



- (51) **International Patent Classification:**
G06F 11/20 (2006.01)
- (21) **International Application Number:**
PCT/US2019/026989
- (22) **International Filing Date:**
11 April 2019 (11.04.2019)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/667,174 04 May 2018 (04.05.2018) US
16/360,862 21 March 2019 (21.03.2019) US
- (71) **Applicant:** MICROSOFT TECHNOLOGY LICENSING, LLC [US/US]; Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) **Inventors:** SEELA, Pradeep; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). ZHANG, Pei; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). LOU, Yongfu; Microsoft Techno-

gy Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). **DEVARAJU, Jegan**; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). **VARADARAJAN, Krishnan**; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). **SAH, Maneesh**; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). **LANGE, Kristopher T.**; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). **MAINALI, Shane K.**; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(74) **Agent:** MANNAVA, Ashok K.; (USOC - Mannava & Kang, P.C.), Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

(54) **Title:** DATA REPLICATION IN A DISTRIBUTED STORAGE SYSTEM

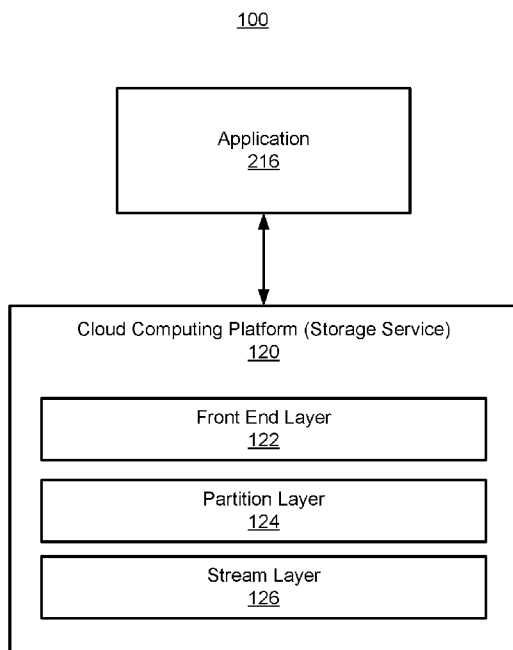


FIG. 1

(57) **Abstract:** Data is replicated between primary and secondary storage systems using a data pulling process or a log shipping process. If data is to be replicated through the log shipping process, a transaction and its associated data get stored in a transaction log in the secondary storage system. Subsequently, when the transaction gets executed, actual data gets persisted from the transaction log to a data log in the secondary storage system. If the data is to be replicated through the data pulling process, the transaction is stored in the transaction log. However, the associated data for the transaction is retrieved from the first primary storage system and stored directly in the data log in the secondary system.



(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).

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*
- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

- *with international search report (Art. 21(3))*

DATA REPLICATION IN A DISTRIBUTED STORAGE SYSTEM

BACKGROUND

[0001] A distributed storage system can support several different types of applications, from social networking to managing medical records. The distributed storage system can be implemented with user files, structured storage, and message delivery in order to support storage and workflow of the applications. Specifically, a distributed storage system can be configured to receive requests to perform transactions that facilitate the storage and workflow of the applications.

[0002] To provide failover capability, a distributed storage system may replicate data. For example, a cloud provider implementing the distributed storage system may allow tenants to replicate databases. A database may be replicated in the same data center or in a different datacenter which may be in a different geographical region. The replicated database may be available for querying or failover. For example, if the primary database fails, or simply needs to be taken offline, failover can be initiated to the secondary database so the customer can continue to access their data.

[0003] To replicate data, the secondary database is created. Once created, the secondary database is populated with the data copied from the primary database, which is commonly referred to as seeding. After the secondary database has been created and seeded, updates to the primary database may be replicated to the secondary database. The replication of the

updates can cause performance issues. For example, the replication of the updates may include committing transactions to the primary database before they are replicated to the secondary database. Accordingly, there can be a delay in updating the secondary database and providing the tenant with access

5 to their data via the second database.

BRIEF DESCRIPTION OF DRAWINGS

[0004] Embodiments and examples are described in detail in the following description with reference to the following figures. The embodiments are illustrated by examples shown in the accompanying figures in which like
5 reference numerals indicate similar elements.

[0005] Figure 1 shows an architecture of a distributed storage system, according to an embodiment;

[0006] Figure 2 shows components of a front end layer and a partition layer of the distributed storage system, according to an embodiment;

10 **[0007]** Figure 3 shows a data flow diagram for a write request, according to an embodiment;

[0008] Figure 4 shows a flow chart that may be performed by the distributed storage system, according to an embodiment;

[0009] Figure 5 shows geo replication between storage stamps,
15 according to an embodiment;

[0010] Figure 6A shows a high level flow diagram for geo replication, according to an embodiment;

[0011] Figure 6B shows a high level flow diagram for a data pulling process, according to an embodiment;

20 **[0012]** Figure 7 shows a detailed flow diagram for the data pulling process, according to an embodiment;

[0013] Figure 8 shows a flow chart for geo replication, according to an embodiment; and

[0014] Figure 9 shows an example of a computing environment, according to an embodiment.

DETAILED DESCRIPTION

[0015] For simplicity and illustrative purposes, the principles of the present disclosure are described by referring mainly to embodiments and examples thereof. In the following description, numerous specific details are set forth in order to provide an understanding of the embodiments and examples. It will be apparent, however, to one of ordinary skill in the art, that the embodiments and examples may be practiced without limitation to these specific details. In some instances, well known methods and/or structures have not been described in detail so as not to unnecessarily obscure the description of the embodiments and examples. Furthermore, the embodiments and examples may be used together in various combinations.

[0016] According to an embodiment of the present disclosure, a distributed storage system may store data in a plurality of nodes. The system may include a primary data structure that stores data, such as pointers to data objects stored in the system, and metadata for the data together. The system may also include a secondary data structure, referred to as a blocks table, that stores data and meta data for certain data objects that may be striped across partitions in the system. Certain criteria may be used to determine whether to stripe certain data objects across partitions in the system. In an example, a size threshold of a data object may be used to select data objects for striping. In an example, large data objects, having a size greater than a threshold (e.g., a 1 megabyte (MB) threshold, or a 4 MB threshold, etc.), are striped across multiple partitions to improve throughput for accessing those data objects, and the blocks table is used to store information for those data objects. In other

examples, multiple heuristics or criteria are used to select data objects for striping. For example, size, location, type of data storage (e.g., flash versus disk), etc., may be used to determine whether to stripe a data object and store it in the blocks table. In certain situations, the objects table can grow to terabytes and can become a bottleneck for accessing data objects in the system. By storing certain data objects in the blocks table, the bottleneck may be minimized and throughput for accessing those objects is greatly increased.

[0017] According to an embodiment of the present disclosure, data replication in the distributed storage system is implemented through a hybrid combination of log shipping and data pulling. For example, small data objects are replicated using log shipping, however, larger data objects are replicated by pulling the data objects from a primary location to a secondary location. For example, a large block of data, which may have a size greater than a predetermined threshold, may be striped across multiple partitions as is discussed above. To stripe the block of data, the block of data may be divided into smaller subblocks, which may have a predetermined size. These smaller subblocks are immutable. As a result, the subblocks may be pulled from the primary location to the secondary replication for data replication, because ordering of transactions maintained by log shipping may not be needed. This provides a variety of time savings for replication. For example, the subblocks may be pulled from a plurality of partition servers in parallel, reducing the amount of time to collect the data at the secondary location. Also, the data pulling reduces input/output (IO) transactions, as the log shipping systems persist logs on the secondary location before replaying them.

[0018] For purposes of a detailed discussion below, embodiments and examples of the distributed storage system are described with reference to a storage service supported by a cloud computing platform, and further by way of example, a cloud computing platform having a front end layer, partition layer, and stream layer. However, the embodiments and examples described herein can be performed in different types of operating environments having alternate configurations of the functional components described herein. As such, the embodiments and examples may be extended to other implementation contexts.

[0019] The distributed storage system can be implemented as a cloud computing platform that runs cloud services across different data centers and geographic regions. The cloud computing platform can implement a fabric controller component for provisioning and managing resource allocation, deployment/upgrade, and management of cloud services. Typically, a cloud computing system acts to store data or run service applications in a distributed manner. The service-application components (e.g., tenant infrastructure or tenancy) of the cloud computing system may include nodes (e.g., computing devices, processing units, or blades in a server rack) that are allocated to run one or more portions of a tenant's service applications.

[0020] Nodes of the distributed storage system may include virtual machines or physical machines. The nodes can concurrently run separate service applications, respectively, in individualized computing environments that support the resources and/or operating system specific to each service application. Further, each service application may be divided into functional portions such that each functional portion is able to run on a separate virtual

machine. In cloud computing systems, multiple servers may be used to run service applications and perform data storage operations in a cluster. In particular, the servers may perform data operations independently but exposed as a single device referred to as a cluster. Each server in the cluster may be referred to as a node.

[0021] A storage service on the cloud computing platform can be a service supported using the fabric controller component. The storage service can be responsible for managing the replication and data placement across disks and load balancing the data and the application traffic with storage clusters. The storage service can be responsible for managing access to a high volume of storage. The storage service can implement a storage stamp as a cluster of N racks of storage nodes and a location service that manages the storage stamps. Specifically, the location service can allocate location stamps and manage them across the storage stamps for disaster recovery and load balancing.

[0022] Figure 1 shows an architecture for a distributed storage system 100 on a cloud computing platform 120, according to an embodiment of the present disclosure. The system 100 includes a storage service including a front-end (FE) layer 122, a partition layer 124, and a stream layer 126 (or distributed file system layer). The FE layer 122 may be configured for receiving incoming requests from an application 216 or client. Upon receiving a request, the FE layer 122 can authenticate and authorize the request. The FE layer 122 can route the request to a partition server in the partition layer. A partition server is also referred to as a table server or TS. A partition map is maintained

that keeps track of partitions and which partition server is serving which partition. The FE layer 122 may access the partition map to determine the partition server serving a partition to be accessed.

- [0023]** The partition layer 124 is responsible for managing and understanding high level abstractions of user files, structured storages, and messaging. The partition layer 124 also provides a scalable index, transaction ordering, and storing consistency for objects. The partition layer 124 can specifically support storing object data on top of the stream layer 126. In an example, the partition layer partitions data objects within a storage stamp.
- 5
- [0024]** The partition layer 124 includes the primary data structure, referred to as the objects table (OT). The OT indexes blocks stored in the distributed storage system. The OT can grow to be a massive table, such as several petabytes. The OT is dynamically broken up into range partitions, referred to herein as partitions, and spread across partition servers of a storage stamp. A partition may include a contiguous range of rows in the OT from a given low-key to a high-key. For example, the storage space in the storage stamp is represented as a key space divided amongst a plurality of partitions. In an example, each row in the OT has a primary key determined from properties of the data object in the row, and the keys may be used for indexing and sorting the rows in the OT and for performing look-ups. Each key may be defined by one or more properties of the data objects stored in the system 100. For example, the primary key may include one or more properties of a data object, such as Account Name, Partition Name and Object Name. In an example, properties of data objects may be selected to generate keys for each
- 10
- 15
- 20

row in the OT table. For example, properties may be selected to generate a partition key and a row key for each row. An account name, a table name and partition key together may identify the partition within the storage service where the data object is stored. The row key may uniquely identify each row in a
5 partition in the OT. As with keys in a relational database table, partition key and row key values are indexed to create a clustered index that enables fast look-ups. The OT may be distributed across partition servers serving requests for the partitions.

[0025] Also, the partition layer 124 includes the blocks table. The blocks
10 table may also be partitioned and distributed across partition servers. The blocks table indexes blocks stored in the partition servers. The blocks table may include one or more keys defined by one or more properties of the data objects stored in the blocks table. When a binary large object (blob) or another type of data object is striped across partitions, a blocklink (e.g., a pointer) may
15 be created for each subblock, and stored in the blocks table along with other metadata. The blocklink is also stored in the OT for the blob, and may be used to retrieve the subblock for the blob. The blocks table and the OT are further discussed below.

[0026] In an embodiment, the blob data is stored in the same logs and
20 checkpoints as typical row data, but in different structures. In an additional exemplary embodiment, the blob data is stored in separate logs and checkpoints from the row data. This allows for compaction/merging of the row checkpoints without having to recopy blob data, and it allows for a single-write blob insert/update. As previously mentioned, the blob data is typically large,

therefore, abstaining from the recopying of blob data provides efficiencies. By storing the blob data in different logs and checkpoints, the expense of recopying the data can be minimized because blob data can be large (relative to traditional row data). Additionally, the separate logs and checkpoints can be stored in the same file or stream, or in the alternative, the blob data can be stored in a separate stream from the row data.

[0027] The stream layer 126 stores bits on disk or other types of storage devices and supports replicating the data across many servers to keep data durable within a storage stamp. The stream layer 126 supports files called streams which are ordered lists of large storage chunks called extents. The stream layer stores and replicates extents. The data stored in the stream layer is accessible from the partition layer 124. The stream layer 126 may provide a file system namespace and an Application Programming Interface (API) for the partition layer 124 to perform writes and reads. Writes can be append-only. The interface between the stream layer 126 and partition layer 124 allows a client to open, close, delete, rename, append to, and concatenate streams. A stream can refer to an ordered list of extent points and an extent is a sequence of append blocks. An extent can be a seal extent, such that, the extent can no longer be appended to. An application can read data from extents to access block contents of the extents.

[0028] The stream layer 126 can include a stream manager and an extent node. The stream manager is responsible for tracking the stream namespace, what extents are in each stream and the extent allocation across extent nodes. The stream manager performs lazy re-replication of extent

replicas that are lost due to hardware failures or unavailability. Each extent node maintains the storage for a set of replicas assigned to the corresponding extent by the stream manager. Each extent node contains a view about the extents it owns and where the peer replicas are for a given extent.

5 **[0029]** The distributed storage system 100 can support several different types of applications, from social networking to managing medical records. The distributed storage system 100 can specifically be implemented with user files, structured storage, and message delivery in order to support storage and workflow of the applications. The distributed storage system 100 can be
10 configured to receive transactions that facilitate the storage and workflow of the applications.

[0030] Figure 2 illustrates an example of components of the distributed storage system 100, including components of the FE layer 122 and the partition layer 124. Many of the components described herein are functional entities that
15 may be implemented as discrete or distributed components or in conjunction with other components, and in any suitable combination and location. Various functions described herein as being performed by one or more entities may be carried out by hardware, firmware, and/or software. For instance, functions and operations performed by the layers and components described herein may be
20 carried out by a processor executing instructions stored in memory.

[0031] The components of the partition layer 124 may include master 202 (e.g. partition master 202), servers 204, 206, 208, and 210 (e.g. partition servers 204, 206, 208, and 210), and FE 212 (e.g. partition FE 212). Each may reside on any type of computing device. The components may communicate

with each other over a network, which may include, without limitation, one or more local area networks (LANs) and/or wide area networks (WANs). Although a single master, four servers, and a single FE are shown in figure 2, any number of masters, servers, and FEs may be employed within the distributed storage system 100 within the scope of implementations of the present disclosure.

[0032] Stored data of the distributed storage system 100 is divided amongst a plurality of partitions. For example, a key space representing stored data objects is divided amongst the plurality of partitions. The partitions may include range partitions in the OT, as is mentioned above and as is further discussed below. Master 202 is configured to assign the partitions to partition servers 204, 206, 208, and 210, and/or other servers of a cloud computing system. Master 202 can also be configured to determine when a partition of the partitions is not hosted by any server due to a failure and reassign the partition to a different server. Master 202 is further configured to control load balancing of the partitions on servers 204, 206, 208, 210, and/or other servers of the distributed storage system 100 not shown. Additionally, master 202 is configured to monitor resource utilization with respect to any of the partitions and/or servers 204, 206, 208, and 210, and/or other servers of the distributed storage system 100 not shown. It is further contemplated that master 202 is configured to support integrated resource allocation and load balancing in accordance with embodiments of the present invention, as discussed in more detail below.

[0033] The servers, such as partition servers 204, 206, 208, and 210, are utilized to store and provide access to the data objects of the distributed storage system 100. Master 202 is configured to manage the servers. Furthermore, front end (FE) 212 is configured to provide applications, such as application 5 216, access to the distributed storage system 100. Each server, such as partition servers 204, 206, 208, and 210, can be responsible for providing read and write access from zero to many partitions assigned to the server. In the example shown in FIG. 2, server 204 is hosting partitions P1 and P4, server 206 is hosting partitions P2 and P7, server 208 is hosting partition P3, and server 10 210 is hosting partitions P5 and P6.

[0034] FE 212, which may be part of the front-end layer 122 shown in figure 1, is linked into an application, such as application 216. FE 212 may be configured to issue commands to some of the partitions (e.g. partitions P1, P2, P3, P4, P5, P6, and P7) hosted by partition servers 204, 206, 208, and 210 15 and/or other servers. Also, in some implementations, FE 212 may communicate with the application indirectly, for example, through a virtual IP and software load balancer or other means that directs communication requests. The FE 212 can utilize a partition map, such as partition map 218, to determine which of the servers is hosting (mapped to) which of the partitions 20 and can send commands to those servers. Results of the commands can be received back from the servers and can be passed to the application 216. The partition map 218 stores mappings between the partitions and the servers they are assigned to, and can typically be maintained by a master, such as partition master 202.

[0035] Data objects stored in the distributed storage system 100 may include binary large objects referred to as blobs. The distributed storage system 100 is functional to provide scalable, available, and durable blob storage in a structured storage environment. A blob is a collection of binary data stored
5 as a single entity. As a matter of example, blobs may be images, audio, records, text, or other multimedia objects. Blobs can also be executable code or other collections of binary data.

[0036] In an example, the OT is in the partition layer 124 shown in figure 1, and can store data and metadata for blobs. For example, the OT includes
10 row data for each of the blobs. The row data includes metadata corresponding to each blob, and the row data includes a pointer to a storage location for the blob. In an example, a schema for the OT may include account name, container name, blob name, and other meta data describing blobs. The schema may also include one or more pointers to the blob that describe the location of
15 the blob in the stream layer, such as a stream position. The OT may also include a blocklink for its subblocks if they are striped across partitions.

[0037] In an example, the blocks table is in the partition layer 124, and can store data and metadata for subblocks of a blob that are striped across a plurality of partitions. For example, the blocks table includes row data for each
20 subblock. The row data includes metadata corresponding to each subblock, and the row data includes a pointer to a storage location for the subblock. In an example, a schema for the blocks table may include a blocklink, account name, container name, blob name, block ID and other meta data describing the subblock. The schema may also include a pointer to the subblock. In an

example, the blocklink includes a globally unique ID (GUID) for the subblock, and the blocklink is also stored in the OT for the corresponding blob and is used as a table key linking the row for the subblock in the blocks table to the row for the blob in the OT.

5 **[0038]** According to an example of the present disclosure, a block may be striped across multiple partitions instead of being stored on a single partition.

Figure 3 shows an example of a data flow for a write request that may be striped across multiple partition servers. In this example, the write request is a put block request. Although the data flow is shown for a write request, the
10 system 100 is operable to perform read (e.g., get request) and write requests for a block which may or may not be striped across multiple partition servers.

[0039] Referring to figure 3, at (1), a FE of the distributed storage system 100, such as FE 212 shown in figure 2, receives a put block request from a user. The put block request may alternatively come from an application, such
15 as application 216 shown in figures 1 and 2. The put block request may include a blob name, shown as BNAME, assuming the block is part of a blob, and a block identifier (ID), shown as BID1, of the block to be put in the distributed storage system 100. A block is a data object. In an example, the block is for a collection of data objects that may be stored or managed as a single entity. The
20 block may be one of a plurality of blocks of a blob.

[0040] The FE 212 determines whether to stripe the incoming put block request across multiple partitions instead of storing in a single partition. For example, a pool of dedicated partitions for striping, shown as partition pool 301, may be maintained. In an example, the partition pool 301 may not be directly

connected to the FE 212. For example, the partition pool 301 is maintained by the partition layer 124 shown in figure 1, and the FE 212, which may be part of the front-end layer 122, communicates with the partition layer 124 to store or retrieve data to or from a partition from the partition pool 301.

- 5 **[0041]** Referring back to figure 3, the FE 212 determines whether the block for the put block request meets predetermined criteria for striping. For example, assume that the FE 212 determines the block size of the block in the put block request meets predetermined criteria, such as being greater than a predetermined size, and then in response to the determination, the FE 212
- 10 decides to stripe the block across a plurality of partitions from the partition pool 301. The FE 212 divides the block into subblocks. A default size may be used for the subblocks. The subblocks determined by the FE 212 are immutable. The FE 212 generates a blocklink for each of the subblocks. The blocklink includes a GUID for each of the subblocks. The GUID for each subblock
- 15 uniquely identifies the subblock across the partitions that may store the subblocks for the striped block. For example, each GUID uniquely identifies a subblock in the partition pool 301. The GUID may include a prefix that is random and is used to select a partition from the partition pool 301 to store each subblock. For example, a goal of the FE 212 may be to evenly distribute the
- 20 load of the partition pool 301. Each of the partitions in the partition pool 301 may be pre-assigned a prefix value in a range. To store each subblock, a prefix value in the range may be randomly generated to select a partition from the partition pool 301 to store the subblock to evenly distribute the load among the partitions in the partition pool 301.

[0042] At (2), the FE 212 sends putblock requests for the subblocks to the partition servers serving the applicable partitions from the partition pool 301 after determining to stripe the block. For example, the FE 212 sends a sequence of put subblock requests to the partition servers serving requests for partitions that are selected from the partition pool 301 according to the GUID prefixes selected for the subblocks. The request for each subblock may identify the GUID generated for the subblock, the block ID, the block name, and may specify other information. For each subblock request, a row in the blocks table may store the information for the request. For example, a portion of the first row of the blocks table is shown, and may include G1, which is the GUID for the first subblock, BID1, which is the block ID for the block containing the subblock, and a subblock map which specifies the partition server to host the subblock. A row may be created in the blocks table for each subblock. At (3), the FE 212 receives a status of each of the put subblock requests from the corresponding partition servers, which indicates the information has been written to the corresponding partition servers associated with the partitions from the partition pool 301 that were selected to store the subblocks.

[0043] At (4), the FE 212 sends a PutBlockLink request to the applicable partition server after getting confirmation the subblocks have been written. For example, the PutBlockLink request is sent to the partition server serving the partition associated with the blob, which was initially identified by its blob name in the putblock request at (1). For example, the FE 212 sends a PutBlockLink command to the applicable partition server storing the OT associated with the block ID. The PutBlockLink command may include the block name, blocklink,

blockID and other applicable inputs, and this information is stored in a row of the objects table of the partition server. Accordingly, the blocklink, e.g., G1, stored in the OT is a key to the blocks table to identify the location of the subblocks.

[0044] At (5), the FE 212 receives from the partition server serving the
5 OT, an indication that the row has been written for the subblock to the OT. A row may be written to the OT in the partition server for each blocklink. At (6), a status is sent to the user/application 216 that the put block request is completed.

[0045] Figure 4 illustrates a flow chart 400, according to an embodiment.
10 At 401, a new table is created, such as the blocks table, to serve as an index hosting different subblocks of a blob (or other data object) across multiple partitions. The OT may already exist. The blocks table can host subblocks of blobs from multiple accounts of a tenant of the cloud distributed storage system. At 402, the blocks table is divided into "X" partitions. In examples, "X" is
15 percentage of a total partition count. Each partition is associated with a prefix of a GUID. A pool of the partitions may be created and the partitions are served by a plurality of partition servers.

[0046] At 403, a write request is received, e.g., put block request. The write request may be for a blob. At 404, a determination is made as to whether
20 a data object to be written satisfies predetermined criteria. For example, a determination is made as to whether the block to be written has a size greater than a predetermined threshold. Other criteria may be used for the determination and may be dynamic. If no (e.g., size is less than or equal to threshold), then the block is not striped, and the block's metadata is stored in

the OT at 405. For example, the block and its metadata are written to a partition serving the OT. At 406, if yes (e.g., size is greater than the threshold), the block is divided into subblocks, and the subblocks are striped across partitions and written to the blocks tables associated with the partitions. The subblocks, their
5 blocklinks including GUIDs, and other metadata are stored in the blocks tables of the partitions. At 407, the blocklinks are stored in the OT associated with the original block, such as a blob, which is divided into the subblocks. Then, if a read request for the block is received, the FE 212 may retrieve the subblocks from the blocks tables based on the blocklinks in the OT. In an example, to
10 retrieve a striped block, a block identifier of the block is determined. Then, an entry in the OT associated with the block identifier of the block is identified. The pointer, including GUIDs, which identifies entries in the blocks table for the subblocks of the block is determined from the entry. A lookup is executed in the blocks table according to the pointer to identify the partitions storing the
15 subblocks for retrieving the subblocks.

[0047] The block tables may be transparent to the user or application 216 making the request that is received at 403. By dividing the block into subblocks and storing them in different partition servers, the subblocks may be written or read in parallel, and may improve throughput of the requests. Also, bottlenecks
20 caused by only using the OT to read and write blocks, such as blobs, are avoided.

[0048] The distributed storage system 100 may replicate a tenant's data to make available a redundant copy of the data to the tenant for disaster recovery or for other purposes. Figure 5 shows a high-level block diagram of

replication that may be performed in the distributed storage system 100. The data replication may be from a primary location 501a to a secondary location 501b. The locations 501a-b may be different geographic locations. In an example, the locations may be data centers provided in different geographic
5 locations. The different geographic locations may be different regions, so that if a catastrophe, such as a natural disaster, happens in one data center, it should not affect the other data center.

[0049] In an example, the data replication may be performed between partition layer 124a in a storage stamp in the primary location 501a and partition
10 layer 124b in a storage stamp in the secondary location 501b. Other layers, such as a front-end layers (e.g., including front end 212a and 212b) and stream layers, in the storage stamps may also be involved in the replication. A location service of the distributed storage system 100 may keep track of the primary and secondary locations 501a-b for a tenant's account, and the storage stamps in
15 the locations that are assigned to the account. For example, when provisioning the account, the location service chooses a stamp in each location and registers the AccountName with both stamps such that the US South stamp P (e.g., primary, P) takes live traffic and the US North stamp S (e.g., secondary, S) takes inter-stamp replication (also called geo-replication) traffic from stamp P for
20 the account.

[0050] When a write comes into stamp P for the account, the change is replicated within that stamp using intra-stamp replication at the stream layer then success is returned to the client. After the update has been committed in stamp P, the partition layer 124a in stamp P asynchronously geo-replicates the

change to the secondary stamp S using inter-stamp replication. When the change arrives at stamp S, the transaction is applied in the partition layer 124b and this update fully replicates using intra-stamp replication within stamp S. Since the inter-stamp replication is done asynchronously, recent updates that
5 have not been inter-stamp replicated can be lost in the event of disaster.

[0051] Inter-stamp replication is used for both account geo-replication and migration across stamps. For disaster recovery, an abrupt failover may need to be immediately performed, and recent changes may be lost. For migration a clean failover may be performed so there is no data loss. In both
10 failover scenarios, the location service makes the secondary stamp S for the account the new primary and switches the primary location for the account to the secondary stamp S.

[0052] According to an embodiment, inter-stamp replication, such as replication between storage stamps P and S in primary and secondary locations
15 501a-b, is performed using log shipping and data pulling. Log shipping may be used for replicating small data objects, such as data objects having a size less than or equal to a predetermined threshold. The threshold may be the same that is used for purposes of determining whether to stripe a data block, such as described above. In an example, log shipping may be used for data that is not
20 striped. The striping is described above with respect to figures 3-4. For larger data objects, such as a blob that is striped and has its subblocks written to the blocks table, the subblocks may be replicated using a data pulling process as is further discussed below.

[0053] Figures 6A-6B show flow diagrams for geo replication, which may include asynchronous inter-stamp replication for stamps located in different geographic regions. In particular, figure 6A shows steps for log shipping, and figure 6B shows steps for the data pulling process. One or more of the

5 components performing the geo replication may include machine readable instructions stored on a storage medium and executed by a processor to perform the operations of the component. Components, such as logs, tables, log streams, etc., may be stored on a storage medium.

[0054] Referring to figure 6A, at the primary location 501a, a client (e.g.,

10 a user or application) may communicate data to a stamp at the primary location 501a. The stamp may be associated with an account of the client. The stamp may be comprised of a plurality of partitions, such as multiple partition servers. A partition server may store memory tables and logs. For example, a partition

15 may be comprised of an update log, a block log, a page log, and/or a geo message log. In an example, a log is located within a stream layer of a distributed computing environment such that the log is a discrete data stream that is append only.

[0055] A log stream may be relied upon to re-establish data of a storage stamp following a failure of a portion of a storage stamp (e.g., a storage node).

20 For example, data may be committed to a storage stamp at a particular partition, but following the failure of the partition, the state of the partition is recreated, at least in part, by replaying one or more logs associated with that partition. Data may not be "committed" to a particular storage stamp until it is replayed into (e.g., written into) one or more logs of the storage stamp. A

distinction exists between data that is merely stored or written to a stamp and data that is committed to the stamp. For example, data may be written to a log of a storage stamp, but that data is not accessible to a client until the data is committed to the storage stamp. Committal of the data from a log, in an
5 exemplary embodiment, may occur by the playing (or replaying) of the data from the log in order, which may not be strictly sequential, but instead merely in order. Data may be received in a non-sequential manner at the secondary location 501b. However, it may be desired to commit the data in a sequential manner. The non-sequentially received data may be stored in a log, such as
10 data log 623, until a portion of the data can be committed to the storage stamp in a sequential manner.

[0056] Referring to figure 6A, a client (e.g., a user or application) may communicate data to a stamp at the primary location 501a, and the data may be committed to a partition at the stamp, and data for the transaction is written
15 to log 609. Also, a transaction that was executed to store the data in the stamp is written to log 609 or another log.

[0057] After the data is committed at the primary location 501a, geosender 610 obtains replication information from the log 609, at 631, and transmits the replication information, at 632, to a georeceiver 620 at the
20 secondary location 501b. The replication information may include the transaction executed at the primary location 501a and may include the data committed at the primary location 501a, such as for the log shipping process as is discussed in more detail below. For example, the geosender 610 communicates a geo message to georeceiver 620 including the replication

information. The replication information in the geo message may include a collection of transactions which have been batched to form a geo message for a particular geo location, stamp, and/or partition. A geo message may include a message ID that uniquely identifies that geo message (e.g., epoch # and
5 sequence #). Further yet, a geo message may be annotated to include a commit ID from the primary storage system. There can be many log entries making up a transaction, with a final commit log entry for the transaction. Information may indicate what the original transaction was, e.g., Put Block, Put Block List, etc., which is used to determine how to correctly replay the
10 transaction at the destination, and include additional information to re-execute the transaction at the destination (e.g., Last Modified Time, for Put Block List the final block list indicating if the blocks came from the committed or uncommitted list, etc.).

[0058] The replication information is written to the transaction log 621
15 from the georeceiver 620 at 633. In an example, a geo message is written almost immediately, upon receipt by the georeceiver 620 at a secondary data stamp at the secondary location 501b, in a transaction log 621. This near immediate writing of the geo message may allow for the secondary data store at the secondary location 501b to provide an acknowledgement back to the
20 supplying primary data store at the primary location 501a that the data has been written durably (but may not have yet been committed) at the secondary data store. In this example, the primary location 501a does not have to wait for acknowledgement of the re-execution of some of the geo-transactions (e.g. portions of the geo message) at the secondary location 501b. Consequently,

the geo message is written to the transaction log 621, and then replayed at a later time.

[0059] The transaction log 621, also referred to as a geomessage log (GML), at the secondary location 501b stores the replication information, including transactions, logs, data, etc. For example, for log shipping, the georeceiver 620 writes block data for each of the transactions and the transactions to the transaction log 621. This is illustrated by the below example of putblock transactions received from the client, which are executed on a primary stamp at the primary location 501a. The transactions are transmitted in a geo message to the secondary location 501b by the geosender 610, and received at the georeceiver 620, whereby the georeceiver 620 then writes the following to the transaction log 621:

PutSubBlock(Gx,0,...) → Write to .GML

PutSubBlock(Gx,1,...) → Write to .GML

15 PutSubBlock(Gx,2,...) → Write to .GML

PutSubBlockList(Gx,..) → Write PutSubBlockList to .GML.

[0060] For log shipping, the transaction and the block data are written to the transaction log 621 at 633, and the block data is subsequently written to data log 623. For example, the geo replayer 622 retrieves the block data from the transaction log 621 and stores the block data in the data log 623, such as when the transaction is dispatched as is further discussed below. In an example, transactions are received and the order of execution of the transactions is maintained in the transaction log 621. Then, the transactions are

replayed according to their order. Then, the geo replayer 622 can execute the transaction from the transaction log 621 along with persisting the block data in the data log 623 (i.e., committing the data to the data log 623) according to the stored order of execution. For example, at 635 the geo replayer 622 obtains a
5 transaction and the associated block data from the transaction log 621, and persists the data for the transaction in the data log 623 at 636 when the transaction is executed. The georeplayer 622 executes the transaction or multiple transactions, whereby the executed transactions are shown as dispatched transactions 624, and the dispatched transactions 624 should bring
10 data objects stored at the secondary location 501b to a same state of the data objects stored at the primary location 501a. For example, when a transaction is replayed from the transaction log 621, it is replayed into a corresponding Update/Block/Page logs, and the data is committed at a secondary stamp at the secondary location 501b.

15 **[0061]** As was discussed above, log shipping is a replication technique for maintaining a copy of data. A transaction log may include a file that contains a record of the changes that were made to data at the primary location. All changes are recorded in the transaction log files before the changes are written to the database at the primary location. If a database shuts down unexpectedly,
20 unfinished transactions can be restored by replaying the transaction log files into the database. Log shipping writes the log first to the transaction log 621 and then replays the transactions in the logs. For example, logs are received and the order of the transactions in the logs are maintained in the transaction log 621. Then, the transactions are replayed according to their order.

[0062] In contrast to log shipping, for the data pulling process, the steps of writing the block data to the transaction log 621 and then writing data to the data log 623 are bypassed, because the block data is pulled from the stamp at the primary location 501a and is written to a blocks table, such as blocks table 630b, at the secondary location 501b without first writing the data to the transaction log 621. This is illustrated in figure 6B. For example, as represented by the dashed line, for the block links persisted/stored in the log 609 and blocks table 630a at the primary location 501a, the data is pulled from the blocks table 630a and stored in the blocks table 630b at the secondary location 501b. Thus, the extra steps of writing data from the transaction log 621 to the appropriate locations in the secondary location 501b are eliminated. The transactions, which are written to the transaction log at 633, are executed at 635 using the data previously stored in the blocks table 630b through the data pulling process.

[0063] Figure 7 shows a data flow for the data pulling process for geo replication. The data pulling process may be applied to subblocks stored in a blocks table at the primary location 501a, whereby storing subblocks in the blocks table is described above with respect to figures 3 and 4, and the data is pulled to a blocks table at the secondary location 501b. Generally, a pull process instead of a push process is used to write data from a blocks table at the primary location 501a to a blocks table at the secondary location 501b for geo replication, such as shown in figure 6B. In the pull process, before writing metadata, such as the blocklink (i.e., pointer referencing subblock in blocks table), to the transaction log 621, the actual data is pulled from partition servers

on the primary location 501a and written to partition servers on the secondary location 501b.

[0064] Referring to figure 7, at (10), at the primary location 501a, new PutBlockLink transaction records get packed and sent as part of a geo message, which is shown as PUTGEOMESSAGE (MSG1). The new PutBlockLink transaction records that are packed into the geo message may include the “N” PUTSUBBLOCK transactions discussed with respect to step 2 of figure 3.

[0065] In figure 7, at (11), at the secondary location 501b, FE 212b receives the geo message PUTGEOMESSAGE (MSG1), and sends a put geo message command, shown as PUTGEOMSG (MSG1), to put the geo message in an OT at the secondary location 501b. For example, PUTGEOMSG (MSG1) is sent to a partition server 502, at the secondary location 501b, which serves the applicable partition of the OT, shown as objects table 503. The transaction records in MSG1 are written to the transaction log 621, e.g., GML. The data for the transactions in MSG1 is pulled from the primary location 501a as is further discussed below.

[0066] To pull the data for the transaction records, such as the “N” PUTSUBBLOCK transactions, at (12), the partition server 502 sends a command to the FE 212b to get the subblocks for the “N” PUTSUBBLOCK transactions identified in the PUTGEOMESSAGE (MSG1). The command is shown as FETCHBLOCKDATA (G1, BID1, BNAME). The FETCHBLOCKDATA may identify the GUIDs and other information for the subblocks to be pulled from the primary location 501a. At (13), the FE 212b transmits a request (e.g.,

GETBLOCKDATA (G1, BNAME)) to FE 212a at the primary location 501a for pulling the actual data that was persisted at the primary location 501a under the “N” PUTSUBBLOCK transactions, which are identified in the PUTGEOMESSAGE (MSG1). The request identifies the subblocks in the
5 blocks table, for example, using the GUIDs and the block name.

[0067] At (14)-(16), the FE 212a gets the subblocks from the blocks table 630a. For example, as shown at (14) and (15), a GETBLOCKDATA command identifying the subblock is sent to the blocks table 630a, and the requested subblock is received from the blocks table 630a. At (16), the FE 212a sends the
10 requested subblock to the FE 212b. This may be performed for all the subblocks for the “N” PUTSUBBLOCK transactions.

[0068] At (17)-(18), the FE 212b receives the subblocks, and follows a similar workflow to the one performed on the primary location side, such as by sending N PutSubBlocks to the blocks table partitions serving the subblocks.
15 Upon receiving a completion notification, the partition server serving the applicable blocks table partition writes the data for the N PutSubBlocks to the data log 623.

[0069] At (19), the partition server 502 sends a response to the FE 212b that the write for the original PutGeoMessageCommand from step 10 is
20 completed, and, at (20), the FE 212b sends a status to the partition server at the primary location 501a that the original PutGeoMessageCommand from step 10 is completed. Now the transaction log 621 has the transactions and the actual data for the transactions is stored in the data log 623, and the georeplayer 622 in figure 6B can execute the transactions. In an example, the data log 623 may

include the blocks table 630b. In this case, step 18 may be bypassed because the data for the transactions is already stored in its appropriate location at (17), and the georeplayer 622 execute the transactions whose associated block data is already persisted in the blocks table 630b.

5 **[0070]** There are technical advantages to the data pulling process as opposed to the log shipping process for geo replication. For log shipping, on the geo receiver/replayer side, there is a redundant write for the actual block data. For example, for log shipping, data is first written to the GML at the secondary location side, and then during replay, the data is read and written to
10 a blob log. For the data pulling process, the initial data write to the GML of the OT that is performed for log shipping is avoided as the actual data is written only once in the blocks table log stream. Also, the data read from the GML during replay is avoided as replay of block links involve only metadata. In addition, the geo sender/receiver throughput is increased, because pull
15 requests can be simultaneously issued to different FEs. Also, faster geo replayer throughput is achieved as there is no actual data write that happens during the replay phase from the GML.

[0071] Figure 8 illustrates a method for geo replication. At 801, data, such as data for a blob, is striped and stored in the blocks table on the primary
20 location side, such as described with respect to figures 3 and 4. At 802, a message is received at the FE on the secondary location side to put a transaction in the GML. The FE determines whether the message is for a transaction that includes data stored in the blocks table. For example, the message includes a GUID of the data in the blocks table if it is stored in the

blocks table. If the data is in the blocks table, at 803, the data pulling process is performed. For example, at 804, the FE on the secondary location side sends a message to the FE of the primary location side to retrieve the data from the blocks table on the primary location side. The message, for example, includes
5 the GUID of the data in the blocks table and other metadata for the data. At 805, the FE on the secondary location side receives the requested data from the FE on the primary location side, and at 806, the data for the transactions is written to the data log. Also, the transactions are written to the transaction log. The georeplayer can then execute the transactions in the transaction log using
10 the data in the data log, such as shown in figure 6B. If it is determined at 802 that the data is not in the blocks table, then the transaction and data for the transaction are written to the transaction log according to a log shipping process at 807, such as shown in figure 6A.

[0072] Figure 9 shows an example of a suitable computing and
15 networking environment on which the examples and embodiments described above may be implemented. The computing system environment is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the examples and embodiments described above. The examples and embodiments described
20 above is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with the examples and embodiments described above include, but are not limited to: personal computers, server computers, hand-held or laptop

devices, tablet devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

5 **[0073]** The examples and embodiments described above may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, and so forth, which perform particular tasks or implement particular abstract data types. The

10 examples and embodiments described above may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in local and/or remote computer storage media including memory storage devices.

15 **[0074]** Figure 9 shows an example of a system for implementing various aspects of the example and embodiments described herein and may include a computing device in the form of a computer 900. Components may include, but are not limited to, various hardware components, such as processing unit 901, data storage 902, such as a system memory, and system bus 903 that couples

20 various system components including the data storage 902 to the processing unit 901. The system bus 903 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus,

Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

[0075] The computer 900 typically includes a variety of non-transitory
5 computer-readable media 904. Computer-readable media 904 may be any
available media that can be accessed by the computer 900 and includes both
volatile and nonvolatile media, and removable and non-removable media, but
excludes propagated signals. By way of example, and not limitation, computer-
readable media 904 may comprise computer storage media and communication
10 media. Computer storage media includes volatile and nonvolatile, removable
and non-removable media implemented in any method or technology for
storage of information such as computer-readable instructions, data structures,
program modules or other data. Computer storage media includes, but is not
limited to, RAM, ROM, EEPROM, flash memory or other memory technology,
15 CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic
cassettes, magnetic tape, magnetic disk storage or other magnetic storage
devices, or any other medium which can be used to store the desired
information and which can be accessed by the computer 900. Communication
media typically embodies computer-readable instructions, data structures,
20 program modules or other data in a modulated data signal such as a carrier
wave or other transport mechanism and includes any information delivery
media. The term “modulated data signal” means a signal that has one or more
of its characteristics set or changed in such a manner as to encode information
in the signal. By way of example, and not limitation, communication media

includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media.

Combinations of the any of the above may also be included within the scope of computer-readable media. Computer-readable media may be embodied as a
5 computer program product, such as software stored on computer storage media.

[0076] The data storage or system memory 902 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) and random access memory (RAM). A basic input/output
10 system (BIOS), containing the basic routines that help to transfer information between elements within computer 900, such as during start-up, is typically stored in ROM. RAM typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 901. By way of example, and not limitation, data storage 902 holds an
15 operating system, application programs, and other program modules and program data.

[0077] Data storage 902 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, data storage 902 may be a hard disk drive that reads from or writes to
20 non-removable, nonvolatile magnetic media, a magnetic disk drive that reads from or writes to a removable, nonvolatile magnetic disk, and an optical disk drive that reads from or writes to a removable, nonvolatile optical disk such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary

operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The drives and their associated computer storage media, described above and illustrated in FIG. 9, provide storage of computer-
5 readable instructions, data structures, program modules and other data for the computer 900.

[0078] A user may enter commands and information through a user interface 909 or other input devices such as a tablet, electronic digitizer, a microphone, keyboard, and/or pointing device, commonly referred to as mouse,
10 trackball or touch pad. Other input devices may include a joystick, game pad, satellite dish, scanner, or the like. Additionally, voice inputs, gesture inputs using hands or fingers, or other natural user interface (NUI) may also be used with the appropriate input devices, such as a microphone, camera, tablet, touch pad, glove, or other sensor. These and other input devices are often connected
15 to the processing unit 901 through a user interface 909 that is coupled to the system bus 903, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 906 or other type of display device is also connected to the system bus 903 via an interface, such as a video interface. The monitor 906 may also be integrated
20 with a touch-screen panel or the like. Note that the monitor and/or touch screen panel can be physically coupled to a housing in which the computing device 900 is incorporated, such as in a tablet-type personal computer. In addition, computers such as the computing device 900 may also include other peripheral

output devices such as speakers and printer, which may be connected through an output peripheral interface or the like.

[0079] The computer 900 may operate in a networked or cloud-computing environment using logical connections to one or more remote devices, such as a remote computer. The remote computer may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 900. The logical connections depicted in FIG. 9 include one or more local area networks (LAN) and one or more wide area networks (WAN), but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0080] When used in a networked or cloud-computing environment, the computer 900 may be connected to a public or private network through a network interface 907. In some embodiments, a modem or other means for establishing communications over the network. The modem, which may be internal or external, may be connected to the system bus 903 via the network interface 907 or other appropriate mechanism. A wireless networking component such as comprising an interface and antenna may be coupled through a suitable device such as an access point or peer computer to a network. In a networked environment, program modules depicted relative to the computer 900, or portions thereof, may be stored in the remote memory storage device. It may be appreciated that the network connections shown are

exemplary and other means of establishing a communications link between the computers may be used.

[0081] The computer 900 may be a server or another type of computer.

The computer 900 may include a platform for hosting one or more components
5 of the partition layer or the FE layer.

[0082] Embodiments and examples are described above, and those skilled in the art will be able to make various modifications to the described embodiments and examples without departing from the scope of the embodiments and examples.

What is claimed is:

1. A computing device in a distributed storage system operable to perform data replication from a primary storage system to a secondary storage system, the computing device comprising:

at least one memory containing machine-readable instructions;

and

at least one processor to execute the machine-readable instructions to:

10 determine whether a block to be replicated from the primary storage system to the secondary storage system is to be replicated via a data pulling process or a log shipping process;

in response to determining the block is to be replicated through the log shipping process, store a transaction and data for the block in a transaction log, and subsequently execute the transaction from the transaction log to cause the data to be written from the transaction log to a data log in the secondary storage system; and

15 in response to determining the block is to be replicated through the data pulling process, store the transaction in the transaction log, retrieve the data for the block from the primary storage system, and store the data in the data log.

2. The computing device of claim 1, wherein for the data pulling process, the data is stored in the data log without previously storing the data in the transaction log.
- 5 3. The computing device of claim 1, wherein the transaction is performed on the primary storage system, and the at least one processor is to receive a message from a computer in the primary storage system to replicate the transaction on the secondary storage system.
- 10 4. The computing device of claim 1, wherein to determine whether to replicate via a data pulling process or a log shipping process, the at least one processor is to:
- determine whether the data for the block is stored in a blocks table in the primary storage system, wherein the data is stored in the blocks
- 15 table in response to the block being striped across a plurality of partitions in the primary storage system.
5. The computing device of claim 4, wherein the block is striped across the plurality of partitions in the primary storage system if a size of the
- 20 block is greater than a predetermined threshold.

6. The computing device of claim 4, wherein to stripe the block across the plurality of partitions, the primary storage system is to:
- receive a write request for a block at the primary storage system;
 - determine whether the block satisfies predetermined criteria;
 - 5 in response to the block satisfying the predetermined criteria, striping the block across a plurality of partitions in a pool of partitions,
 - wherein striping the block across the plurality of partitions comprises storing subblocks and metadata for the subblocks, including globally unique identifiers, in a partitioned blocks tables associated with the plurality of
 - 10 partitions; and
 - storing a link to entries for the blocks table in an objects table associated with the block.
- 15 7. The computing device of claim 6, wherein a portion of each of the globally unique identifiers is randomly selected for load distribution across the plurality of partitions.
8. The computing device of claim 6, wherein in response to the block
- 20 not satisfying predetermined criteria, storing the block and its metadata in an objects table.

9. The computing device of claim 8, wherein for the log shipping process the data is retrieved from the objects table in the primary storage system.

5 10. A computer-implemented method for replicating data from a first storage system to a second storage system, the method comprising:

determining whether a block to be replicated from the first storage system to the second storage system is to be replicated via a data pulling process or a log shipping process;

10 in response to determining the block is to be replicated through the log shipping process, storing a transaction and data for the block in a transaction log; and

in response to determining the block is to be replicated through the data pulling process, storing the transaction in the transaction log, retrieving
15 the data for the block from the first storage system, and storing the data in a data log.

11. The method of claim 10, wherein for the data pulling process, the data is stored in the data log without previously storing the data in the
20 transaction log.

12. The method of claim 10, wherein the transaction is performed on the first storage system, and the method comprises receiving a message from a computer in the first storage system to replicate the transaction on the second storage system.

5

13. The method of claim 10, wherein determining whether to replicate via a data pulling process or a log shipping process comprises:

determining whether the data for the block is stored in a blocks table in the first storage system, wherein the data is stored in the blocks table in

10 response to the block being striped across a plurality of partitions in the first storage system.

14. The method of claim 13, wherein to stripe the block across the plurality of storage partitions, comprises:

15 receiving a write request for a block at the first storage system;

determining whether the block satisfies predetermined criteria;

in response to the block satisfying the predetermined criteria,

striping the block across a plurality of partitions in a pool of partitions,

20 wherein striping the block across the plurality of partitions comprises storing subblocks and metadata for the subblocks, including a

globally unique identifiers, in a partitioned blocks tables associated with the plurality of partitions; and

storing a link to entries for the blocks table in an objects table associated with the block.

5

15. The method of claim 14, wherein a portion of each of the globally unique identifiers is randomly selected for load distribution across the plurality of partitions.

1/10

100

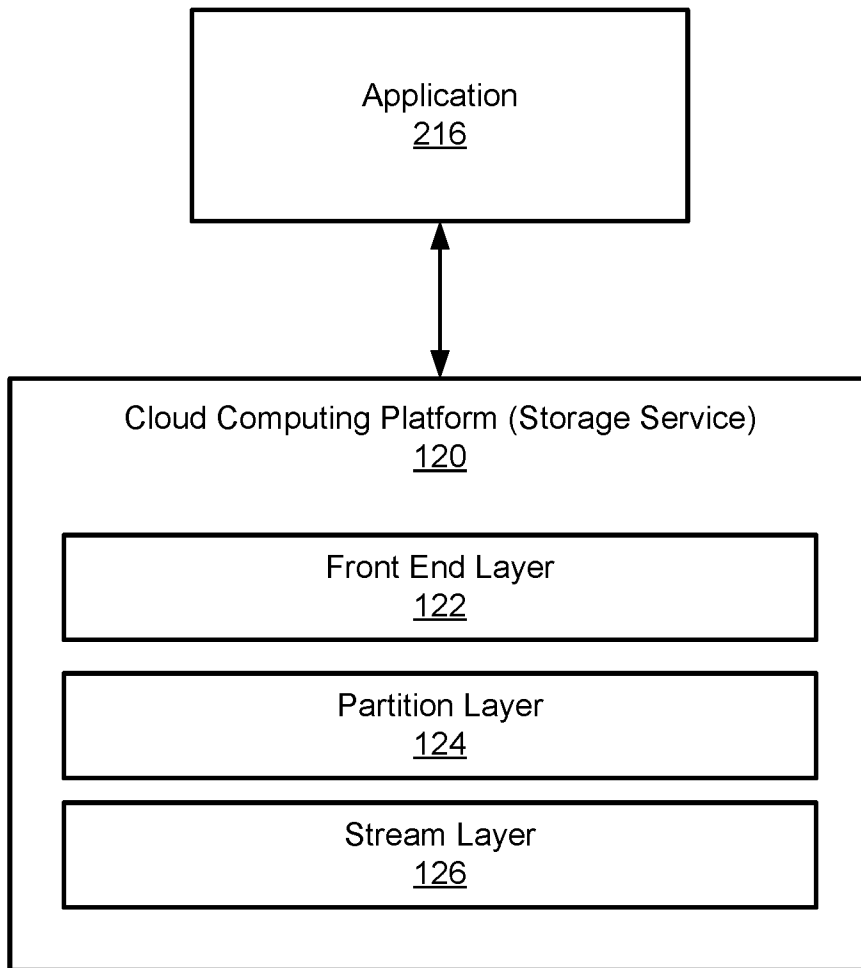


FIG. 1

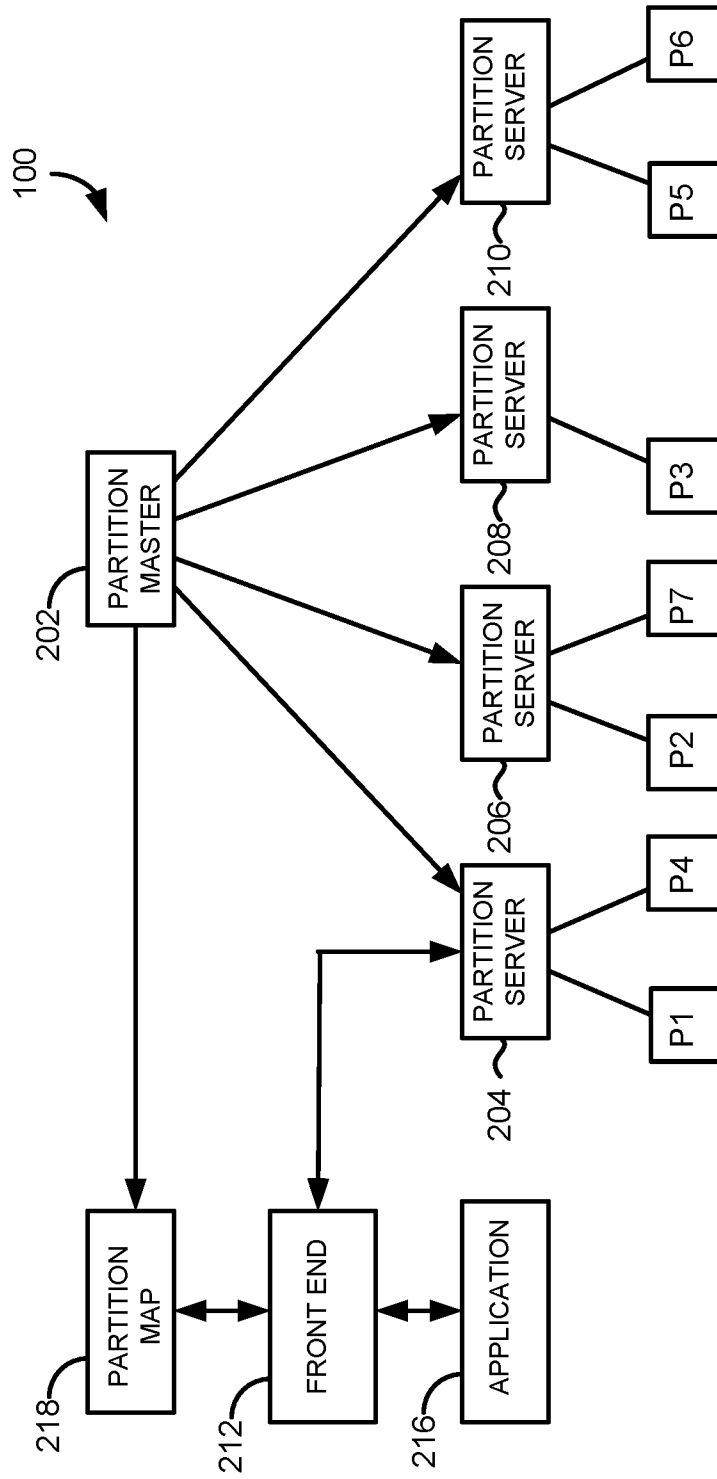


FIG. 2




FIG. 3

4/10

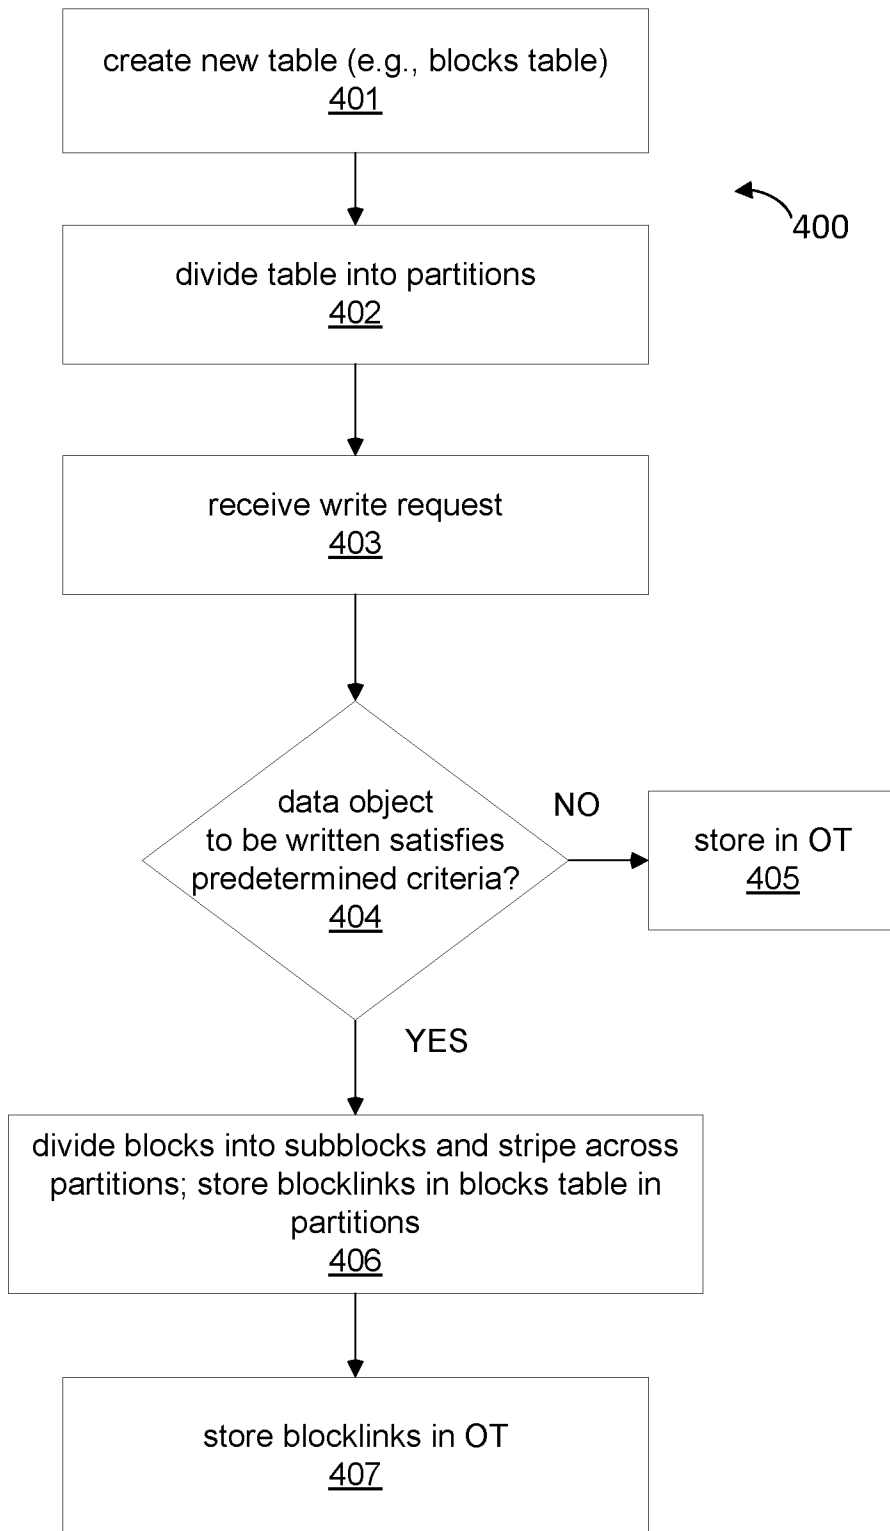


FIG. 4

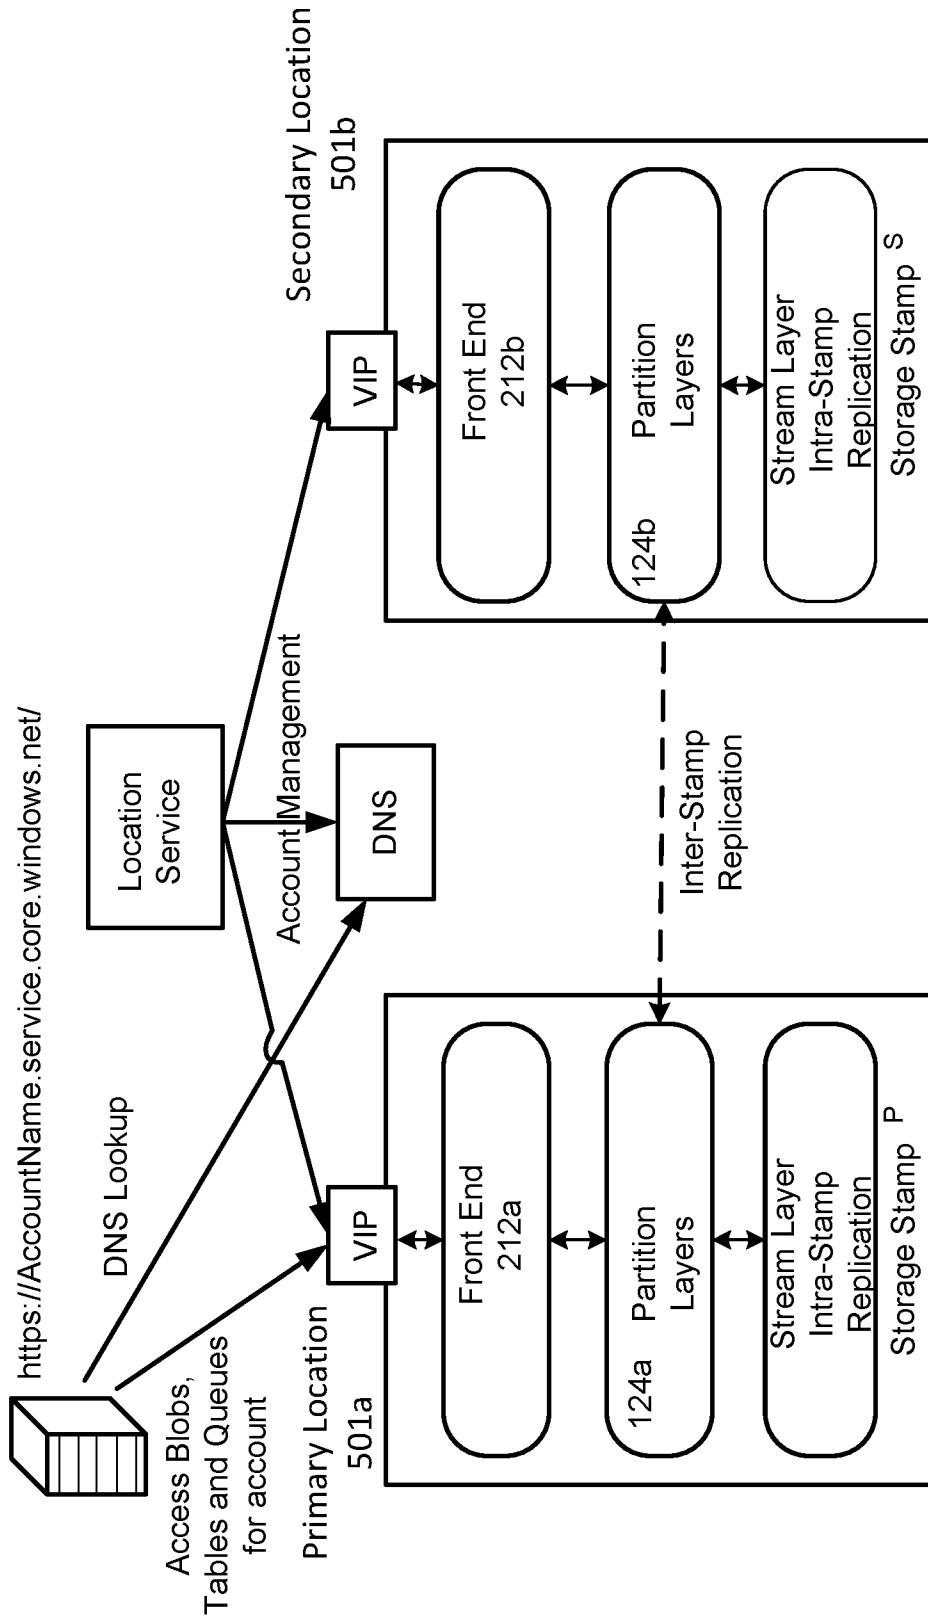


FIG. 5




FIG. 6A




FIG. 6B




FIG. 7

9/10

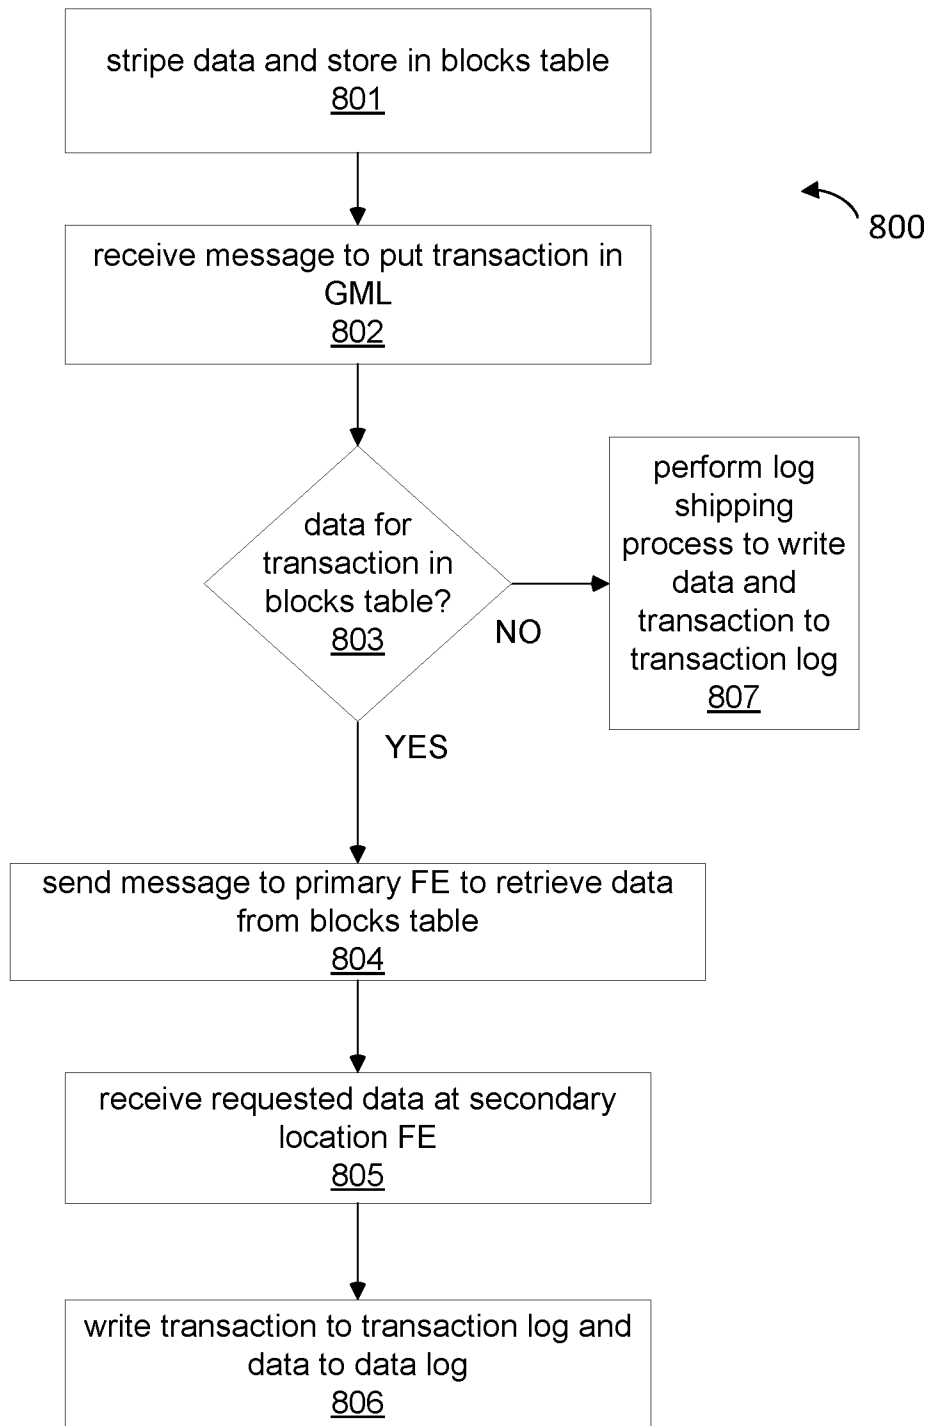


FIG. 8




FIG. 9

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/026989

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F11/20
ADD.
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)
G06F
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
EPO-Internal, WPI Data, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2012/084260 A1 (CHERKAUER KEVIN J [US] ET AL) 5 April 2012 (2012-04-05) abstract paragraphs [0001] - [0004], [0013] - [0018], [0028] - [0057] figures 1A, 2-4, 6	1-15
Y	US 2018/046551 A1 (BOURBONNAIS SERGE [US] ET AL) 15 February 2018 (2018-02-15) abstract paragraphs [0002], [0005] - [0007], [0018] - [0046], [0052] - [0063], [0082] - [0096] claims 1, 9 figures 1, 2, 4, 5	1-15

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	"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
"E" earlier application or patent but published on or after the international filing date	"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
"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)	"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
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 1 July 2019	Date of mailing of the international search report 10/07/2019
---	---

Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Johansson, Ulf
--	---

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/026989

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2011/066592 A1 (NEWPORT WILLIAM T [US] ET AL) 17 March 2011 (2011-03-17) abstract paragraphs [0004] - [0008], [0015], [0025] - [0032] figures 1, 3, 4 -----	1-15
A	EP 1 400 899 A2 (QUANTUM CORP [US]) 24 March 2004 (2004-03-24) paragraphs [0003] - [0006], [0009] - [0014], [0017] - [0047] figures 3A, 3B, 3C, 4B, 4C, 5C, 6A, 6B -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2019/026989

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2012084260 A1	05-04-2012	US 2012084260 A1 US 2018101558 A1	05-04-2012 12-04-2018

US 2018046551 A1	15-02-2018	US 2018046551 A1 US 2018113766 A1	15-02-2018 26-04-2018

US 2011066592 A1	17-03-2011	JP 5496839 B2 JP 2011060292 A KR 20110029071 A US 2011066592 A1 US 2013041869 A1	21-05-2014 24-03-2011 22-03-2011 17-03-2011 14-02-2013

EP 1400899 A2	24-03-2004	EP 1400899 A2 JP 2004118837 A US 2004059869 A1 US 2006206665 A1	24-03-2004 15-04-2004 25-03-2004 14-09-2006
