(12) UK Patent (19) GB (11) 2510589 (13) B

(45) Date of B Publication 22.07.2020

(54) Title of the Invention: Data replication

(51) INT CL: ***G06F 11/20*** (2006.01)

(21) Application No: **1302203.3**

(22) Date of Filing: **07.02.2013**

(43) Date of A Publication **13.08.2014**

(56) Documents Cited:
US 6381677 B1 **US 20090043979 A1**

(58) Field of Search:
As for published application 2510589 A viz:
INT CL **G06F**
Other: **Online: WPI, EPODOC**
updated as appropriate

Additional Fields
Other: **None**

(72) Inventor(s):
John Fawcett
Adam Shepherd

(73) Proprietor(s):
Metaswitch Networks Ltd.
(Incorporated in the United Kingdom)
100 Church Street, ENFIELD, Middlesex, EN2 6BQ,
United Kingdom

(74) Agent and/or Address for Service:
EIP
Fairfax House, 15 Fulwood Place, LONDON,
WC1V 6HU, United Kingdom

GB 2510589 B

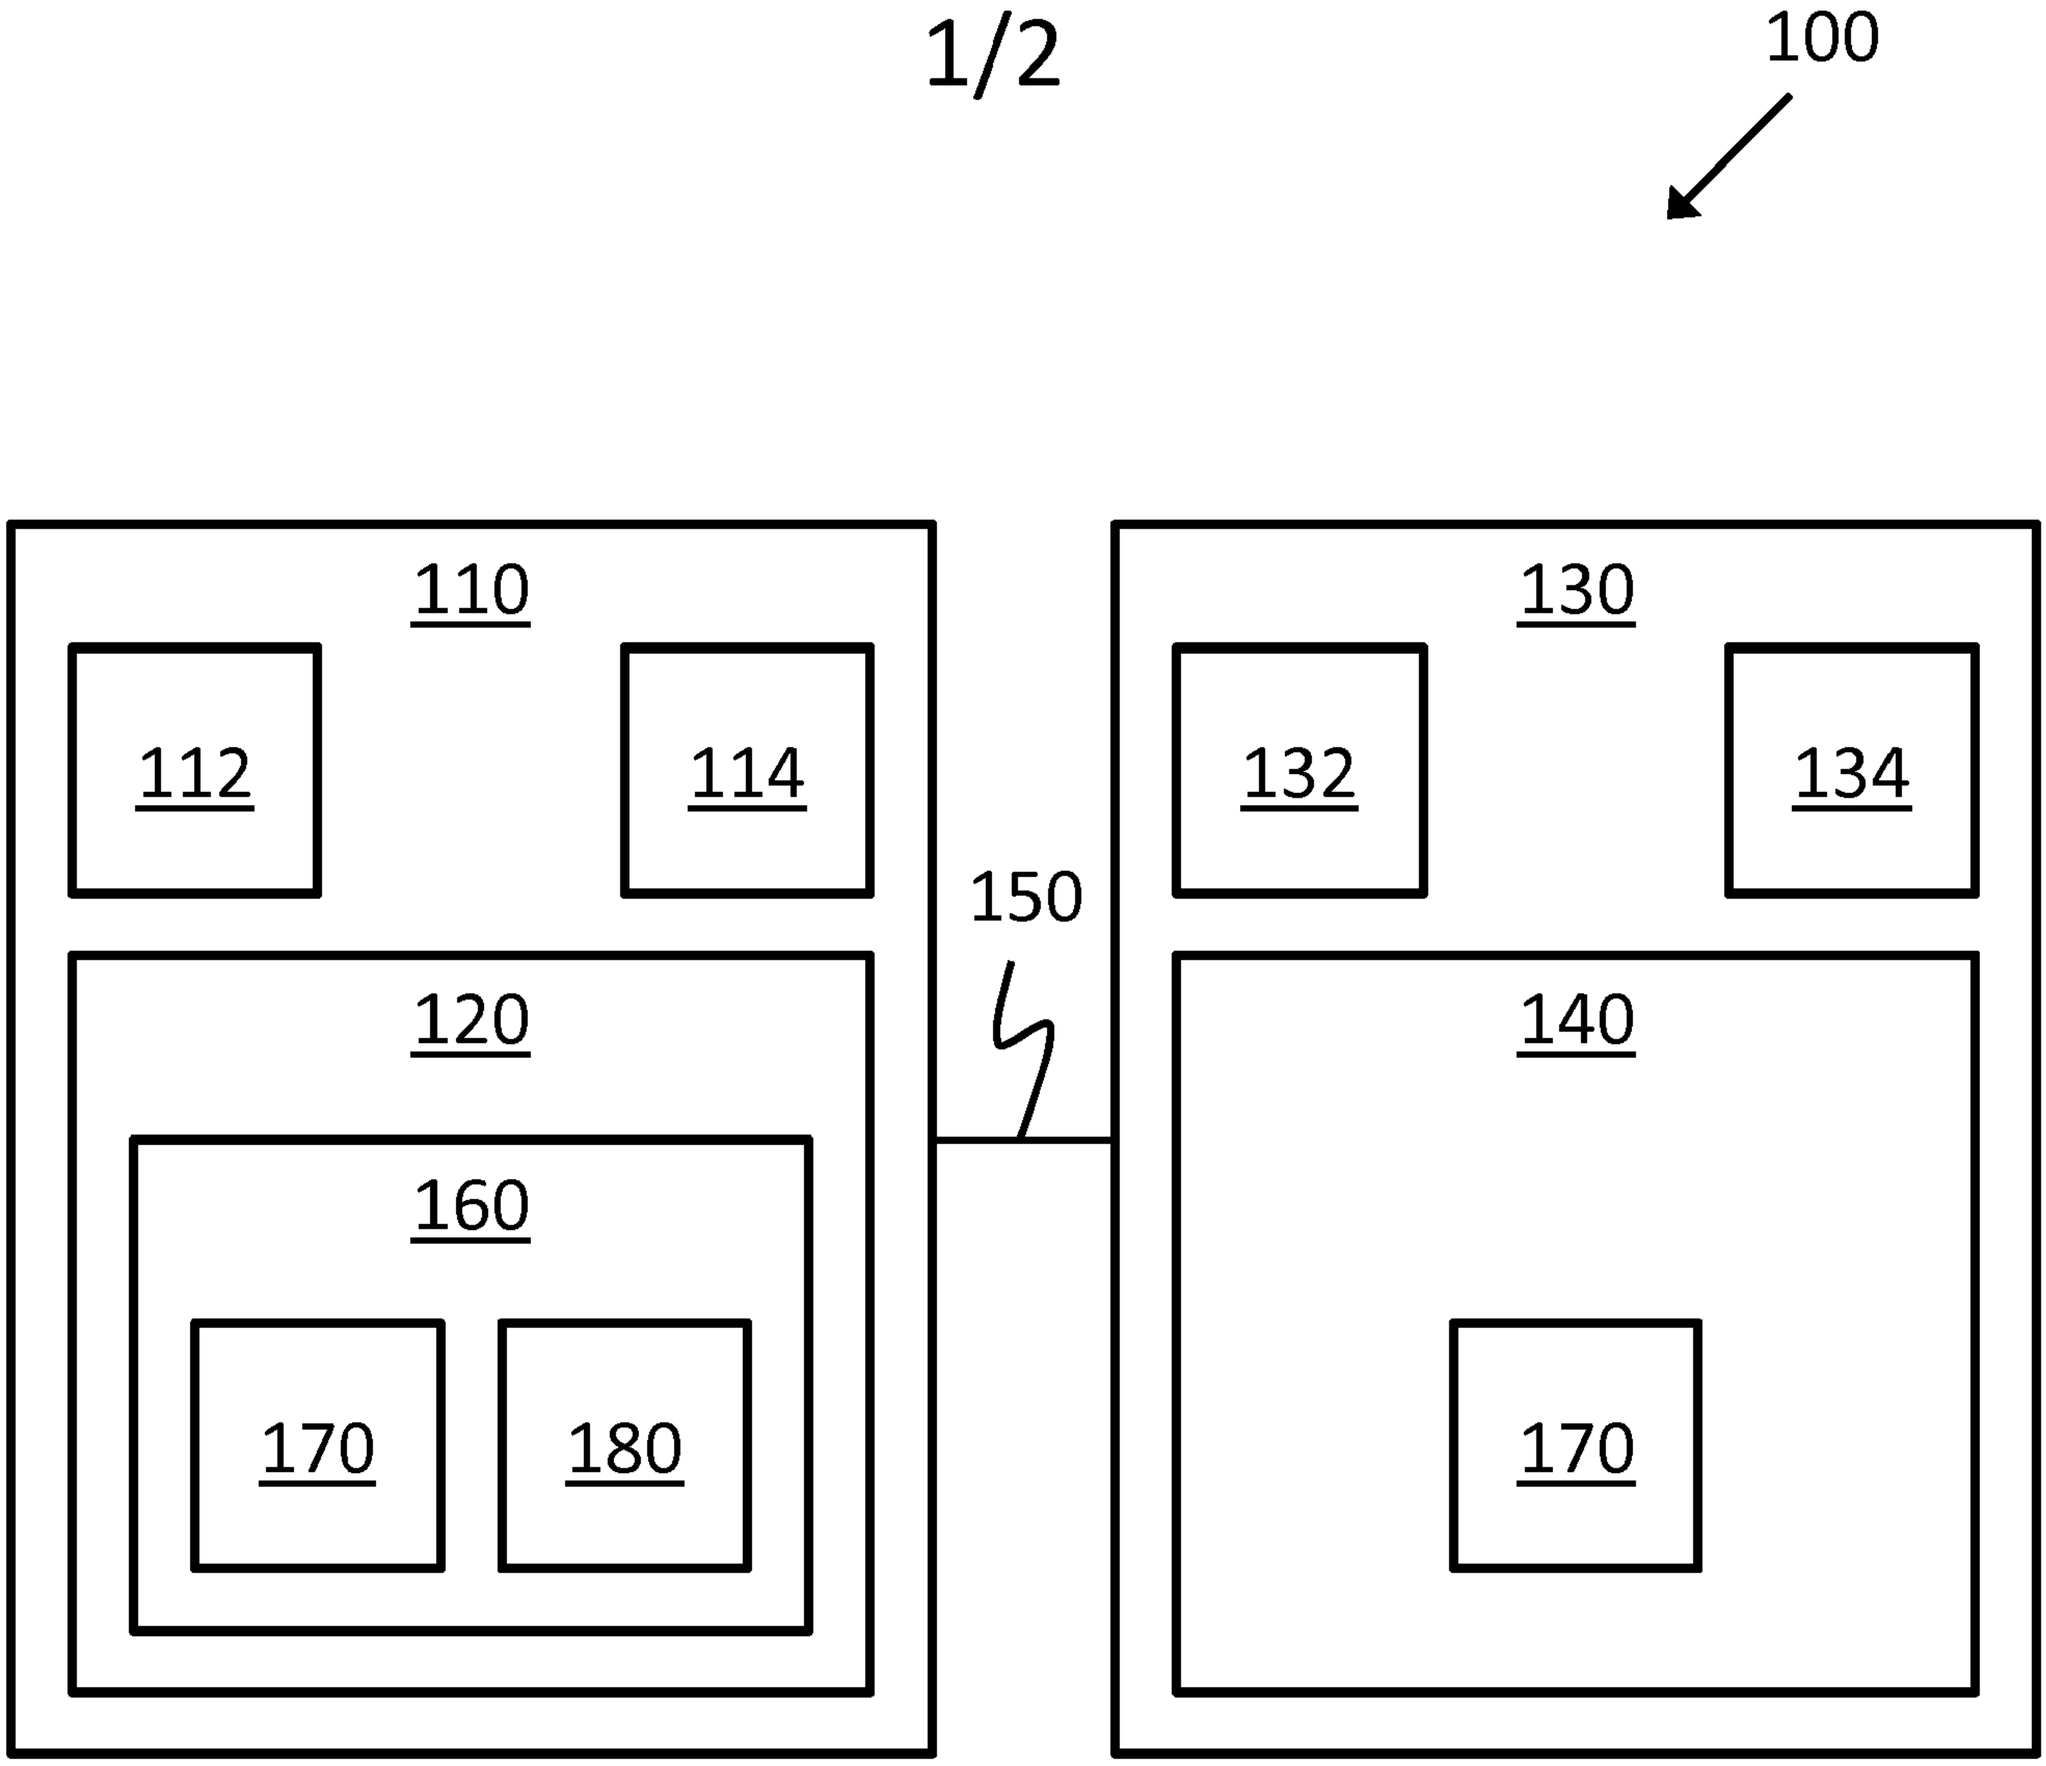

FIG. 1

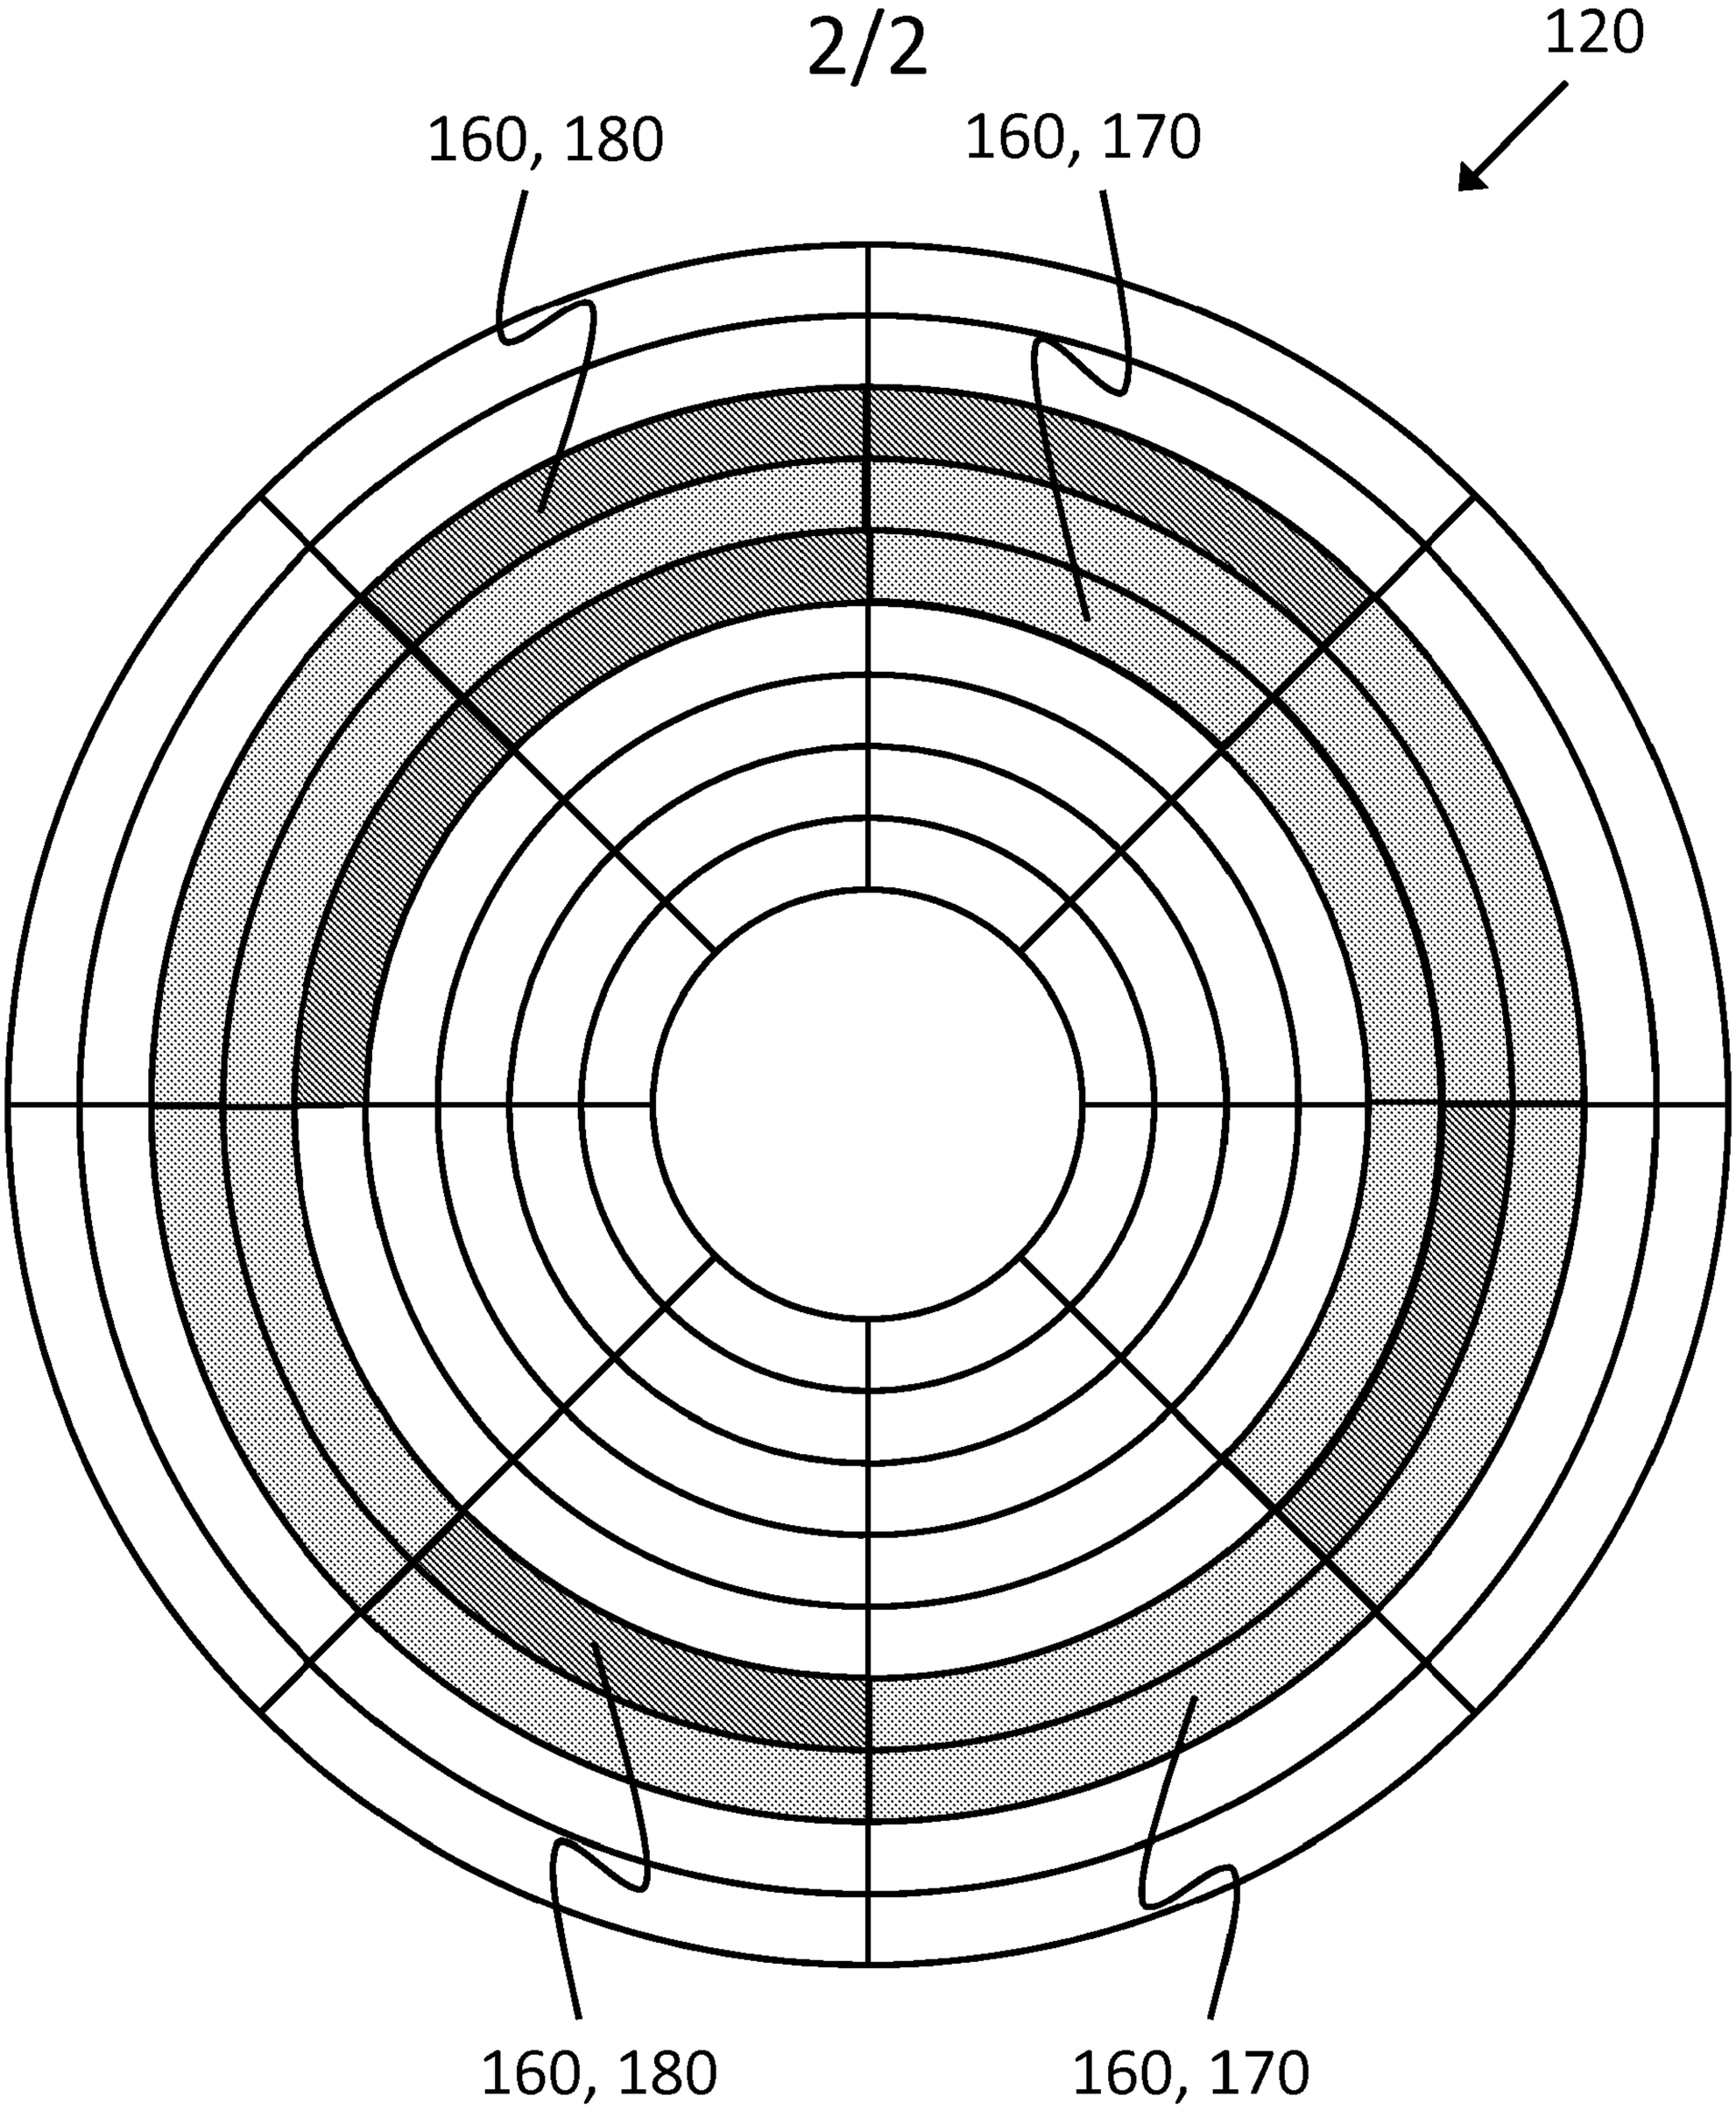

FIG. 2

The following terms are registered trade_marks and should be read as such wherever they occur in this document:

Microsoft

DATA REPLICATION

Technical Field

The present invention relates to data replication.

Background

It is known to deploy server-based applications in an active-standby configuration with continuous, real-time data replication from the active server to the standby server. If the standby server is unable to receive changes to the data set, for example if it is offline for maintenance or is disconnected from the active server as a result of a temporary network condition, the active server accumulates a change set of data to be synchronized to the standby server when the standby server reconnects to the active server; so-called "catch-up" replication.

However, when those changes in the change set are resynchronized with the standby server, the active server's disk storing the changed data has to read a scattered collection of disk blocks corresponding to the changed data; a workload demanding considerable track-to-track seeking. This significantly decreases input/output (I/O) performance (in other words, the number of bytes per second of data that can be written to or read from the disk) to the detriment of the active server's ability to handle live load for applications that are sensitive to disk I/O latency or jitter in disk access time.

For example, Microsoft (RTM) sync framework is a synchronisation platform which maintains complete logs of every local change that has not yet been successfully replicated. This increases storage requirements and disk I/O (both are required to store the additional information), and then causes significant disk head seeking which is detrimental to latency-sensitive apps.

An enterprise class magnetic disk may be able to deliver one hundred megabytes per second sustained transfer rate when dealing with files of several megabytes, dropping to around twenty megabytes per second or less transfer rate if the I/O requests are dominated by track-to-track seeking.

Almost all known redundant storage solutions avoid the problem of catch-up replication, rather than seeking to solve it. Avoiding the problem comes at a cost, however, in requiring specialist hardware. Battery back-up is one common requirement, together with a requirement that power be restored within a certain number of hours or data consistency is not guaranteed.

Although application code can be written to minimise disk seeking, out-of-band I/O workload caused by maintenance operations, backup, and in particular the resynchronization, may result in the worst case workload that limits application scalability.

Some existing data replication solutions attempt to throttle a resynchronization rate until the problem of disk I/O activity becomes acceptable. However, this prolongs the catch-up process considerably and may even be divergent if the replication process cannot keep up with the pace of change of data to be replicated.

Summary

According to a first aspect of the present invention, there is provided a method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

retrieving a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmitting at least part of the retrieved data set for replication at the at least one further data storage device.

As such, the at least some data that has not been replicated to the at least one further data storage device can be replicated to the at least one further data storage device while reducing the workload of the replication on apparatus performing the replication. Where the replication-performing apparatus is providing services that are

sensitive to the available workload, the impact of replication on such services may thereby be reduced.

In some embodiments, at least part of the data set is stored sequentially in the first data storage device and said retrieving the data set comprises sequentially retrieving the at least part of the data set from the first data storage device. It will be appreciated that references herein to at least part of the data set being stored sequentially do not imply the actual step of storing the at least part of the data set sequentially.

In some embodiments, the data set is stored sequentially in the first data storage device and said retrieving the data set comprises sequentially retrieving the data set from the first data storage device. It will be appreciated that references herein to the data set being stored sequentially do not imply the actual step of storing the data set sequentially.

Retrieval of the sequentially stored data set is associated with a lower workload than retrieval of non-sequentially stored data. The impact of replication on apparatus performing the replication may thereby be reduced.

In some embodiments, at least part of the at least some data that has not been replicated to the at least one further data storage device is stored non-sequentially in the first data storage device. It will be appreciated that references herein to at least part of the at least some data that has not been replicated to the at least one further data storage device being stored non-sequentially do not imply the actual step of storing the at least some data that has not been replicated to the at least one further data storage device.

In some embodiments, the at least some data that has not been replicated to the at least one further data storage device is stored non-sequentially in the first data storage device. It will be appreciated that references herein to the at least some data that has not been replicated to the at least one further data storage device being stored non-sequentially do not imply the actual step of storing the at least some data that has not been replicated to the at least one further data storage device.

Retrieval of the non-sequentially stored at least part of the at least some data that has not been replicated to the at least one further data storage device is associated

with a higher workload than retrieval of sequentially stored data. The data set, which is associated with a lower retrieval workload, may be retrieved such that the impact of replication on apparatus performing the replication may thereby be reduced.

Some embodiments comprise retrieving the data set from the first data storage device in response to occurrence of at least one trigger event.

In some embodiments, the at least one trigger event comprises at least one of:

detecting a failure associated with the at least one further data storage device; and

detecting at least partial degradation of connectivity associated with the at least one further data storage device.

Some embodiments comprise providing one or more workload-sensitive services while retrieving the data set from the first data storage device.

In some embodiments, the first data storage device is comprised in active server apparatus and the at least part of the retrieved data set is transmitted to standby server apparatus.

In some embodiments, the at least one further data storage device is comprised in standby server apparatus.

In some embodiments, the at least some data that has been replicated to the at least one further data storage device comprises data that has been transmitted to the least one further data storage device and for which replication of the data at the at least one further data storage device has been confirmed.

In some embodiments, the at least some data that has not been replicated to the at least one further data storage device comprises data that has been transmitted to the least one further data storage device but for which replication of the data at the at least one further data storage device has not been confirmed.

In some embodiments, the first data storage device comprises a rotating data storage medium.

According to a second aspect of the present invention, there is provided a method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

receiving at least part of a data set, the data set having been retrieved from the first data storage device and comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device having been less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

storing at least some of the received at least part of the retrieved data set in the at least one further data storage device.

Some embodiments comprise storing a previously received data set comprising at least the at least some data that has been replicated to the at least one further data storage device separately from the received at least part of the received data set in the at least one further data storage device.

Some embodiments comprise restoring the previously received data set in response to determining that not all of the retrieved data set has been received.

Some embodiments comprise discarding the previously received data set from the at least one further data storage device in response to determining that all of the retrieved data set has been received.

According to a third aspect of the present invention, there is provided apparatus for replicating data between a first data storage device and at least one further data storage device, the apparatus being arranged to:

retrieve a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmit at least part of the retrieved data set for replication at the at least one further data storage device.

In some embodiments, the apparatus comprises the first data storage device.

In some embodiments, the apparatus comprises or is comprised in active server apparatus and the at least part of the data set is received from active server apparatus.

According to a fourth aspect of the present invention, there is provided apparatus for replicating data between a first data storage device and at least one further data storage device, the apparatus being arranged to:

receive at least part of a data set, the data set having been retrieved from the first data storage device and comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device having been less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

store at least some of the received at least part of the retrieved data set in the at least one further data storage device.

In some embodiments, the apparatus comprises the at least one further data storage device.

In some embodiments, the apparatus comprises or is comprised in standby server apparatus.

According to a fifth aspect of the present invention, there is provided a computer program adapted to perform a method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

retrieving a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmitting at least part of the retrieved data set for replication at the at least one further data storage device.

According to a sixth aspect of the present invention, there is provided a computer program adapted to perform a method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

receiving at least part of a data set, the data set having been retrieved from the first data storage device and comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device having been less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

storing at least some of the received at least part of the retrieved data set in the at least one further data storage device.

According to a seventh aspect of the present invention, there is provided an active-standby server system for replicating data between a first data storage device and at least one further data storage device, the active-standby server system comprising at least:

active server apparatus arranged to:

retrieve a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmit at least part of the retrieved data set for replication at the at least one further data storage device; and

standby server apparatus arranged to:

receive at least part of the transmitted at least part of the retrieved data set; and

store at least some of the received at least part of the transmitted at least part of the retrieved data set in the at least one further data storage device.

Further features and advantages of the invention will become apparent from the following description of preferred embodiments of the invention, given by way of example only, which is made with reference to the accompanying drawings.

Brief Description of the Drawings

Figure 1 shows a schematic diagram of a data communications network in accordance with some embodiments.

Figure 2 shows a schematic diagram of a data storage device in accordance with some embodiments.

Detailed Description

Figure 1 shows a schematic diagram of a data communications network 100 in accordance with some embodiments.

The data communications network 100 includes a first server system 110 which comprises one or more servers which may be geographically co-located or may be remotely located from each other. In embodiments, the first server system 110 comprises at least one processor 112 and at least one memory 114, for example volatile memory such as Random Access Memory (RAM) which may include one or more computer programs. The at least one processor 112 may execute the one or more computer programs to cause the cause first server system 110 to perform, amongst other things, data replication as will be described below. In some embodiments, the first server system 110 comprises one or more data storage devices, for example a first data storage device 120. In some embodiments, the first data storage device 120 comprises non-volatile storage such as a hard disk drive.

The data communications network 100 also includes a second server system 130 which comprises one or more servers which may be geographically co-located or

may be remotely located from each other. In embodiments, the second server system 130 comprises at least one processor 132 and at least one memory 134, for example volatile memory such as RAM which may include one or more computer programs. The at least one processor 132 may execute the one or more computer programs to cause the cause second server system 130 to perform, amongst other things, data replication as will be described below. In some embodiments, the second server system 130 comprises one or more data storage devices, for example at least one further data storage device 140. In some embodiments, the at least one further data storage device 140 comprises non-volatile storage such as one or more hard disk drives.

The first and second server systems 110, 130 may be geographically co-located or may be remotely located from each other.

The first and second server systems 110, 130 communicate with each other via one or more connections 150. The first and second server systems 110, 130 may be directly connected to each other, for example if they are geographically co-located. Alternatively, the first and second server systems 110, 130 may be indirectly connected to each other via one or more intermediate entities or nodes, for example if they are remotely located and are connected via one or more data communication networks.

In some embodiments, the first server system 110 comprises one or more active (or 'primary') servers. Thus, in such embodiments, the first data storage device 120 may be comprised in active server apparatus. In some embodiments, the second server system 130 comprises one or more standby (or 'backup') servers. Thus, in such embodiments, the at least one further data storage device 140 may be comprised in standby server apparatus.

In some embodiments, the first data storage device 120 and/or the at least one further data storage device 140 comprises a rotating data storage medium, for example a hard disk drive (HDD), Compact Disc-ReWritable (CD-RW), Digital Versatile Disc-ReWritable (DVD-RW), laser disk, minidisk or the like.

Embodiments will now be described that relate to replication of data between the first data storage device 120 and the at least one further data storage device 140.

In these embodiments, the first data storage device 120 is a disk drive comprised in the first server system 110 and will therefore be referred to as a first disk drive 120. The first server system 110 has a role of an active server in an active-standby server system and will therefore be referred to as such in describing these embodiments.

Furthermore, in these embodiments, there is one further data storage device 140 which is also a disk drive and will therefore be referred to as the second disk drive 140. The second server system 130 has a role of a standby server in the active-standby server system and will therefore be referred to as such in describing these embodiments.

It will be appreciated, however, that other embodiments for replicating data between the first data storage device 120 and the at least one further data storage device 140 in accordance with embodiments described herein are envisaged.

Some embodiments described below for replicating data are applicable to various different types of system and software application. However, one such class of software application is I/O-bound, as opposed to storage-bound, software applications; in other words, software applications where the platform bottleneck is the number of disk operations that can be performed per second, rather than the total amount of data that can be stored on a disk. Disk storage is increasing rapidly and becoming cheaper per unit data storage, but I/O rates are remaining roughly the same.

Some embodiments described below indirectly increase the scalability of disk I/O-bound applications deployed on an active-standby hardware configuration, in particular but not exclusively to applications that are sensitive to disk I/O latency, or jitter in disk access time, when data replication occurs at the disk block level. By reducing the variance in disk access time, the impact of background data synchronisation on interactive application performance is also reduced.

Some embodiments described below use full synchronisations or replications instead of catch-up replications to minimise the effect of the replication on disk I/O-latency sensitive software applications.

Some embodiments described below also provide a safety net of a moment-in-time copy of a previous version of replicated data in case the active server 110 fails during a full resynchronization.

In some embodiments described below, the additional catch-up workload presented to the first disk drive 120 is sequential, rather than non-sequential, disk operations which can reduce a data retrieval workload associated with performing the catch-up.

In these embodiments, the active server 110 retrieves a data set 160 from the first disk drive 120.

In some embodiments, some or all of the data set 160 is stored sequentially in the first disk drive 120. Such embodiments therefore comprise sequentially retrieving the some or all of the data set 160 from the first disk drive 120. Retrieval of data that is stored sequentially is more readily tolerated by I/O-latency sensitive application software than retrieval of non-sequentially stored data.

Some embodiments comprise retrieving the data set 160 from the first disk drive 120 in response to occurrence of at least one trigger event. The at least one trigger event may comprise at least one of (a) detecting a failure associated with the second disk drive 140 and/or the standby server 130; and (b) detecting at least partial (for example partial or full) degradation of connectivity with the second disk drive 140 and/or the standby server 130.

The data set 160 comprises at least some data, generally referred to herein as "already replicated data" and denoted using reference 170, that has been replicated to the second disk drive 140. The already replicated data 170 may comprise data that has been transmitted to second disk drive 140 and for which replication of the data at the second disk drive 140 has been confirmed.

The data set 160 also comprises at least some data, generally referred to herein as "data to be replicated" and denoted using reference 180, that has not been replicated to the second disk drive 140.

There are various reasons why the data to be replicated 180 may not already have been replicated to the second disk drive 140. For example, the data to be replicated 180 may comprise data that has been transmitted to second disk drive 140

but for which replication of the data at the second disk drive 140 has not been confirmed.

In some embodiments, the data to be replicated 180 comprises a "change set" of data that comprises data that has changed since a previous successful replication between the active server 110 and the standby server 130.

In some embodiments, some or all of the data to be replicated is stored non-sequentially in the first disk drive 120.

The active server 110 transmits at least part of the retrieved data set 160 for replication at the second disk drive 140. For example, the active server 110 may transmit only a subset of the retrieved data set 160, for example only the data to be replicated 180, or the transmission of the data set 160 may be only partly successful, for example if the connection 150 between the active server 110 and the standby server 130 fails during transfer of the data set 160.

In some embodiments, the active server 110 replicates the entire data set 160, not just the data to be replicated 180 to the standby server 130.

As explained above, the workload for retrieving the data set 160 from the first disk drive 120 is less than the workload for retrieving the data to be replicated 180 from the first disk drive 120. As such, retrieval of the data set 160 from the first disk drive 120 is more readily tolerated by I/O-latency sensitive application software than retrieval of just the data to be replicated 180 to the standby server 130. The workload may be lower where, for example, I/O performance or data retrieval rates are improved such that, for example, resources required for controlling or processing the retrieval of data are reduced.

In some embodiments, the active server 110 handles a live load while retrieving the data set 160 from the first disk drive 120. Replication of the data set 160 in accordance with some embodiments described herein may take longer than would be the case if only the data set to be replicated 180 were replicated to the standby server 130. However, replication in accordance with embodiments described herein has a reduced effect on the ability of the active server 110 to handle a live load.

The standby server 130 receives at least part of the data set 160 transmitted by the active server 110. In other words, the standby server 130 may receive all or part of the data set 160 transmitted by the active server 110.

As explained above, the data set 160 has been retrieved from the first disk drive 120. The data set 160 comprises at least some data that has been replicated to the second disk drive 140; the already replicated data 170. The data set 160 also comprises at least some data that has not been replicated to the second disk drive 140; the data to be replicated 180. The workload for retrieving the data set 160 from the first disk drive 120 is less than the workload for retrieving the data to be replicated 180 from the first disk drive 120.

The standby server 130 stores at least some of the received at least part of the retrieved data set 160 in the second disk drive 140. In other words, the standby server 130 may store all or part of received at least part of the retrieved data set 160 in the second disk drive 140.

In some embodiments, the standby server 130 stores a previously received data set comprising at least the already replicated data 170 separately from the received at least part of the received data set in the second disk drive 140 as a back-up copy.

In some embodiments, the standby server 130 creates an empty disk partition and saves off the previously received data set as a moment-in-time copy of the previously received data set in the newly created disk partition. Although this may increase disk storage requirements, this may not be a significant concern where disk storage is relatively inexpensive and readily available.

If the active server 110 fails before the standby server 130 has fully caught up with (in other words has received all of the data set 160 from) the active server 110, the standby server 130 may discard the partly received data set 160 and restore the stored previously received data as the standby server 130 assumes the role of an active server system. The previously received data constitutes a valid filesystem, albeit a slightly outdated version.

In some embodiments, the standby server 130 restores the previously received data set in response to determining that not all of the retrieved data set has been received from the active server 130.

An alternative approach may be to use, for example, Logical Volume Manager (LVM) snapshots instead of saving off the data set 160 and to a brand new disk partition. This may reduce the storage overhead since LVM copy-on-write would only use additional space where data had been changed. However, the standby server 130 reads, compares and (if different) write twice, which LVM does not handle particularly well.

In some embodiments, the standby server 130 discards the stored previously received data set in response to determining that all of the retrieved data set has been received from the active server 110. The standby server 130 however continues to receive further data to be replicated, such as data that has changed since a previously successful replication, from the active server 110. This continued replication may be real-time replication. At this stage, normal replication operation is resumed.

During catch-up replication, data is transferred in an order that is most efficient for the I/O subsystem of the active server 110 to handle. As a result, the filesystem on the standby server 130 is not necessarily flowing through the same set of consistent states that the active server 110 experienced; rather it is not a valid filesystem at all until the synchronisation completes and all of the data set 140 has been transferred. As such, as described above, the standby server 130 saves a snapshot of the previously received data set in case the active server 110 fails and the synchronisation is unable to complete.

There are various mechanisms that may be used to replicate data between the active server 110 and the standby server 130. One such mechanism is Distributed Replicated Block Device (DRBD). DRBD is an open source project that allows active-standby and active-active Linux systems to replicate data in real time at the disk block level.

Figure 2 shows a schematic diagram of a data storage device in accordance with some embodiments.

As described above with reference to Figure 1, the first disk drive 120 stores a data set 160. The data set 160 comprises at least some already replicated data 170 and at least some data to be replicated 180.

As depicted in Figure 2, the data set 160 is stored sequentially in the first disk drive 120. At least some of the already replicated data 170, which is depicted in Figure 2 using light shading, is stored non-sequentially in the first disk drive 120. However, at least some of the already replicated data 170 is also stored sequentially in the first disk drive 120. At least some of the data to be replicated 180, which is depicted in Figure 2 using dark shading, is stored non-sequentially in the first disk drive 120. However, at least some of the data to be replicated 180 is also stored sequentially in the first disk drive 120.

As explained above, retrieval of data that is stored sequentially is more readily tolerated by I/O-latency sensitive application software than retrieval of non-sequentially stored data. As such, where the replication-performing apparatus is providing services that are sensitive to the available workload, the impact of replication on such services may thereby be reduced.

The above embodiments are to be understood as illustrative examples of the invention. Further embodiments of the invention are envisaged.

For example, although embodiments have been described above in which the first data storage device 120 is a local data storage device in the sense that it is local to the first server system 110, embodiments are envisaged in which the first data storage 120 is not a local data storage device. For example, it might be located remotely from the last server system 110.

Embodiments have been described above in which the first server system 110 comprises an active server and the second server system 130 comprises a standby server and wherein the active server replicates data to the standby server. Other embodiments are envisaged in which the first server system 110 comprises an active server and the second server system 130 also comprises an active server and wherein the active server in the first server system 110 replicates data to the active server in the second server system 130.

It is to be understood that any feature described in relation to any one embodiment may be used alone, or in combination with other features described, and may also be used in combination with one or more features of any other of the embodiments, or any combination of any other of the embodiments. Furthermore, equivalents and modifications not described above may also be employed without departing from the scope of the invention, which is defined in the accompanying claims.

Claims

1. A method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

retrieving a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmitting at least part of the retrieved data set for replication at the at least one further data storage device.

2. A method according to claim 1, wherein at least part of the data set is stored sequentially in the first data storage device and wherein said retrieving the data set comprises sequentially retrieving the at least part of the data set from the first data storage device.

3. A method according to claim 1 or 2, wherein the data set is stored sequentially in the first data storage device and wherein said retrieving the data set comprises sequentially retrieving the data set from the first data storage device.

4. A method according to any preceding claim, wherein at least part of the at least some data that has not been replicated to the at least one further data storage device is stored non-sequentially in the first data storage device.

5. A method according to any preceding claim, wherein the at least some data that has not been replicated to the at least one further data storage device is stored non-sequentially in the first data storage device.

6. A method according to any preceding claim, comprising retrieving the data set from the first data storage device in response to occurrence of at least one trigger event.

7. A method according to claim 6, wherein the at least one trigger event comprises at least one of:

detecting a failure associated with the at least one further data storage device; and

detecting at least partial degradation of connectivity associated with the at least one further data storage device.

8. A method according to any preceding claim, comprising providing one or more workload-sensitive services while retrieving the data set from the first data storage device.

9. A method according to any preceding claim, wherein the first data storage device is comprised in active server apparatus.

10. A method according to any preceding claim, wherein the at least one further data storage device is comprised in standby server apparatus.

11. A method according to any preceding claim, wherein the at least some data that has been replicated to the at least one further data storage device comprises data that has been transmitted to the least one further data storage device and for which replication of the data at the at least one further data storage device has been confirmed.

12. A method according to any preceding claim, wherein the at least some data that has not been replicated to the at least one further data storage device comprises data that has been transmitted to the least one further data storage device

but for which replication of the data at the at least one further data storage device has not been confirmed.

13. A method according to claim 1, wherein the first data storage device comprises a rotating data storage medium.

14. A method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

receiving at least part of a data set, the data set having been retrieved from the first data storage device and comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device having been less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

storing at least some of the received at least part of the retrieved data set in the at least one further data storage device.

15. A method according to claim 14, comprising storing a previously received data set comprising at least the at least some data that has been replicated to the at least one further data storage device separately from the received at least part of the received data set in the at least one further data storage device.

16. A method according to claim 15, comprising restoring the previously received data set in response to determining that not all of the retrieved data set has been received.

17. A method according to claim 15 or 16, comprising discarding the previously received data set from the at least one further data storage device in response to determining that all of the retrieved data set has been received.

18. Apparatus for replicating data between a first data storage device and at least one further data storage device, the apparatus being arranged to:

retrieve a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmit at least part of the retrieved data set for replication at the at least one further data storage device.

19. Apparatus according to claim 18, comprising the first data storage device.

20. Apparatus according to claim 18 or 19, wherein the apparatus comprises or is comprised in active server apparatus and said at least part of the retrieved data set is transmitted to standby server apparatus.

21. Apparatus for replicating data between a first data storage device and at least one further data storage device, the apparatus being arranged to:

receive at least part of a data set, the data set having been retrieved from the first data storage device and comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device having been less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

store at least some of the received at least part of the retrieved data set in the at least one further data storage device.

22. Apparatus according to claim 21, comprising the at least one further data storage device.

23. Apparatus according to claim 21 or 22, wherein the apparatus comprises or is comprised in standby server apparatus and said at least part of the data set is received from active server apparatus.

24. A computer program adapted to perform a method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

retrieving a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmitting at least part of the retrieved data set for replication at the at least one further data storage device.

25. A computer program adapted to perform a method of replicating data between a first data storage device and at least one further data storage device, the method comprising:

receiving at least part of a data set, the data set having been retrieved from the first data storage device and comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device having been less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

storing at least some of the received at least part of the retrieved data set in the at least one further data storage device.

26. An active-standby server system for replicating data between a first data storage device and at least one further data storage device, the active-standby server system comprising at least:

active server apparatus arranged to:

retrieve a data set from the first data storage device, the data set comprising at least some data that has been replicated to the at least one further data storage device and at least some data that has not been replicated to the at least one further data storage device, a workload for retrieving the data set from the first data storage device being less than a workload for retrieving the at least some data that has not been replicated to the at least one further data storage device from the first data storage device; and

transmit at least part of the retrieved data set for replication at the at least one further data storage device; and

standby server apparatus arranged to:

receive at least part of the transmitted at least part of the retrieved data set; and

store at least some of the received at least part of the transmitted at least part of the retrieved data set in the at least one further data storage device.