



- (51) **International Patent Classification:**  
G06F 12/16 (2006.01) G06F 17/30 (2006.01)
- (21) **International Application Number:**  
PCT/US2013/063120
- (22) **International Filing Date:**  
2 October 2013 (02.10.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
13/645,409 4 October 2012 (04.10.2012) US
- (71) **Applicant:** DEPHIX CORP. [US/US]; 275 Middlefield Road, Suite 50, Menlo Park, CA 94025 (US).
- (72) **Inventors:** STEWART, Michael; C/o Delphix Corp., 275 Middlefield Road, Suite 50, Menlo Park, CA 94025 (US). SRIHARI, Vinay; C/o Delphix Corp., 275 Middlefield Road, Suite 50, Menlo Park, CA 94025 (US). SINHA, Subhadeep; C/o Delphix Corp., 275 Middlefield Road, Suite 50, Menlo Park, CA 94025 (US). LEVENTHAL, Adam; C/o Delphix Corp., 275 Middlefield Road, suite 50, Menlo Park, CA 94025 (US). AHRENS, Matthew; C/o Delphix Corp., 275 Middlefield Road, Suite 50, Menlo Park, CA 94025 (US).
- (74) **Agents:** PANWAR, Rajendra, B. et al.; Fenwick & West LLP, Silicon Valley Center, 801 California Street, Mountain View, CA 94041 (US).

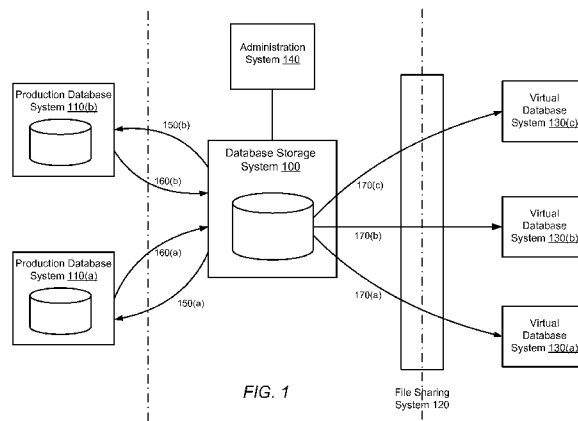
(81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP,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

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

**Published:**

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) **Title:** CREATING VALIDATED DATABASE SNAPSHOTS FOR PROVISIONING VIRTUAL DATABASES



(57) **Abstract:** Virtual databases are provisioned using point-in-time copies of a source database. Transaction logs received from the source database are applied to a point-in-time copy of the source database to create recovery milestones comprising database blocks that can be used to provision virtual databases. The recovery milestones are created in advance of receiving a request to provision a virtual database. Each recovery milestone is validated to determine that a consistent database snapshot can be created using the recovery milestone. If a problem is identified that prevents creation of a consistent database snapshot, the problem is fixed before creating the next recovery milestone. In response to receiving a request to provision a virtual database corresponding to a point in time, a recovery milestone is identified and transaction logs applied to the recovery milestone to create a database snapshot corresponding to the requested point in time.

WO 2014/055685 A2

## **CREATING VALIDATED DATABASE SNAPSHOTS FOR PROVISIONING VIRTUAL DATABASES**

### **BACKGROUND**

**[0001]** This invention relates generally to databases, and in particular to provisioning of virtual databases in advance of a request by creating validated database snapshots.

**[0002]** Databases store data that is critical to an organization and thus form an important part of an organization's information technology infrastructure. As the information available in an organization grows, so does the complexity of the infrastructure required to manage the databases that store the information. The increased complexity of the infrastructure increases the resources required to manage the databases and the applications that depend on the databases. These increased costs may include the costs associated with hardware for managing the databases as well as the costs associated with additional personnel needed to maintain the hardware. The increased complexity of the infrastructure also affects the maintenance operations associated with the databases, for example, causing backup and recovery operations to take significantly longer.

**[0003]** In a typical organization's infrastructure environment, production database servers run applications that manage the day-to-day transactions of the organization. Changes to production databases or to applications that depend on the production databases are tested on copies of the databases to protect the production environment. Copies of the production databases may be required for several stages in the lifecycles of workflows associated with the production database and applications that depend on the production databases. For example, the stages in the lifecycle of a change incorporated in a production database may include a development stage, a tuning stage, a testing stage, a quality assurance stage, a certification stage, a training stage, and a staging stage. Making copies of the production databases for each stage requires redundant and expensive hardware infrastructure as well as the time overhead required to copy the data, which may take significant amount of time.

### **SUMMARY**

**[0004]** To address the needs unmet by existing database technologies, embodiments of the invention pre-provision validated database snapshots in advance of receiving a request. A database storage system receives different point-in-time copies of a source database comprising database blocks of the source database and stores them on a storage system. The database storage system receives transaction logs from the source database corresponding to transactions executed by the source database. The database storage system creates multiple

recovery milestones, each recovery milestone comprising database blocks to which transaction logs are applied to create a consistent snapshot corresponding to a point in time version of the source database. At least some of the stored database blocks are associated with multiple recovery milestones. Each recovery milestone is created by applying the transaction logs to either a point-in-time copy of the source database or to a previous recovery milestone. The database storage system validates the recovery milestones created to determine whether a consistent snapshot of the source database can be created using the recovery milestone. If the database storage system receives a request to create a virtual database, the database storage system identifies a recovery milestone and provisions the requested virtual database based on the recovery milestone.

**[0005]** In an embodiment, if the database storage system determines that a recovery milestone is invalid, the database storage system executes steps to identify the cause of creation of the invalid recovery milestone and takes measures to create a corresponding valid recovery milestone.

**[0006]** The features and advantages described in this summary and the following detailed description are not all-inclusive. Many additional features and advantages will be apparent to one of ordinary skill in the art in view of the drawings, specification, and claims.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

**[0007]** FIG. 1 is diagram illustrating how information is copied from a production database to a database storage system and provisioned as virtual databases using a file sharing system, in accordance with an embodiment of the invention.

**[0008]** FIG. 2 is a schematic diagram of the architecture of a system that makes storage efficient copies of information from a production database and provisions virtual databases, in accordance with an embodiment of the invention.

**[0009]** FIG. 3 illustrates a process of provisioning a virtual database upon request from a point-in-time copy of the production database, in accordance with an embodiment of the invention.

**[0010]** FIG. 4 illustrates a process of pre-provisioning recovery milestones to provision a validated database snapshot, in accordance with an embodiment of the invention, in accordance with an embodiment of the invention.

**[0011]** FIG. 5 is a flowchart of the process for pre-provisioning validated snapshots of a database, in accordance with an embodiment of the invention.

**[0012]** The figures depict various embodiments of the present invention for purposes of illustration only. One skilled in the art will readily recognize from the following discussion

that alternative embodiments of the structures and methods illustrated herein may be employed without departing from the principles of the invention described herein.

#### **DETAILED DESCRIPTION**

**[0013]** A database comprises data stored in a computer for use by computer implemented applications. A database server is a computer program that can interact with the database and provides database services, for example, access to the data stored in the database. Database servers include commercially available programs, for example, database servers included with database management systems provided by ORACLE, SYBASE, MICROSOFT SQL SERVER, IBM DB2, MYSQL, POSTGRES, and the like. A database may be implemented using a database model, for example, a relational mode, object model, hierarchical mode or network model. The term “production database” is used in particular examples to illustrate a useful application of the technology; however, it can be appreciated that the techniques disclosed can be used for any database, regardless of whether the database is used as a production database. Furthermore, embodiments can create a virtual database using storage level snapshots of production databases or clones of production databases instead of a live production database. The virtual databases are “virtual” in the sense that the physical implementation of the database files is decoupled from the logical use of the database files by a database server. Virtual databases are described in U.S. Patent Application No. 12/603,541 filed on October 21, 2009, and issued as U.S. patent 8,150,808, each of which is incorporated by reference herein in its entirety.

**[0014]** In one embodiment, information from the production database is copied to a database storage system at various times, such as periodically. A virtual database created for a point in time is stored as a set of files that contain the information of the database as available at that point in time. Each file includes a set of database blocks and the data structures for referring to the database blocks. The files corresponding to the virtual database are made available to the database server using a file sharing mechanism, which links the virtual database to the appropriate database blocks stored on the storage system. The process of making the virtual database available to a database server is called “provisioning” the virtual database. In some embodiments, provisioning the virtual database includes managing the process of creating a running database server based on virtual database. The database server on which a virtual database has been provisioned can then read from and write to the files stored on the storage system. A database block may be shared between different files, each file associated with a different virtual database. In particular, a database block is shared

if the corresponding virtual database systems 130 are only reading the information in the database block and not writing to the database block.

**[0015]** The database storage system receives transaction logs corresponding to transactions executed in the production database since the last point-in-time copy of the production database was obtained. These transaction logs are applied to the point-in-time copy to create a database corresponding to a state of production database in between two point-in-time copies. If a request for provisioning a database corresponding to a time in between two point-in-time copies is obtained, the transaction logs are applied to the most recent point-in-time copy occurring before the time for which the virtual database is requested. The most time consuming step in provisioning a virtual database is applying the transaction logs to a point-in-time copy. Furthermore, if there is any corruption in the data used to create the virtual database, the database storage system may not be able to apply the transaction logs without taking steps to fix the problem causing the corruption of the transaction logs. The steps taken to fix any problem may be slow, for example, a step may require human intervention. As a result, provisioning the virtual database can take significant amount of time since a request for provisioning is received, causing the requestor to wait.

**[0016]** Embodiments reduce the amount of time between receiving a request for a virtual database and provisioning the requested virtual database. The database storage system creates multiple recovery milestones after receiving a point-in-time copy of the production database. Each recovery milestone corresponds to a state of the production database at a point in time. Each recovery milestone is created by applying transaction logs to the database blocks received as a point-in-time copy of the production database. The database storage system also performs validation of recovery milestones to make sure that a valid database snapshot can be created using the recovery milestone. If a problem is detected, the database storage system may take appropriate steps to fix the problem. Fixing these problems in advance of receiving any request for provisioned databases reduces the time that a requestor has to wait if the corruption of data or a bug or defect in a software module is identified after receiving a request for provisioning a virtual database.

**[0017]** Furthermore, since the transaction logs stored in the database storage system are applied to the point-in-time copies, the rate at which the point-in-time copies are received from the production database can be reduced compared to a system in which the recovery milestones are not created. Besides significant time savings can be achieved if multiple virtual databases are provisioned corresponding to the same point in time, for example, different virtual databases provisioned for different platforms.

### Overall System Environment

[0018] FIG. 1 illustrates one embodiment for how information may be copied from a production database to a database storage system and provisioned as virtual databases using a file sharing system. The production database systems 110 manage data for an organization. In some embodiments information may be copied from storage level snapshots of production databases or clones of production databases instead of a live production database. The database storage system 100 retrieves data associated with databases from one or more production database systems 110 and stores the data in an efficient manner, further described below. A database administrator user interface allows a database administrator to perform various actions supported by the database storage system 100.

[0019] In response to a request from the administrator system 140, or based on a predefined schedule, the database storage system 100 may send a request 150 for data to a production database system 110. The production database system 110 responds by sending information stored in the production database as a stream of data 160. The request 150 is sent periodically and the production database system 110 responds by sending information representing changes of data stored in the production database since the last response 160 sent by the production database system 110. The database storage system 100 receives the data 160 sent by the production database system 110 and stores the data. The database storage system 100 may analyze the data 160 received to determine whether to store the information or skip the information if the information is not useful for reconstructing the database at previous time points. The database storage system 100 stores the information efficiently, for example, by keeping versions of database blocks that have changed and reusing database blocks that have not changed.

[0020] To create a virtual database, the database storage system 100 creates files that represent the information corresponding to the production database system 110 at a given point in time. The database storage system 100 exposes 170 the corresponding files to a virtual database system 130 using a file sharing system 120. The virtual database system 130 runs a database server that can operate with the files exposed 170 by the database storage system 100. Hence, a virtual copy of the production database is created for the virtual database system 130 for a given point in time in a storage efficient manner.

### System Architecture

[0021] FIG. 2 shows a high level block diagram illustrating a system environment suitable for making storage efficient copies of information from a production database and provisioning one or more virtual databases using that information. The system environment

comprises one or more production database systems 110, a database storage system 100, an administration system 140, and one or more virtual database systems 130. Systems shown in FIG. 2 can communicate with each other if necessary via a network.

**[0022]** A production database system 110 is typically used by an organization for maintaining its daily transactions. For example, an online bookstore may save all the ongoing transactions related to book purchases, book returns, or inventory control in a production database system 110. The production database system 110 includes a database server 245 and a production DB data store 250. In alternative configurations, different and/or additional modules can be included in a production database system 110.

**[0023]** The production DB data store 250 stores data associated with a database that may represent for example, information representing daily transactions of an enterprise. The database server 245 is a computer program that provides database services and application programming interfaces (APIs) for managing data stored on the production DB data store 250. Mounting the production DB data store 250 on the database storage system 100 allows transfer of information stored on the production database system 110 to the database storage system 100.

**[0024]** The data stored in the storage system data store 290 can be exposed to a virtual database system 130 allowing the virtual database system 130 to treat the data as a copy of the production database stored in the production database system 110. The database storage system 100 includes a point-in-time copy manager 210, a transaction log manager 220, a storage allocation manager 265, a file sharing manager 270, a virtual database manager 275, a recovery milestone manager 280, a validation module 285, and a storage system data store 290. In alternative configurations, different and/or additional modules can be included in the database storage system 100.

**[0025]** The point-in-time copy manager 210 interacts with the production database system 110 by sending a request to retrieve information representing a point-in-time copy (also referred to as a “PIT copy”) of a database stored in the production DB data store 250. The point-in-time copy manager 210 stores the data obtained from the production database system 110 in the storage system data store 290. The data retrieved by the point-in-time copy manager 210 corresponds to database blocks (or pages) of the database being copied from the production DB data store 250. After a first PIT copy request to retrieve information from production DB data store 250, a subsequent PIT copy request may need to retrieve only the data that changed in the database since the previous request. The data collected in the first request can be combined with the data collected in a second request to reconstruct a copy of

the database corresponding to a point in time at which the data was retrieved from the production DB data store 250 for the second request.

**[0026]** The transaction log manager 220 sends request to the production database system 110 for retrieving portions of the transaction logs stored in the production database system 110. The data obtained by the transaction log manager 220 from the production database system 110 is stored in the storage system data store 290. In one embodiment, a request for transaction logs retrieves only the changes in the transaction logs in the production database system 110 since a previous request for the transaction logs was processed. The database blocks retrieved by a point in time copy manager 210 combined with the transaction logs retrieved by the transaction log manager 220 can be used to reconstruct a copy of a database in the production system 110 corresponding to times in the past in between the times at which point-in-time copies are made.

**[0027]** The storage allocation manager 265 provides the functionality of saving data retrieved from the production database system 110. For example, the point-in-time copy manager 210 may call APIs of storage allocation manager to save blocks of data retrieved from the production database system 110. The storage allocation manager 265 keeps track of the various versions of each block of data that may be obtained from the production database system 110. For a given time point, the storage allocation manager 265 can be requested to provide the latest version of a block of data obtained before the given time point. The storage allocation manager 265 can also be used for making copies of blocks of data. If a block of data is copied for read-only purposes, the storage allocation manager 265 allocates only sufficient storage to keep a pointer of reference to the existing block of data. However, if an attempt to write to the copied block of data is made, the storage allocation manager 265 allocates sufficient storage to make an actual copy of the block of data to avoid updating the original block of data.

**[0028]** The file sharing manager 270 allows files stored in the storage system data store 290 to be shared across computers that may be connected with the database storage system 100 over the network. The file sharing manager 270 uses the file sharing system 120 for sharing files. An example of a system for sharing files is a network file system (NFS). A system for sharing files may utilize fiber channel Storage area networks (FC-SAN) or network attached storage (NAS) or combinations and variations thereof. The system for sharing files may be based on small computer system interface (SCSI) protocol, internet small computer system interface (iSCSI) protocol, fiber channel protocols or other similar and related protocols. In some embodiments, the database storage system 100 may utilize a

logical volume manager. Sharing a file stored in the storage system data store 290 using the file sharing manager 270 allows a remote computer, for example, the virtual database systems 130 to access the data in the shared file. A remote system may be able to read and write from/to the file shared by the storage system data store 290. In an embodiment, files are organized in a format emulating a given file system disk layout, such as the file system of WINDOWS operating system called NTFS or the UNIX file system (UFS).

**[0029]** The recovery milestone manager 280 creates recovery milestones after a point-in-time copy is received for a production database. Each recovery milestone is created by applying transaction logs to the database blocks received as a point-in-time copy of the production database. The recovery milestone manager 280 invokes the transaction log manager 220 to apply transaction logs to database blocks. The recovery milestone may be stored as a set of files storing the set of database blocks. In an embodiment, the implementation of the recovery milestone is similar to (or same as) the implementation of a point-in-time copy of the database.

**[0030]** The validation module 285, verifies if a recovery milestone can be used to create a consistent snapshot. If the recovery milestone cannot be used to create a valid database snapshot, the validation module 285 executes steps to fix the error. For example, the validation module 285 executes steps to receive appropriate data from the production database. If the problem was caused by a bug in a software module, the validation module 285 may take steps to fix the software module, for example, by obtaining a corrected version of the software module from the appropriate vendor. In an embodiment, the validation module 285 provides appropriate information to a system administrator to take the required steps. In an embodiment, the data blocks stored in the storage system for a recovery milestone are same as the data blocks stored on the storage system for a virtual database created based on the recovery milestone. However, the virtual database may include additional information required to start up a database server based on the recovery milestone, for example, certain configuration information.

**[0031]** The virtual database manager 275 receives requests for creation of a virtual database for a virtual database system 130. The request for creation of a virtual database may be sent by a database administrator using the administration system 140 and identifies a production database system 110, a virtual database system 130, and includes a past point-in-time corresponding to which a virtual database needs to be created. The virtual database manager 275 creates the files corresponding to the virtual database being created and shares the files with the virtual database system 130. The database administrator for a virtual

database system 130 may be different from a database administrator for the production database system 110. To create a virtual database from a recovery milestone, virtual database manager 275 computes from metadata, the minimum set of logs that need to be applied to a recovery milestone to make it consistent.

**[0032]** A virtual database system 130 includes a database server 260. The database server 260 is similar in functionality to the database server 245 and is a computer program that provides database services and application programming interfaces (APIs) for managing data stored on a production DB data store 250. The data managed by the database server 260 may be stored on the storage system data store 290 that is shared by the database storage system 100 using a file sharing system 120. In alternative configurations, different and/or additional modules can be included in a virtual database system 130.

**[0033]** It should also be appreciated that in practice at least some of the components of the database storage system 100 may be distributed over multiple computers, communicating over a network. For convenience of explanation, however, the components of the database storage system 100 are discussed as though they were implemented on a single computer. In another embodiment, certain components may be located on a separate system that is coupled to the database storage system 100 by a network. The database storage system 100 may also include one or more input/output devices that allow data to be input and output to and from the system. It will be understood that embodiments of the database storage system 100 also include standard software and hardware components such as operating systems and the like and further include standard hardware components (e.g., network interfaces, storage devices, etc.) not shown in the figure for clarity of example.

**[0034]** FIG. 3 illustrates a process of provisioning a virtual database upon request from a point-in-time copy of the production database. The point-in-time copy manager 210 receives 305 a point-in-time copy of the production database at time T1. The virtual database manager 275 provisions 320 a virtual database upon request from a system administrator. The request for provisioning a virtual database may be received at any time after the point-in-time copy of the production database is received at time T1. The transaction log manager 220 retrieves transaction logs and applies the logs to the point-in-time copy of the production database received to create 310 a DB snapshot. The transaction logs may be received periodically by the transaction manager. However the transaction logs are applied to the point-in-time copy after the request to provision the virtual database is received. Accordingly, a significant amount of transaction logs may have to be applied to the point-in-time copy to provision the virtual database. In an embodiment, the request for provisioning

the virtual database specifies a point in time for which the virtual database needs to be provisioned. The transaction logs are provisioned up to the point in time specified in the request. In some embodiments, all transaction logs up to the current time at which the request is received may be applied.

**[0035]** Since transaction logs corresponding to a long time interval are applied to the point-in-time copy of the database, occasionally some of the logs that need to be applied may be corrupted. For example, some logs may be missing or overwritten by other data. The logs may be corrupted due to a bug or defect in the software processing the logs or possibly due to loss of data during transfer. The transaction log manager 220 may recognize that there is corruption of data after applying the logs to the point-in-time copy. For example, the transaction log manager 220 may recognize that applying the logs to the point-in-time copy does not provide a consistent snapshot of the database. In this situation, certain procedure may have to be performed to recover from the invalid data. The procedure may involve human involvement or may be automatic. However, the procedure may require time before a valid state of data is achieved and the requested virtual database can be provisioned. All this delay is experienced after the request for provisioning the virtual database is received. Accordingly, the end user may have to wait for significant amount of time before the end user receives the requested virtual database.

**[0036]** FIG. 4 illustrates a process of pre-provisioning recovery milestones to provision a validated database snapshot, in accordance with an embodiment of the invention. The process illustrated in FIG. 4 shows a time line indicating various steps that are performed to create multiple recovery milestones and to provision a virtual database based on a recovery milestone. At time T<sub>0</sub>, database blocks corresponding to a point-in-time copy of the production database are received. These database blocks correspond to the database blocks of the production database that changed since the last point-in-time copy was received.

**[0037]** A first recovery milestone is created 410a at time T<sub>1</sub> right after the point-in-time copy of the production database is received once enough transaction logs are received to make create a consistent snapshot. Subsequently recovery milestones 410b, 410c, 410d, may be created based on a predefined schedule, for example, on a regular basis. The frequency at which the recovery milestones are created is typically higher than the frequency at which point-in-time copies of databases are received from the production database. For example, if a point-in-time copy of the production database is received daily, the recovery milestone may be created on an hourly basis or once every 30 minutes. Creating a recovery milestone comprises applying 430 transaction logs received from the production database to the

previous recovery milestone created or to the point-in-time copy of the database received in the case of the first recovery milestone created after receiving the point-in-time copy of the database.

**[0038]** The validation module 285 validates each recovery milestone created to determine whether a valid database snapshot can be created based on the recovery milestone. If the validation module 285 identifies 420 a problem, the validation module 285 performs any required procedure to correct the problem. In an embodiment, the validation module 285 may create a virtual database for purposes of validation. For example, the validation may be performed by executing a test query on a virtual database based on the recovery milestone. Alternatively, the process of starting up a database server based on the virtual database may indicate whether the virtual database is valid.

**[0039]** When a request to provision a virtual database corresponding to time TN is received, the virtual database manager 275 identifies the most recent recovery milestone that was created before the time TN (e.g., time T4 as shown in FIG. 4) and creates 310b a database snapshot based on the identified recovery milestone. In an embodiment, if a virtual database is created by applying transaction logs to the recent most recovery milestone, a new recovery milestone is created corresponding to the requested virtual database. The new recovery milestone can be used to create subsequent recovery milestones.

**[0040]** FIG. 5 is a flowchart of the process for pre-provisioning validated snapshots of a database, in accordance with an embodiment of the invention. The point-in-time copy manager 210 receives 510 a point-in-time copy of the production database. The point-in-time copy may be received based on a predefined schedule or on demand based on a user command. For example, the point-in-time copy manager 210 may receive the point-in-time copy from the production database periodically, for example, every day at a particular time or once every few days. The transaction log manager 220 periodically retrieves transaction logs from the production database. Typically the rate at which the transaction log manager 220 retrieves the transaction logs is higher than the rate at which the point-in-time copy of the production database is retrieved. The point-in-time copy manager 210 waits 535 for a scheduled delay before receiving the next point-in-time copy. The point-in-time copy of the production database comprises the database blocks of the production database that changed since the previous point-in-time copy was obtained.

**[0041]** The recovery milestone manager 280 creates recovery milestones on a periodic basis. The frequency at which the recovery milestone manager 280 creates recovery milestones is higher than the frequency at which the point-in-time copy is retrieved. As a

result, several recovery milestones are typically created corresponding to each point-in-time copy of the production database that is retrieved. In an embodiment, the recovery milestone manager 280 creates a recovery milestone after getting the point-in-time copy of the production database right after the transaction log manager 220 retrieves logs that are needed for creating a consistent snapshot. In an embodiment, the recovery milestone manager 280 invokes an application programming interface (API) for executing a process for determining if the logs retrieved can create a consistent snapshot. The meta data associated with the logs is analyzed to determine if they contain the changes needed to make the snapshot consistent. The recoverability is computed a priori to determine whether required logs are missing.

**[0042]** For example, as shown in FIG. 4, a recovery milestone is created 410a at time T1 right after the point in time copy is received 305b at time T0. Subsequently, the recovery milestones are created 515 based on a predefined schedule. For example, as shown in FIG. 5, the recovery milestone manager 280 waits 540 for a scheduled delay before creating the next recovery milestone.

**[0043]** The recovery milestone manager 280 applies the database logs obtained by the transaction log manager 220 to the point-in-time copy of the production database received 510 to obtain a recovery milestone. Subsequently, the recovery milestone manager 280 applies the database logs obtained by the transaction log manager 220 to the most recent recovery milestone. The validation module 285 performs 520 validation of the data obtained from the production database. For example, the validation module 285 may determine whether there is any corruption in the transaction logs obtained by the transaction log manager 220. If the transaction logs retrieved are corrupted, the recovery milestone manager 280 is unable to obtain a consistent snapshot of a database using the transaction logs. The corruption in the transaction logs may occur because a portion of the transaction logs is missing or overwritten by some other data.

**[0044]** If the validation module 285 determines 525 that the data retrieved is valid and a consistent database snapshot can be created using the retrieved data, the recovery milestone manager waits 540 for the scheduled delay before creating the next recovery milestone. If the validation module 285 determines 525 that the data obtained from the production database is not valid, the validation module 285 may perform 530 certain procedure to fix the cause of the invalid data. The validation module 285 may request the transaction logs from the production system again and verify whether the subsequent copy of the transaction logs is valid. In an embodiment, the validation module 285 presents required information to a system administrator to perform procedures to recover from any error that may have

occurred. It is possible that the corruption in data occurred due to some defect or bug in the software used during one of the above steps, for example, the software executing the production database system. In this situation, the system administrator may have to perform certain updates to the software to fix the problem. Once the software has been properly upgraded, the system may be able to proceed with the steps of the flowchart illustrated in FIG. 5. These steps can take significant time in order to perform 530 the procedure to fix certain problems. In an embodiment, if the system determines that a milestone cannot be created because of missing logs, the system requests the missing logs or provides information to a system administrator enabling the system administrator to retrieve the missing logs.

**[0045]** The above process allows problems with the data to be fixed in advance of a request to provision a virtual database. For example, without the above process it is possible that the problem is recognized after a request for provisioning a virtual database is received. In this situation, the requestor of the virtual database may have to wait for significant time while the procedure for fixing any problem is fixed before the virtual database is made available to the requestor.

**[0046]** The steps illustrated in the flowchart shown in FIG. 5 may be performed concurrently or in a different order than that shown in FIG. 5. For example, certain validation steps may be performed 520 before the recovery milestone is created 515. Also, certain steps may be performed by different modules instead of the modules described above. The creation of recovery milestones can be policy driven or on-demand based on requests.

**[0047]** In some embodiments, the production database may be set up so that logging is turned off for one or more tables. If logging is completely turned off for all tables, the pre-provisioning of virtual databases may not be performed by applying logs since logs are not available. However, in some situations, logging may be turned off for some tables but available for other tables. In these situations, the pre-provisioning can be performed such that the user can query only tables for which logging is available in the virtual database that is created. If the user attempts to query tables for which logging is turned off, the user may not get latest information or even a consistent snapshot for these tables. Also, the user may not be able to execute queries that join tables with logging on and tables with logging off in a virtual database generated from a pre-provisioned recovery milestone. For tables that have logging turned off, latest information describing the table is obtained by receiving database blocks from the production database rather than by receiving transaction logs.

Additional Configuration Considerations

**[0048]** Throughout this specification, plural instances may implement components, operations, or structures described as a single instance. Although individual operations of one or more methods are illustrated and described as separate operations, one or more of the individual operations may be performed concurrently, and nothing requires that the operations be performed in the order illustrated. Structures and functionality presented as separate components in example configurations may be implemented as a combined structure or component. Similarly, structures and functionality presented as a single component may be implemented as separate components. These and other variations, modifications, additions, and improvements fall within the scope of the subject matter herein.

**[0049]** Certain embodiments are described herein as including logic or a number of components, modules, or mechanisms. Modules may constitute either software modules (e.g., code embodied on a machine-readable medium or in a transmission signal) or hardware modules. A hardware module is tangible unit capable of performing certain operations and may be configured or arranged in a certain manner. In example embodiments, one or more computer systems (e.g., a standalone, client or server computer system) or one or more hardware modules of a computer system (e.g., a processor or a group of processors) may be configured by software (e.g., an application or application portion) as a hardware module that operates to perform certain operations as described herein.

**[0050]** In various embodiments, a hardware module may be implemented mechanically or electronically. For example, a hardware module may comprise dedicated circuitry or logic that is permanently configured (e.g., as a special-purpose processor, such as a field programmable gate array (FPGA) or an application-specific integrated circuit (ASIC)) to perform certain operations. A hardware module may also comprise programmable logic or circuitry (e.g., as encompassed within a general-purpose processor or other programmable processor) that is temporarily configured by software to perform certain operations. It will be appreciated that the decision to implement a hardware module mechanically, in dedicated and permanently configured circuitry, or in temporarily configured circuitry (e.g., configured by software) may be driven by cost and time considerations.

**[0051]** Accordingly, the term “hardware module” should be understood to encompass a tangible entity, be that an entity that is physically constructed, permanently configured (e.g., hardwired), or temporarily configured (e.g., programmed) to operate in a certain manner or to perform certain operations described herein. As used herein, “hardware-implemented module” refers to a hardware module. Considering embodiments in which hardware modules

are temporarily configured (e.g., programmed), each of the hardware modules need not be configured or instantiated at any one instance in time. For example, where the hardware modules comprise a general-purpose processor configured using software, the general-purpose processor may be configured as respective different hardware modules at different times. Software may accordingly configure a processor, for example, to constitute a particular hardware module at one instance of time and to constitute a different hardware module at a different instance of time.

**[0052]** Hardware modules can provide information to, and receive information from, other hardware modules. Accordingly, the described hardware modules may be regarded as being communicatively coupled. Where multiple of such hardware modules exist contemporaneously, communications may be achieved through signal transmission (e.g., over appropriate circuits and buses) that connect the hardware modules. In embodiments in which multiple hardware modules are configured or instantiated at different times, communications between such hardware modules may be achieved, for example, through the storage and retrieval of information in memory structures to which the multiple hardware modules have access. For example, one hardware module may perform an operation and store the output of that operation in a memory device to which it is communicatively coupled. A further hardware module may then, at a later time, access the memory device to retrieve and process the stored output. Hardware modules may also initiate communications with input or output devices, and can operate on a resource (e.g., a collection of information).

**[0053]** The various operations of example methods described herein may be performed, at least partially, by one or more processors that are temporarily configured (e.g., by software) or permanently configured to perform the relevant operations. Whether temporarily or permanently configured, such processors may constitute processor-implemented modules that operate to perform one or more operations or functions. The modules referred to herein may, in some example embodiments, comprise processor-implemented modules.

**[0054]** Similarly, the methods described herein may be at least partially processor-implemented. For example, at least some of the operations of a method may be performed by one or processors or processor-implemented hardware modules. The performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but deployed across a number of machines. In some example embodiments, the processor or processors may be located in a single location (e.g., within a

home environment, an office environment or as a server farm), while in other embodiments the processors may be distributed across a number of locations.

**[0055]** The one or more processors may also operate to support performance of the relevant operations in a “cloud computing” environment or as a “software as a service” (SaaS). For example, at least some of the operations may be performed by a group of computers (as examples of machines including processors), these operations being accessible via a network (e.g., the Internet) and via one or more appropriate interfaces (e.g., application program interfaces (APIs).)

**[0056]** The performance of certain of the operations may be distributed among the one or more processors, not only residing within a single machine, but deployed across a number of machines. In some example embodiments, the one or more processors or processor-implemented modules may be located in a single geographic location (e.g., within a home environment, an office environment, or a server farm). In other example embodiments, the one or more processors or processor-implemented modules may be distributed across a number of geographic locations.

**[0057]** Some portions of this specification are presented in terms of algorithms or symbolic representations of operations on data stored as bits or binary digital signals within a machine memory (e.g., a computer memory). These algorithms or symbolic representations are examples of techniques used by those of ordinary skill in the data processing arts to convey the substance of their work to others skilled in the art. As used herein, an “algorithm” is a self-consistent sequence of operations or similar processing leading to a desired result. In this context, algorithms and operations involve physical manipulation of physical quantities. Typically, but not necessarily, such quantities may take the form of electrical, magnetic, or optical signals capable of being stored, accessed, transferred, combined, compared, or otherwise manipulated by a machine. It is convenient at times, principally for reasons of common usage, to refer to these signals using words such as “data,” “content,” “bits,” “values,” “elements,” “symbols,” “characters,” “terms,” “numbers,” “numerals,” or the like. These words, however, are merely convenient labels and are to be associated with appropriate physical quantities.

**[0058]** Unless specifically stated otherwise, discussions herein using words such as “processing,” “computing,” “calculating,” “determining,” “presenting,” “displaying,” or the like may refer to actions or processes of a machine (e.g., a computer) that manipulates or transforms data represented as physical (e.g., electronic, magnetic, or optical) quantities within one or more memories (e.g., volatile memory, non-volatile memory, or a combination

thereof), registers, or other machine components that receive, store, transmit, or display information.

**[0059]** As used herein any reference to “one embodiment” or “an embodiment” means that a particular element, feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment. The appearances of the phrase “in one embodiment” in various places in the specification are not necessarily all referring to the same embodiment.

**[0060]** Some embodiments may be described using the expression “coupled” and “connected” along with their derivatives. It should be understood that these terms are not intended as synonyms for each other. For example, some embodiments may be described using the term “connected” to indicate that two or more elements are in direct physical or electrical contact with each other. In another example, some embodiments may be described using the term “coupled” to indicate that two or more elements are in direct physical or electrical contact. The term “coupled,” however, may also mean that two or more elements are not in direct contact with each other, but yet still cooperate or interact with each other. The embodiments are not limited in this context.

**[0061]** As used herein, the terms “comprises,” “comprising,” “includes,” “including,” “has,” “having” or any other variation thereof, are intended to cover a non-exclusive inclusion. For example, a process, method, article, or apparatus that comprises a list of elements is not necessarily limited to only those elements but may include other elements not expressly listed or inherent to such process, method, article, or apparatus. Further, unless expressly stated to the contrary, “or” refers to an inclusive or and not to an exclusive or. For example, a condition A or B is satisfied by any one of the following: A is true (or present) and B is false (or not present), A is false (or not present) and B is true (or present), and both A and B are true (or present).

**[0062]** In addition, use of the “a” or “an” are employed to describe elements and components of the embodiments herein. This is done merely for convenience and to give a general sense of the invention. This description should be read to include one or at least one and the singular also includes the plural unless it is obvious that it is meant otherwise.

**[0063]** Upon reading this disclosure, those of skill in the art will appreciate still additional alternative structural and functional designs for a system and a process for creating virtual databases from point-in-time copies of production databases stored in a storage manager. Thus, while particular embodiments and applications have been illustrated and described, it is to be understood that the disclosed embodiments are not limited to the precise construction

and components disclosed herein. Various modifications, changes and variations, which will be apparent to those skilled in the art, may be made in the arrangement, operation and details of the method and apparatus disclosed herein without departing from the spirit and scope defined in the appended claims.

What is claimed is:

1. A method for pre-provisioning validated database snapshots, the method comprising:
  - receiving a plurality of database blocks corresponding to the source database;
  - storing the database blocks on a storage system;
  - receiving transaction logs from the source database corresponding to transactions executed in the source database;
  - creating multiple recovery milestones without having received a request for provisioning a virtual database, each recovery milestone comprising database blocks from which a consistent snapshot of source database corresponding to a point in time can be created, wherein at least some of the database blocks are associated with multiple recovery milestones, the creating of each recovery milestone comprising:
    - applying a set of transaction logs to either a plurality of database blocks received from the source database or a plurality of database blocks of a previous recovery milestone; and
    - validating the recovery milestone created to determine whether the recovery milestone can be used to create a consistent snapshot of the source database; and
  - responsive to receiving a request to provision a virtual database corresponding to a point in time, identifying a recovery milestone corresponding to the point in time, and provisioning the requested virtual database based on the recovery milestone.
2. The method of claim 1, wherein responsive to the point in time of the requested virtual database matching the point in time of a recovery milestone, provisioning the requested virtual database comprising the database blocks of the matching recovery milestone.
3. The method of claim 1, wherein creating each recovery milestone comprises:
  - responsive to the point in time of the requested virtual database not matching point in times of any of the recovery milestones, selecting a recovery milestone with a point in time closest to and occurring prior to the requested point in time; and

applying transaction logs to the database blocks of the selected recovery milestone to create a database snapshot corresponding to the requested point in time; and provisioning the requested virtual database comprising the database blocks of the database snapshot created.

4. The method of claim 1, wherein the multiple recovery milestones are created based on a predefined schedule.
5. The method of claim 1, wherein validating the recovery milestone comprises: responsive to determining that a set of transaction logs corresponding to the recovery milestone are corrupted, determining that the recovery milestone is invalid.
6. The method of claim 1, wherein validating the recovery milestone comprises: responsive to the determining that the recovery milestone is invalid, executing steps to identify the defect resulting in creation of the invalid recovery milestone and executing steps to fix the defect.
7. The method of claim 6, wherein validating the recovery milestone comprises: responsive to the validating determining that an invalid recovery milestone was created, receiving a recent version of a vendor library to create a corresponding valid recovery milestone.
8. The method of claim 1, wherein provisioning the requested virtual database based on the recovery milestone comprises applying a first set of transaction logs to the recovery milestone, wherein the first set of transaction logs is smaller than a set of transaction logs applied to the point-in-time copy of the source database to obtain the same recovery milestone.
9. The method of claim 1, further comprising: determining a rate at which point-in-time copies of the source database are received when recovery milestones are created as lower than a rate at which point-in-time copies of the source database are received if no recovery milestones are created.
10. A method for pre-provisioning validated database snapshots, the method comprising:  
receiving a plurality of database blocks corresponding to the source database;  
storing the database blocks on a storage system;

receiving transaction logs from the source database corresponding to transactions executed in the source database; and  
creating multiple recovery milestones without having received a request for provisioning a virtual database, each recovery milestone comprising database blocks from which a consistent snapshot of source database corresponding to a point in time can be created, wherein at least some of the database blocks are associated with multiple recovery milestones, the creating of each recovery milestone comprising:  
applying a set of transaction logs to either a plurality of database blocks received from the source database or a plurality of database blocks of a previous recovery milestone.

11. The method of claim 10, further comprising:

responsive to receiving a request to provision a virtual database corresponding to a point in time, identifying a recovery milestone corresponding to the point in time, and provisioning the requested virtual database based on the recovery milestone.

12. The method of claim 11, wherein provisioning the requested virtual database based on the recovery milestone comprises applying a first set of transaction logs to the recovery milestone, wherein the first set of transaction logs is smaller than a set of transaction logs applied to the point-in-time copy of the source database to obtain the same recovery milestone.

13. The method of claim 10, further comprising:

receiving a request to provision a virtual database corresponding to a point in time;  
responsive to the point in time of the requested virtual database not matching point in times of any of the recovery milestones, selecting a recovery milestone with a point in time closest to and occurring prior to the requested point in time;  
applying transaction logs to the database blocks of the selected recovery milestone to create a database snapshot corresponding to the requested point in time;  
and  
provisioning the requested virtual database comprising the database blocks of the database snapshot created.

14. The method of claim 10, wherein creating each recovery milestone further comprises:
- validating the recovery milestone created to determine whether the recovery milestone can be used to create a consistent snapshot of the source database;
15. The method of claim 14, wherein validating the recovery milestone comprises: responsive to determining that a set of transaction logs corresponding to the recovery milestone are corrupted, determining that the recovery milestone is invalid.
16. The method of claim 14, wherein validating the recovery milestone comprises: responsive to the determining that the recovery milestone is invalid, executing steps to identify the defect resulting in creation of the invalid recovery milestone and executing steps to fix the defect.
17. The method of claim 10, further comprising:
- determining a rate at which point-in-time copies of the source database are received when recovery milestones are created as lower than a rate at which point-in-time copies of the source database are received if no recovery milestones are created.
18. A computer program product having a non-transitory computer-readable storage medium storing computer program modules configured to execute on the computer processor, the computer program modules comprising:
- a point-in-time copy manager configured to:
    - receive a plurality of database blocks corresponding to the source database and storing the database blocks on a storage system;
  - a transaction log manager configured to:
    - receive transaction logs from the source database corresponding to transactions executed in the source database;
  - a recovery milestone manager configured to:
    - create multiple recovery milestones in advance of receiving a request for provisioning a virtual database, each recovery milestone comprising database blocks from which a consistent snapshot of source database corresponding to a point in time can be created, wherein at least some of the database blocks are

associated with multiple recovery milestones, the creating of each recovery milestone comprising:

applying a set of transaction logs to either a plurality of database blocks received from the source database or a plurality of database blocks of a previous recovery milestone; and

validating the recovery milestone created to determine whether the recovery milestone can be used to create a consistent snapshot of the source database; and

a virtual database manager configured to:

responsive to receiving a request to provision a virtual database corresponding to a point in time, identify a recovery milestone corresponding to the point in time, and provision the requested virtual database based on the recovery milestone.

19. The computer program product of claim 18, wherein creating each recovery milestone comprises:

selecting a recovery milestone with a point in time closest to and occurring prior to the requested point in time; and

applying transaction logs to the database blocks of the selected recovery milestone to create a database snapshot corresponding to the requested point in time; and

provisioning the requested virtual database comprising the database blocks of the database snapshot created.

20. The computer program product of claim 18, wherein the multiple recovery milestones are created based on a predefined schedule.

21. The computer program product of claim 18, wherein validating the recovery milestone comprises:

responsive to determining that a set of transaction logs corresponding to the recovery milestone are corrupted, determining that the recovery milestone is invalid.

22. The computer program product of claim 18, wherein validating the recovery milestone comprises:

responsive to the determining that the recovery milestone is invalid, executing steps to identify the defect resulting in creation of the invalid recovery milestone and executing steps to fix the defect.

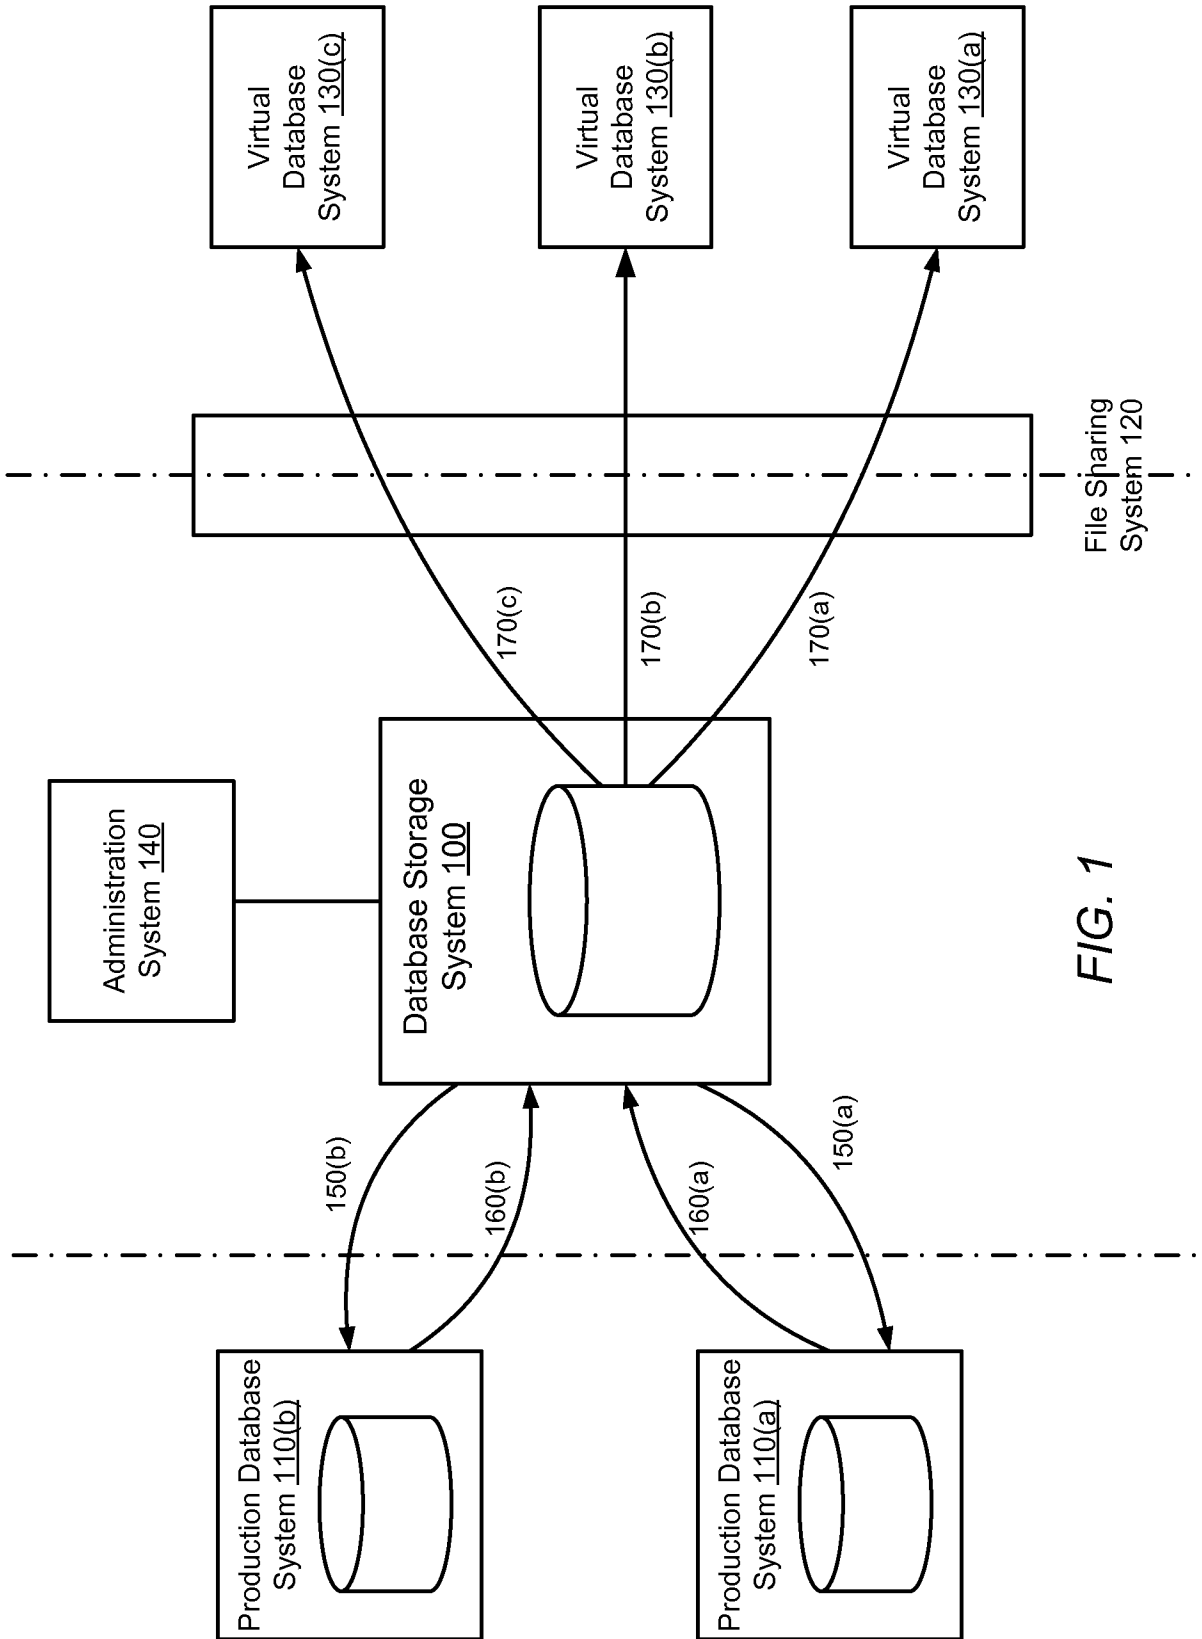


FIG. 1

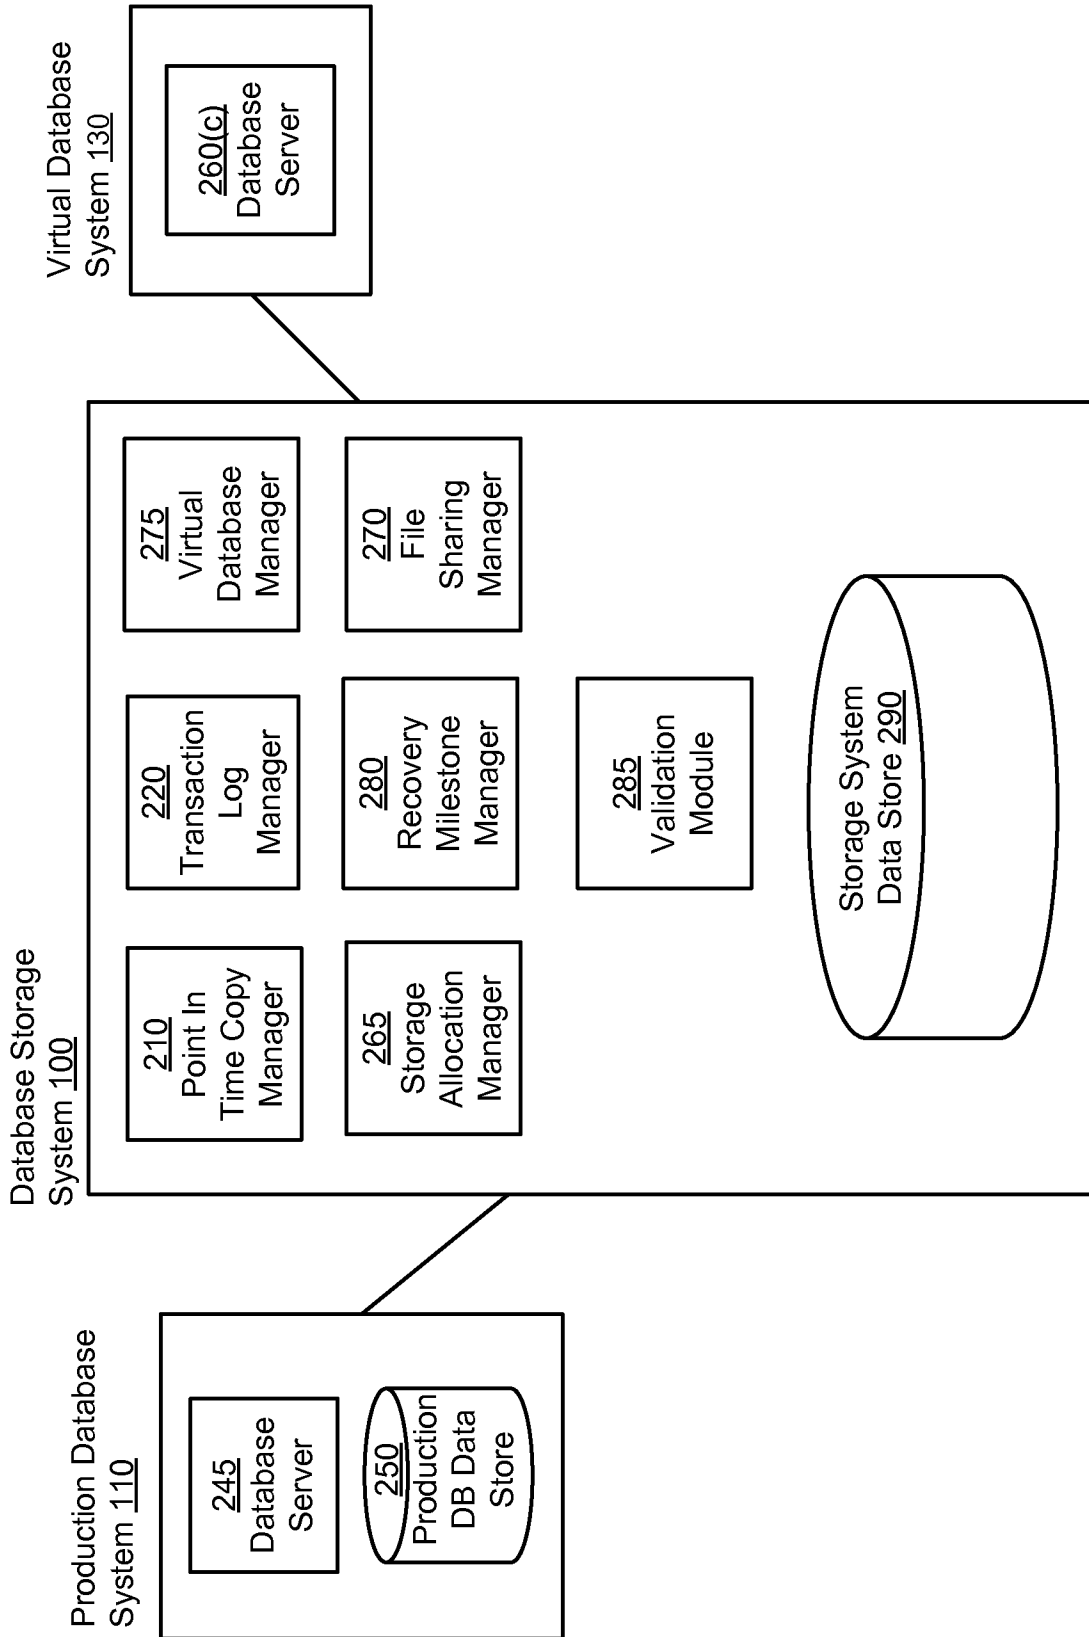


FIG. 2

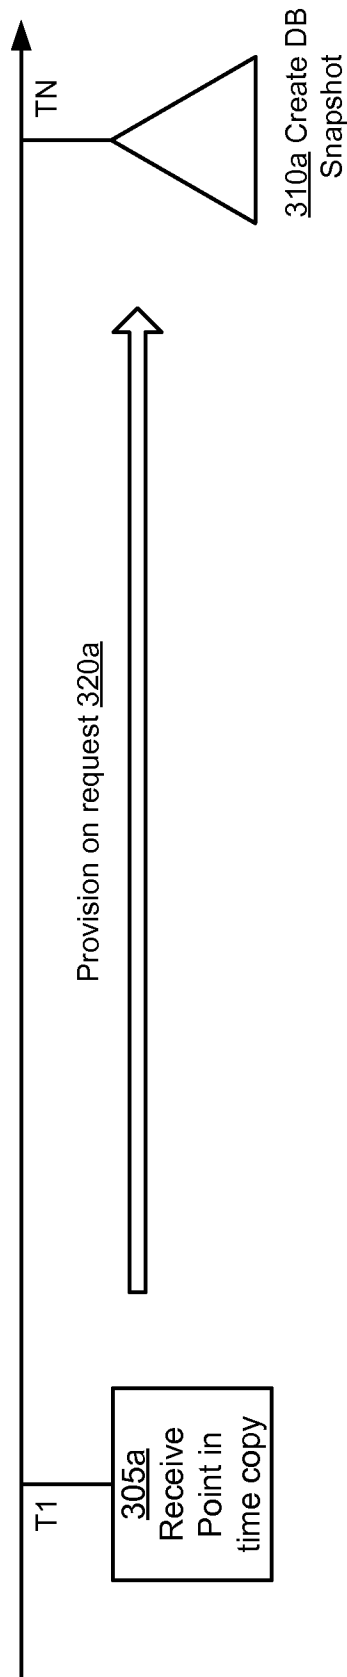


FIG. 3

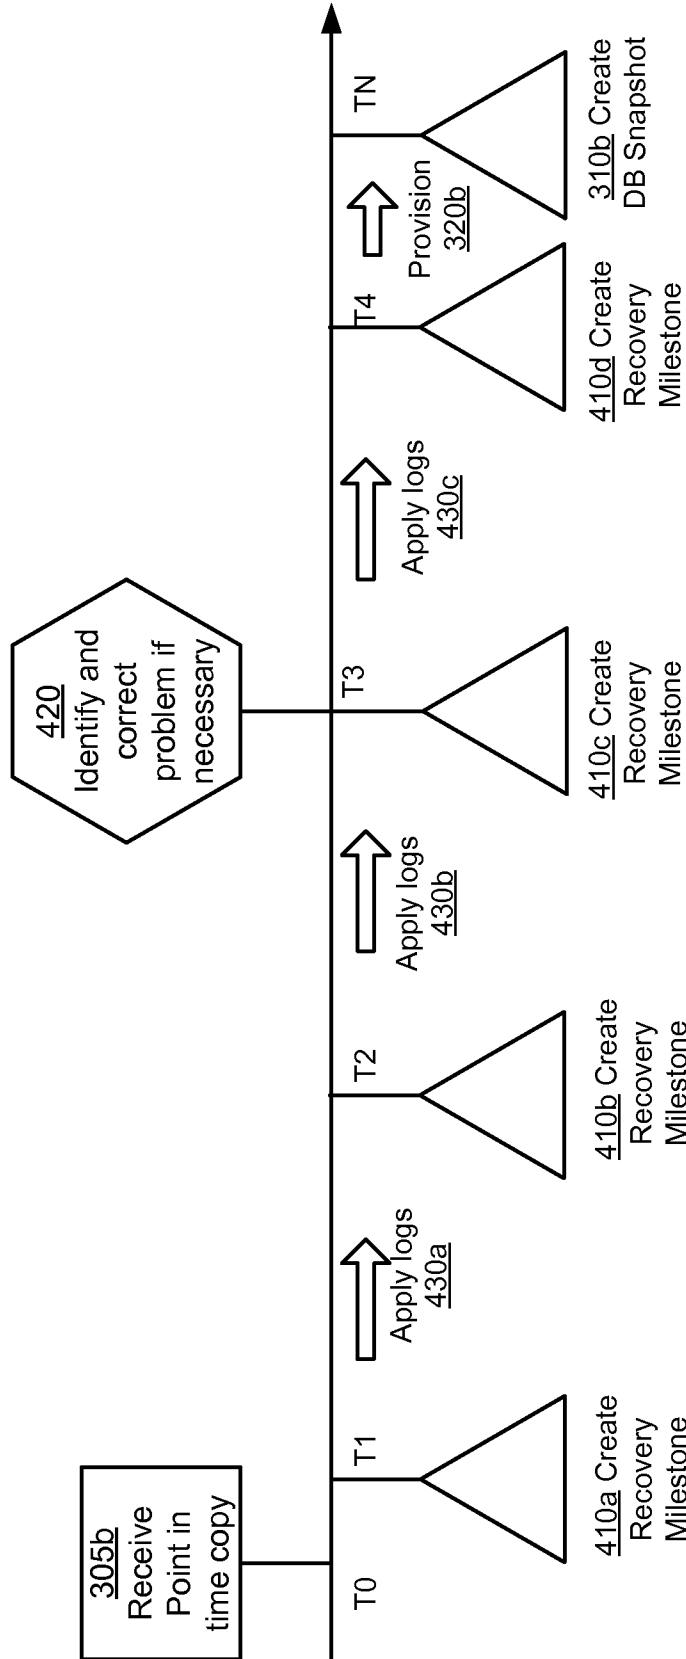


FIG. 4

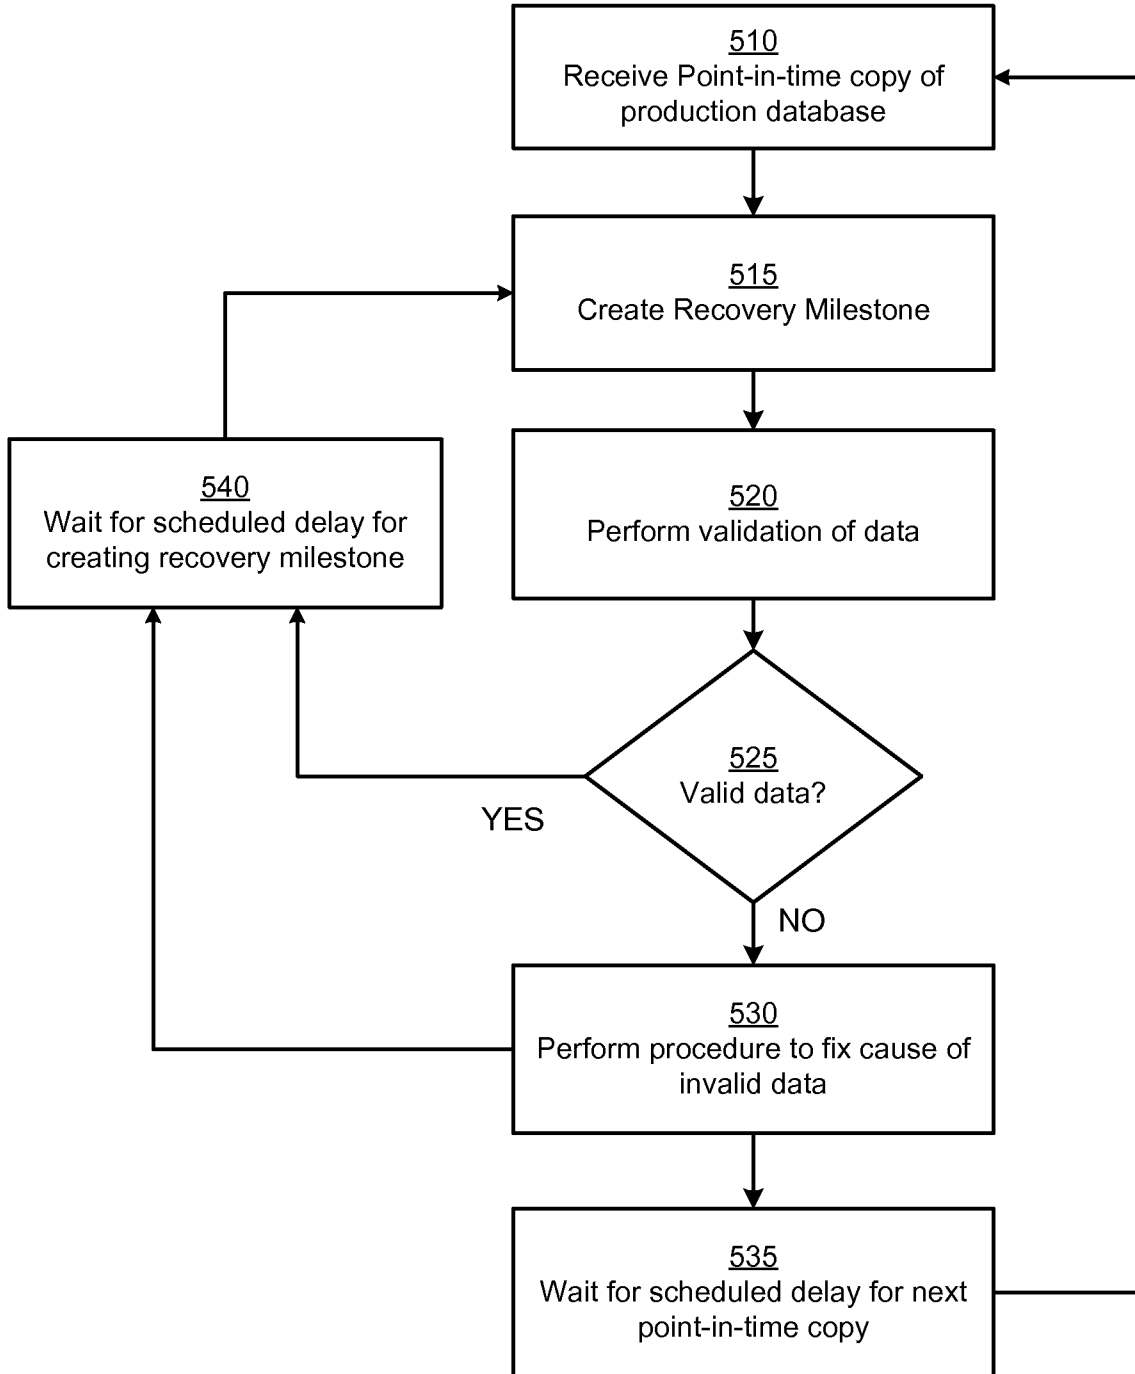


FIG. 5