



(51) International Patent Classification:

G06F 7/00 (2006.01)

G06Q 10/06 (2012.01)

G06F 17/00 (2019.01)

(21) International Application Number:

PCT/US2019/059798

(22) International Filing Date:

05 November 2019 (05.11.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/182,112

06 November 2018 (06.11.2018) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:

US

16/182,112 (CON)

Filed on

06 November 2018 (06.11.2018)

(71) Applicant: SNOWFLAKE INC. [US/US]; 450 Concar Dr., San Mateo, California 94402 (US).

(72) Inventors: CSERI, Istvan; 450 Concar Dr., San Mateo, California 94402 (US). GRABS, Torsten; 450 Concar Dr., San Mateo, California 94402 (US). DAGEVILLE, Benoit; 450 Concar Dr., San Mateo, California 94402 (US).

(74) Agent: SCHEER, Bradley W. et al.; P.O. Box 2938, Minneapolis, Minnesota 55402 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, 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, ST, 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:

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

(54) Title: TRACKING CHANGES IN DATABASE DATA




FIG. 8

(57) Abstract: Systems, methods, and devices for tracking changes to database data. A method includes determining a change to be executed on a micro-partition of a table of a database and executing the change on the table by generating a new micro-partition that embodies the change. The method includes updating a table history that includes a log of changes made to the table, wherein each change in the log of changes includes a timestamp, and wherein updating the table history includes inserting the change into the log of changes.



TRACKING CHANGES IN DATABASE DATA

CROSS-REFERENCE TO RELATED APPLICATION

[0001] This application claims priority to U.S. Patent Application Serial No. 16/182,112, filed on November 6, 2018, the contents of which are incorporated herein in their entirety.

TECHNICAL FIELD

[0002] The present disclosure relates systems, methods, and devices for databases and more particularly relates to tracking changes to database data.

BACKGROUND

[0003] Databases are widely used for data storage and access in computing applications. Databases may include one or more tables that include or reference data that can be read, modified, or deleted using database operations usually referred to as DML (Data Manipulation Language) statements. Databases can store anywhere from small to extremely large sets of data within one or more tables. This data can be accessed by various users in an organization or even be used to service public users, such as via a website or an application program interface (API). Both computing and storage resources, as well as their underlying architecture, can play a significant role in achieving desirable database performance.

[0004] Database data can be modified by various commands, including insert, delete, and update commands that modify one or more rows in a database table. It can be costly to track such modifications and to determine delta information between a first set of database data and a second set of database data. Systems, methods, and devices for efficient tracking of changes made to database data are disclosed herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] Non-limiting and non-exhaustive embodiments of the present disclosure are described with reference to the following figures, wherein like reference numerals refer to like parts throughout the various figures unless otherwise specified.

[0006] FIG. 1 is a block diagram illustrating modifications made to table versions, according to one embodiment;

[0007] FIG. 2 is a block diagram illustrating a delete command performed on a micro-partition of a database, according to one embodiment;

[0008] FIG. 3 is a block diagram illustrating an insert command performed on a micro-partition of a database, according to one embodiment;

[0009] FIG. 4 is a block diagram illustrating an update command performed on a micro-partition of a database, according to one embodiment;

[0010] FIG. 5 is a block diagram illustrating a database system having a database service manager, according to one embodiment;

[0011] FIG. 6 is a block diagram illustrating components of a database service manager, according to one embodiment;

[0012] FIG. 7 is a block diagram illustrating components of a change tracking manager, according to one embodiment;

[0013] FIG. 8 is a schematic flow chart diagram illustrating a method for tracking changes to database data, according to one embodiment;

[0014] FIG. 9 is a schematic flow chart diagram illustrating a method for tracking changes to database data, according to one embodiment; and

[0015] FIG. 10 is a block diagram depicting an example computing device consistent with at least one embodiment of processes and systems disclosed herein.

DETAILED DESCRIPTION

[0016] Systems, methods, and devices for tracking changes to database data and for returning a delta for a database table are disclosed. Database change tracking can return a summary of what rows have changed in a database table and how those rows have changed. The delta for a database table indicates which rows have changed and how they have changed between two timestamps and does not include information on all intermediate changes that occurred between the two timestamps. Change tracking auditing indicates a complete history of row modifications in a specified time period.

[0017] Change tracking is historically very costly and requires significant resources and storage capacity. Database systems are traditionally built using a transaction log that embodies all changes made in the database. The transactional log is normally used to make transactional changes atomic and durable and may further be used in recover in an instance where the database system crashes and information must be restored. The transactional includes all changes that have occurred on the database and may be used to reconstruct those changes. However, the transactional log approach requires significant system resource because items must be stored using a scan of the log because there is no random-access storage structure to speed up the process of reconstructing the database. Further, the transactional log itself is physical and must be parsed and applied to a page structure to provide whole row images and changes.

[0018] In an alternative approach known in the art, temporal information about data changes may be stored in two transactional timestamp columns within a database table. The transactional timestamp columns indicate transactional times when a row was valid. An internal history table may be generated that stores updated data information with the timestamps from the transactional timestamp columns. Updates and deletes copy the values of changed rows into the internal history table and update the transactional timestamp columns to mark the valid

transactional lifetime of rows. When the historical or changed data set is required, a current table and the internal history table must be joined to gather the delta or historical information between two transactional time points. The maintenance of the transactional timestamp columns and the internal history table adds significant overhead to the database transactions. Additionally, delta queries must join the information from two tables when requested.

[0019] Change tracking information can be stored as metadata in a database. This metadata describes the data that is stored in database tables of customers but is not actually the stored table data. Metadata can get very large, especially if there are large database tables of many customers. Current database systems have severe limitations handling large amounts of metadata. Current database systems store metadata in mutable storage devices and services, including main memory, file systems, and key-value stores. These devices and services allow the metadata to be updated data in-place. If a data record changes, it may be updated with the new information and the old information is overwritten. This allows databases to easily maintain mutable metadata by updating metadata in-place.

[0020] However, these mutable storage devices and services have limitations. The limitations are at least two-fold. First, mutable storage devices such as main memory and file systems have a hard limit in terms of storage capacity. If the size of the metadata exceeds these limits, it is impossible to store more metadata there. Second, mutable storage services such as key-value stores perform poorly when reading large volumes of metadata. Reading data is performed using range scans, which take a long time to finish. In practice, range scans can take many minutes or even approach an hour to complete in large scale deployments.

[0021] These limitations make it impossible to store large amounts of metadata in existing mutable storage devices and services. Systems, methods, and devices disclosed herein provide for improved metadata storage and management that includes storing metadata in immutable (non-mutable) storage such as micro-partitions. As used herein, immutable or non-mutable

storage includes storage where data cannot or is not permitted to be overwritten or updated in-place. For example, changes to data that is located in a cell or region of storage media may be stored as a new file in a different, time-stamped, cell or region of the storage media. Mutable storage may include storage where data is permitted to be overwritten or updated in-place. For example, data in a given cell or region of the storage media can be overwritten when there are changes to the data relevant to that cell or region of the storage media.

[0022] In one embodiment, metadata is stored and maintained on non-mutable storage services in the cloud. These storage services may include, for example, Amazon S3®, Microsoft Azure Blob Storage®, and Google Cloud Storage®. Many of these services do not allow to update data in-place (i.e., are non-mutable or immutable). Data files may only be added or deleted, but never updated. In one embodiment, storing and maintaining metadata on these services requires that, for every change in metadata, a metadata file is added to the storage service. These metadata files may be periodically consolidated into larger “compacted” or consolidated metadata files in the background.

[0023] In an embodiment, all data in tables is automatically divided into an immutable storage device referred to as a micro-partition. The micro-partition may be considered a batch unit where each micro-partition has contiguous units of storage. By way of example, each micro-partition may contain between 50 MB and 500 MB of uncompressed data (note that the actual size in storage may be smaller because data may be stored compressed). Groups of rows in tables may be mapped into individual micro-partitions organized in a columnar fashion. This size and structure allow for extremely granular pruning of very large tables, which can be comprised of millions, or even hundreds of millions, of micro-partitions. Metadata may be automatically gathered about all rows stored in a micro-partition, including: the range of values for each of the columns in the micro-partition; the number of distinct values; and/or additional properties used for both optimization and efficient query processing. In one embodiment,

micro-partitioning may be automatically performed on all tables. For example, tables may be transparently partitioned using the ordering that occurs when the data is inserted/loaded.

[0024] A detailed description of systems and methods consistent with embodiments of the present disclosure is provided below. While several embodiments are described, it should be understood that this disclosure is not limited to any one embodiment, but instead encompasses numerous alternatives, modifications, and equivalents. In addition, while numerous specific details are set forth in the following description to provide a thorough understanding of the embodiments disclosed herein, some embodiments may be practiced without some or all these details. Moreover, for the purpose of clarity, certain technical material that is known in the related art has not been described in detail to avoid unnecessarily obscuring the disclosure.

[0025] In an embodiment, a method for tracking changes to database data is disclosed. The method includes determining a change to be executed on a micro-partition of a table of a database, wherein the change may be any command and may include, for example, a delete command, an update command, a merge command, and/or an insert command. The method includes executing the change on the table by generating a new micro-partition that embodies the change and may further include removing the (original) micro-partition. The method includes updating a table history that comprises a log of changes made to the table, wherein each change in the log of changes comprises a timestamp, and wherein updating the table history comprises inserting the change into the log of changes. The method includes determining a delta for the table between a first timestamp and a second timestamp based on the table history. In an embodiment, the delta indicates what rows have changed and how those rows have changed between the first timestamp and the second timestamp and does not include information on intermediate changes made to those rows between the first timestamp and the second timestamp.

[0026] In an embodiment, the table history includes one or more change tracking columns stored within the micro-partition and/or the new micro-partition. The one or more change tracking columns may include information indicating one or more of: a prior micro-partition name for a row value, a prior row identification for a row value, an action taken on a row value, whether a row value was updated, a log of changes made to a row or row value, a lineage for a row indicating one or more changes made to the row over a time period, a listing of transactions that have initiated a modification to a row, and any other suitable information that may indicate change made to a row or micro-partition or a historical listing of changes that have been made to a row or micro-partition in a table. In an embodiment, the table history is stored in a metadata column that includes a micro-partition lineage storing a lineage of micro-partition names over a time period. The metadata column may further include a row lineage storing a lineage of ordinal numbers for a given row over a time period.

[0027] In an embodiment, a delta for a table is determined based on metadata stored in a table history that may include one or more change tracking columns. In such an embodiment, prior versions of a row or micro-partition may be removed from the database and the table history provides historical information on changes made to the row or micro-partition. Such an embodiment may provide change tracking information in an efficient manner that does not require significant storage capacity for storing historical versions of rows or micro-partitions.

[0028] In an embodiment, file metadata is stored within metadata storage. The file metadata contains table versions and information about each table data file. The metadata storage may include mutable storage (storage that can be overwritten or written in-place), such as a local file system, system, memory, or the like. In one embodiment, the micro-partition metadata consists of two data sets: table versions and file information. The table versions data set includes a mapping of table versions to lists of added files and removed files. File information consists of information about each micro-partition, including micro-partition path,

micro-partition size, micro-partition key id, and summaries of all rows and columns that are stored in the micro-partition, for example. Each modification of the table creates new micro-partitions and new micro-partition metadata. Inserts into the table create new micro-partitions. Deletes from the table remove micro-partitions and potentially add new micro-partitions with the remaining rows in a table if not all rows in a micro-partition were deleted. Updates remove micro-partitions and replace them with new micro-partitions with rows containing the updated records.

[0029] In one embodiment, metadata may be stored in metadata micro-partitions in immutable storage. In one embodiment, a system may write metadata micro-partitions to cloud storage for every modification of a database table. In one embodiment, a system may download and read metadata micro-partitions to compute the scanset. The metadata micro-partitions may be downloaded in parallel and read as they are received to improve scan set computation. In one embodiment, a system may periodically consolidate metadata micro-partitions in the background. In one embodiment, performance improvements, including pre-fetching, caching, columnar layout and the like may be included. Furthermore, security improvements, including encryption and integrity checking, are also possible with metadata files with a columnar layout.

[0030] FIG. 1 illustrates a schematic block diagram of a timeline 100 of table versions. The example timeline 100 illustrates three table versions, namely table version 1, table version 2, and table version 3. Table version 1 includes data in the form of three micro-partitions (MPs), namely micro-partition 1 (MP1), micro-partition 2 (MP2), and micro-partition 3 (MP3). A change 102 is executed on table version 1 to generate table version 2. The change 102 includes deleting rows in MP2 to generate a new micro-partition 4 (MP4) and deleting the original MP2. The change 102 executed on table version 1 generates table version 2 which includes the original MP1 and MP3 along with the newly generated MP4. As a result of the change 102, MP2 has been removed from the table as reflected in table version 2. A change 104 is executed

on table version 2 to generate table version 3. The change 104 includes inserting new rows such that micro-partition 5 (MP5) is generated and MP3 is removed from the table. Table version 3 includes the original MP1, the MP4 generated as a result of the change 102, and MP5 generated as a result of the change 104. The MP2 was removed as a result of the change 102 and the MP3 was removed from the table as a result of the change 104.

[0031] As illustrated in FIG. 1, a database table may store database data in one or more micro-partitions, wherein the micro-partitions constitute immutable storage devices. When a change or modification is executed on the table, the affected micro-partitions are removed, and new micro-partitions are created that reflect the change. In an embodiment, the original unmodified micro-partition is not removed but is also stored with the new micro-partition. The change may include any command that impacts one or more rows in the table, including for example, a delete command, an insert command, an update command, and/or a merge command.

[0032] In an embodiment, a delta for the table may be determined that indicates all changes that have been made between, for example, table version 1 and table version 3. A delta for the embodiment illustrated in FIG. 1 may indicate that rows were deleted in micro-partition 2 and rows were inserted in micro-partition 3. The delta may further indicate that micro-partition 4 was generated based on micro-partition 2 in response to a delete transaction, and that micro-partition 5 was generated based on micro-partition 3 in response to an insert transaction.

[0033] The delta may be determined based on a table history that includes a log of changes made to the table. Each change in the log of changes may include a timestamp indicating when the change was made and further indicating what transaction initiated the change. The log of changes may include historical data indicating what rows were changed, where rows were originally stored (*e.g.*, in what micro-partition a row was originally stored), what transactions caused changes to certain rows, prior row identifiers for a row, prior values for a row, an

indication of whether a row was updated, and so forth. In an implementation, log of changes is stored as one or more change tracking columns stored within the table itself.

[0034] In an embodiment, the table history is updated only if a transaction is fully completed. Therefore, if a transaction is initiated but is not fully completed, the table history will not be updated to include that transaction. For example, if a transaction is initiated to delete certain rows of an original micro-partition, a new micro-partition will be generated that includes all original rows in the original micro-partition except for those that should be deleted based on a delete command. If the transaction is not completed, *i.e.* if the new micro-partition is not fully generated, then the table history will not be updated to indicate that the transaction occurred.

[0035] In an embodiment, the table history is stored in a metadata column that includes a micro-partition lineage storing a lineage of micro-partition names over a time period. The metadata column may further include a row lineage storing a lineage of ordinal numbers for a given row over a time period. Such values are NULL when a table is first created or when a new row is generated within the table. When a new row is generated, the stored values will be NULL, and the actual logical value may be derived from the name of the micro-partition that the row is stored within and the rank of the row within that micro-partition.

[0036] In an embodiment, when a row is copied into a new micro-partition as part of a DML query execution, and if the value for the row is NULL, the lineage of the row is materialized by replacing the value with the applicable micro-partition name and the applicable rank of the row within that micro-partition. When a row is copied into a new micro-partition and the value of the row is not NULL, the name of the row is copied over into the new micro-partition.

[0037] FIGS. 2-4 illustrate exemplary embodiments of delete, insert, and update commands that may be executed on a database table. It should be appreciated that the table

schemas illustrated in FIGS. 2-4 are illustrative and include simple values to represent rows and columns that may be included in a database table.

[0038] FIG. 2 illustrates a block diagram of an example delete command 200 and a resulting delta 210 that may be returned after the delete command 200 is complete. Micro-partition 2 (MP2) as illustrated in FIG. 2 includes four columns. Column 1 includes entries for numerical ids that are primarily used for identification purposes for the user data. Column 2 includes entries for row values that may include any value depending on the subject or purpose of the database table. The metadata name column includes table history information about which micro-partition the data originated from or was last located within. The metadata row column includes table history information about which row the data originated from or was located within.

[0039] As illustrated in FIG. 2, the delete command 200 is performed on MP2 and deletes rows 2 and 3 from MP2 to generate the new MP4. As an example, as illustrated in FIG. 2, MP2 includes four rows – namely rows 1, 2, 3, and 4. It should be appreciated that a micro-partition may include any number of rows and may often include thousands of rows in a single micro-partition. The values for each of the rows in MP2 are listed as value1, value2, value3, and value4 for the four rows by way of example but it should be appreciated the value may include any suitable value as pertinent to the database. In the original and unmodified MP2, the metadata name for each of the four rows is “NULL (MP2)” indicating the data is original to that micro-partition and does not yet have any change history. Similarly, the metadata row column for MP2 is NULL and indicates the original row number because the data is original to MP2 and does not yet have a change tracking history.

[0040] MP4 is generated based on the delete command 200 performed on MP2 that deleted rows 2 and 3 as illustrated in FIG. 2. MP4 now only includes rows 1 and 4 having values value1 and value4, respectively. The metadata name for each of rows 1 and 4 is “MP2” indicating the

row data originated or was last located within MP2. The metadata row for each of rows 1 and 4 is 1 and 4, respectively, indicating where the rows were last located.

[0041] A delta 210 may be determined after the delete command 200 is performed on a table. In an embodiment, a timestamp is attached to each transaction that occurs on the table. If the transaction is fully completed, then the timestamp is further attached to the change tracking history for that transaction. Attaching the timestamp to the change tracking history enables the system to know when a table was changed by a certain transaction and when a certain change occurred on any of a plurality of rows in the table.

[0042] The delta 210 illustrated in FIG. 2 determines a difference or a change that occurred between MP2 and MP4. In various embodiments the delta 210 may determine a total change that has occurred between any two timestamps, even if many transactions have occurred on the data between those two timestamps and the data has been changed multiple times. The delta 210 provides an indication of a total change between two timestamps without providing information on any incremental changes that occurred between a first timestamp and a second timestamp.

[0043] The delta 210 includes four columns, namely a column 1 and column 2 (similar to those shown in MP2 and MP4) along with a metadata action column and a metadata is update column. Column 1 indicates the numerical ids are 2 and 3, respectively. Column 2 indicates that the values of rows 2 and 3 are value2 and value3, respectively. In various embodiments, where the values of rows 2 and 3 may have changed one or more times between the first timestamp and the second timestamp, Column1 and Column 2 may indicate the original value or the most recent value. The metadata action column indicates that rows 2 and 3 underwent a DELETE command. The metadata is update column indicates whether the metadata was updated. In the delta 210 illustrated in FIG. 2, the metadata is update column returns a FALSE because the rows did not undergo an update (but were instead deleted).

[0044] In an embodiment, a table history is generating that includes all transactions that have been initiated against the table. Such transactions may include, for example, data manipulation language (DML) commands such as delete, insert, merge, or update commands initiated on a micro-partition. The table history may be dynamically updated to reflect all deleted or inserted rows on table for an interval of transaction time. The table history may include a list of DML statements sorted by transaction time, where each transaction includes a timestamp. In an embodiment, it is assumed that all DML statements will delete, insert, and/or update rows at the same time.

[0045] In an embodiment, a table history may be determined by retrieving a list of added and removed micro-partitions between two transactions. For each given micro-partition, a lineage sequence of dependencies may be generated that indicates which rows have been updated, how those rows were updated, and what transaction caused each update. A delta may be determined by requesting a changeset between two timestamps in the lineage. The delta may return a listing of all added micro-partitions and all removed micro-partitions between the two timestamps. The listing may be consolidated by removing those micro-partitions that appear on both the list of added micro-partitions and on the list of removed micro-partitions.

[0046] In an embodiment, the row granularity changes between added and removed micro-partitions is determined. In such an embodiment, a full outer join is performed between a rowset of added micro-partitions and a rowset of removed micro-partitions on a metadata partition and metadata row number columns. When the resulting joined rowset includes NULL values for the metadata columns in the data from the added rowset, this indicates those rows represent DELETES. If the values for the metadata columns in the data from the inserted rowset are NULL, this indicates those rows represent INSERTs. If the metadata columns originating from the added and removed rowsets are both NULL, this indicates the rows were potentially

updated, and comparing the original data columns will result in the information indicating whether the rows were actually modified.

[0047] Further in an embodiment, a row granularity list of changes may be determined between any two transaction times for a given table. A side-by-side representation may be generated that may be easily used in a merge statement by checking which part of data is present. A table valued function may be utilized to query the table history for a given table (or a materialized view). Further, a SQL statement may be utilized by referring the INSERTED or DELETED columns in a changeset to return an indication of which rows in the table have been inserted or deleted.

[0048] FIG. 3 illustrates a block diagram of an example insert command 300 and a resulting delta 310 that may be returned after the insert command 300 is complete. FIG. 3 begins with an exemplary micro-partition 3 (MP3) that undergoes an insert command 300 to generate micro-partition 5 (MP5). The insert command 300 inserts rows 17 and 18 into MP3. As an example,, MP3 includes three rows, namely rows 21, 22, and 23 having values of value21, value22, and value23, respectively. The metadata name is NULL (MP3) for each of the three rows because there is not yet a change tracking history for the rows that indicates where the rows originated or were last stored. The NULL() notation indicates that the values for the metadata columns are NULL when rows are first inserted into the table. The NULL() notation can reduce overhead. When values for a row are copied into a new micro-partition, the rank of the row is notated in the NULL() notation, such that the first row is NULL(1), the second row is NULL(2), the third row is NULL(3), and so forth.

[0049] MP5 is generated based on the insert command 300 and now includes rows 17 and 18 - we are referring to these rows by using the content of the first column. The values for second column are value17 and value18, respectively, because rows 17 and 18 were inserted into MP5 and those are the assigned values for the rows. The values for rows 21, 22, and 23

have not changed. The metadata name information for rows 21, 22, and 23 is “MP3” because the data originated from or was last stored in micro-partition 3. The metadata row information for rows 21, 22, and 23 is 1, 2, and 3, respectively, because rows 21, 22, and 23 were originally or last stored in rows 1, 2, and 3 in micro-partition 3. The metadata name information and the metadata row information for rows 17 and 18 is “NULL” because the rows originated in MP5 and do not yet have any change tracking history information.

[0050] The delta 310 for the insert command 300 illustrates the total change made between a first timestamp and a second timestamp. As illustrated in FIG. 2, the delta 310 illustrates the change that occurred between MP5 and MP3. It should be appreciated that in alternative embodiments or implementations, a delta may indicate a total change or modification that occurred on a table between any two timestamps without indicating incremental changes that occurred on the table.

[0051] The delta 310 includes rows 17 and 18 having value17 and value18, respectively because rows 17 and 18 were added to MP3 because of the insert command 300. The metadata action is “INSERT” for rows 17 and 18 because an insert command 300 was the transaction that caused a modification to the rows. The metadata is update information is “FALSE” for rows 17 and 18 because the rows were not updated but were instead inserted.

[0052] FIG. 4 illustrates a block diagram of an example update command 400 and a resulting delta 410 that may be returned after the update command 400 is complete. In the example embodiment illustrated in FIG. 4, micro-partition 78 (MP78) is updated to generate micro-partition 91 (MP91). The update command 400 updates rows 1 and 4 to new values. MP78 includes rows 1, 2, 3, and 4 having values of value1, value2, value3, and value4, respectively. The metadata name information is “NULL (MP78) for each of the rows because there is not yet change tracking history for the rows indicating where the rows were last stored.

The metadata row information for each of the rows is NULL because there is not yet change tracking history for the rows indicating which row the values were last stored.

[0053] MP91 includes rows 1, 2, 3, and 4. However, due to the update command 400, row 1 now has a value of VALUE11 and row 4 now has a value of VALUE44. The metadata name information for each of rows 1, 2, 3, and 4 is “MP78” because the values originated from or were last stored in MP78. The metadata row information for row 1 is “1” because that value was last stored in row 1 in MP78. Similarly, for rows 2, 3, and 4, the metadata row information is “2”, “3”, and “4”, respectively. When rows 1, 2, 3, and 4 were copied over to the new MP91 the ordinal position (also referred to as “rank”) of the rows was also copied over from the original MP78.

[0054] The delta 410 indicates the change between a first timestamp and a second timestamp. As illustrated in FIG. 4, the delta 410 indicates a total change between MP78 and MP91 due to the update command 400. The delta 410 indicates that rows 1 and 4 that had a value of “value1” and “value2”, respectively, were deleted. The delta 410 indicates that rows 1 and 4 that have a value of “VALUE11” and “VALUE44”, respectively, were inserted. The metadata is update information is “TRUE” for all rows because an update command 400 was performed on the rows. As indicated in the delta 410, when an update command is performed, the original row is deleted, and a new row is inserted to carry out the command.

[0055] Turning to FIG. 5, a block diagram is shown illustrating a processing platform 500 for providing database services, according to one embodiment. The processing platform 500 includes a database service manager 502 that is accessible by multiple users 504, 506, and 508. The database service manager 502 may also be referred to herein as a resource manager or global services. In some implementations, database service manager 502 can support any number of users desiring access to data or services of the processing platform 500. Users 504-508 may include, for example, end users providing data storage and retrieval queries and

requests, system administrators managing the systems and methods described herein, software applications that interact with a database, and other components/devices that interact with database service manager 502. In a particular embodiment as illustrated herein, the users 504-508 may initiate changes to database data and may request a delta for a database table.

[0056] The database service manager 502 may provide various services and functions that support the operation of the systems and components within the processing platform 500. Database service manager 502 has access to stored metadata associated with the data stored throughout data processing platform 500. The database service manager 502 may use the metadata for optimizing user queries. In some embodiments, metadata includes a summary of data stored in remote data storage systems as well as data available from a local cache (e.g., a cache within one or more of the clusters of the execution platform 512). Additionally, metadata may include information regarding how data is organized in the remote data storage systems and the local caches. Metadata allows systems and services to determine whether a piece of data needs to be processed without loading or accessing the actual data from a storage device.

[0057] As part of the data processing platform 500, metadata may be collected when changes are made to the data using a data manipulation language (DML), which changes may be made by way of any DML statement. Examples of manipulating data may include, but are not limited to, selecting, updating, changing, merging, and inserting data into tables. As part of the processing platform 500, micro-partitions may be created, and the metadata may be collected on a per file and a per column basis. This collection of metadata may be performed during data ingestion or the collection of metadata may be performed as a separate process after the data is ingested or loaded. In an implementation, the metadata may include a number of distinct values; a number of null values; and a minimum value and a maximum value for each file. In an implementation, the metadata may further include string length information and ranges of characters in strings.

[0058] In one embodiment, at least a portion of the metadata is stored in immutable storage such as a micro-partition. For example, the metadata may be stored on the storage platform 514 along with table data. In one embodiment, the same or separate cloud storage resources that are used for table data may be allocated and used for the metadata. In one embodiment, the metadata may be stored in local immutable storage. In one embodiment, information about the metadata in immutable storage, or information about metadata files stored in immutable storage, is stored in mutable storage 510. The information about metadata may be referenced for locating and accessing the metadata stored in immutable storage. In one embodiment, systems with metadata storage may be restructured such that the metadata storage is used instead to store information about metadata files located in immutable storage.

[0059] Database service manager 502 is further in communication with an execution platform 512, which provides computing resources that execute various data storage and data retrieval operations. The execution platform 512 may include one or more compute clusters. The execution platform 512 is in communication with one or more data storage devices 516, 518, and 520 that are part of a storage platform 514. Although three data storage devices 516, 518, and 520 are shown in FIG. 5, the execution platform 512 is capable of communicating with any number of data storage devices. In some embodiments, data storage devices 516, 518, and 520 are cloud-based storage devices located in one or more geographic locations. For example, data storage devices 516, 518, and 520 may be part of a public cloud infrastructure or a private cloud infrastructure, or any other manner of distributed storage system. Data storage devices 516, 518, and 520 may include hard disk drives (HDDs), solid state drives (SSDs), storage clusters, or any other data storage technology. Additionally, the storage platform 514 may include a distributed file system (such as Hadoop Distributed File Systems (HDFS), object storage systems, and the like.

[0060] In some embodiments, the communication links between database service manager 502 and users 504-508, mutable storage 510 for information about metadata files (i.e., metadata file metadata), and execution platform 512 are implemented via one or more data communication networks and may be assigned various tasks such that user requests can be optimized. Similarly, the communication links between execution platform 512 and data storage devices 516-520 in storage platform 514 are implemented via one or more data communication networks. These data communication networks may utilize any communication protocol and any type of communication medium. In some embodiments, the data communication networks are a combination of two or more data communication networks (or sub-networks) coupled to one another. In alternative embodiments, these communication links are implemented using any type of communication medium and any communication protocol.

[0061] The database service manager 502, mutable storage 510, execution platform 512, and storage platform 514 are shown in FIG. 5 as individual components. However, each of database service manager 502, mutable storage 510, execution platform 512, and storage platform 514 may be implemented as a distributed system (e.g., distributed across multiple systems/platforms at multiple geographic locations) or may be combined into one or more systems. Additionally, each of the database service manager 502, mutable storage 510, the execution platform 512, and the storage platform 514 may be scaled up or down (independently of one another) depending on changes to the requests received from users 504-508 and the changing needs of the data processing platform 500. Thus, in the described embodiments, the data processing platform 500 is dynamic and supports regular changes to meet the current data processing needs.

[0062] FIG. 6 illustrates a block diagram depicting components of database service manager 502, according to one embodiment. The database service manager 502 includes an

access manager 602 and a key manager 604 coupled to a data storage device 606. The access manager 602 handles authentication and authorization tasks for the systems described herein. The key manager 604 manages storage and authentication of keys used during authentication and authorization tasks. A request processing service 608 manages received data storage requests and data retrieval requests. A management console service 610 supports access to various systems and processes by administrators and other system managers.

[0063] The database service manager 502 also includes an SQL compiler 612, an SQL optimizer 614 and an SQL executor 616. SQL compiler 612 parses SQL queries and generates the execution code for the queries. SQL optimizer 614 determines the best method to execute queries based on the data that needs to be processed. SQL executor 616 executes the query code for queries received by database service manager 502. A query scheduler and coordinator 618 sends received queries to the appropriate services or systems for compilation, optimization, and dispatch to an execution platform 612. A virtual warehouse manager 620 manages the operation of multiple virtual warehouses.

[0064] Additionally, the database service manager 502 includes a change tracking manager 628, which manages the information related to the data stored in the remote data storage devices and in the local caches. A monitor and workload analyzer 624 oversees the processes performed by the database service manager 502 and manages the distribution of tasks (e.g., workload) across the virtual warehouses and execution nodes in the execution platform 512. Change tracking manager 628 and monitor and workload analyzer 624 are coupled to a data storage device 626. In one embodiment, the configuration and metadata manager 622 collects, stores, and manages metadata in an immutable storage resource. In one embodiment, updates to metadata result in new files and are not updated in place.

[0065] Metadata files, as discussed herein, may include files that contain metadata of modifications (e.g., each modification) to any database table in a data warehouse. A

modification of a database table may generate one or more metadata files, often just a single metadata file. In one embodiment, metadata files contain the following information: information about a metadata file, including a version number; a list of all added table data files; a list of deleted table data files; and information about each added table data file, including file path, file size, file key id, as well as summaries of all rows and columns that are stored in the table data file.

[0066] In one embodiment, the contents of metadata files may vary over time. If format or content of a metadata file changes, the version number of the metadata file may be incremented. In one embodiment, the metadata store (or other mutable data storage resource) only stores information about metadata files (which are stored in immutable storage), not about table data files. In practice, information about metadata files stored in the metadata store (or other mutable storage) is very limited and may contain data for thousands of metadata files. In one embodiment, information for up to 30,000 metadata files may be stored in metadata store or other mutable storage. This dramatically reduces the amount of storage needed in the metadata store or other mutable storage.

[0067] In one embodiment, a system writes metadata files to cloud storage for every modification of a database table (e.g., modification of table data files). In addition to adding and deleting files, every modification to a database table in the data warehouse also generates one or more metadata files. Typically, a modification creates a single metadata file. However, if the modification to the table is large (e.g., an insert into a table that produces very many files), it may result in the creation of multiple metadata files. Further operation of the change tracking manager 628 will be discussed further in relation to FIGS. 6-12.

[0068] The database service manager 502 also includes a change tracking manager 628, which manages the generation of change tracking history such as one or more change tracking columns stored in a table. The change tracking manager 628 further manages the generation of

a delta report indicating a total change that has occurred on a database table between a first timestamp and a second timestamp. Because multiple users/systems may access the same data simultaneously, changes to the data may be synchronized to ensure that each user/system is working with the current version of the data and has access to a change tracking history for the data.

[0069] Security improvements are also implemented in some embodiments. In one embodiment, metadata files and change tracking information is encrypted using individual file keys. Within a micro-partition, columns may be encrypted individually using AES-CTR mode with different start counters. This allows a database system to read an individual column from a micro-partition because it can be decrypted without needing to decrypt the whole micro-partition at once. Encryption improves security because nobody can read the micro-partition without having the proper file key.

[0070] For verification that a micro-partition has not been altered, the system may store hashes of columns for each column within the micro-partition. Before decrypting the data, the system compares the hash of the encrypted column with the stored hash of the column of this micro-partition. If the hashes do not match, the micro-partition must have been altered. This improves security because all altering of micro-partitions is detected by the database system.

[0071] FIG. 7 is a schematic block diagram illustrating components of a change tracking manager 628, according to one embodiment. The change tracking manager 628 may collect, store, and manage metadata about table data files (*i.e.* micro-partitions) as well as metadata about metadata files. Such metadata includes change tracking information or a table history including a log of changes that have occurred on a table. The change tracking manager 628 includes a table data component 702, a metadata component 704, a change tracking history component 706, an encryption component 708, and a delta component 710. The components 702-710 are given by way of illustration only and may not all be included in all embodiments.

In fact, some embodiments may include only one or any combination of two or more of the components 702-710. For example, some of the components may be located outside or separate from the change tracking manager 628, such as within a database service manager 502 or processing platform 500. Furthermore, the components 702-710 may comprise hardware, computer readable instructions, or a combination of both to perform the functionality and provide the structures discussed herein.

[0072] The table data component 702 stores table data for a database, the table data includes information about rows and columns of one or more database tables. The table data component 702 may store table data in micro-partitions within a storage resource. Example storage resources include cloud storage and/or immutable storage. In one embodiment, the storage resources for storage of table data files may be dynamically allocated to accommodate increases or decreases in storage requirement. The table data component 702 may manage and store table data by causing the data to be stored or updated in a remote resource, such as a cloud storage resource or service.

[0073] The metadata component 704 stores metadata on immutable storage such as a micro-partition. The metadata may include information about or describing the table data for the database stored by the table data component 702. In one embodiment, the metadata files may include metadata such as an indication of added or deleted table data files. The metadata may include file information for table data files, the file information including one or more of a file name and a storage location. In one embodiment, the metadata may be stored in files on the same cloud storage resources as the table data. In one embodiment, metadata component 704 may cause the metadata to be stored within metadata files in a column-by-column format in remote cloud storage.

[0074] The metadata component 704 may also collect and manage storage of metadata within metadata files on the immutable storage. The metadata component 704 may create, in

response to a change in the table data, a new metadata file in the immutable storage without modifying previous metadata files. The new metadata file may include metadata indicating the change in the table data. In one embodiment, the metadata in the new metadata file indicates an addition or a deletion of a table data file comprising the table data. The metadata component 704 may also delete expired metadata files. Expired metadata files may include those older than a specific age and that are not referenced in metadata information stored by the change tracking history component 706.

[0075] The change tracking history component 706 stores and manages information about changes made to the table data. The change tracking history may be stored in local mutable storage and/or within the table data as a change tracking column, for example. In one embodiment, the change tracking history may be stored and updated in-place. In one embodiment, the change tracking history is stored within an immutable micro-partition in a change tracking column indicating a most recent change that occurred on a row and/or a log of changes that have occurred on the table. In an embodiment, the change tracking history component 706 secures a timestamp to a transaction that occurred on the table and further secures a timestamp to each change or modification that occurs on one or more rows of the table.

[0076] The encryption component 708 is configured to encrypt table data and metadata. In one embodiment, the encryption component 708 encrypts the metadata column-by-column to allow for independent decryption and reading of metadata for a specific column.

[0077] The delta component 710 generates a delta for a table between a first timestamp and a second timestamp. The delta component 710 determines a total change or modification that has occurred on one or more rows in a table without providing information on incremental changes that occurred between the first timestamp and the second timestamp.

[0078] FIG. 8 is a schematic flow chart diagram illustrating an example method 800 for determining a change made to a database table. The method 800 may be performed by a change tracking manager 628, database service manager 502, processing platform 500, and/or other service or platform.

[0079] The method 800 begins and a table data component 702 determines at 802 to be executed on a micro-partition of a table of a database. The table data component 702 executes at 804 the change on the table by generating a new micro-partition that reflects the change. A change tracking history component 706 updates at 806 a table history that comprises a log of changes made to the table, wherein each change in the log of changes comprises a timestamp, and wherein updating the table history comprises inserting the change into the log of changes.

[0080] FIG. 9 is a schematic flow chart diagram illustrating an example method 900 for determining a change made to a database table. The method 900 may be performed by a change tracking manager 628, database service manager 502, processing platform 500, and/or other service or platform.

[0081] The method 900 begins and a table data component 702 determines at 902 a change to be executed on a micro-partition of a table of a database. The table data component 702 executes at 904 the change on the table by generating a new micro-partition that reflects the change and by removing the micro-partition. A change tracking history component 706 updates at 906 one or more change tracking columns in the new micro-partition, wherein the one or more change tracking columns indicate one or more of: a prior micro-partition associated with a row in the new micro-partition, a prior row identification associated with a row in the new micro-partition, or a prior transaction associated with a row in the new micro-partition. A request processing service 608 receives at 908 a request to determine a delta for the micro-partition and the new micro-partition between a first timestamp and a second timestamp. A

delta component 710 responds to the request by querying at 910 the one or more change tracking columns to determine the delta between the first timestamp and the second timestamp.

[0082] FIG. 10 is a block diagram depicting an example computing device 1000. In some embodiments, computing device 1000 is used to implement one or more of the systems and components discussed herein. For example, computing device 1000 may include or be part of a change tracking manager 628, a database service manager 502, a processing platform 500, and/or any other components or systems discussed herein. As another example, the components, systems, or platforms discussed herein may include one or more computing devices 1000. Further, computing device 1000 may interact with any of the systems and components described herein. Accordingly, computing device 1000 may be used to perform various procedures and tasks, such as those discussed herein. Computing device 1000 can function as a server, a client or any other computing entity. Computing device 1000 can be any of a wide variety of computing devices, such as a desktop computer, a notebook computer, a server computer, a handheld computer, a tablet, and the like.

[0083] Computing device 1000 includes one or more processor(s) 1002, one or more memory device(s) 1004, one or more interface(s) 1006, one or more mass storage device(s) 1008, and one or more Input/Output (I/O) device(s) 1010, all of which are coupled to a bus 1212. Processor(s) 1002 include one or more processors or controllers that execute instructions stored in memory device(s) 1004 and/or mass storage device(s) 1008. Processor(s) 1002 may also include various types of computer-readable media, such as cache memory.

[0084] Memory device(s) 1004 include various computer-readable media, such as volatile memory (e.g., random access memory (RAM)) and/or nonvolatile memory (e.g., read-only memory (ROM)). Memory device(s) 1004 may also include rewritable ROM, such as Flash memory.

[0085] Mass storage device(s) 1008 include various computer readable media, such as magnetic tapes, magnetic disks, optical disks, solid state memory (e.g., Flash memory), and so forth. Various drives may also be included in mass storage device(s) 1008 to enable reading from and/or writing to the various computer readable media. Mass storage device(s) 1008 include removable media and/or non-removable media.

[0086] I/O device(s) 1010 include various devices that allow data and/or other information to be input to or retrieved from computing device 1000. Example I/O device(s) 1010 include cursor control devices, keyboards, keypads, microphones, monitors or other display devices, speakers, printers, network interface cards, modems, lenses, CCDs or other image capture devices, and the like.

[0087] Interface(s) 1006 include various interfaces that allow computing device 1000 to interact with other systems, devices, or computing environments. Example interface(s) 1006 include any number of different network interfaces, such as interfaces to local area networks (LANs), wide area networks (WANs), wireless networks, and the Internet.

[0088] Bus 1212 allows processor(s) 1002, memory device(s) 1004, interface(s) 1006, mass storage device(s) 1008, and I/O device(s) 1010 to communicate with one another, as well as other devices or components coupled to bus 1212. Bus 1212 represents one or more of several types of bus structures, such as a system bus, PCI bus, IEEE 1394 bus, USB bus, and so forth.

Examples

[0089] The following examples pertain to further embodiments.

[0090] Example 1 is a method for tracking changes to database data. The method includes determining a change to be executed on a micro-partition of a table of a database and executing the change on the table by generating a new micro-partition that embodies the change. The method includes updating a table history that comprises a log of changes made to the table,

wherein each change in the log of changes comprises a timestamp, and wherein updating the table history comprises inserting the change into the log of changes. The method includes determining a delta for the table between a first timestamp and a second timestamp based on the table history.

[0091] Example 2 is a method as in Example 1, wherein the change comprises deleting a row in a first micro-partition and wherein executing the change comprises generating a second micro-partition not including the row and deleting the first micro-partition.

[0092] Example 3 is a method as in any of Examples 1-2, wherein the change comprises adding a row to a first micro-partition and wherein executing the change comprises generating a second micro-partition comprising the row and deleting the first micro-partition.

[0093] Example 4 is a method as in any of Examples 1-3, wherein the change comprises updating a row in a first micro-partition and wherein executing the change comprises generating a second micro-partition comprising the update and deleting the first micro-partition.

[0094] Example 5 is a method as in any of Examples 1-4, wherein executing the change on the table further comprises deleting the micro-partition.

[0095] Example 6 is a method as in any of Examples 1-5, wherein the change is applied to a plurality of micro-partitions of the table such that executing the change on the table comprises generating a plurality of new micro-partitions and deleting the plurality of micro-partitions.

[0096] Example 7 is a method as in any of Examples 1-6, wherein the change comprises one or more of a delete, an insert, a merge, or an update to one or more rows of the table.

[0097] Example 8 is a method as in any of Examples 1-7, wherein the delta for the table comprises an indication of rows and micro-partitions of the table that have changed between the first timestamp and the second timestamp, wherein the first timestamp and the second timestamp comprise any times in the table history.

[0098] Example 9 is a method as in any of Examples 1-8, further comprising generating a row identification for each row of a plurality of rows of the table, wherein the row identification is stored in the table history in association with an applicable timestamp in the log of changes made to the table.

[0099] Example 10 is a method as in any of Examples 1-9, wherein generating the row identification comprises: for an unmodified row, deriving a row identification based on how data in the unmodified row is stored; or for a modified row that is modified by executing the change on the table, generating a row identification for the modified row in the new micro-partition that embodies the change.

[0100] Example 11 is a method as in any of Examples 1-10, wherein determining the delta for the table comprises comparing a first set of data associated with the first timestamp and a second set of data associated with the second timestamp, wherein the comparing comprises deriving tuple granularity changes between the first set of data and the second set of data.

[0101] Example 12 is a method as in any of Examples 1-11, wherein the table history comprises a change tracking column stored in the table and wherein the micro-partition is an immutable storage device.

[0102] Example 13 is a method as in any of Examples 1-12, wherein generating the new micro-partition that embodies the change comprises: copying unmodified rows from the micro-partition along with change history information stored in association with the unmodified rows; generating new modified rows based on the change; and generating a row identification for each of the new modified rows.

[0103] Example 14 is a method as in any of Examples 1-13, wherein the first timestamp is associated with a first set of micro-partitions and the second timestamp is associated with a second set of micro-partitions, and wherein determining the delta for the table between the first timestamp and the second timestamp comprises: comparing a first set of rules for the first set

of micro-partitions and a second set of rules for the second set of micro-partitions; deriving applicable change tracking information from the table history for each of the first timestamp and the second timestamp; and determining one or more rows that have been added, deleted, or updated between the first set of micro-partitions and the second set of micro-partitions.

[0104] Example 15 is a method as in any of Examples 1-14, wherein the log of changes in the table history comprises: a listing of one or more transactions that initiated a modification to the table, wherein the modification comprises one or more of: an added row, a deleted row, or an updated row; and a timestamp for each transaction in the listing of one or more transactions.

[0105] Example 16 is a method as in any of Examples 1-15, wherein the log of changes in the table history comprises a lineage for each of a plurality of rows in the table, wherein the lineage comprises a log of modifications with a timestamp associated with each modification, and wherein the lineage enables tracking each of the plurality of rows across a plurality of micro-partitions in the table.

[0106] Example 17 is a system for tracking changes to database data. The system includes means for determining a change to be executed on a micro-partition of a table of a database and means for executing the change on the table by generating a new micro-partition that embodies the change. The system includes means for updating a table history that comprises a log of changes made to the table, wherein each change in the log of changes comprises a timestamp, and wherein updating the table history comprises inserting the change into the log of changes. The system includes means for determining a delta for the table between a first timestamp and a second timestamp based on the table history.

[0107] Example 18 is a system as in Example 17, wherein the change comprises deleting a row in a first micro-partition and wherein the means for executing the change is configured to generate a second micro-partition not including the row and deleting the first micro-partition.

[0108] Example 19 is a system as in any of Examples 17-18, wherein the change comprises adding a row to a first micro-partition and wherein the means for executing the change is configured to generate a second micro-partition comprising the row and deleting the first micro-partition.

[0109] Example 20 is a system as in any of Examples 17-19, wherein the change comprises updating a row in a first micro-partition and wherein the means for executing the change is configured to generate a second micro-partition comprising the update and deleting the first micro-partition.

[0110] Example 21 is a system as in any of Examples 17-20, wherein the means for executing the change on the table is further configured to delete the micro-partition.

[0111] Example 22 is a system as in any of Examples 17-21, wherein the change is applied to a plurality of micro-partitions of the table such that executing the change on the table comprises generating a plurality of new micro-partitions and deleting the plurality of micro-partitions.

[0112] Example 23 is a system as in any of Examples 17-22, wherein the change comprises one or more of a delete, an insert, a merge, or an update to one or more rows of the table.

[0113] Example 24 is a system as in any of Examples 17-23, wherein the delta for the table comprises an indication of rows and micro-partitions of the table that have changed between the first timestamp and the second timestamp, wherein the first timestamp and the second timestamp comprise any times in the table history.

[0114] Example 25 is a system as in any of Examples 17-24, further comprising means for generating a row identification for each row of a plurality of rows of the table, wherein the row identification is stored in the table history in association with an applicable timestamp in the log of changes made to the table.

[0115] Example 26 is a system as in any of Examples 17-25, wherein the means for generating the row identification is configured to: for an unmodified row, derive a row identification based on how data in the unmodified row is stored; or for a modified row that is modified by executing the change on the table, generate a row identification for the modified row in the new micro-partition that embodies the change.

[0116] Example 27 is a system as in any of Examples 17-26, the means for determining the delta for the table is configured to compare a first set of data associated with the first timestamp and a second set of data associated with the second timestamp, wherein the comparing comprises deriving tuple granularity changes between the first set of data and the second set of data.

[0117] Example 28 is a system as in any of Examples 17-27, wherein the table history comprises a change tracking column stored in the table and wherein the micro-partition is an immutable storage device.

[0118] Example 29 is a system as in any of Examples 17-28, wherein the means for generating the new micro-partition that embodies the change is configured to: copy unmodified rows from the micro-partition along with change history information stored in association with the unmodified rows; generate new modified rows based on the change; and generate a row identification for each of the new modified rows.

[0119] Example 30 is a system as in any of Examples 17-29, wherein the first timestamp is associated with a first set of micro-partitions and the second timestamp is associated with a second set of micro-partitions, and wherein determining the delta for the table between the first timestamp and the second timestamp comprises: comparing a first set of rules for the first set of micro-partitions and a second set of rules for the second set of micro-partitions; deriving applicable change tracking information from the table history for each of the first timestamp

and the second timestamp; and determining one or more rows that have been added, deleted, or updated between the first set of micro-partitions and the second set of micro-partitions.

[0120] Example 31 is a system as in any of Examples 17-30, wherein the log of changes in the table history comprises: a listing of one or more transactions that initiated a modification to the table, wherein the modification comprises one or more of: an added row, a deleted row, or an updated row; and a timestamp for each transaction in the listing of one or more transactions.

[0121] Example 32 is a system as in any of Examples 17-31, wherein the log of changes in the table history comprises a lineage for each of a plurality of rows in the table, wherein the lineage comprises a log of modifications with a timestamp associated with each modification, and wherein the lineage enables tracking each of the plurality of rows across a plurality of micro-partitions in the table.

[0122] Example 33 is non-transitory computer readable storage media storing instructions that, when executed by one or more processors, cause the one or more processors to: determine a change to be executed on a micro-partition of a table of a database; execute the change on the table by generating a new micro-partition that embodies the change; update a table history that comprises a log of changes made to the table, wherein each change in the log of changes comprises a timestamp, and wherein updating the table history comprises inserting the change into the log of changes; and determine a delta for the table between a first timestamp and a second timestamp based on the table history.

[0123] Example 34 is non-transitory computer readable storage media as in claim 33, wherein the change comprises deleting a row in a first micro-partition and wherein the instructions cause the one or more processors to execute the change by generating a second micro-partition not including the row and deleting the first micro-partition.

[0124] Example 35 is non-transitory computer readable storage media as in any of Examples 33-34, wherein the change comprises adding a row to a first micro-partition and wherein the instructions cause the one or more processors to execute the change by generating a second micro-partition comprising the row and deleting the first micro-partition.

[0125] Example 36 is non-transitory computer readable storage media as in any of Examples 33-35, wherein the change comprises updating a row in a first micro-partition and wherein the instructions cause the one or more processors to execute the change by generating a second micro-partition comprising the update and deleting the first micro-partition.

[0126] Example 37 is non-transitory computer readable storage media as in any of Examples 33-36, wherein the instructions further cause the one or more processors to execute the change on the table by deleting the micro-partition.

[0127] Example 38 is non-transitory computer readable storage media as in any of Examples 33-37, wherein the change is applied to a plurality of micro-partitions of the table such that the instructions cause the one or more processors to execute the change on the table by generating a plurality of new micro-partitions and deleting the plurality of micro-partitions.

[0128] Example 39 is non-transitory computer readable storage media as in any of Examples 33-38, wherein the change comprises one or more of a delete, an insert, a merge, or an update to one or more rows of the table.

[0129] Example 40 is non-transitory computer readable storage media as in any of Examples 33-39, wherein the delta for the table comprises an indication of rows and micro-partitions of the table that have changed between the first timestamp and the second timestamp, wherein the first timestamp and the second timestamp comprise any times in the table history.

[0130] Example 41 is non-transitory computer readable storage media as in any of Examples 33-40, wherein the instructions further cause the one or more processors to generate a row identification for each row of a plurality of rows of the table, wherein the row

identification is stored in the table history in association with an applicable timestamp in the log of changes made to the table.

[0131] Example 42 is non-transitory computer readable storage media as in any of Examples 33-41, wherein the instructions cause the one or more processors to generate the row identification by: for an unmodified row, deriving a row identification based on how data in the unmodified row is stored; or for a modified row that is modified by executing the change on the table, generating a row identification for the modified row in the new micro-partition that embodies the change.

[0132] Example 43 is non-transitory computer readable storage media as in any of Examples 33-42, wherein the instructions cause the one or more processors to determine the delta for the table by comparing a first set of data associated with the first timestamp and a second set of data associated with the second timestamp, wherein the comparing comprises deriving tuple granularity changes between the first set of data and the second set of data.

[0133] Example 44 is non-transitory computer readable storage media as in any of Examples 33-43, wherein the table history comprises a change tracking column stored in the table and wherein the micro-partition is an immutable storage device.

[0134] Example 45 is non-transitory computer readable storage media as in any of Examples 33-44, wherein the instructions cause the one or more processors to generate the new micro-partition that embodies the change by: copying unmodified rows from the micro-partition along with change history information stored in association with the unmodified rows; generating new modified rows based on the change; and generating a row identification for each of the new modified rows.

[0135] Example 46 is non-transitory computer readable storage media as in any of Examples 33-45, wherein the first timestamp is associated with a first set of micro-partitions and the second timestamp is associated with a second set of micro-partitions, and wherein the

instructions cause the one or more processors to determine the delta for the table between the first timestamp and the second timestamp by: comparing a first set of rules for the first set of micro-partitions and a second set of rules for the second set of micro-partitions; deriving applicable change tracking information from the table history for each of the first timestamp and the second timestamp; and determining one or more rows that have been added, deleted, or updated between the first set of micro-partitions and the second set of micro-partitions.

[0136] Example 47 is non-transitory computer readable storage media as in any of Examples 33-46, the log of changes in the table history comprises: a listing of one or more transactions that initiated a modification to the table, wherein the modification comprises one or more of: an added row, a deleted row, or an updated row; and a timestamp for each transaction in the listing of one or more transactions.

[0137] Example 48 is non-transitory computer readable storage media as in any of Examples 33-37, wherein the log of changes in the table history comprises a lineage for each of a plurality of rows in the table, wherein the lineage comprises a log of modifications with a timestamp associated with each modification, and wherein the lineage enables tracking each of the plurality of rows across a plurality of micro-partitions in the table.

[0138] Example 49 is an apparatus including means to perform a method as in any of Examples 1-17.

[0139] Example 50 is a machine-readable storage including machine-readable instructions that, when executed, implement a method or realize an apparatus of any of Examples 1-17.

[0140] The flow diagrams and block diagrams herein illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present disclosure. In this regard, each block in the flow diagrams or block diagrams may represent a module, segment, or portion of code, which comprises one or more executable instructions for implementing the specified

logical function(s). It will also be noted that each block of the block diagrams and/or flow diagrams, and combinations of blocks in the block diagrams and/or flow diagrams, may be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions. These computer program instructions may also be stored in a computer-readable medium that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable medium produce an article of manufacture including instruction means which implement the function/act specified in the flow diagram and/or block diagram block or blocks.

[0141] The systems and methods described herein provide a flexible and scalable data warehouse using new data processing platforms, methods, systems, and algorithms. In some embodiments, the described systems and methods leverage a cloud infrastructure that supports cloud-based storage resources, computing resources, and the like. Example cloud-based storage resources offer significant storage capacity available on-demand at a low cost. Further, these cloud-based storage resources may be fault-tolerant and highly scalable, which can be costly to achieve in private data storage systems. Example cloud-based computing resources are available on-demand and may be priced based on actual usage levels of the resources. Typically, the cloud infrastructure is dynamically deployed, reconfigured, and decommissioned in a rapid manner.

[0142] In the described systems and methods, a data storage system utilizes an SQL (Structured Query Language)-based relational database. However, these systems and methods are applicable to any type of database using any data storage architecture and using any language to store and retrieve data within the database. The systems and methods described herein may also provide a multi-tenant system that supports isolation of computing resources

and data between different customers/clients and between different users within the same customer/client.

[0143] Various techniques, or certain aspects or portions thereof, may take the form of program code (i.e., instructions) embodied in tangible media, such as floppy diskettes, CD-ROMs, hard drives, a non-transitory computer readable storage medium, or any other machine readable storage medium wherein, when the program code is loaded into and executed by a machine, such as a computer, the machine becomes an apparatus for practicing the various techniques. In the case of program code execution on programmable computers, the computing device may include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device. The volatile and non-volatile memory and/or storage elements may be a RAM, an EPROM, a flash drive, an optical drive, a magnetic hard drive, or another medium for storing electronic data. One or more programs that may implement or utilize the various techniques described herein may use an application programming interface (API), reusable controls, and the like. Such programs may be implemented in a high-level procedural or an object-oriented programming language to communicate with a computer system. However, the program(s) may be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language, and combined with hardware implementations.

[0144] It should be understood that many of the functional units described in this specification may be implemented as one or more components, which is a term used to more particularly emphasize their implementation independence. For example, a component may be implemented as a hardware circuit comprising custom very large-scale integration (VLSI) circuits or gate arrays, off-the-shelf semiconductors such as logic chips, transistors, or other discrete components. A component may also be implemented in programmable hardware

devices such as field programmable gate arrays, programmable array logic, programmable logic devices, or the like.

[0145] Components may also be implemented in software for execution by various types of processors. An identified component of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions, which may, for instance, be organized as an object, a procedure, or a function. Nevertheless, the executables of an identified component need not be physically located together but may comprise disparate instructions stored in different locations that, when joined logically together, comprise the component and achieve the stated purpose for the component.

[0146] Indeed, a component of executable code may be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices. Similarly, operational data may be identified and illustrated herein within components and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set or may be distributed over different locations including over different storage devices, and may exist, at least partially, merely as electronic signals on a system or network. The components may be passive or active, including agents operable to perform desired functions.

[0147] Reference throughout this specification to “an example” means that a particular feature, structure, or characteristic described in connection with the example is included in at least one embodiment of the present disclosure. Thus, appearances of the phrase “in an example” in various places throughout this specification are not necessarily all referring to the same embodiment.

[0148] As used herein, a plurality of items, structural elements, compositional elements, and/or materials may be presented in a common list for convenience. However, these lists

should be construed as though each member of the list is individually identified as a separate and unique member. Thus, no individual member of such list should be construed as a de facto equivalent of any other member of the same list solely based on its presentation in a common group without indications to the contrary. In addition, various embodiments and examples of the present disclosure may be referred to herein along with alternatives for the various components thereof. It is understood that such embodiments, examples, and alternatives are not to be construed as de facto equivalents of one another but are to be considered as separate and autonomous representations of the present disclosure.

[0149] Although the foregoing has been described in some detail for purposes of clarity, it will be apparent that certain changes and modifications may be made without departing from the principles thereof. It should be noted that there are many alternative ways of implementing both the processes and apparatuses described herein. Accordingly, the present embodiments are to be considered illustrative and not restrictive.

[0150] Those having skill in the art will appreciate that many changes may be made to the details of the above-described embodiments without departing from the underlying principles of the disclosure. The scope of the present disclosure should, therefore, be determined only by the following claims.

CLAIMS

What is claimed is:

1. A method for tracking changes to database data, the method comprising:
determining a change to be executed on a micro-partition of a table of a database;
executing the change by generating a new micro-partition that embodies the change;
and
updating a table history that comprises a log of changes made to the table, each change in the log of changes comprising a timestamp, the updating of the table history comprising inserting the change into the log of changes.
2. The method of claim 1, further comprising receiving a request for a delta for the table between a first timestamp and a second timestamp, and responsively determining the delta by querying the table history.
3. The method of either claim 1 or claim 2, wherein executing the change further comprises removing the micro-partition from the database.
4. The method of either claim 1 or claim 2, wherein the change is applied to a plurality of micro-partitions of the table such that executing the change comprises generating a plurality of new micro-partitions and removing the plurality of micro-partitions from the database.
5. The method of either claim 1 or claim 2, wherein the change comprises one or more of a delete, an insert, a merge, and an update to one or more rows of the table.

6. The method of claim 2, wherein the delta comprises an indication of at least one of: one or more rows and one or more micro-partitions of the table that changed between the first timestamp and the second timestamp.

7. The method of either claim 1 or claim 2, further comprising generating a row identification for each row of a plurality of rows of the table, wherein the row identification is stored in the table history in association with an applicable timestamp in the log of changes made to the table.

8. The method of claim 7, wherein generating the row identification for each row of the plurality of rows of the table comprises:

for any unmodified row of the plurality of rows of the table, deriving the row identification for the unmodified row based on a micro-partition and a row position in which row data in the unmodified row is stored; and

for any modified row of the plurality of rows of the table that is modified by executing the change, generating the row identification for the modified row in the new micro-partition that embodies the change.

9. The method of claim 2, wherein determining the delta comprises comparing a first set of data associated with the first timestamp and a second set of data associated with the second timestamp, wherein the comparing comprises deriving tuple granularity changes between the first set of data and the second set of data.

10. The method of either claim 1 or claim 2, wherein the table history comprises one or more change tracking columns stored in a micro-partition, and wherein the one or more change tracking columns comprise data indicating one or more of: an original micro-

partition identification for a value, an original row identification for a value, a change action that occurred on a row, and an indication of whether a value was updated.

11. The method of either claim 1 or claim 2, wherein generating the new micro-partition that embodies the change comprises:

copying one or more unmodified rows from the micro-partition along with change history information stored in association with the one or more unmodified rows;
generating one or more new modified rows based on the change; and
generating a row identification for each of the one or more new modified rows.

12. The method of claim 2, wherein the first timestamp is associated with a first set of micro-partitions and the second timestamp is associated with a second set of micro-partitions, and wherein determining the delta comprises:

comparing a first set of rules for the first set of micro-partitions and a second set of rules for the second set of micro-partitions;
deriving applicable change tracking information from the table history for each of the first timestamp and the second timestamp; and
determining one or more rows that have been one or more of added, deleted, and updated between the first set of micro-partitions and the second set of micro-partitions.

13. The method of either claim 1 or claim 2, wherein the log of changes in the table history comprises:

a listing of one or more transactions that initiated a modification to the table, wherein each such modification comprises: an added row, a deleted row, or an updated row; and
a timestamp for each transaction in the listing of one or more transactions.

14. The method of either claim 1 or claim 2, wherein the log of changes in the table history comprises a lineage for each of a plurality of rows in the table, wherein the lineage comprises a log of modifications with a timestamp associated with each modification, and wherein the lineage enables tracking each of the plurality of rows across a plurality of micro-partitions in the table.

15. A system for tracking changes to database data, the system comprising:
one or more processors; and
data storage containing instructions executable by the one or more processors to perform operations comprising:
determining a change to be executed on a micro-partition of a table of a database;
executing the change by generating a new micro-partition that embodies the change; and
updating a table history that comprises a log of changes made to the table, each change in the log of changes comprising a timestamp, the updating of the table history comprising inserting the change into the log of changes.

16. The system of claim 15, the operations further comprising receiving a request for a delta for the table between a first timestamp and a second timestamp, and responsively determining the delta by querying the table history.

17. The system of either claim 15 or claim 16, wherein the change comprises one or more of a delete, an insert, a merge, and an update to one or more rows of the table.

18. The system of claim 16, wherein the delta comprises an indication of at least one of: one or more rows and one or more micro-partitions of the table that changed between the first timestamp and the second timestamp.

19. The system of either claim 15 or claim 16, the operations further comprising generating a row identification for each row of a plurality of rows of the table, wherein the row identification is stored in the table history in association with an applicable timestamp in the log of changes made to the table, and wherein generating the row identification for each row of the plurality of rows of the table comprises:

for any unmodified row of the plurality of rows of the table, deriving the row identification for the unmodified row based on where the data in the unmodified row is stored; and

for any modified row of the plurality of rows of the table that is modified by executing the change, generating the row identification for the modified row in the new micro-partition that embodies the change.

20. The system of claim 16, wherein determining the delta comprises comparing a first set of data associated with the first timestamp and a second set of data associated with the second timestamp, wherein the comparing comprises deriving tuple granularity changes between the first set of data and the second set of data.

21. The system of either claim 15 or claim 16, wherein the table history comprises one or more change tracking columns stored in a micro-partition, and wherein the one or more change tracking columns comprise data indicating one or more of: an original micro-partition identification for a value, an original row identification for a value, a change action that occurred on a row, and an indication of whether a value was updated.

22. The system of either claim 15 or claim 16, wherein generating the new micro-partition that embodies the change comprises:

- copying one or more unmodified rows from the micro-partition along with change history information stored in association with the one or more unmodified rows;
- generating one or more new modified rows based on the change; and
- generating a row identification for each of the one or more new modified rows.

23. The system of claim 16, wherein the first timestamp is associated with a first set of micro-partitions and the second timestamp is associated with a second set of micro-partitions, and wherein determining the delta comprises:

- comparing a first set of rules for the first set of micro-partitions and a second set of rules for the second set of micro-partitions;
- deriving applicable change tracking information from the table history for each of the first timestamp and the second timestamp; and
- determining one or more rows that have been one or more of added, deleted, and updated between the first set of micro-partitions and the second set of micro-partitions.

24. The system of either claim 15 or claim 16, wherein the log of changes in the table history comprises:

- a listing of one or more transactions that initiated a modification to the table, wherein each such modification comprises: an added row, a deleted row, or an updated row; and
- a timestamp for each transaction in the listing of one or more transactions.

25. The system of either claim 15 or claim 16, wherein the log of changes in the table history comprises a lineage for each of a plurality of rows in the table, wherein the lineage comprises a log of modifications with a timestamp associated with each modification,

and wherein the lineage enables tracking each of the plurality of rows across a plurality of micro-partitions in the table.

26. Computer readable storage media containing instructions executable by a processor for causing the processor to perform operations comprising:

determining a change to be executed on a micro-partition of a table of a database;

executing the change by generating a new micro-partition that embodies the change;

and

updating a table history that comprises a log of changes made to the table, each change in the log of changes comprising a timestamp, the updating of the table history comprising inserting the change into the log of changes.

27. The computer readable storage media of claim 26, the operations further comprising receiving a request for a delta for the table between a first timestamp and a second timestamp, and responsively determining the delta by querying the table history.

28. The computer readable storage media of either claim 26 or claim 27, wherein the change comprises one or more of a delete, an insert, a merge, and an update to one or more rows of the table.

29. The computer readable storage media of claim 27, wherein the delta comprises an indication of at least one of: one or more rows and one or more micro-partitions of the table that changed between the first timestamp and the second timestamp.

30. The computer readable storage media of either claim 26 or claim 27, the operations further comprising generating a row identification for each row of a plurality of

rows of the table, wherein the row identification is stored in the table history in association with an applicable timestamp in the log of changes made to the table, and wherein generating the row identification for each row of the plurality of rows of the table comprises:

for any unmodified row of the plurality of rows of the table, deriving the row identification for the unmodified row based on how data in the unmodified row is stored; and

for any modified row of the plurality of rows of the table that is modified by executing the change, generating the row identification for the modified row in the new micro-partition that embodies the change.

31. The computer readable storage media of claim 27, wherein determining the delta comprises comparing a first set of data associated with the first timestamp and a second set of data associated with the second timestamp, wherein the comparing comprises deriving tuple granularity changes between the first set of data and the second set of data.

32. The computer readable storage media of either claim 26 or claim 27, wherein the table history comprises one or more change tracking columns stored in a micro-partition, and wherein the one or more change tracking columns comprise data indicating one or more of: an original micro-partition identification for a value, an original row identification for a value, a change action that occurred on a row, and an indication of whether a value was updated.

33. The computer readable storage media of either claim 26 or claim 27, wherein generating the new micro-partition that embodies the change comprises:

copying one or more unmodified rows from the micro-partition along with change history information stored in association with the one or more unmodified rows;

generating one or more new modified rows based on the change; and

generating a row identification for each of the one or more new modified rows.

34. The computer readable storage media of claim 27, wherein the first timestamp is associated with a first set of micro-partitions and the second timestamp is associated with a second set of micro-partitions, and wherein determining the delta comprises:

comparing a first set of rules for the first set of micro-partitions and a second set of rules for the second set of micro-partitions;

deriving applicable change tracking information from the table history for each of the first timestamp and the second timestamp; and

determining one or more rows that have been one or more of added, deleted, and updated between the first set of micro-partitions and the second set of micro-partitions.

35. The computer readable storage media of either claim 26 or claim 27, wherein the log of changes in the table history comprises:

a listing of one or more transactions that initiated a modification to the table, wherein each such modification comprises: an added row, a deleted row, or an updated row; and

a timestamp for each transaction in the listing of one or more transactions.

36. The computer readable storage media of either claim 26 or claim 27, wherein the log of changes in the table history comprises a lineage for each of a plurality of rows in the table, wherein the lineage comprises a log of modifications with a timestamp associated with each modification, and wherein the lineage enables tracking each of the plurality of rows across a plurality of micro-partitions in the table.




FIG. 1

2/9




FIG. 2

3/9




FIG. 3

4/9




FIG. 4

5/9




FIG. 5




FIG. 6

7/9




FIG. 7




FIG. 8

8/9




FIG. 9




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US19/59798

A. CLASSIFICATION OF SUBJECT MATTER

IPC - G06F 7/00, 17/00; G06Q 10/06 (2020.01)

CPC - G06F 7/00, 16/22, 16/214, 16/219, 16/24535, 16/24539, 16/24544, 16/24568, 16/2379, 16/254, 16/2358; G06Q 10/06

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

See Search History document

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

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2015/0347401 A1 (ORACLE INTERNATIONAL CORPORATION) 03 December 2015; paragraphs [0039], [0044], [0070], [0087]	1-8, 10-13, 15-19, 21-24, 26-30, 32-35 ----- 9, 14, 20, 25, 31, 36
Y	US 2014/0095471 A1 (ORACLE INTERNATIONAL CORPORATION) 03 April 2014; paragraphs [0056], [0301]	9, 20, 31
Y	US 2006/0047622 A1 (FOLKERT, N et al.) 02 March 2006; paragraphs [0026]-[0027], [0034]	14, 25, 36

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

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

"D" document cited by the applicant in the international application

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

03 January 2020 (03.01.2020)

Date of mailing of the international search report

29 JAN 2020

Name and mailing address of the ISA/US

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

Facsimile No. 571-273-8300

Authorized officer

Shane Thomas

Telephone No. PCT Helpdesk: 571-272-4300