

(43) International Publication Date
31 March 2016 (31.03.2016)(51) International Patent Classification:
G06F 9/445 (2006.01) G06F 9/50 (2006.01)

(74) Agents: MEYER, Sheldon, R. et al.; Tucker Ellis LLP, One Market Plaza, Steuart Tower, Suite 700, San Francisco, California 94105 (US).

(21) International Application Number:

PCT/US2015/052060

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

(22) International Filing Date:
24 September 2015 (24.09.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/055,482 25 September 2014 (25.09.2014) US
62/076,159 6 November 2014 (06.11.2014) US


(71) Applicant: ORACLE INTERNATIONAL CORPORATION [US/US]; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).

(72) Inventors: SEOVIC, Aleksandar; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US). RAJA, Harvey; 500 Oracle Parkway, M/S 5op7, Redwood Shores, California 94065 (US).

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR SUPPORTING DYNAMIC DEPLOYMENT OF EXECUTABLE CODE IN A DISTRIBUTED COMPUTING ENVIRONMENT

FIGURE 2A

(57) Abstract: A system and method supports dynamic deployment of executable code in a distributed computing environment. A server node in the distributed computing environment can receive a class definition from a client to execute, and generate and load into memory an instance of a class using said class definition without restarting or redeploying the server node. The class definition can define a new or updated class otherwise unavailable to the server node. Classes are identified with unique class identities which enables determination of whether a class is new or updated. The class identity can be used to determine the need for transmitting a class definition to a server node and also to ensure that a correct version of a class is implemented. In a particular case the new or updated class definition implements a lambda expression.



SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, **Published:**
GW, KM, ML, MR, NE, SN, TD, TG).

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

SYSTEM AND METHOD FOR SUPPORTING DYNAMIC DEPLOYMENT OF EXECUTABLE CODE IN A DISTRIBUTED COMPUTING ENVIRONMENT

5 Copyright Notice:

[0001] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights 10 whatsoever.

Field of Invention:

[0002] The present invention is generally related to computer systems, and is particularly related to a distributed computing environment.

15

Summary:

[0003] The present disclosure describes systems and methods that can support dynamic deployment of executable code in a distributed computing environment. A client generates a class definition for one or more class or updated class. The class definition is 20 serialized and transmitted to a server in the distributed computing environment. The server in the distributed computing environment receives the serialized class definition, deserializes the class definition and generates the one or more class or updated class on the server for execution. The class or updated class may represent a lambda expression. Furthermore, the class or updated class is provided with a unique identity to ensure that a correct class or class 25 version is executed and that a class definition that is unavailable on the server is serialized and distributed. Dynamic deployment of executable code enables execution of an agent on the server node where the agent depends on a new class definition otherwise unavailable to the server node.

[0004] In an embodiment, the present disclosure describes a method for supporting 30 lambda expressions in a distributed computing environment. The method includes receiving, via a server, a class definition implementing a lambda expression. The client generates said class definition implementing a lambda expression, serializes it and sends it to the server. The server defines a lambda class based on the received class definition. If the lambda expression is updated on the client side, a new class definition representing the updated lambda 35 expression is sent to the server.

[0005] These and other objects and advantages of the present invention will become apparent to those skilled in the art from the following description of the various embodiments, when read in light of the accompanying drawings.

Brief Description of the Figures:

[0006] **Figure 1** illustrates a distributed data grid, in accordance with an embodiment of the invention.

[0007] **Figure 2A** illustrates a distributed computing environment supporting dynamic deployment of executable code, in accordance with an embodiment of the invention.

[0008] **Figure 2B** illustrates a method supporting dynamic deployment of executable code, in accordance with an embodiment of the invention.

[0009] **Figure 3A** illustrates a distributed computing environment supporting dynamic deployment of executable code implementing lambda expressions, in accordance with an embodiment of the invention.

[00010] **Figure 3B** illustrates a method supporting dynamic deployment of executable code implementing lambda expressions, in accordance with an embodiment of the invention.

Detailed Description:

[00011] Described herein are systems and methods which support dynamic deployment of executable code in a distributed computing environment. A server node in the distributed computing environment can receive a class definition from a client to execute and generate and load into memory an instance of a class using said class definition without restarting or redeploying the server node. The class definition can implement a new or updated class definition otherwise unavailable to the server node. Classes are identified with unique class identities which enable determination of whether a class is new or updated. The class identity can be used to determine the need for transmitting the class definition to a server node and also to ensure that a correct version of a class is implemented. Dynamic deployment of executable code enables execution of an agent on the server node where the agent depends on a new class definition otherwise unavailable to the server node.

[00012] In a particular embodiment, the present disclosure describes systems and methods that can support dynamic deployment of class definitions implementing lambda expressions in a distributed computing environment. A server in the distributed computing environment can receive from a client one or more class definition, wherein the client generates said one or more class definition to implement a lambda expression. Then, the server can define a lambda class on the server side based on the received class definition. Furthermore, when the lambda expression is updated on the client side a new class definition for the updated lambda class can be prepared and sent from the client to the server.

[00013] The systems and methods for supporting dynamic deployment of executable code including dynamic deployment of class definitions implementing lambdas as described herein and illustrated in **Figures 2A, 2B, 3A, and 3B** have particular utility in the distributed data grid described below with respect to **Figure 1**. The systems and methods for supporting

dynamic deployment of executable code including dynamic deployment of class definitions implementing lambdas may also be applied in wide variety of alternative distributed computing environments.

[00014] In the following description, the invention will be illustrated by way of example and

5 not by way of limitation in the figures of the accompanying drawings. References to various embodiments in this disclosure are not necessarily to the same embodiment, and such references mean at least one. While specific implementations are discussed, it is understood that this is provided for illustrative purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without departing from the 10 scope and spirit of the invention.

[00015] Furthermore, in certain instances, numerous specific details will be set forth to provide a thorough description of the invention. However, it will be apparent to those skilled in the art that the invention may be practiced without these specific details. In other instances, well-known features have not been described in as much detail so as not to obscure the 15 invention.

[00016] The present invention is described with the aid of functional building blocks illustrating the performance of specified functions and relationships thereof. The boundaries of these functional building blocks have often been arbitrarily defined herein for the convenience of the description. Thus functions shown to be performed by the same elements may in 20 alternative embodiments be performed by different elements. And functions shown to be performed in separate elements may instead be combined into one element. Alternate boundaries can be defined so long as the specified functions and relationships thereof are appropriately performed. Any such alternate boundaries are thus within the scope and spirit of the invention.

[00017] Common reference numerals are used to indicate like elements throughout the drawings and detailed description; therefore, reference numerals used in a figure may or may not be referenced in the detailed description specific to such figure if the element is described elsewhere. The first digit in a three digit reference numeral indicates the series of figures in which the element first appears.

30

Distributed Data Grid

[00018] A distributed data grid is a system in which a collection of computer servers work together in one or more clusters to manage information and related operations, such as computations, within a distributed or clustered environment. A distributed data grid can be used 35 to manage application objects and data that are shared across the servers. A distributed data grid provides low response time, high throughput, predictable scalability, continuous availability and information reliability. As a result of these capabilities, a distributed data grid is well suited

for use in computational intensive, stateful middle-tier applications. In particular examples, distributed data grids, such as e.g., the Oracle® Coherence data grid, store information in-memory to achieve higher performance, and employ redundancy in keeping copies of that information synchronized across multiple servers, thus ensuring resiliency of the system and continued availability of the data in the event of failure of a server.

5 **[00019]** In the following description, an Oracle® Coherence data grid having a partitioned cache is described. However, one of ordinary skill in the art will understand that the present invention, described for example in the summary above, can be applied to any distributed data grid known in the art without departing from the scope of the invention.

10 Moreover, although numerous specific details of an Oracle® Coherence distributed data grid are described to provide a thorough description of the invention, it will be apparent to those skilled in the art that the invention may be practiced in a distributed data grid without these specific details. Thus, a particular implementation of a distributed data grid embodying the present invention can, in some embodiments, exclude certain features, and/or include 15 different, or modified features than those of the distributed data grid described below, without departing from the scope of the invention.

15 **[00020]** **Figure 1** illustrates and example of a distributed data grid 100 which stores data and provides data access to clients 150. A "data grid cluster", or "distributed data grid", is a system comprising a plurality of computer servers (e.g., 120a, 120b, 120c, and 120d) which 20 work together in one or more cluster (e.g., 100a, 100b, 100c) to store and manage information and related operations, such as computations, within a distributed or clustered environment. While distributed data grid 100 is illustrated as comprising four servers 120a, 120b, 120c, 120d, with five data nodes 130a, 130b, 130c, 130d, and 130e in a cluster 100a, the distributed data 25 grid 100 may comprise any number of clusters and any number of servers and/or nodes in each cluster. The distributed data grid can store the information in-memory to achieve higher performance, and employ redundancy in keeping copies of that information synchronized across multiple servers, thus ensuring resiliency of the system and continued availability of the data in the event of server failure. In an embodiment, the distributed data grid 100 implements the present invention, described for example in the summary above and the detailed 30 description below.

35 **[00021]** As illustrated in **Figure 1**, a distributed data grid provides data storage and management capabilities by distributing data over a number of servers (e.g., 120a, 120b, 120c, and 120d) working together. Each server of the data grid cluster may be a conventional computer system such as, for example, a "commodity x86" server hardware platform with one to two processor sockets and two to four CPU cores per processor socket. Each server (e.g., 120a, 120b, 120c, and 120d) is configured with one or more CPU, Network Interface Card (NIC), and memory including, for example, a minimum of 4GB of RAM up to 64 GB of RAM or

more. Server 120a is illustrated as having CPU 122a, Memory 124a and NIC 126a (these elements are also present but not shown in the other Servers 120b, 120c, 120d). Optionally each server may also be provided with flash memory – e.g. SSD 128a – to provide spillover storage capacity. When provided the SSD capacity is preferably ten times the size of the RAM.

5 The servers (e.g., 120a, 120b, 120c, 120d) in a data grid cluster 100a are connected using high bandwidth NICs (e.g., PCI-X or PCIe) to a high-performance network switch 120 (for example, gigabit Ethernet or better).

[00022] A cluster 100a preferably contains a minimum of four physical servers to avoid the possibility of data loss during a failure, but a typical installation has many more servers

10 Failover and fallback are more efficient the more servers that are present in each cluster and the impact of a server failure on a cluster is lessened. To minimize communication time between servers, each data grid cluster is ideally confined to a single switch 102 which provides single hop communication between servers. A cluster may thus be limited by the number of ports on the switch 102. A typical cluster will therefore include between 4 and 96

15 physical servers.

[00023] In most Wide Area Network (WAN) configurations of a distributed data grid 100, each data center in the WAN has independent, but interconnected, data grid clusters (e.g., 100a, 100b, and 100c). A WAN may, for example, include many more clusters than shown in

Figure 1. Additionally, by using interconnected but independent clusters (e.g., 100a, 100b,

20 100c) and/or locating interconnected, but independent, clusters in data centers that are remote from one another, the distributed data grid can secure data and service to clients 150 against simultaneous loss of all servers in one cluster caused by a natural disaster, fire, flooding, extended power loss and the like. Clusters maintained throughout the enterprise and across geographies constitute an automatic ‘backup store’ and high availability service for enterprise

25 data.

[00024] One or more nodes (e.g., 130a, 130b, 130c, 130d and 130e) operate on each server (e.g., 120a, 120b, 120c, 120d) of a cluster 100a. In a distributed data grid the nodes may be for example, software applications, virtual machines, or the like and the servers may comprise an operating system, hypervisor or the like (not shown) on which the node operates.

30 In an Oracle® Coherence data grid, each node is Java virtual machine (JVM). A number of JVM/nodes may be provided on each server depending on the CPU processing power and memory available on the server. JVM/nodes may be added, started, stopped, and deleted as required by the distributed data grid. JVMs that run Oracle® Coherence automatically join and cluster when started. JVM/nodes that join a cluster are called cluster members or cluster

35 nodes.

[00025] In an Oracle® Coherence data grid cluster members communicate using Tangosol Cluster Management Protocol (TCMP). TCMP is an IP-based protocol that is used

to discover cluster members, manage the cluster, provision services, and transmit data between cluster members. The TCMP protocol provides fully reliable, in-order delivery of all messages. Since the underlying UDP/IP protocol does not provide for either reliable or in-order delivery, TCMP uses a queued, fully asynchronous ACK and NACK-based mechanism for reliable delivery of messages, with unique integral identity for guaranteed ordering of messages in queues associated with the JVMs operating on a server. The TCMP protocol requires only three UDP/IP sockets (one multicast, two unicast) and six threads per JVM/node, regardless of the cluster size.

[00026] The functionality of a data grid cluster is based on services provided by cluster nodes. Each service provided by a cluster node has a specific function. Each cluster node can participate in (be a member of) a number of cluster services, both in terms of providing and consuming the cluster services. Some cluster services are provided by all nodes in the cluster whereas other services are provided by only one or only some of the nodes in a cluster. Each service has a service name that uniquely identifies the service within the data grid cluster, and a service type, which defines what the service can do. There may be multiple named instances of each service type provided by nodes in the data grid cluster (other than the root cluster service). All services preferably provide failover and fallback without any data loss.

[00027] Each service instance provided by a cluster node typically uses one service thread to provide the specific functionality of the service. For example, a distributed cache service provided by a node is provided by single service thread of the node. When the schema definition for the distributed cache is parsed in the JVM/node, a service thread is instantiated with the name specified in the schema. This service thread manages the data in the cache created using the schema definition. Some services optionally support a thread pool of worker threads that can be configured to provide the service thread with additional processing resources. The service thread cooperates with the worker threads in the thread pool to provide the specific functionality of the service.

[00028] In an Oracle® Coherence data grid, the cluster service (e.g., 136a, 136b, 136c, 136d, 136e) keeps track of the membership and services in the cluster. Each cluster node always has exactly one service of this type running. The cluster service is automatically started to enable a cluster node to join the cluster. The cluster service is responsible for the detection of other cluster nodes, for detecting the failure (death) of a cluster node, and for registering the availability of other services in the cluster. The proxy service (e.g., 138c) allows connections (e.g. using TCP) from clients that run outside the cluster. The invocation Service (e.g., 134d) allows application code to invoke agents to perform operations on any node in the cluster, or any group of nodes, or across the entire cluster. Although shown on only one node each, the invocation service and proxy service can be configured on any number up to all of the nodes of the distributed data grid.

[00029] In an Oracle® Coherence data grid, the distributed cache service (e.g., 132a, 132b, 132c, 132d, 132e) is the service which provides for data storage in the distributed data grid and is operative on all nodes of the cluster that read/write/store cache data, even if the node is storage disabled. The distributed cache service allows cluster nodes to distribute (partition) data across the cluster 100a so that each piece of data in the cache is managed primarily (held) by only one cluster node. The distributed cache service handles storage operation requests such as put, get, etc. The distributed cache service manages distributed caches (e.g., 140a, 140b, 140c, 140d, 140e) defined in a distributed schema definition and partitioned among the nodes of a cluster.

5 **[00030]** A partition is the basic unit of managed data in the distributed data grid and stored in the distributed caches (e.g., 140a, 140b, 140c, 140d, and 140e). The data is logically divided into primary partitions (e.g., 142a, 142b, 142c, 142d, and 142e), that are distributed across multiple cluster nodes such that exactly one node in the cluster is responsible for each piece of data in the cache. Each cache (e.g., 140a, 140b, 140c, 140d, and 140e) can hold a

10 number of partitions. Each partition (e.g., 142a, 142b, 142c, 142d, 142e) may hold one datum or it may hold many. A partition can be migrated from the cache of one node to the cache of another node when necessary or desirable. For example, when nodes are added to the cluster, the partitions are migrated so that they are distributed among the available nodes including newly added nodes. In a non-replicated distributed data grid there is only one active copy of

15 each partition (the primary partition). However, there is typically also one or more replica/backup copy of each partition (stored on a different server) which is used for failover. Because the data is spread out in partition distributed among the servers of the cluster, the responsibility for managing and providing access to the data is automatically load-balanced across the cluster.

20 **[00031]** The distributed cache service can be configured so that each piece of data is backed up by one or more other cluster nodes to support failover without any data loss. For example, as shown in **Figure 1**, each partition is stored in a primary partition (e.g., dark shaded squares 142a, 142b, 142c, 142d, and 142e) and one or more synchronized backup copy of the partition (e.g., light shaded squares 144a, 144b, 144c, 144d, and 144e). The backup copy

25 of each partition is stored on a separate server/node than the primary partition with which it is synchronized. Failover of a distributed cache service on a node involves promoting the backup copy of the partition to be the primary partition. When a server/node fails, all remaining cluster nodes determine what backup partitions they hold for primary partitions on failed node. The cluster nodes then promote the backup partitions to primary partitions on whatever cluster node they are held (new backup partitions are then created).

30 **[00032]** A distributed cache is a collection of data objects. Each data object/datum can be, for example, the equivalent of a row of a database table. Each datum is associated with a

unique key which identifies the datum. Each partition (e.g., 142a, 142b, 142c, 142d, 142e) may hold one datum or it may hold many and the partitions are distributed among all the nodes of the cluster. In an Oracle® Coherence data grid each key and each datum is stored as a data object serialized in an efficient uncompressed binary encoding called Portable Object Format (POF).

5 [00033] In order to find a particular datum, each node has a map, for example a hash map, which maps keys to partitions. The map is known to all nodes in the cluster and is synchronized and updated across all nodes of the cluster. Each partition has a backing map which maps each key associated with the partition to the corresponding datum stored in the 10 partition. An operation associated with a particular key/datum can be received from a client at any node in the distributed data grid. When the node receives the operation, the node can provide direct access to the value/object associated with the key, if the key is associated with a primary partition on the receiving node. If the key is not associated with a primary partition on the receiving node, the node can direct the operation directly to the node holding the primary 15 partition associated with the key (in one hop). Thus, using the hash map and the partition maps, each node can provide direct or one-hop access to every datum corresponding to every key in the distributed cache.

20 [00034] In some applications, data in the distributed cache is initially populated from a database 110 comprising data 112. The data 112 in database 110 is serialized, partitioned and distributed among the nodes of the distributed data grid. Distributed data grid 100 stores data 25 objects created from data 112 from database 110 in partitions in the memory of servers 120a, 120b, 120c, 120d such that clients 150 and/or applications in data grid 100 can access those data objects directly from memory. Reading from and writing to the data objects in the distributed data grid 100 is much faster and allows more simultaneous connections than could be achieved using the database 110 directly. In-memory replication of data and guaranteed 30 data consistency make the distributed data grid suitable for managing transactions in memory until they are persisted to an external data source such as database 110 for archiving and reporting. If changes are made to the data objects in memory the changes are synchronized between primary and backup partitions and may subsequently be written back to database 110 using asynchronous writes (write behind) to avoid bottlenecks.

35 [00035] Although the data is spread out across cluster nodes, a client 150 can connect to any cluster node and retrieve any datum. This is called location transparency, which means that the developer does not have to code based on the topology of the cache. In some embodiments, a client might connect to a particular service e.g., a proxy service on a particular node. In other embodiments, a connection pool or load balancer may be used to direct a client to a particular node and ensure that client connections are distributed over some or all the data nodes. However connected, a receiving node in the distributed data grid receives tasks from a

client 150, and each task is associated with a particular datum, and must therefore be handled by a particular node. Whichever node receives a task (e.g. a call directed to the cache service) for a particular datum identifies the partition in which the datum is stored and the node responsible for that partition, the receiving node, then directs the task to the node holding the requested partition for example by making a remote cache call. Since each piece of data is managed by only one cluster node, an access over the network is only a "single hop" operation. This type of access is extremely scalable, since it can use point-to-point communication and thus take optimal advantage of a switched fabric network such as InfiniBand.

[00036] Similarly, a cache update operation can use the same single-hop point-to-point approach with the data being sent both to the node with the primary partition and the node with the backup copy of the partition. Modifications to the cache are not considered complete until all backups have acknowledged receipt, which guarantees that data consistency is maintained, and that no data is lost if a cluster node were to unexpectedly fail during a write operation. The distributed cache service also allows certain cluster nodes to be configured to store data, and others to be configured to not store data.

[00037] In some embodiments, a distributed data grid is optionally configured with an elastic data feature which makes use of solid state devices (e.g. SSD 128a), most typically flash drives, to provide spillover capacity for a cache. Using the elastic data feature a cache is specified to use a backing map based on a RAM or DISK journal. Journals provide a mechanism for storing object state changes. Each datum/value is recorded with reference to a specific key and in-memory trees are used to store a pointer to the datum (a tiny datum/value may be stored directly in the tree). This allows some values (data) to be stored in solid state devices (e.g. SSD 128a) while having the index/memory tree stored in memory (e.g. RAM 124a). The elastic data feature allows the distributed data grid to support larger amounts of data per node with little loss in performance compared to completely RAM-based solutions.

[00038] A distributed data grid such as the Oracle® Coherence data grid described above can improve system performance by solving data operation latency problems and by caching and processing data in real time. Applications cache data in the data grid, avoiding expensive requests to back-end data sources. The shared data cache provides a single, consistent view of cached data. Reading from the cache is faster than querying back-end data sources and scales naturally with the application tier. In memory performance alleviates bottlenecks and reduces data contention, improving application responsiveness. Parallel query and computation is supported to improve performance for data-based calculations. The distributed data grid is fault-tolerant, providing for data reliability, accuracy, consistency, high availability, and disaster recovery. The distributed data grid enables applications to scale linearly and dynamically for predictable cost and improved resource utilization. For many applications, a distributed data grid offers a valuable shared data source solution.

Code Execution In A Distributed Data Grid

[00039] In embodiments, the distributed data grid supports the use of agents to perform operations on the nodes of the distributed data grid. For partitioned data, the agent can be configured to execute on the node (or nodes) that owns the data to execute against. Queuing, 5 concurrency management, agent execution, data access by the agent and data modification by the agent all occur on that grid node. (Only the synchronous backup of the resultant data modification, if any, requires additional network traffic.) For many processing purposes, it is much more efficient to provide the agent to the node than to handle distributed concurrency 10 control, coherency and data updates. In other words, the Data Grid determines the location to execute the agent based on the configuration for the data topology, moves the agent there, executes the agent (automatically handling concurrency control for the item while executing the agent), backs up the modifications if any, and returns a result.

[00040] An entry processor is one example of an agent used to perform an operation on the nodes of the distributed data grid. Other examples of agents/executable code include, 15 but are not limited to: aggregators, filters, value extractors, and invocables. For example an Oracle® Coherence data grid supports a lock-free programming model through the EntryProcessor API. Advantageously, an entry processor performs an implicit low-level lock on the entries it is processing, thereby allowing the client to place processing code in an entry processor without having to worry about concurrency control. An entry processor can be 20 invoked on a specific key, on a collection of keys, or on a Filter (the Filter is executed against the cache entries and each entry that matches the Filter criteria has the entry processor executed against it). Entry processors are executed in parallel across the cluster (on the nodes that own the individual entries.) This provides a significant advantage over having a client lock all affected keys, pull all required data from the cache, process the data, place the data back 25 in the cache, and unlock the keys. The processing occurs in parallel across multiple computers (as opposed to serially on one computer) and the network overhead of obtaining and releasing locks is eliminated. The code of the entry processor for execution against the chosen entries is transmitted into the distributed data grid. For many transaction types, the use of entry processors minimizes contention and latency and improves system throughput, without 30 compromising the fault-tolerance of data operations.

[00041] For example, an agent can be created on a client 150 of the distributed data grid 100 shown in Figure 1. In order to execute the agent, the agent is serialized and transmitted to one or more node in the distributed data grid 100. To serialize an object means to convert its state to a byte stream so that the byte stream can be reverted back (deserialized) 35 into a copy of the object. Deserialization is the process of converting the serialized form of an object back into a copy of the object. The Java platform specifies a default way by which serializable objects are serialized. A (Java) class can override this default serialization and

define its own way of serializing objects of that class. When an object is serialized, information that identifies its class is recorded in the serialized stream. However, the class's definition ("class file") itself is typically NOT recorded. It is typically the responsibility of the system (for example the server) that is deserializing the object to determine how to locate and load the necessary class files.

[00042] A problem may arise where the agent utilizes a class definition ("class file") which is not available on a node of the distributed data grid upon which it is desired to realize the agent. In this situation, the agent will fail to operate on the node because the necessary class cannot be found by the node on which the agent is attempting to operate. Another problem may arise where the class definition is different on the client than it is on the node. For example, a class definition may have been modified/updated on the client but not modified/updated on the nodes of the distributed data grid. In this situation, the agent may operate because the necessary class is found on the node, however, the agent may not operate as intended because it utilizes a different class definition than intended by the creator of the agent. In prior embodiments, where new or updated class files were deployed on a client, it was necessary redeploy and restart nodes in the distributed data grid to support those new or updated class files before they could be utilized by agents operating in the distributed data grid. This is very cumbersome and time consuming for large clusters.

[00043] In embodiments, the distributed data grid 100, overcomes these problems by implementing a system and method for supporting dynamic deployment of new or updated class definitions from the client 150 to nodes of the distributed data grid 100 for execution in agents. Additionally, the distributed data grid 100 implements version control for class definitions to ensure that a class called by an agent (such as an entry processor) has the same class definition as the class on the client 150 where the agent was prepared. Utilizing the system and method for supporting dynamic deployment of class definitions and version control for class definitions increases the utility of the distributed data grid because it allows new or updated class definitions to be correctly be executed on nodes of the distributed data grid. Furthermore, avoiding the need to redeploy and restart nodes in the distributed data grid to support new or updated class definitions improves performance of the distributed data grid.

Dynamic Deployment of Class Definitions

[00044] **Figure 2A**, illustrates a system for dynamic deployment of a new class definition from the client 150 to a node 130a of the distributed data grid to enable an agent which utilizes the new class definition to be realized on the node. Although only one client and one node are shown, the same functionality can be applied on all clients and nodes of the distributed data grid. As shown in Figure 2A, node 130a is implemented as a Java Virtual Machine (JVM). For example, an Oracle® JRockit JVM, IBM® JVM, Oracle® JVM or another vendor

implementation of JVM. Node 130a, includes execution engine 210, native method interface 212, native method libraries 214, and class loader 217. Node 130a has a JVM memory area 220, which includes method area 212, heap 214, JVM language stacks 216, PC registers 218, and native method stacks 230. Upon startup of node 130a, class loader 217 loads and 5 prepares and initializes classes based on class definitions found in the class path for the JVM. Client 150 is provided with remote class support 262 and node 130a is provided with remote class support 232. The remote class support functionality allows for dynamic deployment of new class definitions from the client 150 to a node 130a of the distributed data grid using a remote class constructor 264 transmitted over-the-wire from the client 150 to the node 130a 10 as described below.

[00045] As shown in Figure 2A, a user of a client system 150 can prepare agent code 250 for an agent to be operated on node 130a. Agent code can include a new class definition 252, which is either a definition of a new class having a new class name or an updated definition for a class (i.e. the class is preexisting). Note that even where the class definition is 15 changed it is technically a new class with a new class definition from the client's perspective. The preexisting class on the client cannot be modified – instead a new class including any changes is implemented. When agent code 250 is attempted to be sent for execution on node 130a, remote class support 262 examines the agent code 250 to determine whether it includes any new class definition 252. Where one or more new class definition 252 is found, remote 20 class support 262 prepares remote class constructor 264 for new class definition 252. Remote class constructor 264 includes new class definition 266 and the captured agent state 268 required to construct an instance of the agent on the node 130a. The remote class constructor 264 is used for wire transport for the new class and enables automatic serialization of the new class definition for transmission over-the-wire to node 130a.

[00046] Remote class support 262 utilizes a class/ID table 260. Class/ID table 260 includes a class identifier for each preexisting class based on the class name and a hash function (for example an MD5 hash) of the class definition. This ensures that the identifier for each class is unique. Moreover, where a class definition is updated, the hash ensures that a new identifier will be assigned and thus the updated class will have a new class identity. Thus, 30 using the class/ID table 260, remote class support 262 can determine whether agent code 250 includes any new class definition 252.

[00047] Remote class support 262 thus determines, using the class/ID table 260, whether any new class definition 252 is utilized in agent code 250. If all classes are preexisting on the node, remote class support just returns the identity of the class (which can be executed on node 130a). However, if a new class definition 252 is utilized by the agent code 250, (i.e. the class is unavailable on the node), then a new class definition 252 is prepared and a new 35 identity corresponding to the new class definition 252 is stored in Class/ID table 260. A remote

class constructor 264 is prepared representing the new class definition 266 and the captured agent state 268 required to realize the agent on the node 130a. The remote class constructor instance is automatically serialized and transmitted over-the-wire to node 130a.

[00048] The remote class constructor 264 is received by the node 130a and processed

5 by remote class support 232. Remote class constructor 264 is automatically deserialized when it is received over-the-wire from client 150. Node 130a also maintains its own Class/ID Table 260 for validating class identities. Using the deserialized remote class constructor 264 and the remote class support 232, node 130a prepares a class loader for new class definition 266, loads the class into JVM memory 220, and initializes an instance of the new class in the JVM.

10 Node 130a, is thus enabled to execute code utilizing the class instance based on the new class definition. Remote class support 232 is thus enabled to realize agent 238 for execution on node 130a.

[00049] **Figure 2B** illustrates a method 270 for dynamic deployment of new or updated class definitions from the client 150 to a node 130a of the distributed data grid. At step 272,

15 the client attempts to send an agent for execution on a node. At step 274, the client determines an identity for the class. The identity is unique and thus the client can determine at step 276 whether the class already exists by comparing the identity to identities previously stored in the Class/ID table. If the identity does not match a class in the Class/ID table, the class must be a new or updated class.

20 **[00050]** At step 278, if the class is a new or updated class (identity not present in the Class/ID table), the client generates a new class definition for the class. A remote constructor is prepared on the client. The remote constructor includes the new class definition and the captured state of the agent required to realize the agent on the node. The remote class constructor instance is automatically serialized and transmitted over-the-wire to the node in step 280.

25 **[00051]** At step 282 the node receives the remote constructor over-the-wire from the client. At step 284, the remote constructor is deserialized and executed on the node. At step 286, the new class definition and the captured agent state required to realize the agent is loaded into JVM memory to create an instance of the new or updated class on the node. At step 288, the instance of the new or updated class is made available for execution on the node in order to realize the agent on the node. Note that where the class definition is changed on the client it is technically a new class with a new class definition from the client's perspective. The preexisting class on the node cannot be modified – instead a new class (having a new class identity) including the changes is implemented. At step 290, the agent is executed on the node.

35 **[00052]** Accordingly, the method allows dynamic deployment of a new or updated class definition from the client to the node in order to permit implementation of such new or updated

class without requiring redeployment of the node/JVM. This dynamic deployment of new or updated classes without requiring redeployment of the node/JVM enhances the utility of the distributed data grid to execute agents and eliminates the need to restart nodes when new or updated class definitions are required.

5

Dynamic Deployment Of Class Definitions Implementing Lambdas

[00053] In an embodiment, the system and method for dynamic deployment of class definitions can be used for class definitions associated with lambda expressions. In recent years, object-oriented programming languages such as Java™ have become increasingly popular for a wide variety of applications running on many different types of platforms. Some programming languages may support “lambda expressions”. In object-oriented programming languages, lambda expressions may be used as a lightweight mechanism for defining objects that have just one method and no state. The lightweight syntax for the lambda expression can be optimized for the common usage in various programming language environments. The Java™ programming language is an example of a lexically-scoped imperative programming language. Support for lambda expressions was introduced in the Java™ 8 programming language. Accordingly code for agents to be executed on nodes of a distributed data grid may in some cases include lambda expressions.

[00054] Lambda expressions are a lightweight way to describe functions that can be treated as values by a programming language. Lambda expressions can be anonymous methods, which are aimed at addressing the “vertical problem”, or bulky syntax, by replacing the machinery of anonymous inner classes with a syntactically lighter-weight mechanism. Lambda expressions can be viewed as providing support for “code as data” in that methods can accept lambdas (functions) as arguments and return them as a result. A lambda expression is like a method: it provides a list of formal parameters and a body—an expression or block—expressed in terms of those parameters. Lambda expressions may be expressed using syntax such as the following:

(argument list)→body.

[00055] Aspects of the use of lambda expressions in the Java™ programming language are described, for example in the following U.S. Patent Applications which are incorporated herein by reference: U.S. Patent Application Ser. No. 13/779,536 entitled “SYSTEM AND METHOD FOR SUPPORTING COMPATIBILITY CHECKING FOR LAMBDA EXPRESSION” filed February 27, 2014; U.S. Patent Application Ser. No. 13/677,001 entitled “TARGET TYPING OF OVERLOADED METHOD AND CONSTRUCTOR ARGUMENTS” filed November 14, 2012; U.S. Patent Application Ser. No. 13/677,000 entitled “SYSTEM AND METHOD FOR INFERRING IMMUTABILITY OF PROGRAM VARIABLES” filed November 14, 2012; U.S.

Patent Application Ser. No. 13/779557 entitled "SYSTEM AND METHOD FOR COMPILED LAMBDA EXPRESSION IN A PROGRAMMING LANGUAGE ENVIRONMENT" filed February 27, 2013; and U.S. Patent Application Ser. No. 13/028111 entitled "BINDING-BASED CHARACTERIZATION OF CLOSURES IN SOFTWARE PROGRAMS" filed February 15, 5 2011.

[00056] Lambda expressions are "anonymous" in that they are not named. Thus, Lambda expressions enable definition of anonymous methods, and enable treating such anonymous methods as instances of functional interfaces. The use of functional interfaces with anonymous inner classes is a common pattern in the Java™ programming language, thus the 10 Java™ programming language leverages functional interfaces to implement lambda expressions. In the Java™ 8 programming language, a lambda expression can be described as a stripped-down form of the anonymous class declaration in which the parameter of the lambda expression is the same as the parameter to the anonymous class declaration's method, and the body of the lambda expression is the same as the body of the anonymous 15 class declaration's method (but is able to elide the 'return' keyword because the only thing this particular lambda expression's body does is return a value). The expression-body form eliminates the need for a return keyword, which could otherwise represent a substantial syntactic overhead relative to the size of the expression. Thus, in the Java™ 8 programming language, lambda expressions are implemented in terms of class definitions for functional 20 interfaces.

[00057] Described herein are systems and methods that can support dynamic deployment of classes implementing lambdas in a distributed computing environment including, for example, the distributed data grid of Figure 1. As described above, the distributed data grid supports performance of transaction using an agent operating on the nodes of the 25 distributed data grid. The code for such agents may include lambda expressions. However, in prior implementations, classes used in an agent had to be available (pre-existing) in the class path for each cluster node – this precluded the use of lambdas in agents unless the class implementing the lambda had previously been deployed to the nodes of the distributed data grid (or the nodes were redeployed). For the server to actually process the lambda expression 30 required the same compiled code in both the client and server class path. That is the server must have a class corresponding to the lambda expression in the server class path. Any changes or introduction of new lambdas on the client required redeployment and restart of both the client and the server in order to deploy a new class corresponding to the lambda expression to the server. This is very cumbersome and time consuming for large clusters.

[00058] In order to overcome this limitation, in embodiments of the present invention, the client transmits not only the metadata describing the lambda expression for use in an agent, but also, the actual class definition for the class implementing the lambda expression. Dynamic

deployment of classes implementing lambda expressions can be performed as a particular case of the dynamic deployment of class definition described above. At the server, the class definition received from the client is parsed and a new lambda class generated from it for use in the agent/entry processor. The server defines a lambda class based on the class definition received from the client and executes it. This allows modification of the existing lambda expressions or the introduction of new lambda expressions without the need for redeployment or server restart to place the compiled code for the lambda expression in the class path.

5 [00059] **Figure 3A** shows an illustration of a system supporting dynamic lambdas in a distributed computing environment, in accordance with an embodiment of the invention. As shown in Figure 3, the distributed computing environment 300, such as a distributed data grid, includes a client 301 and a server 302. In an embodiment, the distributed computing environment may be a distributed data grid 100 as shown in Figure 1, in which client 301 represents one or more of clients 150 and server 302 represents one or more of nodes 130a, 130b, 130c, 130d, and 130e.

10 [00060] As described above, in prior embodiments without dynamic deployment of class definitions, only the metadata 305, which describes the lambda expression 303, was sent over from the client 301 to the server 302. Then, the server 302 can take advantage of the lambda expression 303 on the client side based on the received metadata 306. It was the responsibility of the server 302 deserializing the received metadata 306 to determine how to locate and load 15 the necessary class files for the static lambda. The class files therefore need to have already been deployed on the server 302 and also be the same as on the client 301 (i.e. not changed/updated on the client). Thus, any changes or introduction of new lambdas on the client side require a redeployment and/or a restart of both the client 301 and the server 302 before they could be invoked on the server 302 by the client.

20 [00061] On the other hand, in embodiments of the present invention which enable the use of dynamic deployment of class definitions, when the new or updated class definitions are used to implement lambdas, the client 301 generates one or more class definition 307 to execute 25 from the lambda expression 303, and sends the class definition 307 to the server 302 in the distributed computing environment 300. After the server 302 receives the class definition 308 to execute, the server 302 defines a lambda instance 304 on the server side, based on the received class definition 308, and then executes the lambda instance 304 defined by the received class definition 308. Additionally, if the client 301 updates the lambda expression 303 on the client side, e.g. when the lambda expression 303 is modified or is introduced with a new behavior, a new class definition 307 can be prepared and transmitted to the server 302 and a 30 new lambda instance can be created and executed having the new behavior.

35 [00062] Figure 3B illustrates a method for supporting dynamic lambdas in a distributed computing environment. At step 320, the client generates said class definition from a lambda

expression. At step 322, the server receives the class definition to execute from a client. At step 324, the server defines and loads a lambda class on the server-side based on the received class definition. At step 326, the server creates and executes the lambda instance defined on the server-side based on the received class definition. If the lambda expression is updated on the client side the process is repeated from steps 320 to 326 using a new class definition based on the updated lambda expression and resulting in a new lambda class definition and instance on the server side.

[00063] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation.

10 It will be apparent to persons skilled in the relevant art that various changes in form and detail can be made therein without departing from the spirit and scope of the invention.

[00064] It is also noted that, while for purposes of illustration, the above description is provided largely in the context of using the Java™ programming language (or a language closely related to Java™), the techniques described may be used for any programming language that 15 supports Lambda expressions, overloading of invocation operations such as methods, constructors, or functions. The techniques may be used for object-oriented languages in some embodiments, or for non-object-oriented languages (such as various procedural languages) in other embodiments.

[00065] Many features of the present invention can be performed in, using, or with the 20 assistance of hardware, software, firmware, or combinations thereof. The present invention may be conveniently implemented using one or more conventional general purpose or specialized digital computer, computing device, machine, or microprocessor, including one or more processors, memory and/or computer readable storage media programmed according to the teachings of the present disclosure. Features of the invention may also be implemented in 25 hardware using, for example, hardware components such as application specific integrated circuits (ASICs) and programmable logic device. Implementation of the hardware state machine so as to perform the functions described herein will be apparent to persons skilled in the relevant art.

[00066] Features of the present invention can be incorporated in software and/or firmware 30 for controlling the hardware of a processing system, and for enabling a processing system to interact with other mechanisms utilizing the results of the present invention. Such software or firmware may include, but is not limited to, application code, device drivers, operating systems and execution environments/containers. Appropriate software coding can readily be prepared by skilled programmers based on the teachings of the present disclosure, as will be apparent 35 to those skilled in the software art.

[00067] In some embodiments, the present invention includes a computer program product which is a storage medium or computer readable medium (media) having instructions stored

thereon/in which can be used to program a computer to perform any of the processes of the present invention. The storage medium or computer readable medium can include, but is not limited to, any type of disk including floppy disks, optical discs, DVD, CD-ROMs, microdrive, and magneto-optical disks, ROMs, RAMs, EPROMs, EEPROMs, DRAMs, VRAMs, flash 5 memory devices, magnetic or optical cards, nanosystems (including molecular memory ICs), or any type of media or device suitable for storing instructions and/or data. In embodiments, the storage medium or computer readable medium can be non-transitory.

10 [00068] The foregoing description of the present invention has been provided for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations will be apparent to the practitioner skilled in the art. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, thereby enabling others skilled in the art to understand the invention for various embodiments and with various modifications that are suited to the particular use contemplated. It is intended that the scope 15 of the invention be defined by the following claims and their equivalents.

Claims:

What is claimed is:

- 5 1. A method for supporting dynamic deployment of executable code from a client to a node in a distributed computing environment, the method comprising:
 - receiving at the node a serialized object from the client, wherein the serialized object comprises a class definition for a new class not present in the class path of the node and a captured state of an agent utilizing said new class;
- 10 10. deserializing the serialized object on the node to retrieve said class definition and said captured state of the agent;
 - loading the class definition and captured state of the agent in memory of the node to generate an instance of the new class on the node; and
 - executing the agent on the node.
- 15 2. The method of claim 1, wherein said class definition implements a lambda expression.
3. The method of any of claims 1 to 2, further comprising:
 - receiving agent code on the client;
- 20 20. determining the class definition for said new class from said agent code on the client;
 - and
 - generating a class identity for said new class based on the class definition.
4. The method of any of claims 1 to 3, further comprising:
 - 25 maintaining a class/identity table on said client;
 - receiving agent code on the client for said agent;
 - determining the class definition for said new class from said agent code on the client;
 - generating a class identity for said new class based on the class definition; and
 - using the class/identity table to determine that said new class is not present in the class
- 30 30. path of the node.
5. The method of any of claims 1 to 3, further comprising:
 - maintaining a class/identity table on said client;
 - receiving agent code on the client for said agent;
- 35 35. determining the class definition for said new class from said agent code on the client;
- generating a class identity for said new class based on the class definition;
- using the class/identity table to determine that said new class is not present in the class

path of the node; and

adding the class identity for said new class to the class/identity table.

6. The method of any of claims 1 to 5, further comprising:

5 receiving agent code on the client for said agent;
determining the class definition for said new class from said agent code on the client;
generating on said client the serialized object comprising a class definition for a new
class not present in the class path of the node and the captured state of an agent utilizing said
new class; and

10 transmitting said serialized object to the node.

7. The method of any of claims 1 to 6, further comprising:

receiving agent code on the client for said agent, wherein said agent code includes a
lambda expression;

15 determining the class definition for said new class from said agent code on the client,
wherein the class definition implements said lambda expression;
generating on said client the serialized object comprising a class definition for a new
class not present in the class path of the node and the captured state of an agent utilizing said
new class; and

20 transmitting said serialized object to the node.

8. The method of any of claims 1 to 7, wherein the distributed computing environment
comprises a data grid cluster and wherein dynamic deployment of executable code enables
execution of the agent on a plurality of nodes in said data grid cluster holding data to be
25 operated on by said agent without redeployment of said plurality of nodes in said data grid
cluster.

9. The method of any of claims 1 to 8, wherein the distributed computing environment
comprises a data grid cluster and wherein said agent is an entry processor for execution
30 against data held in said node.

10. The method of any of claims 1 to 9, wherein the distributed computing environment
comprises a data grid cluster and wherein said agent is one of an aggregator, a filter, a value
extractor, and an invocable, for execution against data held in said node.

35

11. A system for supporting dynamic deployment of executable code in a distributed
computing environment, the system comprising:

a computer system comprising a microprocessor and a memory;
 a node deployed on said computer system as part of said distributed computing environment; and
 wherein said node is configured to,
5 receive a serialized object from a client, wherein the serialized object comprises a class definition for a new class not present in the class path of the node and a captured state of an agent utilizing said new class,
 deserialize the serialized object on the node to retrieve said class definition and said captured state of the agent,
10 load the class definition and captured state of the agent in a portion of said memory allocated to said node to generate an instance of the new class on the node, and
 execute the agent on the node.

15 12. The system of claim 11, wherein said class definition implements a lambda expression.

13. The system of any of claims 11 to 12, further comprising:
 a client computer system comprising a microprocessor and a memory; and
 a client operating on said client computer system, wherein said client is configured to,
20 receive agent code for said agent,
 determine the class definition for said new class from said agent code; and
 generate a class identity for said new class based on the class definition.

14. The system of any of claims 11 to 13, further comprising:
25 a client computer system comprising a microprocessor and a memory; and
 a client operating on said client computer system, wherein said client is configured to,
 maintain a class/identity table,
 receive agent code for said agent,
 determine the class definition for said new class from said agent code,
30 generate a class identity for said new class based on the class definition, and
 use the class/identity table to determine that said new class is not present in the class path of the node.

15. The system of any of claims 11 to 13, further comprising:
35 a client computer system comprising a microprocessor and a memory; and
 a client operating on said client computer system, wherein said client is configured to,
 maintain a class/identity table,

receive agent code for said agent,
determine the class definition for said new class from said agent code,
generate a class identity for said new class based on the class definition,
use the class/identity table to determine that said new class is not present in
5 the class path of the node, and
add the class identity for said new class to the class/identity table.

16. The system of any of claims 11 to 15, further comprising:
a client computer system comprising a microprocessor and a memory; and
10 a client operating on said client computer system, wherein said client is configured to,
maintain a class/identity table,
receive agent code for said agent,
determine the class definition for said new class from said agent code,
generate the serialized object comprising a class definition for a new class not
15 present in the class path of the node and the captured state of an agent utilizing said
new class, and
transmit said serialized object to the node.

17. The system of any of claims 11 to 16, further comprising:
20 a client computer system comprising a microprocessor and a memory; and
a client operating on said client computer system, wherein said client is configured to,
maintain a class/identity table,
receive agent code for said agent wherein said agent code includes a lambda
expression,
25 determine the class definition for said new class from said agent code wherein
the class definition implements said lambda expression,
generate the serialized object comprising a class definition for a new class not
present in the class path of the node and the captured state of an agent utilizing said
new class, and
30 transmit said serialized object to the node.

18. The system of any of claims 11 to 17, wherein the distributed computing environment
comprises a data grid cluster and wherein dynamic deployment of executable code enables
execution of the agent on a plurality of nodes in said data grid cluster holding data to be
35 operated on by said agent without redeployment of said plurality of nodes in said data grid
cluster.

19. The system of any of claims 11 to 18, wherein the distributed computing environment comprises a data grid cluster and wherein said agent is one of an entry processor, an aggregator, a filter, a value extractor, and an invocable, for execution against data held in said node.

5

20. A non-transitory computer readable medium including instruction stored thereon for supporting dynamic deployment of executable code from a client to a node in a distributed computing environment, which instructions, when executed cause the node to perform steps comprising:

10 receiving at the node a serialized object from the client, wherein the serialized object comprises a class definition for a new class not present in the class path of the node and a captured state of an agent utilizing said new class;

deserializing the serialized object on the node to retrieve said class definition and said captured state of the agent;

15 loading the class definition and captured state of the agent in memory of the node to generate an instance of the new class on the node; and
executing the agent on the node.

21. A computer program comprising program instructions for execution on one or more computer systems, wherein the program instructions when executed causes the one or more computer systems to perform the method of any of claims 1 to 10.

20 22. A computer program product comprising the computer program of claim 21 stored in a non-transitory machine readable data storage medium.

25


FIGURE 1


FIGURE 2A


FIGURE 2B



FIGURE 3A

FIGURE 3B

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/052060

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F9/445 G06F9/50
ADD.

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	AL-JAROOLDI J ET AL: "Middleware infrastructure for parallel and distributed programming models in heterogeneous systems", IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, IEEE SERVICE CENTER, LOS ALAMITOS, CA, US, vol. 14, no. 11, 1 November 2003 (2003-11-01), pages 1100-1111, XP011103739, ISSN: 1045-9219, DOI: 10.1109/TPDS.2003.1247671	1,11, 20-22
A	the whole document	2-10, 12-19
A	----- US 2009/144714 A1 (FAN NAN [CN] ET AL) 4 June 2009 (2009-06-04) the whole document	1-22
	----- -/-	

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
23 December 2015	12/01/2016
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Breche, Philippe

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/052060

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	Laurent Baduel ET AL: "Programming, Composing, Deploying for the Grid" In: "Grid Computing: Software Environments and Tools", 1 January 2006 (2006-01-01), Springer London, London, XP055238382, ISBN: 978-1-85233-998-2 pages 1-30, the whole document -----	1-22
2		

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2015/052060

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009144714	A1 04-06-2009	NONE	