



US012007846B2

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

(10) **Patent No.:** **US 12,007,846 B2**

(45) **Date of Patent:** **\*Jun. 11, 2024**

(54) **MANIFEST-BASED SNAPSHOTS IN  
DISTRIBUTED COMPUTING  
ENVIRONMENTS**

(58) **Field of Classification Search**  
CPC ..... G06F 16/27  
See application file for complete search history.

(71) Applicant: **Cloudera, Inc.**, Palo Alto, CA (US)

(56) **References Cited**

(72) Inventors: **Jonathan Ming-Cyn Hsieh**, San Francisco, CA (US); **Matteo Bertozzi**, Palo Alto, CA (US)

U.S. PATENT DOCUMENTS

5,825,877 A 10/1998 Dan et al.  
5,905,990 A 5/1999 Inglett

(Continued)

(73) Assignee: **CLOUDERA, INC.**, Santa Clara, CA (US)

OTHER PUBLICATIONS

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 589 days.

Babaoglu, O., and K. Marzullo, "Consistent global states of distributed systems: Fundamental concepts and mechanisms," Distributed Systems, 53 pages, Jan. 1993.

(Continued)

This patent is subject to a terminal disclaimer.

*Primary Examiner* — Khanh B Pham

(21) Appl. No.: **16/943,776**

(74) *Attorney, Agent, or Firm* — Perkins Coie LLP

(22) Filed: **Jul. 30, 2020**

(57) **ABSTRACT**

(65) **Prior Publication Data**

US 2020/0356448 A1 Nov. 12, 2020

Scalable architectures, systems, and services are provided herein for creating manifest-based snapshots in distributed computing environments. In some embodiments, responsive to receiving a request to create a snapshot of a data object, a master node identifies multiple slave nodes on which a data object is stored in the cloud-computing platform and creates a snapshot manifest representing the snapshot of the data object. The snapshot manifest comprises a file including a listing of multiple file names in the snapshot manifest and reference information for locating the multiple files in the distributed database system. The snapshot can be created without disrupting I/O operations, e.g., in an online mode by various region servers as directed by the master node. Additionally, a log roll approach to creating the snapshot is also disclosed in which log files are marked. The replaying of log entries can reduce the probability of causal consistency in the snapshot.

**Related U.S. Application Data**

(63) Continuation of application No. 15/605,205, filed on May 25, 2017, now Pat. No. 10,776,217, which is a (Continued)

(51) **Int. Cl.**

**G06F 16/30** (2019.01)

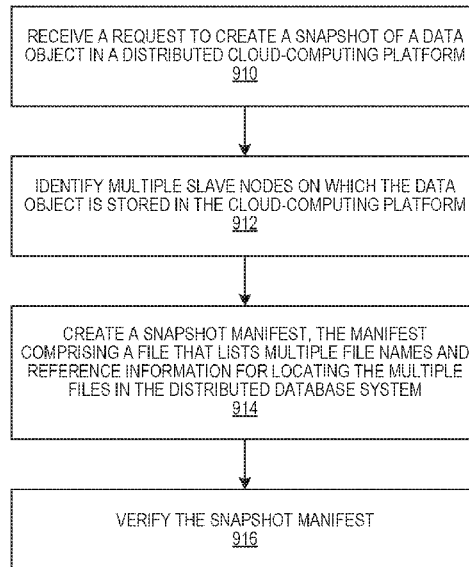
**G06F 11/14** (2006.01)

**G06F 16/27** (2019.01)

(52) **U.S. Cl.**

CPC ..... **G06F 11/1464** (2013.01); **G06F 16/27** (2019.01); **G06F 11/1456** (2013.01); **G06F 2201/84** (2013.01)

**16 Claims, 12 Drawing Sheets**



**Related U.S. Application Data**

continuation of application No. 14/527,563, filed on Oct. 29, 2014, now Pat. No. 9,690,671.

(60) Provisional application No. 61/898,955, filed on Nov. 1, 2013.

**References Cited**

U.S. PATENT DOCUMENTS

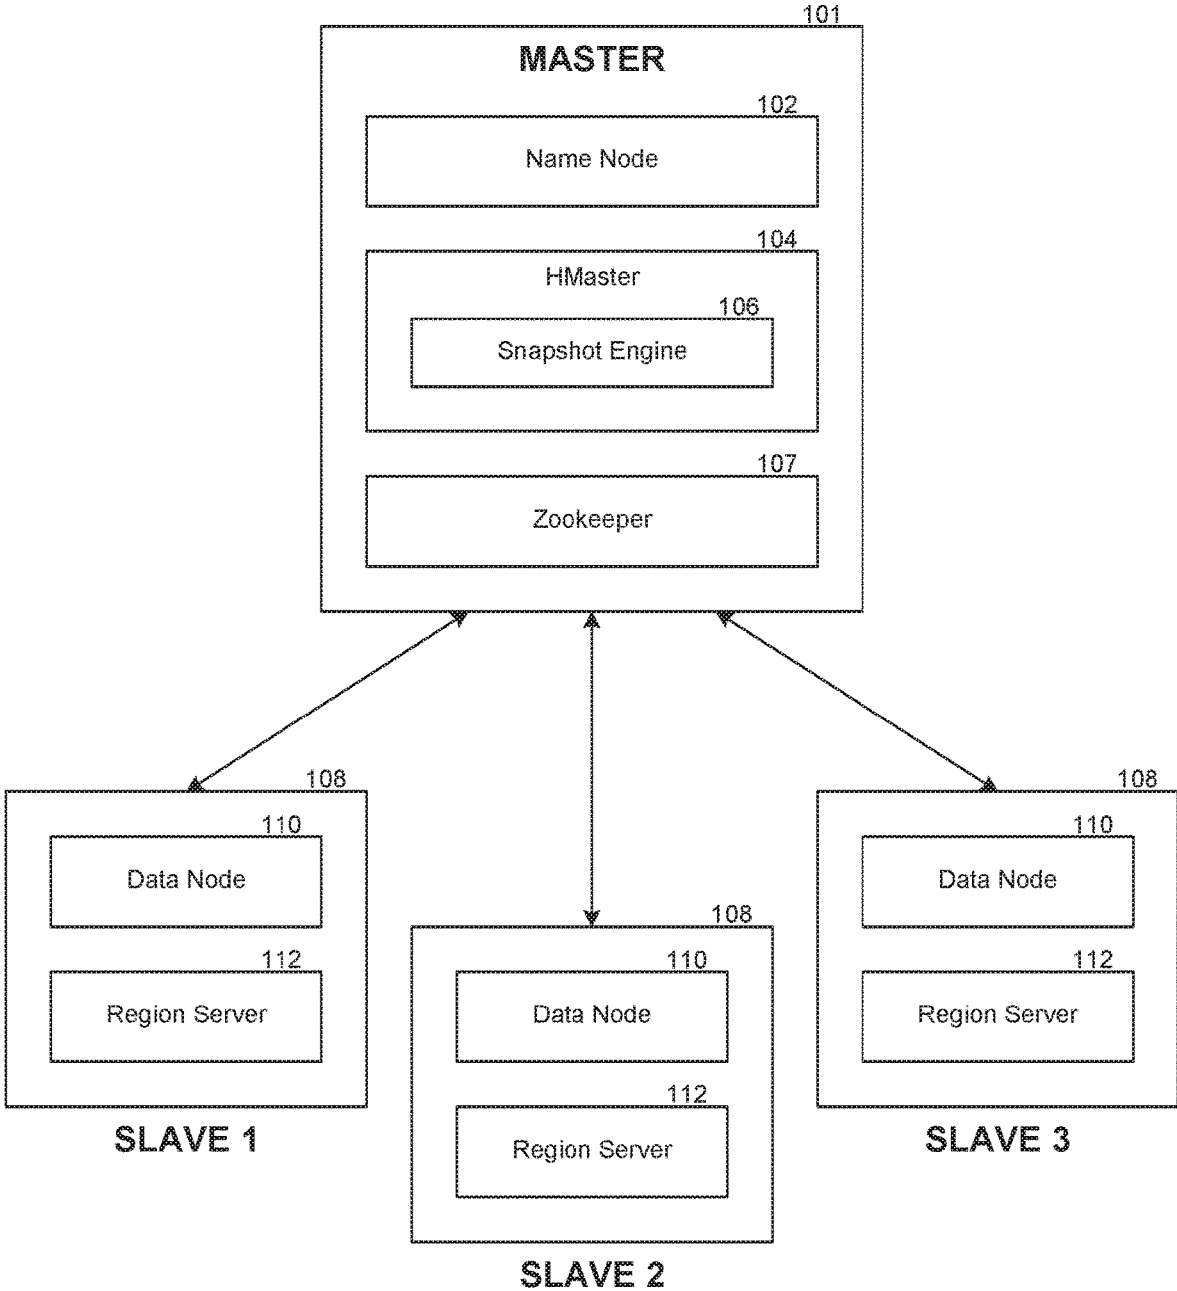
6,678,828 B1 1/2004 Pham et al.  
 6,931,530 B2 8/2005 Pham et al.  
 7,120,572 B1 10/2006 Liang  
 7,143,288 B2 11/2006 Pham et al.  
 7,246,369 B1 7/2007 Duan et al.  
 7,389,313 B1 6/2008 Hsieh et al.  
 8,024,560 B1 9/2011 Alten  
 8,042,063 B1 10/2011 Lin-Hendel  
 2,306,919 A1 11/2012 Sakamura et al.  
 8,306,919 B2 11/2012 Sakamura et al.  
 2,667,267 A1 3/2014 Garcia et al.  
 8,667,267 B1 3/2014 Garcia et al.  
 8,788,815 B1 7/2014 Garcia et al.  
 8,789,208 B1 7/2014 Sundaram et al.  
 8,821,602 B2 9/2014 McAlister  
 9,081,855 B1 7/2015 Abeloe et al.  
 9,448,893 B1 9/2016 Whitehead et al.  
 9,635,132 B1\* 4/2017 Lin ..... H04L 67/1097  
 9,690,671 B2 6/2017 Hsieh et al.  
 9,740,583 B1 8/2017 Brandwine  
 10,474,483 B2 11/2019 Kottomtharayil et al.  
 11,010,011 B2 5/2021 Varadharajan et al.  
 2002/0073322 A1 6/2002 Park et al.  
 2003/0101321 A1 5/2003 Ohran  
 2005/0171983 A1 8/2005 Deo et al.  
 2005/0182749 A1 8/2005 Matsui  
 2005/0246397 A1 11/2005 Edwards et al.  
 2006/0031747 A1 2/2006 Wada et al.  
 2006/0050877 A1 3/2006 Nakamura  
 2006/0143453 A1 6/2006 Imamoto et al.  
 2006/0156018 A1 7/2006 Lauer et al.  
 2007/0067450 A1 3/2007 Malloy et al.  
 2007/0177737 A1 8/2007 Jung et al.  
 2007/0180255 A1 8/2007 Hanada et al.  
 2007/0186112 A1 8/2007 Perlin et al.  
 2007/0226488 A1 9/2007 Lin et al.  
 2007/0255943 A1 11/2007 Kern et al.  
 2009/0259838 A1 10/2009 Lin  
 2009/0300416 A1 12/2009 Watanabe et al.  
 2009/0307783 A1 12/2009 Maeda et al.  
 2010/0008509 A1 1/2010 Matsushita et al.  
 2010/0011178 A1 1/2010 Feathergill  
 2010/0161556 A1 6/2010 Anderson et al.  
 2010/0198972 A1 8/2010 Umbehoeker  
 2010/0241614 A1 9/2010 Shaull et al.  
 2010/0296652 A1 11/2010 Nakayama et al.  
 2010/0306286 A1 12/2010 Chiu et al.  
 2010/0325713 A1 12/2010 Kurita et al.  
 2011/0029713 A1 2/2011 Wade et al.  
 2011/0055578 A1 3/2011 Resch  
 2011/0078549 A1 3/2011 Thueringer et al.  
 2011/0246964 A1 10/2011 Cox, III et al.  
 2011/0276495 A1 11/2011 Varadarajan et al.

2011/0302417 A1 12/2011 Whillock et al.  
 2012/0036357 A1 2/2012 Struik  
 2012/0054152 A1 3/2012 Adkins et al.  
 2012/0130874 A1 5/2012 Mane et al.  
 2012/0131341 A1 5/2012 Mane et al.  
 2012/0150819 A1 6/2012 Lindahl et al.  
 2012/0233522 A1 9/2012 Barton et al.  
 2012/0323854 A1 12/2012 Schreter  
 2013/0054976 A1 2/2013 Brown et al.  
 2013/0085988 A1 4/2013 Hiraguchi et al.  
 2013/0151884 A1\* 6/2013 Hsu ..... G06F 16/1827  
 714/E11.073  
 2013/0191589 A1\* 7/2013 Smith ..... G06F 11/2733  
 711/114  
 2013/0198120 A1 8/2013 Vezina  
 2013/0218840 A1 8/2013 Smith et al.  
 2013/0282650 A1 10/2013 Zhang et al.  
 2014/0032593 A1 1/2014 Libenzi et al.  
 2014/0040575 A1 2/2014 Horn  
 2014/0040669 A1 2/2014 Davis et al.  
 2014/0089264 A1\* 3/2014 Talagala ..... G06F 11/1448  
 707/649  
 2014/0188805 A1 7/2014 Vijayan et al.  
 2014/0188825 A1\* 7/2014 Muthukkaruppan .....  
 G06F 16/178  
 707/694  
 2014/0196115 A1 7/2014 Pelykh  
 2014/0201753 A1 7/2014 He et al.  
 2014/0214773 A1\* 7/2014 Zuckerman ..... G06F 16/1865  
 707/684  
 2014/0229422 A1 8/2014 Jain et al.  
 2014/0330784 A1 11/2014 Sundaram et al.  
 2015/0081243 A1 3/2015 Ganai et al.  
 2015/0127608 A1 5/2015 Hsieh et al.  
 2020/0356447 A1 11/2020 Hsieh et al.

OTHER PUBLICATIONS

Cheng, Security Attack Safe Mobil and Cloud-Based One-Time Password Tokens Using Rubbing Encryption Algorithm, ACM, Jun. 2011, pp. 304-336.  
 Corbett et al., "Spanner: Google's Globally Distributed Database," Transactions on Computer Systems (TOCS), vol. 31, No. 3, 14 pages, Aug. 2013.  
 Du et al., "A Pairwise Key Predistribution Scheme for Wireless Sensor Networks," ACM, May 2005, vol. 8 Issue 2, pp. 228-258.  
 Kim et al., Design and Implementation of a Private and Public Key Crypto Processor and Its Application to a Security System, IEEE, Feb. 2004, vol. 50, Issue 1, pp. 214-224.  
 Ko et al., "A Study of Encryption Algorithm for RFID tag (SEED: 8 Rounds X 64 bit block)," IEEE, 2008, pp. 672-677.  
 Lamport, L., "Time, clocks, and the ordering of events in a distributed system," Communications of the ACM, vol. 21, No. 7, pp. 558-565, Jul. 1978.  
 Stoller, S.D., "Detecting global predicates in distributed systems with clocks," Distributed Computing, vol. 13, No. 2, pp. 85-98, Feb. 2000.  
 Non-Final Office Action for U.S. Appl. No. 13/362,698, dated Apr. 29, 2013, 23 pages.  
 Non-Final Office Action for U.S. Appl. No. 13/362,695, dated Apr. 29, 2013, 23 pgs.

\* cited by examiner



*FIG. 1*

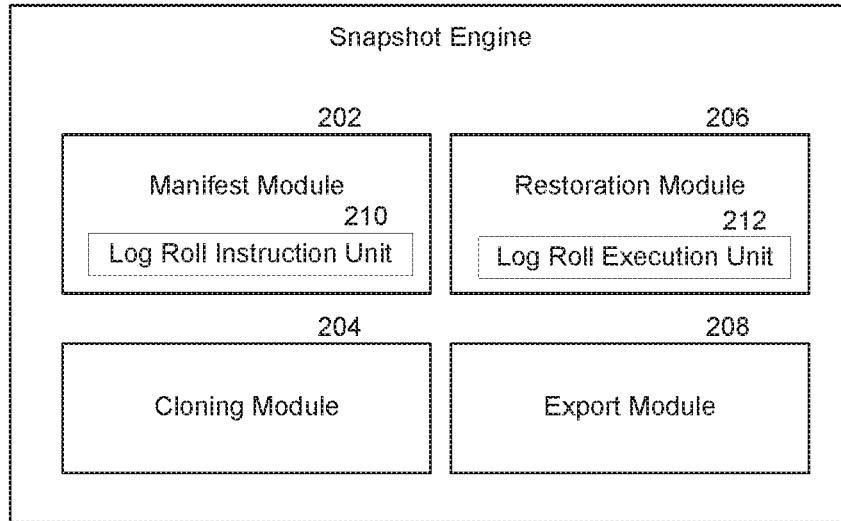


FIG. 2




FIG. 3




FIG. 4A




FIG. 4B




FIG. 4C




FIG. 5A




FIG. 5B




FIG. 6A




FIG. 6B




FIG. 6C

604

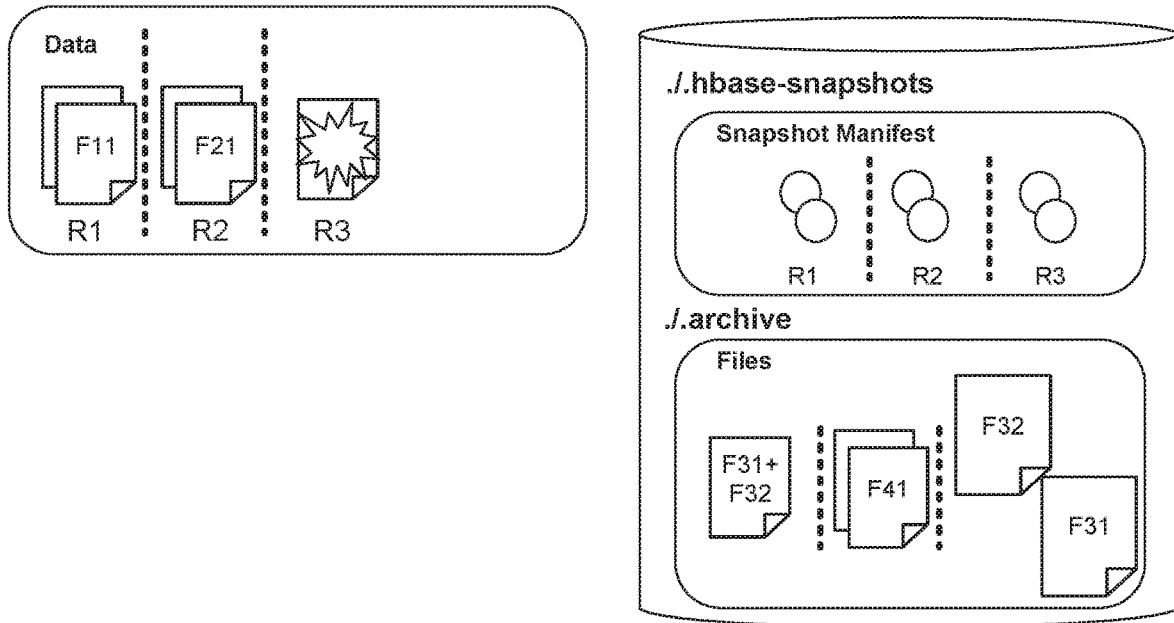


FIG. 6D




FIG. 6E






FIG. 7A




FIG. 7B




FIG. 7C




FIG. 8A

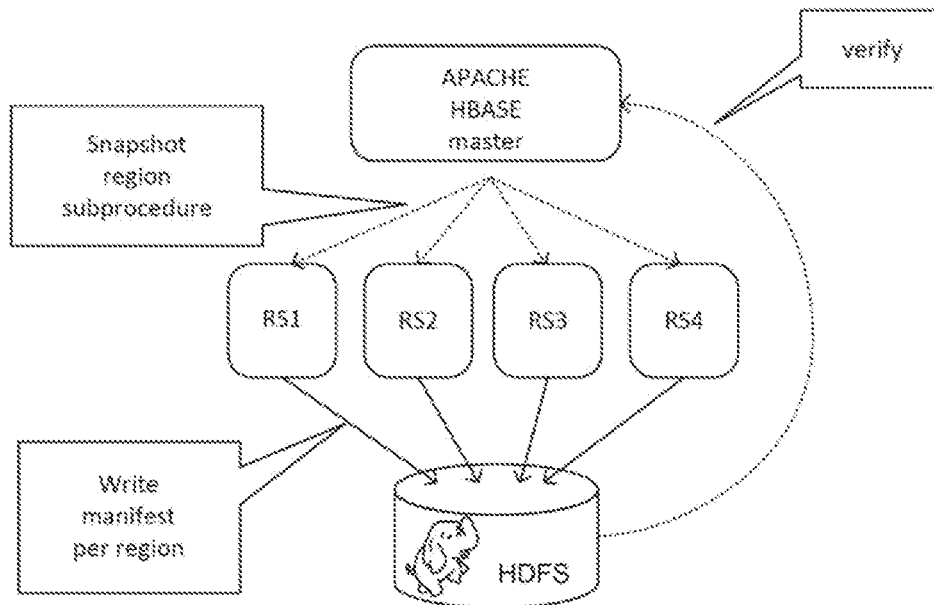
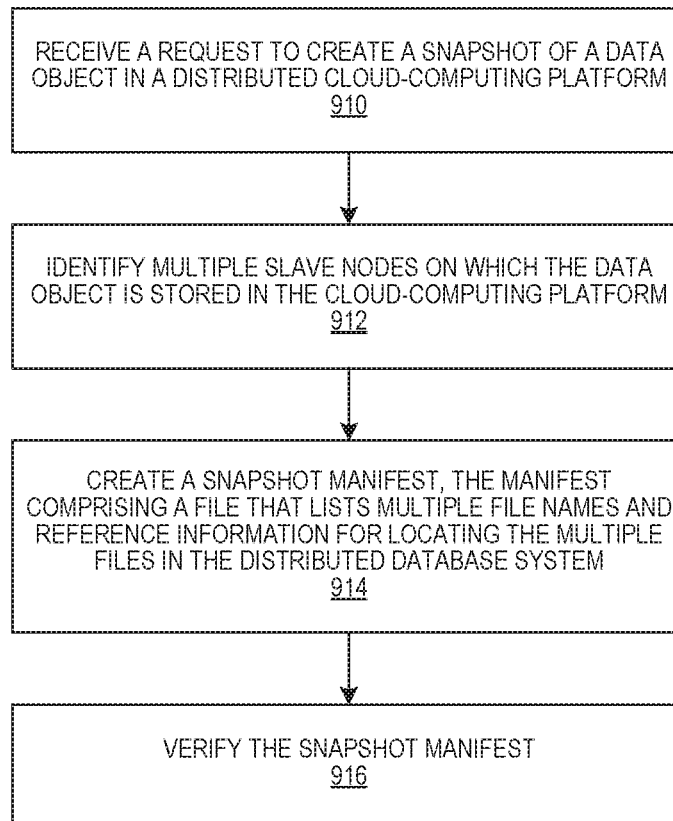
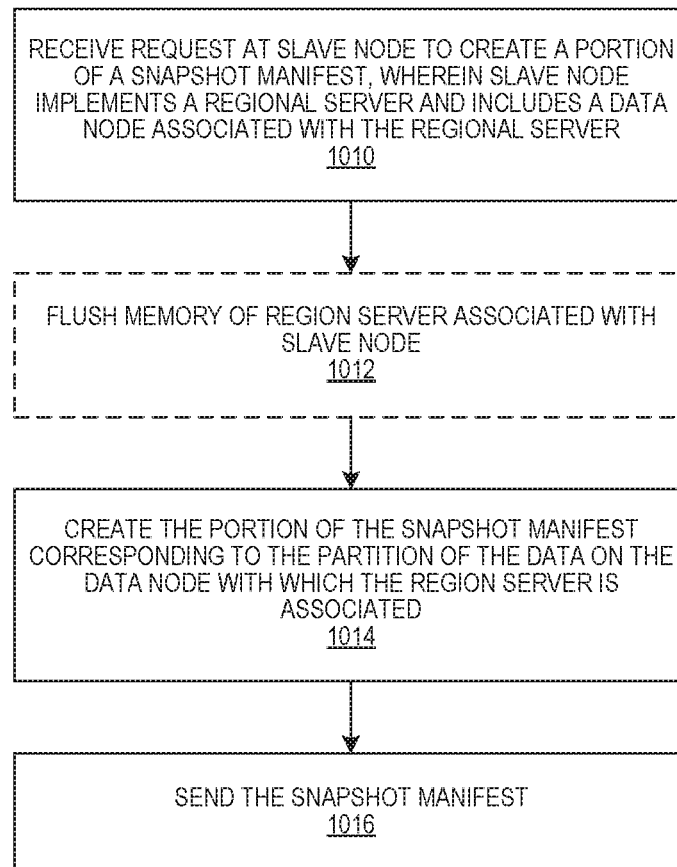


FIG. 8B



**FIG. 9**

**FIG. 10**

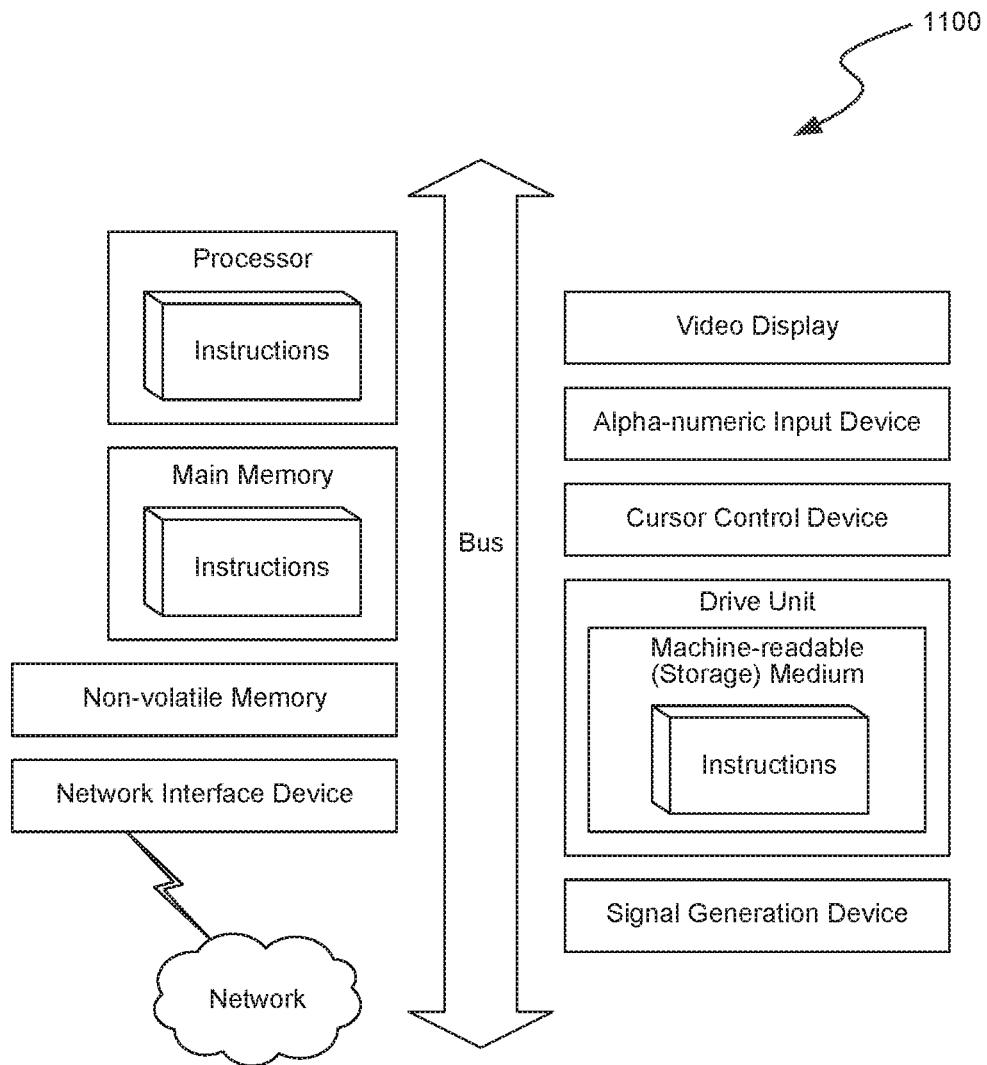


FIG. 11

## MANIFEST-BASED SNAPSHOTS IN DISTRIBUTED COMPUTING ENVIRONMENTS

### CROSS-REFERENCE TO RELATED APPLICATION(S)

This application is a continuation of U.S. patent application Ser. No. 15/605,205, filed May 25, 2017, entitled “MANIFEST-BASED SNAPSHOTS IN DISTRIBUTED COMPUTING ENVIRONMENTS,” which is a continuation of U.S. patent application Ser. No. 14/527,563, filed Oct. 29, 2014, entitled “MANIFEST-BASED SNAPSHOTS IN DISTRIBUTED COMPUTING ENVIRONMENTS (now U.S. Pat. No. 9,960,671),” which claims benefit of and priority to U.S. Provisional Patent Application No. 61/898,955 entitled “SNAPSHOTS WITH CAUSAL CONSISTENCY FOR EXPORTING, CLONING AND RESTORING DISTRIBUTED DATA,” filed on Nov. 1, 2013, the contents of all these applications being incorporated by reference herein.

### BACKGROUND

The Apache Hadoop project (hereinafter “Hadoop”) is an open-source software framework for developing software for reliable, scalable and distributed processing of large data sets across clusters of commodity machines. A Hadoop cluster typically comprises a name node and multiple data nodes. Hadoop implements a distributed file system, known as Hadoop Distributed File System (HDFS). HDFS provides a unified file system for the cluster, with the name node managing the name space of the unified file system, by linking together file systems on the data nodes. Hadoop also includes a MapReduce function that provides a programming framework for job scheduling and cluster resource management.

Hadoop is supplemented by other Apache projects, including Apache Zookeeper (hereinafter “Zookeeper”) and Apache HBase (hereinafter “HBase”). ZooKeeper is a centralized service for maintaining configuration information and naming. ZooKeeper also provides distributed synchronization and group services. HBase is a scalable, distributed Not-only Structured Query Language (Not-only SQL) or No Structured Query Language (NoSQL) database or data store that supports structured storage of large tables. Generally, an HBase installation includes a region server associated with each of the data nodes and depends on a ZooKeeper to coordinate the region servers. Each of the region servers works with data files called HFiles underlying the large tables, write ahead logs called HLogs, and other metadata in a data directory on the data node.

HBase supports several approaches of batch backups. One approach requires three MapReduce jobs, first to dump each table on a source cluster into a sequence file (export), second to copy a directory of files from the source cluster to a target cluster (dist cp), and third to save a sequence file into a table on the target cluster (import). Another approach requires one MapReduce job to read data from one table on a source cluster and write the data to another table on a target cluster (copy table). Unfortunately, these approaches all involve the execution of table manipulation commands that incur high latency and substantially impact existing workloads. Accordingly, it would be beneficial to have a more efficient approach for backup and related purposes, such as cloning and restoration.

## BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 contains a diagram illustrating an example environment in which a snapshot engine may be deployed.

FIG. 2 is a block diagram illustrating example components of a snapshot engine.

FIG. 3 is a diagram illustrating an example snapshot for a table.

FIGS. 4A-4C illustrate an example process in which a snapshot manifest is updated.

FIGS. 5A and 5B illustrate an example process of cloning a table from a snapshot.

FIGS. 6A-6E are diagrams illustrating an example process of restoring a table using a snapshot.

FIGS. 7A-7C are diagrams illustrating an example process in which creating a snapshot in an online mode results in causal inconsistency.

FIGS. 8A-8B are diagrams illustrating various components of a distributed cloud-computing platform involved in offline and online creation of a manifest-based snapshot of a data object in the distributed cloud-computing platform, respectively.

FIG. 9 is a flow diagram illustrating creation of a manifest-based snapshot of a data object in a distributed cloud-computing platform, according to an embodiment.

FIG. 10 is a flow diagram illustrating creation of a portion of a manifest-based snapshot of a data object in a distributed cloud-computing platform, according to an embodiment.

FIG. 11 depicts a diagrammatic representation of a machine, in the example form, of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed.

### DETAILED DESCRIPTION

The following description and drawings are illustrative and are not to be construed as limiting. Numerous specific details are described to provide a thorough understanding of the disclosure. However, in certain instances, well-known or conventional details are not described in order to avoid obscuring the description. References to one or an embodiment in the present disclosure can be, but not necessarily are, references to the same embodiment and such references mean at least one of the embodiments.

Reference in this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the disclosure. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment, nor are separate or alternative embodiments mutually exclusive of other embodiments. Moreover, various features are described which may be exhibited by some embodiments and not by others. Similarly, various requirements are described which may be requirements for some embodiments but no other embodiments.

The terms used in this specification generally have their ordinary meanings in the art, within the context of the disclosure, and in the specific context where each term is used. Certain terms that are used to describe the disclosure are discussed below, or elsewhere in the specification, to provide additional guidance to the practitioner regarding the description of the disclosure. For convenience, certain terms may be highlighted, for example using italics and/or quotation marks. The use of highlighting has no influence on the

scope and meaning of a term; the scope and meaning of a term is the same, in the same context, whether or not it is highlighted. It will be appreciated that same thing can be said in more than one way.

Consequently, alternative language and synonyms may be used for any one or more of the terms discussed herein, nor is any special significance to be placed upon whether or not a term is elaborated or discussed herein. Synonyms for certain terms are provided. A recital of one or more synonyms does not exclude the use of other synonyms. The use of examples anywhere in this specification including examples of any terms discussed herein is illustrative only, and is not intended to further limit the scope and meaning of the disclosure or of any exemplified term. Likewise, the disclosure is not limited to various embodiments given in this specification.

Without intent to further limit the scope of the disclosure, examples of instruments, apparatus, methods and their related results according to the embodiments of the present disclosure are given below. Note that titles or subtitles may be used in the examples for convenience of a reader, which in no way should limit the scope of the disclosure. Unless otherwise defined, all technical and scientific terms used herein have the same meaning as commonly understood by one of ordinary skill in the art to which this disclosure pertains. In the case of conflict, the present document, including definitions, will control.

As referred to herein, snapshot is a collection of metadata required to reconstitute the data near a particular point in time. Embodiments of the present disclosure include an HBase-based snapshot engine for Hadoop. Embodiments of the present disclosure further describe systems and methods for managing snapshots, as well as exporting, cloning and restoring cluster data using the snapshots to recover from user errors and bootstrap data replication.

Among other benefits, the technologies disclosed in this application enable efficient and effective data management for a distributed, cloud computing environment. For example, the creation, maintenance and usage of snapshots facilitate export, cloning, restoration and other data operations to increase data quality and availability with minimal impact on existing workloads.

Additionally, technology is disclosed for creating manifest-based snapshots in distributed computing environments. In some embodiments, responsive to receiving a request to create a snapshot of a data object, a master node identifies multiple slave nodes on which a data object is stored in the cloud-computing platform and creates a snapshot manifest representing the snapshot of the data object. The snapshot manifest comprises a file including a listing of multiple file names in the snapshot manifest and reference information for locating the multiple files in the distributed database system. The snapshot can be created without disrupting I/O operations, e.g., in an online mode by various region servers as directed by the master node. Additionally, a log roll approach to creating the snapshot is also disclosed in which log files are marked. Replaying of log entries ensures causal consistency in the snapshot.

FIG. 1 contains a diagram illustrating an example environment in which a snapshot engine may be deployed. In one embodiment, the environment includes a Hadoop cluster having a master machine **101** and multiple slave machines **108A-N**. As part of the HBase installation, the master machine may run the name node **102**, an HBase HMaster **104** and a Zookeeper server **107**. Likewise, the slave machines **108A-N** may run data nodes **110A-N** and an HBase region server **112A-N**, respectively. As illustrated in

the example of FIG. 1, the HMaster **104** supports a snapshot engine **106**. It is appreciated that the system can operate as described herein with any number of slaves, as few as one slave.

As discussed above, Hadoop implements a distributed file system, known as Hadoop Distributed File System (HDFS). HDFS provides a unified file system for the cluster, with the name node managing the name space of the unified file system, by linking together file systems on the data nodes. As described herein, the distributed file system can logically be described as including the name node and the data nodes.

FIG. 2 contains a block diagram illustrating example components of the snapshot engine. In one embodiment, the snapshot engine comprises a manifest module **202**, a cloning module **204** a restoration module **206** and an export module **208**. The manifest module **202** manages a snapshot for a table, including creation and modification. The cloning module **204** handles cloning a table or creating an additional handle to the table using a snapshot for the table. The restoration module **206** handles restoring a table to a previous state using a snapshot for the table. The export module **208** handles exporting a table from one cluster to another.

FIG. 3 contains a diagram illustrating an example snapshot for a data object, e.g., a table. HBase can represent changes to data existing in the distributed file system as addendums without modifying the existing data. In the example of FIG. 3, data **302** represents data stored on the distributed file system. More specifically, the data **302** represents the data stored on several name nodes corresponding to a table of several regions, including R1, R2 and R3. For example, HFiles including F1-1 and any addenda (e.g., F1-2-F1-*n*) would be considered to be in the region R1, HFiles including F21 and any addenda would be considered to be in the region R2, and HFiles including F31 and any addenda would be considered to be in the region R3. The manifest module **202** of FIG. 2 directs each of the respective region servers associated with the data nodes to create a snapshot manifest **304** representing a portion of the snapshot manifest, e.g., fsnapshot, of a region of the table and save that portion in a snapshot directory, for example, on the name node. The snapshot manifest **304** contains the name of those HFiles and a reference to the H-File, which enables the retrieval of the HFile. When each of the portions of the snapshot manifest **304** are successfully saved in the snapshot directory, the manifest module **202** may move the snapshot directory to make it available for clone/restore operations.

FIGS. 4A-4C illustrate an example process in depicting updating of a snapshot manifest, according to an embodiment. Generally, HBase does not modify an HFile except that it may merge the HFile with another HFile or split the HFile into two HFiles. In FIG. 4A, a snapshot originally contains references to the HFile F31 and an addendum F32 in the region R3. In FIG. 4B, it may be determined that the addendum F32 should be merged into the HFile F31, which would, however, break the references in the snapshot manifest invalidating the snapshot. In FIG. 4C, to preserve the validity of the snapshot, the manifest module **202** creates copies of the HFile F31 and the addendum F32 and saves them in an archive directory, for example, on the name node, typically before the merge takes place. The manifest module **202** also updates the snapshot manifest to replace the references to the original HFile F31 and the addendum F32 in the region R3 by references to the new HFile F31 and addendum F32. Similarly, when an HFile is split into two, the manifest module **202** would create a copy of the original HFile and save it in the archive directory. It would also



update the snapshot manifest to replace the reference to the original HFile by references to the two new HFiles.

The manifest module 202 may create a snapshot at any time or according to any predetermined schedule. As one example, it may create a snapshot periodically. As another example, it may create a snapshot when a sufficient amount of new data is created on the data nodes. In certain embodiments, the manifest module 202 may create a snapshot in an offline mode or an online mode. In the offline mode, the manifest module 202 would disable access to a table, examine the file system namespace to determine which HFiles form the table, and create the manifest. In the online mode, the manifest module 202 would communicate with each of the region servers in turn to capture the state of the data in real time. In one embodiment, in response to the communication, a region server would flush the data currently in memory into the table, and the manifest module 202 would include references to the HFiles created by the region server. In the offline mode, the snapshot for a table would be a true reflection of the table at a point in time. In the online mode, however, that may not be the case, as will be discussed in detail later.

Among other features, a snapshot enables easy backup of data at a particular point in time. In one embodiment, the export module 208 submits a MapReduce job for a table to copy all the corresponding HFiles, HLogs, and other related data in the data directories on the data nodes, the snapshot manifest for the table in the snapshot directory, and the table metadata on the name node, from a source cluster to a target cluster, which may have any number of data nodes. The table metadata contains information on how the HBase tables are mapped to the file system on the target cluster. The table metadata also contains settings about how the data is stored (compressed, encoded, etc.) on the target cluster. This process bypasses normal table manipulation commands and entails low-latency workloads.

It is possible to clone a table based on a snapshot for the table. FIGS. 5A and 5B illustrate an example process of cloning a table from a snapshot. In FIG. 5A, the cloning module 204 first creates a copy of the table metadata. It then “copies” the relevant HFiles on each data node according to the snapshot into a new directory on the data node corresponding to the name of the newly cloned table. However, each copy of a relevant HFile would not contain actual data but would constitute a link to the corresponding HFile, which basically never changes. For example, the newly created HFiles 502 in the region R3 are links to the original HFiles F31 and F32. A link is different from a reference contained in a snapshot manifest, though, in that a link would not break even as the actual data is moved. In FIG. 5B, even if the addendum F32 has been merged into the HFile F31 and, as a result, these HFiles have effectively been moved to the archive directory on the name node, the newly created HFiles 502 would remain valid by linking to the HFiles F31 and F32 in the archive directory instead.

With a snapshot available, it is also possible to roll back a table to a previous state. FIG. 6 contains five diagrams FIGS. 6A-6E illustrating an example process of restoring a table according to a snapshot. FIG. 6A shows the data when the snapshot was taken, while FIG. 6B shows the data at present. In one embodiment, to restore the table back to how it was when the snapshot was taken, the restoration module 206 disables access to the table and compares the snapshot with the table at present. In FIG. 6C, upon determining that the HFiles in the region R4 were not present back then, the restoration module 206 requests that the appropriate region server updates the appropriate table to remove those HFiles

with 602 but saves copies of those HFiles in the archive directory on the name node with 604. In FIG. 6D, upon determining that the HFiles in the region R3, including the combination of the H-Files F31 and F32, were not present back then, the restoration module 206 takes similar actions. In FIG. 6E, upon determining that copies of those HFiles that were in the region R3 back then are saved in the archive directory on the name node, the restoration module 206 requests that the appropriate region server creates HFiles that constitute links to those copies in the archive directory. It is possible that the restoration fails due to network issues, server issues, etc. In one embodiment, the restoration module 206 will restart the restoration process.

To create a snapshot in the online mode, the manifest module 202 would contact the region servers at slightly different times, which then each capture a part of a table in real time. The difference in times is typically very small, in the order of 10 to 20 seconds. However, it is theoretically possible to have inconsistency in the captured data.

FIGS. 7A-7C illustrate an example depicting creation of a snapshot in the online mode which may result in causal inconsistency. In FIG. 7A, with 702, the name node requests that the region server RS3 perform a memory flush, following which the name node builds the snapshot for the part of the table in the region R3. As a result, that part of the table now contains additional data 710 flushed out from the memory, and the snapshot contains references 712 to the HFiles which form that part of the table. In FIG. 7B, with 704, the region server RS3 inputs data A to the memory and with 706, the region server RS2 inputs data B which is related to data A to the memory, upon client requests. In FIG. 7C, with 708, the name node finally requests that the region server RS2 perform a memory flush, following which the name node builds the snapshot for the part of the table in the region R2. As a result, that part of the table now contains additional data 714 flushed out from the memory, which includes the data B, and the snapshot contains references 716 to the HFiles which form that part of the table. As can be seen, however, at this point in time, both data A and data B have been input, but only data B is accounted for in the snapshot, leading to causal inconsistency.

To reduce such causal inconsistencies for restoration purposes, in one embodiment, to establish a snapshot for a table, the snapshot module 202 may request that each region server enhance the HLogs. In response to the request, each region server would add an entry to an HLog for a memory flush. Next, it would place a marker in the HLog, by rolling or terminating the HLog and starting a separate one, for example, to indicate the taking of the snapshot. As adding one or two entries to a log takes very little time, at least compared to flushing data in a memory to a table, the snapshot module 202 may quickly move from one region server to the next, and all the markers would be placed in the respective logs at about the same time.

When it is time to restore the table according to the snapshot, the restoration module 206 would first clear the table. Next, it would replay the HLogs or reenact the log entries, such as adding a piece of data to memory or flushing it to the table up to where a marker is placed. With these log roll approaches, all the events up to a single point in time would be accounted for, ensuring causal consistency.

FIGS. 8A-8B are diagrams illustrating various components of a distributed cloud-computing platform involved in creation of a manifest-based snapshot of a data object in the distributed cloud-computing platform. More specifically, FIG. 8A illustrates offline creation of a manifest-based snapshot of a data object in the distributed cloud-computing

platform and FIG. 8B illustrates online creation of a manifest-based snapshot of a data object in the distributed cloud-computing platform, according to an embodiment. In the examples of FIGS. 8A-8B, the distributed cloud-computing platform is a Hadoop framework, although alternative configurations are possible.

Referring first to the example of FIG. 8A, when an Apache HBase master receives a request to create a manifest-based snapshot of a data object in the distributed cloud-computing platform the Apache HBase master creates the snapshot. As discussed above, in order to create a snapshot, the Apache HBase master typically creates one empty file for each file in a source table directory and copies region information corresponding to each compute region, e.g., a .regioninfo file for each region. Additionally, the Apache HBase master also copies a table descriptor file and a snapshot file, e.g., snapshotInfo file. However, as described herein, the Apache HBase master advantageously generates a manifest file (or H-file) that contains a list of the files in the manifest. In some embodiments, the Apache HBase master may create one manifest file per region. Alternatively, or additionally, in some embodiments, the Apache HBase master may create one manifest file per region server or slave node.

For instance, in the offline snapshot of example of FIG. 8A, the Apache HBase master writes a snapshot manifest per region and subsequently verifies the snapshot files. The verification process can include traversing the filesystem, e.g., `fs.listFiles()` and accessing the .regioninfo files to obtain the relevant information for the snapshot (e.g., file names and/or location information). Advantageously, the use of the manifest file(s) (or H-file(s)) that contain lists of files reduce utilization of the NameNode during snapshot and/or restore operations.

The example of FIG. 8B is similar to the example of FIG. 8A in the use of manifest file(s) (or H-file(s)). However, FIG. 8B illustrates an example online snapshot operation whereby region servers RS1-RS4 are each directed by the Apache HBase master to create the manifest files for data stored on corresponding data nodes (not shown). Again, as illustrated in FIG. 1, each slave node can include a region server and a data node. As shown in the example of FIG. 8B, the Apache HBase master directs and/or otherwise requests the regions servers RS1-RS4 to write the manifests via snapshot region subprocedures that are generated by the Apache HBase master and sent to each region server. The example of FIG. 8B includes four region servers, although it is appreciated that the distributed cloud-computing platform could include any number of region servers.

As described herein, creation of a snapshot while online allows the distributed database system to continue to accept I/O requests, e.g., reads and/or writes from clients during creation of the snapshot.

FIG. 9 is a flow diagram illustrating a set of operations 900 for creation of a manifest-based snapshot of a data object in a distributed cloud-computing platform, according to an embodiment. More specifically, the set of operations 900 illustrate example creation of a manifest-based snapshot. A master machine (or master node) such as, for example, master machine 101 of FIG. 1 can, among other functions, perform the set of operations 900.

To begin, at operation 910, the master machine receives a request to create a snapshot of a data object in a distributed cloud-computing platform. The request can be generated internally responsive to some automated and/or periodic triggering, manually by an administrator, by an external system, or in any other possible way known in the art.

At operation 912, the master machine identifies multiple slave nodes on which the data object is stored in the cloud-computing platform. As discussed herein, the distributed database system includes the master node and the multiple slave nodes. In some embodiments, each slave node implements a region server and includes a data node associated with the region server. In some embodiments, the data object is stored in partitions on multiple of the data nodes.

At operation 914, the master machine creates a snapshot manifest. As described herein, the manifest can include a file that lists multiple file names and/or reference information for locating the multiple files in the distributed database system. In some embodiments, the creation of the snapshot manifest can include directing, by the master machine, each region server to create a portion of the snapshot manifest corresponding to the partition of the data on the data node with which the region server is associated. In some embodiments, each portion of the snapshot manifest represents a single computing region. In some embodiments, each portion of the snapshot manifest comprises a single file including a listing of multiple file names in the portion of the snapshot manifest and reference information for locating the multiple files on corresponding data nodes.

In some embodiments, directing each region server to create the portion of the snapshot manifest corresponding to the partition of the data on the data node can further comprise requesting that the region server flush memory prior to creation of the portion of the snapshot manifest. In some embodiments, directing each region server to create the portion of the snapshot manifest corresponding to the partition of the data on the data node can further comprise directing each region server to drop a marker in a log file indicating a position of the log file associated with the snapshot manifest.

In some embodiments, directing each region server to create the portion of the snapshot manifest corresponding to the partition of the data on the data node with which the region server is associated includes generating, by the master node, a request to each region server for the respective portions of the snapshot manifest corresponding to the partition of the data on the data node with which the region servers are associated, and sending the requests to the corresponding region servers. The requests may be the snapshot region subprocedures discussed with reference to FIG. 8B.

In some embodiments, the master machine (or machine on which HDFS is controlled, e.g., NameNode) receives the respective portions of the snapshot manifest from the slave nodes and combines the respective portions of the snapshot manifest to form the snapshot manifest. This can include aggregating the portions of the snapshot manifest received from each of the slave nodes. Alternatively, these portions can be combined into a single file.

Lastly, at operation 916, the master machine verifies the snapshot manifest. For example, the master machine can traverse the files to verify that the snapshot is successfully created as discussed with reference to FIGS. 8A-8B. In some embodiments, the distributed database system continues to accept reads and writes from clients during creation of the snapshot, e.g., online mode.

As discussed herein, in some embodiments, the distributed cloud-computing platform is a Hadoop framework. In such cases, the master node can be an HBase master node, the slave nodes can be HBase slave nodes, and/or the data object can be a Hadoop table.

FIG. 10 is a flow diagram illustrating a set of operations 1000 for creation of a portion of a manifest-based snapshot

of a data object in a distributed cloud-computing platform, according to an embodiment. More specifically, the set of operations **1000** illustrate creation of an example manifest-based snapshot in an online mode of operation. A slave machine such as, for example, slave machine **108** of FIG. **1** can, among other functions, perform the set of operations **1000**.

To begin, at operation **1010**, the slave machine (or slave node) receives a request to create a portion of a snapshot manifest. As described herein, the slave machine (or slave node) can implement a region server and include the data node associated with the region server.

At operation **1012**, the slave machine (or slave node) optionally flushes the region server memory. For example, a log can be kept in memory until it is of a certain size at which point it is written out as an addenda as described herein. In some embodiments, flushing the memory results in the memory being written out to a persistent memory or log.

At operation **1014**, the slave machine (or slave node) creates the portion of the snapshot manifest corresponding to the partition of the data on the data node with which the region server is associated. In some embodiments, creating the portion of the manifest-based snapshot corresponding to the partition of the data on the data node with which the region server is associated further comprises dropping a marker in a log file indicating a position of the log file associated with the snapshot manifest.

In some embodiments, the portion of the snapshot manifest comprises a single file including a listing of multiple file names in the portion of the snapshot manifest and reference information for locating the multiple files on the data node. In some embodiments, the portion of the snapshot manifest represents a single region.

Lastly, at operation **1016**, the slave machine (or slave node) sends the snapshot manifest to another machine such as, for example, the master machine (or machine on which HDFS is controlled, e.g., NameNode).

As discussed herein, in some embodiments, the distributed cloud-computing platform is a Hadoop framework. In such cases, the master node can be an HBase master node, the slave nodes can be HBase slave nodes, and/or the data object can be a Hadoop table.

FIG. **11** shows a diagrammatic representation of a machine in the example form of a computer system within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed.

In the example of FIG. **11**, the computer system **1100** includes a processor, memory, non-volatile memory, and an interface device. Various common components (e.g., cache memory) are omitted for illustrative simplicity. The computer system **1100** is intended to illustrate a hardware device on which any of the components depicted in the example of FIG. **1** (and any other components described in this specification) can be implemented. The computer system **1100** can be of any applicable known or convenient type. The components of the computer system **1100** can be coupled together via a bus or through some other known or convenient device.

The processor may be, for example, a conventional microprocessor such as an Intel Pentium microprocessor or Motorola power PC microprocessor. One of skill in the relevant art will recognize that the terms “machine-readable (storage) medium” or “computer-readable (storage) medium” include any type of device that is accessible by the processor.

The memory is coupled to the processor by, for example, a bus. The memory can include, by way of example but not limitation, random access memory (RAM), such as dynamic RAM (DRAM) and static RAM (SRAM). The memory can be local, remote, or distributed,

The bus also couples the processor to the non-volatile memory and drive unit. The nonvolatile memory is often a magnetic floppy or hard disk, a magnetic-optical disk, an optical disk, a read-only memory (ROM), such as a CD-ROM, EPROM, or EEPROM, a magnetic or optical card, or another form of storage for large amounts of data. Some of this data is often written, by a direct memory access process, into memory during execution of software in the computer **1100**. The non-volatile storage can be local, remote, or distributed. The non-volatile memory is optional because systems can be created with all applicable data available in memory. A typical computer system will usually include at least a processor, memory, and a device (e.g., a bus) coupling the memory to the processor.

Software is typically stored in the non-volatile memory and/or the drive unit. Indeed, for large programs, it may not even be possible to store the entire program in the memory. Nevertheless, it should be understood that for software to run, if necessary, it is moved to a computer readable location appropriate for processing, and for illustrative purposes, that location is referred to as the memory in this paper. Even when software is moved to the memory for execution, the processor will typically make use of hardware registers to store values associated with the software, and local cache that, ideally, serves to speed up execution. As used herein, a software program is assumed to be stored at any known or convenient location (from nonvolatile storage to hardware registers) when the software program is referred to as “implemented in a computer-readable medium”. A processor is considered to be “configured to execute a program” when at least one value associated with the program is stored in a register readable by the processor.

The bus also couples the processor to the network interface device. The interface can include one or more of a modem or network interface. It will be appreciated that a modem or network interface can be considered to be part of the computer system. The interface can include an analog modem, ISDN modem, cable modem, token ring interface, satellite transmission interface (e.g. “direct PC”), or other interfaces for coupling a computer system to other computer systems. The interface can include one or more input and/or output devices, The I/O devices can include, by way of example but not limitation, a keyboard, a mouse or other pointing device, disk drives, printers, a scanner, and other input and/or output devices, including a display device. The display device can include, by way of example but not limitation, a cathode ray tube (CRT), liquid crystal display (LCD), or some other applicable known or convenient display device. For simplicity, it is assumed that controllers of any devices not depicted in the example of FIG. **11** reside in the interface.

In operation, the computer system **1100** can be controlled by operating system software that includes a file management system, such as a disk operating system. One example of operating system software with associated file management system software is the family of operating systems known as Windows® from Microsoft Corporation of Redmond, Washington, and their associated file management systems. Another example of operating system software with its associated file management system software is the Linux operating system and its associated file management system. The file management system is typically stored in

the nonvolatile memory and/or drive unit and causes the processor to execute the various acts required by the operating system to input and output data and to store data in the memory, including storing files on the non-volatile memory and/or drive unit.

Some portions of the detailed description may be presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise, as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as “processing” or “computing” or “calculating” or “determining” or “displaying” or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system’s registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the methods of some embodiments. The required structure for a variety of these systems will appear from the description below. In addition, the techniques are not described with reference to any particular programming language, and various embodiments may, thus, be implemented using a variety of programming languages.

In alternative embodiments, the machine operates as a standalone device or may be connected (e.g., networked) to other machines. In a networked deployment, the machine may operate in the capacity of a server or a client machine in a client-server network environment, or as a peer machine in a peer-to-peer (or distributed) network environment.

The machine may be a server computer, a client computer, a personal computer (PC), a tablet PC, a laptop computer, a set-top box (STB), a personal digital assistant (PDA), a cellular telephone, an iPhone, a Blackberry, a processor, a telephone, a web appliance, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine,

While the machine-readable medium or machine-readable storage medium is shown in an exemplary embodiment to be a single medium, the term “machine-readable medium” and “machine-readable storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and

servers) that store the one or more sets of instructions. The term “machine-readable medium” and “machine-readable storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instructions for execution by the machine and that cause the machine to perform any one or more of the methodologies of the presently disclosed technique and innovation.

In general, the routines executed to implement the embodiments of the disclosure may be implemented as part of an operating system or a specific application, component, program, object, module or sequence of instructions referred to as “computer programs.” The computer programs typically comprise one or more instructions set at various times in various memory and storage devices in a computer, and that, when read and executed by one or more processing units or processors in a computer, cause the computer to perform operations to execute elements involving the various aspects of the disclosure,

Moreover, while embodiments have been described in the context of fully functioning computers and computer systems, those skilled in the art will appreciate that the various embodiments are capable of being distributed as a program product in a variety of forms, and that the disclosure applies equally regardless of the particular type of machine or computer-readable media used to actually effect the distribution.

Further examples of machine-readable storage media, machine-readable media, or computer-readable (storage) media include but are not limited to recordable type media such as volatile and non-volatile memory devices, floppy and other removable disks, hard disk drives, optical disks (e.g., Compact Disk Read-Only Memory (CD ROMS), Digital Versatile Disks, (DVDs), etc.), among others, and transmission type media such as digital and analog communication links.

Unless the context clearly requires otherwise, throughout the description and the claims, the words “comprise,” “comprising,” and the like are to be construed in an inclusive sense, as opposed to an exclusive or exhaustive sense; that is to say, in the sense of “including, but not limited to.” As used herein, the terms “connected,” “coupled,” or any variant thereof, means any connection or coupling, either direct or indirect, between two or more elements; the coupling of connection between the elements can be physical, logical, or a combination thereof. Additionally, the words “herein,” “above,” “below,” and words of similar import, when used in this application, shall refer to this application as a whole and not to any particular portions of this application. Where the context permits, words in the above Detailed Description using the singular or plural number may also include the plural or singular number respectively. The word “or,” in reference to a list of two or more items, covers all of the following interpretations of the word: any of the items in the list, all of the items in the list, and any combination of the items in the list.

The above detailed description of embodiments of the disclosure is not intended to be exhaustive or to limit the teachings to the precise form disclosed above. While specific embodiments of, and examples for, the disclosure are described above for illustrative purposes, various equivalent modifications are possible within the scope of the disclosure, as those skilled in the relevant art will recognize. For example, while processes or blocks are presented in a given order, alternative embodiments may perform routines having steps, or employ systems having blocks, in a different order, and some processes or blocks may be deleted, moved, added, subdivided, combined, and/or modified to provide

alternative or subcombinations. Each of these processes or blocks may be implemented in a variety of different ways. Also, while processes or blocks are, at times, shown as being performed in series, these processes or blocks may instead be performed in parallel or may be performed at different 5 times. Further, any specific numbers noted herein are only examples: alternative implementations may employ differing values or ranges.

The teachings of the disclosure provided herein can be applied to other systems, not necessarily the system described above. The elements and acts of the various embodiments described above can be combined to provide further embodiments. 10

Any patents and applications and other references noted above, including any that may be listed in accompanying filing papers, are incorporated herein by reference. Aspects of the disclosure can be modified, if necessary, to employ the systems, functions, and concepts of the various references described above to provide yet further embodiments of the disclosure. 15

These and other changes can be made to the disclosure in light of the above Detailed Description. While the above description describes certain embodiments of the disclosure, and describes the best mode contemplated, no matter how detailed the above appears in text, the teachings can be practiced in many ways. Details of the system may vary considerably in its implementation details, while still being encompassed by the subject matter disclosed herein. As noted above, particular terminology used when describing certain features or aspects of the disclosure should not be taken to imply that the terminology is being redefined herein to be restricted to any specific characteristics, features, or aspects of the disclosure with which that terminology is associated. In general, the terms used in the following claims should not be construed to limit the disclosure to the specific 20 embodiments disclosed in the specification, unless the above Detailed Description section explicitly defines such terms. Accordingly, the actual scope of the disclosure encompasses not only the disclosed embodiments, but also all equivalent ways of practicing or implementing the disclosure under the claims. 25

While certain aspects of the disclosure are presented below in certain claim forms, the inventors contemplate the various aspects of the disclosure in any number of claim forms. For example, while only one aspect of the disclosure is recited as a means-plus-function claim under 35 U.S.C. § 112, ¶6, other aspects may likewise be embodied as a means-plus-function claim, or in other forms, such as being embodied in a computer-readable medium. (Any claims intended to be treated under 35 U.S.C. § 112, ¶6 will begin 30 with the words “means for”.) Accordingly, the applicant reserves the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the disclosure.

What is claimed is:

1. A computer-implemented method for operating a distributed computing platform, the distributed computing platform including a plurality of computing nodes, the method comprising: 35

implementing, by each of a plurality of first computing nodes of the plurality of computing nodes, a corresponding region server associated with a data node; receiving, by each of the plurality of first computing nodes, from a second computing node of the plurality of computing nodes, a request to create a corresponding portion of a snapshot manifest; and 40 in response to receiving the request:

creating, by each of the plurality of first computing nodes, the corresponding portion of the snapshot manifest corresponding to a partition of a data object stored on the data node with which the corresponding region server is associated with, wherein the snapshot manifest comprises a file including a listing of multiple file names in the snapshot manifest and reference information for locating multiple files corresponding to the multiple file names in the plurality of computing nodes, and wherein the second computing node is configured to form the snapshot manifest by combining the corresponding portion of the snapshot manifest of each of the plurality of first computing nodes; and 45

transmitting, by each of the plurality of first computing nodes, to the second computing node, the corresponding portion of the snapshot manifest.

2. The method of claim 1, further comprising: flushing memory of the corresponding region server before creating the corresponding portion of the snapshot manifest. 50

3. The method of claim 1, wherein creating the corresponding portion of the snapshot manifest includes: dropping a marker in a log file indicative of a position of the log file associated with the snapshot manifest. 55

4. The method of claim 1, further comprising: copying the partition of the data object into a new directory on the associated data node.

5. The method of claim 4, wherein a copy of the partition of the data object comprises a link to the partition of the data object but not actual data associated with the partition of the data object. 60

6. The method of claim 4, further comprising: receiving, by a first computing node of the plurality of first computing nodes, from the second computing node of the plurality of computing nodes, a second request to copy relevant partitions of the data object, the request associated with an operation to clone a table based on the snapshot; wherein the partition of the data object is copied by the first computing node in response to receiving the second request. 65

7. The method of claim 1, wherein the plurality of computing nodes include a master node and a plurality of slave nodes.

8. The method of claim 7, wherein each of the plurality of first computing nodes is one of the plurality of slave nodes, and wherein the second computing node is the master node.

9. The method of claim 7, wherein the master node comprises an HBASE master node, the slave nodes comprise HBASE slave nodes, and the data object comprises a HADOOP table.

10. The method of claim 1, further comprising: avoiding modification of the partition of the data object except when merging or splitting the partition of the data object.

11. The method of claim 1, further comprising: receiving, by a first computing node of the plurality of first computing nodes, from the second computing node of the plurality of computing nodes, a second request to restore a data object to a previous state based on a snapshot of the data object at the previous state; and in response to receiving the second request: copying a current state of the partition of the data object into an archive directory on the name node; and 65

15

updating the partition of the data object from the current state to the previous state.

12. The method of claim 1, further comprising: receiving, by a first computing node of the plurality of first computing nodes, from the second computing node of the plurality of computing nodes, a second request indicative of detection of a causal inconsistency; and in response to receiving the second request, adding an entry to a log file for a memory flush.

13. A distributed computer system comprising: a master node; and a plurality of slave nodes communicatively coupled to the master node, each of the plurality of slave nodes configured to:

implement a corresponding region server associated with a data node;

receive, from the master node, a request to create a corresponding portion of a snapshot manifest; and in response to receiving the request:

create the corresponding portion of the snapshot manifest corresponding to a partition of a data object stored on the data node with which the corresponding region server is associated with, wherein the snapshot manifest comprises a file including a listing of multiple file names in the snapshot manifest and reference information for locating multiple files corresponding to the multiple file names in the plurality of slave nodes, and wherein the master node is configured to form the snapshot manifest by combining the corresponding portion of the snapshot manifest of each of the plurality of slave nodes; and transmit, to the master node, the corresponding portion of the snapshot manifest.

16

14. The distributed computer system of claim 13, wherein at least one of the plurality of slave nodes is further configured to:

flush memory of the region server before creating the portion of the snapshot manifest.

15. The distributed computer system of claim 13, wherein creating the corresponding portion of the snapshot manifest includes:

dropping a marker in a log file indicative of a position of the log file associated with the snapshot manifest.

16. A non-transitory computer readable storage medium storing instructions, execution of which in a computer system, causes the computer system to:

implement a plurality of region servers each associated with a corresponding data node;

receive, from a master node of a distributed computing platform, a request to create a corresponding portion of a snapshot manifest for each of the plurality of region servers; and

in response to receiving the request:

create the corresponding portion of the snapshot manifest corresponding to a partition of a data object stored on the corresponding data node with which each of the plurality of region servers is associated with, wherein the snapshot manifest comprises a file including a listing of multiple file names in the snapshot manifest and reference information for locating multiple files corresponding to the multiple file names in a distributed database system, and wherein the master node is configured to form the snapshot manifest by combining the corresponding portion of the snapshot manifest of each of the plurality of regions; and

transmit, to the master node, the corresponding portion of the snapshot manifest.

\* \* \* \* \*