

METHODS, SYSTEMS AND COMPUTER PROGRAM PRODUCTS FOR COORDINATED DISASTER RECOVERY

Field of the Invention

The present invention relates in general to computers, and more particularly to coordinated disaster recovery production takeover operations.

Background of the invention

In today's society, computer systems are commonplace. Computer systems may be found in the workplace, at home, or at school. Computer systems may include data storage systems, or disk storage systems, to process and store data. Data storage systems, or disk storage systems, are utilized to process and store data. A storage system may include one or more disk drives. These data processing systems typically require a large amount of data storage. Customer data, or data generated by users within the data processing system, occupies a great portion of this data storage. Many of these computer systems include virtual storage components.

Virtual storage components are found in a variety of computing environments. A typical virtual storage component is the magnetic tape cartridge used via a magnetic tape drive. Multiple tape drives may be contained in a tape library, along with several slots to hold tape cartridges. Such data storage systems utilize storage components (usually direct access storage, such as disk arrays) to virtually present tape libraries or tape drives. Both types of technologies are commonly used for backup and recovery purposes. Virtual tape libraries, which integrate with existing backup software and existing backup and recovery processes, enable typically faster backup and recovery operations. It is often required that such data storage entities be replicated from their origin site to remote sites. Replicated data systems may externalize various logical data storage entities, such as files, data objects, backup images, data snapshots or virtual tape cartridges.

Replicated data entities enhance fault tolerance abilities and availability of data. Thus, it is critical to create disaster recovery (DR) plans for these massive computer systems, particularly in today's global economy. DR plans are required by variable sized companies and by governments in most of the western world. Most modern standards denote a 3-4 sites (many-to-many) topology group for replicating data between the storage systems in order to maintain 3 to 4 copies of the data in the storage systems.

Summary of the Described Embodiments

As previously mentioned, modern standards typically denote a 3-4 sites (many-to-many) topology group for replicating data between the storage systems in order to maintain three to four copies of the data in the storage systems. Within the many-to-many topologies, challenges arise in assuring takeover processes, which are apart of the disaster recovery (DR) plan, avoid creating situations that reduce productivity and efficiencies. Such challenges include preventing possible data corruption scenarios, particularly when involving synchronization processes between multiple interlaced systems, and/or situations where users end up with wrong cartridges at a particular production site. Such inefficiencies reduce performance and may compromise the integrity of maintaining copies of data within a storage system.

Accordingly, and in view of the foregoing, various exemplary method, system, and computer program product embodiments for coordinated disaster recovery are provided. In one embodiment, by way of example only, a reconciliation process is performed for resolving intersecting and non-intersecting data amongst disaster recovery systems for takeover operations. An ownership synchronization process is coordinated for replica cartridges via the reconciliation process at the disaster recovery systems. The disaster recovery systems continue as a replication target for source systems and as a backup target for local backup applications. Additional system and computer program product embodiments are disclosed and provide related advantages.

In addition to the foregoing exemplary method embodiment, other exemplary system and computer product embodiments are provided and supply related advantages. The foregoing

summary has been provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. The claimed
5 subject matter is not limited to implementations that solve any or all disadvantages noted in the background.

Brief Description of the Drawings

10 In order that the advantages of the invention will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments that are illustrated in the appended drawings. Understanding that these drawings depict embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional
15 specificity and detail through the use of the accompanying drawings, in which:

Fig. 1 illustrates an exemplary computing environment in which aspects of the present invention may be implemented;

20 Fig. 2 illustrates an exemplary computing device including a processor device in a computing environment in which aspects of the present invention may be implemented;

Fig. 3 is a flowchart illustrating an exemplary method for coordinating disaster recovery production takeover operations in many-to-many topology;

25 Fig. 4 is a flowchart illustrating an exemplary method for announcing a system offline;

Fig. 5 is a flowchart illustrating an exemplary method for coordinating an ownership synchronization process for replica cartridges via a reconciliation process;

30 Fig. 6 is a flowchart illustrating an exemplary method for performing a reconciliation process amongst disaster recovery systems for a takeover operation;

Fig. 7 illustrates an exemplary block diagram of the types of mutuality between source data sets distributed to different remote systems;

Fig. 8A illustrates an exemplary block diagram of many-to-many system (four systems) for replication with system #3 being a source system and replicating to all other remote systems;

Fig. 8B illustrates an exemplary block diagram of the remote system before a disaster recovery takeover with the source system #3 no longer available;

Fig. 8C illustrates an exemplary block diagram demonstrating the takeover operation performed by the first disaster recovery system #1 and consulting disaster recovery system #2;

Fig. 8D illustrates an exemplary block diagram demonstrating the takeover operation performed by the second disaster recovery system #2;

Fig. 8E illustrates an exemplary block diagram demonstrating the takeover operation performed by the second disaster recovery system #4; and

Fig. 8F illustrates an exemplary block diagram demonstrating each of the disaster recovery systems exiting the disaster recovery mode and continuing to work as normal.

Detailed Description of the Drawings

Throughout the following description and claimed subject matter, the following terminology, pertaining to the illustrated embodiments, is described.

A “cartridge ownership ” is intended to refer to an attribute of a cartridge indicating the cartridge’s ability to be written at a certain system. A cartridge may be write-enabled on its owner system. A “disaster recovery (DR) mode ” is intended to refer to an indication at a remote system that a certain remote system is now used as DR for a certain source system. The DR mode may cause replication communication from the source system to be blocked

in order to protect replicated data. A “replication” is intended to refer to a process of incrementally copying deduplicated data between systems, which reside in the same replication grid. A “replication grid” is intended to refer to a logical group, which provides context in which replication operation may be established between different physically connected members. A “replication grid manager” is intended to refer to a component (such as a software component operated by a processor device) in charge of replication and changing ownership activity in a grid’s context. A “VTL” or “virtual tape library” is intended to refer to a virtual tape library – computer software emulating a physical library. A “cartridge” may include the term data storage entity, data storage entities, replicated data storage entity, replicated data storage entities, files, data objects, backup images, data snapshots, virtual tape cartridges, and other known art commonly known in the industry as a cartridge in a computer environment. Also, a source system site may refer to a first storage system, first storage site, and primary storage system. A remote system site may be referred to as a *secondary storage site*, a *secondary storage system*, and a *remote storage system*. Also, a remote system site may also be referred to as a disaster recovery system when the remote system is operating in disaster recovery mode.

The many-to-many topology may create problems for one-to-one and many-to-one topologies. When different data sets or multiple intersecting data sets are being replicated from a source site to different destinations, a normal disaster recovery process should recover from multiple sites, and in case of intersection, should be recovered only on one of the destinations (the one that has its backup environment production ownership). A disaster recovery solution should prevent a shutdown of the DR system for a number of source systems that may be in the midst of replication and prevent potential data loss/corruption and/or prolonged RPO (Recovery Point Objective). The current state of the art fails to address these issues thereby reducing performance and efficiency may be reduced.

In contrast, and to address the inefficiencies and performance issues previously described, the mechanisms of the illustrated embodiments serve to coordinate disaster recovery production takeover processes in a many-to-many topology in a more effective manner, for example, in a many-to-many topology for deduplication virtual tape library (VTL) systems. Within the many-to-many topologies, multiple systems may act as a disaster recovery (DR)

system and move to a DR mode. The production environment may also be moved to the proper DR systems' sites. The temporary production sites may create new cartridges and/or write on old cartridges while still being a target for multiple other source systems. In order to allow production to move permanently to the DR sites (because the production site is permanently declared terminated and no replacement site is planned), coordinated ownership synchronization processes may occur within a replication grid at the DR sites so that ownership over source system cartridges may be changed to the DR sites (new production sites). The entire coordination process may occur while concurrently receiving replication data from other source systems.

In an alternative embodiment, the mechanisms are configured for performing a reconciliation process for resolving intersecting and non-intersecting data amid multiple disaster recovery systems for a takeover operation. The ownership synchronization process for replica cartridges are coordinated via the reconciliation process at several disaster recovery systems. The disaster recovery systems continue to be a replication target for multiple source systems (that may not be offline) and a backup target for local backup applications.

Turning now to Fig. 1, exemplary architecture 10 of data storage systems (e.g., virtual tape systems) in a computing environment is depicted. Architecture 10 provides storage services to local hosts 18 and 20, and replicate data to remote data storage systems as shown. A local storage system server 12 in communication with a storage device 14 is connected to local hosts 18 and 20 over a network including components such as Fibre channel switch 16. Fibre channel switch 16 is capable, for example, of executing commands (such as small computer systems interface (SCSI) commands) for tape devices. The skilled artisan will appreciate that architecture 10 may include a variety of storage components. For example, storage devices 14 may include conventional hard disk drive (HDD) devices, or may include solid state drive (SSD) devices.

Local storage system server 12 is connected over network 22 to a remote storage system server 24. Remote server 24 communicates with a locally connected disk storage device 26, and with locally connected hosts 30 and 32 via another network and network component 28

such as Fibre channel switch 28. Network 22 may include a variety of network topologies, such as a wide area network (WAN), a local area network (LAN), a storage area network (SAN), and other configurations. Similarly, switches 16 and 28 may include other types of network devices.

5

Architecture 10, as previously described, provides local storage services to local hosts, and provides replicate data to the remote data storage systems (as denoted by data replication functionality using arrow 34). As will be described, various embodiments of the present invention and claimed subject matter may be implemented on architectures such as architecture 10.

10

Fig. 2 illustrates a portion 200 of an exemplary computer environment that can be used to implement embodiments of the present invention. A computer 202 comprises a processor 204 and a memory 206, such as random access memory (RAM). In one embodiment, storage system server 12 (Fig. 1) may include components similar to those shown in computer 202. The computer 202 is operatively coupled to a display 219, which presents images such as windows to the user on a graphical user interface 218. The computer 202 may be coupled to other devices, such as a keyboard 216, a mouse device 220, a printer 228, etc. Of course, those skilled in the art will recognize that any combination of the above components, or any number of different components, peripherals, and other devices, may be used with the computer 202.

15

20

25

30

Generally, the computer 202 operates under control of an operating system (OS) 208 (e.g. z/OS, OS/2, LINUX, UNIX, WINDOWS, MAC OS) stored in the memory 206, and interfaces with the user to accept inputs and commands and to present results, for example through a graphical user interface (GUI) module 232. In one embodiment of the present invention, the OS 208 facilitates the backup mechanisms. Although the GUI module 232 is depicted as a separate module, the instructions performing the GUI functions can be resident or distributed in the operating system 208, the application program 210, or implemented with special purpose memory and processors. OS 208 includes a replication module 240 and disaster recovery module 242 which may be adapted for carrying out various processes and mechanisms in the exemplary embodiments described below, such as performing the

coordinated disaster recovery production takeover operation functionality. The replication module 240 and disaster recovery module 242 may be implemented in hardware, firmware, or a combination of hardware and firmware. In one embodiment, replication module 240 may also be considered a replication grid manager or replication manager for performing and/or managing the replication and change ownership activity in a replication grid's context as further described. Moreover, the replication module 242 may perform all of the replication type events and/or processes needed to execute the mechanisms of the illustrated embodiments while simultaneously performing and functioning as a replication grid manager. In one embodiment, the replication module 240 and disaster recovery module 242 may be embodied as an application specific integrated circuit (ASIC). As the skilled artisan will appreciate, functionality associated with the replication module 240 and disaster recovery module 242 may also be embodied, along with the functionality associated with the processor 204, memory 206, and other components of computer 202, in a specialized ASIC known as a system on chip (SoC). Further, the functionality associated with the replication module and disaster recovery module 242 (or again, other components of the computer 202) may be implemented as a field programmable gate array (FPGA).

As depicted in Fig. 2, the computer 202 includes a compiler 212 that allows an application program 210 written in a programming language such as COBOL, PL/1, C, C++, JAVA, ADA, BASIC, VISUAL BASIC or any other programming language to be translated into code that is readable by the processor 204. After completion, the computer program 210 accesses and manipulates data stored in the memory 206 of the computer 202 using the relationships and logic that was generated using the compiler 212. The computer 202 also optionally comprises an external data communication device 230 such as a modem, satellite link, Ethernet card, wireless link or other device for communicating with other computers, e.g. via the Internet or other network.

Data storage device 222 is a direct access storage device (DASD) 222, including one or more primary volumes holding a number of datasets. DASD 222 may include a number of storage media, such as hard disk drives (HDDs), solid-state devices (SSD), tapes, and the like. Data storage device 236 may also include a number of storage media in similar fashion to device 222. The device 236 may be designated as a backup device 236 for holding

5 backup versions of the number of datasets primarily stored on the device 222. As the skilled artisan will appreciate, devices 222 and 236 need not be located on the same machine. Devices 222 may be located in geographically different regions, and connected by a network link such as Ethernet. Devices 222 and 236 may include one or more volumes, with a
5 corresponding volume table of contents (VTOC) for each volume.

In one embodiment, instructions implementing the operating system 208, the computer
program 210, and the compiler 212 are tangibly embodied in a computer-readable medium,
e.g., data storage device 220, which may include one or more fixed or removable data
10 storage devices 224, such as a zip drive, floppy disk, hard drive, DVD/CD-ROM, digital
tape, flash memory card, solid state drive, etc., which are generically represented as the
storage device 224. Further, the operating system 208 and the computer program 210
comprise instructions which, when read and executed by the computer 202, cause the
computer 202 to perform the steps necessary to implement and/or use the present invention.
15 For example, the computer program 210 may comprise instructions for implementing the
grid set manager, grid manager and repository manager previously described. Computer
program 210 and/or operating system 208 instructions may also be tangibly embodied in the
memory 206 and/or transmitted through or accessed by the data communication device 230.
As such, the terms "article of manufacture," "program storage device" and "computer
20 program product" as may be used herein are intended to encompass a computer program
accessible and/or operable from any computer readable device or media.

Embodiments of the present invention may include one or more associated software
application programs 210 that include, for example, functions for managing a distributed
25 computer system comprising a network of computing devices, such as a storage area
network (SAN). Accordingly, processor 204 may comprise a storage management processor
(SMP). The program 210 may operate within a single computer 202 or as part of a
distributed computer system comprising a network of computing devices. The network may
encompass one or more computers connected via a local area network and/or Internet
30 connection (which may be public or secure, e.g. through a virtual private network (VPN)
connection), or via a fibre channel SAN or other known network types as will be understood

by those skilled in the art. (Note that a fibre channel SAN is typically used only for computers to communicate with storage systems, and not with each other.)

As previously mentioned, the mechanisms of the present invention provide for coordinating replica cartridges' ownership synchronization process at remote systems while they are in a disaster recovery (DR) mode and while still being replication targets for other source systems and backup targets for local backup applications. The remote systems that are declared to be in the DR mode may become part of a disaster recovery system(s). The declaration of going into DR mode may be performed by the remote systems' administrators within their own systems and may be specific for the system that has gone down. The outcome of a DR mode may be complete blockage of all replication communication from a specific source system, such as the source system that is offline or gone down and is no longer available. In order to exit the DR mode the user may choose to run a takeover operation to synchronize ownership over the source system cartridges in coordination with other possible destinations (e.g., various remote systems or other source systems) of the source system.

As will be described below, the mechanisms of the present invention seek to provide the ability of an inherent and coordinated synchronization process for a virtual tape (VT) system in order to restore a replication group state to its original state prior to a disaster. Thus, the mechanisms allow for seamless production site switching to a number of disaster recovery (DR) sites, which include a replica baseline. Also, synchronization processes for the replication and coordination may work in parallel to normal replication in order to provide a DR capability to single or multiple sets of source systems while allowing the remaining source systems to replicate as normal.

Fig. 3 is a flowchart illustrating an exemplary method 300 for coordinating disaster recovery production takeover operations in many-to-many topology within a computing environment. The method 300 begins (step 302) by performing a reconciliation process for resolving intersecting and non-intersecting data amid multiple disaster recovery systems for a takeover operation(s) (step 304). The ownership synchronization process for replica cartridges are coordinated via the reconciliation process at several disaster recovery systems (step 306).

The disaster recovery systems continue to be a replication target for multiple source systems and a backup target for local backup applications (step 308). The method 300 ends (step 310).

5 In one embodiment, the mechanisms may announce a source system offline. The user decides to announce his source system offline in order to allow the DR systems to takeover the offline source systems data/cartridges. The source system that was selected to go offline may be checked to have already left the replication grid prior to the takeover operation. The announcement of the source system going offline and/or leaving the replication grid may be
10 distributed among all the replication grid systems.

Fig. 4 is a flowchart illustrating an exemplary method 400 for announcing a system offline. The method 400 commences (step 402) by declaring a source system offline (step 404). Allow disaster recovery systems to perform the takeover operation (step 406). A replication
15 grid is checked (step 408). The method 400 determines if the offline source system has exited the replication grid, (step 410). If no, then the method 400 ends (step 414). If yes, then the method 400 will notify all of the replication grid systems that the source system is offline (step 412). The method 400 ends (step 414).

20 Fig. 5 is a flowchart illustrating an exemplary method 500 for coordinating an ownership synchronization process for replica cartridges via the reconciliation process. The method 500 begins (step 502) and determines if non-intersecting datasets are distributed among disaster recovery systems (step 504). If yes, the method 500 will perform the takeover operation separately on each of the disaster recovery systems (step 506). If no, the method
25 500 will determine if overlapping datasets are distributed among the disaster recovery systems (step 508). If yes, the method 500 will execute the takeover operation first, by one of the disaster recovery systems, to change the ownership of each of the cartridges (step 509). If no, the method 500 will determine if intersecting datasets are distributed among the disaster recovery systems (step 510). If no, the method 500 will end (step 522). If yes, the
30 method 500 will determine the ownership of cartridges based on the order of performing the takeover operation by the plurality of disaster recovery systems (step 512). For determining ownership of the cartridges based on the order of performing the takeover operation by the

disaster recovery systems, the method 500 will determine if the disaster recovery systems is first to perform the takeover operation (step 514). If yes, the method 500 will acquire the ownership of each of the cartridges that intersect (step 516). If no, the method 500 will determine if the disaster recovery system(s) is a subsequent disaster recovery system(s) to perform the takeover operation (step 518). If no, the method 500 will end (step 522). If yes, the method 500 will acquire the ownership of the intersecting cartridges intersecting between the subsequent performing disaster recovery systems that is performing the takeover operation (meaning itself) and the disaster recovery systems yet to have performed the takeover operation (step 520). For example, there may be four disaster recovery systems in a grid so the method 500 may perform the takeover operation on the first disaster recovery system, as mentioned above, and then perform the takeover operations for the subsequent disaster recovery systems. The takeover operations may be iteratively performed for the first, second, third, and fourth disaster recovery system, depending on which datasets are intersecting. The method 500 will check and determine if there are additional intersecting datasets existing between the remaining disaster recovery systems (step 521) (this algorithm may converge to the disjointed form). If yes, the method 500 will return and determine the ownership of cartridges based on the order of performing the takeover operation by the plurality of disaster recovery systems (step 512) and repeat the subsequent steps, as mentioned above. If no, the method 500 ends (step 522).

Fig. 6 is a flowchart illustrating an exemplary method 600 for a reconciliation process amongst disaster recovery systems for a takeover operation. The method 600 begins (step 602) by activating a disaster recovery (DR) mode in a remote system of the source system (step 604). The disaster recovery mode may be initiated automatically by a failure that occurs at the source system thereby rendering the source system offline (unavailable) and/or by declaring the source system offline (unavailable) by an administrator's preference/choice. The disaster recovery systems may be allowed to sequentially perform the takeover operation (step 606). Each disaster recover system may take a turn to perform the takeover operation. The method 600 will determine if the takeover operation may be performed for an offline source system (step 608). If no, the method 600 ends (step 622). If yes, the method 600 will send a request within a replication grid via a replication grid manager for a first list from the offline source system indicating ownership of the cartridges by replication

grid members (step 610). A second list is built from each of the cartridges needed for the takeover operations (step 612). A cartridge is identified as a candidate for the ownership of the cartridge to be taken over by comparing the first list with the second list, (step 614). The second list is transferred to the disaster recover systems (step 616). Ownership of the cartridge(s) is switched (step 618). The method 600 will create and continue part of a production site at each of the disaster recovery systems of the offline source system (step 620). The method 600 ends (step 622).

As mentioned, the DR mode may be activated at each of the DR systems for a source system, for example, a source system that is offline. The DR mode may be entered in order to protect replicas (cartridges/data) and in order to allow takeover operation. Each remote user (disaster recovery systems) may choose to sequentially run (e.g., run the takeover process in turn) the takeover operation. The mechanisms check if the takeover operation may be run for a specific chosen source system (e.g., for an offline source system). The DR systems check that the source system is announced offline. A request is sent in the replication grid via a replication arid manager asking for a list of cartridges from the offline source that are already owned by a different replication gird member. The replication grid manager requests from each replication grid member that has obtained ownership over the offline source's cartridges to send a list of the replication grid members own list of owned cartridges (data). The replication grid manager builds a single list and transfers the list to the DR system(s). The mechanisms build a list of all the cartridges needed for takeover. The needed cartridges may have an ownership stamp from the offline source. The mechanisms compare the lists and identify the specific cartridges that are candidates for ownership takeover. The mechanisms switch ownership of all candidate cartridges to the specific DR systems. The switching of ownership may be performed iteratively and asynchronously. The offline source system's production site may be partially created and continued at each DR site according to the specific cartridges being taken over. By allowing each remote user to choose to run the takeover operation in turn and by partially creating and continuing the production site at each DR site, the present invention provides for switch ownership of the cartridges iteratively and/or in parallel for each remote DR system, particularly where the order of execution of the grid's cartridge list creation operation is a decisive factor for which

DR system gets ownership of which cartridges and also depending on the intersection of datasets between different DR systems.

Fig. 7 is an exemplary block diagram 700 of the type of mutuality between source data sets distributed to different remote systems. When dealing with disjointed datasets 720 distributed over to different DR systems, the takeover operations may be performed separately on each system with no existing danger to the data. When dealing with completely overlapping datasets distributed over to different DR systems, the first takeover operation in any of the DR systems may result in changing cartridge ownership for all the cartridges, so that subsequent takeover operations from other DR systems will return without any results. When dealing with intersecting datasets 710 distributed over to different DR systems, the order of the takeover operation determines which of the different DR system acquires ownership of the cartridges. For example, the first DR system running takeover will acquire ownership of the intersecting cartridges for all the DR systems and also acquire ownership of the first DR system running takeover's unique cartridges. The second DR system running takeover will acquire ownership of the intersecting cartridges between itself (the second DR system running takeover) and DR systems, which have not yet run the takeover operation. Such operations may be performed until no intersecting datasets exists between the remaining DR systems. (The calculations/algorithm may then converge to the disjointed form. Each remote user (disaster recovery systems) exits DR mode for the specific source system.

To illustrate the reconciliation process for ownership synchronization processes for the replica cartridges, the following figures serve to illustrate exemplary embodiments of the mechanisms of the present invention. As previously mentioned, the many-to-many topology may create problems for one-to-one and many-to-one topologies. When different data sets or multiple intersecting data sets are being replicated from a source site to different destinations, such as disaster recovery systems, the systems may suffer prolonged failure resulting in failure to pass/replicate a particular cartridge to a desired destination. To demonstrate such failure and disaster recovery takeover processes, figures 8A-8F are shown to illustrate the mechanisms of the present invention.

Turning first to Fig. 8A, an exemplary block diagram 800 of many-to-many system (four systems) for replication with system #3 being a source system and replicating to all other remote systems. In Fig. 8A, the system #3 (shown in Fig. 8 as 810A) is a source system 810. System #3 810A contains three cartridges for replicating, cartridge 3, 4, and 7. System #3 810A is shown as suffering a prolonged failure (large X being displayed to show the failure). Cartridge 3 has passed/replicated from the source system 810A fully to all of the disaster recovery (DR) systems 812 (shown in Fig. 8 as 812A, 812B, and 812C) within the many-to-many systems. Cartridge 7 completely passed from the source system #3 810A to the disaster recovery system #1 812A, but failed to completely pass/replicate after only replicating some data to system #2 812C. Cartridge 4 was replicated only to the destination of the disaster recovery system #2 812C. The remote systems 812 (disaster recovery systems) working as production sites have now created cartridges 6 and 4 seen with the darker shades (or X shaped lines as seen in 812A and 812B). The darker shaded cartridges indicate the ownership of the cartridges within the systems. The lighter shaded cartridges (or cartridges shown with diagonal lines or speckled dots) indicate only replica cartridges.

Fig. 8B is an exemplary diagram 830 illustrating the source system #3 810A as no longer available (e.g., offline). All remote systems' users are in DR mode for source system #3 810A and therefore may not receive replication from source system #3 810A as illustrated by the blocks 820. The other available source systems continue working normally and the DR systems keep backing up local data. The DR state on source #3 810A may only be temporary. If the DR mode is cancelled, without performing the takeover operation, *ownership synchronization of the some/all cartridges when moving production may be lost.*

Fig. 8C is an exemplary diagram 840 illustrating the takeover operation performed by the first DR system #1 812A. All remote systems' users are in DR mode for source system #3 810A and therefore may not receive replication from source system #3 810A as illustrated by the blocks 820. The DR system user runs an offline announcement process and states that the source system #3 810A may be out of the replication grid manager 820 permanently. Cartridges 3, 7 will change ownership to the DR system #1 812A after checking source system #3 810A cartridges are still owned by the source 810A and not another DR system.

Fig. 8D is an exemplary diagram 850 illustrating the takeover operation performed by the first DR system #2 812C. All remote systems' users are in DR mode for source system #3 810A and therefore may not receive replication from source system #3 810A, as illustrated by the blocks 820. DR system #2 812C requests a list of available cartridges for takeover from the replication grid manager 820. The replication grid manager 820 consults and retrieves a list of all of source system #3's "owned by others" cartridges (in this case ownership had changed only in the first takeover operation to DR system #1). Cartridges 3, 7 will not change ownership since they are already owned by an online system in the grid. Cartridge 4 will change ownership to DR system #2 812C after checking source system #3's 810A cartridge is still owned by the source and not another DR system.

Fig. 8E is an exemplary diagram 860 illustrating the takeover operation performed by the first DR system #4 812B. All remote systems' users are in DR mode for source system #3 810A and therefore may not receive replication from source system #3 810A, as illustrated by the blocks 820. The DR system #4 812B requests a list of available cartridges for takeover from the replication grid manager 820. The replication grid manager 820 consults and retrieves a list of all of source system #3's 810A "owned by others" cartridges (in this case ownership had changed only in the first and second takeover operations to DR systems #1 812A and #2 812C). Cartridges 3 will not change ownership since it is already owned by an online system in the grid. No further operation will be pursued.

Fig. 8F is an exemplary diagram 870 illustrating each of the DR systems (812A-812C). All remote systems' users are in DR mode for source system #3 810A and therefore may not receive replication from source system #3 810A, as illustrated by the blocks 820. Each of the DR systems (812A-812C) may continue to work as normal with each of its production data backed up on the respective DR systems (812A-812C), which may contain data of the newly owned cartridges).

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an

embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

5

10

15

20

25

30

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wired, optical fiber cable, RF, etc., or any suitable combination of the foregoing. Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The program code may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network

(LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention have been described above with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks. The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the above figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the

functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

While one or more embodiments of the present invention have been illustrated in detail, the skilled artisan will appreciate that modifications and adaptations to those embodiments may be made without departing from the scope of the present invention as set forth in the following claims.

CLAIMS

1. A method for coordinated disaster recovery by a processor device in a computing storage environment, the method comprising:

performing a reconciliation process for resolving intersecting and non-intersecting data amongst a plurality of disaster recovery systems for a takeover operation; and

coordinating an ownership synchronization process for a plurality of cartridges via the reconciliation process at the plurality of disaster recovery systems, wherein the plurality of disaster recovery systems continue as at least one of a replication target for a plurality of source systems and as a backup target for a plurality of local backup applications;

wherein the takeover operation includes one of:

activating a disaster recovery (DR) mode for at least one remote system of the source system, wherein the at least one remote system of the at least one of the plurality of source systems declared offline becoming part of the plurality of disaster recovery systems,

allowing the plurality of disaster recovery systems to sequentially perform the takeover operation,

determining the takeover operation may be performed for the at least one of the plurality of source systems declared offline,

sending a request in a replication grid via a replication grid manager for a first list from the at least one of the plurality of source systems declared offline indicating ownership of the plurality of cartridges by a plurality of replication grid members,

building a second list of each of the plurality of cartridges needed for the takeover operations,

identifying at least one of the plurality of cartridges as a candidate for taking over of the at least one of the plurality of cartridges for ownership by comparing the first list with the second list,

transferring the second list to the plurality of disaster recover systems,

switching the ownership of the at least one of the plurality of cartridges,

and

creating and continuing at least a portion of a production site at each of a plurality of disaster recovery systems of the at least one of the plurality of source systems declared offline.

2. The method of claim 1, further including performing at least one of:
 - declaring at least one of the plurality of source systems offline to allow for the plurality of disaster recovery systems to perform the takeover operation, wherein a replication grid is checked to determine if the at least one of the plurality of source systems declared offline has exited the replication grid, and
 - notifying a plurality of replication grid systems of the at least one of the plurality of source systems declared offline.
3. The method of claim 1, further including performing the allowing as one of iteratively and in parallel for each of the plurality of disaster recovery systems.
4. The method of claim 1, wherein the reconciliation process includes performing at least one of:
 - if non-intersecting datasets are distributed among the plurality of disaster recovery systems, performing the takeover operation separately on each of the plurality of disaster recovery systems,
 - if overlapping datasets are distributed among the plurality of disaster recovery systems, executing the takeover operation first by one of the plurality of disaster recovery systems to change the ownership of each of the plurality of cartridges, and
 - if intersecting datasets are distributed among the plurality of disaster recovery systems, determining ownership of the plurality of cartridges based on the order of performing the takeover operation by the plurality of disaster recovery systems.
5. The method of claim 4, wherein determining ownership of the plurality of cartridges based on the order of performing the takeover operation by the plurality of disaster recovery systems includes:
 - for the first one of the plurality of disaster recovery systems to perform the takeover operation, acquiring the ownership of each of the plurality of cartridges that intersect, and

for the second one of the plurality of disaster recovery systems to perform the takeover operation, acquiring the ownership of the plurality of cartridges that intersect between the second one of the plurality of disaster recovery systems and the plurality of disaster recovery systems yet to have performed the takeover operation.

6. The method of claim 1, further including exiting a disaster recovery mode by the plurality of disaster recovery systems for the at least one of the plurality of source systems declared offline.

7. A system comprising means adapted for carrying out all the steps of the method according to any preceding method claim

8. A computer program comprising instructions for carrying out all the steps of the method according to any preceding method claim, when said computer program is executed on a computer system.