

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
26 April 2007 (26.04.2007)

PCT

(10) International Publication Number  
**WO 2007/047346 A2**

(51) International Patent Classification:  
**G06F 11/00** (2006.01) **G06F 12/16** (2006.01)

(21) International Application Number:  
PCT/US2006/039857

(22) International Filing Date: 13 October 2006 (13.10.2006)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:  
60/726,187 14 October 2005 (14.10.2005) US

(71) Applicant (for all designated States except US): **REVIVIO, INC.** [US/US]; 12 Hartwell Avenue, Lexington, MA 02421 (US).

(72) Inventors: **PASSERINI, Ronald, Peter**; 41 Princeton Street, Somerville, MA 02144 (US). **PERRY, Robert, Warren**; 14 Pine Grove Avenue, Leominster, MA 01453 (US). **ROCCA, Christopher, Angelo**; 13 Westwood Street, Burlington, MA 01803 (US). **ANTHONY, Michael, Daniel**; 9 Elm Street, Wilmington, MA 01887 (US).

(74) Agents: **ANDERSON, Thomas, E.** et al.; INTELLECTUAL PROPERTY DEPARTMENT, HUNTON & WILLIAMS LLP, 1900 K Street, N.W., Suite 1200, Washington, DC 20006-1109 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: TECHNIQUE FOR TIMELINE COMPRESSION IN A DATA STORE



(57) Abstract: A technique for timeline compression in a data store is disclosed. In one particular exemplary embodiment, the technique may be realized as a method for timeline compression in a storage system, wherein digital content of the storage system is backed up to enable restoration of the digital content to one or more points in a timeline. The method may comprise selecting a time interval in the timeline. The method may also comprise identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more sets of backup data represent at least a portion of old data overwritten during the selected time interval. The method may further comprise discarding other backup data recorded for the selected time interval, thereby reducing a granularity level of the timeline in the selected time interval.



WO 2007/047346 A2

**TECHNIQUE FOR TIMELINE COMPRESSION IN A DATA STORE****CROSS-REFERENCE TO RELATED APPLICATIONS**

This patent application claims priority to U.S.  
5 Provisional Patent Application No. 60/726,187, filed October 14, 2005, which is hereby incorporated by reference herein in its entirety.

This patent application is related to U.S. Patent Application No. 10/924,652, filed August 24, 2004, which is a  
10 continuation-in-part of U.S. Patent Application No. 10/668,833, filed September 23, 2003, each of which is hereby incorporated by reference herein in its entirety.

**FIELD OF THE DISCLOSURE**

15 The present disclosure relates generally to data storage and, more particularly, to a technique for timeline compression in a data store.

**BACKGROUND OF THE DISCLOSURE**

20 In related U.S. Patent Application No. 10/924,652 and U.S. Patent Application No. 10/668,833, a time-dependent data storage and recovery technique is disclosed. Embodiments of such a technique provide a solution for continuous data protection (CDP) wherein write commands directed to a storage  
25 system are intercepted by a storage management system having a current store and a time store. The current store may maintain or have access to a current (or mirror) copy of the storage system's digital content. The time store may record information associated with each intercepted write command,  
30 such as new data in the write command's payload or old data to be overwritten in the current store in response to the write command. Recordation of the new or old data in response to a write command may be referred to as a copy-on-write (COW) operation, and the new or old data recorded may be referred to

as COW data. The time store may also record other information (i.e., metadata) associated with an intercepted write command and/or the corresponding COW operation, such as, for example, a timestamp, an original location in the current store where  
5 the old data are overwritten, and a destination location in the time store to which the COW data are copied. Each COW operation typically backs up one or more blocks of COW data, thereby creating one set of COW data and corresponding metadata. Over a period of time, multiple sets of COW data  
10 and corresponding metadata (including timestamps) may be accumulated as a collection of historical records of what have been written or overwritten in the current store or the storage system. The content of the time store may be indexed based on the metadata to facilitate efficient access to the  
15 COW data.

With a current copy of the storage system's digital content in the current store and the historical records in the time store, the storage management system adds a new dimension, i.e., time, to the storage system. Assuming the  
20 storage management system has been operatively coupled to the storage system since a past time, the storage management system may quickly and accurately restore any addressable content in the storage system to any point in time between the past time and a present time.

25 Ideally, it might be desirable to maintain such a data recovery capability for as long a timeline as possible. However, to accommodate an extended timeline, a significant amount of storage space is needed to store the COW data and corresponding metadata for every write command in that  
30 timeline. Even more storage space is needed if the storage system sees a relatively high write rate (i.e., number of write operations per unit time). One temporary solution may be to simply increase storage capacity of the time store. However, apart from a higher cost, a simple storage increase

may not scale well with the rest of the system and tends to create a deluge of other problems, such as a performance degradation due to difficulties of parsing through an additional amount of data. Without an infinite storage capacity, most storage systems have to settle for the reality that only a finite length of timeline (e.g., ten days or two weeks) can be maintained. In conventional data protection systems, it is typical to keep a few days' worth of backup data and completely discard the backup data that are more than a few days old. In these systems, data recovery capabilities are limited to the past few days for which backup data are available. Alternatively, the backup data that are more than a few days old may be moved off site on a regular basis. Such a brute-force solution can be costly and disruptive, not to mention its slow response to data recovery requests where off-site data are needed.

In view of the foregoing, it would be desirable to provide a solution for data storage management which overcomes the above-described inadequacies and shortcomings.

#### SUMMARY OF THE DISCLOSURE

A technique for timeline compression in a data store is disclosed. In one particular exemplary embodiment, the technique may be realized as a method for timeline compression in a storage system, wherein digital content of the storage system is backed up to enable restoration of the digital content to one or more points in a timeline. The method may comprise selecting a time interval in the timeline. The method may also comprise identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more sets of backup data represent at least a portion of old data overwritten during the selected time interval. The method may further comprise discarding other backup data recorded for the selected time interval,

thereby reducing a granularity level of the timeline in the selected time interval.

5 In accordance with other aspects of this particular exemplary embodiment, the digital content of the storage system may be backed up through copy-on-write operations into a plurality of sets of copy-on-write data and corresponding metadata, and the step of identifying may further comprise identifying one or more sets of copy-on-write data and corresponding metadata recorded for the selected time  
10 interval.

In accordance with further aspects of this particular exemplary embodiment, a length of the time interval may be selected based at least in part on a desired granularity level of the timeline.

15 In accordance with additional aspects of this particular exemplary embodiment, the step of identifying may further comprise: determining whether a storage unit in the storage system has been overwritten more than once during the selected time interval; if the storage unit has been overwritten once  
20 during the selected time interval causing a sole set of copy-on-write data and corresponding metadata to be recorded, selecting the sole set; and if the storage unit has been overwritten more than once during the selected time interval causing multiple sets of copy-on-write data and corresponding  
25 metadata to be recorded, selecting one of the multiple sets.

In accordance with another aspect of this particular exemplary embodiment, if the storage unit has been overwritten more than once during the selected time interval, the selected set of copy-on-write data and corresponding metadata may be  
30 the earliest set recorded for the selected time interval.

In accordance with yet another aspect of this particular exemplary embodiment, the method may further comprise coalescing metadata in the one or more identified sets of copy-on-write data and corresponding metadata.

In accordance with still another aspect of this particular exemplary embodiment, the method may further comprise: identifying copy-on-write data that correspond to the coalesced metadata; and replacing all sets of copy-on-write data and corresponding metadata previously recorded for the selected time interval with a new set comprising the identified copy-on-write data and the coalesced metadata.

In accordance with a further aspect of this particular exemplary embodiment, the method may further comprise: selecting multiple time intervals in a portion of the timeline based on a desired granularity level for the portion of the timeline; and repeating the steps of identifying and discarding for the selected multiple time intervals.

In accordance with a yet further aspect of this particular exemplary embodiment, the storage system may comprise a plurality of storage devices and the method may further comprise: repeating the steps of identifying and discarding for one or more of the plurality of storage devices to cause the plurality of storage devices to have a consistent granularity level of the timeline with respect to one another.

In accordance with a still further aspect of this particular exemplary embodiment, the steps of selecting, identifying and discarding may be triggered when one or more of the following conditions are met: a predetermined storage capacity for the timeline has been consumed; a predetermined amount of data have been accumulated for a granularity level of the timeline; granularity levels of the timeline for at least two storage devices in the storage system are inconsistent; an instruction to reduce the granularity of the timeline is received; and a scheduled time for reducing the granularity of the timeline is reached.

In accordance with another aspect of this particular exemplary embodiment, the method may further comprise scanning the storage system for a storage device for which the

granularity of the timeline can be reduced.

In another particular exemplary embodiment, the techniques may be realized as at least one signal embodied in at least one carrier wave for transmitting a computer program of instructions configured to be readable by at least one processor for instructing the at least one processor to execute a computer process for performing the method as recited above.

In yet another particular exemplary embodiment, the techniques may be realized as at least one processor readable carrier for storing a computer program of instructions configured to be readable by at least one processor for instructing the at least one processor to execute a computer process for performing the method as recited above.

In still another particular exemplary embodiment, the techniques may be realized as a system for timeline compression in a storage system, wherein digital content of the storage system is backed up to enable restoration of the digital content to one or more points in a timeline. The system may comprise means for selecting a time interval in the timeline. The system may also comprise means for identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more sets of backup data represent at least a portion of old data overwritten during the selected time interval. The system may further comprise means for discarding other backup data recorded for the selected time interval, thereby reducing a granularity level of the timeline in the selected time interval.

In a further particular exemplary embodiment, the techniques may be realized as a system for timeline compression in a storage system, wherein digital content of the storage system is backed up to enable restoration of the digital content to one or more points in a timeline. The system may comprise a storage medium for storing instructions.

The system may also comprise at least one processor for: selecting a time interval in the timeline; identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more sets of backup data represent at least a portion of old data overwritten during the selected time interval; and discarding other backup data recorded for the selected time interval, thereby reducing a granularity level of the timeline in the selected time interval.

The present disclosure will now be described in more detail with reference to exemplary embodiments thereof as shown in the accompanying drawings. While the present disclosure is described below with reference to exemplary embodiments, it should be understood that the present disclosure is not limited thereto. Those of ordinary skill in the art having access to the teachings herein will recognize additional implementations, modifications, and embodiments, as well as other fields of use, which are within the scope of the present disclosure as described herein, and with respect to which the present disclosure may be of significant utility.

#### BRIEF DESCRIPTION OF THE DRAWINGS

In order to facilitate a fuller understanding of the present disclosure, reference is now made to the accompanying drawings, in which like elements are referenced with like numerals. These drawings should not be construed as limiting the present disclosure, but are intended to be exemplary only.

Figure 1a shows a timeline maintained for a storage system based on a traditional method.

Figure 1b shows a timeline maintained for a storage system in accordance with an embodiment of the present disclosure.

Figure 2 shows a flow chart illustrating an exemplary timeline compression method in accordance with an embodiment



of the present disclosure.

Figure 3 shows a state diagram illustrating an exemplary method for timeline compression in accordance with an embodiment of the present disclosure.

5        Figure 4 shows an exemplary timeline for three related LUs in accordance with an embodiment of the present disclosure.

10        Figure 5 shows another exemplary timeline for three related LUs in accordance with an embodiment of the present disclosure.

Figure 6 shows major objects involved in an exemplary program for timeline compression in accordance with embodiments of the present disclosure.

#### 15                    **DETAILED DESCRIPTION OF EXEMPLARY EMBODIMENTS**

As used herein, "backup data" refers generally to data that have been recorded and/or organized (or even re-organized) with a purpose of restoring or recovering digital content of a storage system.

20        "Copy-on-write data" (or "COW data") refers to substantive data (e.g., new data to be written or old data to be overwritten in response to a write command) that have been recorded in a copy-on-write operation. New data to be written in response to a write command are sometimes referred to as  
25        "after image data" or "after image," while old data to be overwritten in response to a write command are sometimes referred to as "before image data" or "before image." The copy-on-write operation may be an actual operation performed in response to an actual write command. Or, the copy-on-write  
30        operation may be a virtual operation that includes the collective effect of multiple copy-on-write operations that occur during a selected time interval.

"Corresponding metadata" refers to informational data (e.g., timestamps) regarding the associated COW data in a

copy-on-write operation. Typically, one copy-on-write operation causes one set of COW data and corresponding metadata to be created. Despite their correlation, COW data and corresponding metadata may be stored in separate storage devices or segments. In a time store, COW data may be organized in one or more timestamped "data chunks."

"Raw data" refers to one or more sets of COW data and corresponding metadata that have been recorded in response to actual write commands and have not been coalesced or otherwise modified since their recordation. In some circumstances, "COW data" and "corresponding metadata" may refer to COW data and corresponding metadata, respectively, that have been coalesced, re-organized or otherwise modified in a timeline compression process, wherein a resulting set of COW data and corresponding metadata may be considered as originating from a virtual copy-on-write operation in response to one or more write commands during a selected time interval. In other words, "COW data" and "corresponding metadata" may sometimes refer to backup data that are not on the raw data level.

"Granularity level" of a timeline refers to a time scale (e.g., weekly, daily, hourly, by the second, or by the millisecond) with which digital content of a storage system can be restored to a point in the timeline. The granularity level of a timeline is typically determined by the specific mechanism employed to back up digital content, how complete backup data are kept, and how the backup data are organized.

A typical "storage system" may comprise one or more storage devices which may be physical, virtual or logical devices or a combination thereof. According to one embodiment, a storage system may comprise a storage area network (SAN) having one or more datasets, wherein each dataset may comprise one or more nodes, and wherein one or more logical units (LUs) may be coupled to each node. Hereinafter, for ease of illustration, the term "storage

system" may refer to an entire storage system or a portion (e.g., dataset or node) thereof. Typically, a timeline may be maintained for all LUs in a same dataset.

"Timeline storage" refers to a storage space for backup data in a time store. Timeline storage is typically organized in terms of quota groups, wherein each quota group allocates a predetermined storage space for a timeline associated with a corresponding dataset.

Embodiments of the present disclosure provide a technique known as "timeline compression" that allows a more extended timeline to be maintained for a storage system (or a dataset) without any substantial increase in timeline storage capacity or complete discarding of older backup data. This may be achieved by selectively decreasing a granularity level of the timeline as backup data are aging. One or more older portions of raw data backed up for a storage system may be coalesced and/or re-organized into one or more data chunks that reflect write operations in the storage system on a coarser level of granularity (e.g., hourly or daily) than the raw data normally would reflect. Such reduction in the granularity level of the timeline may offer a flexible, user-definable tradeoff wherein the timeline storage is economized without sacrificing older data entirely. As a result, a much longer timeline may be maintained for a storage system without any significant impact on its data protection or data recovery capabilities. The coalescence and/or re-organization process of backup data from one granularity level to another may be referred to as a "timeline rollup" or "rollup."

In the detailed description that follows, references will be made to embodiments of the time-dependent data storage and recovery technique as disclosed in U.S. Patent Application No. 10/924,652 and U.S. Patent Application No. 10/668,833. It should be appreciated that embodiments of the present disclosure are easily adaptable to other data protection

methods or systems that maintain historical records of digital content of a storage system.

Referring to Figure 1a, there is shown a timeline maintained for a storage system based on a traditional method.

5 In this example, present day may be Monday of Week 0. The timeline may have been continuously maintained for a storage system for a few weeks (i.e., Week -1, Week -2, Week -3, Week -4 and so on). If several weeks' worth of backup data were all stored in the form of raw data, a large amount of storage  
10 space in a time store would be required. If, for example, there is only enough storage space to store fourteen days' worth of raw data, then, according to traditional approaches, those raw data that are more than fourteen days old must be completely discarded. That is, by the beginning of the  
15 present day (i.e., Monday of Week 0), all raw data recorded for Week -3 and earlier may have already been discarded. As the present day goes on and new raw data are accumulated, raw data recorded for Monday of Week -2 may have to be sacrificed in order to make room for the newly recorded raw data.  
20 Therefore, at any given time, digital content of the storage system is backed up only for the previous two weeks, while no historical record is available beyond that two-week period.

Figure 1b shows a timeline maintained for a storage system in accordance with an embodiment of the present  
25 disclosure. This timeline may be generated by subjecting raw data to a timeline compression process that selectively reduces a granularity level of the timeline. It is recognized that, as time goes by, the oldest backup data are the least likely to be needed on the finest granularity level.  
30 Therefore, it might suffice to keep only a few days' worth of raw data in order to be able to restore the storage system to any point in time in the past few days. For older backup data, the granularity level of the timeline may be progressively reduced. As shown in Figure 1b, on Monday of

Week 0, for example, backup data for the past three days (i.e., Friday, Saturday and Sunday of Week -1) are kept in the form of raw data. In a second time period, that is, more than three days ago and up to the beginning of Week -2, for example, the backup data may be kept in the form of hourly data. That is, the original raw data may be selectively coalesced and/or discarded, as will be described in more detail below, such that only enough backup data is kept to be able to restore digital content of the storage system to any hour during the second time period. In a third time period that spans Week -3 and Week -4, for example, the backup data may be kept in the form of daily data. Prior to the beginning of Week -4, for example, only weekly data may be kept available. As a result of this exemplary timeline compression scheme, it may be possible to maintain a much longer timeline than the traditional approach illustrated in Figure 1a. Even if there is only enough storage space for fourteen days' worth of raw data, a timeline much longer than two weeks, maybe a few months, may be maintained. As a result, digital content of the storage system that is several weeks old may still be restored. The only tradeoff is that such restoration of the older content may be available on a coarser level of granularity than what raw data can facilitate. For example, instead of being able to pick and choose a data recovery point by the second or by the millisecond, a user may only be able to select recovery points on an hourly, daily or weekly scale.

A timeline compression functionality in accordance with embodiments of the present disclosure may be implemented in any type of storage systems, preferably in connection with a storage management system having a current store and a time store. A set of parameters, known as Timeline Lifecycle Profile (TLP), may be configured by a user to control the timeline compression process. According to one embodiment, the TLP may specify four levels of backup data and a user

configurable amount of time to keep each level of backup data before that level of backup data may be rolled up to a next level. The TLP may also specify a minimum amount of timeline storage used before a rollup takes place.

5 Referring again to Figure 1b, the timeline illustrated therein reflects one exemplary TLP relating to four levels of backup data. Level-0 TLP may specify conditions that must be fulfilled before raw data can be rolled up to the next level. The conditions may be defined as an amount of raw data, in  
10 terms of time length and/or timeline storage capacity, to keep before some raw data may be rolled up to the next level. For example, the Level-0 TLP may require that at least 3 days' worth of raw data using 40% of the timeline storage capacity be accumulated before raw data older than 3 days or beyond the  
15 40% storage limit may be rolled up to the next level. According to this configuration, even if there are more than 3 days of raw data, the raw data older than 3 days will not be rolled up until the raw data have also used up at least 40% of the timeline storage capacity. This configuration may also  
20 require that, even if 40% of the timeline storage capacity is occupied by raw data, a rollup is not performed unless more than 3 days' worth of raw data have been accumulated. In the Level-0 TLP, a default value for the time length may be infinite, which means that all other levels may be ignored and  
25 there will be no rollup of Level-0 raw data. A default limit for timeline storage capacity may be 0%, which means the amount of timeline storage in use will not be considered in determining whether to initiate a rollup of the raw data.

Level-1 TLP may specify a reduced granularity level  
30 (e.g., hourly) of the timeline compared with Level-0, as well as conditions that must be met before Level-1 data (e.g., hourly data) can be rolled up to the next level. Similar to Level-0 TLP, the conditions may be defined as an amount of Level-1 data, in terms of time length and/or timeline storage

capacity, to keep before Level-1 data may be rolled up to the next level (e.g., as shown in Figure 1b, 11 days' worth of hourly data using 20% of the timeline storage capacity). In the Level-1 TLP, a default value for the time length may be  
5 infinite, which means that all other levels may be ignored and rollups do not continue beyond Level-1 data. A default limit for timeline storage capacity may be 0%, which means the amount of timeline storage in use will not be considered in determining whether to initiate a rollup of the Level-1 data.

10 Level-2 TLP may specify a further reduced granularity level (e.g., daily) of the timeline as well as conditions that trigger a rollup of the Level-2 data. For example, the Level-2 TLP may require that at least 14 days' worth of daily data using 20% of the timeline storage capacity be accumulated  
15 before the older daily data may be rolled up to the next level.

Level-3 TLP may specify an even further reduced granularity level (e.g., weekly) of the timeline as well as conditions that trigger a rollup of the Level-3 data to the  
20 next level (e.g., monthly data). For example, the Level-3 TLP may require that at least 12 weeks' worth of weekly data using 10% of the timeline storage capacity be accumulated before the older weekly data may be rolled up to the next level.

According to embodiments of the present disclosure, a  
25 user typically does not explicitly create a TLP since a default TLP may already exist when a quota group is created in a storage system. Upon creation of the quota group, the user may have the option of modifying the default parameters of the TLP. Similarly, a user typically does not explicitly delete a  
30 TLP. The TLP may be deleted when the associated quota group is deleted. A user may modify a TLP in order to change the desired behavior of timeline compression. The TLP may be modified at any time without any immediate effect on the timeline. The default TLP, upon creation of a quota group,

may specify that no rollups take place. A user may also modify a TLP when modifying attributes of a quota group. When an LU is added to a dataset, the LU may inherit an existing TLP for the corresponding quota group. The effect on the timeline may be that the start of the timeline may shift forward. This behavior may be the same as when there are no rollups defined. However, rollups may continue at the current rollup level for the entire dataset, and the new LU's current rollup state may be set to reflect that of the rest of the LUs in the dataset. When an LU is removed from a dataset, the start of remaining timeline may shift backward. This may be the same behavior exhibited when there are no rollups defined.

Referring to Figure 2, there is shown a flow chart illustrating an exemplary timeline compression method in accordance with an embodiment of the present disclosure.

In step 202, a rollup may be started for a storage device in a storage system. The method steps in Figure 2 illustrate a simplest rollup operation of backup data recorded for a portion of a timeline from one level to a next level, wherein it is assumed that the timeline is maintained for a particular storage device only. Timeline compression involving multiple storage devices will be described separately below.

In step 204, a time interval may be selected for a rollup operation. Selection of the time interval is typically based on conditions specified in the TLP. For example, if it is a rollup from raw data to hourly data, the time interval may be one hour long and selected from a portion of the timeline where hourly data are desired. This portion of the timeline may span multiple hours. Thus, the method steps 204 through 212 may be repeated for every hour in this portion of the timeline. Similarly, if it is a rollup from hourly data to daily data, the time interval may be 24 hours long and selected from a portion of the timeline where daily data are desired.



In step 206, COW data and corresponding metadata may be identified to represent backup data recorded for the selected time interval. For example, a unit of storage (e.g., a block) in the storage device may have been overwritten one or more times during the selected time interval. If a block has been overwritten only once during the selected time interval, the resultant set of COW data and corresponding metadata may be identified to represent backup data for this block. If a block has been overwritten multiple times during the selected time interval, a set of COW data and corresponding metadata that results from the earliest write operation may be identified to represent backup data for this block. Alternatively, a set of COW data and corresponding metadata that results from the latest write operation during the selected time interval may be identified. Additional or alternative criteria may be used to identify the representative backup data for the storage device during the selected time interval.

In step 208, other backup data that have been recorded for the selected time interval but are not selected in step 206 may be discarded or simply ignored. For example, metadata or other indexing data for the unselected backup data may be deleted or erased such that the unselected backup data are effectively deleted from the timeline.

In step 210, the COW data and corresponding metadata identified or selected in step 206 may be coalesced. According to one embodiment, these COW data and/or corresponding metadata are preferably coalesced into fixed size allocation units known as "buckets." For example, when rolling up raw data into hourly data, one hour's worth of raw data may be coalesced into one 512 KB fixed-size hourly bucket. Later on, 24 hourly buckets of backup data may be coalesced into one 12 MB fixed-size daily bucket. Coalescence of the selected COW data and corresponding metadata may be

carried out in a number of ways. According to one embodiment, the metadata for the selected time interval may be coalesced first. Then, the COW data (pointed to by the coalesced metadata) may be coalesced and copied into memory. The  
5 coalesced COW data may then be stored as new, higher level COW data for the selected time interval. Accordingly, the coalesced metadata may be modified and stored as corresponding metadata for the new COW data.

In step 212, it may be determined whether the rollup has  
10 been done for this storage device. If there are additional time intervals to roll up, the process may loop back to step 204 to repeat the coalescence of backup data until the rollup for all appropriate time intervals has been completed according to the requirements specified in the TLP. Then, in  
15 step 214, the rollup may end for this storage device. A timeline rollup operation may be configured (e.g., as a thread or sub-routine) to start or restart upon initiation by a user or upon triggering of one or more events defined in the TLP. For example, the rollup thread may be activated whenever a  
20 pre-determined percentage of the timeline storage has been used, before or after a deprecation of timeline, and/or on a periodic basis. A rollup thread may also be awoken when a second rollup thread on another node determines that it must rollup data for an LU on that node.

Referring to Figure 3, there is shown a state diagram  
25 illustrating an exemplary method for timeline compression in accordance with an embodiment of the present disclosure. The state diagram shows five phases in a timeline rollup operation. As mentioned above, the rollup operation may be  
30 preferably a thread running on or in coordination with a storage management system. Once started, the rollup thread may coalesce and/or re-organize backup data for a storage system by going through a Determination Phase, a Scan Phase, a Coalescence Phase, a Copy Phase, and a Replacement Phase.

In the Determination Phase, it may be determined what work needs to be done, such as, for example, what level of backup data to perform the rollup for and for which LU. A typical environment for implementation of a timeline compression may be a storage management system coupled to a storage system having a plurality of storage devices (e.g., LUs). One typical requirement for timeline compression involving multiple LUs is consistency of granularity levels of the timeline across all LUs. Accordingly, there may be a need to determine which level of backup data to roll up and which LU to start with. The level of backup data to perform the rollup for, also known as a "rollup level," may be determined based on rollup state data (described in detail below) that have been maintained for the LUs. In order to determine which LU to perform the rollup for, the LUs may be examined one at a time, typically by analyzing the rollup state data and/or other informational data associated with the LUs. For each LU, the oldest backup data in the timeline may be identified for the current rollup level. If, for example, the LU in question is owned by a local node for which the rollup thread is running, and if there are enough valid, sealed backup data for the time range in question (e.g., at least two data chunks), then the rollup operation may proceed to the next phase (i.e., Scan Phase). Otherwise, another LU may be chosen and analyzed. If no LU is found that meets the criteria defined in the TLP, then the rollup operation may end. If an LU is found that meets the criteria to perform a rollup, then all related LUs (e.g., those associated with the same quota group) may be rolled up as well. If it is determined that a rollup must occur, the rollup thread on other nodes may be awoken using a remote message.

In the Scan Phase, for backup data that can be rolled up, a scan of the metadata may be performed in order to identify the time period defined by the TLP and to retrieve relevant

metadata. The scan may be a query of the metadata on a slice by slice basis in time order. According to embodiments of the present disclosure, each LU may be divided into a plurality of fixed-size logical partitions (e.g., 16 Gigabytes (GB) each) for ease of management and for load balancing purposes, wherein each fixed-size logical partition may be referred to as one "slice." When a slice has been scanned, the identified metadata may be coalesced in the Coalescence Phase, after which a next slice may be scanned in the Scan Phase. The slice by slice cycle may repeat until all slices have been exhausted for the LU and time period in question.

In the Coalescence Phase, the metadata that have been identified in the Scan Phase may be coalesced into fixed size allocation units. Coalescence of the metadata may be achieved in any of a variety of ways. According to one embodiment, the identified metadata may first be stored in memory in a time-ascending order according to their timestamps. Then, starting from the oldest metadata, the identified metadata may be inserted into a binary tree (B-tree) that is indexed by starting Logical Block Addresses (LBA's) recorded in the metadata. The B-tree may be first scanned for LBA overlap before the coalesced metadata are inserted. The newer COW data containing LBA overlaps may be discarded. Splits may also be done on COW data where needed to achieve coalescence. For example, for a particular time interval, if metadata for Blocks 1-5 are already stored in the B-tree, then another metadata row associated with Blocks 3-8 may be split into two portions, one associated with Blocks 3-5 and the other associated with Blocks 6-8, where the former may be discarded while the latter may be entered into the B-tree. The resulting metadata may be coalesced into relatively large fixed sized allocation units (e.g., 512 KB), each allocation unit representing a single COW operation. Statistics kept as records may be coalesced and inserted into the B-tree. These

statistics may be kept on an LU basis so that a determination can be made in the Copy Phase as to whether to actually copy the COW data into a rollup chunk or simply modify the corresponding (coalesced) metadata. The statistics kept may include, for example, the number of original blocks, the number of duplicate blocks (block savings), the number of original COW operations, and the number of coalesced COW operations (metadata savings).

Each resulting metadata row may be that of the earliest original operation in a particular slice. For example, if Block 128 was written twice in a row, once at time T1 and again at time T2, the resulting coalesced metadata may only reflect the write operation at time T1. Once the metadata for a slice has been coalesced, the next slice may be processed for the same LU in the Scan Phase. This procedure continues until all slices for the LU have been processed, whereupon the Copy Phase may be executed.

In the Copy Phase, the coalesced metadata may be read from memory, and the COW data the coalesced metadata point to may be copied into rollup chunks. Whether or not to actually copy the COW data may depend on an evaluation of the statistics collected in the Coalescence Phase. General, an actual copy of the COW data is performed only when there is some saving of storage space.

After the evaluation has passed and all of the metadata have been coalesced for an LU, the COW data may be copied into resulting rollup chunks on a slice by slice basis. The COW data may be copied from a time store (in the storage management system) to the same time store using pre-allocated non-replicated buffers. In order to accomplish the data copy, an event chain may be created and a RollupTimeStoreMove event may be pushed onto the event chain. The RollupTimeStoreMove object may be derived from an SGIO and a WaitableEvent object and may be handled by an IO execution context. The event may

have two extents, one for the read of the original COW data and one for the write to the rollup chunk in the time store. The Copy Phase may wait for the event to complete and then update an in-memory array with a data structure representing a new indexing operation. This array may be used in the Replacement Phase to update the coalesced metadata. Once all of the COW data have been copied, the rollup thread may proceed to the Replacement Phase.

In the Replacement Phase, the relevant portion of old COW data may be replaced with the newly created rollup chunks, the corresponding metadata may be replaced or modified to reflect the new chunks and coalesced COW data contained therein, and the rollup state may be updated. The Replacement Phase is typically carried out atomically and may be rolled back should a catastrophic failure event occur. The LU whose metadata are being updated may be locked to keep configuration from changing and to prevent timeline deprecation from occurring during this phase. The old COW data coalesced in the previous phases may be freed, and the new rollup chunks may be added to the appropriate place in the timeline. The original metadata may be deleted and replaced by the new, coalesced metadata. Finally, the LU may be unlocked, and a next LU that meets the Determination Phase criteria may be processed.

According to one embodiment, the old COW data may be transformed into new rollup chunks without an actual copy taking place. The original data chunks may simply be updated to reflect the rollup timestamps, and their rollup level may be updated accordingly.

The rollup operation runs through a number of phases one of which involves copying COW data to newly rolled-up chunks, which may take a significant amount of time. As a result, other critical components of the system that interfere with rollups or change the state of the timeline, such as Instant Restore, Time Images, or Timeline Deprecation, may cause an

in-progress rollup to be canceled. If a rollup is canceled during processing of a particular level, all other levels may be cancelled as well. Therefore, in the state diagram shown in Figure 3, there is a Cancel route from almost every phase to End. The Replacement Phase may not be arbitrarily cancelled and may proceed while the LU in question is locked. To avoid excessive interruption of the rollups, the timeline compression process may be preferably scheduled to start when resources in the storage system and/or the storage management system are in low demand.

It may be desirable to track the state of rollups that have occurred on an LU by LU basis. This way, other components in the storage system can determine what backup data have been rolled up, and what LUs have not yet had their backup data rolled up. For example, the determination of a common timeline start across multiple LUs requires rollup state information associated with the LUs of interest. The state of rollups may also be needed for future rollups in order to determine where to start a rollup.

The rollup state may be stored in a database table in a global database called "rollup\_state," one example of which is shown in Table 1.

Table 1: Rollup State Table

Field	Type	Description
LU	UInt32	Logical Unit (unique index)
GenerationNumber	UInt32	Counter used to uniquely identify the rollup.
LastRollupLevel	UInt32	The last rollup level that completed for this LU.
LastRollupRangeStart	UInt64	The last start range time a rollup ran for this LU.
LastRollupRangeEnd	UInt64	The last end range time a rollup ran for this LU.

Each entry in Table 1 may have a generation number that uniquely identifies the rollup level for each LU. The generation number may be updated whenever a rollup is completed for an LU.

5 In addition to rollup state on an LU, each LU entry in the global database may have a field describing the last rollup level for it. This may be used to determine the timeline for an LU within a rollup level.

10 There may be cases where multiple rollups for the same rollup level and LU may exist in a same rollup chunk. In order to return valid image times for this case, a history of rollups that have taken place may be kept. The start and end time ranges for each rollup may be stored in a table called "rollup\_history," one example of which is shown as Table 2.

15

Table 2: Rollup History Table.

Field	Type	Description
LU	Uint32	Logical Unit
RollupLevel	Uint32	The rollup level that completed for this LU.
RollupRangeStart	Uint64	The start range time this rollup ran for this LU.
RollupRangeEnd	Uint64	The end range time this rollup ran for this LU.

20 This information may be used to determine valid image times for an LU and rollup level. Once a rollup level is rolled up (e.g. from Level-1 to Level-2) or a rollup level has been deprecated, the entry representing that particular rollup may be deleted.

25 Since an LU is typically related to or grouped with other LUs, the result of a rollup on one LU may affect how the timeline for a group of LUs is represented. When a timeline rollup involves a group of LUs (e.g., in a dataset or quota group), the start of a timeline for the group of LUs may be



the earliest common timeline start (ECTS) among all the LUs. However, when there are multiple levels of rollup data, there may be points in time in the timeline that become invalid for the start of an image time. Figure 4 illustrates two different rollup levels across three related LUs. Each tick on the timelines may represent a data chunk in a time store. The start of the timeline for hourly data may be time t3 since it is the ECTS among the three LUs. In this example, valid image times are t3, t4, and t5 for the hourly data. The times t0, t1, and t2 can only be chosen through dissociation with the other related LUs. Since times t3, t4, and t5 are hourly buckets of backup data and are essentially single-point-in-time (SPIT) images, image times chosen in between these times may be invalid. The Level-0 data in this example runs from time t6 through a present time. Any Level-0 time between t6 and the present time may be valid image times.

When rollup levels among multiple LUs are inconsistent, the valid image times may differ from those when the rollup levels are consistent among the LUs. If, in the above example, Level-1 rollup is only completed for LU-1, then points t3 and t4 may not be valid image times. Figure 5 illustrates this scenario. Although the ECTS remains the same, t5 may be the only valid image time across all LUs for Level-1 data. Any image time between t6 and the present time may still be valid.

According to embodiments of the present disclosure, a few primitive application program interfaces (APIs) may be provided to represent a timeline that accommodates transient rollup inconsistencies. For example, primitive APIs may be provided to perform the following functions: (a) retrieve start and end of timeline for an LU across all rollup levels, wherein the start and end time range, as well as the generation number for the last rollup, may be returned for the LU; (b) retrieve start and end time of a timeline for an LU

within each rollup level for an LU (e.g. start and end of Level-1 data for LU-1), wherein the start and end time range, as well as the generation number for the last rollup, may be returned for Level-0, Level-1, Level-2 and Level-3 data; (c)

5 retrieve the next and previous image times based on an input image time for an LU, wherein the start time before and the start time after the image time may be returned for each of the two times. Using these primitive APIs, a user or client program that has knowledge about the relationship among the  
10 LUs may query the information for each LU and determine the start and end of the timeline for each level.

In the exemplary timeline shown in Figure 4, the primitive APIs may be used to find the start and end of the Level-1 timeline for the related LUs as follows.

15 (1) The ECTS may be found for the set of LUs based on what the latest timeline start is for the set of LUs, in this case time t3.

(2) The Level-1 data with earliest start time past the ECTS, in this case t5, may be the start of the Level-1  
20 timeline.

(3) The Level-1 data with the latest end time past the ECTS, in this case t5, may be the end of the Level-1 timeline.

(4) All image times chosen between the start and end  
25 times may be validated by using API(c) above for each LU. Finding the closest time before and after the image time may yield up to two valid data points for an image time.

Similarly, the primitive APIs may be used to find the start of the Level-0 timeline for the related LUs in Figure 4.

30 (1) The ECTS may be found for the set of LUs based on what the latest timeline start is for the set of LUs, in this case time t3.

(2) The Level-0 data with the earliest start time past the latest end time of the Level-1 data, in this case

time t6, may be the start of the Level-0 timeline.

(3) Any image time from this point forward may be valid for Level-0 data.

For a Time Image or Instant Restore operation, an image  
5 time may be chosen in a two-stage process wherein the image  
time is validated (or confirmed) by each storage device before  
the requested Time Image or Instant Restore operation takes  
place. In that two-stage process, there may be a window of  
10 opportunity for a rollup to change the timeline in such a way  
that makes a chosen image time invalid. To close this window,  
the validation phase in that two-stage process may cancel a  
currently running rollup and prevent other rollups from  
running. In addition, if a rollup is in a state that cannot  
be canceled, then the validation phase may complete in error  
15 to force a new query of the timeline. Configuration events in  
the validate phase for Time Image and Instant Restore may also  
pass in the rollup generation number queried when the timeline  
was queried. If the generation number has changed since the  
query, the validation phase may also complete in error to  
20 force a new query of the timeline.

In a storage system, each quota group may have an  
allocated timeline storage space for rollup purposes.  
According to one embodiment, for example, the amount of  
storage provisioned may be based on the following formula:  
25 (Maximum number of LUs in the storage system × Maximum number  
of rollup levels possible) + 100 working chunks for doing a  
rollup.

In order to have an active Level-1, Level-2, and Level-3  
rollup chunk for each LU, and to have 100 chunks for  
30 performing the active rollup, approximately 20 GB of storage  
space may be allocated as rollup quota per quota group.

There may not be enough space to accommodate a Copy Phase  
for a given rollup. This situation may be detected using

statistics generated in the Coalescence Phase, and the rollup for that rollup level may be cut short. The rollup state information may then be updated to reflect that the rollup has not yet completed and the next LU chosen by the determination  
5 phase may be the same LU, so that the rollup may resume.

Rollup chunks may be essentially transformed into new COW data in the Replacement Phase. The old data chunks may be freed as they are replaced. The rollup quota may be updated to reflect changes in the amount of space used. This  
10 guarantees that a rollup does not cause a deprecation of timeline and that there is space for at least 100 rollup chunks available at the start of any given rollup.

When a rollup thread is woken on a node due to a deprecation of timeline or a threshold of the timeline storage  
15 capacity being crossed, the rollup thread on other nodes may be woken as well, so that LUs belonging to the same dataset may be processed. The mechanism used to awaken the remote thread may be a simple spread message.

Starting a capture mode, adding LUs to a capturing  
20 dataset, and removing LUs from a capturing dataset all have the same effect on rollup chunks as these events have on normal time store data chunks. Every capturing LU may have an active rollup chunk for every level configured in the TLP. When an LU has capture mode turned on a rollup chunk may be  
25 activated for each level configured.

Timeline deprecation typically involves discarding older backup data in a timeline once the timeline storage is approaching its quota. If a deprecation of timeline is about to take place, any rollup currently running may be immediately  
30 cancelled. When the deprecation is completed, another rollup may be immediately scheduled.

Rolled-up data may typically be deprecated first in the timeline since they may be the oldest data in the timeline. One anomaly in terms of deprecation may be that it is possible

for an active rollup chunk to be deprecated. Multiple rollup chunks may contain duplicate start times, much like data chunks copied back during an Instant Restore. The same deprecation rules apply to these rollup chunks as do the copy  
5 back chunks. Rollup chunks with duplicate start times may be all deprecated regardless of how much space is needed.

Time images may be short lived entities. As a result, they may be deleted when deprecation of timeline takes place. One goal of timeline compression may be to avoid this  
10 deprecation of timeline based on the premise that backup data at a higher level of granularity may be better than no data at all. This being the case, there may be two options:

1. Let rollups behave the same way as deprecation behaves and delete the time image when a rollup runs into  
15 a time image.

2. Stop rollups when encountering a time image. The danger being that deprecation may delete the time image anyway.

A global option may be set up by a user to control how rollups  
20 behave when encountering a time image. For example, any time image creation may cancel a running rollup.

An instant restore may cause the state of a timeline to change drastically and quickly. As a result, an active rollup may be immediately canceled when an instant restore is  
25 initiated.

Much like time images, user-defined timeline annotations may be invalidated when timeline is deprecated. The same holds true for rollups.

On failover, any running rollup may be cancelled.  
30 Leftover data for rollups that may have been in progress on a remote node may be cleaned up by a surviving node. Rollup chunks may be scrubbed the same way COW data may be scrubbed on failover.

Figure 6 outlines major objects involved in an exemplary

program for timeline compression in accordance with  
embodiments of the present disclosure. A RollupManager object  
may be an interface for starting, canceling, maintaining  
synchronization, and querying state for rollups. In addition,  
5 it may be responsible for maintaining TLP objects. The  
RollupManager may contain an instance of a TimelineRollup  
object. The TimelineRollup object may be an active object  
responsible for performing all phases of a rollup. There may  
also be an object called RollupState which may be an interface  
10 for querying and setting state for rollups. It may maintain  
transient states related to the rollup process as well as the  
persistent states for LUs which may be stored in a global  
database. A RegionMaps/Indexing object may be responsible for  
building B-trees as described above. A tsAlloc object may  
15 allocate time store cache for rollup chunks. A tsQuotaGroup  
object may allocate timeline storage space for a dataset. A  
TimelineProfile object may record a TLP for a dataset.

At this point it should be noted that the technique for  
timeline compression in accordance with the present disclosure  
20 as described above typically involves the processing of input  
data and the generation of output data to some extent. This  
input data processing and output data generation may be  
implemented in hardware or software. For example, specific  
electronic components may be employed in a storage area  
25 network (SAN) or similar or related circuitry for implementing  
the functions associated with timeline compression in  
accordance with the present disclosure as described above.  
Alternatively, one or more processors operating in accordance  
with stored instructions may implement the functions  
30 associated with timeline compression in accordance with the  
present disclosure as described above. If such is the case,  
it is within the scope of the present disclosure that such  
instructions may be stored on one or more processor readable  
carriers (e.g., a magnetic disk), or transmitted to one or

more processors via one or more signals.

The present disclosure is not to be limited in scope by the specific embodiments described herein. Indeed, other various embodiments of and modifications to the present  
5 disclosure, in addition to those described herein, will be apparent to those of ordinary skill in the art from the foregoing description and accompanying drawings. Thus, such other embodiments and modifications are intended to fall within the scope of the present disclosure. Further, although  
10 the present disclosure has been described herein in the context of a particular implementation in a particular environment for a particular purpose, those of ordinary skill in the art will recognize that its usefulness is not limited thereto and that the present disclosure may be beneficially  
15 implemented in any number of environments for any number of purposes. Accordingly, the claims set forth below should be construed in view of the full breadth and spirit of the present disclosure as described herein.

## CLAIMS

1. A method for timeline compression in a storage system, wherein digital content of the storage system is backed up to  
5 enable restoration of the digital content to one or more points in a timeline, the method comprising:

selecting a time interval in the timeline;

identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more  
10 sets of backup data represent at least a portion of old data overwritten during the selected time interval; and

discarding other backup data recorded for the selected time interval, thereby reducing a granularity level of the timeline in the selected time interval.

15 2. The method according to claim 1, wherein the digital content of the storage system is backed up through copy-on-write operations into a plurality of sets of copy-on-write data and corresponding metadata, and wherein the step of  
20 identifying further comprises identifying one or more sets of copy-on-write data and corresponding metadata recorded for the selected time interval.

3. The method according to claim 2, wherein a length of the  
25 time interval is selected based at least in part on a desired granularity level of the timeline.

4. The method according to claim 2, wherein the step of identifying further comprises:

30 determining whether a storage unit in the storage system has been overwritten more than once during the selected time interval;

if the storage unit has been overwritten once during the selected time interval causing a sole set of copy-on-write



data and corresponding metadata to be recorded, selecting the sole set; and

if the storage unit has been overwritten more than once during the selected time interval causing multiple sets of copy-on-write data and corresponding metadata to be recorded, selecting one of the multiple sets.

5. The method according to claim 4, wherein, if the storage unit has been overwritten more than once during the selected time interval, the selected set of copy-on-write data and corresponding metadata is the earliest set recorded for the selected time interval.

6. The method according to claim 2, further comprising:  
coalescing metadata in the one or more identified sets of copy-on-write data and corresponding metadata.

7. The method according to claim 6, further comprising:  
identifying copy-on-write data that correspond to the coalesced metadata; and  
replacing all sets of copy-on-write data and corresponding metadata previously recorded for the selected time interval with a new set comprising the identified copy-on-write data and the coalesced metadata.

8. The method according to claim 2, further comprising:  
selecting multiple time intervals in a portion of the timeline based on a desired granularity level for the portion of the timeline; and  
repeating the steps of identifying and discarding for the selected multiple time intervals.

9. The method according to claim 2, wherein the storage system comprises a plurality of storage devices, the method

further comprising:

repeating the steps of identifying and discarding for one or more of the plurality of storage devices to cause the plurality of storage devices to have a consistent granularity level of the timeline with respect to one another.

10. The method according to claim 2, wherein the steps of selecting, identifying and discarding are triggered when one or more of the following conditions are met:

10 a predetermined storage capacity for the timeline has been consumed;

a predetermined amount of data have been accumulated for a granularity level of the timeline;

granularity levels of the timeline for at least two storage devices in the storage system are inconsistent;

15 an instruction to reduce the granularity of the timeline is received; and

a scheduled time for reducing the granularity of the timeline is reached.

20 11. The method according to claim 2, further comprising:

scanning the storage system for a storage device for which the granularity of the timeline can be reduced.

25 12. At least one signal embodied in at least one carrier wave for transmitting a computer program of instructions configured to be readable by at least one processor for instructing the at least one processor to execute a computer process for performing the method as recited in claim 1.

30 13. At least one processor readable carrier for storing a computer program of instructions configured to be readable by at least one processor for instructing the at least one processor to execute a computer process for performing the

method as recited in claim 1.

14. A system for timeline compression in a storage system, wherein digital content of the storage system is backed up to enable restoration of the digital content to one or more points in a timeline, the system comprising:

means for selecting a time interval in the timeline;

means for identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more sets of backup data represent at least a portion of old data overwritten during the selected time interval; and

means for discarding other backup data recorded for the selected time interval, thereby reducing a granularity level of the timeline in the selected time interval.

15. A system for timeline compression in a storage system, wherein digital content of the storage system is backed up to enable restoration of the digital content to one or more points in a timeline, the system comprising:

a storage medium for storing instructions; and

at least one processor for:

selecting a time interval in the timeline;

identifying one or more sets of backup data recorded for the selected time interval, wherein the identified one or more sets of backup data represent at least a portion of old data overwritten during the selected time interval; and

discarding other backup data recorded for the selected time interval, thereby reducing a granularity level of the timeline in the selected time interval.

Figure 1a






Figure 1b





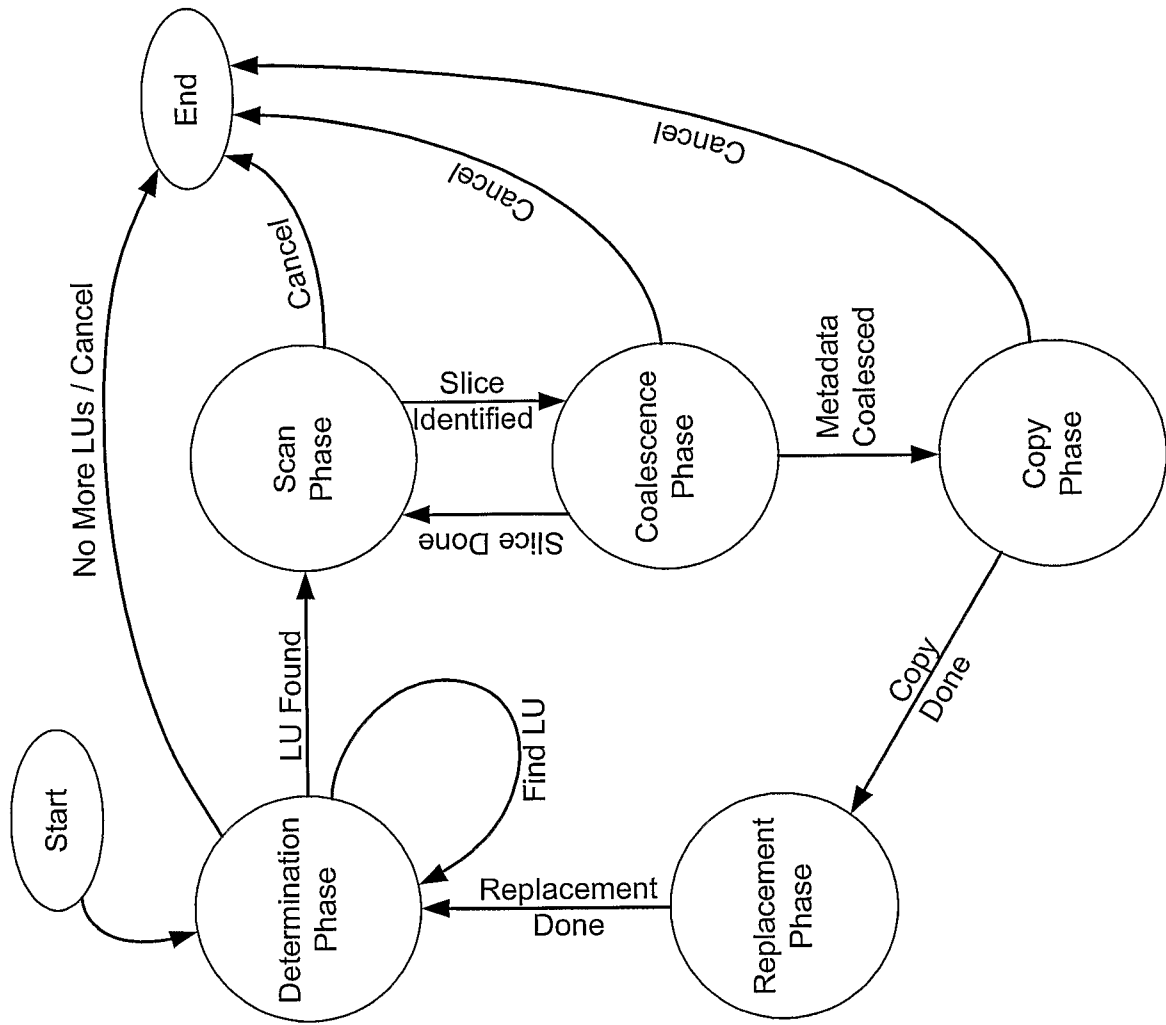


Figure 3

Figure 4




Figure 5






Figure 6

