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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
8 January 2004 (08.01.2004)

PCT

(10) International Publication Number
WO 2004/003783 A1

(51) International Patent Classification[7]: G06F 17/00, G09G 5/00

(21) International Application Number: PCT/US2003/017608

(22) International Filing Date: 4 June 2003 (04.06.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/392,779 1 July 2002 (01.07.2002) US

(71) Applicant *(for all designated States except US)*: **THOMSON LICENSING S.A.** [FR/FR]; 46 Quai A. Le Gallo, F-92648 Boulogne (FR).

(71) Applicants and
(72) Inventors: **CROWTHER, David, Aaron** [US/US]; 4817 SE Stewart Court, Hillsboro, CA 97123 (US). **REDDY, Ravi** [IN/US]; 55 S. W. Hortop Way, Beaverton, OR 97006 (US). **ADKINS, Andy, Eugene** [US/US]; 10297 SW Eastridge, Portland, OR 97225 (US). **CAO, Nanyu** [US/US]; 13446 NW Yvonne Lane, Portland, OR 97229 (US).

(74) Agents: **TRIPOLI, Joseph, S.** et al.; Thomson Licensing, Inc., 2 Independence Way, Suite 2, Princeton, NJ 08543 (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, DE, DK, DM, DZ, EC, EE, ES, 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, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, 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), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, 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*

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:** HETEROGENEOUS DISK STORAGE MANAGEMENT TECHNIQUE

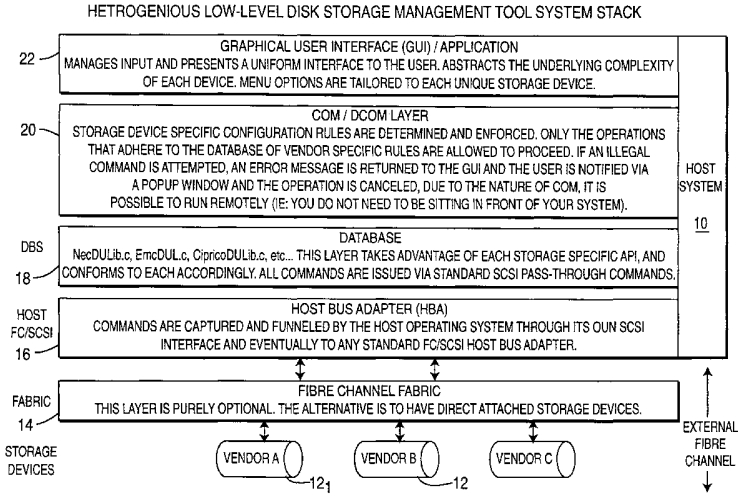

(57) **Abstract:** Management of one or more storage devices (12_1-12) is achieved by first identifying each device by its make and model, and then establishing a database (18) containing information about that device, such as (a) operational rules, (b) commands and (c) processing routines. From the information obtained from the database (18), a Graphical User Interface is created for display to afford a user at least one menu option for selection. Upon selecting the one menu option, the user obtains at least one of (a) a display of information associated with the identified storage device, and (b) execution of at least one process to control, at least in part, the operation of the identified storage device. The user's selection is processed and the Graphical User Interface is automatically updated in response to the processing of the user's selected menu option.

WO 2004/003783 A1

HETEROGENEOUS DISK STORAGE MANAGEMENT TECHNIQUE

CROSS REFERENCE TO RELATED APPLICATION

This application claims priority under 35 U.S.C. 119(e) to U.S. Provisional Patent Application Serial No. 60/392,779, filed July 1, 2002, the teachings of which are incorporated herein.

TECHNICAL FIELD

This invention relates to a technique for managing one or more storage devices.

BACKGROUND ART

Advances in the design and manufacture of magnetic storage devices have significantly reduced the cost of such devices, as measured in terms of the cost per gigabyte of storage capacity. The relatively low cost of high storage capacity devices has resulted in the proliferation of large arrays of such devices, such as a Redundant Array of Independent Disks (RAID). Various manufacturers of storage devices commonly offer such RAIDs with proprietary software for controlling the individual storage devices within the RAID.

Storage devices from different manufactures often have different features. Combining storage devices from different manufacturer affords the ability to take advantage of the best features of each device. Unfortunately, no technique presently exists for low-level management of heterogeneous storage devices, such as a set of heterogeneous storage devices within a RAID. While software presently exists for managing RAIDs from different vendors across a network, such software lacks the ability to manage one or more storage devices from separate vendors in a single RAID, or to effectively manage a stand-alone storage device.

Thus, there is a need for a technique for managing one or more storage devices.

BRIEF SUMMARY OF THE INVENTION

Briefly, in accordance with present principles, there is provided a method for managing at least one storage device, such as a magnetic disk drive, magneto-optical device, or a tape storage device, and preferably, a plurality of such storage devices. The method commences by identifying each device, typically by its manufacturer. For each identified storage device, a database is created containing information about that device. For example, the information within such a database will include the working characteristics of that storage device, specific interface protocols and operational rules. From the information about each storage device obtained from the corresponding database, a Graphical User Interface (GUI) is created for display to provide a user with at least one menu option for selection. Upon selecting the one menu option, the user obtains at least one of (a) a display of information associated with the identified storage device, and (b) execution of at least one process to control, at least in part, the operation of the identified storage device. The user's selection is processed and the Graphical User Interface is automatically updated in response to the processing of the user's selected menu option. Thus, for example, if the user had selected a menu option to cause the identified storage device to perform a particular operation, the Graphical User Interface provides an updated display that reflects the results of that operation performed by the storage device.

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 depicts a block schematic diagram of a system in accordance with present principles for managing at least one storage device;

FIGURE 2 depicts in flow chart form the steps of a method executed by the system of FIG. 1 to manage the at least one storage device; and

FIGURE 3 depicts an opening screen of a Graphical User Interface displayed to a user.

DETAILED DESCRIPTION

FIGURE 1 depicts a block schematic diagram of a host system 10 for managing at least one, and preferably a plurality of storage devices, exemplified by storage devices 12_1, 12_2 and 12_3. Each of the storages devices 12_1, 12_2 and 12_3 can take the form of a single magnetic or optical disk drive, a magnetic tape drive, or an array of individual storage devices, such as a

Redundant Array of Independent Disks (RAID). In practice, the storage devices 12_1-12_3 originate from different manufacturers. Thus, each storage device often possesses slightly different operating characteristics that must be taken into account when managing that device.

A fibre channel fabric 14 couples the storage devices 12_1-12_3 to a host bus adapter 16 that interfaces the devices to the host system 10. The host bus adapter 16 typically has its own Small Computer System Interface (SCSI) to funnel information between the host system 10 and the storage devices 12_1-12_3 through the fibre channel fabric 14. The host bus adapter 16 can include other interfaces for disk storage devices that do not use SCSI. Note that the storage devices 12_1-12_3 could each enjoy a direct connection with the host bus adapter 16, thereby obviating the need for the fibre channel fabric 14. Thus, the storage devices 12_1-12_3 are not limited to fibre channel storage devices but could comprise any type of storage device responsive to SCSI commands.

The host system 10 typically comprises a general-purpose computer such as a well know personal computer or a mini computer. Associated with the host system 10 is at least one database (DBS) 18 that contains one or more libraries. Each library within the database 18 contains information specific to a corresponding one of the storage devices 12_1-12_3. In particular, the library associated with each of the storage devices 12_1-12_3 typically includes the operating characteristics associated with that storage device. The device operating characteristics includes at least one of (a) operational rules, (b) commands, and (c) processing routines. While the embodiment of FIG. 1 depicts a single database 18 that contains the associated library for each of the storage devices 12_1-12_3, the library associated with each storage device could reside on a separate database (not shown). In practice, each library in the database 18 depends only on SCSI commands. Thus, the host bus adapter 16 need not be limited to a Fibre Channel (FC)/SCSI host bus adapter because regardless of the lower level protocol (e.g., fibre channel fabric in this case), the commands in the library don't need to rely on the specific host bus adapter protocol.

The host system 10 includes a Communication Object Model (COM)/ Distributed Communications Object Model (DCOM) layer 20, which, in practice, takes the form of a memory that includes specific configuration rules for each of the storage devices 12_1-12_3. In practice, the configuration rules for each storage device are determined from the information stored in the corresponding library for each such device. The COM/DCOM layer 20 also includes one or more applets, programs and/or sub-routines (collectively referred as "programs") typically in the form of C++ objects that are exposed via multiple interfaces and can be instantiated via multiple languages across multiple machines in one or more networks. The

programs in the COM/DCOM layer 20 monitor each requested operation for each of the storage devices 12_1-12_3 to ensure that the requested operation complies with the rules for that storage device as established by its manufacturer. Any requested operation of a storage device not in compliance with the rule(s) associated with that device will be blocked and an error message will be generated. Execution by the host system 10 of the programs resident in the COM/DCOM layer 20 can occur without the need for the user to remain at a terminal (not shown).

The host system 10 also executes a Graphical User Interface (GUI) application 22 that manages input information received from each of the storage devices 12_1-12_3 as well as input information received from one or more users. From such input information, the GUI application 22 generates a graphical interface for display to the user on a display device (not shown). The overall format of the information displayed by the GUI 22 to the user for each storage drive has a uniform appearance, but different menu options are tailored for each storage device to account for its particular features.

FIGURE 2 illustrates in flow chart form the steps of a method executed by the host system 10 to manage the storage devices 12_1-12_3 in accordance with the present principles. The method of FIG. 2 commences upon execution of step 100 during which the host system 10 of FIG. 1 undertakes discovery of which storage devices are attached thereto. In other words, during step 100, the host system 10 launches a query to determine the identity of the attached devices 12_1-12_3, and particularly, the make (i.e., the manufacturer) and the model number of each device. In addition to discovering the make and model of each device, the host system 10 also ascertains during step 100 whether each device is directly attached to the host bus adapter 16 or is attached through the fibre channel fabric 14. During step 100, the host system 10 will assign an identifier to each discovered storage devices. The assigned identifier can comprise a fibre channel fabric identifier assigned when the device is attached to the fibre channel fabric 14 of FIG. 1. Alternatively, the identifier could comprise a Logical Unit Number (LUN) or vendor identification number assigned to the device.

Following step 100, execution of step 110 occurs at which time, the host system 10 categorizes each discovered storage device by its make and model. During step 110, the host system 10 establishes a library for each storage device discovered during step 100, with each library residing in a single database (e.g., database 18 of FIG. 1) or in separate databases. As discussed previously, each library associated with a storage device will include the operational rules, commands and processing routines associated with that device. Execution of step 120

follows step 110 at which time, the host system 10 executes the Graphical User Interface (GUI) application 22 of FIG. 1 to create and display the GUI to the user.

FIGURE 3 depicts an illustrative screen display 200 that comprises part of the GUI provided to the user. The screen display 200 includes system pane 202 typically appearing at the left-hand side of the screen display. The system pane 202 provides a hierarchical display of the various devices and sub-elements found by the host system 10 during step 100 of FIG. 2, with each device and sub-element appearing in the system pane 202 of FIG. 3 as a corresponding icon. In the exemplary embodiment of FIG. 3, the icon 206 appearing in the system pane 202 corresponds to the presence of a Vendor Unique Storage Device (not shown) found by the host system 10 during step 100 of FIG. 2. The icon 208 appearing in the system pane 202 of FIG. 3 corresponds to a controller node (not shown) within the Vendor Unique Storage Device represented by the icon 206. The icon 210 within the system pane 202 signifies the existence of one or more bound LUNs (not shown) associated with the controller node represented by icon 206. The icon 210 identifies a corresponding LUN within the Bound LUNs represented by the icon 210. Lastly, the icon 212 within the system pane 202 of FIG. 3 signifies the presence of one or more unbound disks (not shown). One or more of the icons 206-212 can have an associated drop-down menus to allow a user to obtain one of (a) a display of information associated with the identified storage device, and (b) execution of at least one process to control, at least in part, the operation of the identified storage device.

Still referring to FIG. 3, the screen display 200 includes a sub-system display pane 214 in addition to the display pane 202. In practice, the sub-system display pane 214 displays specific information related to the selection by a user (via a mouse or the like) of a particular one of the icons 206-212. In the exemplary embodiment, highlighting the controller node icon 208 causes the sub-system display pane 214 to display more specific information about the controller node.

Referring to FIG. 2, following creation and display of the GUI during step 120, the host system 10 then awaits input from a user who typically selects one or more menu options. Such menu options can include (but are not limited to) the following:

- bind a group of physical storage devices into a Logical Unit Number (LUN)
- unbind a set of LUNs
- designate one or more storage devices as hot spares
- set each storage device system clock
- load vendor specific firmware to individual controllers and/or storage devices
- check/set vendor specific controller and/or storage device settings

- blink the storage device Light Emitting Diodes (LEDs) if present for identification purposes
- identify state information
- establish fibre channel fabric IDs for each storage device
- start or stop a RAID rebuild operation
- start or stop one or more consistency check operations
- initialize a file system in a discrete storage system

After receipt of by a user of one or more menu selections, then the step 124 of FIG 2 occurs, at which time the host system 10 of FIG. 1 processes the received menu selection. In practice, the host system 10 processes the menu selection through an appropriate software stack until the specific storage device to which the menu selection was directed executes the selection. Following step 124, step 126 occurs, whereupon the host system 10 automatically updates the GUI previously displayed to the user to reflect changes that occurred as a result of execution of the user-selected menu option. In this way, the user can visually observe the results of the selected menu option. Following step 126, program execution branches back to step 122 to wait for user selection of another menu option.

The foregoing describes a technique for managing one or more storage devices having different operating characteristics.

CLAIMS

1. A method for managing at least one storage device, comprising the steps of:

(a) identifying the one storage device;

(b) establishing a database containing information about the identified device;

(c) providing to the user a graphical user interface in accordance with the information in the database, the graphical interface displaying at least one menu option for the identified device for selection by the user to request at least one of (i) display of information and (ii) execution of at least one process to control the operation of the one storage device;

(d) processing the selected menu option; and

(e) automatically updating the graphical user interface in response to the processing of the selected menu option.

2. The method according to claim 1 wherein the step of identifying the one storage device comprises the step of ascertaining a manufacturer and model number for the one storage device.

3. The method according to claim 1 wherein the step of establishing a database comprises the step of obtaining information that includes at least one of (a) operational rules, (b) commands and (c) processing routines of the one storage device.

4. The method according to claim 3 wherein the processing step includes

determining if the requested execution of the at least one process complies with the operational rules for the identified device, and if not,

blocking execution of the at least one process; and

generating an error message.

5. The method according to claim 1 wherein the step of providing a graphical user interface includes the steps of:

displaying icons representative of each identified devices in a hierarchical fashion in a first pane; and

displaying information in a second pane related a corresponding one of the icons displayed in a first pane.

6. The method according to claim 1 wherein steps (d) and (e) are repeated following receipt of each subsequent menu selection made by a user.

7. Apparatus for managing at least one storage device, comprising:

a database containing information about potential storage devices; and

a processor coupled to at least one storage device for (a) identifying the one storage device; (b) providing to a user a graphical user interface in accordance with the information in the database, the graphical interface displaying at least one menu option for the identified device for selection by the user to request at least one of (i) display of information and (ii) execution of at least one process to control the operation of the one storage device; (d) processing the selected menu option; and (e) automatically updating the graphical user interface in response to the processing of the selected menu option.

8. The apparatus according to claim 7 wherein the processor identifies the one storage device by ascertaining its manufacturer and model number.

9. The apparatus according to claim 7 database contains information that includes at least one of (a) operational rules, (b) commands and (c) processing routines of the one storage device.

10. The apparatus according to claim 7 wherein the processor process the selected menu option by (a) determining if the requested execution of the at least one process complies with the operational rules for the identified device, and if not, (b) blocking execution of the at least one process; and (c) generating an error message.

11. The apparatus according to claim 7 wherein the processor provides the graphical user interface by (a) displaying icons representative of each identified devices in a hierarchical fashion in a first pane; and (b) displaying information in a second pane related a corresponding one of the icons displayed in a first pane.

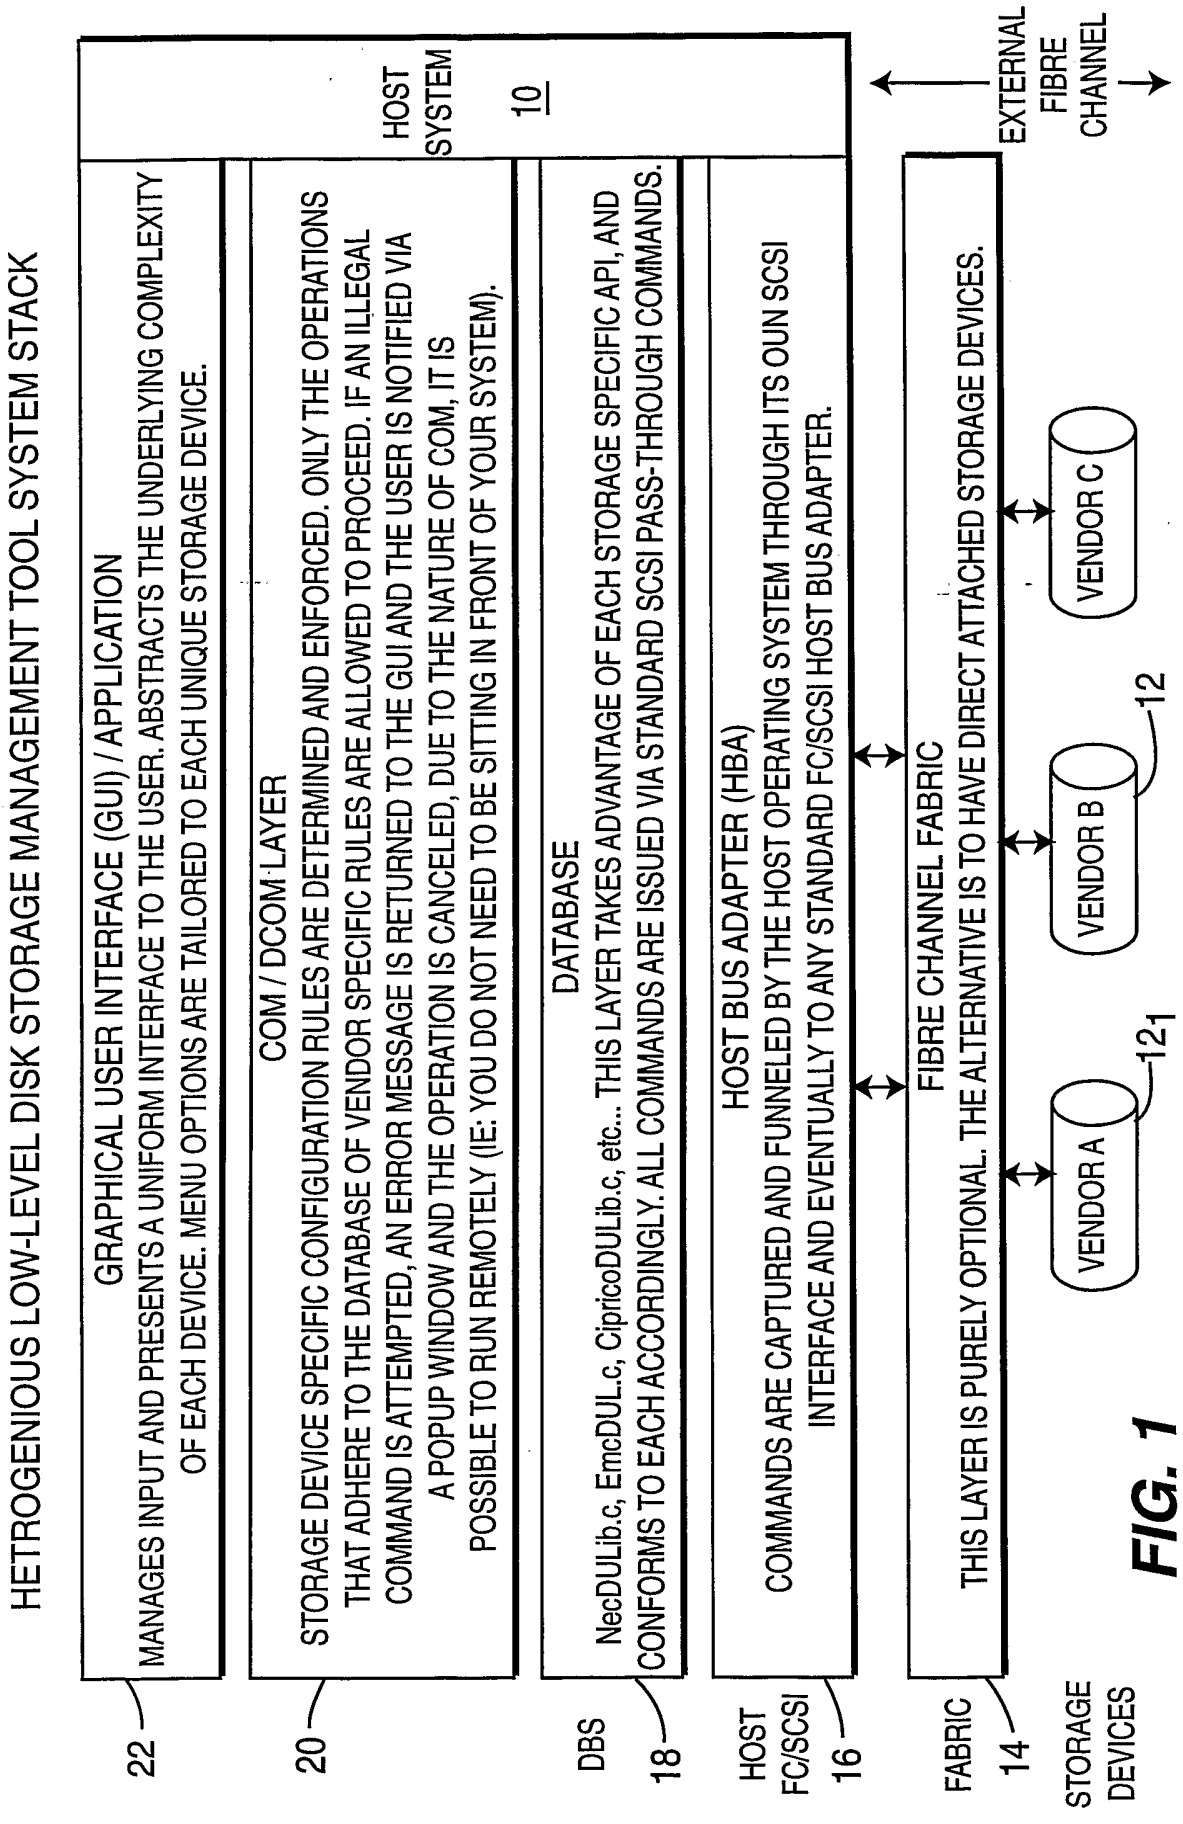
HETROGENIOUS LOW-LEVEL DISK STORAGE MANAGEMENT TOOL SYSTEM STACK
HOST SYSTEM 10
GRAPHICAL USER INTERFACE (GUI) / APPLICATION
MANAGES INPUT AND PRESENTS A UNIFORM INTERFACE TO THE USER. ABSTRACTS THE UNDERLYING COMPLEXITY OF EACH DEVICE. MENU OPTIONS ARE TAILORED TO EACH UNIQUE STORAGE DEVICE.
22
COM / DCOM LAYER
STORAGE DEVICE SPECIFIC CONFIGURATION RULES ARE DETERMINED AND ENFORCED. ONLY THE OPERATIONS THAT ADHERE TO THE DATABASE OF VENDOR SPECIFIC RULES ARE ALLOWED TO PROCEED. IF AN ILLEGAL COMMAND IS ATTEMPTED, AN ERROR MESSAGE IS RETURNED TO THE GUI AND THE USER IS NOTIFIED VIA A POPUP WINDOW AND THE OPERATION IS CANCELED, DUE TO THE NATURE OF COM, IT IS POSSIBLE TO RUN REMOTELY (IE: YOU DO NOT NEED TO BE SITTING IN FRONT OF YOUR SYSTEM).
20
DBS
DATABASE
NecDULib.c, EmcDUL.c, CipricoDULib.c, etc... THIS LAYER TAKES ADVANTAGE OF EACH STORAGE SPECIFIC API, AND CONFORMS TO EACH ACCORDINGLY. ALL COMMANDS ARE ISSUED VIA STANDARD SCSI PASS-THROUGH COMMANDS.
18
HOST FC/SCSI
HOST BUS ADAPTER (HBA)
COMMANDS ARE CAPTURED AND FUNNELED BY THE HOST OPERATING SYSTEM THROUGH ITS OUN SCSI INTERFACE AND EVENTUALLY TO ANY STANDARD FC/SCSI HOST BUS ADAPTER.
16
FABRIC
FIBRE CHANNEL FABRIC
THIS LAYER IS PURELY OPTIONAL. THE ALTERNATIVE IS TO HAVE DIRECT ATTACHED STORAGE DEVICES.
14
EXTERNAL FIBRE CHANNEL
STORAGE DEVICES
VENDOR A
121
VENDOR B
12
VENDOR C

FIG. 1

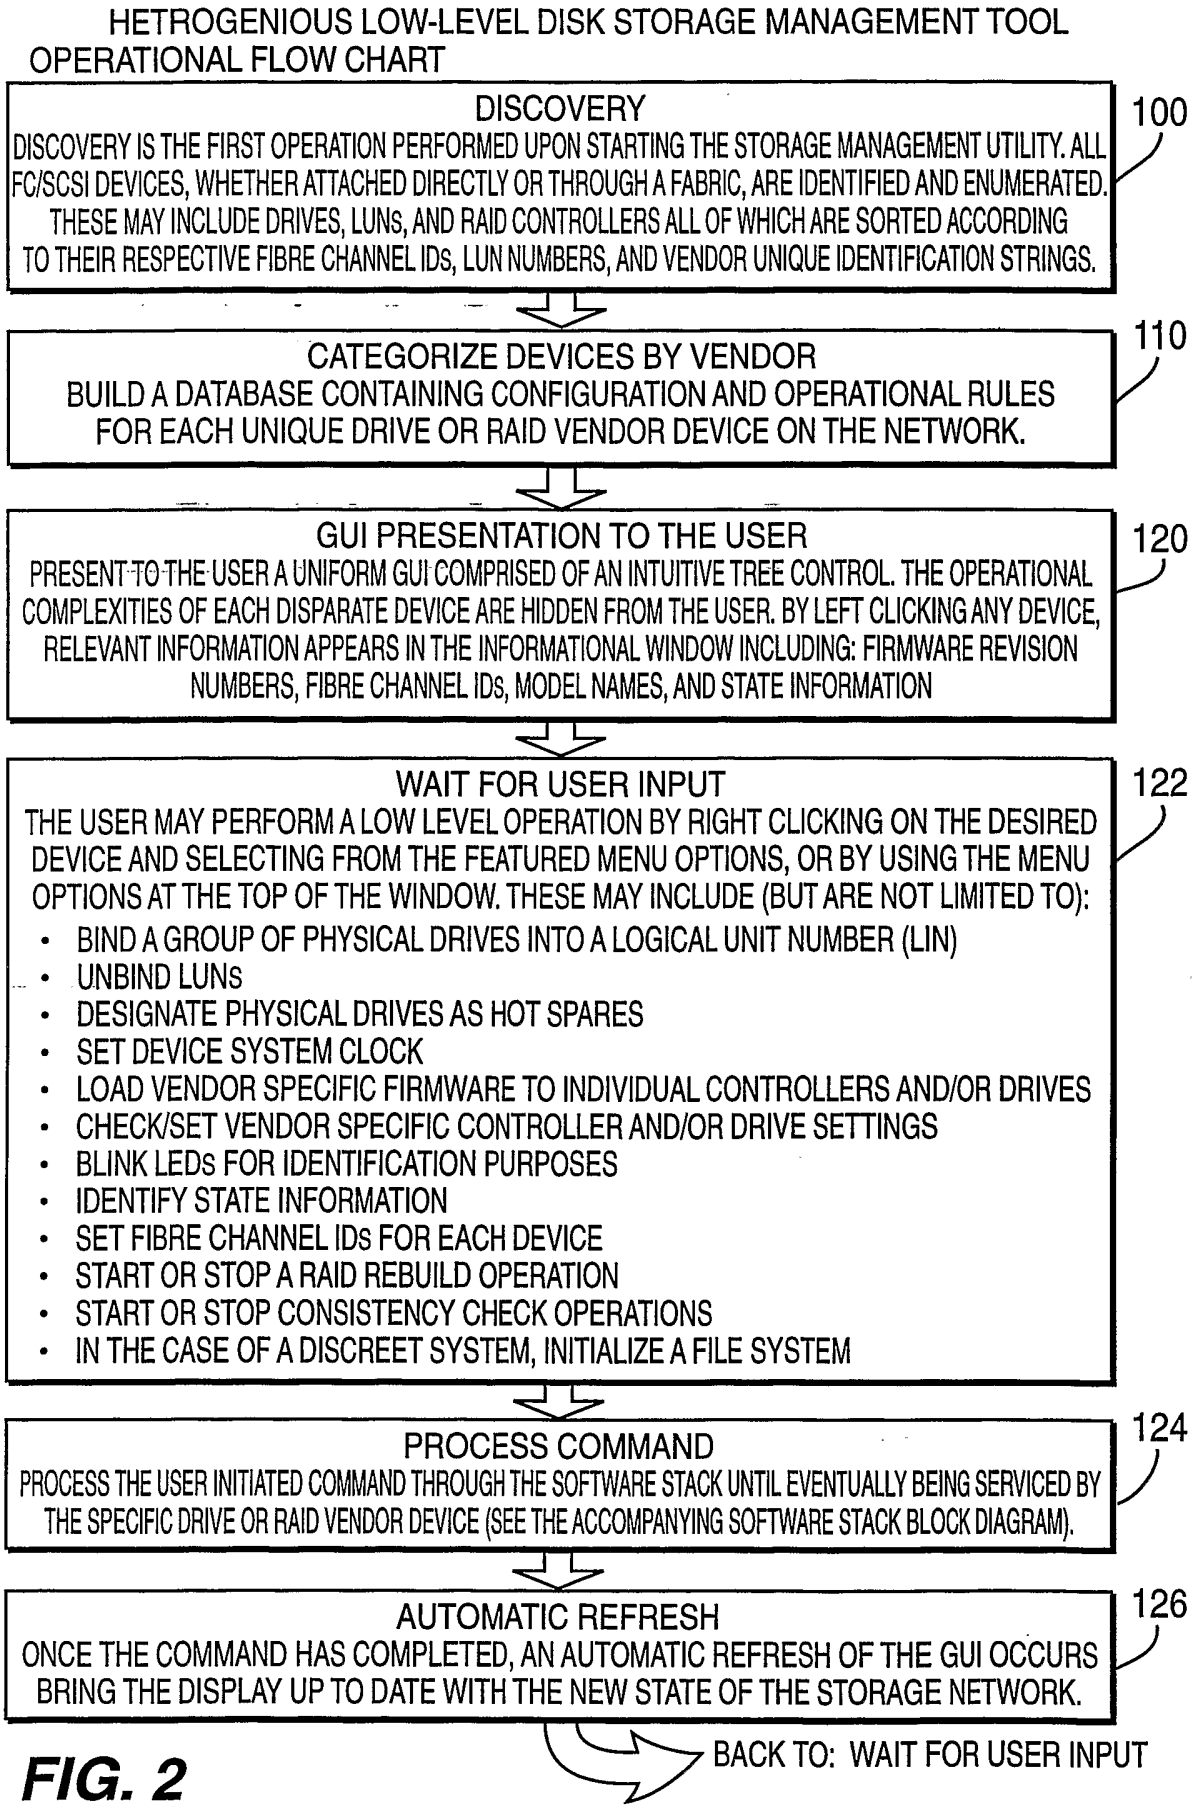
HETROGENIOUS LOW-LEVEL DISK STORAGE MANAGEMENT TOOL
OPERATIONAL FLOW CHART
DISCOVERY
DISCOVERY IS THE FIRST OPERATION PERFORMED UPON STARTING THE STORAGE MANAGEMENT UTILITY. ALL FC/SCSI DEVICES, WHETHER ATTACHED DIRECTLY OR THROUGH A FABRIC, ARE IDENTIFIED AND ENUMERATED. THESE MAY INCLUDE DRIVES, LUNs, AND RAID CONTROLLERS ALL OF WHICH ARE SORTED ACCORDING TO THEIR RESPECTIVE FIBRE CHANNEL IDs, LUN NUMBERS, AND VENDOR UNIQUE IDENTIFICATION STRINGS.
100
CATEGORIZE DEVICES BY VENDOR
BUILD A DATABASE CONTAINING CONFIGURATION AND OPERATIONAL RULES FOR EACH UNIQUE DRIVE OR RAID VENDOR DEVICE ON THE NETWORK.
110
GUI PRESENTATION TO THE USER
PRESENT TO THE USER A UNIFORM GUI COMPRISED OF AN INTUITIVE TREE CONTROL. THE OPERATIONAL COMPLEXITIES OF EACH DISPARATE DEVICE ARE HIDDEN FROM THE USER. BY LEFT CLICKING ANY DEVICE, RELEVANT INFORMATION APPEARS IN THE INFORMATIONAL WINDOW INCLUDING: FIRMWARE REVISION NUMBERS, FIBRE CHANNEL IDs, MODEL NAMES, AND STATE INFORMATION
120
WAIT FOR USER INPUT
THE USER MAY PERFORM A LOW LEVEL OPERATION BY RIGHT CLICKING ON THE DESIRED DEVICE AND SELECTING FROM THE FEATURED MENU OPTIONS, OR BY USING THE MENU OPTIONS AT THE TOP OF THE WINDOW. THESE MAY INCLUDE (BUT ARE NOT LIMITED TO):
• BIND A GROUP OF PHYSICAL DRIVES INTO A LOGICAL UNIT NUMBER (LIN)
• UNBIND LUNs
• DESIGNATE PHYSICAL DRIVES AS HOT SPARES
• SET DEVICE SYSTEM CLOCK
• LOAD VENDOR SPECIFIC FIRMWARE TO INDIVIDUAL CONTROLLERS AND/OR DRIVES
• CHECK/SET VENDOR SPECIFIC CONTROLLER AND/OR DRIVE SETTINGS
• BLINK LEDs FOR IDENTIFICATION PURPOSES
• IDENTIFY STATE INFORMATION
• SET FIBRE CHANNEL IDs FOR EACH DEVICE
• START OR STOP A RAID REBUILD OPERATION
• START OR STOP CONSISTENCY CHECK OPERATIONS
• IN THE CASE OF A DISCREET SYSTEM, INITIALIZE A FILE SYSTEM
122
PROCESS COMMAND
PROCESS THE USER INITIATED COMMAND THROUGH THE SOFTWARE STACK UNTIL EVENTUALLY BEING SERVICED BY THE SPECIFIC DRIVE OR RAID VENDOR DEVICE (SEE THE ACCOMPANYING SOFTWARE STACK BLOCK DIAGRAM).
124
AUTOMATIC REFRESH
ONCE THE COMMAND HAS COMPLETED, AN AUTOMATIC REFRESH OF THE GUI OCCURS BRING THE DISPLAY UP TO DATE WITH THE NEW STATE OF THE STORAGE NETWORK.
126
BACK TO: WAIT FOR USER INPUT

FIG. 2

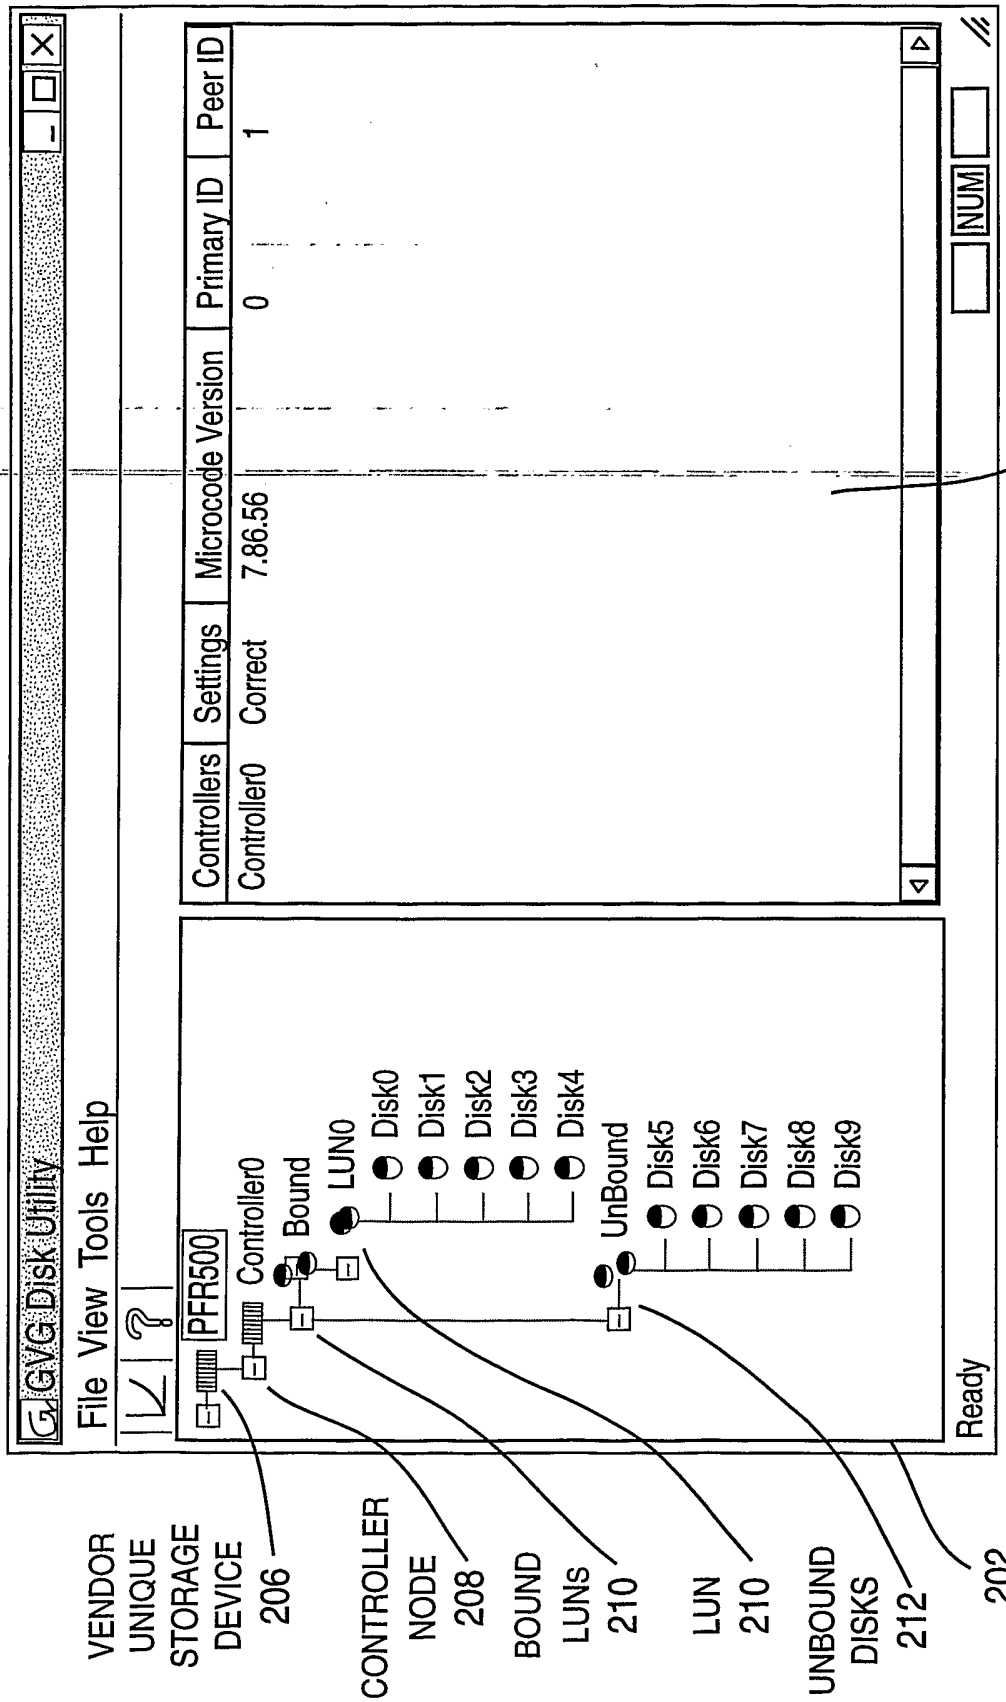
GVG Disk Utility
File View Tools Help
PFR500
Controller0
Bound
LUN0
Disk0
Disk1
Disk2
Disk3
Disk4
UnBound
Disk5
Disk6
Disk7
Disk8
Disk9
Ready
Controllers
Settings
Microcode Version
Primary ID
Peer ID
Controller0
Correct
7.86.56
0
1
NUM
214
200
VENDOR UNIQUE STORAGE DEVICE 206
CONTROLLER NODE 208
BOUND LUNs 210
LUN 210
UNBOUND DISKS 212
202

FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US03/17608

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 17/00; G09G 5/00
US CL : 707/1, 102; 345/810, 835, 848-855

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. : 707/1, 102; 345/810, 835, 848-855

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)
EAST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 6,346,954 B1 (CHU et al.) 12 February 2002 (12.02.2002), column 6, column 7.	1-11
Y, P	US 6,519,678 B1 (BASHAM et al.) 11 February 2003 (11.02.2003), column 7, column 8.	1-11
Y	US 6,266,740 B1 (DON et al.) 24 July 2001 (24.07.2001), column 5, lines 11-64, column 6, lines 28-49, column 8, line 51 to column 9, line 3.	1-11
Y	US 6,101,559 A (SCHULTZ et al.) 08 August 2000 (08.08.2000), column 5, line 37 to column 6, line 27.	1-11

☐ 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
05 August 2003 (05.08.2003)

Date of mailing of the international search report
22 AUG 2003

Name and mailing address of the ISA/US
Mail Stop PCT, Attn: ISA/US
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450
Facsimile No.

Authorized officer
Greta Robinson JM
James R. Matthews

Telephone No. 703-308-7565

Form PCT/ISA/210 (second sheet) (July 1998)