

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
6 November 2008 (06.11.2008)

PCT

(10) International Publication Number
WO 2008/134159 A1

- (51) **International Patent Classification:**
G06F 15/16 (2006.01)
- (21) **International Application Number:**
PCT/US2008/058145
- (22) **International Filing Date:** 25 March 2008 (25.03.2008)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
11/740,737 26 April 2007 (26.04.2007) US
- (71) **Applicant (for all designated States except US):** MICROSOFT CORPORATION [US/US]; One Microsoft Way, Redmond, Washington 98052-6399 (US).
- (72) **Inventors:** SEDUKHIN, Igor; One Microsoft Way, Redmond, Washington 98052-6399 (US). ESHNER, Daniel; One Microsoft Way, Redmond, Washington 98052-6399 (US). SWARTZ, Steve; One Microsoft Way, Redmond, Washington 98052-6399 (US). BHASKARA, Udaya Kumar; One Microsoft Way, Redmond, Washington 98052-6399 (US). PANDITHARADHYA, Nagalinga Durga Prasad Sripathi; One Microsoft Way, Redmond, Washington 98052-6399 (US). KULKARNI, Amol Sudhakar; One Microsoft Way, Redmond, Washington 98052-6399 (US). WU, Haoran Andy; One Microsoft

Way, Redmond, Washington 98052-6399 (US). **BORSA, Mariusz Gerard**; One Microsoft Way, Redmond, Washington 98052-6399 (US). **NEARY, Michael Oliver**; One Microsoft Way, Redmond, Washington 98052-6399 (US).

(81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, 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, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, 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, MT, NL, NO, PL, PT, RO, SE, SI, SK, 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))

[Continued on next page]

(54) **Title:** DISTRIBUTED BEHAVIOR CONTROLLED EXECUTION OF MODELED APPLICATIONS




FIG. 1

(57) **Abstract:** The present invention extends to methods, systems, and computer program products for distributed behavior controlled execution of modeled applications. Embodiments of the invention facilitate the interoperation of a central data store along with various peer-to-peer functionalities. Thus, distributed applications can be executed in an environment that utilizes advantages of both a central data store and peer-to-peer messaging. A read-only portion of a centralized data store can be used to provide some behavior control over various processing systems during execution of a distributed software application. Thus, embodiments of the invention can be used to form a policy-driven collective of nodes forming a distributed, decentralized run-time for model-based applications. Utilizing a cooperating collective of nodes that comply with centralized policies and execute applications over a peer-to-peer fabric permits a distributed application runtime to accommodate distribution and decentralization on a large scale.

WO 2008/134159 A1



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

Published:

— *with international search report*

DISTRIBUTED BEHAVIOR CONTROLLED EXECUTION OF MODELED APPLICATIONS

BACKGROUND

1. Background and Relevant Art

5 [0001] Computer systems and related technology affect many aspects of society. Indeed, the computer system's ability to process information has transformed the way we live and work. Computer systems now commonly perform a host of tasks (e.g., word processing, scheduling, accounting, etc.) that prior to the advent of the computer system were performed manually. More recently, computer systems have
10 been coupled to one another and to other electronic devices to form both wired and wireless computer networks over which the computer systems and other electronic devices can transfer electronic data. Accordingly, the performance of many computing tasks are distributed across a number of different computer systems and/or a number of different computing components.

15 [0002] Various different network architectures can be utilized to permit computer systems to interact with one another. Network architectures can be classified to some extent according to the functional relationships that exist between computer systems of the network architecture. For example, network architectures can be classified as client-server architectures, peer-to-peer architectures, etc.

20 [0003] A client-server architecture separates client and server functionalities between different computer systems. The purpose of different client-server architectures may vary and/or be applied in a variety of ways. However generally, the relationships between computer systems within a client-server architecture are very similar across all client-server architectures. Typically a plurality of client
25 computer systems is configured to send requests (e.g., for data) to a server computer system (or in some environments a server farm). The server computer system (or sever farm) receives and processes requests (e.g., accessing requested data) and sends an appropriate response (e.g., requested data) back to a client.

[0004] One common client-server relationship is that of the World Wide Web
30 ("WWW"). Web servers (the servers) maintain and control access to content (e.g.,

Web pages) that can be requested from Web browsers (the clients) over the Internet.

[0005] Client-server architectures are advantageous for a number of reasons, due primarily to data being stored at a central location. For example, since data is stored at a central location security and access control with respect to the data can be more easily implemented. Further, since there is a clear divide between the functionality of a client and a server, a client can be easily upgrade without impacting operation of the server and vice versa.

[0006] Unfortunately, client-server architectures also have some number of disadvantages. One primary disadvantage is a lack of robustness, which also results from data being stored in a central location. That is, if a server malfunctions, there may be no way to fulfill client requests. Servers can also suffer from traffic congestion, when a number of received requests is so large that it overloads server resources. Further, the resources for establishing a server in many environments typically exceeds the resources of most home computer systems, thereby preventing many users from being able to create a server.

[0007] Peer-to-peer network architectures do not have the notion of servers and clients. Each computer system in a peer-to-peer network is viewed as an equal peer node. Thus, peer-to-peer network architectures rely on the computing power and bandwidth of all the computer systems on the network, as opposed to concentrating resources in a small number of servers (or even in a single server). Peer-to-peer network architectures are often used to connect computer systems together to form *ad hoc* networks, such as, for example, to share files or exchange real-time data.

[0008] Peer-to-peer network architectures are advantageous for a number of reasons, primarily due to resources being shared among computer systems. For example, when a new computer system joins a peer-to-peer network, the resource capacity of the network is increased, since the resources of the joining computer are available for use by other computer systems on the peer-to-peer network. In contrast, resource capacity in a client-server network does not increase when a new client sends a request to a server. Further, since resources are shared among computer systems, peer-to-peer networks are typically more robust than client-

server networks. That is, if one computer system with access to data fails, there is some likelihood that one or more other computer systems also have access to the same data.

[0009] However, peer-to-peer networks also have a variety of disadvantages.

5 Since each computer system is viewed equally, it can be difficult for any one computer system control the behavior of other computer systems, such as, for example, to implement access control or other security measures. This lack of central control makes peer-to-peer networks significantly more vulnerable to malicious attacks (e.g., poisoning attacks, polluting attacks, denial-of-service attacks, spamming etc.) than client-server networks.

[0010] Further there is essentially no restriction on when a computer system can join or leave a peer-to-peer network. Thus, to maintain robustness with respect to requestable data, the same portion of data may be transferred from computer system to computer system across a network (e.g., the Internet) multiple times. For example, when a computer system leaves a peer-to-peer network it can transfer requestable data to another computer system so that the requestable data remains available. Accordingly, each time a computer system leaves a peer-to-peer network, some portion of data is (potentially for a second, third, etc., time) transferred to a remaining computer system.

20 [0011] Enterprise software has typically been developed for use in client-server network architectures (and thus suffers from the problems inherent to client-server architectures). Enterprise software is often proprietary software used to automate business processes, such as, for example, accounting, sales, etc. As such, the companies that develop enterprise software typically desire the enterprise software to be controlled for use in their corporate environment. Enterprise software is frequently very complex and not well adapted for use outside of the intended operating environment. Accordingly, enterprise software is often prevented from operating in a Web based environment and thus, from taking advantage of the simplicity and dynamisms of the Worldwide Web (“WWW”).

30 [0012] Somewhat conversely, Web based applications can view the entire Internet as a single (but uncontrolled) data center and can easily access data from locations

through out the WWW. Based on accessed data, Web based applications can easily be remixed and run with new settings to facilitate changing developer desires. However, since data is made accessible to and from a variety of potentially uncontrolled locations, rapidly developed and remixed Web based applications (to
5 some extent suffer from the problems inherent to peer-to-peer architectures and) often have an undesirable low level of associated security and reliability.

BRIEF SUMMARY

[0013] The present invention extends to methods, systems, and computer **program** products for distributed behavior controlled execution of modeled applications.
10 Embodiments of the invention include a system configured to execute distributed software applications. The system includes a plurality of processing systems. Configurable distributed logic for executing distributed software applications is distributed across the plurality of processing systems.

[0014] A data stored is connected to and shared among the distributed logic. The
15 data store includes read-only data and read-write data. The read-only data prescribes the behavior of the distributed logic when executing a distributed software application. The plurality of processing systems use the read-write data to store and share distributed software application data when executing a distributed software application.

[0015] A message infrastructure is connected to and shared among the distributed
20 logic. The message infrastructure is configured to exchange messages among processing systems to a) discover the presence of configurable distributed logic at processing systems for use in executing distributed software applications and b) coordinate processing systems to implement the behavior prescribed in the read-
25 only data when executing a distributed software application.

[0016] In some embodiments, a distributed processing system executes a model-
based distributed software application. The distributed processing system includes a group of cooperative executive services collectively representing a distributed application runtime. An executive service of the distribute runtime receives a
30 command to execute a model-based software application stored in a read-only portion of a common data repository. The executive service discovers other

executive services of the distributed runtime via a specified kind of messaging infrastructure identified as service bus. Service bus allows executive services to signal each other. The executive service and the other executive services cooperatively group together to provide distributed logic for executing the model-based software application.

[0017] The group of executive services coordinates via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository. Executing the model-based software application includes: a) using the read-write portion of the common data repository to reflect the state of the model-based software application and b) using the read-write portion of the common data repository to reflect operational data corresponding to the model-based software application.

[0018] This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter.

[0019] Additional features and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by the practice of the invention. The features and advantages of the invention may be realized and obtained by means of the instruments and combinations particularly pointed out in the appended claims. These and other features of the present invention will become more fully apparent from the following description and appended claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] In order to describe the manner in which the above-recited and other advantages and features of the invention can be obtained, a more particular description of the invention briefly described above will be rendered by reference to specific **embodiments** thereof which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the

invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

5 [0021] Figure 1 illustrates an example computer architecture that facilitates distributed behavior controlled execution of modeled applications.

[0022] Figure 2 illustrates an example computer architecture that facilitates executing an application across distributed execution logic in accordance with behavior execution policies.

10 [0023] Figure 3 illustrates a flow chart of an example method for executing an application across distributed execution logic in accordance with behavior execution policies.

DETAILED DESCRIPTION

[0024] The present invention extends to methods, systems, and computer program products for distributed behavior controlled execution of modeled applications. Embodiments of the invention include a system configured to execute distributed software applications. The system includes a plurality of processing systems. Configurable distributed logic for executing distributed software applications is distributed across the plurality of processing systems.

20 [0025] A data store is connected to and shared among the distributed logic. The data store includes read-only data and read-write data. The read-only data prescribes the behavior of the distributed logic when executing a distributed software application. The plurality of processing systems use the read-write data to store and share distributed software application data when executing a distributed software application.

25 [0026] A message infrastructure is connected to and shared among the distributed logic. The message infrastructure is configured to exchange messages among processing systems to a) discover the presence of configurable distributed logic at processing systems for use in executing distributed software applications and b) coordinate processing systems to implement the behavior prescribed in the read-
30 only data when executing a distributed software application.

[0027] In some embodiments, a distributed processing system executes a model-based distributed software application. The distributed processing system includes a group of cooperative executive services collectively representing a distributed application runtime. An executive service of the distribute runtime receives a command to execute a model-based software application stored in a read-only portion of a common data repository. The executive service discovers other executive services of the distributed runtime via a specified kind of messaging infrastructure identified as service bus. Service bus allows executive services to signal each other. The executive service and the other executive services cooperatively group together to provide distributed logic for executing the model-based software application.

[0028] The group of executive services coordinates via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository. Executing the model-based software application includes: a) using the read-write portion of the common data repository to reflect the state of the model-based software application and b) using the read-write portion of the common data repository to reflect operational data corresponding to the model-based software application.

[0029] Embodiments of the present invention may comprise a special purpose or general-purpose computer including computer hardware, as discussed in greater detail below. Embodiments within the scope of the present invention also include computer-readable media for carrying or having computer-executable instructions or data structures stored thereon. Such computer-readable media can be any available media that can be accessed by a general purpose or special purpose computer. By way of example, and not limitation, computer-readable media can comprise physical (or recordable type) computer-readable storage media, such as, RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.

[0030] In this description and in the following claims, a “network” is defined as one or more data links that enable the transport of electronic data between computer systems and/or modules. When information is transferred or provided over a network or another communications connection (either hardwired, wireless, or a combination of hardwired or wireless) to a computer, the computer properly views the connection as a computer-readable medium. Thus, by way of example, and not limitation, computer-readable media can also comprise a network or data links which can be used to carry or store desired program code means in the form of computer-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer.

[0031] Computer-executable instructions comprise, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing device to perform a certain function or group of functions. The computer executable instructions may be, for example, binaries, intermediate format instructions such as assembly language, or even source code. Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the described features or acts described above. Rather, the described features and acts are disclosed as example forms of implementing the claims.

[0032] Those skilled in the art will appreciate that the invention may be practiced in network computing environments with many types of computer system configurations, including, personal computers, desktop computers, laptop computers, message processors, hand-held devices, multi-processor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, mobile telephones, PDAs, pagers, and the like. The invention may also be practiced in distributed system environments where local and remote computer systems, which are linked (either by hardwired data links, wireless data links, or by a combination of hardwired and wireless data links) through a network, both perform tasks. In a distributed system environment, program modules may be located in both local and remote memory storage devices.

[0033] Figure 1 illustrates an example computer architecture 100 that facilitates distributed behavior controlled execution of modeled applications. Depicted in computer architecture 100 are data store 101, processing systems 104, 105, and 106, and message infrastructure 108. Each of the depicted components can be
5 connected to one another over a network, such as, for example, a Local Area Network ("LAN"), a Wide Area Network ("WAN"), and even the Internet. Accordingly, data store 101, processing systems 104, 105, and 106, and message infrastructure 108, as well as any other connected components, can create message related data and exchange message related data (e.g., Internet Protocol ("IP")
10 datagrams and other higher layer protocols that utilize IP datagrams, such as, Transmission Control Protocol ("TCP"), Hypertext Transfer Protocol ("HTTP"), Simple Mail Transfer Protocol ("SMTP"), etc.) over the network.

[0034] Processing systems 104, 105, and 106 include logic 104L, 105L, and 106L respectively. Collectively logic 104L, 105L, and 106L represent distributed logic
15 107. Distributed logic 107 is configurable for executing distributed software applications that are distributed across processing systems 104, 105, and 106.

[0035] Data store 101 is connected to and shared among distributed logic 107. As depicted, data store 101 includes read-only data 102 and read-write data 103. Read-only data 102 may be data that is initially written to data store (write once)
20 and then is read-only after the initial write. Read-only data can prescribe the behavior of distributed logic 107 when executing distributed software applications. Processing systems 104, 105, and 106 (as well as contained logic 104L, 105L and 106L respectively) can use read-write data 103 to store and share distributed software application data when executing a distributed software application.

[0036] A series of three periods (an ellipsis) to the left of processing system 104
25 and to the right of processing system 106 represents that other processing systems can be included in computer architecture 100. These other processing systems can contribute logic to distributed logic 107.

[0037] Message infrastructure 108 is connected to and shared among distributed
30 logic 107. Message infrastructure 108 is configured to exchange messages among processing systems 104, 105, and 106. Message infrastructure 108 can be used to

discover the presence of configurable distributed logic, such as, for example, logic 104L, 105L, and 106L, at processing systems 104, 105, and 106 for use in executing distributed applications. Message infrastructure 108 can also be used to coordinate processing systems 104, 105, and 106 to implement the behavior
5 prescribed in read-only data 102 when executing a distributed application. In some embodiments, message infrastructure 108 is a peer-to-peer discovery and messaging (signaling) infrastructure.

[0038] Accordingly, embodiments of the invention facilitate the interoperation of a central data store along with various peer-to-peer functionalities. Thus,
10 distributed applications can be executed in an environment that utilizes advantages of both a central data store and peer-to-peer messaging. The read-only portion of a centralized data store can be used to provide some behavior control over various processing systems during execution of a distributed software application.

[0039] Figure 2 illustrates an example computer architecture 200 that facilitates
15 executing an application across distributed execution logic in accordance with behavior execution policies. Depicted in computer architecture 200 are repository 201, executive services 204, 205, and 206, and service bus 208. Each of the depicted components can be connected to one another over a network, such as, for example, a Local Area Network ("LAN"), a Wide Area Network ("WAN"), and
20 even the Internet. Accordingly, repository 201, executive services 204, 205, and 206, and service bus 208, as well as any other connected components, can create message related data and exchange message related data (e.g., Internet Protocol ("IP") datagrams and other higher layer protocols that utilize IP datagrams, such as, Transmission Control Protocol ("TCP"), Hypertext Transfer Protocol ("HTTP"),
25 Simple Mail Transfer Protocol ("SMTP"), etc.) over the network.

[0040] Executive services 204, 205, and 206 can be dispersed across the network. Executive services 204, 205, and 206 are responsible for running and managing distributed applications. A series of three periods (an ellipsis) to the left of executive service 204 and to the right of executive service 206 represents that other
30 executive services can be included in computer architecture 200.

[0041] Repository 201 includes read-only (or write once) data 202 and read-write data 203. Read-only data 202 further includes application models 211 and policies 212. Application models 211 and policies 212 are shared among executive services, such as, for example, executive services 204, 205, and 206. Policies 212 are declarative statements defining the collective behavior of executive services. For example, a policy can define a maximum number of messages (e.g., twenty) that are permitted to be sent from an executive service per second.

[0042] Accordingly, each executive service in computer architecture 200 interprets and complies with policies 212. Executive services can be permitted some freedom in local behavior as defined in policies 212. For example, executive services can be permitted to perform local error recovery, such as, for example, restarting a process. However, other more serious and/or complex errors, such as, for example, disk failure, may impact multiple executive services for appropriate recovery (e.g., redeployment of application components). Thus, recovery from these more serious and/or complex errors can include collective behavior of a number of executive services. Accordingly, an individual executive service can be prohibited from initiating recovery for these more serious and/or complex errors.

[0043] Application models 211 include application models that can be consumed by a collective of executive services to implement a distributed application. For example, an application model can be a declarative program run and managed by a collection of executive services.

[0044] Read-write data 203 further includes application state 213 and operational data 214. Application state 213 can reflect the state of distributed application (e.g., an application model form application models 211), such as, for example, configured, deployed, running, etc. Operational data 214 can include information, such as, for example, run time metrics, used to determine compliance with policies 212. For example, operational data 214 can include an indication of how many messages an executive service is sending per second. An executive service can write operational data into operational data 214 based on runtime behavior of distributed applications.

[0045] Executive services can use service bus 208 (e.g., a peer-to-peer fabric, such as, for example, a distributed publish/subscribe infrastructure) to discover each other and implement cooperative behavior. Service bus 208 can include a transitively federated namespace service (permitting executive services to federate with one another) and a message bus. For example, through communication via service bus 208, executive services 204, 205, and 206 can collectively form a cluster to host a Web site as indicated in a model of a distributed application (e.g., contained in application models 211).

[0046] Other cooperative behavior can be implemented when a catastrophic failure of one executive service occurs. One or more healthy “buddy” executive services can signal one another via service bus 208 and select another executive service (e.g., through a vote) to take over and recover from the error.

[0047] As depicted, executive services 204, 205, and 206 include controllers 204C, 205C, and 206C respectively. Generally, controllers interpret policies (e.g., policies 212) and enforce cooperative behavior among executive services. For example, controllers can enforce looking after “buddy” executive services and enforce suppressing negative local behavior, such as, for example, attempting to shut down a local service (e.g., Internet Information Services (“IIS”) or SQL server) that contains part of a running distributed application.

[0048] Thus, embodiments of the invention include environments where every executive service has the same view on policies and application models and can discover and signal other executive services. Accordingly, it is possible to interact with one executive service to operate a collective group of (or all of the) executive services in an environment, such as, for example, computer architecture 200.

[0049] Figure 3 illustrates a flow chart of an example method 300 for executing an application across distributed execution logic in accordance with behavior execution policies. The method 300 will be described with respect to the components and data in computer architecture 200.

[0050] Method 300 includes an act of an executive service of a distributed runtime receiving a command to execute a model-based software application stored in the read-only portion of a common data repository (act 301). For example,

executive service 204 can receive execute command 221 to execute application 211A.

[0051] Method 300 includes an act of executive service discovering other executive services of the distributed run-time via the service bus (act 302). For example, executive service 204 can discover executive services 205 and 206 through message exchanges 222 and 223 respectively, via service bus 208.

[0052] Method 300 includes an act of the executive service and the other executive services cooperatively grouping together to provide distributed logic for executing the model-based software application (act 303). For example, executive service 204 and executive services 205 and 206 can cooperatively group together (e.g., in accordance with policies 212) to provide distributed logic for executing application 211A.

[0053] Method 300 includes an act of the group of executive services coordinating via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository (act 304). For example, the group of executive services 204, 205, and 206 can coordinate via service bus 208 to execute application 211A in accordance with policies 212.

[0054] Coordinating via the service bus to execute the model-based software application in accordance with the policies includes an act of using the read-write portion of the common data repository to reflect the state of the model-based software application (act 305). For example, executive services 204, 205, and 206 can write data to and read data from application state 213 to reflect the state (e.g., configured, deployed, running, etc.) of application 211.

[0055] Coordinating via the service bus to execute the model-based software application in accordance with the policies includes an act of using the read-write portion of the common data repository to reflect operational data corresponding to the model-based software application (act 306). For example, executive services 204, 205, and 206 can write data to and read data from operational data 214 to reflect runtime metrics (e.g., average response time, messages per second, etc.)

corresponding to application 211. The run time metrics can be used to determine compliance with policies 212.

[0056] Thus, embodiments of the invention can be used to form a policy-driven collective of nodes forming a distributed, decentralized run-time for model-based applications. Controllers at each node can enforce policies and coordinate with other nodes. Utilizing a cooperating collective of nodes that comply with centralized policies and execute applications over a peer-to-peer fabric permits a distributed application runtime to accommodate distribution and decentralization on a large scale. For example, a large scale WAN or even the Internet can be viewed as a data center.

[0057] The present invention may be embodied in other specific forms without departing from its spirit or essential characteristics. The described embodiments are to be considered in all respects only as illustrative and not restrictive. The scope of the invention is, therefore, indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

CLAIMS

What is claimed:

1. A system (100) configured to execute distributed software applications, the system comprising:
 - 5 a plurality of processing systems (104, 105,106), configurable distributed logic for executing distributed software applications distributed across the plurality of processing systems;
 - a data store (101) connected to and shared among the distributed logic, the data store including:
 - 10 read-only data (102), the read-only data prescribing the behavior of the distributed logic when executing a distributed software application; and
 - read-write data (103), the plurality of processing systems using the read-write data to store and share distributed software application data when executing a distributed software application; and
 - 15 a message infrastructure (108) connected to and shared among the distributed logic, the message infrastructure configured to exchange messages among processing systems to:
 - 20 discover the presence of configurable distributed logic (104L, 105L, 106L) at processing systems for use in executing distributed software applications; and
 - coordinate processing systems to implement the behavior prescribed in the read-only data when executing a distributed software application.
 - 25 2. The system as recited in claim 1, wherein the plurality of processing systems form a distributed, decentralized run-time for model-based applications.
 3. The system as recited in claim 1, wherein each of the plurality of processing systems includes a controller that monitors the behavior of the processing system.
 4. The system as recited in claim 1, wherein the read-only data comprises
30 policies that processing system are to comply with when executing distributed applications.

5. The system as recited in claim 1, further comprising second read-only data describing one or more model-based software applications configured for execution by the distributed logic.

6. The system as recited in claim 1, wherein the read-write data comprises data used to describe the application state of a distributed application, wherein the application state is selected from among configured, deployed, and running.

7. The system as recited in claim 1, wherein the read-write data comprises run time metrics used to determine compliance with prescribed behavior indicated in the read-only data.

8. The system as recited in claim 1, wherein the message infrastructure is a distributed publish/subscribe infrastructure.

9. The system as recited in claim 1, wherein the message infrastructure is a peer-to-peer messaging infrastructure.

10. In a distributed processing system (200), the distributed processing system including a group of cooperative executive services (204, 205, 206) collectively representing a distributed runtime, distributed logic for executing model-based software applications distributed across the group of cooperative executive services, the distributed processing system including a common data repository (201) having a read-only portion (202) and a read-write portion (203), the read-only portion storing one or more model-based software applications (2114) and storing policies (212) prescribing the behavior of the distributed runtime when executing a model-based software application, the read-write portion usable by executive services to maintain application state and application operational data, the distributed system including a service bus (208), executive services using the service bus to discover and cooperatively communicate with one another to execute model-based software applications, a method for executing a model-based distributed software application, the method comprising:

an act of an executive service of the distributed run-time receiving a command (221) to execute a model-based software application stored in the read-only portion of the common data repository;

an act of the executive service discovering (222) other executive services of the distributed run-time via the service bus;

an act of the executive service and the other executive services cooperatively grouping together to provide distributed logic for executing the model-based software application;

an act of the group of executive services coordinating via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository, including:

an act of using the read-write portion of the common data repository to reflect the state (213) of the model-based software application; and

an act of using the read-write portion of the common data repository to reflect operational data (214) corresponding to the model-based software application.

11. The method as recited in claim 10, wherein the act of the executive service discovering other executive services of the distributed run-time via the service bus comprises the executive service discovering other executive services of the distributed run-time via a peer-to-peer fabric.

12. The method as recited in claim 10, wherein the act of the group of executive services coordinating via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository comprises an act of a controller at each executive service monitor the behavior of the executive service for compliance with the policies.

13. The method as recited in claim 10, wherein the act of the group of executive services coordinating via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository comprises an act of the group of executive services coordinating via a peer-to-peer infrastructure to execute the model-based software application in accordance with policies stored in a central repository.

14. The method as recited in claim 10, wherein the act of the group of executive services coordinating via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository comprises utilizing the policies to suppress local behavior
5 at an executive service that negatively affects execution of a distributed application at one or more other executive services.

15. The method as recited in claim 10, wherein the act of using the read-write portion of the common data repository to reflect operational data corresponding to the model-based software application comprises the act of using the read-write
10 portion of the common data repository to maintain runtime metrics that can be evaluated to determine if executive services are complying with the polices.

16. A computer program product for use in a distributed processing system (200), the distributed processing system including a group of cooperative executive services (204, 205, 206) collectively representing a distributed runtime, distributed
15 logic for executing model-based software applications distributed across the group of cooperative executive services, the distributed processing system including a common data repository (201) having a read-only portion (202) and a read-write portion (203), the read-only portion storing one or more model-based software applications (2114) and storing policies (212) prescribing the behavior of the
20 distributed runtime when executing a model-based software application, the read-write portion usable by executive services to maintain application state and application operational data, the distributed system including a service bus (208), executive services using the service bus to discover and cooperatively communicate with one another to execute model-based software applications, the computer
25 program product for implementing a method for executing a model-based distributed software application, the computer program product including one or more computer-readable media having stored thereon computer-executable instructions that, when executed, cause an executive service of the distributed processing system to perform the method, including the following:

30 receive a command (221) to execute a model-based software application stored in the read-only portion of the common data repository;

discover (222) other executive services of the distributed run-time via the service bus;

group together cooperatively with the other executive services to provide distributed logic for executing the model-based software application;

coordinate via the service bus to execute the model-based software application in accordance with the policies stored in the read-only portion of the common data repository, including:

use the read-write portion of the common data repository to reflect the state (213) of the model-based software application; and

use the read-write portion of the common data repository to reflect operational data (214) corresponding to the model-based software application.

17. The computer program product as recited in claim 16, wherein the computer-executable instructions that, when executed, cause an executive service to coordinate via the service bus to execute the model-based software application comprise computer-executable instructions that, when executed, monitor the behavior of the executive service for compliance with the policies.

18. The computer program product as recited in claim 16, wherein the computer-executable instructions that, when executed, cause an executive service to coordinate via the service bus to execute the model-based software application comprise computer-executable instructions that, when executed, cause an executive service to coordinate via a peer-to-peer infrastructure to execute the model-based software application in accordance with policies stored in a central repository.

19. The computer program product as recited in claim 16, wherein the computer-executable instructions that, when executed, cause an executive service to coordinate via the service bus to execute the model-based software application comprise computer-executable instructions that, when executed, cause the executive service to utilize the policies to suppress local behavior at another executive service that negatively affects execution of a distributed application at one or more other executive services.

20. The computer program product as recited in claim 16, wherein the computer-executable instructions that, when executed, cause an executive service to use the read-write portion of the common data repository to reflect operational data corresponding to the model-based software application comprise computer-executable instructions that, when executed, cause an executive service to use the
5 read-write portion of the common data repository to maintain runtime metrics that can be evaluated to determine if executive services are complying with the policies.

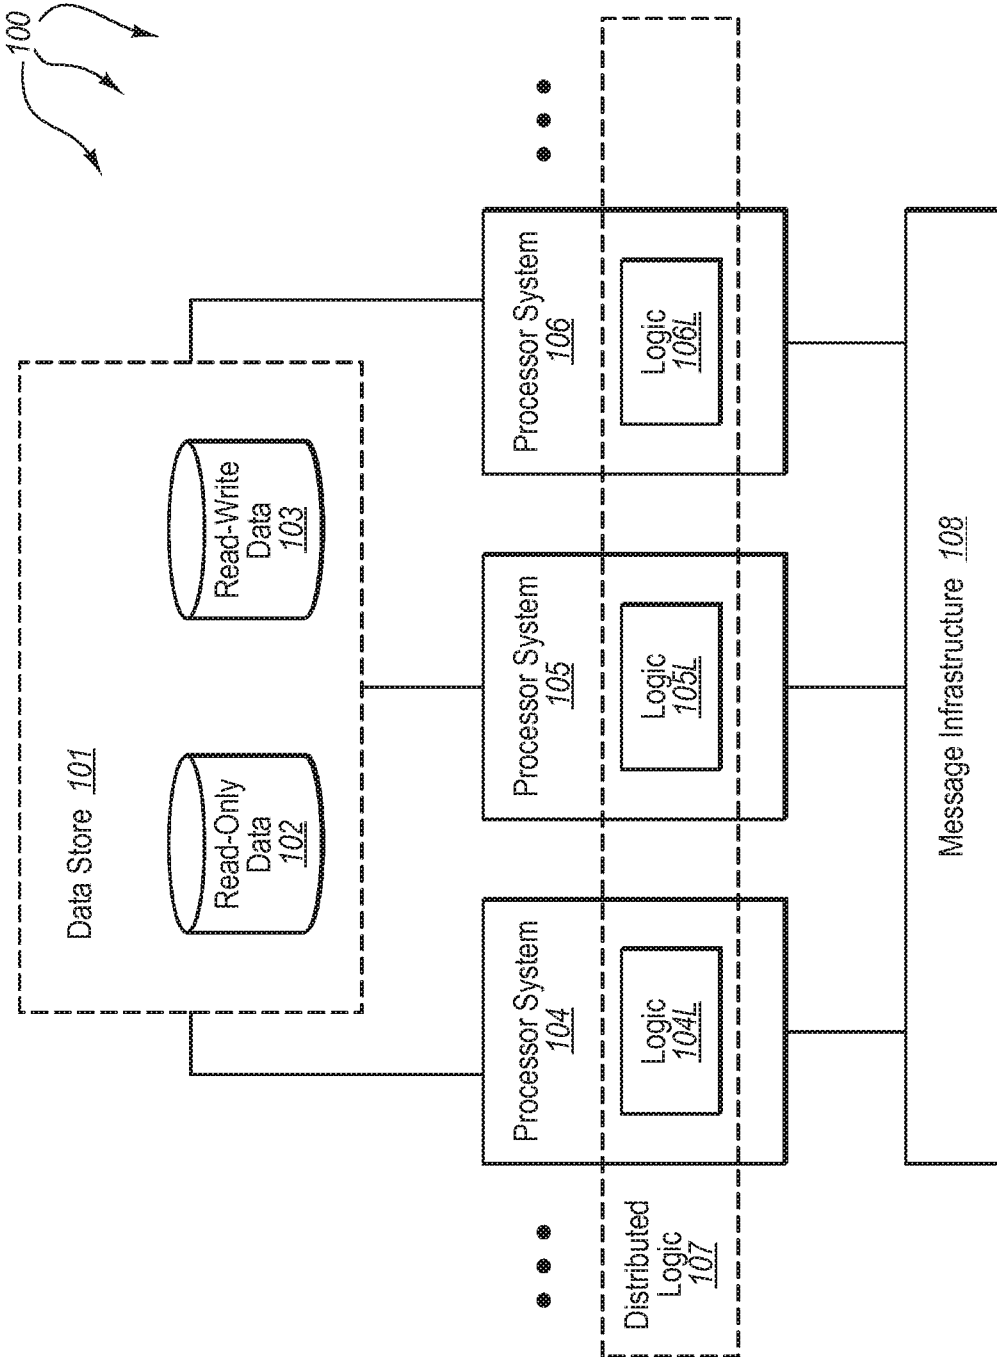


FIG. 1




FIG. 2

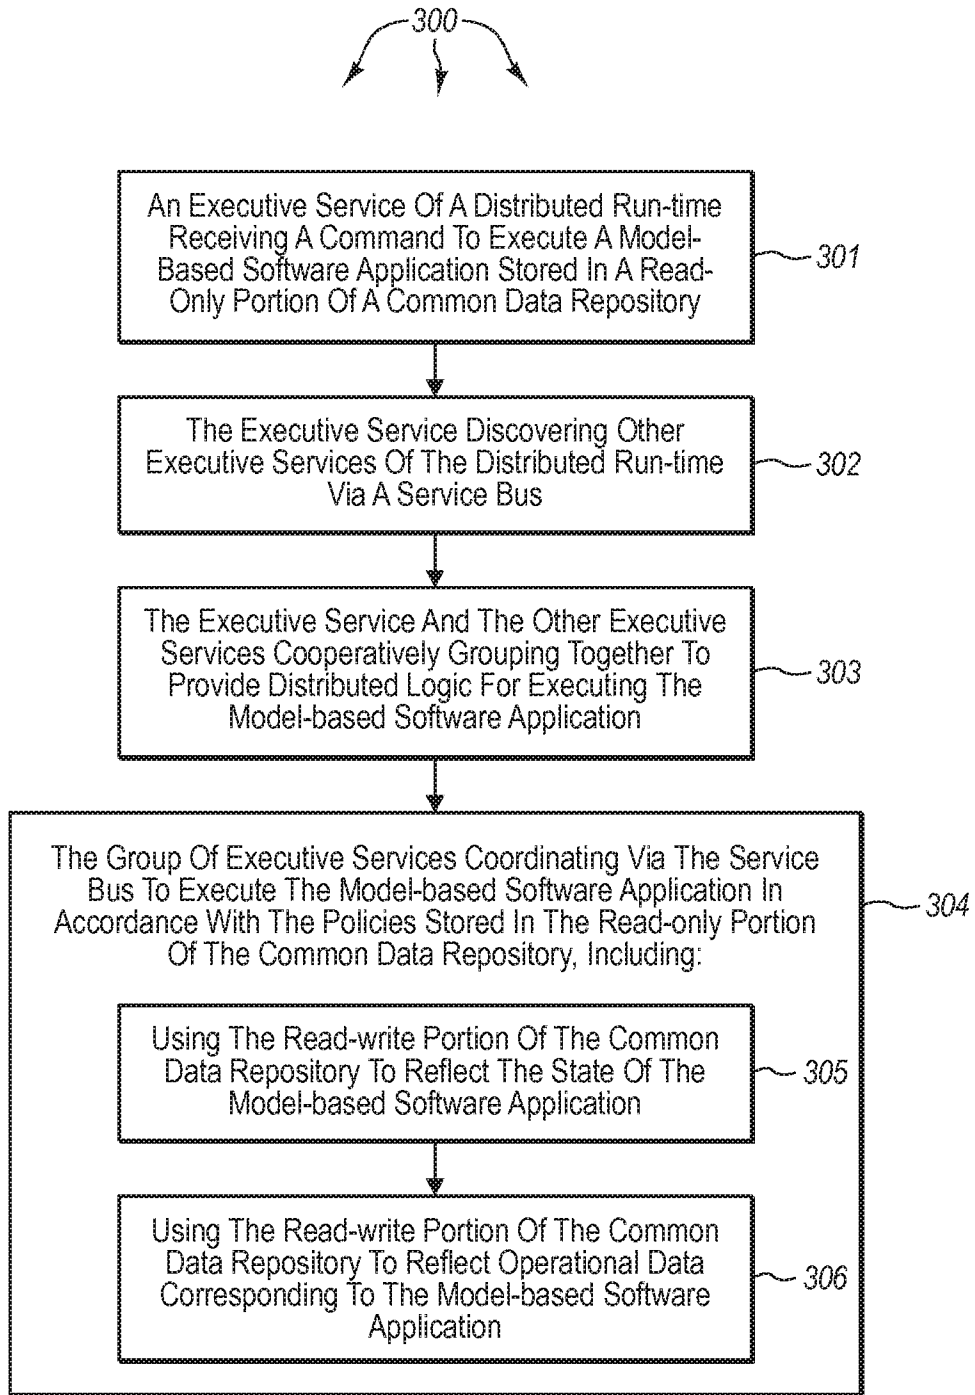




FIG. 3

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2008/058145

A. CLASSIFICATION OF SUBJECT MATTER		
<i>G06F 15/16(2006.01)i</i>		
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) IPC 8 : G06F		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched Korean Utility models and applications for Utility models since 1975 Japanese Utility models and application for Utility models since 1975		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) eKIPASS(KIPO internal) "configurable" "logic" "distributed application"		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 6640241 B1 (RAYMOND E. OZZLE et al.) 28 October 2003 See column 3, line 61 - column 24, line 36.	1-20
A	US 6654783 B1 (EDWARD A. HUBBARD) 25 November 2003 See column 2, line 64 - column 19, line 61.	1-20
A	US 2006/0206537 A1 (YING-HSIN ROBERT CHIANG) 14 September 2006 See page 2, [0026] - page 9, [0104].	1-20
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
<p>* Special categories of cited documents:</p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier application or patent but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"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</p> <p>"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</p> <p>"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</p> <p>"&" document member of the same patent family</p>		
Date of the actual completion of the international search 22 AUGUST 2008 (22.08.2008)		Date of mailing of the international search report 22 AUGUST 2008 (22.08.2008)
Name and mailing address of the ISA/KR  Korean Intellectual Property Office Government Complex-Daejeon, 139 Seonsa-ro, Seo-gu, Daejeon 302-701, Republic of Korea Facsimile No. 82-42-472-7140		Authorized officer KWON, Oh Seong Telephone No. 82-42-481-8526 

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2008/058145

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6640241 B1	28.10.2003	CA 2380148 A1	25.01.2001
		EP 1328882 A2	23.07.2003
		EP 1328882 B1	23.04.2008
		IL 147704 A0	14.08.2002
		JP 2003-536123 T2	02.12.2003
		KR 10-2002-0042805 A	07.06.2002
		WO 2001-06365 A2	25.01.2001
		WO 2001-06365 C1	16.08.2001
		WO 2001-06365 A3	15.05.2003
		US 6654783 B1	25.11.2003
CA 2403895 A1	04.10.2001		
EP 1360580 A2	12.11.2003		
US 2001-039497 A1	08.11.2001		
US 2001-027467 A1	04.10.2001		
US 2002-013832 A1	31.01.2002		
US 2003-033543 A1	13.02.2003		
US 2003-149765 A1	07.08.2003		
US 2004-103139 A1	27.05.2004		
US 2004-148336 A1	29.07.2004		
US 2004-215829 A1	28.10.2004		
US 2005-010664 A1	13.01.2005		
US 2005-149532 A1	07.07.2005		
US 2008-117062 A1	22.05.2008		
JP 2004-513406 A	30.04.2004		
US 6847995 B1	25.01.2005		
US 6891802 B1	10.05.2005		
US 6963897 B1	08.11.2005		
US 7003547 B1	21.02.2006		
US 7020678 B1	28.03.2006		
US 7039670 B1	02.05.2006		
US 7082474 B1	25.07.2006		
US 7092985 BB	15.08.2006		
US 7254607 BB	07.08.2007		
WO 2001-73545 A2	04.10.2001		
WO 2001-73545 A3	28.08.2003		
US 2006-206537 A1	14.09.2006		