

Description

This invention generally relates to improvements in data storage devices for the input and output of information to a data processing system and, more particularly, to the retention of status and device identification information on a status track of each of the data storage devices and another location to provide global identification and management of interchangeable data storage devices.

The past several years have seen the growth of online workstations, the evolution of distributed processing, and the acceleration of information technology in many new application areas. The result has been an increase in the use of online data base systems and a growth in the requirement for storage capacity and increased reliability and flexibility in data storage devices.

The concepts of self-test, redundancy, cross-check verification between various sources of information and the like are well known in the art. Particularly, with the advent of complex digital computer systems used for applications such as process control or supervisory control. An example of such a system is illustrated in U. S. Patent 4,032,757 to Eccles, which uses a pair of channels to continually compare the events occurring in each computer. The constant comparison allows the additional processor to quickly take over control of the process if the other processor fails. The problem with this approach is the time that the extra processor needs to begin processing after a failure. In critical processes such as nuclear power plant control, any time lag could be unacceptable. A better approach is presented in U. S. Patent 4,270,168 to Murphy et al., which uses a plurality of processors, self checks and joint answer checking to assure that each processor can assume real time processing on behalf of another processor in the event of its failure. The increased reliability presented in these systems is a memory resident application approach that fails to handle a large data base application spanning many data storage devices.

The general configuration of a data processing system typically comprises a host processor or processors, a memory and various peripheral units. The peripheral units include terminals, printers, communications devices and data storage devices (DASD). We are concerned with the control that provides information from DASD to a data base application residing in the host processor memory. A good example of prior art approaches to this type of processing is presented in US Patent 3,999,163 to Levy et al., U.S. Patent 4,067,059 to Derchak and U.S. Patent 4,189,769 to Cook et al. These patents present various ways to enable a

host to process information residing on DASD. While these patents describe production systems that readily lend themselves to data base applications, the arrangements they describe are vulnerable in that they do not maintain high availability of DASD information when a power-off occurs.

EP 59838 discloses a data processing apparatus having a peripheral system including a plurality of channel paths. These channel paths are named as members of path groups. Each path group has one or more channel paths. Within each path group, communications between the data processing system and the peripheral system are multipathed among the paths in such path groups. Devices within the peripheral system are selectively assignable to access groups for exclusive use within one or more of such path groups. A given path group selectively provides communication paths for each device is an access group. Each access group includes one or more path groups. Temporary suspension and reassignment facilities are disclosed. Both devices and channel paths can be either in a grouped or ungrouped state. The assignment function to access groups is particularly useful in interleaved channel transfers for ensuring subsystem integrity.

Accordingly, the present invention provides a method for managing a data processing system having a host system, an attached controller comprising a set of storage paths attached to a cache, a data storage device controller, and a nonvolatile storage, and a memory containing an array structure comprising a status table containing the status records of the data storage devices, the system generating a global status record combining the status records of the controller and of the attached data storage devices, and the data storage device controller comprising attaching means for a plurality of strings of data storage devices each of the strings comprising a plurality of data storage devices having a plurality of records, a plurality of tracks, and a device status track containing the status information of the data storage device, characterised in that:

said device status track also comprises a reconfiguration count for recording the number of times the device has been reassigned, a device identification number and a controller identification number recorded thereon;

said status table in the array structure also comprises a reconfiguration count, a device identification number, and a controller identification number for each of the data storage devices;

said global status record also comprises the device identification number and the controller identification number; and

the global identification and management of the data storage devices is assured by:

checking each of the plurality of data storage devices to verify the device identification number, the controller identification number and the status records;

generating the global status record;

writing the device identification number to each of the status records on each of the data storage devices; and

writing the global status record to the status records of a first data storage device and a second data storage device.

Put another way, the present invention provides a data processing system for assuring the global identification and management of a plurality of data storage devices comprising:

a host system containing a plurality of channels;

a controller attached to at least one of the channels and comprising a set of storage paths, each of the storage paths having access to a cache, a data storage device controller, a non-volatile storage, and a memory comprising an array structure having a status table containing the status records of the data storage devices;

a global status record combining the status records of the controller and of the attached data storage devices, which is stored in the status table of the array in the controller;

a plurality of data storage devices attached to the data storage device controller, each of the data storage devices comprising a plurality of tracks, and a device status track containing the status information of the data storage device;

characterised in that:

said device status track also comprises a re-configuration count for recording the number of times the device has been reassigned, a device identification number and a controller identification number recorded thereon;

said status table in the array structure also comprises a reconfiguration count, a device identification number, and a controller identification number for each of the data storage devices;

said global status record also comprises the device identification number and the controller identification number; and

the system includes:

means for checking each of the plurality of data storage devices to verify the device identification number, the controller identification number and the status records;

means for writing the device identification number to each of the status records on each of the data storage devices; and

means for writing the global status record to the status records of a first data storage device and a second data storage device.

The present invention will be described further by way of example with reference to an embodiment thereof as well as an example of the prior art and illustrated in the accompanying drawings, in which:

Fig. 1 is a block diagram showing a typical prior art data storage system, namely that of U.S. Patent 4,207,609; and Fig. 2 is a highly stylised block diagram showing the basic elements of one form of data storage system according to the present invention.

The following disclosure describes one form of improved management of data storage devices and the data stored on them. Status and Identification information is retained on the data storage devices as well as on a backup medium to assure the retention of information when a power out or other disablement of the devices occur and to verify that the correct physical devices are still attached to the controllers when the data storage device resumes operation.

Referring now to the drawings, and more particularly to Fig. 1, there is depicted a typical prior art data storage system. Fig. 1 shows a multi-CPU and shared DASD configuration according to U. S. Patent 4, 207, 609 assigned to the assignee of this invention. Three CPUs 1, 3, and 4 hereinafter designated by Roman numerals respectively CPU I CPU II and CPU III, are suitably cross-connected to a pair of control units 21 and 23, hereinafter designated CU I and CU II, over counterpart channels. In this regard, CPU I attaches four channels 5, 7, 8 and 10; CPU II attaches two channels 9, 12, and CPU III attaches two channels 14 and 16. Channels 1 and 2 of CPU I terminate in CU I ports a and b, while channels 3 and 4 of CPU I terminate in CU II ports e and f. Channel 1 of CPU II terminates in CU I port c with channel 2 of CPU II attaching CU II at port g. Lastly channel 1 of CPU III drives port d of CU I with channel 2 of CPU III driving port h of CU II. CU I and CU II attach devices 53 through DASD controllers 27 and 29 over a demand/response interface embracing tag and data lines. The interface between CU I and DASD controller 27 includes tag lines 65 and data/in and data/out lines 57 and 59. Likewise, tag lines 67 and data/in and data/out paths 61 and 63 couple CU II to DASD controller 29.

The demand/response type interface is one in which the tag out identifies and validates the information on the data lines out and in. In this regard, each CPU/control unit/device operates asynchronously with respect to each other and directional control is asserted from the top down. Active connection between elements for executing operations is necessary for extended periods only in respect of the movement of data from device to channel or from channel to device. Otherwise,

tasks or operations may be performed in disconnect mode.

Devices 53 are accessed either through DASD controller 27 or DASD controller 29. DASD controller 27 includes attachment logic 39 and sequence controller 41 coupling devices 53 over path 55. Similarly, DASD controller 29 includes attachment logic 47 and sequence controller 49 coupling devices 53 over path 51. This memory contains the map of network topology and the necessary context information utilised by the CUs in managing the connection between CPUs to devices. Memory 25' comprises a pair of random access memories (RAM) 43 and 45 which can be commonly accessed by either CU through the appropriate attachment logic. CU I can access, for reading, writing or updating purposes, RAMs 43 and 45 through attachment logic 39 and path 67. CU II can commonly access the RAM's via attachment logic 47 and path 65. The operations concerned with accessing of devices 53 involving the dispatching and positioning of electro-mechanical elements such as the head disk arm assembly on direct access storage devices 0-4 is regulated by the commands and parameters sent from the control unit through the attachment logic and sequence controller.

An example of the host CPU is the IBM 3090 processor. The IBM 3090 is a compatible growth system which uses either Multiple Virtual Storage/Extended Architecture (MVS/XA) or Virtual Machine (VM) as an operating system. The operating system controls the flow of data to/from main storage and provides an application environment which is System/370 compatible. The IBM 3090 processor is described in more detail in IBM publication, 3090 Processor Complex Installation Manual - Physical Planning, GC22-7074. An example of DASD controllers 27 and 29 is the IBM 3880 Model 3 DASD Controller. The IBM 3880 Model 23 contains a high speed cache and two storage directors 41 and 49, each with a single storage path for independent control of attached DASD devices 53. The IBM 3880 controller is described in IBM publication, IBM 3880 Storage Control Model 23 Description, GC32-0082. A good example of a DASD device 0-3 is the IBM 3380, which provides a fixed-medium head and disk assembly (HDA) that contains the heads, disks and access mechanisms or actuators within a sealed enclosure. Each 3380 unit contains two HDAs. There are two independent, movable actuators within each HDA, effective on a set of disks mounted on a spindle common to both actuators. Each actuator has its own address which is selectable from the CPUs. A string of DASDs attached to a 3880 storage director can contain four units which implies sixteen actuators. A more detailed description of the IBM 3380 DASD is found in IBM publication, 3380 Description and

User's Guide, GA26-1664.

To better understand the embodiment of the present invention hereinafter disclosed, it will be necessary to discuss some of the details of this prior art DASD system operation. When the Host CPU requires information, it requests it by requesting a specific channel, controller and DASD actuator. The channel is identified by a number which is a one byte hexadecimal number, the controller by a four bit hexadecimal number and the DASD actuator number by a four bit hexadecimal number. For example, if the host CPU were to send a two byte hexadecimal address 0111 in a start input/output operation (I/O), then actuator one, attached to controller one, attached to channel one would prepare for I/O and send the host a ready signal. Since the mapping down and back is the same, the original transaction header can be used to determine the path back to the host. Any host program accessing the DASD could send the header described above to uniquely define the DASD actuator to access.

Fig. 2 shows the basic elements of one data storage sub-system of the embodiment, which, although similar to that shown in Fig. 1, is shown in a very generalised block form. The assumption that has been made in Fig. 2, is that the maximum possible integers are physically connected and working. It must be remembered that the invention is intended to deal with circumstances in which this becomes progressively untrue so that there will be times in the following discussion when, for example, it will appear that only two HDA's are connected, despite what is shown in Fig. 2. It is immaterial from the point of the discussion as to whether this is because the remainder of the potential items are there but incorrectly, are there but have failed or just are not connected or not yet connected.

To return to the general situation depicted in Fig. 2, a host CPU 110 is attached via its channels 120 to a DASD controller 125, incorporating its two directors 180, to each of which is attached a respective DASD string 70 and 75. Any or all HDA's may belong to both strings. Information is passed through whichever path is specified by the host. A major difference between the described prior art and the embodiment lies in the handling and use of status information to manage the data storage system and the non-volatile storage 220. The DASD controller 125 also contains an EPROM vital product data area 95. The vital product data area contains pointers to where global status is stored.

A joint array structure (JAS) 200 is used to conserve the host to DASD interface and implement dual copy, high availability functions. The dual copy function provides enhanced data availability and reliability by maintaining two identical

copies, a primary copy and a secondary copy, of volumes, called a duplex pair, on two physically separated DASD units, for example, a primary data storage device in string 70 and a secondary data storage device in string 75. How these are used in the environment of Fig. 2 will now be described.

The data storage devices in strings 70 and 75 each have a dedicated status track for the retention of status information. The dedicated status track is unusable by non-system applications and retains two copies of the global subsystem status for the controller and the first and second data storage devices. The storage subsystem is the set consisting of the controller and the attached data storage devices. In the example, this consists of the single controller 125 and the two data storage devices 70 and 75.

Status information about the subsystem is retained across power downs by writing a single record of status and identification information twice on the status track of each device. Global subsystem status is also kept on the first and second data storage devices 70 and 75. Device identification is kept on all the devices in the subsystem, but global subsystem status is only kept on the two global data storage devices. The global subsystem status portion of the record is formatted but not used if the device is not one of the two global subsystem status devices.

The global status pointer contains address information about the two data storage devices and a count of how many times the devices have been reassigned, which data is used at initial program load to determine which is the most current copy of the global status pair. A more detailed description of the global status pointer is provided below.

The global subsystem status consists of the following information:

- the global status pointer - which is a four byte field including the one-byte address of each of the data storage devices including the first 70 and second 75 global data storage devices and

- a two-byte number which indicates the number of times the definition of the data storage devices have changed.

Global status contains the status of the subsystem components, a cache store 145 and a non-volatile storage 220, and devices as established by failure or system commands.

Information on the status track is divided into device identification information and status information. Device identification information is used to insure that the device attached is the one that the subsystem expects. The status information is divided into two types, subsystem status and device status. Subsystem status contains information that pertains to the entire subsystem rather than to a specific device such as the global status pointer,

whether the subsystem storage is available or not and information to help insure the integrity of the non-volatile store. The device status contains information that pertains to a specific DASD device such as the channel address that the device will respond to, whether the device is active for cache or fast write or not, and the duplex pair status of the device. Other information is included with the status track, the above are examples of the information identified.

The global subsystem status information discussed above is kept on two devices for availability reasons. However, each device has its device number written. If one device should fail, then the other will be used to keep the information and the failed device will be replaced by another operational device.

Each controller has vital product data that is maintained and used to identify and establish information concerning the hardware characteristics of the attached devices. The vital product data is stored in EPROM in each controller as depicted at 95 and 96 of Fig. 2. The information includes:

- Model Type
- Configuration
- System Adapter and System Address
- Subsystem Storage Identification
- Storage Control Serial Numbers
- Subsystem Storage Size
- Non-volatile Storage Size
- Global Status Pointer
- Controller Type For Each String
- Maximum Number of Drives Configured for Each String.

This information is used to initialise the JAS 200 in the controller 25 to reflect the current system's configuration. In our example, the JAS 200 contains the addresses that data storage device 70 and data storage device 75 are known to the host 110.

When the subsystem is initialised, the following steps are performed:

The devices undergo an initial microprogram load (IML) to establish an operational data path between the attaching channels and the attached devices. The Storage Path (SP) microcode is loaded from the IML device by the support facility in the data storage device. A SP 140 is an addressable, independent processor that routes data to a data storage device 70 or 75. A basic assurance test routine is run to insure that the storage path is operational. This process is coordinated for all of the SPs 140 in the subsystem residing on the control 125. Each SP 140 can complete IML independent of the other SPs. The location of a global status device is required before access is permitted. If an IML device pointer is available, the most recent pointer must agree with the available

device or devices. If one IML device pointer is available, it must point to one device with identical pointers. The IML device pointer does not have to be identical to the device pointer.

Then, the global subsystem status array is read from each of the two global subsystem data storage devices 75 and 70. When valid and current global subsystem status is identified, the global subsystem status is written to both data storage devices to insure the information is consistent on both of the devices. As devices become available, the device number and subsystem identification number of each data storage device is read and checked against the global status array to ensure that the correct hardware is attached and no movement has occurred. When this process is completed, the host is notified of the available SPs 140 and normal processing commences. Device identification is verified to assure that data left in the non-volatile store for a device is not applied to the wrong device because of a device reconfiguration as a result of device replacement, for instance. The controller can establish a pair of status data storage devices as long as a pair of operational devices can be found to hold the status, or one of the two previous status devices is one of the previous pair and one initialised device can be updated, or both initialised devices can be updated to show the new pair. If the controller cannot establish the two status data storage devices, then it will continue to retry and return an error code for any attempts to access the failing device.

If the current device availability allows a better global subsystem status device configuration, the devices will be reconfigured following these rules:

- 1) put the devices on separate strings;
- 2) put the devices on separate units; and
- 3) put the devices on separate spindles.

The global subsystem status devices are also reassigned by the controller when the controller can only access one of a pair of devices. If it is necessary for any of the reasons described above to reconfigure the system, the order of update will be as follows:

- 1) Increment the reconfiguration count in the global status pointer to show a reconfiguration has taken place.
- 2) Global subsystem status will be written to the new global status device. Its status will reflect the new pair.
- 3) The global subsystem status will be written to the old but still working global status device (in the case of the malfunctioning unit).
- 4) The global status pointer will be updated in both EPROM vital product areas 95 and 96.

The reconfigured system will be optimally tuned and ready for use when the operations described above are completed.

In operation, if either device fails, the controller will replace the failing device and copy the global status to another storage device. The assignment of a new data storage device is based on the following preference, if two strings of data storage devices are available, as in our case with 70 and 75, then the first available data storage device on each string will be used; however, if only one string is available, then the operational devices in separate physical units will be used.

If new devices are added, or devices are replaced, the devices will be adapted and the status track written as long as the device status on the global subsystem status shows that the device position was not assigned to a fast write or dual copy volume. At installation or initialisation, the global status pointers will be 'zeros'. The subsystem will format and write device IDs in all device status tracks, assign global status devices and write global status and write global status pointers to the EPROM.

While the invention has been described in terms of a preferred embodiment in a specific operating system environment, those skilled in the art will recognise that the invention can be practiced, with modification, in other and different operating systems.

Claims

1. A method for managing a data processing system having a host system (110), an attached controller (125) comprising a set of storage paths attached to a cache, a data storage device controller (160), and a nonvolatile storage (220), and a memory containing an array structure (200) comprising a status table containing the status records of the data storage devices, the system generating a global status record combining the status records of the controller and of the attached data storage devices, and the data storage device controller comprising attaching means for a plurality of strings of data storage devices (75, 70) each of the strings comprising a plurality of data storage devices having a plurality of records, a plurality of tracks, and a device status track containing the status information of the data storage device, characterised in that:

said device status track also comprises a reconfiguration count for recording the number of times the device has been reassigned, a device identification number and a controller identification number recorded thereon;

said status table in the array structure also comprises a reconfiguration count, a device identification number, and a controller identification number for each of the data storage



device or devices. If one IML device pointer is available, it must point to one device with identical pointers. The IML device pointer does not have to be identical to the device pointer.

Then, the global subsystem status array is read from each of the two global subsystem data storage devices 75 and 70. When valid and current global subsystem status is identified, the global subsystem status is written to both data storage devices to insure the information is consistent on both of the devices. As devices become available, the device number and subsystem identification number of each data storage device is read and checked against the global status array to ensure that the correct hardware is attached and no movement has occurred. When this process is completed, the host is notified of the available SPs 140 and normal processing commences. Device identification is verified to assure that data left in the non-volatile store for a device is not applied to the wrong device because of a device reconfiguration as a result of device replacement, for instance. The controller can establish a pair of status data storage devices as long as a pair of operational devices can be found to hold the status, or one of the two previous status devices is one of the previous pair and one initialised device can be updated, or both initialised devices can be updated to show the new pair. If the controller cannot establish the two status data storage devices, then it will continue to retry and return an error code for any attempts to access the failing device.

If the current device availability allows a better global subsystem status device configuration, the devices will be reconfigured following these rules:

- 1) put the devices on separate strings;
- 2) put the devices on separate units; and
- 3) put the devices on separate spindles.

The global subsystem status devices are also reassigned by the controller when the controller can only access one of a pair of devices. If it is necessary for any of the reasons described above to reconfigure the system, the order of update will be as follows:

- 1) Increment the reconfiguration count in the global status pointer to show a reconfiguration has taken place.
- 2) Global subsystem status will be written to the new global status device. Its status will reflect the new pair.
- 3) The global subsystem status will be written to the old but still working global status device (in the case of the malfunctioning unit).
- 4) The global status pointer will be updated in both EPROM vital product areas 95 and 96.

The reconfigured system will be optimally tuned and ready for use when the operations described above are completed.

In operation, if either device fails, the controller will replace the failing device and copy the global status to another storage device. The assignment of a new data storage device is based on the following preference, if two strings of data storage devices are available, as in our case with 70 and 75, then the first available data storage device on each string will be used; however, if only one string is available, then the operational devices in separate physical units will be used.

If new devices are added, or devices are replaced, the devices will be adapted and the status track written as long as the device status on the global subsystem status shows that the device position was not assigned to a fast write or dual copy volume. At installation or initialisation, the global status pointers will be 'zeros'. The subsystem will format and write device IDs in all device status tracks, assign global status devices and write global status and write global status pointers to the EPROM.

While the invention has been described in terms of a preferred embodiment in a specific operating system environment, those skilled in the art will recognise that the invention can be practiced, with modification, in other and different operating systems.

Claims

1. A method for managing a data processing system having a host system (110), an attached controller (125) comprising a set of storage paths attached to a cache, a data storage device controller (160), and a nonvolatile storage (220), and a memory containing an array structure (200) comprising a status table containing the status records of the data storage devices, the system generating a global status record combining the status records of the controller and of the attached data storage devices, and the data storage device controller comprising attaching means for a plurality of strings of data storage devices (75, 70) each of the strings comprising a plurality of data storage devices having a plurality of records, a plurality of tracks, and a device status track containing the status information of the data storage device, characterised in that:

said device status track also comprises a reconfiguration count for recording the number of times the device has been reassigned, a device identification number and a controller identification number recorded thereon;

said status table in the array structure also comprises a reconfiguration count, a device identification number, and a controller identification number for each of the data storage



devices;

said global status record also comprises the device identification number and the controller identification number; and

the global identification and management of the data storage devices is assured by:

checking each of the plurality of data storage devices to verify the device identification number, the controller identification number and the status records;

generating the global status record;

writing the device identification number to each of the status records on each of the data storage devices; and

writing the global status record to the status records of a first data storage device and a second data storage device.

2. A method as claimed in claim 1, in which a failing data storage device is replaced with a functioning one of the data storage devices and information identifying the functioning data storage device is updated as the replacement for the failed data storage device.

3. A method as claimed in either preceding claim, in which an initial memory load of a data storage device is prevented from successfully completing unless the controller identification number and the device identification number in the status table of the controller match the device and controller identification numbers on the device status track of the data storage device.

4. A method as claimed in any preceding claim wherein the data storage devices are configured to distribute information flow over each of the strings by:

connecting each of the data storage devices to the string with the least attached data storage devices; and

connecting the plurality of strings to more than one of the data storage device controller attached means.

5. A method as claimed in any preceding claim, wherein the system response to a global status data storage device failure comprises:

incrementing the reconfiguration count in the status table of the joint array structure to indicate that a reconfiguration has taken place; and

writing the incremented reconfiguration count onto the device status track of the functioning data storage device.

6. A data processing system for assuring the global identification and management of a plurality of data storage devices comprising:

a host system (110) containing a plurality of channels;

a controller (125) attached to at least one of the channels and comprising a set of storage paths, each of the storage paths having access to a cache, a data storage device controller (160), a nonvolatile storage (220), and a memory comprising an array structure (200) having a status table containing the status records of the data storage devices;

a global status record combining the status records of the controller and of the attached data storage devices, which is stored in the status table of the array in the controller;

a plurality of data storage devices (75, 70) attached to the data storage device controller, each of the data storage devices comprising a plurality of tracks, and a device status track containing the status information of the data storage device;

characterised in that:

said device status track also comprises a reconfiguration count for recording the number of times the device has been reassigned, a device identification number and a controller identification number recorded thereon;

said status table in the array structure also comprises a reconfiguration count, a device identification number, and a controller identification number for each of the data storage devices;

said global status record also comprises the device identification number and the controller identification number; and

the system includes:

means for checking each of the plurality of data storage devices to verify the device identification number, the controller identification number and the status records;

means for writing the device identification number to each of the status records on each of the data storage devices; and

means for writing the global status record to the status records of a first data storage device and a second data storage device.

Patentansprüche

1. Ein Verfahren zur Verwaltung eines Datenverarbeitungssystems mit einem Hauptrechnersystem (110), einem angeschlossenen Steuerteil (125), welches einen Satz von an einen Cache angeschlossenen Speicherpfaden umfaßt, einen Steuerteil für die Datenspeicherungseinrichtung (160) und einen nicht flüchtigen Spei-



- cher (220) enthält, und einen Speicher mit einer Matrix-Struktur (200), der eine Statustabelle umfaßt, welche die Statussätze der Datenspeichereinrichtungen enthält, wobei das System einen globalen Statussatz erzeugt, der die Statussätze des Steuerteils und der angeschlossenen Datenspeichereinrichtungen verbindet, und der Steuerteil der Datenspeichereinrichtung Anschlußmittel für eine Vielzahl von Ketten von Datenspeichereinrichtungen (75, 70) enthält, wobei jede Kette eine Vielzahl von Datenspeichereinrichtungen mit einer Vielzahl von Sätzen, einer Vielzahl von Spuren und einer Einrichtungsstatusspur enthält, welche die Statusinformationen der Datenspeichereinrichtung enthält, dadurch gekennzeichnet, daß:
- die genannte Einrichtungsstatusspur auch eine Neukonfigurationszählung umfaßt, die aufzeichnet, wie oft die Einrichtung neu zugewiesen wurde, eine Einrichtungs-Identifizierungsnummer und eine darauf aufgezeichnete Steuerteil-Identifizierungsnummer;
- die genannte Statustabelle in der Matrix-Struktur auch eine Neukonfigurationszählung, eine Einrichtungs-Identifizierungsnummer und eine Steuerteil-Identifizierungsnummer für jede Datenspeichereinrichtung umfaßt;
- der genannte globale Statussatz auch die Einrichtungs-Identifizierungsnummer und die Steuerteil-Identifizierungsnummer umfaßt; und
- die globale Identifizierung und Verwaltung der Datenspeichereinrichtungen sichergestellt wird durch:
- Prüfen jeder Datenspeichereinrichtung der Vielzahl von Datenspeichereinrichtungen zur Feststellung der Einrichtungs-Identifizierungsnummer, der Steuerteil-Identifizierungsnummer und der Statussätze;
- Erzeugung des globalen Statussatzes;
- Schreiben der Einrichtungs-Identifizierungsnummer auf jeden Statussatz auf jeder Datenspeichereinrichtung; und
- Schreiben des globalen Statussatzes auf die Statussätze einer ersten Datenspeichereinrichtung und einer zweiten Datenspeichereinrichtung.
2. Ein Verfahren gemäß Anspruch 1, in welchem eine gestörte Datenspeichereinrichtung gegen eine funktionsfähige Datenspeichereinrichtung ausgetauscht wird, und Informationen, die die funktionsfähige Datenspeichereinrichtung identifizieren, als Ersatz für die gestörte Datenspeichereinrichtung aktualisiert werden.

3. Ein Verfahren gemäß jedem der vorhergehenden Ansprüche, in welchem eine erste Speicheraladung einer Datenspeichereinrichtung solange verhindert wird, bis die Steuerteil-Identifizierungsnummer und die Einrichtungs-Identifizierungsnummer in der Statustabelle des Steuerteils mit den Einrichtungs- und Steuerteil-Identifizierungsnummern auf der Einrichtungsstatusspur der Datenspeichereinrichtung übereinstimmen.
4. Ein Verfahren gemäß jedem der vorhergehenden Ansprüche, in welchem die Datenspeichereinrichtungen so konfiguriert sind, daß sie den Informationsfluß über jede Kette folgendermaßen verteilen:
- durch Verbinden jeder Datenspeichereinrichtung mit der Kette, an die die wenigsten Datenspeichereinrichtungen angeschlossen sind; und
- durch Verbinden der Vielzahl von Ketten mit mehr als einem an den Steuerteil der Datenspeichereinrichtung angeschlossenen Mittel.
5. Ein Verfahren gemäß jedem der vorhergehenden Ansprüche, in dem das Ansprechen des Systems auf eine Störung einer Globalstatus-Datenspeichereinrichtung folgendes umfaßt:
- Hochzählen der Neukonfigurationszählung in der Statustabelle der gemeinsamen Matrix-Struktur, um anzuzeigen, daß eine Neukonfiguration stattgefunden hat; und
- Schreiben der hochgezählten Neukonfigurationszählung auf die Einrichtungsstatusspur der funktionsfähigen Datenspeichereinrichtung.
6. Ein Datenverarbeitungssystem zur Sicherstellung der globalen Identifizierung und Verwaltung einer Vielzahl von Datenspeichereinrichtungen, welches folgendes umfaßt:
- ein Hauptrechnersystem (110) mit einer Vielzahl von Kanälen;
- einen Steuerteil (125), der an mindestens einen Kanal angeschlossen ist, und eine Reihe von Speicherpfaden umfaßt, wobei jeder Speicherpfad Zugang zu einem Cache, einem Steuerteil für eine Datenspeichereinrichtung (160), einem nicht flüchtigen Speicher (220) und einem Speicher mit einer Matrix-Struktur (200) hat, der eine Statustabelle mit den Statusätzen der Datenspeichereinrichtungen enthält;
- einen globalen Statussatz, der die Statussätze des Steuerteils und der angeschlossenen Datenspeichereinrichtungen verbindet, der



in der Statustabelle der Matrix in dem Steuer-
teil gespeichert ist;

eine Vielzahl von Datenspeicherungsein-
richtungen (75, 70), die an den Steuer-
teil für die Datenspeicherungseinrichtungen ange-
schlossen sind, wobei jede Datenspeiche-
rungseinrichtung eine Vielzahl von Spuren um-
faßt, und eine Einrichtungsstatusspur, welche
die Statusinformation der Datenspeicherung-
einrichtung enthält;

dadurch gekennzeichnet, daß:

die genannte Einrichtungsstatusspur auch
eine Neukonfigurationszählung umfaßt, die auf-
zeichnet, wie oft die Einrichtung neu zugewie-
sen wurde, eine Einrichtungs-Identifizierungs-
nummer und eine darauf aufgezeichnete
Steuer-Identifizierungsnummer;

die genannte Statustabelle in der Matrix-
Struktur auch eine Neukonfigurationszählung,
eine Einrichtungs-Identifizierungsnummer und
eine Steuer-Identifizierungsnummer für jede
Datenspeicherungseinrichtung umfaßt;

der genannte globale Statussatz auch die
Einrichtungs-Identifizierungsnummer und die
Steuer-Identifizierungsnummer umfaßt; und

das System folgendes enthält:

Mittel zur Prüfung jeder Datenspeiche-
rungseinrichtung einer Vielzahl von Datenspei-
chungseinrichtungen zur Feststellung der
Einrichtungs-Identifizierungsnummer, der
Steuer-Identifizierungsnummer und der Sta-
tussätze;

Mittel zum Schreiben der Einrichtungs-
Identifizierungsnummer auf jeden der Status-
sätze in jeder der Datenspeicherungseinrich-
tungen; und

Mittel zum Schreiben des globalen Status-
satzes auf die Statussätze einer ersten Daten-
speicherungseinrichtung und einer zweiten Daten-
speicherungseinrichtung.

Revendications

1. Procédé de gestion d'un système de traite-
ment de données ayant un système hôte
(110), un contrôleur rattaché (125) comprenant
un ensemble de chemins de mémoire ratta-
chés à une antémémoire, un contrôleur de
dispositifs d'emménagement de données (160)
et une mémoire permanente (220), et une mé-
moire contenant une structure en tableau (200)
comportant une table d'états contenant les
enregistrements d'état des dispositifs d'emma-
gasinage de données, le système engendrant
un enregistrement d'état global combinant les
enregistrements d'état du contrôleur et des
dispositifs d'emménagement de données ratta-
chés, et le contrôleur de dispositifs d'emma-

sinage de données comportant des moyens
de rattachement d'une pluralité de chaînes de
dispositifs d'emménagement de données (75,
70), chacune des chaînes comportant une plu-
ralité de dispositifs d'emménagement de don-
nées ayant une pluralité d'enregistrements,
une pluralité de pistes, et une piste d'état de
dispositif contenant les informations d'état du
dispositif d'emménagement de données, carac-
térisé en ce que:

Ladite piste d'état de dispositif comprend
également un compte de reconfigurations pour
l'enregistrement du nombre de fois qu'un dis-
positif a été ré-attribué, un numéro d'identifica-
tion de dispositif et un numéro d'identification
de contrôleur enregistrés dessus;

ladite table d'état dans la structure en ta-
bleau comprend également un compte de re-
configurations, un numéro d'identification de
dispositif, et un numéro d'identification de
contrôleur pour chacun des dispositifs d'emma-
gasinage de données;

ledit enregistrement d'état global com-
prend également le numéro d'identification de
dispositif et le numéro d'identification de
contrôleur, et

L'identification et la gestion globales des
dispositifs d'emménagement de données sont
assurées par:

le contrôle de chacun de la pluralité des
dispositifs d'emménagement de données pour
vérifier le numéro d'identification de dispositif,
le numéro d'identification de contrôleur et les
enregistrements d'état;

la génération de l'enregistrement d'état
global;

l'écriture du numéro d'identification de dis-
positif sur chacun des enregistrements d'état
de chacun des dispositifs d'emménagement de
données, et

l'écriture de l'enregistrement d'état global
sur les enregistrements d'état d'un premier
dispositif d'emménagement de données et d'un
deuxième dispositif d'emménagement de don-
nées.

2. Procédé selon la revendication 1, dans lequel
un dispositif d'emménagement de données dé-
faillant est remplacé par un dispositif en fonc-
tionnement des dispositifs d'emménagement de
données, et l'information identifiant le dispositif
d'emménagement de données en fonctionne-
ment est mise à jour lors du remplacement du
dispositif d'emménagement de données défail-
lant.
3. Procédé selon l'une ou l'autre des revendica-
tions précédentes, dans lequel un chargement



- en mémoire initial d'un dispositif d'emmagasina-
ge de données est empêché d'être mené à
bien à moins que le numéro d'identification de
contrôleur et le numéro d'identification de dis-
positif dans la table d'états du contrôleur ne
coïncide avec les numéros d'identification de
dispositif et de contrôleur sur la piste d'état de
dispositif du dispositif d'emmagasina-
ge de données.
4. Procédé selon l'une quelconque des revendica-
tions précédentes, dans lequel les dispositifs
d'emmagasina-ge de données sont configurés
de manière à répartir le débit des informations
sur chacune des chaînes en:
- connectant chacun des dispositifs d'em-
magasina-ge de données à la chaîne ayant le
moins de dispositifs d'emmagasina-ge de don-
nées rattachés; et
- connectant la pluralité de chaînes à plus
d'un des moyens rattachés au contrôleur de
dispositifs d'emmagasina-ge de données.
5. Procédé selon l'une quelconque des revendica-
tions précédentes, dans lequel la réponse du
système à une défaillance de dispositif d'em-
magasina-ge de données d'état global com-
prend:
- l'incrémentation du compte de reconfigura-
tions dans la table d'états de la structure en
tableau conjointe pour indiquer qu'une reconfi-
guration a eu lieu; et
- l'écriture du compte de reconfigurations
incrémenté sur la piste d'état de dispositif du
dispositif d'emmagasina-ge de données en
fonctionnement.
6. Système de traitement de données pour assu-
rer l'identification et la gestion globales d'une
pluralité de dispositifs d'emmagasina-ge de
données, comprenant:
- un système hôte (110) contenant une plu-
ralité de canaux;
- un contrôleur (110) rattaché à au moins un
des canaux et comprenant un ensemble de
chemins de mémoire, chacun des chemins de
mémoire ayant accès à une antémémoire, un
contrôleur de dispositifs d'emmagasina-ge de
données (160), une mémoire permanente
(220), et une mémoire comprenant une structu-
re en tableau (200) ayant une table d'états
contenant les enregistrements d'état des dis-
positifs d'emmagasina-ge de données;
- un enregistrement d'état global combinant
les enregistrements d'état du contrôleur et des
dispositifs d'emmagasina-ge de données ratta-
chés, qui est emmagasiné dans la table d'états
du tableau du contrôleur;

une pluralité de dispositifs d'emmagasina-
ge de données (70, 75) rattachés au contrôleur
de dispositifs d'emmagasina-ge de données,
chacun des dispositifs d'emmagasina-ge de
données comprenant une pluralité de pistes, et
une piste d'état de dispositif contenant les
informations d'état du dispositif d'emmagasina-
ge de données;

caractérisé en ce que:

ladite piste d'état de dispositif comprend
en outre un compte de reconfigurations de
l'enregistrement du nombre de fois qu'un dis-
positif a été ré-attribué, un numéro d'identifica-
tion de dispositif et un numéro d'identification
de contrôleur enregistrés dessus;

ladite table d'états dans la structure en
tableau comprend également un compte de
reconfigurations, un numéro d'identification de
dispositif, et un numéro d'identification de
contrôleur pour chacun des dispositifs d'em-
magasina-ge de données;

ledit enregistrement d'état global com-
prend également le numéro d'identification de
dispositif et le numéro d'identification de
contrôleur; et

le système comprend:

des moyens pour contrôler chacun de la
pluralité des dispositifs d'emmagasina-ge de
données afin de vérifier le numéro d'identifica-
tion de dispositif, le numéro d'identification de
contrôleur, et les enregistrements d'état;

des moyens pour écrire le numéro d'iden-
tification de dispositif sur chacun des enregis-
trements d'état de chacun des dispositifs
d'emmagasina-ge de données; et

des moyens pour écrire l'enregistrement
d'état global sur les enregistrements d'état
d'un premier dispositif d'emmagasina-ge de
données et d'un deuxième dispositif d'em-
magasina-ge de données.