitions;

   the storage device is adapted to obtain an IP address for a partition at least in part by providing the name associated with the partition to the address server; and

   the storage device adapted to support multicast spanning and mirroring at least in part by being adapted to allow an external device to use a single multicast IP to access at least two partitions of the storage device.

2. A storage device having a first type of media logically split into a first partition directly addressed by a first IP address and a second partition directly addressed by a second IP address that is distinct from the first IP address, wherein the storage device is adapted to utilize a datagram protocol in which data blocks are atomically mapped to a target device.

3. A storage device having a first type of media logically split into a first partition directly addressed by a first IP address and a second partition directly addressed by a second IP address that is distinct from the first IP address, wherein the storage device

is adapted to associate a first name provided to it by an external device with the first partition and to associate a second name provided to it by an external device with the second partition, wherein the names are partition identifiers that are not IP addresses.

4. A storage device adapted to support multicast spanning at least in part by being adapted to allow an external device to use a single multicast IP address to access at least two partitions of the storage device.

5. A storage device adapted to support mirroring at least in part by being adapted to allow an external device to use a single multicast IP address to access at least two partitions of the storage device, and by being adapted to cause the contents of one of the at least two partitions to mirror the contents of another of the at least two partitions.

6. A storage system adapted to support multicast spanning at least in part by being adapted to allow an external device to use a single multicast IP address to access at least two partitions located on separate storage devices.

7. A storage system adapted to support mirroring at least in part by being adapted to allow an external device to use a single multicast IP address to access at least two partitions located on separate storage devices, and by being adapted to cause the contents of one of the at least two partitions to mirror the contents of another of the at least two partitions.

8. A storage system comprising a set of spanned or mirrored partitions wherein at least two of the partitions are located on devices separated by a proxy server.

9. The device of claim 2 wherein the device is adapted to receive and process datagrams comprising a command filed, logical block address field, a data block field, a token field, and no more than 30 additional bytes encapsulated within the data portion of an IP packet.

10. The device of claim 9 wherein the storage device has a set block size and the data block field is equal to the storage device set block size.

11. The storage device of claim 3 wherein the device is adapted to obtain the first and second IP addresses from an address server upon initialization of the storage device,

and is adapted to allow the obtained IP addresses to change whenever the storage device is initialized while maintaining the associations between provided names and partitions.

12. The storage device of claim 11 wherein the storage device is adapted to obtain an IP address for a partition at least in part by providing the name associated with the partition to the address server.

13. The storage device of claim 1, 2, or 3, wherein the first partition is also directly addressed by a third IP address distinct from both the first and second IP addresses.

14. The storage device of claim 1, 2, or 3, further comprising a third partition that is directly addressed by a third IP address.

15. The storage device of claim 14, wherein the third partition comprises a second type of media that is different from the first type of media.

16. The storage device of claim 14, wherein the third IP address is distinct from the first and second IP addresses.

17. The storage device of claim 14, wherein the third IP address is the same as the first IP address.

18. The storage device of claim 1, 2, or 3, wherein the first type of media comprises a rotating media.

19. The storage device of claim 1, 2, or 3, wherein the first type of media comprises a non-rotating memory.

20. The storage device of claim 1, 2, or 3, wherein the first type of media comprises an optical memory.

21. The storage device of claim 1, 2, or 3, wherein the first partition is formatted differently from the second partition.

22. The storage device of claim 1, 2, or 3, further comprising a connector that connects to a hardware bus.

23. The storage device of claim 1, 2, or 3, further comprising a wireless IP connection.

24. The storage device of claim 1, 2, or 3 wherein the storage device adapted to support multicast spanning and or mirroring at least in part by being adapted to allow an external device to use a single multicast IP to access at least two partitions of the storage device.

25. A storage system comprising the storage device of claim 1, 2, or 3, and another storage device that is directly addressed by the first IP address.

26. A storage system comprising the storage device of claim 1, 2, or 3, and another storage device having a partition that is directly addressed by the first IP address.

27. A storage system comprising the storage device of claim 1, 2, or 3, and another storage device, substantially the entire storage area of which is directly addressed by the first IP address.

28. A storage system comprising the storage device of claim 1, 2, or 3, and a second storage device having additional partitions that mirror data stored in the first and second partitions.

29. A storage system comprising the storage device of claim 28, wherein the additional partitions of the second storage device are directly addressed by IP addresses that are distinct from the first and second IP addresses.

30. A storage system comprising the storage device of claim 29, wherein the storage devices participating in a mirrored set discontinue processing requests successfully completed by another member of the mirrored set.

Fig 1
(prior art)

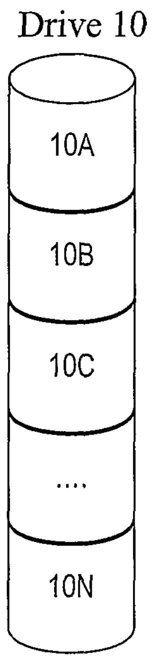

Fig 2
(Prior Art)

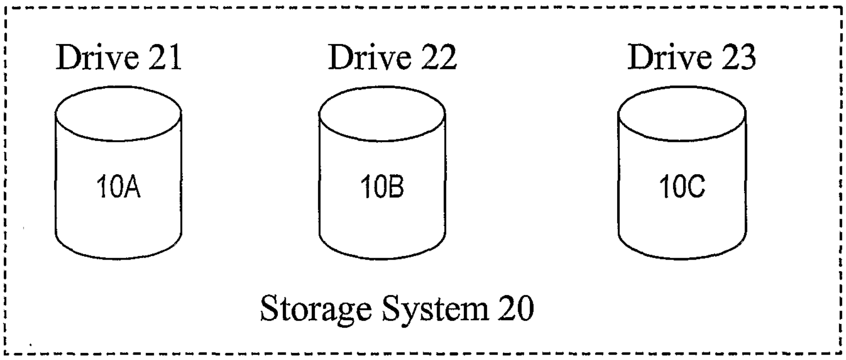

Fig 3

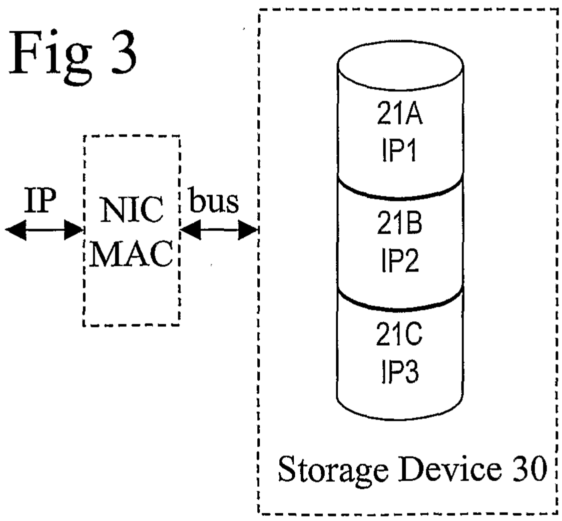

Fig 4

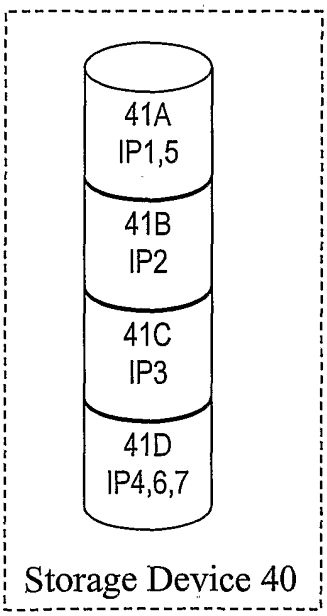

Fig 5

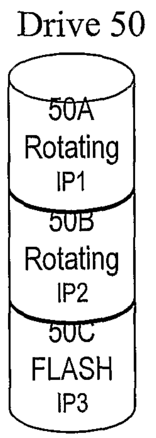

Fig 6

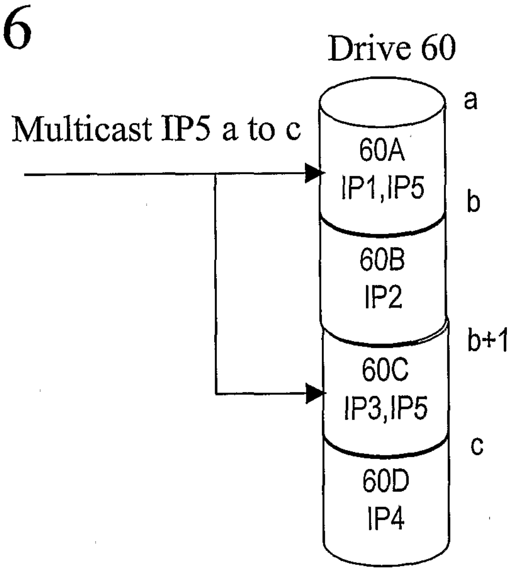

Fig 7

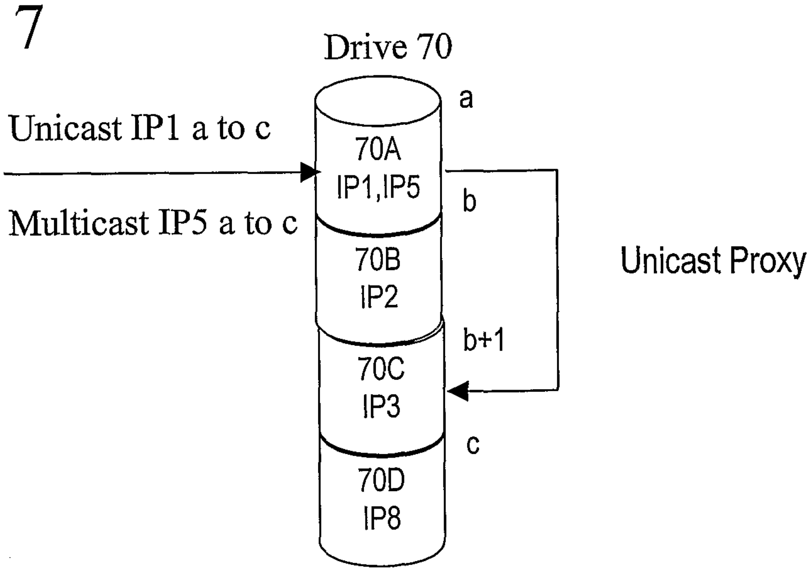

Fig 8

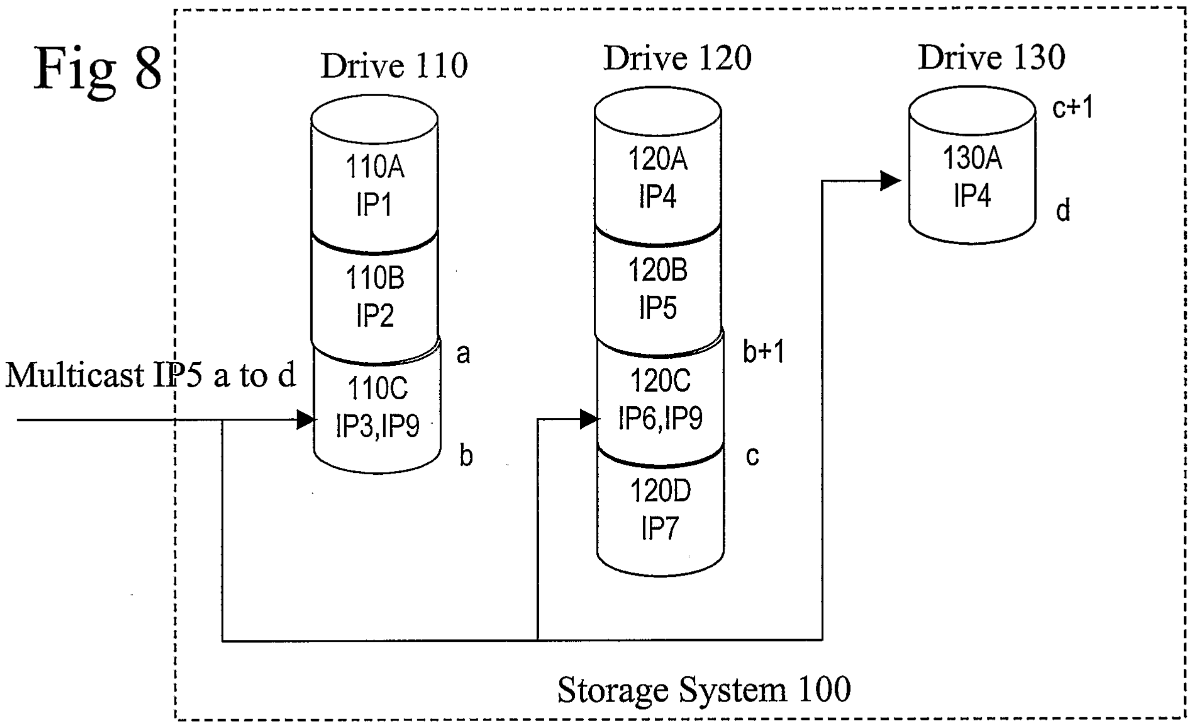
Drive 110
Drive 120
Drive 130
110A
IP1
110B
IP2
110C
IP3,IP9
120A
IP4
120B
IP5
120C
IP6,IP9
120D
IP7
130A
IP4
c+1
d
a
b
b+1
c
Multicast IP5 a to d
Storage System 100

Fig 9

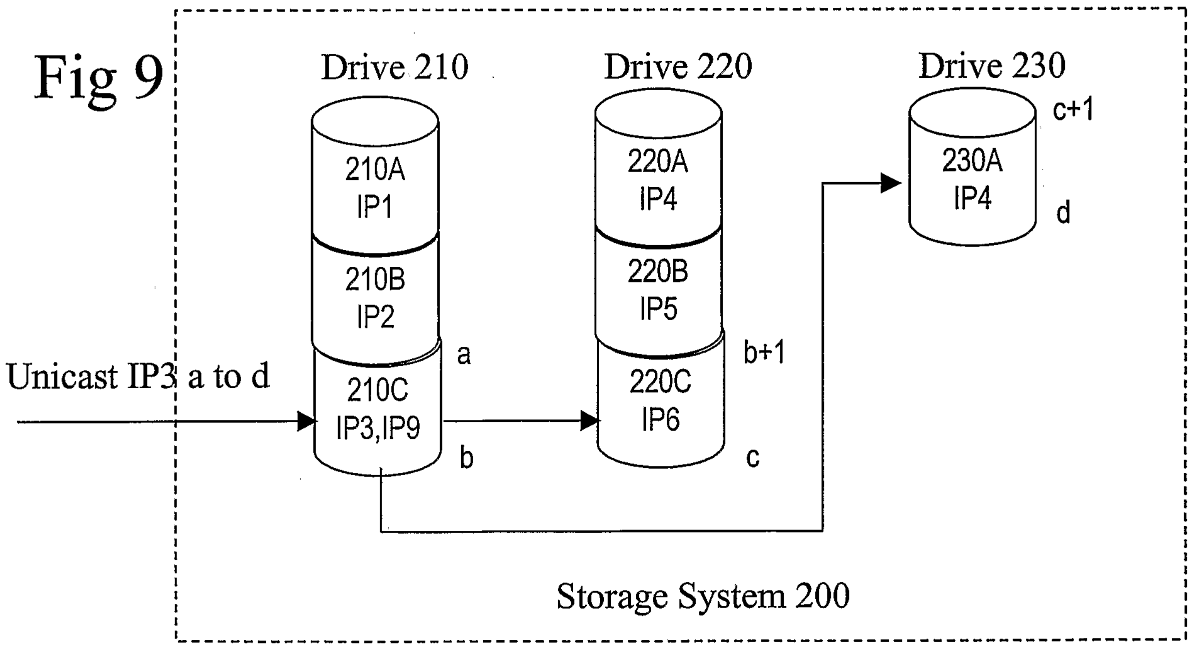
Drive 210
Drive 220
Drive 230
210A
IP1
210B
IP2
210C
IP3,IP9
220A
IP4
220B
IP5
220C
IP6
230A
IP4
c+1
d
a
b
b+1
c
Unicast IP3 a to d
Storage System 200

Fig 10

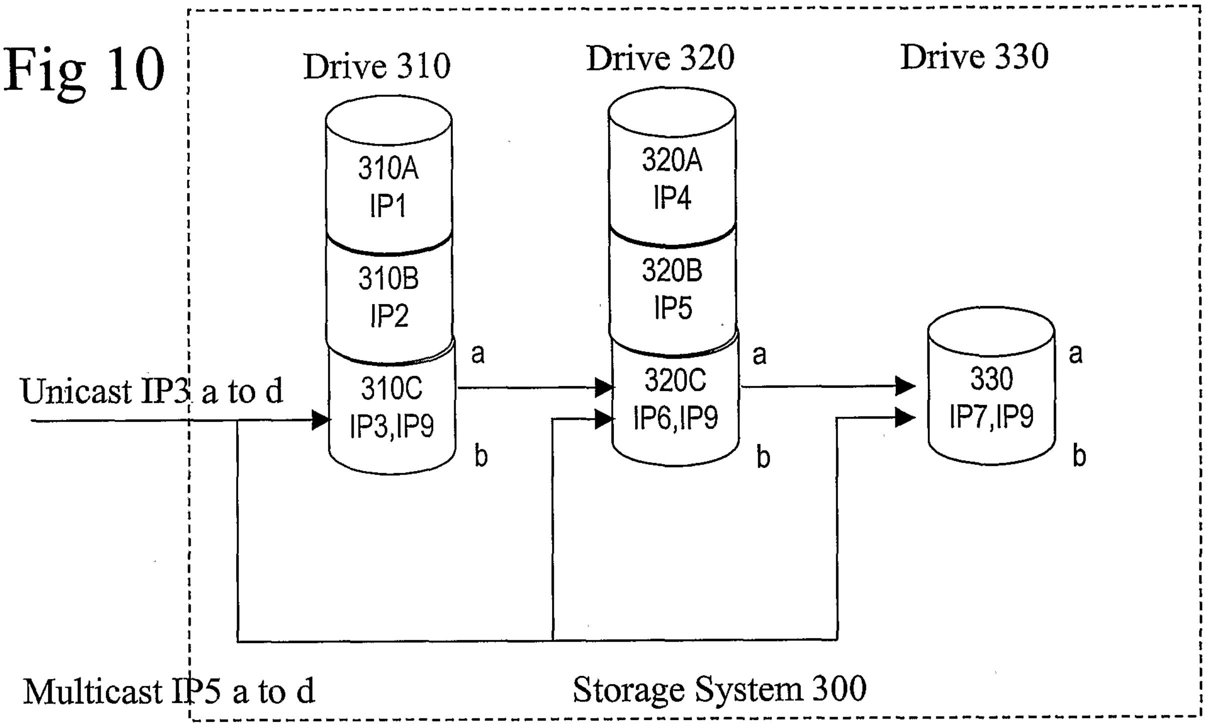
Drive 310
Drive 320
Drive 330
310A
IP1
310B
IP2
310C
IP3,IP9
320A
IP4
320B
IP5
320C
IP6,IP9
330
IP7,IP9
a
b
a
b
a
b
Unicast IP3 a to d
Multicast IP5 a to d
Storage System 300

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US02/40199

A. CLASSIFICATION OF SUBJECT MATTER
IPC(7) : G06F 12/06
US CL : 709/245; 711/153, 173
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
U.S. : 709/245; 711/153, 173

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
WEST: USPT

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,634,111 A (OEDA et al.) 27 May 1997 (27.05.1997), Fig. 6, column 10, lines 5-57	1-21

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search 21 April 2003 (21.04.2003)	Date of mailing of the international search report 08 MAY 2003
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230	Authorized officer Dave Wiley Peggy Harrod Telephone No. 703 305-9700