

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
8 April 2010 (08.04.2010)

(10) International Publication Number
WO 2010/037794 A2

(51) International Patent Classification:

G06F 17/30 (2006.01)

(21) International Application Number:

PCT/EP2009/062714

(22) International Filing Date:

30 September 2009 (30.09.2009)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/102,408 3 October 2008 (03.10.2008) US

(71) Applicant (for all designated States except US): TELEFONAKTIEBOLAGET LM ERICSSON (PUBL) [SE/SE]; SUECIA, S-164 83 Stockholm (SE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): HENRIKSEN, Dennis [DK/DK]; Brohusgade 20, 2.th, DK-2200 Copenhagen (DK). NEVADO JIMÉNEZ, Jorge [ES/ES]; Prado de la Iglesia, 2, Bj4, Alpedrete, E-28436 Madrid (ES). SAN MARTIN ARRIBAS, Marta [ES/ES]; Cavanilles 60, E-28007 Madrid (ES).

(74) Agent: ELZABURU MARQUEZ, Alberto; Miguel Angel 21, E-28010 Madrid (ES).

(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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

(54) Title: MONITORING MECHANISM FOR A DISTRIBUTED DATABASE




FIG. -5-

(57) Abstract: The present invention faces the issue of data replication in different database nodes of a geographically distributed database wherein clients cannot always perform any database-related operation in the closest database node. Thus, the present invention provides for an enhanced distributed database system with a plurality of nodes, each node arranged for storing a replica of at least one partition of data, and a method of handling said distributed database system comprising the steps of: partitioning data into a number of partitions; replicating each partition into a number of replicas; for each partition, distributing the number of replicas amongst database nodes; activating more than one node; monitoring at each active node events of: latest updating of each replica, replica status, status of local resources in charge of each replica, and connectivity status of each replica; upon activation or deactivation of a node, determining which node is considered current master node for each partition in charge of current master replica; for any request received in a node to read/write data, determining the current master node in charge of the current master replica, and routing said request to said current master node.

Monitoring mechanism for a distributed database

TECHNICAL FIELD

[0001] The present invention relates to distributed databases and, more particularly, to a geographically distributed database usable as a common database for Telecom networks. More specifically, the invention pertains to an enhanced distributed database system as well as a method of handling such distributed database system.

BACKGROUND

[0002] The present invention addresses the issue of having a common centralized database for a diversity of applications in the field of Telecom networks. Most of Telecom networks, supporting a variety of different generations of telecommunication systems, being wireline or wireless systems, conventionally make use of one or more common centralized databases to store subscription and subscriber data as well as service data for a variety of applications residing in the Telecommunication network or in a 3rd party service network but accessible to subscribers of said Telecom network.

[0003] As the Telecom networks grow, newer generations of telecommunication systems turn up and the existing common centralized databases are not always adaptable to, or suitable to fit, the needs of all telecommunication systems in a Telecom network. Nevertheless, Telecom networks share quite similar requirements to be fulfilled by any particular database system usable therein.

[0004] A conventional and centralized Telecom database is generally required to support, at least, the following characteristics: resiliency and high availability; consistency; high performance and low latency; high capacity; scalability; geographical redundancy; flexible deployment and data model; single point of access (one in each geographical location); and no single point of failure.

[0005] In this respect, geographical redundancy for a conventional and centralized Telecom database has generally been understood as having a replicated node in addition to a main node, the main node being the operative one and the replicated node being on stand-by to resume the operations in case the main node goes down for any reason.

[0006] Nowadays, a purely centralized database, amongst other drawbacks, is quite costly in terms of signalling transport; and, in practice, the usage of its resources is not well balanced, since the variety of application data makes some data to be accessed more frequently than others. Regarding the geographical location where the centralized 5 database is located, the choice may be quite significant on the costs of signalling transport, both on economical terms as well as on load and risk of failure terms. Obviously, the longer the signalling path is between a database client and the database itself the more risk exists of encountering a node unable to further submit the signalling towards its final destination. Likewise, the longer this signalling path is the more load 10 the Telecom network supports and the longer execution times. Apart from that, the Telecom networks are nowadays spread through different and far territories often communicated through different access networks belonging to other network operators; in this scenario, the longer this signalling path is the more network operators may be involved and more costs are derivable thereof.

15 [0007] On the other hand, as exploring other technologies, a distributed database may be regarded as a plurality of databases physically or logically distributed, likely under control of a central database management system, and wherein storage devices are not all necessarily attached to a common processing unit. Thus, the distributed database might be built up with multiple computers located in the same physical location, or may 20 be dispersed over a network of interconnected computers.

[0008] Generally speaking, the distribution of database instances is not necessarily a consequence of data distribution itself but is beneficial for the purpose of data replication in order to obtain high available systems and for geographical redundancy. More particularly, a so-called Data Layer Architecture technology, which replaces the 25 conventional centralized database system by a number of Telecom database front-ends accessing to a common back-end, the back-end being centralized or distributed, is an exemplary applicability of a distributed database fulfilling the above requirements for a Telecom database and usable for data distribution and replication.

[0009] The replication of data in different database instances of a distributed 30 database requires a complex management to keep the existent replicas up-to-date in the different database instances. Moreover, clients of conventionally distributed databases in

the field of Telecom networks cannot always be able to perform any database-related operation in a closer database instance, thus not always minimizing the signalling path required for said database-related operation. Furthermore, clients of conventionally distributed databases in the field of Telecom networks may experience trouble to access
5 data where any database instance is unavailable.

SUMMARY

[0010] The present invention is aimed to at least minimize the above drawbacks and provides for an enhanced distributed database system with a plurality of nodes, wherein each node is arranged for storing a replica of at least one partition of data, and a
10 method of handling said distributed database system.

[0011] In accordance with a first aspect of the present invention, there is provided a new method of handling a distributed database system with a plurality of nodes, wherein each node is arranged for storing a replica of at least one partition of data.

[0012] This method comprises the steps of: partitioning data to be stored into a
15 number 'p' of partitions; replicating each partition into a number 'r' of replicas; for each partition, distributing the number 'r' of replicas amongst corresponding number 'r' of nodes selected from the plurality of nodes; configuring each node with a list of identifiers of other nodes usable to address each other; activating more than one node amongst the plurality of nodes; monitoring at each active node at least one event
20 selected from: latest updating of each replica, replica status, status of local resources in charge of each replica, and connectivity status of each replica; upon activation or deactivation of a node amongst the plurality of nodes, determining which node amongst the active nodes is considered current master node for each partition and in charge of current master replica for said partition; for any request received in a node to read/write
25 data in the distributed database system, determining the partition said data belongs to and the current master node in charge of the current master replica for said partition, and routing said request to said current master node.

[0013] Independently of different embodiments further discussed, on determining which node amongst the active nodes is considered the master node for each partition,
30 and thus being in charge of current master replica for said partition, the method may

include in at least one node amongst the active nodes: a step of collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica; a step of applying pre-configured rules for 5 prioritizing each replica in the active nodes depending on the collected events for each replica; and a step of selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

[0014] The present invention provides for two main embodiments, namely modes 10 of operation, in determining which node amongst the active nodes is considered master node for each partition and in charge of current master replica for the partition.

[0015] In a first mode of operation, and since all the nodes are twin nodes from a processing point of view, all nodes can process the same information and can arrive to determine, for each partition, the same master node in charge of the current master 15 replica. Under this embodiment, not only the above at least one active node, but rather each node in the distributed database system may be arranged to carry out the steps of: collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica; applying pre-configured rules for 20 prioritizing each replica in the active nodes, depending on the collected events for each replica; selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

[0016] In a second mode of operation, upon activation or deactivation of any node 25 amongst the plurality of nodes, the method may further comprise a step of determining the order in which the active nodes have been activated. Where this is the case, all the nodes are arranged for determining the order in which the active nodes were activated, so that the active node firstly activated is considered to be a so-called System Master Monitor node in charge of carrying out the steps of: collecting from each active node 30 information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status

of each replica; applying pre-configured rules for prioritizing each replica in the active nodes, depending on the collected events for each replica; selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for 5 the partition; and informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

[0017] Particularly useful where applying the pre-configured rules, depending on the collected events, produces a same priority for more than one replica, the step in this method of distributing the number 'r' of replica amongst corresponding number 'r' of 10 nodes, for each partition, may include a step of configuring each replica with a default replica priority to be applied where other criteria produce the same replica priorities. Where this is the case, the step of determining which node amongst the active nodes is considered master node for each partition, and in charge of current master replica for said partition, may include in the at least one active node the steps of: collecting from at 15 least one active node information of being configured with a given default replica priority for a partition; and selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition. In particular, where the method operates in accordance with the above first mode of operation, not only the at 20 least one node, but rather each node in the distributed database system may be arranged to carry out these steps.

[0018] However, where the method operates in accordance with the above second mode of operation, the active node firstly activated is considered a System Master Monitor and the step of determining which node amongst the active nodes is considered master node for each partition, and in charge of current master replica for said partition, may include in said System Master Monitor node the steps of: collecting from at least 25 one active node information of being configured with a given default replica priority for a partition; selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition; and informing the other active nodes about the master replica selected for each partition and the master node holding said 30 master replica.

[0019] Generally speaking, all read/write operations are carried out on the master replica for the relevant partition. Consequently, the contents of the master replica and the contents of other replicas for the same partition may vary at a certain point in time. In order to preserve the consistency amongst the different replicas for each partition, the 5 method may further comprise a step of copying for each partition at each active node the contents of the current master replica from the current master node in charge of the current master replica for said partition. Where this step of copying takes place, the method may further comprise a step of marking for each replica copied at each active node at least one of: the latest updating made, replica status, status of local resources in 10 charge of the replica and connectivity status of the replica. This is especially advantageous in order to select another master replica in the future if the current master node in charge of said master replica fails, gets down or unavailable, or becomes an inactive node.

[0020] On the other hand, in accordance with a second aspect of the present 15 invention, there is provided an enhanced distributed database system with a plurality of nodes, wherein each node is arranged for storing a replica of at least one partition of data. In this distributed database system each node includes: a data storage for storing a replica of at least one data partition of data to be stored and for storing identifiers of other nodes usable to address each other; an input/output unit for communicating with 20 other nodes of the distributed database system and with clients requesting read/write operations in the distributed database system; a monitoring unit for monitoring at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica, and connectivity status of each replica; and a processing unit, in cooperation with the data storage, the monitoring unit and the 25 input/output unit, for determining which node amongst active nodes of the distributed database system is considered current master node for each partition and in charge of current master replica for said partition; for determining, for any request received to read/write data in the distributed database system, the partition said data belongs to and the current master node in charge of the current master replica for said partition; and for 30 routing said request to said current master node.

[0021] Aligned with the above method, and irrespective of whether the distributed database system operates in accordance with the first or with the second mode of

operation, the processing unit, the monitoring unit, the data storage and the input/output unit of each node may be arranged for: collecting, from each active node of the distributed database system, information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each

5 replica and connectivity status of each replica; applying pre-configured rules for prioritizing each replica in the active nodes depending on the collected events for each replica; and selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

10 [0022] In particular, and also advantageously aligned with the above method where applying the pre-configured rules, depending on the collected events, produces a same priority for more than one replica, the data storage of each node may be arranged for storing an indicator configured per replica basis to indicate a default replica priority to be applied where other criteria produce the same replica priorities. Where this is the

15 case, the processing unit, the monitoring unit, the data storage and the input/output unit of each node may further be arranged for: collecting from at least one active node information of being configured with a given default replica priority for a partition; and selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be

20 considered the master node for the partition.

[0023] Also aligned with the above method, and especially applicable where the distributed database system operates in accordance with the second mode of operation, the processing unit, the monitoring unit, the data storage and the input/output unit of each node may be arranged for collecting, from each active node of the distributed

25 database system, information to determine the order in which the active nodes have been activated.

[0024] Where this is the case, the active node firstly activated may be considered a System Master Monitor, and the processing unit, the monitoring unit, the data storage and the input/output unit of the System Master Monitor may be arranged for: collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica

and connectivity status of each replica; applying pre-configured rules for prioritizing each replica in the active nodes depending on the collected events for each replica; selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered
5 the master node for the partition; and informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

[0025] Moreover, where the default priority commented above is relevant for prioritizing the replicas for a partition under this second mode of operation, the processing unit, the monitoring unit, the data storage and the input/output unit of the
10 System Master Monitor may further be arranged for: collecting from at least one active node information of being configured with a given default replica priority for a partition; selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition; and informing the other active nodes about
15 the master replica selected for each partition and the master node holding said master replica.

[0026] In order to preserve the consistency amongst the different replicas for each partition, and thus aligned with the above method, the processing unit, the monitoring unit, the data storage and the input/output unit of each node are further arranged for
20 copying for each partition at each active node the contents of the current master replica from the current master node in charge of the current master replica for said partition. Moreover, in order to further selected another master node in charge of the master replica where the present master replica becomes inactive for any reason, the processing unit, the monitoring unit, the data storage and the input/output unit of each node may
25 further be arranged for marking for each replica copied at each active node the latest updating made, replica status, status of local resources in charge of the replica and connectivity status of the replica.

[0027] On the other hand, the invention may be practised by a computer program, in accordance with a third aspect of the invention, the computer program being loadable
30 into an internal memory of a computer with input and output units as well as with a processing unit, and comprising executable code adapted to carry out the above method

steps. In particular, this executable code may be recorded in a carrier readable in the computer.

BRIEF DESCRIPTION OF THE DRAWINGS

[0028] The features, objects and advantages of the invention will become apparent 5 by reading this description in conjunction with the accompanying drawings, in which:

[0029] FIG. 1A illustrates a simplified view of the sequence of actions to be performed for partitioning data to be stored into a number 'p' of partitions and replicating each partition into a number 'r' of replica.

[0030] FIG. 1B illustrates a simplified view of the sequence of actions to be 10 performed for distributing for each partition the number 'r' of replica amongst a corresponding number 'r' of nodes selected from the plurality of nodes and for configuring each node with a list of identifiers of other nodes usable to address each other.

[0031] FIG. 1C illustrates a simplified view of the sequence of actions to be 15 performed, as a continuation of the actions illustrated in FIG. 1A and FIG. 1B, for carrying out a method of handling a distributed database system with a plurality of nodes, wherein each node is arranged for storing a replica of at least one partition of data.

[0032] FIG. 2 shows a simplified view of an exemplary configuration of a plurality 20 of nodes in the distributed database with useful data to describe some embodiments of the invention.

[0033] FIG. 3 illustrates an exemplary implementation of a node amongst the plurality of nodes included in the distributed database.

[0034] FIG. 4 shows a simplified view of an exemplary configuration of a node 25 amongst the plurality of nodes in the distributed database with useful data arranged as clusters to describe some embodiments of the invention.

[0035] FIG. 5 illustrates a simplified view of the sequence of actions to be performed for routing any request received in a node to read/write data in the distributed

database towards the current master node in charge of the current master replica for the partition which such data belongs to.

5 [0036] FIG. 6 illustrates an exemplary state machine provided for in accordance with an embodiment of the invention to determine which node amongst the active nodes is considered a Controller System Monitor in charge of coordinating the other nodes and deciding which the master node is for each replica.

10 [0037] FIG. 7 illustrates an exemplary sequence of actions to be carried out with support of the state machine shown in Fig. 6 in order to determine which node amongst the active nodes is considered a Controller System Monitor upon activation of a number of nodes in the distributed database system.

[0038] FIG. 8 shows still another exemplary sequence of actions to be carried out with support of the state machine shown in Fig. 6 in order to determine which node amongst the active nodes is considered a Controller System Monitor upon inactivity of a previously considered Controller System Monitor.

15 DETAILED DESCRIPTION

[0039] The following describes currently preferred embodiments of an enhanced distributed database system with a plurality of nodes, each node storing a replica of at least one partition of data, and a method of handling said distributed database system.

20 [0040] A Telecom database system may include several geographically distributed nodes, wherein each node may include several data storage units and wherein each data storage unit in each node may allocate a particular replica of a subset of the data, namely a replica of a partition. As illustrated in Fig. 1A and 1B, an exemplary distribution of a data set 10 amongst data storage units of nodes 1-4 may be carried out by following a number of steps provided for in accordance with the invention.

25 [0041] As shown in Fig. 1A, the data set 10 is partitioned during a step S-005 into a number of partitions 11-14, each partition comprising a particular subset of the data set 10. Then, for each partition a number of replicas are generated during a step S-010. The number of replicas for each partition is not required to be the same for all the partitions. Thus, as exemplary illustrated in Fig. 1A, four replicas 111-114 are generated

for the partition 11, whereas three replicas 121-123 and 141-143 are respectively generated for partitions 12 and 14, and only two replicas 131-132 are generated for the partition 13.

5 [0042] As shown in Fig. 1B, these replicas may be grouped per partition basis during a preliminary step S-015 of determining the required geographical distribution. Apart from that, as determining the geographical distribution of nodes that the database system consists of, each node is assigned during a step S-017 an identifier usable for addresses purposes. The exemplary illustration in Fig. 1B shows a distributed database system consisting of four nodes 1-4 with respective identifiers N-1 ID, N-2 ID, N-3 ID
10 and N-4 ID.

15 [0043] Then, during a step S-020, the replicas generated for each partition may be distributed amongst the nodes that the database system consists of. As exemplary illustrated in Fig. 1B, the first replica 111 of the first partition is stored in node 1, the second replica 112 of the first partition is stored in node 2, the third replica 113 of the first partition is stored in node 3, the fourth replica 114 of the first partition is stored in node 4, the first replica 121 of the second partition is stored in node 3, the second replica 122 of the second partition is stored in node 1, the third replica 123 of the second partition is stored in node 2, the first replica 131 of the third partition is stored in node 4, the second replica 132 of the third partition is stored in node 1, the first replica 141 of
20 the fourth partition is stored in node 3, the second replica 142 of the fourth partition is stored in node 4, and the third replica 143 of the fourth partition is stored in node 2. Thus, not every node is required to store a replica for each partition and not every partition must have a replica in each node of the distributed data system provided for in accordance with an aspect of the present invention.

25 [0044] Apart from that, each node may also be configured during this step with identifiers of the other nodes. Thus, the node 1 stores identifiers 151 identifying the nodes 2, 3, 4; the node 2 stores identifiers 152 identifying the nodes 1, 3, 4; the node 3 stores identifiers 153 identifying the nodes 1, 2, 4; and the node 4 stores identifiers 154 identifying the nodes 1, 2, 3.

30 [0045] In operation, and depending on certain events further discussed, one particular replica in a certain node may acquire a highest priority and is thus determined

to be the master replica for a partition whereas the certain node is considered the master node for the partition. However, there may be situations where the events produce the same priority for different replicas of a partition in different nodes so that a master replica can not be determined. In order to save the ambiguity where more than one 5 replica is given the same priority, the present invention provides for configuring each replica with a default replica priority to be applied where other criteria, namely the result of processing the above events, produce the same replica priorities. As shown in Fig. 2, each node 1-4 respectively includes a replica 1101, 2101, 3101, and 4101 for a number of partitions, and a default priority per replica basis 1102, 2102, 3102, and 10 4102.

[0046] Once each node of the distributed database system has been configured as described above, the distributed database system is ready for entering in operation, node by node or as simultaneously as wanted by the operator.

[0047] Fig. 1 C illustrates subsequent sequence of actions to carry out the method 15 of handling said distributed database system in accordance with another aspect of the invention. Even though all the nodes behave in a similar manner, the complete sequence of actions at each node may depend on the order in which the different nodes are activated, as further described in accordance with embodiments of the invention. Thus, Fig. 1C exemplary illustrates a scenario where the node 2 is the firstly activated one, 20 during a step S-030, followed by the activation of the node 3 during a step S-035, then the node 4 during a step S-060, and finally the node 1 during a step S-070.

[0048] The activation of each node 2, 3, 4, 1 may be followed by respective steps 25 S-040, S-045, S-080, S-075 of determining at each active node a start time when the said node has been activated. This optional step is useful to further determine the order in which the active nodes have been activated where the database system is operating with a node acting as a System Master Monitor in charge of coordinating the other nodes and deciding which is the master replica for each partition as further commented. In this respect, as Fig. 2 illustrates, each node 1-4 includes a respective indication 1104, 30 2104, 3104, 4104 indicating the working time that the node has been active since the start time.

5 [0049] Irrespective of whether the step of determining the start time for each node is carried out, the activation of each node 2, 3, 4, 1 is generally followed by respective steps S-050, S-055, S-095, S-090 of monitoring at each active node 2, 3, 4, 1 at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica, and connectivity status of each replica.

10 [0050] To this end, as exemplary illustrated in Fig. 2, each node 1-4 respectively includes an indicator of the latest updating per replica basis 1103, 2103, 3103, 4103, along with the replica 1101, 2101, 3101, 4101 for a number of partitions, and the default priority per replica basis 1102, 2102, 3102, 4102 already commented above. Apart from the data exemplary shown in Fig. 2, other data may be stored per replica basis at each node 1-4 such as those illustrated in Fig. 4. Thus, as exemplary illustrated in Fig. 4, the node 2 respectively includes an indicator 312, 323, 343 of connectivity status per replica basis along with the replica 112, 123, 143 for a number of partitions, and a replica status per replica basis 212, 223, 243 indicating whether the replica is considered the master replica, an active replica being up and running but not being the master, or a stand-by replica on which configuration is taking place so that cannot be considered up and running. In particular, each node may have a further storage field per replica, not illustrated in any drawing, to include indicators of status of local resources per replica basis so that the process of supervising the local resources for each replica may be isolated from the step of monitoring such event. Alternatively, the above step of monitoring at each active node at least one event of the choice may include a step of supervising the local resources for each replica and a step of determining the status of each local resource.

25 [0051] Still with reference to Fig. 4, each node 1-4 of this distributed database system may be configured with a number of clusters wherein each cluster includes: a replica of a partition, an indicator of connectivity status for the replica, a replica status for the replica, an indicator of the latest updating per replica basis, and a default priority per replica basis. Thus, as exemplary illustrated in Fig. 4, the node 2 comprises clusters 160, 170, 180, each cluster respectively including a replica 112, 123, 143 of a partition 30 along with data and indicators for the replica.

[0052] Back to the sequence of actions illustrated in Fig. 1C, and once the step of monitoring events have been carried out at each node, each node initiates during respective steps S-065, S-085, S-100, S-110 the sending of so-called 'Alive' messages towards all the other nodes known to each node.

5 **[0053]** As anyone skilled in the art may understand, some actions in node 1 occur before corresponding actions in node 4, whereas the node 4 was initiated before than the node 1. This is generally possible since the processor load in one node may be higher than in another node, thus resulting in slower performance in the former, and it may also be due to different signalling delays through different network paths.

10 **[0054]** Particularly illustrated in Fig. 1C by way of example, the node 2 had sent during the step S-065 the so-called 'Alive' messages before the node 1 had been activated. As the node 2 receives the 'Alive' messages from the node 1 during the step S-100, the node 2 may be aware that the node 1 had not received the original 'Alive' message and the node 1 does not have the complete information from all the nodes. In 15 these situations, a node like the node 2 in the present exemplary situation carries out a step S-105 of monitoring again the events and a step S-115 of submitting again the 'Alive' messages towards all the nodes known to the node 2. Moreover, as also shown in Fig. 1C, the 'Alive' message is received at the node 2 from the node 3 during a step S-110 intermediate between the step S-105 of monitoring the events at the node 2 and 20 the step S-115 of submitting the 'Alive' messages, so that the latest 'Alive' message submitted from the node 2 may be considered up to date with each node being aware of all other active nodes.

[0055] These 'Alive' messages are exchanged between nodes of the distributed database system and are helpful on determining, upon activation or deactivation of a 25 node of the distributed database system, which node amongst the active nodes is considered the current master node for each partition and in charge of the current master replica for said partition, as further discussed.

[0056] As commented above, certain events determined in operation are taken into consideration on deciding the master replica for a partition and the master node in 30 charge of said master replica. In order to decide which is the master replica for each partition, the following information may be taken into account: which replica for each

partition has the most recently updated contents with complete information, namely, with the updating level; the replica status of each replica for each partition since, obviously, only up and running replicas are eligible; the connectivity status of each replica for each partition; and the default priority configured for each replica of a 5 partition. Moreover, this default priority may be configured to overrule the results of the previous criteria or may be configured to be just applicable where the previous criteria produce the same results for more than one replica of a partition.

[0057] Furthermore, pre-configured rules may be applied for prioritizing each replica in the active nodes depending on these events determined in operation. For 10 example, the pre-configured rules may be such that the connectivity status takes priority over the updating level, or such that the default priority is taken into account immediately after than the replica status, or any other criteria on the events for prioritizing replicas.

[0058] Regarding the contents of each replica for a partition, as well as the 15 updating of such contents, conventional routines may be provided, whereby each replica updates contents at certain times from the master replica. Thus, not all the replicas for a partition update contents at the same time, and not all the replicas for the partition progress with the updating at the same ratio. During the updating, each node in the distributed database system has to monitor how far the updating has progressed so that 20 the information exchanged with the other nodes in this respect takes due account on whether the replica contents can be considered a complete information or not, and the point in time when the updating was carried out.

[0059] In order to take into account the above events on determining the master 25 replica for a partition and the master node in charge of said master replica, the events monitored at each node are communicated towards the other nodes in the distributed database system. To this end, each node 1-4 may be configured with a virtual IP address of the other nodes, or make use of the respective node identifiers 151-154 to identify and address each other, and each node 1-4 may periodically send an 'Alive' message to the other nodes, for example, after expiry of successive delay times. Advantageously in 30 an embodiment of the invention, the 'Alive' messages may be sent with the known TCP

protocol instead of making use of the UDP for some kind of heartbeat such as to avoid the possibility of unidirectional link failures, which are easily detected with TCP.

[0060] More precisely, each 'Alive' message includes a node identifier N-1 ID, N-2 ID, N-3 ID, N-4 ID identifying the sender node to the receiver nodes and may include, 5 for each replica of a partition, at least one of: an identifier of the partition that the replica belongs to, the replica status, updating level, updating time, connectivity status, and default priority.

[0061] Additionally, and especially useful to further determine the order in which the active nodes have been activated, each 'Alive' message sent from any node 1-4 10 towards the other nodes in the distributed database system may respectively include an indication of the working time 1104, 2104, 3104, 4104 that the sender node has been active since its start time.

[0062] The present invention provides for two main embodiments, namely modes of operation, in determining which node amongst the active nodes is considered master 15 node for each partition and in charge of current master replica for the partition.

[0063] In a first mode of operation, and since all the nodes are twin nodes from a processing point of view, each node receives 'Alive' messages from all other nodes with monitoring information determined therein, so that all nodes can process the same information and can arrive to determine, for each partition, the same master node in 20 charge of the current master replica. Under this embodiment, each node in the distributed database system may be arranged to carry out the steps of: collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica; applying pre-configured rules for prioritizing each 25 replica in the active nodes, depending on the collected events for each replica; selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

[0064] Especially advantageous where these events do not result in a highest 30 replica priority for any particular replica under this first mode of operation, each node in

the distributed database system may be arranged to further carry out the steps of: collecting from at least one active node information of being configured with a given default replica priority for a partition; and selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

[0065] In a second mode of operation, all the nodes process the 'Alive' messages to determine the order in which the active nodes were activated so that the active node firstly activated is considered to be a so-called System Master Monitor node in charge of carrying out the steps of: collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica; applying pre-configured rules for prioritizing each replica in the active nodes, depending on the collected events for each replica; selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition; and informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

[0066] Especially advantageous where these events do not result in a highest replica priority for any particular replica under this second mode of operation, the System Master Monitor node in the distributed database system may be arranged to further carry out the steps of: collecting from at least one active node information of being configured with a given default replica priority for a partition; selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition; and informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

[0067] In order to carry out the above method, the present invention generally provides for a distributed database system with a plurality of nodes, such as nodes 1-4, wherein each node is arranged for storing a replica of at least one partition of data. As exemplarily illustrated in Fig. 3 for the node 2, each node comprises a data storage 15 for storing a replica 2101 of at least one data partition 112, 123, 143 of data to be stored,

and for storing identifiers 152 of other nodes usable to address each other; an input/output unit 30 for communicating with other nodes 1, 3, 4 of the distributed database system; a monitoring unit 60 for monitoring at least one event selected from: latest updating 2103 of each replica, replica status 212, 223, 243, status of local 5 resources in charge of each replica, and connectivity status 312, 323, 343 of each replica; and a processing unit 20, acting in cooperation with the data storage, the monitoring unit and the input/output unit, for determining which node amongst active nodes of the distributed database system is considered current master node 2105 for each partition and in charge of current master replica for said partition.

10 [0068] Moreover, in this distributed database system, the processing unit 20, the monitoring unit 60, the data storage 15 and the input/output unit 30 of each node are arranged for: collecting, from each active node of the distributed database system, information about the at least one event selected from: latest updating of each replica 1103, 2103, 3103, 4103, replica status 212, 223, 243, status of local resources in charge 15 of each replica and connectivity status 312, 323, 343 of each replica; applying the pre-configured rules commented above for prioritizing each replica in the active nodes depending on the collected events for each replica; and selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node 2105, 4105 for 20 the partition.

[0069] Furthermore, and especially advantageous where the previous events do not result in a highest priority for any replica in this distributed database system, the processing unit 20, the monitoring unit 60, the data storage 15 and the input/output unit 30 of each node are arranged for: collecting from at least one active node information of 25 being configured with a given default replica priority for a partition; and selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

[0070] In order to operate the distributed database system in accordance with the 30 second mode of operation, the processing unit 20, the monitoring unit 60, the data storage 15 and the input/output unit 30 of each node are arranged for collecting, from

each active node of the distributed database system, information 1104, 2104, 3104, 4104 to determine the order in which the active nodes have been activated, so that the active node firstly activated is considered a System Master Monitor.

5 [0071] In particular, where a node is considered the System Master Monitor node of the distributed database system, the processing unit 20, the monitoring unit 60, the data storage 15 and the input/output unit 30 of said System Master Monitor may further be arranged for informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

10 [0072] The above method of handling the distributed database system may require a particular discussion on how clients of the distributed database system may access the information stored therein and, more particularly, for the purpose of read or write operations.

15 [0073] In principle, clients of the distributed database system, such as a Home Location Register, an Authentication Centre, or a Home Subscriber Server of a Telecom network might be, can access any data from any node of the distributed database system. However, in order to preserve data consistency amongst the different replicas, only one of the replicas would receive a request for reading and/or writing, and this would be the master replica. The data in the master replica, as already commented above, is from time to time updated into the other replicas.

20 [0074] The whole contents of the distributed database system can thus be accessible from any node that the distributed database system consists of. To this end, each node may include one or more access gateways (hereinafter AG), wherein the AG is an entity in charge of forwarding the requests to read/write data towards the node where the master replica is located; and, since the database protocol may be different from the accessing protocol, the AG is in charge of receiving the request from a client to read/write data and accessing the other nodes in the distributed database system. Since there may be provided more than one AG in each node in order to maximize performance, a Load Balancer (hereinafter LB) may also be provided to distribute traffic amongst the more than one AG. However, such LB is not required in a node 25 configuration having just one AG.

[0075] As Fig. 5 illustrates, the node 2 may be provided with an LB 190a and three AG's 191a-193a, and be in charge of a replica 112 for a first partition 11, a replica 123 for a second partition 12, and a replica 143 for a fourth partition 14, wherein the replica 112 is the master replica for the first partition and the node 2 is the master node 5 in charge of said master replica for the first partition; whereas the node 3 may be provided with an LB 190b and three AG's 191b-193b, and be in charge of a replica 113 for a first partition 11, a replica 121 for a second partition 12, and a replica 141 for a fourth partition 14, wherein the replica 121 is the master replica for the second partition and the node 3 is the master node in charge of said master replica for the second 10 partition.

[0076] Thus, this method of handling the distributed database system includes, for any request received in a node to read/write data in the distributed database system, a step of determining the partition that said data belongs to and the current master node in charge of the current master replica for said partition, and a step of routing said request 15 to said current master node.

[0077] As exemplary illustrated in Fig. 5, requests from a client 5 to read/write data in the distributed database system may be received at any node, such as the node 2. Fig. 5 exemplary illustrates a request to read/write data received at the LB 190a during a step S-150. This request may be assigned during a step S-151 to AG 193a, which 20 determines that the data to read/write belongs to the second partition 12, and this AG also determines that the current master node in charge of said partition is the node 3. Then, the AG 193a routes the request towards the node 3 during a step S-152. This request may be received in the LB 190b of the node 3, if more than one AG exists in said node 3, or may be received in a unique AG if the node 3 only includes one AG, or 25 may be received in a particular AG 191b if such AG of the node 3 is known to the AG 193a by configuration means, as exemplary depicted in this Fig. 5. The AG 191b of the node 3 receiving the request accesses during a step S-152 the master replica 121 for the second partition 12 to read/write data therein in accordance with the request.

[0078] On the other hand, Fig. 5 also illustrates the case where the request is 30 received in a master node holding the master replica for the partition which the data

belongs to. In this exemplary case, a request to read/write data is received at the LB 190a of the node 2 during a step S-160.

5 [0079] This request may be assigned during a step S-161 to AG 191a, which determines that the data to read/write belongs to the first partition 11, and this AG also determines that the current master node in charge of said partition is this node 2. Then, the AG 191a internally routes the request to access during a step S-162 the master replica 112 for the first partition 11 to read/write data therein in accordance with the request.

10 [0080] The LB of any node of the distributed database system, such as node 2, may be constructed with an input/output unit 50 dedicated for communicating with the client 5, as illustrated in Fig. 3, and with resources of the processing unit 20 arranged for selecting the appropriate AG to balance the workload and performance of the node. This input/output unit 50 may be an integral part of the input/output unit 30 that each node of the distributed database system comprises, or may be a separate module included in said 15 input/output unit 30.

20 [0081] The AG of any node of the distributed database system, such as node 2, may be constructed with an input/output unit 40 dedicated for communicating with other nodes 1, 3, 4, as illustrated in Fig. 3, and with resources of the processing unit 20 arranged for determining the partition that the data to read/write belongs to, for determining whether the master node in charge of the master replica for said partition is the current node or another node of the distributed database system, and for accessing the data storage 15 to access said data, where the master node in charge of the master replica for said partition is the current node, or for routing the request towards another node determined to be the current master node for the partition, otherwise.

25 [0082] In accordance with embodiments of the invention, the monitoring unit 60 may include a unique unit to monitor and compile the above events, a so-called Local System Monitor (hereinafter LSM) in the instant specification, or might include an active LSM and a stand-by LSM so that the latter could take over operations in case the former fails. In the following, any reference to the monitoring unit or to the LSM is 30 constructed as meaning the active LSM in the node under reference.

[0083] In particular, where the invention operates in accordance with the above second mode of operation, the monitoring unit 60 of the so-called System Master Monitor node is considered to be a Controller System Monitor (hereinafter CSM) whereas each monitoring unit of the other nodes in the distributed database system is 5 still referred to as LSM. The CSM may thus decide which the master replica is, by taking into account the events information received from the LSM of each node in the distributed database system, and by applying the pre-configured rules for prioritizing each replica in the active nodes. The CSM communicates each LSM in the other nodes what the master replica for each partition is.

10 [0084] In determining which node amongst the active nodes is considered the System Master Monitor in the second mode of operation the exemplary state-machine illustrated in Fig. 6 may be applied, wherein transitions between states are due either to the reception of an 'Alive' message from other node or the expiration of a timer. For the sake of simplicity, rather than referring to the System Master Monitor node versus the 15 other nodes of the distributed database system, we refer in this discussion to their respective monitoring units 60, namely, to the CSM of the System Master Monitor node versus each LSM of said other nodes.

[0085] In an embodiment of the invention, illustrated in Fig. 6, the possible states and transitions may be:

20 – **Inactive.** It is the state when each LSM starts. If the LSM is in this state, the node does not answer requests, and it can host no master replicas. Transition to this state makes the LSM start sending alive messages to the rest of the nodes.

– **Active.** In active state, each LSM is listening to the 'Alive' messages with information about the other nodes and, particularly, listening to 'Alive' messages 25 from the CSM, if any, with the master replicas. Each LSM is also sending information about its own node, distributing information about the master replicas internally, in its own node to any local AG; and, optionally, forwarding the information to the rest of the nodes in the 'Alive' messages.

– **Potential CSM.** If an LSM stays in this state for a certain configurable time, a so-called DELAY_TIME in the instant specification, such LSM becomes the 30

CSM. It also happens if it receives the ‘Alive’ messages from all the rest of the nodes in the system before the `DELAY_TIME` expires.

- **CSM.** A node reaching this state decides which replica is the master replica for every partition in the distributed database system.

5 [0086] Nevertheless, other embodiments of states, transitions, or both are foreseeable as well.

[0087] Regarding the transitions between the states shown in Fig. 6, wherein transitions from any state to the inactive state are not shown for the sake of simplicity, any LSM starts in inactive state and, as soon as the node owning the LSM sends ‘Alive’ 10 messages to all the other nodes in the distributed database system, this LSM goes to active state. That is, the transition ST-1 from the inactive state to active state is the sending of ‘Alive’ messages from the node owning the LSM in inactive state towards all the other nodes in the distributed database system.

[0088] In accordance with an optional embodiment of the invention provided for 15 in order to solve consistency in split brain situations, by avoiding more than one sub-network to be active, the transition to inactive state may happen when one or more nodes fail, and there are less than $(n+1)/2$ nodes up and running, including the CSM, or where there are less than $n/2+1$, not including the CSM, being ‘n’ the total number of 20 nodes in the distributed database system. For instance, if we have a distributed database system consisting of three nodes and the node hosting the CSM becomes isolated, its sub-network has only one node, that is less than $(3+1)/2=2$. So, this node hosting the CSM goes to inactive state. The other sub-network has two nodes, that means that it has $2/2+1=2$ nodes up, so they remain active and a new CSM is elected. In this optional 25 configuration, each LSM detects that one or more nodes are down if after a configurable period of time, a so-called `INACTIVE_TIME` in the instant specification, each LSM receives no ‘Alive’ messages from said one or more nodes. The transition to inactive makes any LSM resetting its execution time, thus becoming “young” again. That prevents a previous CSM that was isolated to become CSM again and send deprecated information.

5 [0089] This optional embodiment may be carried out by appropriately configuring the distributed database system. If such embodiment is not wanted, a suitable configuration parameter can be reset so that the account of number of nodes isolated from the system is not relevant for letting the nodes go to the active state, and letting the system to operate with more than one sub-network even if isolated between them.

10 [0090] Back to the transitions between the states shown in Fig. 6, and assuming that the optional embodiment commented above is configured to operate, any node with the LSM in the active state can go back to inactive state if not enough 'Alive' messages from other nodes are received, namely, less than $(n+1)/2$ nodes up including the CSM, or less than $n/2+1$ not including the CSM.

15 [0091] Otherwise, where enough 'Alive' messages are received from other nodes and the information in the received 'Alive' messages indicate that the sender nodes started later, and no more nodes remain to send messages, the LSM in the receiver node becomes CSM. That is, the transition ST-2.1 from the active state to CSM state is the reception of enough 'Alive' messages from the other nodes, and no more nodes remain to send messages in the distributed database system. Particularly in a two-node system, DELAY_TIME may be gone with no messages received, and the node becomes CSM as well. However, where 'Alive' messages were received from other nodes indicating that said nodes started later, but there are still nodes from which 'Alive' messages are 20 expected to be received, the transition ST-2.2 takes place and the LSM in the receiver node becomes Potential CSM.

25 [0092] From the Potential CSM state, a node can go back to the inactive state, for the same reasons as commented for previous transitions. Otherwise, where a still expected 'Alive' message is received from a node indicating its LSM is a consolidated CSM, or where a still expected 'Alive' message is received from a node showing that other LSM started earlier, the transition ST-3.1 illustrated in Fig. 6 takes place and the LSM in Potential CSM state goes to active state in the node receiving any of these expected 'Alive' messages. On the other hand, where the so-called DELAY_TIME has passed without receiving any further expected 'Alive' message showing a older node, or 30 where 'Alive' messages were received from the rest of the nodes showing that no one started earlier, the transition ST-3.2 illustrated in Fig. 6 takes place and the LSM in

Potential CSM state goes to CSM state in the node where the DELAY_TIME has expired or receiving any of these expected ‘Alive’ messages.

[0093] As already commented above, there may be two types of timers: the so-called DELAY_TIME timer in the instant specification, which is the time an LSM waits 5 for ‘Alive’ messages of the rest of the nodes informing about an older LSM before proclaiming itself as CSM; and the so-called INACTIVE_TIME timer in the instant specification, which is the time an LSM waits for ‘Alive’ messages from a node before concluding such node is down and unavailable.

[0094] Once the above process is finished, the node with a monitoring unit 10 60 consolidated as CSM starts sending ‘Alive’ messages including the list of master replicas.

[0095] Further to the handling of the state-machine illustrated in Fig. 6, the Fig. 7 shows an exemplary sequence of actions amongst nodes of a distributed database system to determine which node amongst 3 exemplary nodes is the node with a monitoring unit 15 60 consolidated as CSM. As already commented above, the CSM is the LSM that is activated first. In order to determine that, each LSM at start up sends ‘Alive’ messages with the working time 2104, 4104, as exemplary shown in Fig. 2, to the rest of the nodes. The nodes that are part of the system are known by configuration. After sending the ‘Alive’ messages, the LSM waits for a certain period of time, DELAY_TIME, to 20 receive ‘Alive’ messages from other nodes. The phase of establishing the CSM finishes when that time is gone, or before that, if ‘Alive’ messages from sufficient other nodes are received. During that phase, a potential CSM can be assigned, according to the information received until that time, but it won’t be confirmed till this phase is over.

[0096] Thus, as illustrated in Fig. 7, the first node to be activated is node 1 during 25 a step S-200. As soon as the LSM starts, namely, the monitoring unit 60 of the node 1, it sends ‘Alive’ messages during the step S-205 to the other two nodes 1 and 2 of the distributed database system, and also starts a timer of DELAY_TIME seconds. Since the other nodes are not running yet each corresponding LSM do not receive such ‘Alive’ messages.

[0097] Then, the LSM of node 2 starts running during a step S-210, it sends 'Alive' messages to nodes 1 and 3 during a step S-220 and starts its own timer of DELAY_TIMER seconds. After the LSM of the node 2 has started, but before it sent the 'Alive' messages, the LSM of the node 3 starts during a step S-215 and starts its own timer of DELAY_TIMER seconds.

[0098] When the node 1 receives the 'Alive' message from the node 2 appoints itself as Potential CSM during a step S-225 since, with the information it has, it is the LSM that has started earlier whilst waiting for information from the node 3. At this stage, the node 3 is able to send 'Alive' messages towards nodes 1 and 2 during a step 10 S-230.

[0099] When the node 2 receives the 'Alive' message from the node 3, it appoints itself as Potential CSM during a step S-240 since, with the information it has, the LSM of the node 2 has started earlier than the LSM of the node 3. Likewise, when the node 3 receives the 'Alive' message from the node 2, it appoints node 2 as Potential CSM 15 during a step S-245 since, with the information it has, the LSM of the node 2 has started earlier, whilst both nodes 1 and 2 are awaiting information from the node 1, still unknown to the node 3.

[0100] When the node 1 receives the 'Alive' message from the node 3, it appoints itself as consolidated CSM during a step S-235 since the node 1 has already received 20 'Alive' messages from all the nodes in the system and, with that information, the LSM of the node 1 is the one that was activated earlier. Then, the node 1 sends 'Alive' messages during a step S-250 to the other nodes informing them that the node 1 is the consolidated CSM. The nodes 2 and 3 eventually receiving the 'Alive' messages from the node 1 realize and thus mark during respective steps S-255 and S-260 that the 25 monitoring unit of the node 1 is the CSM, and no other node can assume this role in the current situation.

[0101] The 'Alive' messages are sent from each node to the rest of the nodes periodically, in order to be able to reassign the CSM role in case of failures. In this way, any change in the configuration of the distributed database system, like adding or 30 removing nodes, can be immediately known in all the nodes.

[0102] Thus, in accordance with both embodiments of the invention, namely with both modes of operation, each node, based on the information received from the other nodes, may determine which node is considered the CSM, i.e. the one that has a longer working time amongst the active nodes. In order for a node to become a CSM, such 5 node could wait at least the `DELAY_TIME` for assuring there has been some time for receiving remote 'Alive' messages from any other potential CSM. The node reaching the CSM state may communicate its decision to the remaining nodes in order to assure there are no ambiguities, though this communication is not necessary where all the nodes behave in accordance with the first mode of operation. This communication may 10 be TCP-based in order to assure reliability. In order to reduce reconfiguration times, this behaviour may be indirectly achieved when the CSM communicates the replica configuration to the other nodes. In this respect, as the replica status is included in the 'Alive' messages, once elected, the CSM knows which replica status configuration is suitable for working. In fact, this information may be known to all nodes, but they may 15 wait for CSM to confirm under the second mode of operation.

[0103] Each node could maintain a list with the current nodes which it has received an 'Alive' message from, along with the replica status received within. Each time the 'Alive' message is received, the sender node may be set as active. Master nodes for partitions may only be elected within active nodes. A node may be set as not 20 available when there has been a period, a so-called `INACTIVE_PERIOD`, without receiving messages from it. This time could be two or three times the mean time of the receiving messages period, which could be set initially to the aforementioned `DELAY_TIME`. The node may also be set as unavailable if the 'Alive' messages sent to it do not get through. In this way node availability is detected very fast.

25 **[0104]** Generally speaking, the node with monitoring unit 60 considered the CSM may submit towards the other nodes in the distributed database system its node identifier, working time, replica status for each replica, master node for each replica, node status from the state-machine, list of active nodes from where 'Alive' messages have been received (even if offline), current Master Replica Information (hereinafter 30 MRI), update time and node Id (including exec time) for the CSM node that set the MRI.

5 [0105] In particular, a so-called MRI message can include a list of replicas with their hosting node. In an embodiment of the invention, the MRI message might also be sent also to offline nodes, in order to receive confirmation for the MRI. In this respect, an offline node could be the link between a real CSM node and a so-called sub-CSM node, the latter being an LSM running in a node that believes that it is the master, but it is not so. Anyhow, as stated before, no replica of an offline node could be set as master. Therefore, in order to avoid this problem, the CSM process may wait the DELAY_TIME before sending the MRI.

10 [0106] Generally speaking, the CSM election mechanism is such that all monitoring processes are synchronized. As explained before the idea is that the oldest process becomes the CSM. In determining which the oldest LSM is, several embodiments are foreseeable. In a first embodiment, the present invention takes into account the working time, that is, each process determines, with its local time, how many seconds it has been working from the start, and sends that info in the 'Alive' 15 message. There may be a subjacent drawback in this case, as probably latency has an important role to play. The receiving node will see the working time sent, but no the latency, so it would difficult to accurately determine whether the receiving node is younger or older than the sending node. Latency times could be measured, via pings message, and a mean could be established.

20 [0107] In a second embodiment, the present invention takes into account the Start-up time. Under this embodiment, all processes send their start-up time in the 'Alive' messages. This time will be relative to local machine time at each node. Therefore, there may be a need to synchronize all machines in the system. Linux systems, and operating systems in general, have solved this problem many years ago using NTP.

25 [0108] Under these first and second embodiments, and during the CSM election process, whenever an LSM reaches the inactive state for the exemplary state-machine illustrated in Fig. 6, the respective working time or start-up time might need to be reset in order to avoid that a degraded CSM reaches again the CSM state when it recovers.

30 [0109] As already commented above, upon activation or deactivation of any node of the distributed database system, there is a step of determining amongst the active nodes which one is the current master node in charge of the master replica for each

partition. In determining which this master node is, the present invention provides for two embodiments: a first embodiment where all the nodes work independently to determine which the master node is; and a second embodiment wherein a CSM is determined to decide which the master node is for each partition.

5 [0110] Fig. 8 illustrates an exemplary situation wherein the currently considered CSM, which exemplary is the node 3, goes down and becomes unavailable for the other nodes, namely, for nodes 1 and 2.

[0111] As shown in Fig. 8, the node 3 acting as CSM sent a latest set of 'Alive' messages during a step S-300 towards the other nodes 1 and 2. Upon respective 10 reception of such messages, both nodes 1 and 2 trigger a so-called Inactivity Period which is not reset until new 'Alive' messages are received from the node 3. Since the Inactivity Period expires at both nodes 1 and 2 without having received further 'Alive' messages from the node 3, both nodes 1 and 2 mark the node 3 as unavailable during respective steps S-320 and S-330. In this exemplary case, the node 2 was older than the 15 node 1, and this information was known to node 2 since it had received the latest periodic 'Alive' message from the node 1 during the step S-310. Therefore, the node 2 goes itself to the Potential CSM state during a step S-340, and in accordance with the state-machine illustrated in Fig. 6. Then, the node 2 sends during a step S-350 its own periodic 'Alive' messages towards nodes 1 and 3, the latter being unavailable to receive 20 such message. The node 1 receiving the 'Alive' message from the node 2 is aware that the node 2 is older and arrives to the conclusion during a step S-360 that the node 2 is the current CSM. After having expired at the node 2 the so-called *DELAY_TIME*, the monitoring unit 60 of the node 2 is consolidated as CSM during a step S-370, and the node 2 is the current System Master Monitor.

25 [0112] Further, where the node 3 recovers again, which is not shown in any drawing, the node 3 sends its own periodic 'Alive' messages towards nodes 1 and 2. Upon receipt of such messages at nodes 1 and 2, they both mark the node 3 as available again. As already commented above, the node 3 has reset its working time or the start-up time, as the case may be, so that the situation does not change and, after being in a 30 Potential CSM state during a *DELAY_TIME*, the current CSM is still the monitoring

unit of the node 2, which reaches the CSM state in accordance with the state-machine shown in Fig. 6.

[0113] Regarding the addition of new nodes never before included in the distributed database system, the situation, from a CSM election perspective, is very 5 similar as the situation where a node has been down and it recovers again. The only difference is that the failing node is known to the CSM because it is present in configuration tables, whereas a completely new node is unknown to the existing nodes. In this respect, the CSM as well as the other monitoring processes do not care about 'Alive' messages received from non-configured nodes. Therefore, an important aspect 10 where introducing a new node is the inclusion of an identifier of the new node into the corresponding configuration tables 151-154 of the already existing nodes, as illustrated in Fig. 1B, and the configuration of the replicas in such new node. Once this has been completed, the monitoring unit of the new node can start. In this way, the new node 15 sends 'Alive' messages to the other already existing active nodes, and the other already existing active nodes will be aware of the new node. As the CSM becomes aware of the new replicas in the new node, it reconfigures the system, if necessary, accordingly. In this respect, where the system operates in accordance with the above first mode of 20 operation, that is, without having a CSM, all nodes respectively process the information received in the 'Alive' messages and arrive to a same conclusion on which is the master node in charge of the master replica for each partition. In principle, the 'Alive' messages comprise the information in the so-called MRI messages. Nevertheless, the MRI 25 messages may be sent separately from the 'Alive' messages.

[0114] On the other hand, the addition of a new partition to a running distributed database system may be an easy task with the behaviour previously described. One may 25 think on replicating the partition on at least two replicas to be respectively added into two existing nodes in order to get higher availability than with just one replica in one node. As before explained for configuration of any node, apart from the replica, all the corresponding indicators explained above per replica basis have to be configured as well. Thus, the elected System Master Monitor node with the CSM under the second 30 mode of operation, or every node under the first mode of operation, sooner or later receives 'Alive' messages from all the nodes, wherein information for the new replicas are received from the nodes hosting said new replicas.

[0115] In order to effectively providing a highly available distributed database system, the monitoring process is expected to be highly available too. Therefore, as already commented above, two LSM may be running in each node: an active LSM and a stand-by LSM. In operation, the active LSM forwards all received messages to the 5 stand-by LSM, particularly, the 'Alive' messages and also any MRI message, if such MRI message is received. In this way, in case the active LSM crashes, the stand-by LSM takes over immediately, so that a node cannot go down simply because the monitoring process has failed.

[0116] For the purpose of the present invention, the monitoring process may be 10 carried out by the monitoring unit 60 likely in cooperation with the processing unit 20.

[0117] Apart from all the above, other specific embodiments are foreseeable. For instance, all dynamic data might be maintained dynamically, instead of being stored along with the configuration data, and updated from time to time. This dynamic data includes the MRI information, as well as the information about the replicas and nodes 15 that is usually sent in the 'Alive' messages. This may be easily achieved by using multicast techniques in a port different than the one used for processes that answer the requests.

[0118] In this respect, all 'Alive' messages received from the other nodes, as well 20 as 'Alive' messages sent to configured nodes, may be sent by multicast, no special treatment is necessary. Likewise, MRI messages received by any LSM of a node, or sent by the CSM of the System Master Monitor, may be received or sent by multicast.

[0119] Since the MRI is a two phase commit, special considerations are necessary:

- If it is an LSM process, and no MRI_ACK has been sent, there is no problem as 25 CSM may reinitiate the connection (anyway, it is very unlikely to happen just in the middle). Nothing to be sent by multicast.
- If it is an LSM process, and the MRI_ACK has been sent, multicast the MRI with a flag indicating "confirmation pending" (for MRI_NACK nothing must be done as it will not be confirmed). This way, the stand-by process may be able to understand the confirmation.

- If it is an LSM process and an MRI_CONFIRM has been received, multicast the MRI with the flag as “confirmed”. Stand-by process may interpret the MRI as the most recent one.
- If it is the CSM process, and the MRI has not been sent to any node, nothing must be multicast (the stand-by process, as getting active, may detect that its current MRI does not match the elaborated one, and may initiate the process).
- If it is the CSM process and the MRI has been sent to any node, multicast it as “pendant of confirmation”. If the stand-by is to take control it may re-initiate connections with all LMs (already opened connections may time-out and remote SM process may ignore first MRI indication) and send the MRI again.
- If it is the CSM process and the MRI_CONFIRM has been sent to any node, multicast the MRI with flag “confirming”. If the stand-by is to take control, it may re-initiate connections to all LSM process to confirm again (already confirmed ones may ignore the messages).
- If it is the CSM process and the MRI_CONFIRM has been sent to all nodes, multicast the MRI as “confirmed”, so the stand-by may interpret it as the newer one, but may not initiate confirmation procedure.

[0120] As explained above, the MRI messages may need for multicast an additional flag for the “confirmation” state at the very end in order to use common code for parsing. Anyway, all multicast messages might also include, in the last part, one byte that represents the weight of the process in order to resolve race conditions between active and stand-by.

[0121] This section describes how the active and stand-by process may detect their state, and what are they supposed to do. Whenever the process starts, it may start listening to receive any multicast message in the distributed database system. If some multicast is received, it may be set as stand-by, and may update any internal data according to the received state (it may not process the state machine as its elected state is received from the Alive message). If the process does not receive any multicast message in **DELAY_TIME** period after starting, or after previous multicast received packet, it may start listening (opening the port) for ‘Alive’ and MRI message, and may

send its current MRI (which could be empty at initial start up) with the process weight (it will be different for any process in the node). In case the active process receives a multicast message with a weight greater than its own and a non empty MRI, it may stop listening on previous port and may degrade to stand-by (this may solve initial race

5 conditions, which are supposed not to happen during normal operation). The purpose of comparing the empty MRI is to avoid a more weighted process to get control if re-started very fast (anyway, it could wait *DELAY_TIME* after restarting, so this could not happen unless it just starts at the same time as current active process send its 'Alive' message, and also at the same time neighbour nodes also send their 'Alive' message,

10 which are also multicast).

[0122] The following table describes a possible implementation of the protocol used between nodes of the above distributed database system in accordance with an embodiment of the invention:

Position (Bytes)	Length (bytes)	Name	Comments
0	1	Msg Type	0x01 will be the Alive message
1	1	Version	
2	2	Node Id	Node Id
4	4	Exec Time	Execution time of current node
8	2	m	Number of replicas local monitored. Only available replicas are sent
10	2	replica master status	Two bytes ORed informing whether sent replicas are configured or not as masters in the node. First node is first bit (relative to 2^0) and bytes are considered in Big Indian order. Example: 0x5 (0..0101) means first and third node sent are master.
12	1	Node status	Node state: inactive, active, potential CSM, CSM (1, 2, 3, 4) according to event

			machine
P=13	2	replica id	Id of first replica sent
P+1	1	Memory usage & Ds status	Percentage of memory used in replica
P+3	2	replica Idgroup	Id of Second replica
P+5	1	Memory usage	Memory usage in second replica
...			Remaining available replicas
P+3*(m-1)	2	replica Id	Id of last replica
P+3*(m-1)+2	1	Memory & St	Percentage for the last Ds.
P+3*m	2	Nodes n	Number of available nodes
Q = P +3*m +2	2	Node Id	Id of fist available node
Q+2	2	Node Id	Id of second available node
...			
Q + 2*(n-1)	2	Node Id	Id of last available node
Q + 2*n	8+5*D	MRI	Master Replicas Information D represents the number of replicas in this field.

[0123] The invention may also be practised by a computer program, loadable into an internal memory of a computer with input and output units as well as with a processing unit. This computer program comprises to this end executable code adapted 5 to carry out the above method steps when running in the computer. In particular, the executable code may be recorded in a carrier readable means in a computer.

[0124] The invention is described above in connection with various embodiments that are intended to be illustrative and non-restrictive. It is expected that those of ordinary skill in this art may modify these embodiments. The scope of the invention is 10 defined by the claims in conjunction with the description and drawings, and all modifications that fall within the scope of the claims are intended to be included therein.

CLAIMS

1. A method of handling a distributed database system with a plurality of nodes, wherein each node is arranged for storing a replica of at least one partition of data, the method comprising the steps of:

5 – partitioning (S-005) data to be stored into a number ‘p’ of partitions;

 – replicating (S-010) each partition into a number ‘r’ of replicas;

 – for each partition, distributing (S-015, S-020) the number ‘r’ of replicas amongst corresponding number ‘r’ of nodes selected from the plurality of nodes;

10 – configuring (S-017, S-020) each node with a list of identifiers of other nodes usable to address each other; and

 – activating (S-030, S-035, S-060, S-070) more than one node amongst the plurality of nodes;

and **characterised in that** it also comprises the steps of:

15 – monitoring (S-050, S-055, S-090, S-095) at each active node at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica, and connectivity status of each replica;

 – upon activation or deactivation of a node amongst the plurality of nodes, determining (S-105, S-115) which node amongst the active nodes is considered current master node (2105) for each partition and in charge of current master replica for said partition; and

20 – for any request received (S-150; S-160) in a node to read/write data in the distributed database system, determining the partition (11, 12) said data belongs to and the current master node (2105) in charge of the current master replica (112, 121) for said partition, and routing (S-151, S-152, S-153; S-161) said request to said current master node.

2. The method of claim 1, wherein the step of determining which node amongst the active nodes is considered the master node for each partition, and in charge of current master replica for said partition, includes in at least one node selected amongst the active nodes the steps of:
 - 5 – collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica;
 - applying pre-configured rules for prioritizing each replica in the active nodes depending on the collected events for each replica; and
 - 10 – selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.
3. The method of claim 2, wherein the step of distributing the number ‘r’ of replica amongst corresponding number ‘r’ of nodes, for each partition, includes a step of configuring each replica with a default replica priority to be applied where other criteria produce the same replica priorities.
 - 15
4. The method of claim 3, wherein the step of determining which node amongst the active nodes is considered master node for each partition, and in charge of current master replica for said partition, includes in at least one node selected amongst the active nodes the steps of:
 - 20 – collecting from at least one active node information of being configured with a given default replica priority for a partition; and
 - selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.
- 25
5. The method of claim 1, wherein upon activation or deactivation of a node amongst the plurality of nodes, the method further comprises a step of determining the order in which the active nodes have been activated.

6. The method of claim 5, wherein the active node firstly activated is considered a System Master Monitor in charge of carrying out the steps of:

- collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica;
- applying pre-configured rules for prioritizing each replica in the active nodes depending on the collected events for each replica;
- selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition; and
- informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

7. The method of claim 5, wherein the active node firstly activated is considered a System Master Monitor in charge of carrying out the steps of:

- collecting from at least one active node information of being configured with a given default replica priority for a partition;
- selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition; and
- informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.

8. The method of claim 1, further comprising a step of copying for each partition at each active node the contents of the current master replica from the current master node (105) in charge of the current master replica (112, 121, 132, 142) for said partition.

9. The method of claim 8, further comprising a step of marking for each replica copied at each active node at least one of: the latest updating made, replica status,

status of local resources in charge of the replica and connectivity status of the replica.

10. A distributed database system with a plurality of nodes (1, 2, 3, 4), wherein each node (2) is arranged for storing a replica of at least one partition of data, and each 5 node includes:

- a data storage (15) for storing a replica (2101) of at least one data partition (112, 123, 143) of data to be stored and for storing identifiers (152) of other nodes usable to address each other;
- an input/output unit (30) for communicating with other nodes (1, 3, 4) of the 10 distributed database system and with clients (5) requesting read/write operations in the distributed database system;
- a monitoring unit (60) for monitoring at least one event selected from: latest updating (2103) of each replica, replica status (212, 223, 243), status of local resources in charge of each replica, and connectivity status (312, 323, 343) of 15 each replica (112, 123, 143); and
- a processing unit (20), in cooperation with the data storage, the monitoring unit and the input/output unit, for determining which node amongst active nodes of the distributed database system is considered current master node for each partition and in charge of current master replica for said partition; for determining, for any request received to read/write data in the distributed 20 database system, the partition said data belongs to and the current master node in charge of the current master replica for said partition; and for routing said request to said current master node.

11. The distributed database system of claim 10, wherein the processing unit (20), the 25 monitoring unit (60), the data storage (15) and the input/output unit (30) of each node are arranged for:

- collecting, from each active node of the distributed database system, information about the at least one event selected from: latest updating of each replica (1103, 2103, 3103, 4103), replica status (212, 223, 243), status of local

resources in charge of each replica and connectivity status (312, 323, 343) of each replica;

- applying pre-configured rules for prioritizing each replica in the active nodes depending on the collected events for each replica; and

5 – selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node (2105, 4105) for the partition.

12. The distributed database system of claim 11, wherein the data storage (15) of each node is arranged for storing an indicator (2102) configured per replica basis to indicate a default replica priority to be applied where other criteria produce the same replica priorities.

10 13. The distributed database system of claim 12, wherein the processing unit (20), the monitoring unit (60), the data storage (15) and the input/output unit (30) of each node are further arranged for:

15 – collecting from at least one active node information of being configured with a given default replica priority for a partition; and

– selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node for the partition.

20 14. The distributed database system of claim 10, wherein the processing unit (20), the monitoring unit (60), the data storage (15) and the input/output unit (30) of each node are arranged for collecting, from each active node of the distributed database system, information (1104, 2104, 3104, 4104) to determine the order in which the active nodes have been activated.

25 15. The distributed database system of claim 14, wherein the active node firstly activated is considered a System Master Monitor, and wherein the processing unit (20), the monitoring unit (60), the data storage (15) and the input/output unit (30) of the System Master Monitor are arranged for:


- collecting from each active node information about the at least one event selected from: latest updating of each replica, replica status, status of local resources in charge of each replica and connectivity status of each replica;
- applying pre-configured rules for prioritizing each replica in the active nodes 5 depending on the collected events for each replica;
- selecting a replica in a specific node for each partition with a highest replica priority, this replica to be considered the master replica and this specific node to be considered the master node (2105, 4105) for the partition; and
- informing the other active nodes about the master replica selected for each 10 partition and the master node holding said master replica.

16. The distributed database system of claim 14, wherein the active node firstly activated is considered a System Master Monitor, and wherein the processing unit (20), the monitoring unit (60), the data storage (15) and the input/output unit (30) of the System Master Monitor are further arranged for:



- collecting from at least one active node information of being configured with a 15 given default replica priority for a partition;
- selecting a replica in a specific node for each partition with a highest default replica priority, this replica to be considered the master replica and this specific node to be considered the master node (2105, 4105) for the partition; 20 and
- informing the other active nodes about the master replica selected for each partition and the master node holding said master replica.


17. The distributed database system of claim 10, wherein the processing unit (20), the monitoring unit (60), the data storage (15) and the input/output unit (30) of each node are further arranged for copying for each partition at each active node the 25 contents of the current master replica from the current master node (2105, 4105) in charge of the current master replica (112, 121, 132, 142) for said partition.

18. The distributed database system of claim 17, wherein the processing unit (20), the monitoring unit (60), the data storage (15) and the input/output unit (30) of each node are further arranged for marking for each replica copied at each active node the latest updating made, replica status, status of local resources in charge of the replica and connectivity status of the replica.
5
19. A computer program, loadable into an internal memory of a computer with input and output units as well as with a processing unit, the computer program comprising executable code adapted to carry out the method steps according to any of claims 1 to 9 when running in the computer.
- 10 20. A carrier medium readable in a computer and including the computer program of claim 19.

**FIG. -1A-**


2 (10)

F16.-1B-




Node-1				Node-2				Node-3				Node-4							
1101		1102		1103		2101		2102		2103		4101		4102		4103		4104	
Replica 1-P1	Priority 2	Replica 2-P1	Priority 2	Replica 1-P2	Priority 3	Replica 2-P2	Priority 3	Replica 1-P3	Priority 2	Replica 2-P3	Priority 2	Replica 4-P1	Priority 4	Replica 1-P4	Priority 2	Replica 2-P4	Priority 1	Replica 3-P1	Priority 3
1104																			
Working Time: 87654				Working Time: 99001				Working Time: 222				Working Time: 65432				Working Time: 4105			
1104				2104				4104				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105				4105				4 (10)			
1104				2105				4105											


5 (10)

**FIG.-3-**

6 (10)

FIG.-4-

7 (10)

**FIG.-5-**

8 (10)






FIG.-6-

9 (10)

**FIG.-7-**

10 (10)

FIG.-8-