

(12) **United States Patent**
Patterson et al.

(10) **Patent No.:** **US 6,668,264 B1**
 (45) **Date of Patent:** **Dec. 23, 2003**

(54) **RESYNCHRONIZATION OF A TARGET VOLUME WITH A SOURCE VOLUME**

(75) Inventors: **Hugo Patterson**, Mountain View, CA (US); **Michael Federwisch**, San Jose, CA (US)

(73) Assignee: **Network Appliance, Inc.**, Sunnyvale, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 333 days.

(21) Appl. No.: **09/825,855**

(22) Filed: **Apr. 3, 2001**

(51) **Int. Cl.**⁷ **G06F 17/30**

(52) **U.S. Cl.** **707/205; 707/204; 707/203**

(58) **Field of Search** **707/202, 204, 707/205, 10, 3, 203**

(56) **References Cited**

U.S. PATENT DOCUMENTS

| | | | |
|---------------|---------|------------------|--------|
| 4,742,450 A | 5/1988 | Duvall et al. | |
| 4,751,635 A * | 6/1988 | Kret | 707/10 |
| 4,875,159 A | 10/1989 | Cary et al. | |
| 4,887,204 A | 12/1989 | Johnson et al. | |
| 4,897,781 A | 1/1990 | Chang et al. | |
| 5,008,786 A | 4/1991 | Thatte | |
| 5,043,876 A | 8/1991 | Terry | |
| 5,208,813 A | 5/1993 | Stallmo | |
| 5,222,217 A | 6/1993 | Blount et al. | |
| 5,276,840 A | 1/1994 | Yu | |
| 5,305,326 A | 4/1994 | Solomon et al. | |
| 5,313,626 A | 5/1994 | Jones et al. | |
| 5,313,646 A | 5/1994 | Hendricks et al. | |
| 5,357,509 A | 10/1994 | Ohizumi | |
| 5,454,095 A | 9/1995 | Kraemer et al. | |
| 5,504,883 A | 4/1996 | Coverston et al. | |
| 5,604,862 A | 2/1997 | Midgely et al. | |
| 5,633,999 A | 5/1997 | Clowes et al. | |
| 5,649,152 A | 7/1997 | Ohran et al. | |
| 5,649,196 A | 7/1997 | Woodhill et al. | |
| 5,668,943 A | 9/1997 | Attanasio et al. | |

| | | | |
|----------------|---------|-------------------|--------|
| 5,721,916 A | 2/1998 | Pardikar | |
| 5,819,292 A | 10/1998 | Hitz et al. | |
| 5,819,310 A | 10/1998 | Vishlitzky et al. | |
| 5,835,953 A | 11/1998 | Ohran | |
| 5,907,672 A | 5/1999 | Matze et al. | |
| 5,996,086 A | 11/1999 | Delaney et al. | |
| 6,000,039 A | 12/1999 | Tanaka et al. | |
| 6,119,244 A | 9/2000 | Schoenthal et al. | |
| 6,131,088 A * | 10/2000 | Hill | 705/27 |
| 6,377,951 B1 * | 4/2002 | Campbell | 707/10 |

FOREIGN PATENT DOCUMENTS

| | | |
|----|----------------|--------|
| EP | 0 497 067 A1 | 8/1992 |
| WO | WO 00/07104 A1 | 2/2000 |

OTHER PUBLICATIONS

Gertrude G. Reusser and Donald C. Loughry, "Hewlett-Packard and the Open Systems Interconnection Reference Model", Hewlett-Packard Journal, Oct. 1986, vol. 37, No. 10.

Bruce Nelson and Yu-Ping-Cheng, "The Myth of Transfer Rate", How and Why SCSI Is better than IPI for NFS, Technical Report 6, Second Edition, Jul. 1992, Auspex.

* cited by examiner


Primary Examiner—Jean M. Corrielus

(74) *Attorney, Agent, or Firm*—Swernofsky Law Group, PC

(57) **ABSTRACT**

An improved method and apparatus for quickly and efficiently updating the original source volume and original target volumes after the original source volume has become temporarily unavailable. The original target volume is characterized as a source volume while the original source volume is temporarily unavailable. Transfer lists of different data blocks are generated. Data blocks not originally found on a source are copied to the target. Data blocks included on a target that were not found on the source are removed. By focusing upon specific data blocks, this technique avoids the use of filer overhead and other computational resources that would be expended if the entire volume were recopied.

28 Claims, 3 Drawing Sheets






FIG. 1

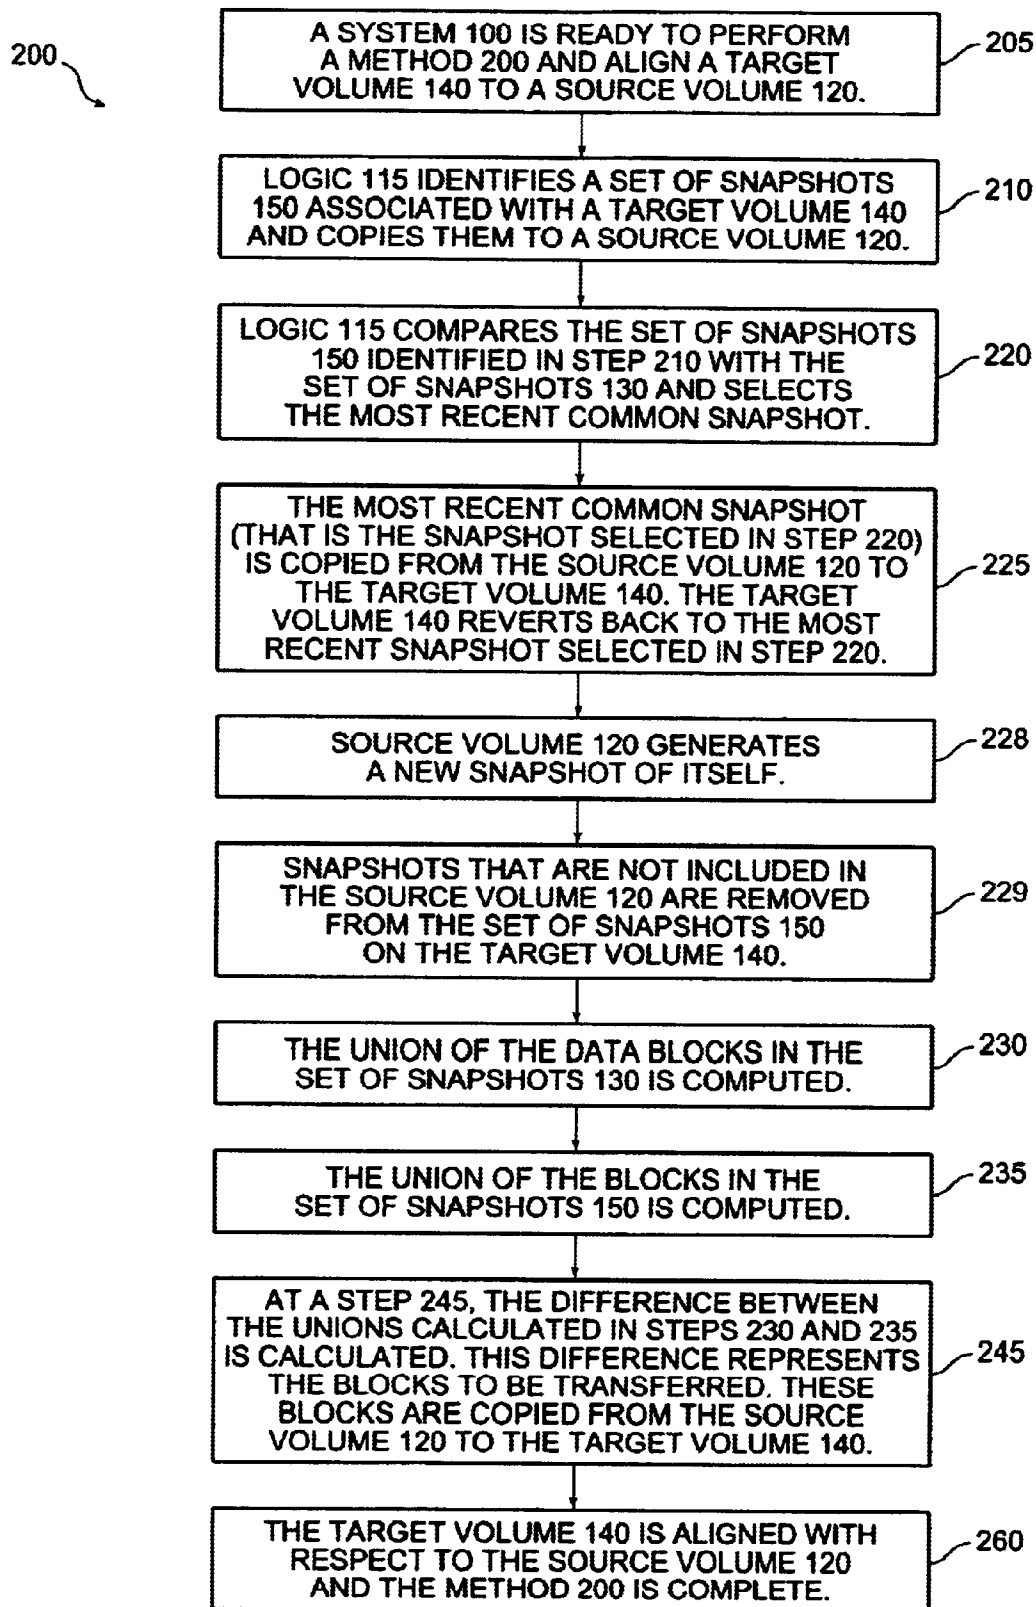


FIG. 2

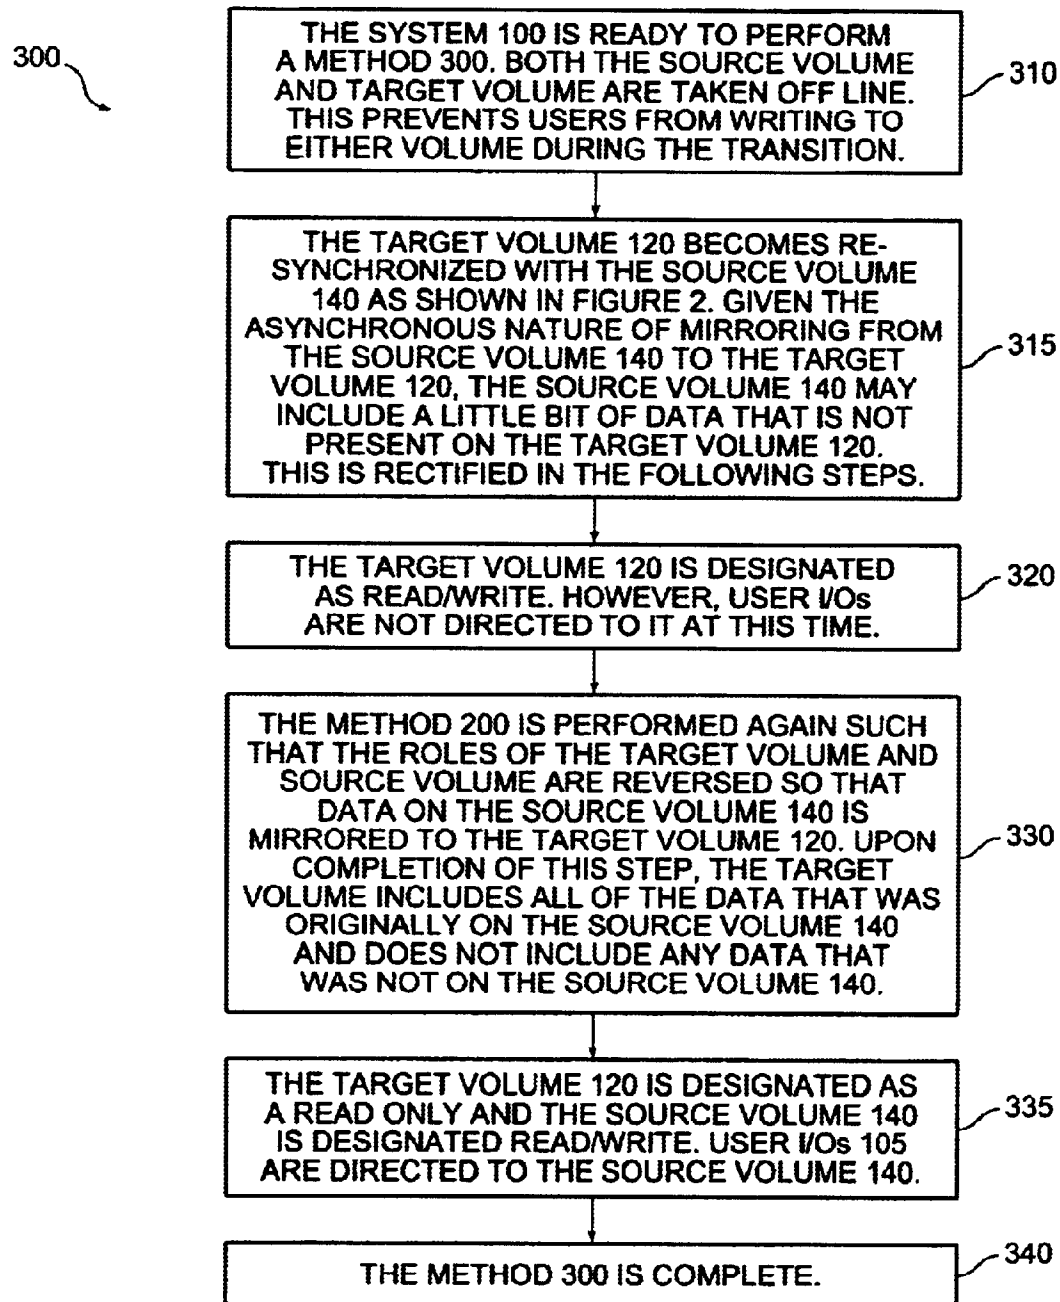


FIG. 3

RESYNCHRONIZATION OF A TARGET VOLUME WITH A SOURCE VOLUME

BACKGROUND OF THE INVENTION

1. Field of Invention

This invention relates to data storage systems. In particular, the invention relates to synchronization of source and target volumes in a mirrored storage system.

2. Related Art

Snapshots and multiple volumes are frequently used to prevent data loss when a data storage drive fails in a file system. Such snapshots “capture” the contents of the files and directories in a volume at a particular point in time in order to recover earlier versions of a file following an unintended deletion or modification. Such snapshots can also be copied to one or more volumes, which then can be used as a mirror or a collection of mirrors and which can provide a back-up copy of the file system. When used in this way, the mirror can be referred to as a target volume. In general, a target volume is a “read-only” volume that contains a set of data that is equivalent to the set of data on an original source volume. Such target volumes can be written to only by the original source volume.

A target volume may be updated periodically with respect to a source volume by looking to the most recent snapshot that the target and source have in common and using that snapshot as a consistency point (CP). The file blocks in the most recent common snapshot and the file blocks of a new snapshot are compared. The set of differences resulting from this comparison are written to the less up-to-date volume. In this way, both source and target volumes maintain equivalent sets of file blocks.

A source volume may become unavailable due to a failure of the source volume or to a failed connection to the source volume. Under such conditions, it is advantageous to temporarily use the target volume as a source volume by designating it as a “read/write” volume. User I/Os are directed to write to the target volume while the original source volume is unavailable.

One problem with writing to a target volume is that it may cause the target volume to contain data not found in the original source volume. A partial solution to this problem involves transferring data from the target to the source once the source is restored. However this is undesirable because it requires diversion of computational resources and filer overhead.

Accordingly, it would be advantageous to provide an improved technique for quickly and efficiently updating source and target volumes after a target volume has been written to. This is achieved in an embodiment of the invention that addresses the foregoing deficiencies.

SUMMARY OF THE INVENTION

The invention provides an improved method and apparatus for quickly and efficiently updating an original source volume and an original target volume after the original target volume has been used as a source volume in a file system. One or more snapshots are used to compare data included in the source and target volume. Instead of transferring the entire volume, only the data that is missing from a source and a target volume is transferred.

In a first aspect of the invention, a target volume becomes synchronized with a source volume after the target has been written to by an entity other than the source. Synchroniza-

tion is a two phase process. In the first phase, a target volume provides the source volume with a list of snapshots and associated snapshot numbers used to determine the sequence of the snapshots on the list. The source volume compares its own list of snapshot numbers with the list of the target volume's snapshot numbers and determines the newest common snapshot. This newest common snapshot is a consistency point between the target volume and the source volume. The source volume sends the target volume a set of snapshot numbers describing the newest common snapshot and the target volume reverts back to this snapshot. In the second phase, file system software identifies all the data blocks contained in any one or more of the snapshots of the source volume. This file system software also identifies all the data blocks in any one or more of the snapshots of the target volume using the data sent by the target volume to the source volume as described supra. A set of data blocks that are included in the source volume and not included in the target volume is generated. This can be accomplished by making a comparison based on logical differences, generating a virtual or actual list or other techniques known in the art.

File system software synchronizes the target volume with the source volume. First, the file system software removes snapshots from a target volume if the snapshots are not included in the source volume's snapshot list. Second, the file system software adds the set of data blocks identified above (that is the set of data blocks that are included in the source volume and not included in the target volume) to its memory. Lastly, the file system software adds snapshots to the target volume if the snapshots are included in the source volume's snapshot list and not in the target volume's snapshot list. At this point, the target volume includes the data blocks that are present on the source volume.

In a second aspect of the invention, the roles of the target volume and source volume are reversed and the process described supra is performed again so as to synchronize source volume with the target volume. This is necessary because the target volume may include data blocks not included in the source volume. After both source and target volumes are synchronized, the target volume stops being written to and the source once again is used as the active file. This is accomplished by 1) designating the target as a read-only volume, 2) designating the source as a read/write volume, and 3) redirecting users'I/O's back to the source volume.

In a preferred embodiment, sources and volumes can be synchronized dynamically, using a WAFL (Write Anywhere File Layout) system using RAID (Redundant Arrays of Independent Disks) architecture. However, various other types of file systems involving redundant copies of data can also be used.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a block diagram of a system for synchronizing a target volume to a source volume.

FIG. 2 shows a flow diagram of a method for synchronizing a target volume to a source volume.

FIG. 3 shows a flow diagram a method for synchronizing a target volume and a source volume to each other.

Lexicography

The following terms are related to aspects of the invention as described below. The general meanings of these terms are exemplary and in no way limiting.

Source volume—in general, the term “source volume” refers to a read/write volume.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.