



- (51) International Patent Classification:
G06F 11/14 (2006.01) *G06F 11/20* (2006.01)
- (21) International Application Number:
PCT/US2016/035949
- (22) International Filing Date:
6 June 2016 (06.06.2016)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/173,670 10 June 2015 (10.06.2015) US
14/804,000 20 July 2015 (20.07.2015) 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: NARAYANAN, Dushyanth; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). DRAGOJEVIC, Aleksandar; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US). CASTRO, Miguel Oom Temudo de; Microsoft Technology Licensing, LLC, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (74) Agents: MINHAS, Sandip et al.; Microsoft Corporation, Attn: Patent Group Docketing (Bldg. 8/1000), One Microsoft Way, Redmond, Washington 98052-6399 (US).

- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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,

[Continued on next page]

(54) Title: RECOVERY IN DATA CENTERS



(57) Abstract: A server at a cluster of servers in a data center is described. The server comprises a memory which is part of a distributed memory of the cluster. The server has at least one processor executing transactions and lock-free reads on software objects stored in regions of the distributed memory, the software objects and details of the transactions being replicated in the distributed memory. The server has a network interface card arranged to receive a message indicating a new configuration of the cluster comprising addition, removal or potential failure of at least one of the other servers. The processor is configured to use a recovery process which enables the lock-free reads and committed ones of the transactions to retain the properties of atomicity, consistency, isolation and durability across configuration changes.



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).

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

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

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

RECOVERY IN DATA CENTERS

BACKGROUND

[0001] Data centers comprising large numbers of servers connected via a network are increasingly used for cloud computing services and are under high demand. As demand increases so do situations where one or more servers in the data center fail and this leads to outages of cloud computing services and other errors. Data center operators often need to add more servers to a data center in order to scale up to meet demand, or to remove servers for maintenance or upgrade.

SUMMARY

[0002] The following presents a simplified summary of the disclosure in order to provide a basic understanding to the reader. This summary is not intended to identify key features or essential features of the claimed subject matter nor is it intended to be used to limit the scope of the claimed subject matter. Its sole purpose is to present a selection of concepts disclosed herein in a simplified form as a prelude to the more detailed description that is presented later.

[0003] A server at a cluster of servers in a data center is described. The server comprises a memory which is part of a distributed memory of the cluster. The server has at least one processor executing transactions and lock-free reads on software objects stored in regions of the distributed memory, the software objects and details of the transactions being replicated in the distributed memory. The processor can receive a message indicating a new configuration of the cluster comprising addition, removal or potential failure of at least one of the other servers. The processor is configured to use a recovery process which enables the lock-free reads and committed transactions to retain the properties of atomicity, consistency, isolation and durability across configurations. Many of the attendant features will be more readily appreciated as the same becomes better understood by reference to the following detailed description considered in connection with the accompanying drawings.

DESCRIPTION OF THE DRAWINGS

[0004] The present description will be better understood from the following detailed description read in light of the accompanying drawings, wherein:

FIG. 1 is a schematic diagram of a data center with a cluster of servers having a new recovery process and showing the data center being used to provide cloud services to end user devices and web servers;

FIG. 2 is a schematic diagram of a data center server such as one of the servers of FIG. 1;

FIG. 3 is a schematic diagram of a configuration of servers and configuration manager at a data center server cluster;

FIG. 4 is a schematic diagram of a transaction protocol;

FIG. 5 is a flow diagram of a recovery process in a data center server cluster;

FIG. 6 is a flow diagram of a method at a server of a data center as part of a recovery process;

FIG. 7 is a flow diagram of a process of transaction state recovery in a data center;

FIG. 8 is a flow diagram of a process at a server of a data center as part of a recovery process;

FIG. 9 is a flow diagram of a process at a server of a data center which is a coordinator of transaction A and of a process at a server of the data center which is a primary of a region participating in transaction A;

FIG. 10 is an exemplary computing-based device in which embodiments of a server of a data center may be implemented, such as the server of FIG. 2.

Like reference numerals are used to designate like parts in the accompanying drawings.

DETAILED DESCRIPTION

[0005] The detailed description provided below in connection with the appended drawings is intended as a description of the present examples and is not intended to represent the only forms in which the present example may be constructed or utilized. The description sets forth the functions of the example and the sequence of steps for constructing and operating the example. However, the same or equivalent functions and sequences may be accomplished by different examples.

[0006] There is an increasing need for applications which run in data centers to provide cloud computing services which aim to be reliable. For example, the aim may be to provide database transactions which have atomicity, consistency, isolation and durability (known as ACID transactions). A transaction may be a single logical operation on a database, such as a distributed memory of a cluster of servers at a data center. For

example, a transaction may be a transfer of data from one data record to another, perhaps involving multiple changes. In the examples described below a database may be a distributed memory logically defined over main memory of a plurality of servers of a data center.

5 **[0007]** Existing data center deployments typically do not provide ACID distributed transactions because previous attempts to do so have led to poor performance and/or poor availability. Data center designers have tended to avoid transactions completely, to weaken consistency guarantees, or to provide single-machine transactions that require programmers to partition their data.

10 **[0008]** In situations where failures in a data center occur, such as one or more servers crashing, it is extremely difficult to automatically recover from the failure situation in a way in which transactions in the data center remain ACID. This also applies where a data center seeks to recover from other configuration changes, such as addition of more servers or removal of one or more servers for maintenance or upgrade. A configuration
15 change is any change in the data center infrastructure, such as a change in the arrangement of servers, or a change in the number of servers which are operational.

[0009] The embodiments described below are not limited to implementations which solve any or all of the disadvantages of known data center technology.

[0010] FIG. 1 is a schematic diagram of a data center 100 comprising a plurality of
20 servers 102 connected via a network 104. The data center implements a new recovery process 124 which is stored at the servers 102 and which enables a cluster of servers at the data center to automatically recover from configuration changes, in a way which facilitates any committed transactions (including lock-free reads) implemented in the data center to be ACID despite the recovery. The new recovery process may also facilitate recovery
25 such that any committed transactions implemented in the data center server cluster retain serializability, and in some cases, strict serializability. Serializability of a set of transactions is a well known term in the art. In summary serializability of a set of transactions means that externally observable effects of executing the transactions are the same as if the transactions were executed in some sequential ordering with no overlap in
30 time, with each transaction executing instantaneously at a single point in time. Strict serializability means this point in time where a transaction appears to execute, is always between the start of transaction execution and the point when transaction completion is reported.

[0011] The recovery process works to facilitate any committed transactions (including lock-free reads) implemented in the data center to be ACID despite the recovery in a variety of ways. Some transactions whose commit phase spans configuration changes wait for the recovery protocol to decide if they commit or abort.

5 These are transactions that are affected by the configuration change. Transactions that are not affected by the configuration change are able to execute a normal commit protocol whereas other transactions wait for the recovery process to decide their outcome.

[0012] Atomicity is the characteristic that a transaction is “all or nothing”. That is, if one part of the transaction fails, the whole transaction fails and the database state is left
10 unchanged. Consistency provides that a transaction brings the database from one valid state to another. Isolation provides that transactions do not interfere with one another. Durability provides that once a transaction has been committed, it will remain so. The recovery protocol acts to enable atomicity even in the case of recovery from power failures, crashes and other configuration changes. The recovery protocol achieves this at
15 least in part by identifying transactions whose commit phase spans configuration changes and using the recovery protocol to decide if they commit or abort so as to achieve atomicity, consistency, isolation, and durability. In some examples, the recovery protocol provides the properties of ACID transactions at least in part by configuring servers in the data center such that replies to one-sided read requests received at the server are refused at
20 least on the basis of a configuration identifier associated with the request. In some examples the recovery protocol provides the property of durability at least in part by using replication in the distributed memory of the data center and using transaction details. In some examples the recovery protocol provides the properties of ACID transactions at least in part by configuring servers in the data center to recover write locks of transactions
25 already being committed at the time of the new configuration, the lock recovery process operating to recover the write locks from replicated transaction logs, which may have been written with one-sided writes.

[0013] A configuration change can result from failure of one or more servers, from addition or removal of servers for maintenance, or other reasons. A lock-free read is a
30 round trip request, from a requesting server to a region at another server storing an object that returns the value of the object. It is performed with a single one-sided read, which means that the processor on the server that has the data is not involved in processing the

request. A lock-free read gives the same outcome as a transaction that reads a single object and writes no objects. .

[0014] The servers 102 of a cluster at the data center 100 execute one or more applications in order to provide cloud services to end user equipment 122, web servers 120
5 or other computing entities via a network 108 of any suitable type such as the internet or an intranet. The end user equipment can include tablet computers 110, personal desktop computers 112, smart phones 114, wearable computers 116, laptop computers 118, and others.

[0015] Six servers (also referred to as machines) 102 are shown in FIG. 1 although
10 in practice many thousands of servers may be present. One or more agreement nodes 106 are provided to enable configurations of servers to be deployed as described in more detail below. Other components such as load balancers may be present in the data center 100 although these are not illustrated in FIG. 1 for clarity.

[0016] The data center 100 comprises a plurality of agreement nodes 106 which
15 together provide an agreement service for agreeing new configurations of the servers 102 as described in more detail below. The agreement service uses data which is replicated over the plurality of agreement nodes 120 so that the agreement service is fault tolerant.

[0017] The servers of the cluster each have a local memory which together are
20 treated as a distributed memory in the examples mentioned below. The memory may be divided into regions as described in more detail below, and the regions are replicated over the cluster of servers. Using replication in this way facilitates recovery because data from regions at a failed server can be recovered from replicas of those regions elsewhere in the cluster.

[0018] FIG. 2 is a schematic diagram of a data center server 120 such as one of the
25 servers of FIG. 1. It comprises one or more central processing units 200, 202, a non-volatile memory 204 and a network interface card 206 or other means for communicating with other entities in the data center via network 104. The non-volatile memory 204 stores data in a manner such that it is retained in the event of the server crashing or losing power. When the server re-starts, data stored in the non-volatile memory 204 becomes available
30 again. This can assist with making transactions implemented in the data center have the properties of ACID transactions even when all servers lose power. For example, the recovery process is able to use data from the non-volatile memory once it becomes available again, although that data may need updating.

[0019] In the examples described below, non-volatile memory is used, which may be implemented using dynamic random access memory (DRAM), a battery and a persistent store. However, it is not essential to use non-volatile memory. The data center server may use only volatile memory such as DRAM in some examples. This is because the replication of the distributed memory regions enables data to be recovered in situations where only volatile memory is used. By using non-volatile memory in combination with replication of the distributed memory regions, recovery is facilitated where there are multiple failures in a cluster of servers at the data center, for example, when all replicas for a software object crash at the same time.

[0020] It is possible for some of the servers of the data center to have only volatile memory and others of the data center servers to have non-volatile memory.

[0021] In some examples the non-volatile memory is implemented using local energy storage such as where batteries are used at each server together with a commodity solid state drive (SSD) or similar persistent storage. When a power failure occurs, the server saves the contents of non-persistent memory at the server (such as dynamic random access memory, DRAM) to the SSD using the energy from the battery. In this way, plentiful, cheap DRAM in data center servers can be used in a way which seeks to persistently store data and which thus facilitates making transactions have ACID properties.

[0022] In other examples non-volatile dual in-line memory modules (NVDIMMs) are used that contain their own private flash memory, controller and supercapacitor. Other types of nonvolatile memory may be used such as phase-change memory and memristors.

[0023] The network interface card 206 may be configured to enable one-sided operations which do not involve the CPU of the server 102. For example, the network interface card 206 may receive a one-sided read request from another of the data center servers over network 104. The network interface card reads the requested data from the non-volatile memory 204 using DMA 208 (direct memory access). The request from the remote server is referred to as a remote direct memory access (RDMA) because it is a request from the memory of one computer into another computer which involves the CPU of the computer initiating the RDMA request but not the CPU of the computer that responds to the RDMA request.

[0024] Using one sided operations in this manner gives the benefit of reducing load on the CPUs. Through empirical investigation the inventors have identified that use

of one-sided RDMA reads and writes where possible to implement a transaction protocol in one or more clusters of servers at a data center gives improved performance. Having said that, it is not essential to use one-sided operations. Two-sided operations in which CPUs are involved may also be used.

5 [0025] FIG. 3 is a schematic diagram of a configuration of servers with a configuration manager at a data center. A configuration comprises a plurality of servers 102 at a data center and where one of the servers is a configuration manager 314. A configuration manager is arranged to allocate new regions of a distributed memory of the plurality of servers. It holds a region mapping 310 of the distributed memory. The distributed memory is illustrated with dotted lines in FIG. 3 as shared global address space 308 to indicate that it is not a shared physical memory but is a memory address space that refers to memory regions which are local at the servers 102. In the example in FIG. 3 memory regions 304, 306 at the local servers are in non-volatile memory 204 but they could be in volatile memory in some examples as mentioned above.

15 [0026] A memory region (of the distributed memory) can be a primary region 304 or a back-up region 306. A primary region is one which stores a primary copy of data and a back-up region is a replica of its counterpart primary region. Back-up regions are typically at different servers from the counterpart primary region. One or more back-up regions per primary region may be used. Regions store data such as software objects of the applications. In this way objects may be replicated in two or more machines.

20 [0027] The configuration manager uses a counter 312 to allocate new primary memory regions as described in more detail below.

[0028] A server 102 of the data center, including the configuration manager 314 which is a type of server 102, has one or more logs 302 which are stores such as buffers, ring buffers, or other stores held in non-volatile memory where the server has non-volatile memory available. The logs 302 store control data for use in implementing the recovery process. For example, version numbers, configuration identifiers, transaction identifiers and other data as described in more detail below.

25 [0029] The data center comprises a mechanism for automatically detecting configuration changes. In some examples this mechanism is a lease scheme used to detect potential server failures. In this case a server 102 stores a lease 300 and the configuration manager 314 stores lease information 300 for each server in the currently active configuration (there is a single active configuration at any point in time). Messages are

sent between the configuration manager 314 and the servers of the configuration to request and renew leases. When a lease expires, a message from the corresponding machine has not been received for the lease period, and so the machine is suspected to have failed. .

[0030] In some examples the data center comprises a mechanism for agreeing a new configuration, for example, if one or more servers fail or if one or more servers are added/removed. The agreement nodes 106 may manage this process. In some examples, the process of agreeing a new configuration uses a configuration protocol known to the servers and the agreement nodes. The configuration protocol may implement precise membership as explained in more detail below. By implementing precise membership the recovery process is able to work even where one-sided reads are involved. By using a configuration protocol which implements precise configuration membership, cluster recovery is facilitated even where one-sided operations are used during execution of applications at the cluster.

[0031] In some examples the network interface cards are configured to support leases with a timer, and to reject requests that arrive after the NIC's lease timer has expired without being renewed. In this case precise membership is not necessary and the recovery process is still able to work even where one-sided reads are involved.

[0032] The servers of the configuration, and the configuration manager 314 have knowledge of the recovery process as mentioned with respect to FIG. 1. The servers 102 and configuration manager 314 also know a transaction protocol and implement this transaction protocol in order to read and write data to and from the distributed memory. One or more applications execute on the servers to provide the cloud service(s). A thread of one of the applications executes a transaction to read and/or write data to the distributed memory and a transaction protocol is used to implement this process. The transaction protocol may be one which facilitates ACID properties of the transaction and which facilitates serializability of transactions and lock-free reads.

[0033] During operation of the data center multiple transactions execute on different threads and different servers and these transactions can be in any phase of execution when a configuration change occurs. This makes recovery from failures, or other configuration changes, extremely difficult.

[0034] The transaction protocol has an execution phase and a commit phase. An individual transaction is executed by a coordinator thread at one of the servers for the life of the transaction unless recovery from a configuration results in a new coordinator thread

for the transaction. It is not necessary to have a single coordinator thread for the whole system. Multiple transactions may run in parallel, either in the execute phase or the commit phase. Different transactions may run in the same coordinator thread, or on a different coordinator thread on the same server, or on a different server. The transaction protocol may obtain a read set of the transaction, produce a tentative write set for the transaction, and commit the transaction once it is ready. A read set is a set of software objects read by the transaction and a write set is a set of software objects tentatively written by the transaction. In some examples the commit phase executes a two phase protocol with two round-trip RPCs (lock and validate) and one one-way RPC (update and unlock). “RPC” denotes a remote procedure call.

[0035] In an example, the transaction protocol is designed for use with one-sided operations. An example of a suitable transaction protocol is given in FIG. 4 but this is one example only and others are possible. In the example of FIG. 4 the transaction is able to read objects using one-sided reads and, in the commit phase, to use one-sided writes. The transaction is able to use validation of read objects using one-sided reads to avoid locking read objects. The transaction protocol may add bits to data objects that enable the detection of concurrent reads. This mechanism enables serializability except in the presence of configuration change. By using the new recovery process, serializability is achieved even when configuration change occurs.

[0036] Transactions use optimistic concurrency control with version numbers in objects. They are executed by a coordinator, which is a thread on any of the servers that executes applications’ transaction code. Coordinators are not replicated and communicate directly with primaries and backups, further reducing latency and message counts.

[0037] Figure 4 shows a timeline for an example transaction protocol. During the execution phase, the coordinator issues reads for local or remote objects, and records their versions. It buffers writes until the transaction tries to commit. Remote objects are read by issuing one-sided RDMA reads to object’s primaries. During the commit phase it uses $Pw(f + 3)$ one-sided RDMA writes where Pw is the number of primaries with objects that are written by the transaction, and Pr onesided RDMA reads where Pr is the number of objects that are read but not written. It is also possible to validate the objects via remote procedure call to the primary, using two one-sided RDMA writes and CPU involvement at the primary. This gives potential efficiencies in the case of primaries that are only read by the transaction and have a large number of objects read by the transaction.

[0038] The commit protocol uses RDMA writes into ring buffers that are stored in non-volatile DRAM at primaries and backups. Conceptually, each ring buffer is written only by the coordinator thread running the transaction and read asynchronously by the peer thread at the primary or backup (the implementation transparently multiplexes ring buffers across threads). Each ring buffer serves both as a message channel and as a transaction log. No other logs are required. When the transaction attempts to commit, the coordinator writes LOCK records to the logs at primaries of objects that were modified by the transaction. These records contain the object versions and the new state for written objects stored on the machine, and the list of all regions containing objects written by the transaction. The primaries use a compare-and-swap on the object version to atomically check if the object versions read by the coordinator are still current and lock the version if they are. They reply to the coordinator by writing an acknowledgement to a ring buffer at the coordinator that indicates whether the transaction should abort or continue.

[0039] Unlike a traditional two-phase commit protocol, readonly participants, which are primaries that store objects read by the transaction but do not store any modified objects, are not involved in this phase of the protocol. Instead, the protocol of FIG. 4 uses an additional read validation phase. The coordinator issues RDMA reads to read the current versions of the objects read by the transaction (these are the Pr one-sided reads referred to above, one per object). If they match the versions originally read by the transaction, they are guaranteed to be current at the serialization point for the transaction, which is the point where all the locks are acquired. If they do not, the coordinator aborts the transaction. This is an important optimization because it ensures that transactions do not consume any CPU time at read-only participants and it eliminates both the need to block on read locks and the need to send messages to lock and unlock them. Latency does increase by two message delays because of it, but this is a good trade-off: RDMA reads take less than 3 microseconds in an unloaded system and the reduction in CPU time offsets the increased delay in a loaded system.

[0040] The overhead incurred due to replication is minimal. The coordinator issues a single RDMA write of a COMMIT-BACKUP record to the non-volatile logs at the backups and then waits for an acknowledgement from the NIC hardware without interrupting the backup's CPU. The COMMIT-BACKUP log record has the same content as LOCK records.

[0041] After receiving the acknowledgements (acks) from all backups, the coordinator writes a COMMIT-PRIMARY record to the non-volatile logs at the primaries. Primaries process COMMIT-PRIMARY records by updating the objects in place, incrementing their versions, and unlocking them, which exposes the writes committed by the transaction. Waiting for acks from the backups before sending COMMIT-PRIMARY is necessary to ensure strict serializability across failures. Otherwise, a primary and the coordinator could fail after some transaction reads the modifications, causing later transactions to miss the modifications.

[0042] If the coordinator is not a primary, it also waits for a hardware ack from at least one primary's NIC before acknowledging to the client. Otherwise the transaction may still abort if all the backups fail causing the system to lose information about the success of read validation.

[0043] Backups and primaries keep the records in their logs until they are truncated. The coordinator truncates logs at primaries and backups lazily after receiving acks from all primaries. It does this by piggybacking identifiers of truncated transactions in other protocol messages. The backups apply the updates to their copies of the objects at truncation time. In traditional two-phase commit protocols, participants reserve the resources to commit a transaction during the prepare phase and refuse to prepare the transaction if they do not have enough resources. A different technique is used in the method of FIG. 4 because backups' CPUs are not involved in the commit protocol. Instead, coordinators reserve space for all commit protocol records including truncate records in primary and backup logs before starting the commit protocol. The reservation is a local operation on the coordinator, as the coordinator is the owner of the log on the remote machine. The reservation is released when the corresponding record is written. Truncation record reservations are also released if the truncation is piggybacked on another message. If the log becomes full, the coordinator uses the reservations to send explicit truncate messages to free up space in the log. This happens rarely but is needed to ensure liveness.

[0044] FIG. 5 is a flow diagram of a recovery process in a cluster of servers at data center. A potential failure, or other type of configuration change is automatically detected 500. For example, using the lease scheme as mentioned above or other methods to monitor behavior of the servers through peer to peer communications between the nodes. The results of the configuration change detection process are not always reliable

since server failures cannot be detected reliably. Consequently, a server *s* that is falsely suspected to be faulty might keep serving read requests even when a majority of servers has declared *s* crashed. Once *s* is declared crashed the recovery process works to reassign its objects to other servers, in the interest of progress. If *s* has not really crashed, the recovery process works to prevent it from serving read requests with object values that are not consistent with the state of the system. More detail about how this is achieved is given below.

[0045] Once a potential configuration change has been detected a configuration protocol is invoked as mentioned above. This results in a new configuration being agreed 502. In some examples the configuration protocol implements precise membership. Precise membership means that all servers know which servers are in the configuration, and they also know that the other servers know that. In this case step 504 of FIG. 5 is carried out to synchronize the servers of the new configuration so they all know which servers are in the new configuration. This prevents a server which does not yet know about the new configuration from observing an inconsistent state by using one-sided requests. For example, it might access a piece of memory on another server *sa* that contained object *o1* in previous configuration, but contains *o2* in the new one. Because the processor on *sa* does not process the RDMA request, it cannot detect that the request is coming from an old configuration. If the server initiating the request does not know about the new configuration, it cannot reject the result either, so it ends up observing an inconsistent state. By making sure that all servers in the configuration are aware of it, a server can be arranged not to issue a lock-free read request to a server in a previous configuration and to reject replies from servers no longer in the configuration. This is now described in more detail with reference to FIG. 6.

[0046] A server in the data center cluster receives a new configuration 600, for example, it receives a message as part of the recovery process which comprises a list of identifiers of servers which are in the new configuration. The server receives 602 a lock-free read reply message (for example, at the network interface card 206). The lock-free read reply message comprises details of the ID of the sending server. The recovery process at one of the server's CPUs 200 checks 604 whether the ID of the sending server is present in the list of identifiers of servers which are in the new configuration. If so it accepts 608 the read reply. If not, the CPU makes more checks 606 in some examples. For example, it checks if the primary for the region being read changed since the read was

issued. Under some circumstances the further checks 606 succeed and the CPU allows the read reply to be accepted 608. If the additional checks 606 fail the read reply is refused 610.

[0047] Returning to FIG. 5, the recovery process may use an alternative process at step 502 for agreeing the new configuration. For example, one-sided operations are not used. When a configuration change occurs this is detected using leases as mentioned above. The servers are not required to immediately agree a new configuration and are able to check if they hold the lease for the new configuration before replying to a read request.

[0048] The recovery process proceeds to commit initial state 506 of transactions (these are either transactions which have been forced to restart as a result of the new configuration, or transactions which have been allowed to continue despite the new configuration). This involves operating the recovery process so that participants of a configuration start a new configuration with consistent state and eventually all regions replicated. The data of the regions does not need to be replicated for transactions to proceed. The data recovery may happen after the transaction state is recovered, and may be done concurrently with resumed execution of the system as described in more detail below. During transaction recovery the information about the transaction log records affecting each region is replicated. During this stage servers acquire new regions and release old regions. Depending on the sequence of server crashes and data replication, it is possible that data is lost if not treated properly. More detail about the process of step 506 is given with reference to FIG. 7 below.

[0049] Once participants of a new configuration have consistent state, transactions are served 508 using the transaction protocol and lock-free reads may go ahead while region state is recovered.

[0050] FIG. 7 is a flow diagram of a process of transaction state recovery in a cluster of servers at a data center. This process comprises recovering 700 locks at backups which have been promoted to primaries, recovering logs 702 at backups, recovering co-ordinator state 704 and deciding 706 whether recovered transactions commit or abort. The process of FIG. 7 may operate in parallel with transactions that start in the new configuration.

[0051] FIG. 8 is a flow diagram of a process at a server of a cluster at a data center as part of a recovery process. The flow diagram of FIG. 8 includes part of the

process of steps 700 and 702 of FIG. 7 from the point of view of a single server of the data center cluster.

[0052] Servers attempt to write the new configuration to the agreement service at the agreement nodes 106. The one that wins becomes the CM and sends a NEW-CONFIG message to all the other servers in the configuration and conceptually to itself. The server receives a new configuration 800. When a machine receives a NEW-CONFIG with a configuration identifier i that is greater than its own, it updates its current configuration identifier and its cached copy of the region mappings, and allocates space to hold any new region replicas assigned to it. From this point, it does not issue new requests to machines that are not in the configuration and it rejects read responses and write acks from those machines. It also uses the process of FIG. 6 with respect to lock-free read requests.

[0053] The server replies to the CM with a NEW-CONFIG-ACK message (not shown in FIG. 8). Once the CM receives NEW-CONFIG-ACK messages from all machines in configuration i , it replies with a NEW-CONFIG-COMMIT message to all of them. This implements the precise membership process mentioned above.

[0054] Then the servers drain their logs 802, that is, they process all complete log records currently in the logs.

[0055] Messages are tagged with the current configuration and receivers reject any messages from previous configurations. This check is insufficient for one-sided RDMA writes into transaction logs because the coordinator acts on acks sent by the network hardware whereas primaries and backups process COMMIT-BACKUP and COMMIT-PRIMARY records asynchronously. To address this problem, transactions have unique identifiers $\langle c, m, t, l \rangle$ assigned at the start of commit that encode a thread-local unique ID l , the thread ID t of the coordinator, the machine ID m of the coordinator, and the configuration c in which the commit started. Primaries and backups cannot simply reject log records for transactions tagged with an old configuration because other primaries may already have committed the transactions and exposed their updates. Therefore, primaries and backups drain their logs by processing all log records when they receive a NEW-CONFIG-COMMIT message and they record the configuration identifier in a variable LastDrained when they are done. Draining the logs is shown in step 802 of FIG. 8. Correctness is ensured by rejecting log records for transactions with configuration identifiers less than or equal to LastDrained.

[0056] A log record is guaranteed to be processed before or during draining if the coordinator received the network hardware ack in the configuration when it started the commit protocol, and the coordinator rejects the ack if it is received in a later configuration.

5 **[0057]** Some transactions whose commit phase spans configuration changes wait for the recovery protocol to decide if they commit or abort. These are transactions that are affected by the configuration change, i.e., either their coordinator failed, the primaries for objects in their read set changed, or any replica for objects in their write set changed. Transactions that are not affected by the configuration change are able to execute the normal commit protocol whereas other transactions wait for the recovery protocol to
10 decide their outcome. All machines agree on whether a transaction is allowed to execute without waiting for the recovery process. The CM guarantees agreement by reading the value of the LastDrained variable at each machine as part of the probe read, a one-sided read that machines attempting to become CM issue to each of the machines in the
15 configuration. For each region r whose mapping has changed since LastDrained, the CM sends two configuration IDs in the NEW-CONFIG message to that machine. These are LastPrimaryChange[r], the last configuration ID when the primary of r changed, and LastReplicaChange[r], the last configuration ID when any replica of r changed. A transaction that started committing in configuration $i-1$ can proceed in configuration i if:
20 for all regions r containing objects modified by the transaction LastReplicaChange[r] $< i$, and for all regions r_0 containing objects read by the transaction LastPrimaryChange[r_0] $< i$, and the coordinator has not been removed from configuration i . Coordinators use this information to decide whether to accept acks for RDMA writes to transaction logs and replicas use it to decide whether to process log records.

25 **[0058]** The server identifies transactions that are affected by the configuration change and the regions modified by those transactions 804. For affected regions where the server is backup, it sends a NEED-RECOVERY message to primaries of those regions 806.

[0059] For each affected region where the server is primary 808 the server
30 prevents 810 reading or modifying of objects in the region until it is done recovering locks for that region. It waits for NEED-RECOVERY messages from backups of the region 812 and recovers transaction state 814 which tells it which objects need to be locked. The server locks 816 objects in the region which are affected by transactions.

[0060] The server then activates the region 818 so that reading or modifying of objects in the region can go ahead (except where those objects are locked).

[0061] The log recover process of FIG. 7, 702 comprises, for each affected region where the server is primary 820, replicating log records to backups if they do not already have them.

[0062] When all primary regions at the server are active, the server sends 822 a REGIONS-ACTIVE message to the configuration manager (CM). When the server receives an ALL-REGIONS-ACTIVE message 824 back from the CM it is triggered to carry out steps 826 and 828 which may occur in parallel. Step 826 comprises beginning allocator recovery. Step 828 comprises data recovery for new backups which involves replicating regions at new backups from their primary counterparts. Data recovery for new backups may proceed in parallel with foreground operation of transaction execution.

[0063] More detail about data recovery for new backups is now given:

[0064] A new backup for a region initially has a freshly allocated and zeroed local region replica. It divides the region across worker threads that recover it in parallel. Each thread issues one-sided operations to read a portion of data at a time from the primary. The portion size is large enough to use the network efficiently but small enough not to impact normal case operation. To reduce impact on foreground performance, recovery is paced by scheduling the next read to start at a random point within an interval after the start of the previous read. Each recovered object is examined before being copied to the backup. If the object was not locked at the primary and it has a version greater than the local version, the backup locks the local version with a compare-and-swap, updates the object state, and unlocks it. Otherwise, the object has been or is being updated by a transaction that created a version greater than or equal to the one recovered, and the recovered state is not applied.

[0065] Allocator recovery 826 is now described.

[0066] An allocator is a component on each machine. It tracks which memory areas within each region are in use (allocated) and which are free and can be returned by an object allocation request within a transaction. This state is recovered to avoid double-allocating the same memory and to avoid leaking a free object by incorrectly marking it allocated. The allocator is arranged to split the region into blocks that are used as SLABs for allocating small objects. It keeps two pieces of meta-data: block headers, which contain the object size, and SLAB free lists. Block headers are replicated to backups when a new block is allocated, meaning that they are available on the new primary after a

failure. Block headers of regions being recovered are re-replicated from the new primary to backups immediately after receiving NEW-CONFIG- COMMIT, as they contain object size information used in data recovery. The primary sends block headers to all backups to avoid inconsistencies that could arise if the old primary failed while replicating the block header.

[0067] The SLAB free lists are kept only in the memory of the primary to reduce overheads of object allocation. Each object has a bit in its header that is set by an allocation and cleared by a free at transaction commit to enable recovering SLAB free lists. The lists are recovered by scanning the objects in the region, which is done in parallel by all threads on the machine. To minimize the impact on transaction lock recovery, allocation recovery starts after ALL-REGIONS-ACTIVE is received.

[0068] FIG. 9 is a flow diagram of a process at a server of a data center cluster which is a coordinator of transaction A and of a process at a server of the data center cluster which is a primary of a region participating in transaction A. These flow diagrams explain how coordinator state is recovered and how committing or aborting a transaction occurs as part of the recovery process.

[0069] Since the coordinator (i.e. the thread coordinating transaction A) is not replicated, the recovery process recovers the coordinator state for a transaction from the logs of the primaries of regions modified by the transaction. Consistent hashing is used to determine the coordinator for a transaction, making sure to that the coordinator remains the same if it is in the new configuration.

[0070] The server which is coordinator of transaction A receives votes 902 from primaries participating in the transaction A. It may also send requests 904 to those primaries for votes. The votes from the primaries are of different types according to log data at the primaries and/or types of messages received at the primaries.

[0071] Depending on the votes it receives the server decides whether to commit or abort the transaction 906. It informs the primaries of its decision by sending COMMIT-RECOVERY or ABORT-RECOVERY messages 908 to those primaries. It waits 910 for acknowledgements from the primaries and sends TRUNCATE-RECOVERY message to those primaries 912.

[0072] A primary which is participating in transaction A 914 carries out the method in the lower part of FIG. 9. For each region the primary has which is affected by

the transaction, it sends a recovery vote message 916 to the coordinator. If a request is received for a vote, it sends the vote after re-replicating log records for the transaction 918.

[0073] More detail about the vote messages is now given. The threads in the primary send RECOVERY-VOTE messages to their peer threads in the coordinator for each recovering transaction that modified the region. The vote is *commit-primary* if any replica has a COMMIT-PRIMARY record in the log or received a COMMIT-RECOVERY message; *abort* if any replica has an ABORT record in the log or received an ABORT-RECOVERY message; *commit-backup* if any replica has a COMMIT-BACKUP record in the log and did not receive an ABORT-RECOVERY message; or *lock* if any replica has a LOCK record in the log and did not receive an ABORT-RECOVERY message. The vote message also includes the configuration identifier, the region identifier, the transaction identifier, and the list of region identifiers modified by the transaction.

[0074] The coordinator waits for RECOVERY-VOTE messages from each primary of a region modified by the transaction. Because some primaries may have never received a log record for the transaction or may already have truncated the log records for the transaction, the coordinator sends explicit requests to vote for a transaction to the primaries of the regions modified by a transaction if it does not receive the votes after a timeout. The REQUEST-VOTE messages include the configuration identifier, the region identifier, and the transaction identifier. If a machine that has the log records for a transaction receives this message, it sends the vote after re-replicating the log records. In this way the recovery process avoids distinguishing a machine that never received a log record for the transaction from a machine that already truncated the log records. In the first case the transaction should be aborted and in the second it may already have committed. To do this, each thread maintains a truncation data structure (referred to as *truncated_ids*) to determine if a transaction has been truncated given its identifier. The transaction data structure may be kept compact by using a lower bound on non-truncated transaction identifiers, which is piggybacked on coordinator messages. If the transaction has been truncated, the thread replies with a vote of truncated. Otherwise, it replies with a vote of unknown.

[0075] More detail about the process of committing/aborting is now given.

[0076] Once the coordinator receives RECOVERY-VOTE messages from each primary of a region modified by the transaction it decides to commit if it received any *commit-primary* vote or if it received a *commit-backup* vote for one region and *lock*,

commit-backup, or *truncate* votes for all the other regions. Otherwise it decides to abort. It then sends back COMMIT-RECOVERY or ABORT-RECOVERY to all participant replicas depending on the decision. Both messages include the configuration identifier and the transaction identifier. Threads that receive these messages treat COMMIT-

5 RECOVERY similarly to a COMMIT-PRIMARY log record if they are the primary and similarly to a COMMIT-BACKUP log record if they are a backup. They treat receiving ABORT-RECOVERY as processing an ABORT record. The coordinator waits for acks from all participant replicas for the decision messages, and then sends a TRUNCATE-RECOVERY message to truncate the state for the transaction.

10 **[0077]** A primary commits and unlocks the objects modified by a transaction only after it processes a COMMIT-PRIMARY log record or receives a COMMIT-RECOVERY message. This ensures that transactions whose writes have been exposed to other transactions are guaranteed to eventually commit at all participants. When a transaction commits, all participant replicas have either LOCK, COMMIT-BACKUP, or COMMIT-
15 PRIMARY log records or they have already truncated the transaction's records. Therefore, if the transaction is recovered the decision will be to commit. The same is true for transactions that the coordinator signaled committed to the client because this is done after receiving an ack from the NIC indicating that the COMMIT-PRIMARY record was written to the log, or an ack message that the COMMIT-RECOVERY message was
20 received.

[0078] More detail about an example configuration manager 314 is now given.

[0079] The distributed memory is organized into regions of configurable size that are typically 2 GB. Regions are the unit of allocation and failure recovery for the cluster.

[0080] Allocation and mapping of regions to replicas is handled by a configuration
25 manager (CM) running on one of the nodes. With 2GB regions, we expect up to 250 regions on a typical server, hence a single CM can handle allocation for thousands of machines. This has several advantages over using consistent hashing. It lets us satisfy fault-tolerance constraints such as always placing replicas in different failure domains (e.g. racks). It makes it easier to balance load across servers and to operate close to capacity.

30 **[0081]** It is also easier to satisfy requests to allocate objects in the same locality group and to keep them in the same locality group when remapping regions after a failure. Finally, remapping on failure can be done such that data movement is minimized.

[0082] To allocate a new region, a node contacts the CM, which assigns a region identifier from a monotonically increasing counter and then selects replicas for the region attempting to balance the number of regions assigned to each machine subject to capacity and (any) locality group constraints. The CM uses a two phase commit protocol to allocate regions to ensure a mapping is replicated at all the replicas for the region before it is used. It starts by sending a prepare message to the selected replicas with the region identifier. If they all succeed in allocating the region, it sends a commit message for the region. Nodes fetch missing region mappings from the CM when needed. Threads cache the mappings, together with RDMA references needed to access the primary, and are thus able to read an object with a single one-sided RDMA read request in the common case.

[0083] The CM also handles failure detection via lease management and reconfiguration after a failure. Co-locating the CM on a node allows the CM to use RDMA-based protocols. The agreement service is used to store only the current configuration $\langle I, S, CM \rangle$ where i is a unique, monotonically increasing 64-bit configuration identifier, S is the set of nodes in the configuration, and CM is the configuration manager. The agreement node provides a service which is invoked only on a configuration change to update the membership of the configuration atomically.

[0084] More detail about an example lease scheme is now given.

[0085] Leases may be used to detect failures. Each non-CM node holds a lease at the CM and the CM holds a lease at every other node. Expiry of any lease triggers failure recovery. Leases are granted via a 3-way handshake. Each machine sends a lease request to the CM and it responds with a message that acts as both a lease grant and a lease request of its own.

[0086] Finally, machines reply with a lease grant. By default, lease renewal is attempted every 1/4 of the lease expiry period as the leases are extremely short which is key to high availability. Under heavy load, the data center cluster can maintain sub-10 ms leases for a 60-machine cluster with no false positives. Significantly larger clusters may require a two-level hierarchy, which in the worst case would double failure detection time.

[0087] Achieving short leases under load required several optimizations. Lease messages use RDMA, but because RDMA queues are shared, they can be delayed due to lock contention and deep queuing behind other messages on the NIC. Dedicating queue pairs to lease management is suboptimal as it drastically increases the number of queue pairs at the CM, degrading NIC performance. Instead, the protocol sends lease messages

using the connectionless unreliable datagram mode of RDMA, with receiver-side polling for low latency.

[0088] Lease renewal must also be scheduled on the CPU in a timely way. The recovery process uses a dedicated lease manager thread that may run at the highest user-space priority.

[0089] More detail about an example configuration change protocol is now given

[0090] When a machine is suspected of failing because of lease expiry, a reconfiguration is initiated by one of the machines which we will call CMnew, which may be the same as the current CM. CMnew first issues an RDMA read to all the machines in the system except the machine that is suspected.

[0091] These read probes (that were mentioned before) detect correlated failures that affect several machines, e.g., power and switch failures, which are all handled by a single reconfiguration. CMnew continues with the reconfiguration if it obtains responses for a majority of the probes. This ensures that if the network is partitioned, the CM will not be in the smaller partition.

[0092] After receiving replies to the read probes, CMnew attempts to update the configuration data stored in Agreement service (comprising nodes 106) from $\langle I, S, CMold \rangle$ to $\langle i + 1, S0, CMnew \rangle$, where $S0$ is the set of machines that replied to probes. The agreement service implements an atomic compare-and-swap that succeeds if the current configuration identifier is still i . For example, the agreement service may be implemented using znode sequence numbers to implement the atomic compare-and-swap. This ensures that only one node can become the CM for configuration $i+1$, as there might be multiple configuration change attempts in parallel.

[0093] A non-CM node initiates reconfiguration only if it suspects that the CM failed. It first asks one of a small number of “backup CMs” to make the change (the k successors of the CM using consistent hashing). If the configuration is unchanged after a timeout period then it attempts the reconfiguration itself. This design avoids a large number of simultaneous requests to the agreement node when the CM fails. The new CM then reassigns regions previously mapped to failed machines to restore the number of replicas to $f + 1$ (to be able to tolerate f failures in each replica group). It tries to balance load and keep locality subject to space constraints. For failed primaries it always promotes a surviving backup to the new primary to reduce the time to recover. If it detects regions that lost all their replicas or there is no space to re-replicate regions, it signals an error.

After remapping regions, the CM sends a NEW-CONFIG message to all the machines in the configuration with the configuration identifier, its own identifier, the identifiers of the other machines in the configuration, and all the new mappings of regions to machines.

NEW-CONFIG also resets the lease protocol if the CM has changed: it grants a lease to all the nodes, which reply granting a lease to the new CM, and all nodes then discard leases from the previous configuration. If the CM is unchanged, lease exchange continues during reconfiguration. External client requests to the CM are blocked between suspecting a failure and NEW-CONFIG-COMMIT, and at other nodes between NEW-CONFIG and NEW-CONFIG-COMMIT.

[0094] FIG. 10 illustrates various components of an exemplary computing-based device 1000 which may be implemented as any form of a computing and/or electronic device, and in which embodiments of a server of a data center cluster may be implemented. The server may be connected via an interconnect 1024 to a cluster of other servers 1026.

[0095] Computing-based device 1000 comprises one or more processors 1002 which may be microprocessors, controllers or any other suitable type of processors for processing computer executable instructions to control the operation of the device in order to recover from configuration changes in a cluster of servers at a data center such that transactions executing in the data center cluster retain ACID properties and may also retain serializability. In some examples, for example where a system on a chip architecture is used, the processors 1002 may include one or more fixed function blocks (also referred to as accelerators) which implement a part of the method of any of FIGs. 4 to 9 in hardware (rather than software or firmware). A part of the method of any of FIGs. 4 to 9 may also be implemented in hardware or software in the network interface 1014.

Platform software comprising an operating system 1004 or any other suitable platform software may be provided at the computing-based device to enable protocols 1008 to be executed on the device such as recovery processes, configuration protocols and transaction protocols.

[0096] The computer executable instructions may be provided using any computer-readable media that is accessible by computing based device 1000. Computer-readable media may include, for example, computer storage media such as memory 1012 and communications media. Computer storage media, such as memory 1012, includes volatile and non-volatile, 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, EPROM, DRAM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other non-transmission medium that can be used to store information for access by a computing device. In contrast, communication media may embody computer readable instructions, data structures, program modules, or other data in a modulated data signal, such as a carrier wave, or other transport mechanism. As defined herein, computer storage media does not include communication media. Therefore, a computer storage medium should not be interpreted to be a propagating signal per se. Propagated signals may be present in a computer storage media, but propagated signals per se are not examples of computer storage media. Although the computer storage media (memory 1012) is shown within the computing-based device 1000 it will be appreciated that the storage may be distributed or located remotely and accessed via a network or other communication link (e.g. using network interface 1014). In some cases the computing based device 1000 comprises a battery and a solid state drive (SSD) and DRAM as described with reference to FIG. 2.

[0097] The computing-based device 1000 also comprises an input/output controller 1016 arranged to output display information to a display device 1018 which may be separate from or integral to the computing-based device 1000. The display information may provide a graphical user interface. The input/output controller 1016 is also arranged to receive and process input from one or more devices, such as a user input device 1020 (e.g. a mouse, keyboard, camera, microphone or other sensor). In some examples the user input device 1020 may detect voice input, user gestures or other user actions and may provide a natural user interface (NUI). This user input may be used to view transaction logs, control parameters of recovery processes or for other purposes. In an embodiment the display device 1018 may also act as the user input device 1020 if it is a touch sensitive display device. The input/output controller 1016 may also output data to devices other than the display device, e.g. a locally connected printing device.

[0098] In an example there is a server at a cluster of servers in a data center, the server comprising:

[0099] a memory being part of a distributed memory of the cluster;

[00100] at least one processor executing transactions and lock-free reads on software objects stored in regions of the distributed memory, the software objects and details of the transactions being replicated in the distributed memory,

5 [00101] the processor arranged to receive a message indicating a new configuration of the cluster comprising addition, removal or potential failure of at least one of the other servers;

[00102] the processor configured to use a recovery process which enables the lock-free reads and committed ones of the transactions to retain the properties of atomicity, consistency, isolation and durability across configuration changes.

10 [00103] For example the processor is configured to use a recovery process which enables serializability of the committed transactions across configuration changes.

[00104] For example the processor is configured such that replies to one-sided read requests received at the server are refused at least on the basis of a configuration identifier sent with the read reply by a sending server.

15 [00105] For example the processor is arranged to receive the message comprising details of identities of all other servers in the new configuration and knowledge that all servers in the new configuration know the identities of all servers in the new configuration, such that one-sided operations on network interface cards of the servers may be performed whilst retaining the property of consistency of the lock-free reads and
20 committed ones of the transactions.

[00106] In some examples the memory comprises non-volatile memory implemented using any of: dynamic random access memory together with a battery and a persistent store to which the contents of the DRAM are written using the battery power if the server crashes, non-volatile dual in-line memory modules, physical memory.

25 [00107] In an example the network interface card is configured to support leases with a timer, and to reject requests that arrive after a lease timer has expired without being renewed.

[00108] In an example there is a method performed at a server in a cluster of servers in a data center, the method comprising:

30 [00109] executing, at a processor, transactions and lock-free reads on software objects stored in regions of a distributed memory of the cluster of servers, the software objects and logs of the transactions being replicated in the distributed memory,

[00110] receiving a message indicating a new configuration of the cluster comprising addition, removal or potential failure of at least one of the servers;

5 [00111] using a recovery process which enables the lock-free reads and committed ones of the transactions to retain the properties of atomicity, consistency, isolation and durability across configuration changes.

[00112] For example, the recovery process comprises a lock recovery process to recover write locks of transactions already being committed at the time of the new configuration, the lock recovery process operating to recover the write locks from the replicated transaction logs, which may have been written with one-sided writes.

10 [00113] For example, the recovery process comprises blocking transactions that are affected by the new configuration until the lock recovery process completes.

[00114] For example, the message indicating the new configuration is received from one of the servers in the cluster which is a configuration manager, the configuration manager having computed a set of the servers to be in the new configuration using at least
15 one agreement node in communication with the cluster, the agreement node configured to decide which of the servers in the cluster is the configuration manager, the configuration manager having sent the computed set of the servers to the servers in the new configuration.

[00115] For example, the recovery process comprises a voting process whereby
20 servers which participate in a transaction affected by the new configuration submit votes to a server acting as coordinator of the transaction, the votes being used to decide whether to commit or abort the transaction.

[00116] For example the votes are determined according to characteristics of logs at the servers and messages received at the servers.

25 [00117] For example the recovery process comprises a process for recovering data to be stored in regions which are new back up regions formed as a result of the new configuration, the process comprising one-sided reads from primary regions holding data to be replicated in the new back up regions.

[00118] For example the recovery process comprises an allocation recovery
30 process, to recover the allocator's free lists used for allocating objects, which occurs after the lock recovery process.

[00119] For example the allocation recovery process comprises scanning object headers of objects stored in regions at a server.

[00120] In an example the recovery process is configured to enable transactions to resume reading objects and committing after the lock recovery process.

[00121] In an example the recovery process comprises reserving space for log records to ensure enough resources are available to commit transactions.

5 [00122] In an example the recovery process allows transactions that are not affected by the configuration change to read objects and commit concurrently with recovery process

[00123] In an example the recovery process enables strict serializability of the committed transactions across configuration changes.

10 [00124] In an example a transaction has a unique identifier.

[00125] In an example a data center cluster comprises:

[00126] a plurality of servers connected via a network;

[00127] a distributed memory comprising local memory at individual ones of the servers;

15 [00128] the servers configured to execute transactions and lock-free reads on software objects stored in regions of the distributed memory, the software objects and details of the transactions being replicated in the distributed memory,

[00129] the servers configured to use a recovery process which enables the lock-free reads and committed ones of the transactions to retain the properties of atomicity, consistency, isolation and durability across configuration changes comprising addition, removal or potential failure of at least one of the other servers.

20 [00130] The term 'computer' or 'computing-based device' is used herein to refer to any device with processing capability such that it can execute instructions. Those skilled in the art will realize that such processing capabilities are incorporated into many different devices and therefore the terms 'computer' and 'computing-based device' each include PCs, servers, mobile telephones (including smart phones), tablet computers, set-top boxes, media players, games consoles, personal digital assistants and many other devices.

25 [00131] The methods described herein may be performed by software in machine readable form on a tangible storage medium e.g. in the form of a computer program comprising computer program code means adapted to perform all the steps of any of the methods described herein when the program is run on a computer and where the computer program may be embodied on a computer readable medium. Examples of tangible storage media include computer storage devices comprising computer-readable media such as

disks, thumb drives, memory etc. and do not include propagated signals. Propagated signals may be present in a tangible storage media, but propagated signals per se are not examples of tangible storage media. The software can be suitable for execution on a parallel processor or a serial processor such that the method steps may be carried out in any suitable order, or simultaneously.

[00132] This acknowledges that software can be a valuable, separately tradable commodity. It is intended to encompass software, which runs on or controls “dumb” or standard hardware, to carry out the desired functions. It is also intended to encompass software which “describes” or defines the configuration of hardware, such as HDL (hardware description language) software, as is used for designing silicon chips, or for configuring universal programmable chips, to carry out desired functions.

[00133] Those skilled in the art will realize that storage devices utilized to store program instructions can be distributed across a network. For example, a remote computer may store an example of the process described as software. A local or terminal computer may access the remote computer and download a part or all of the software to run the program. Alternatively, the local computer may download pieces of the software as needed, or execute some software instructions at the local terminal and some at the remote computer (or computer network). Those skilled in the art will also realize that by utilizing conventional techniques known to those skilled in the art that all, or a portion of the software instructions may be carried out by a dedicated circuit, such as a DSP, programmable logic array, or the like.

[00134] Any range or device value given herein may be extended or altered without losing the effect sought, as will be apparent to the skilled person.

[00135] Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

[00136] It will be understood that the benefits and advantages described above may relate to one embodiment or may relate to several embodiments. The embodiments are not limited to those that solve any or all of the stated problems or those that have any or all of the stated benefits and advantages. It will further be understood that reference to 'an' item refers to one or more of those items.

[00137] The steps of the methods described herein may be carried out in any suitable order, or simultaneously where appropriate. Additionally, individual blocks may be deleted from any of the methods without departing from the spirit and scope of the subject matter described herein. Aspects of any of the examples described above may be
5 combined with aspects of any of the other examples described to form further examples without losing the effect sought.

[00138] The term 'comprising' is used herein to mean including the method blocks or elements identified, but that such blocks or elements do not comprise an exclusive list and a method or apparatus may contain additional blocks or elements.

10 [00139] The term 'subset' is used herein to refer to a proper subset such that a subset of a set does not comprise all the elements of the set (i.e. at least one of the elements of the set is missing from the subset).

[00140] It will be understood that the above description is given by way of example only and that various modifications may be made by those skilled in the art. The above
15 specification, examples and data provide a complete description of the structure and use of exemplary embodiments. Although various embodiments have been described above with a certain degree of particularity, or with reference to one or more individual embodiments, those skilled in the art could make numerous alterations to the disclosed embodiments without departing from the spirit or scope of this specification.

CLAIMS

1. A server at a cluster of servers in a data center, the server comprising:
a memory being part of a distributed memory of the cluster;
at least one processor executing transactions and lock-free reads on software objects stored in regions of the distributed memory, the software objects and details of the transactions being replicated in the distributed memory,
the processor arranged to receive a message indicating a new configuration of the cluster comprising addition, removal or potential failure of at least one of the other servers;
the processor configured to use a recovery process which enables the lock-free reads and committed ones of the transactions to retain the properties of atomicity, consistency, isolation and durability across configuration changes.
2. The server of claim 1 wherein the processor is configured such that replies to one-sided read requests received at the server are refused at least on the basis of a configuration identifier associated with the request.
3. The server of claim 1 or claim 2 wherein the processor is arranged to receive the message comprising details of identities of all other servers in the new configuration and knowledge that all servers in the new configuration know the identities of all servers in the new configuration, such that one-sided operations on network interface cards of the servers may be performed whilst retaining the property of consistency of the lock-free reads and committed ones of the transactions.
4. The server of claim 1 having a network interface card configured to support leases with a timer, and to reject requests that arrive after a lease timer has expired without being renewed.
5. A method performed at a server in a cluster of servers in a data center, the method comprising:
executing, at a processor, transactions and lock-free reads on software objects stored in regions of a distributed memory of the cluster of servers, the software objects and logs of the transactions being replicated in the distributed memory,
receiving a message indicating a new configuration of the cluster comprising addition, removal or potential failure of at least one of the servers;
using a recovery process which enables the lock-free reads and committed ones of the transactions to retain the properties of atomicity, consistency, isolation and durability across configuration changes.

6. The method of claim 5 wherein the recovery process comprises a lock recovery process to recover write locks of transactions already being committed at the time of the new configuration, the lock recovery process operating to recover the write locks from the replicated transaction logs, which may have been written with one-sided writes.
7. The method of claim 6 wherein the recovery process comprises blocking transactions that are affected by the new configuration until the lock recovery process completes.
8. The method of any of claims 5 to 7 wherein the message indicating the new configuration is received from one of the servers in the cluster which is a configuration manager, the configuration manager having computed a set of the servers to be in the new configuration using at least one agreement node in communication with the cluster, the agreement node configured to decide which of the servers in the cluster is the configuration manager, the configuration manager having sent the computed set of the servers to the servers in the new configuration.
9. The method of any of claims 5 to 8 wherein the recovery process comprises a voting process whereby servers which participate in a transaction affected by the new configuration submit votes to a server acting as coordinator of the transaction, the votes being used to decide whether to commit or abort the transaction.
10. The method of claim 9 wherein the votes are determined according to characteristics of logs at the servers and messages received at the servers.
11. The method of any of claims 5 to 10 wherein the recovery process comprises a process for recovering data to be stored in regions which are new back up regions formed as a result of the new configuration, the process comprising one-sided reads from primary regions holding data to be replicated in the new back up regions.
12. The method of claim 5 wherein the recovery process comprises an allocation recovery process, to recover an allocator's free lists used for allocating objects, which occurs after the lock recovery process.
13. The method of claim 12 wherein the allocation recovery process comprises scanning object headers of objects stored in regions at a server.
14. The method of any of claims 5 to 13 wherein the recovery process comprises reserving space for log records to ensure enough resources are available to commit transactions.

15. The method of any of claims 5 to 14 wherein the recovery process allows transactions that are not affected by the configuration change to read objects and commit concurrently with recovery process

1/10




FIG. 1




FIG. 2

3/10




FIG. 3

4/10




FIG. 4

5/10




FIG. 5

6/10




FIG. 6

7/10




FIG. 7

8/10




FIG. 8

9/10




FIG. 9

10/10




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2016/035949

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F11/14 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

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2010/048048 A2 (MICROSOFT CORP [US]) 29 April 2010 (2010-04-29) abstract figures 1-3,4b paragraphs [0006] - [0015], [0028] - [0035] paragraphs [0039] - [0044], [0049] - [0065], [0069] paragraphs [0072] - [0076] -----	1-15
A	WO 99/17203 A1 (SUN MICROSYSTEMS INC [US]) 8 April 1999 (1999-04-08) abstract figures 1-4 page 3, line 29 - page 5, line 24 page 5, line 36 - page 7, line 15 ----- -/--	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

"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

26 August 2016

Date of mailing of the international search report

06/09/2016

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

Weber, Vincent

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/035949

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 2009/217274 A1 (CORBIN SCOTT ROGER [US] ET AL) 27 August 2009 (2009-08-27) abstract figure 1 paragraphs [0013], [0046] - [0059] -----	1-15
A	WO 2013/072451 A2 (UNIV MADRID POLITECNICA [ES]) 23 May 2013 (2013-05-23) abstract page 1, line 20 - page 2, line 4 page 3, lines 1-34 page 6, lines 1-6 page 13, lines 10-34 -----	1-15

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/035949

Patent document cited in search report		Publication date		Patent family member(s)		Publication date
WO 2010048048	A2	29-04-2010	CN	102197389	A	21-09-2011
			EP	2342653	A2	13-07-2011
			JP	2012507075	A	22-03-2012
			US	2010114826	A1	06-05-2010
			WO	2010048048	A2	29-04-2010

WO 9917203	A1	08-04-1999	AT	211836	T	15-01-2002
			AU	744015	B2	14-02-2002
			AU	9782798	A	23-04-1999
			CA	2304715	A1	08-04-1999
			DE	69803476	D1	28-02-2002
			DE	69803476	T2	12-09-2002
			EP	1019822	A1	19-07-2000
			JP	2001518663	A	16-10-2001
			US	6014669	A	11-01-2000
			WO	9917203	A1	08-04-1999

US 2009217274	A1	27-08-2009	CN	102037463	A	27-04-2011
			EP	2260410	A1	15-12-2010
			US	2009217274	A1	27-08-2009
			WO	2009108409	A1	03-09-2009

WO 2013072451	A2	23-05-2013	EP	2780832	A2	24-09-2014
			US	2016179876	A1	23-06-2016
			WO	2013072451	A2	23-05-2013
