



(51) International Patent Classification:
G06F 7/00 (2006.01)

(21) International Application Number:
PCT/US2010/000347

(22) International Filing Date:
4 February 2010 (04.02.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/366,555 5 February 2009 (05.02.2009) US

(71) Applicant (for all designated States except US): **NE-TAPP** [US/US]; 495 East Java Drive, Sunnyvale, CA 94089 (US).

(72) Inventors; and

(71) Applicants : **WU, Grace, Zhanglei** [SG/US]; 495 East Java Drive, Sunnyvale, CA 94089 (US). **MANKOVSKY, Michael, David** [US/US]; 495 East Java Drive, Sunnyvale, CA 94089 (US).

(74) Agent: **ROSSI, Jamie, L.**; ZILKA-KOTAB, PC, P.O. Box 721120, San Jose, CA 95172-1120 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: SYSTEM, METHOD, AND COMPUTER PROGRAM PRODUCT FOR ALLOWING ACCESS TO BACKUP DATA




FIGURE 2

(57) Abstract: A system, method, and computer program product are provided for allowing access to backup data resulting from performance of one or more recovery point backup operations and one or more full backup operations. This is accomplished by associating recovery point backup operations with full backup operations and allowing access to backup data resulting from performance of the recovery point backup operations and the full backup operations. In this way, logs corresponding to the backup operations may be associated with a backup path, such that the logs may be accessed during a restore operation. Furthermore, the recovery point operations may be configured to provide a user defined granularity for data recovery. To this end, in various embodiments, backup data may be stored at a granularity defined by a user and recovered based on a user selection of a recovery point.



Published:

— with international search report (Art. 21(3))

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

-1-

SYSTEM, METHOD, AND COMPUTER PROGRAM PRODUCT FOR ALLOWING ACCESS TO BACKUP DATA

FIELD OF THE INVENTION

[0001] The present invention relates to data backup and, more specifically, to a technique for allowing access to backup data resulting from performance of recovery point backup operations and full backup operations.

BACKGROUND

[0002] In the context of data storage, the term “backup” refers to making copies of data such that the additional copies may be used to restore original data after a data loss event. These additional copies are typically called “backups.” Backups are useful for various purposes. For example, the backups may be used to restore a state following a disaster (i.e. disaster recovery). Additionally, the backups may be used to restore small numbers of files after the files have been accidentally deleted or corrupted.

[0003] To date, data recovery utilizing backups has been limited to accessing backup data corresponding to a single backup path. For example, backing up data recovered using a particular backup may render other backup data not associated with the particular backup useless. By electing to recover data from a particular backup, any other backup data associated with a different backup operation will be rendered useless, as a backup path has already been selected for recovery. Thus, when more than one backup path has been created, backup data may only be recovered from one of the backup paths. This limits the recovery options available once a particular backup is utilized.

[0004] There is thus a need for addressing these and/or other issues associated with the prior art.

SUMMARY

[0005] In order to address issues associated with the prior art, a system, method, and computer program product are provided for allowing access to backup data resulting from performance of one or more recovery point backup operations and one or more full backup operations. This is accomplished by associating recovery point backup operations with full backup operations and allowing access to backup data resulting from performance of the recovery point backup operations and the full backup operations. In this way, logs corresponding to the backup operations may be associated with a backup path, such that the logs may be accessed during a restore operation. Furthermore, the recovery point operations may be configured to provide a user defined granularity for data recovery. To this end, in various embodiments, backup data may be stored at a granularity defined by a user and recovered based on a user selection of a recovery point.

[0006] In various embodiments, the invention may be implemented on a server or client system. For example, the invention may reside in a storage system of a client and/or server system.

[0007] In one exemplary embodiment, data to backup is identified. Additionally, one or more full backup operations are performed on the data. Furthermore, one or more recovery point backup operations are performed on a new portion of the data. Still yet, the one or more recovery point backup operations are associated with at least one of the one or more full backup operations. Moreover, access is allowed to backup data resulting from performance of the one or more recovery point backup operations and the one or more full backup operations. In addition, the backup data is output.

BRIEF DESCRIPTION OF THE DRAWINGS

[0008] Figure 1 illustrates a method for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with one embodiment.

[0009] Figure 2 illustrates a system for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with one embodiment.

[0010] Figure 3A illustrates a method for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with another embodiment.

[0011] Figure 3B shows a method that may be implemented in the context of point-in-time (PIT) restore logic for updating an exclusive-list file under the full backup, in accordance with one embodiment.

[0012] Figure 3C shows a method for retrieving a recovery point for a full backup, in accordance with one embodiment.

[0013] Figure 4 illustrates a backup configuration flow for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with one embodiment.

[0014] Figure 5 illustrates an interface for scheduling a frequent recovery point backup, in accordance with one embodiment.

[0015] Figure 6 illustrates an interface for configuring a frequent recovery point backup, in accordance with one embodiment.

[0016] Figure 7 illustrates an interface for display when a frequent recovery point backup job is configured, in accordance with one embodiment.

[0017] Figure 8 illustrates an interface for authenticating a user prior to allowing the configuration or implementation of frequent recovery point backups, in accordance with one embodiment.

[0018] Figure 9 illustrates an interface for restoring data to a full backup or a frequent recovery point backup state, in accordance with one embodiment.

[0019] Figure 10 illustrates an interface for selecting a recovery point, in accordance with one embodiment.

[0020] Figure 11 is an illustration of a network environment in which an embodiment may be implemented.

[0021] Figure 12 is a block diagram illustrating architecture of a server that may include an embodiment.

[0022] Figure 13 illustrates an example of the operating system of a server according to one possible embodiment.

DETAILED DESCRIPTION

[0023] A system, method, and computer program product are provided for allowing access to backup data resulting from performance of one or more recovery point backup operations and one or more full backup operations. This is accomplished by associating recovery point backup operations with full backup operations and allowing access to backup data resulting from performance of the recovery point backup operations and the full backup operations. In this way, logs corresponding to the backup operations may be associated with a backup path, such that the logs may be accessed during a restore operation. Furthermore, the recovery point operations may be configured to provide a user defined granularity for data recovery. To this end, in various embodiments, backup data may be stored at a granularity defined by a user and recovered based on a user selection of a recovery point.

[0024] In one exemplary embodiment, data to backup is identified. Additionally, one or more full backup operations are performed on the data. Furthermore, one or more recovery point backup operations are performed on a new portion of the data. Still yet, the one or more recovery point backup operations are associated with at least one of the one or more full backup operations. Moreover, access is allowed to backup data resulting from performance of the one or more recovery point backup operations and the one or more full backup operations. In addition, the backup data is output.

[0025] Figure 1 illustrates a method 100 for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with one embodiment. As shown, data to backup is identified. See operation 102.

[0026] The data may include any type of data capable of being backed up (e.g. client data, metadata, etc.). In one embodiment, the data to be backed up may include any type of Exchange data. For example, the data may include database files, transaction log files,

and/or Exchange system files. As an option, the data may include Exchange Mail Store and Public Folders (i.e. user accounts). In another embodiment, the data may include email messages (e.g. storage group log files, etc.).

[0027] Additionally, one or more full backup operations are performed on the data. See operation **104**. In the context of the present description, a full backup operation refers to any data backup operation capable of storing a copy of data in databases, logs (e.g. transaction logs, etc.), and system files. In one embodiment, the full backup operation may include truncating the log files up to a checkpoint.

[0028] Furthermore, one or more recovery point backup operations are performed on a new portion of the data. See operation **106**. In the context of the present description, a recovery point backup operation refers to any backup operation capable of storing transaction logs and data associated with the transaction logs (e.g. metadata, etc.).

[0029] Furthermore, the new portion of the data may include any data that is stored after the most recent full backup operation that occurred before a current recovery point backup operation or any data that is stored after the most recent recovery point backup operation that occurred before the current recovery point backup operation.

[0030] In one embodiment, the recovery point backup operation may only create and store logs associated with new data. In this case, the new data may include any new data present since the performance of a corresponding full backup operation. For example, the new data may include changes, additions, and/or deletions to data corresponding to a full backup operation.

[0031] Additionally, the one or more recovery point backup operations may be performed at predefined time intervals. As an option, the predefined time intervals may be defined by a user. As another option, the one or more recovery point backup operations may be performed automatically, based on changes made to the data. In either

case, the number of the recovery point backup operations to be performed may be defined by a user.

[0032] Still yet, the one or more recovery point backup operations are associated with at least one of the one or more full backup operations. See operation 108. In this case, associating the one or more recovery point backup operations with at least one of the one or more full backup operations may include associating a pass identifier with each of the one or more recovery point backup operations.

[0033] In the context of the present description, a pass identifier refers to any identifier capable of being used to associate a particular recovery point backup operation, or data associated therewith, with one of the full backup operations, or data associated therewith. In one embodiment, the pass identifier may identify a path in a data structure associated with a particular series of backup operations. The pass identifier may indicate one of the full backup operations corresponding to the one or more recovery point backup operations.

[0034] Furthermore, the pass identifier may be utilized to allow access to the backup data. As an option, the pass identifier may include metadata for associating the one or more recovery point backup operations with at least one of the one or more full backup operations.

[0035] Moreover, access is allowed to backup data resulting from performance of the one or more recovery point backup operations and the one or more full backup operations. See operation 110. For example, the backup data may be generated by performing the one or more recovery point backup operations and/or the one or more full backup operations. Subsequently, a recovery operation may be performed as a result of data loss, corruption, or failure. In this case, the performance of the recovery operation may enable access to the backup data by recovering the backup data. Furthermore, the recovery operation may be a user initiated operation. As an option, a user may initiate a recovery operation using a recovery program (e.g. software).

[0036] In one embodiment, this recovery operation may include a point in time recovery operation. In other words, backup data may be accessed based on a desired point in time for recovery. In another embodiment, the recovery operation may allow access to a backup state. For example, access may be allowed to a state corresponding to a full backup, a recovery point backup, and/or any log or data associated therewith.

[0037] In either case, as an option, the backup data to be accessed may be defined by a user. For example, the user may define the data to be accessed (e.g. prior to a recovery operation, etc.) utilizing a user interface. Furthermore, allowing access to the backup data may include allowing access to the first subsequent FRP backup data created before second subsequent FRP backup data, where the first FRP backup data and the second FRP backup data are associated with different full backup operations. For example, the first backup data may be associated with a first full backup operation and the second backup data may be associated with a second full backup operation. In this way, data recovery is not limited to only backup data in a particular series of backup operations.

[0038] It should be noted that allowing access to the backup data may include allowing the data to be recovered. Additionally, allowing access to the backup data may include transferring the backup data from a backup storage location to a storage location that is not necessarily used for backup data. Furthermore, allowing access to the backup data may include allowing the backup data to be read or copied.

[0039] Once the backup data is accessed, the backup data may be output. See operation 112. In this case, the backup data may be output to a memory location, a database, a file, a display, and/or any other location. In one embodiment, at least a portion of the backup data may be output such that the portion is displayed to a user. In one embodiment, outputting the data may include restoring the data. As option, the backup data may be stored in a data structure embodied on a tangible computer readable medium (e.g. a disk, a hard drive, volatile or nonvolatile memory, etc.).

[0040] In one exemplary implementation, a full backup operation may be performed on both databases and transaction logs. Subsequently, one or more recovery point backup operations may be performed only on newly generated transaction logs. In this case, the new data are the logs that are generated after the performance of the full backup operation. Furthermore, the new data may be logs that are generated after a previous recovery point backup operation, if the previous recovery point operation was performed after the full backup operation. Data recovery may then be allowed from the backup data resulting from performance of full backup operation and one or more recovery point backup operations. In this case, data recovery refers to recovering the backup data to a state that the data was in before a backup operation was performed on the data.

[0041] In another exemplary implementation, an association may be generated between one or more recovery point backup operations and at least one full backup operation. The association information may then be stored in a file. Backup data corresponding to a user selected backup point may then be retrieved based on the association information saved in the file. Furthermore, a storage system may be restored in accordance with the user selected backup point by transferring the backup data from a backup storage location to the Exchange live system.

[0042] More illustrative information will now be set forth regarding various optional architectures and features of different embodiments with which the foregoing technique may or may not be implemented, per the desires of the user. It should be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. Any of the following features may be optionally incorporated with or without the other features described.

[0043] Figure 2 illustrates a storage system **200** for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with one embodiment. As an option, the storage system **200** may be implemented to carry out the functionality of Figure 1. Of course, however, the storage system **200** may be implemented in any desired environment. It

should be noted that the definitions provided above equally apply to the present description.

[0044] As shown, the storage system **200** may be in communication with a user system **210**. As an option, the storage system **200** may be in communication with the user system **210** over one or more networks **212**. In this case, the user system **210** may represent any system capable of communicating with the storage system **200** (e.g. a desktop computer, laptop computer, etc.).

[0045] As an option, the user system **210** may be utilized to access and implement functionality associated with the storage system **200**. For example, a user (e.g. an administrator, etc.) may utilize the user system **210** to schedule full data backup operations and frequent recovery point (FRP) backup operations. Furthermore, the storage system **200** may be utilized by the user to store data (e.g. backup and other data, etc.).

[0046] To ensure a high level of data protection, the storage system **200** may utilize software capable of being executed by a processor (e.g. a processor associated with the storage system **200** and/or user system **210**, etc.) to asynchronously replicate newly generated transaction logs to a backup directory **202** at scheduled intervals between the full database backups as part of a frequent recovery point backup operation. In one embodiment, two copies of a log file may be retained. For example, one copy may be stored in a live system directory **204** and one copy may be stored inside the backup directory **202** as a backup if hardlink technology is not used.

[0047] In one embodiment, invoking a full backup or an FRP backup operation may create a folder **206** under the backup directory **202** on the storage system **200**, the same as in a full backup operation, using a unique naming convention. This folder **206** may store all the replicated transaction logs since the last full backup or the last FRP backup up to the next FRP backup. In this way, there may not be any duplicated logs that are part of an FRP backup. Additionally, the storage system **200** may trace which transaction logs to

copy over in each backup utilizing a software application. As an option, a pass identifier may be utilized to trace which transaction logs to copy over in each backup.

[0048] In one embodiment, a task scheduler 208 (e.g. a Windows Task Scheduler, etc.) may be utilized to execute the FRP backup operation. As an option, the FRP backup operation may be configured to run only after a full backup has been created. In this case, if no full backup has ever been created, the FRP backup operation may fail with the ERROR type of event showing a detailed message in a status window, report, and/or a Windows event log.

[0049] In various embodiments, the FRP backup operations may be performed based on different criteria. For example, in one embodiment, the FRP backup operation may be a time driven FRP backup operation. In this case, a user may be provided with an option to fine-tune FRP backup behavior to match business requirements and operating environments. In one embodiment, this option for specifying how often the FRP backup operation will run may be presented on an FRP backup view/wizard interface, as advanced settings, and/or in a command line form.

[0050] In another embodiment, the FRP backup operation may be a data driven FRP backup operation. For example, a data driven FRP backup operation may be triggered by a data change in one or more log files. In this case, the FRP backup may be triggered based on real-time or near real-time changes. As another option, a number of newly closed log files may be a trigger point for the data driven FRP backup operation. In one embodiment, a minimum or maximum frequency of FRP backups may also be defined (e.g. specified by an administrator, etc.) such that FRP backups occur frequent enough but do not occur too frequently. In any of these cases, the FRP backup may be triggered automatically utilizing software. As another option, a user may trigger an FRP backup operation.

[0051] In one embodiment, all the FRP backups between two full database backups may be retained for restoration in order to implement a point-in-time restore ability. To

protect the data loss on replicated transaction logs, a snapshot of the LUN (logic unit numbers) associated with backed logs may be created at the end of an FRP backup and the snapshot created by the previous FRP backup may be deleted. In this case, the FRP backup operation may not truncate any logs such that the logs backed up by FRP backup may still exist in the live system directory **204**. The replicated data in the backup directory **206** may function to provide data protection to the granularity of point-in-time restore to a particular transaction log if LUNs associated with the live logs are corrupt. Each of the FRP backups may use a naming convention that includes date/time information (e.g. 03-14-2007_09.30.00, etc.).

[0052] The FRP backups may enable restoration of data to within the specified time interval of a data loss event. In an up-to-the-minute restore, the storage system **200** may use software to play forward all of the logs in the FRP backups and full backups that were created since the backup to be restored, along with the existing transaction logs in the live system directory **204**.

[0053] As an option, in a point-in-time restore, the log file name and an associated closing time may be displayed to a user. The user (e.g. an administrator, etc.) may then choose to restore databases up to the point of time when the log file was closed. The database may then be restored from the previous full backup and the logs in subsequent FRP backups may be replayed up to that specified time.

[0054] It should be noted that, in most cases, verification (including database integrity verification and transaction log consistent verification) may take a long time and generate heavy I/O for a server. Thus, in one embodiment, the FRP backups may be configured such that log verification is not performed and is deferred to the next full backup. Additionally, as an option, similar to a full database backup, the FRP backup may also generate a report with all types of information regarding the backup.

[0055] By implementing FRP backups, data loss may be minimized. Data may be recovered in the event of data loss or system corruption to the last FRP backup. A user

-13-

may restore the last full database backup along with restoration of all subsequent FRP backups.

[0056] For example, an hourly full backup may be performed at 8:00 a.m. and an FRP backup may be performed every 15 minutes between the full backups (e.g. at 8:15 a.m. and 8:30 a.m., etc.). If at 8:35 a.m., some log files were lost or corrupted, a user may restore the 8:00 a.m. full backup and play forward the logs in 8:15 a.m. and 8:30 a.m. FRP backups to recover the data up to 8:30 a.m.

[0057] Additionally, users may be allowed to restore the database up to a particular log file. For example, the user may restore the last full database backup and play forward all the logs from the subsequent FRP backups within the specified range (e.g. a user specified range, etc.).

[0058] Table 1 shows a data structure for tracing transaction logs, in accordance with one embodiment.

[0059] Table 1

Snapshot Information			
			-EXCH__SNAPMGR-32
			-SG__SG1
			-03-14-2007_08.00.00 ← full backup
			-SnapInfo__03-14-2007_08.00.00.sme
			-Vss__03-14-2007_08.00.00.xml
			-Vss__03-14-2007_08.00.00_{.....}.xml
			-Logs
			-E0000001.log
			-E0000002.log
			-03-14-2007_08:15.00 ← FRP backup
			-Logs
			-E0000003.log
			-E0000004.log
			-E0000005.log

```

|-03-14-2007_08:30.00 ← FRP backup
|   |-Logs
|     |-E0000006.log
|     |-E0000007.log
|-03-14-2007_08:45.00 ← FRP backup
|   |-Logs
|     |-E0000008.log
|
|-03-14-2007_09.00.00 ← full backup
|   |-SnapInfo__03-14-2007_09.00.00.sme
|   |-Vss__03-14-2007_09.00.00.xml
|   |-Vss__03-14-2007_09.00.00_{.....}.xml
|   |-Logs
|     |-E0000009.log
|     |-E000000A.log
|
|-03-14-2007_09:15.00 ← FRP backup
|   |-Logs
|     |-E000000B.log
|     |-E000000C.log
|-03-14-2007_09:30.00 ← FRP backup
|   |-Logs
|     |-E000000D.log
|     |-E000000E.log
|     |-E000000F.log
|-SG__SG2

```

[0060] As shown in Table 1, the hourly full backup was performed at 8:00 a.m., 9:00 a.m., etc. The FRP backup was performed every 15 minutes between the full backups at 8:15 a.m., 8:30 a.m., 8:45 a.m., 9:15 a.m., etc. If at 10:00 a.m. an administrator noticed that a virus was encountered at 8:55 a.m. and invaded the log files, the administrator may choose to recover the database to a state of the database at 8:45 a.m., when the previous FRP backup was created. Furthermore, in the case of a disaster, a user may restore the last full database backup and all the subsequent FRP backups from a mirrored destination server to the recovery server.

[0061] It should be noted that the storage system **200** may include any system capable of storing data and implementing backup operations. Furthermore, the storage

system **200** may represent a client or a server system. Additionally, the storage system **200** may be in communication with one or more other systems and/or storage units over one or more networks.

[0062] Figure **3A** illustrates a method **300** for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with another embodiment. As an option, the method **300** may be implemented in the context of the details of Figures **1-2**. Of course, however, the method **300** may be implemented in any desired environment. Again, the definitions provided above equally apply to the present description.

[0063] As shown, a full backup is to back up databases first. See operation **302**. As part of the full backup, transaction logs are copied from Microsoft Exchange server live system. See operation **304**. At the end of full backup, the committed logs may be truncated from the server live system.

[0064] As part of the full backup, all available logs may be copied. It is then determined whether an FRP backup is to be performed. See operation **306**. New logs generated after the full backup may be backed up in the following FRP backup which may be run at specified interval.

[0065] In this case, it may be determined whether an FRP backup is to be performed based on a time parameter (e.g. a timed FRP backup setting, etc.), a log status parameter (e.g. a flag indicating data has changed, etc.), and/or other parameters. As an option, a user may create a backup schedule for performing FRP backups. Furthermore, each recovery point backup operation may be performed after an associated full backup operation. For example, after the performance of a full backup operation, FRP backups may be performed and may be associated with these full backup operations (e.g. utilizing metadata, etc.).

[0066] If it is determined that an FRP backup is to be performed, any newly generated logs are copied. See operation 308. The newly generated logs may include any logs generated since the last full backup was performed or since a last FRP backup was performed. Furthermore, the copies of the newly generated logs may be associated with a corresponding full backup operation.

[0067] It is then determined whether a full backup is to be performed. See operation 310. Similar to the FRP backup, it may be determined that a full backup is to be performed based on a time parameter (e.g. a user defined time parameter, etc.) and/or based on a data parameter (e.g. a size, time, and/or type of modified data, etc.).

[0068] If it is determined that a full backup is to be performed, all logs are copied as part of the full backup operation. See operation 312. In this case, the logs may include all new committed logs and uncommitted logs. Furthermore, a snapshot of all the storage group databases and transaction logs may be stored as part of the full backup operation. Additionally, as an option and as part of the full backup operation, log files up to a checkpoint may be truncated by Exchange.

[0069] In one embodiment, the full backup may create a snapshot of data using a naming convention including server information, and a timestamp (e.g. "eloginfo__<server>_<timestamp>", "exchsnap__<server>_<timestamp>" etc.). In this case, if the FRP backup uses the same prefix in the same naming convention as full backup, it may be difficult to identify whether the snapshot is created by the full backup or the FRP backup using only the name. This may also render a previous snapshot FRP backup hard to find for deletion.

[0070] Accordingly, in one embodiment, a different prefix (e.g. "efrpingo," etc.) may be used in the snapshot name for the FRP backup. After the copying of all the logs, the FRP backup operation may create a new snapshot of the data using this naming convention and delete the older snapshots. In this way, available recovery points may be

retrieved in the specified date and time frame for data (e.g. a storage group, etc.) based on a selected backup.

[0071] In one embodiment, every time a point-in-time (PIT) restore is performed, an xml file (e.g. "SG__<SGName>_PIT.xml," etc.) including the restored backup name may be created, or updated if it exists, at a storage group level. This file may be used to keep track of the last restored backup of the storage group. Additionally, an exclusive-backup list file (e.g. "PITInfo__<timestamp>.xml," etc.) may be created under the restored full backup. This file may include a list of backups that may be excluded from future PIT restores. In this case, the exclusive-backup list may be used for a full backup.

[0072] Figure 3B shows a method 320 that may be implemented in the context of PIT restore logic for updating an exclusive-list file under the full backup, in accordance with one embodiment.

[0073] As shown, it is determined whether a "Last restored backup" information file exists at the end of executing PIT restore logic for updating an exclusive-list file under the full backup. See operations 322-324. If the "Last restored backup" information file does not exist, the file is created for a current storage group and the current restored backup timestamp is written to the file. See operation 326-328.

[0074] If the "Last restored backup" information file does exist, the "Last restored backup" file is updated with the restored backup timestamp. See operation 330. Additionally, the exclusive-list file is updated under the full backup. See operation 332.

[0075] In this example, during the next backup, the "Last restored backup" information may be retrieved and saved in the exclusive-list file of the new backup. The "Last restored backup" information of a backup may be used to determine if the backup should be included in the list for retrieval of recovery points when an attempt to restore a backup is raised.

[0076] In one embodiment, to retrieve a recovery point for a full backup, all the backups following this backup may be triaged, the backups in the exclusive list may be excluded, and the backups from the same backup path where no PIT restore occurred in between may be included. In this case, if there are common logs in two backups, the backup from the later backup may be utilized.

[0077] Figure 3C shows a method 340 for retrieving a recovery point for a full backup, in accordance with one embodiment. As shown, a variable indicating the last restored backup (e.g. “lastRestoredBackup,” etc.) is set to be equal to the full backup to be restored. See operation 342. An exclusive backup list of the full backup to be restored is then loaded. See operation 344.

[0078] It is then determined if a backup exists in a snapshot information directory. See operation 346. For each backup in the directory, it is determined whether the enumerated backup is in the exclusive backup list. See operation 348. If the enumerated backup is not in the list, it is determined whether the enumerated backups “Last restored backup” indicator is equal to the variable indicating the last restored backup. See operation 350. If it is, this backup is added in a snapshot information list array. See operation 352. It is then determined whether there is another backup in the snapshot information directory. This process occurs for all backups in the snapshot information directory.

[0079] As an example of how point-in-time restore may be utilized to retrieve recovery points, four backups may be considered (e.g. bkup1-bkup4). It may be determined to restore to “bkup2” (e.g. by an administrator, etc.). After the restore to bkup2, a fifth backup, bkup5, may be created. Table 4 shows an illustration of this scenario, in accordance with one embodiment.

[0080] Table 2

bkup1	bkup2	bkup3	bkup4
(1,2)	(3,4)	(5,6,7)	(8,9)
\			bkup5


(3,4,5,...)

[0081] In this case, each of the backups include logs (e.g. logs 1,2,3, etc.). In this example, bkup5 includes the same log file name as in bkup3. However, these logs may be different files (i.e. different log data).

[0082] Once a backup is restored, a new restore path may be created. In this case, all of the backups following the restore point (e.g. bkup3 and bkup4 in this case) may be placed into the exclusive list of the restored backup (i.e. bkup2). The recovery points will be available from the new path. Additionally, bkup3 and/or bkup4 may also be chosen as a restore point. Any other new restore path may be created accordingly.

[0083] Table 3 shows the backup flow illustrated in Table 2 with additional backups and an additional restore path, in accordance with one embodiment.

[0084] Table 3



[0085] Figure 4 illustrates a backup configuration flow 400 for allowing access to backup data resulting from performance of one or more recovery point operations and one or more full backup operations, in accordance with one embodiment. As an option, the backup configuration flow 400 may be viewed in the context of the details of Figures 1-3. Of course, however, the backup configuration flow 400 may be implemented in any desired environment. Further, the definitions provided above equally apply to the present description.

[0086] As shown, a plurality of full backups 402 are performed (e.g. utilizing software executed by a processor, etc.). In this case, the full backups 402 are performed on a timed basis where Full Backup 1 refers to a first full backup that was performed before a second full backup, Full Backup 2, which was performed before a third full backup, Full Backup 3. At some point (e.g. at a time after the performance of three full backups, etc.), a user may desire to restore data associated with the full backups 402. For example, a user may choose to restore to Full Backup 1. After the restore operation, a fourth full backup, Full Backup 4, is created.

[0087] In this case, the Full Backup 4 will include data associated with the Full Backup 1 (e.g. logs 1, 2, etc.) and any new data (e.g. log 3, etc.). After the creation of the Full Backup 4, a fifth full backup, Full Backup 5, may be initiated (e.g. based on a time setting, etc.). The Full Backup 5 will include all untruncated data associated with the Full Backup 4 (e.g. logs 1, 2, 3, etc.) and any new data (e.g. log 4, etc.).

[0088] At this point, the user may desire to again restore the data. In this case, the user may desire to restore the data from the Full Backup 2. After the second restore, a sixth full backup, Full Backup 6, may be initiated. The Full Backup 6 will include all data associated with the Full Backup 2 (e.g. logs 3, 4, etc.) and any new data (e.g. logs 5, 6, etc.). Thus, after restoring data to a first full backup state (e.g. Full Backup 1, etc.) and performing one or more subsequent full and FRP backups, a second full backup state (e.g. Full Backup 2, etc.) or an associated FRP backup, in a path separate from the first full backup state, may be accessed.

[0089] It should be noted that, FRP backups may be performed at any point after a full backup. Furthermore, these FRP backups may each be associated with a full backup. Additionally, although Figure 4 shows logs with the same name at different full backup stages, these logs with the same name may not necessarily include the same data.

[0090] For example, the Full Backup 4 associated with the first restore path may include "Log 3." Furthermore, the Full Backup 2 may include a different "Log 3," where

the Full Backup 2 is located in a separate restore path than the Full Backup 4. Thus, both of the Log 3s may be associated with their respective restore path and/or full backup state to distinguish the logs. In one embodiment, this association may be accomplished utilizing a pass identifier (e.g. meta data, etc.).

[0091] Using the pass identifier, any FRP backup, full backup, and/or data associated therewith may use the pass identifier to maintain an association with an originating path and/or a restore path. For example, Logs 4 and 5 associated with the Full Backup 6 may be logs stored as part of an FRP backup. The FRP backup and the associated logs may utilize a pass identifier to maintain an association with the Full Backup 6.

[0092] Figure 5 illustrates an interface 500 for scheduling a frequent recovery point backup, in accordance with one embodiment. As an option, the interface 500 may be utilized in the context of the details of Figures 1-4. Of course, however, the interface 500 may be implemented in any desired environment. Again, the definitions provided above equally apply to the present description.

[0093] In operation, when a button or text on an interface corresponding to scheduling an FRP backup is selected, the interface 500 may be displayed to a user. Using the interface 500, the user may be presented with an option to display another interface for choosing the frequency of the FRP backup (e.g. from a drop down list, etc.). In this case, the user may be presented with many different frequency options. For example, the frequencies presented to the user may begin at 10 minutes and increment by 5 minutes up to 1 hour. After 1 hour, the frequencies may increment by 1 hour up to 24 hours.

[0094] Figure 6 illustrates an interface 600 for configuring a frequent recovery point backup, in accordance with one embodiment. As an option, the interface 600 may be utilized in the context of the details of Figures 1-5. Of course, however, the interface 600 may be implemented in any desired environment. Further, the definitions provided above equally apply to the present description.

[0095] In operation, the interface **600** may be displayed to a user when an FRP backup operation is not yet scheduled and/or configured. In this case, the user may select a “Create Job” button and choose the data on which the FRP backup is to be performed (e.g. a first storage group, a second storage group, etc.).

[0096] Figure 7 illustrates an interface **700** for display when a frequent recovery point backup job is configured, in accordance with one embodiment. As an option, the interface **700** may be utilized in the context of the details of Figures 1-6. Of course, however, the interface **700** may be implemented in any desired environment. The definitions provided above equally apply to the present description.

[0097] In operation, the interface **700** may be displayed when an FRP backup scheduled job is already created. A user may select an “Update Job” button and the user may specify the data to backup and the frequency of the backup. As an option, the user may be required provide authentication information prior to configuring or implementing the FRP backup jobs/frequencies.

[0098] Figure 8 illustrates an interface **800** for authenticating a user prior to allowing the configuration or implementation of FRP backups, in accordance with one embodiment. As an option, the interface **800** may be utilized in the context of the details of Figure 6 or 7. Of course, however, the interface **800** may be implemented in any desired environment. Again, the definitions provided above equally apply to the present description.

[0099] In operation, once a user inserts the required authentication information, and selects the “OK” button, an FRP backup job with the selected frequency will be created. In one embodiment, the FRP backup job with the selected frequency may be created in a Windows Task Scheduler.

[00100] Figure 9 illustrates an interface 900 for restoring data to a full backup or an FRP backup state, in accordance with one embodiment. As an option, the interface 900 may be utilized in the context of the details of Figures 1-8. Of course, however, the interface 900 may be implemented in any desired environment. Furthermore, the definitions provided above equally apply to the present description.

[00101] As shown, the interface 900 may display a list of data that may be restored. A user may select data to restore from this list. Additionally, the user may select a “Point-In-Time” restore type button or a “Select Recovery Point” button. In this way, backup data may be recovered for a particular point in time, or from a particular backup state (e.g. a particular full backup state, an FRP backup state, etc.).

[00102] Figure 10 illustrates an interface 1000 for selecting a recovery point, in accordance with one embodiment. As an option, the interface 1000 may be utilized in the context of the details of Figures 1-9. Of course, however, the interface 1000 may be implemented in any desired environment. Additionally, the definitions provided above equally apply to the present description.

[00103] As shown, the interface 1000 may present a user with an option to restore from a point in time or from a particular recovery point (e.g. an FRP point, etc.). In one embodiment, a default restore date presented to the user may be configured be the date when the selected full backup was created. As an option, when the user selects a box corresponding to the date, a small calendar window may be displayed. In this case, selecting any date before the default date may be defined as an error.

[00104] Figure 11 is an illustration of a network environment 1100 in which an embodiment may be implemented. As an option, the network environment 1100 may be implemented in the context of the storage system 200 of Figure 2. Of course, however, it should be noted that the network environment 1100 may be used in any desired environment and incorporate any one or more of the features described in the previous figures.

[00105] Of course, the various embodiments described herein are not limited to any particular environment, and may be implemented in various storage processes. In the present illustration, the storage system includes a storage server, such as a server 1120. The server 1120 is coupled with a storage subsystem 1140, which includes a set of mass storage devices 1150, and to a set of clients 1110 through a network 1130, such as a local area network (LAN) or other type of network. Each of the clients 1110 may be, for example, a conventional personal computer (PC), workstation, or any other client.

[00106] The storage subsystem 1140 is managed by the server 1120. For example, the server 1120 may receive and respond to various read and write requests from the clients 1110, directed to data stored in or to be stored in the storage subsystem 1140. The mass storage devices 1150 in the storage subsystem 1140 may be, for example, magnetic disks, optical disks such as compact disks-read only memory (CD-ROM) or digital versatile/video disks (DVD)-based storage, magneto-optical (MO) storage, tape-based storage, or any other type of non-volatile storage devices suitable for storing large quantities of data.

[00107] Although the server 1120 is illustrated as a single unit in Figure 11, it can have a distributed architecture. For example, the server 1120 can be implemented with multiple distributed storage servers. It can also include a physically separate network module (e.g., "N-module") and disk module (e.g., "D-module") (not shown), which communicate with other storage servers over an external interconnect. The N-module acts as a front-end of the server 1120, exporting services to clients; and the D-module acts as the back-end, managing the underlying storage of the server 1120. The N-module and D-module can be contained in separate housings and communicate with each other via network connections. Thus, the N-module and the D-module enable the server 1120 to be physically separated into multiple modules that can be distributed across a network.

[00108] In another embodiment, the server 1120 may have an integrated architecture, where the network and data components are all contained in a single box or unit. The

server **1120** further may be coupled through a switching fabric to other similar storage systems (not shown) that have their own local storage subsystems. In this way, all of the storage subsystems can form a single storage pool, to which any client of any of the storage systems has access.

[00109] In one embodiment, the server **1120** may be equipped with hardware and/or software logic **1160** for carrying out the functionality of Figures 1-10. While, in the present embodiment, such logic **1160** is shown to be a component of the server **1120**, it should be noted that it may reside elsewhere as well (e.g. exterior of the server **1120**, etc.).

[00110] It should be noted that, in one embodiment, the server **1120** may represent the storage system **200** of Figure 2. Furthermore, in another embodiment, the server **1120** and the storage subsystem **1140** may represent the storage system **200** of Figure 2. Thus, the server **1120** and the storage subsystem **1140** may be configured to implement the functionality described in the context of the details of the storage system **200** of Figure 2.

[00111] Figure 12 is a block diagram illustrating architecture of a server **1200** that may include an embodiment. In one embodiment, the server **1200** may represent the server **1120** of Figure 11. Of course, however, it should be noted that the server **1200** may be implemented in any desired environment and incorporate any one or more of the features described in the previous figures.

[00112] The server **1200** includes one or more processors **1205** and memory **1210** coupled to a interconnect **1220**. The interconnect **1220** shown in Figure 12 is an abstraction that represents any one or more separate physical buses, point-to-point connections, or both connected by appropriate bridges, adapters, or controllers. The interconnect **1220**, therefore, may include, for example, a system bus, a Peripheral Component Interconnect (PCI) bus, a HyperTransport or industry standard architecture (ISA) bus, a small computer system interface (SCSI) bus, a universal serial bus (USB),

IIC (I2C) bus, or an Institute of Electrical and Electronics Engineers (IEEE) standard 1394 bus, sometimes referred to as “Firewire”.

[00113] The processor(s) 1205 may include central processing units (CPUs) of the server 1200 and, thus, control the overall operation of the server 1200. In certain embodiments, the processor(s) 1205 accomplish this by executing software stored in memory 1210. The processor(s) 1205 may be, or may include, one or more programmable general-purpose or special-purpose microprocessors, digital signal processors (DSPs), programmable controllers, application specific integrated circuits (ASICs), programmable logic devices (PLDs), or the like, or a combination of such devices.

[00114] The memory 1210 is or includes the primary memory of the server 1200. The memory 1210 represents any form of random access memory (RAM), read-only memory (ROM), flash memory, or the like, or a combination of such devices. In use, the memory 1210 stores, among other things, the operating system 1215 of the server 1200.

[00115] Also connected to the processor(s) 1205 through the interconnect 1220 may be one or more internal mass storage devices 1225, a storage adapter 1230 and a network adapter 1235. The internal mass storage devices 1225 may be or include any medium for storing large volumes of instructions and data 1240 in a non-volatile manner, such as one or more magnetic or optical-based disks. The storage adapter 1230 allows the server 1200 to access a storage subsystem 1245 and may be, for example, a Fibre Channel adapter or a SCSI adapter. The storage adapter 1230 may interface with a D-module portion of the server. The network adapter 1235 provides the server 1200 with the ability to communicate with remote devices, such as clients, over a network 1250 and may be, for example, an Ethernet adapter. The network adapter 1235 may interface with an N-module portion of the server 1200.

[00116] Figure 13 illustrates an example of the operating system 1300 of a server according to one possible embodiment. As an option, the operating system 1300 may be

installed on the server **1200** of Figure 12. Of course, however, it should be noted that the operating system **1300** may be used in any desired environment and incorporate any one or more of the features described in the previous figures.

[00117] As shown, the operating system **1300** includes several modules, or “layers.” These layers include a file system manager **1305**. The file system manager **1305** is software that keeps track of the directory structure (hierarchy) of the data stored in a storage subsystem and manages read/write operations on the data (i.e. executes read/write operations on the disks in response to client requests).

[00118] The operating system **1300** also includes a protocol layer **1310** and an associated network access layer **1315**, to allow a server to communicate over a network to other systems, such as clients. The protocol layer **1310** implements one or more of various higher-level network protocols, such as Network File System (NFS), Common Internet File System (CIFS), Hypertext Transfer Protocol (HTTP) and/or Transmission Control Protocol/Internet Protocol (TCP/IP). The network access layer **1315** includes one or more drivers, which implement one or more lower-level protocols to communicate over the network, such as Ethernet. Interactions between clients and mass storage devices (e.g. disks, etc.) are illustrated schematically as a path, which illustrates the flow of data through the operating system **1300**.

[00119] The operating system **1300** further includes a storage access layer **1320** and an associated storage driver layer **1325** to allow a server to communicate with a storage subsystem. The storage access layer **1320** implements a higher-level disk storage protocol, such as RAID, while the storage driver layer **1325** implements a lower-level storage device access protocol, such as Fibre Channel Protocol (FCP) or SCSI. In one embodiment, the storage access layer **1320** may implement a RAID protocol, such as RAID-4 or RAID-DPT™ (RAID double parity for data protection provided by Network Appliance, Inc.), and therefore may alternatively be referred to as RAID layer.

[00120] The foregoing description has set forth only a few of the many possible implementations. For this reason, this detailed description is intended by way of illustration, and not by way of limitations. Variations and modifications of the embodiments disclosed herein may be made based on the description set forth herein, without departing from the scope and spirit of the present application. Furthermore, the data structures discussed herein can be the same or different data structures in accordance with various embodiments of the invention.

[00121] It is only the following claims, including all equivalents, that are intended to define the scope of the various embodiments. Moreover, the embodiments described above are specifically contemplated to be used alone as well as in various combinations. Accordingly, other embodiments, variations, and improvements not described herein are not necessarily excluded.

CLAIMS

What is claimed is:

1. A method, comprising:
identifying data to backup;
performing one or more full backup operations on the data;
performing one or more recovery point backup operations on a new portion of the data;
associating the one or more recovery point backup operations with at least one of the one or more full backup operations;
allowing access to backup data resulting from performance of the one or more recovery point backup operations and the one or more full backup operations; and
storing the backup data in a data structure embodied on a tangible computer readable medium.
2. The method of claim 1, wherein the data includes log files.
3. The method of claim 2, wherein the one or more full backup operations include truncating the log files.
4. The method of claim 1, wherein the one or more recovery point backup operations are performed at predefined time intervals.
5. The method of claim 4, wherein the predefined time intervals are defined by a user.
6. The method of claim 1, wherein the one or more recovery point backup operations are performed in response to a change in the data.

7. The method of claim 1, wherein associating the one or more recovery point backup operations with at least one of the one or more full backup operations includes associating a pass identifier with each of the one or more recovery point backup operations.
8. The method of claim 7, wherein the pass identifier indicates one of the full backup operations corresponding to the one or more recovery point backup operations.
9. The method of claim 7, wherein the pass identifier is utilized to allow access to the backup data.
10. The method of claim 1, further comprising performing a recovery operation after the performing of the one or more recovery point backup operations.
11. The method of claim 10, wherein the performance of the recovery operation enables access to the backup data by recovering the backup data.
12. The method of claim 10, wherein the recovery operation is a point in time recovery operation.
13. The method of claim 10, wherein the recovery operation allows access to a backup state.
14. The method of claim 1, wherein the backup data to be accessed is defined by a user.
15. The method of claim 1, wherein allowing access to the backup data includes allowing access to first backup data stored before second backup data.
16. The method of claim 15, wherein the first backup data is associated with a first full backup operation of the one or more full backup operations and the second backup

data is associated with a second full backup operation of the one or more full backup operations.

17. The method of claim 1, wherein allowing access to the backup data includes recovering the backup data.

18. The method of claim 1, wherein allowing access to the backup data includes transferring the backup data from a backup storage location.

19. The method of claim 1, wherein allowing access to the backup data includes allowing the backup data to be read and modified.

20. The method of claim 1, wherein a number of the recovery point backup operations to be performed is defined by a user.

21. The method of claim 1, wherein each of the one or more recovery point backup operations is performed after an associated full backup operation.

22. A computer program product embodied on a computer readable medium, comprising:

- computer code for identifying data to backup;
- computer code for performing one or more full backup operations on the data;
- computer code for performing one or more recovery point backup operations on a new portion of the data;
- computer code for associating the one or more recovery point backup operations with at least one of the one or more full backup operations;
- computer code for allowing access to backup data resulting from performance of the one or more recovery point backup operations and the one or more full backup operations; and
- computer code storing the backup data in a data structure.

23. A system, comprising:

one or more processors configured to identify data to backup, perform one or more full backup operations on the data, perform one or more recovery point backup operations on a new portion of the data, associate the one or more recovery point backup operations with at least one of the one or more full backup operations, and allow access to backup data resulting from performance of the one or more recovery point backup operations and the one or more full backup operations.

24. A method, comprising:


generating an association between one or more recovery point backup operations and at least one full backup operation;

storing the association information in a file;


retrieving backup data corresponding to a user selected backup point, the backup data being retrieved based on the association information saved in a file; and

restoring a storage system in accordance with the user selected backup point by transferring the backup data from a backup storage location to the Exchange live system.


1/15

**FIGURE 1**


2/15

**FIGURE 2**

3/15

**FIGURE 3A**

4/15

**FIGURE 3B**

5/15




FIGURE 3C





FIGURE 4

7/15

500

FIGURE 5



8/15

600



SnapManager for Exchange - Backup

Frequent Recovery Point Backup

Existing Job:

Run Every

Job Details:

Storage Groups selected for backup:

Selected Exchange Server: SNAPMGR-32-2K3

☒ First Storage Group

☒ SG2

Select All

Unselect All

Backup Frequency

Every

15

minute(s)

Start at: 12:00:00 AM

☐ Disable scheduled job

Create Recovery Point

Create Job

Delete Job

Help

Cancel

Operation options

☐ Run command after operation

☒ Update SnapMirror after operation

FIGURE 6

9/15

700



SnapManager for Exchange - Backup

Frequent Recovery Point Backup

Existing Job: Run Every

Job Details:

Storage Groups selected for backup:

☒ First Storage Group
 ☒ SG2

Backup Frequency

Start at:

☒ Disable scheduled job

Operation options

☐ Run command after operation
 ☐ Update SnapMirror after operation

Create Recovery Point

Update Job

Delete Job

Help

Cancel

FIGURE 7

10/15

800



Schedule...

Schedule Job Name:
FRPBackup

☒ Replace job if it exists

new-backup -Server 'SNAPMGR-32-2K3' -NoTruncateLogs \$False -StorageGroup 'First Storage Group' 'SG2' -RecoveryPoint -UpdateMirror -BackupCopyRemoteCCRNNode \$False

Run As:

SNAPMGR\administrator

Browse...

Password:

Confirm password:

OK

Cancel

Help

FIGURE 8


SnapManager for Exchange - Restore



Selected Storage Group:

SG2

Selected backup to restore from:

SG2, exchsnap__SNAPMGR-50_05-08-2008_19.16.04, run verification

Choose databases to restore:

	Name	EDB Path
☒	sg2-db1	K:\Program Files\Microsoft\Exchange Server\Mailbox\SG2\sg2-db1.edb
☒	sg2-db2	K:\Program Files\Microsoft\Exchange Server\Mailbox\SG2\sg2-db2.edb

Type of Restore

☐ Up-to the minute restore
☒ Point in time restore

Select Recovery Point

☒ Recover and mount databases after restore

Backup Verification Server:

snapmgr-50

Advanced Options...

Restore...

Test Restore...

Cancel

FIGURE 9

1000

SnapManager for Exchange - Backup



Backup to restore from:
SG2, exchsnap__SNAPMGR-50_05-08-2008_19.08.24, run verification

Selected Exchange Server:
SNAPMGR-50

You can select the date and start time and end time within that date. The available recovery point will be displayed once you click on the Refresh button.

Scope for Recovery Point Selection

Date: 5/ 8/2008 ▼ Start Time: 7:08:24 PM ▲ End Time: 11:59:24 PM ▲

Available Recovery Points

05-08-2008_19.08.58 -- E0100000184
05-08-2008_19.09.03 -- E0100000185
05-08-2008_19.16.53 -- E0100000186
05-08-2008_19.16.59 -- E0100000187

To select a recovery point, please highlight one in the list and click on OK.

Refresh

OK Cancel

FIGURE 10

13/15





FIGURE 11

14/15

**FIGURE 12**

15/15

1300





FIGURE 13

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 10/00347

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 7/00 (2010.01)

USPC - 707/674

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

USPC: 707/674

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
USPC: 707/640, 649, 650, 674, 676, 678; 711/161, 162; 714/15, 20 keyword limited - see search terms below

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

PubWest (PGPB,USPT,USOC,EPAB,JPAB), Google Scholar, Google Patent

Search terms: backup full baseline incremental delta differential changed log history transaction record database periodic frequency regular interval state save include point time data system machine server location site identifier association previous prior entry truncate

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2004/0030852 A1 (COOMBS et al.) 12 February 2004 (12.02.2004), para [0010], [0014]-[0015], [0017]-[0018], [0025], [0028], [0032], [0035]-[0036], [0062], [0082]-[0084]	1-24
A	US 2005/0223277 A1 (BALLARD) 06 October 2005 (06.10.2005),	1-24
A	US 6,785,786 B1 (GOLD et al.) 31 August 2004 (31.08.2004),	1-24

☐ Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

15 May 2010 (15.05.2010)

Date of mailing of the international search report

09 JUN 2010

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774