

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
17 November 2005 (17.11.2005)

PCT

(10) International Publication Number
WO 2005/109195 A2

- (51) International Patent Classification⁷: **G06F 9/46**
- (21) International Application Number:
PCT/EP2005/051804
- (22) International Filing Date: 22 April 2005 (22.04.2005)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
10/841,725 8 May 2004 (08.05.2004) US
- (74) Agent: **LITHERLAND, David, Peter**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).
- (81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (71) Applicant (*for all designated States except US*): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).
- (71) Applicant (*for MG only*): **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour, Portsmouth Hampshire PO6 3AU (GB).
- (72) Inventors; and
- (75) Inventors/Applicants (*for US only*): **BOZEK, James** [US/US]; 21608 Meridian Avenue S, Bothwell, Washington 98021 (US). **STEPHENS, Robert** [US/US]; 7829 Center Blvd #198, Snoqualmie, Washington 98065 (US). **WOOLDRIDGE, James** [US/US]; 5434 324th Place SE, Fall City, Washington 98024 (US).
- (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, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— *without international search report and to be republished upon receipt of that report*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: DYNAMIC MIGRATION OF VIRTUAL MACHINE COMPUTER PROGRAMS

(57) Abstract: A system of an embodiment is disclosed that includes a number of server computing devices and a management server computing device. Each server computing device has a virtual host computer program running thereon to support one or more virtual machine computer programs. Each virtual machine computer program is able to execute an instance of an operating system on which application computer programs are executable. The management server computing device monitors the server computing devices, and causes the virtual machine computer programs supported by the virtual host computer program of a first server computing device to dynamically migrate to the virtual host computer program of a second server computing device, upon one or more conditions being satisfied. The conditions may include the first server being predicted as failure prone, the first server consuming power less than a threshold, and the first server having resource utilization less than a threshold.

Description

DYNAMIC MIGRATION OF VIRTUAL MACHINE COMPUTER PROGRAMS

Technical Field

- [001] The present invention relates to the dynamic migration of virtual machine computer programs from a virtual host computer program running on one server computing device to a virtual host computer program running on another server computing device.

Background Art

- [002] A common type of network that is deployed in organizations is the client-server network. In a client-server network, there may be a number of client computing devices, or clients, which are typically used by end users of an organization, and a number of server computing devices, or servers, which are computing devices that are shared among the clients, and thus the users. Types of servers can include application servers, file servers, intranet servers, e-mail servers, electronic commerce servers, print servers, proxy servers, and web servers, among other kinds of servers.
- [003] To reduce information technology (IT) costs, some types of servers can run virtual host computer programs. A virtual host computer program is a type of supervising computer program, which enables a number of virtual machine computer programs to be run on the server. A virtual machine computer program allows a separate instance of an operating system to be run thereon, such as versions of the Microsoft Windows® operating system, versions of the UNIX® operating system, and versions of the Linux® operating system.
- [004] In an environment within which virtual machine computer programs are run on virtual host computer programs of servers, users may not be required to have full-fledged client computing devices. Rather, each user may be assigned a virtual machine computer program, and have a reduced-capability client computing device that basically handles input from the user to send to the virtual machine computer program, and output from the virtual machine computer program to the user. The server on which the virtual machine computer program runs thus provides the computational support for a number of users, so that expensive clients are not needed for these users.
- [005] In these and other types of uses of virtual machine computer programs, maintaining high availability of the servers is important. For instance, a faulty server can affect a number of users, whereas a faulty client may affect just a single user. Mission-critical application computer programs may run in conjunction with the operating systems of the virtual machine computer programs, regardless of whether the virtual machine computer programs correspond to particular users or whether they run multiple user-

type application computer programs. A faulty server in this instance may affect a number of virtual machine computer programs, and thus a number of mission-critical application computer programs, which can be expensive to the organizations running these mission-critical application programs.

- [006] Furthermore, in most organizations, power conservation has become an issue, both to reduce operational costs of IT, and to prove to the general public that the organizations are good public citizens. In many situations, a larger number of servers than is typically needed are powered on, for peak-utilization situations in which the processing capability of all the servers may be needed. However, the occurrence of such peak-utilization situations may be infrequent, wasting electrical power for the majority of the time in which all of the servers remain powered on.

Disclosure of Invention

- [007] A system of an embodiment of the present invention includes, a number of server computing devices, and a management server computing device. The server computing devices are communicatively coupled to a network. Each server computing device has a virtual host computer program running thereon to support one or more virtual machine computer programs. Each virtual machine computer program is able to execute an instance of an operating system on which application computer programs are executable. The management server computing device is also communicatively coupled to the network. The management server computing device monitors the server computing devices, and causes the virtual machine computer programs supported by the virtual host computer program of one of the server computing devices to migrate to the virtual host computer program of another of the server computing devices, upon one or more conditions being satisfied.

- [008] A management system of another embodiment of the present invention is for a number of server computing devices that have virtual host computer programs running thereon to support virtual machine computer programs that are able to execute instances of operating systems. The management system includes a coordinator server computing device and a management server computing device. The coordinator server computing device communicates with and coordinates the virtual machine computer programs supported by the virtual host computer programs of the server computing devices. The management server computing device causes the virtual machine computer programs supported by the virtual host computer program of a first server computing device to migrate to the virtual host computer program of a second server computing device, upon detecting one or more conditions being satisfied. The management server computing device causes the migration by communicating with the coordinator server computing device.

[009] An article of manufacture of an embodiment of the invention includes a computer-readable medium and means in the medium. The means is for dynamically migrating virtual machine computer programs supported by a virtual host computer program of a first server computing device to a virtual host computer program of a second server computing device, without restarting the virtual machine computer programs, upon one or more conditions being satisfied. The conditions may include one or more of: the first server computing device being predicted as failure prone; the first server computing device being determined as having resource utilization less than a threshold; and, the first server computing device being determined as having power consumption less than a threshold.

[010] A method of varying embodiments of the invention first predicts a failure at a first server computing device, detects resource utilization at the first server computing device as less than a threshold, and/or detects power consumption of the first server computing device as less than a threshold. In response, the method correlates the first server computing device with a virtual host computer program supporting one or more virtual machine computer programs. The method dynamically migrates the virtual machine computer programs from the virtual host computer program of the first server computing device to a virtual host computer program of a second server computing device.

[011] Embodiments of the invention provide for advantages over the prior art. Maximum server availability is provided by predicting failures at servers, and, before these servers can fail, their virtual machine computer programs are dynamically migrated to another server. Dynamic migration means that the operating systems supported by these virtual machine computer programs do not have to be shutdown or restarted during the migration process. Power conservation is provided by monitoring resource utilization of the servers and/or power consumption of the servers. When, for a given server, either falls below a corresponding threshold, this means that the server's processing capability is not being efficiently utilized, such that it should and can be powered down until the capability is needed. Therefore, the virtual machine computer programs running on the server are dynamically migrated to another server, and the former server shut down.

Brief Description of the Drawings

[012] Embodiments of the invention will now be described, by way of example only, with reference to the accompanying drawings in which:

[013] FIG. 1 is a diagram of a system in which there are servers having virtual host computer programs that support virtual machine computer programs, according to an embodiment of the invention;

- [014] FIG. 2 is a diagram of a system in which there are servers having virtual host computer programs that support virtual machine computer programs, according to another embodiment of the invention;
- [015] FIG. 3 is a flowchart of a method for ensuring maximum server availability, according to an embodiment of the invention;
- [016] FIG. 4 is a flowchart of a method for ensuring power conservation by monitoring resource utilization, according to an embodiment of the invention; and
- [017] FIG. 5 is a flowchart of a method for ensuring power conservation by monitoring power consumption, according to an embodiment of the invention.

Mode for the Invention

- [018] In the following detailed description of exemplary embodiments of the invention, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration specific exemplary embodiments in which the invention may be practiced. These embodiments are described in sufficient detail to enable those skilled in the art to practice the invention. Other embodiments may be utilized, and logical, mechanical, and other changes may be made without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined only by the appended claims.

[019] System and overview

- [020] FIG. 1 shows a system 100, according to an embodiment of the invention. The system 100 includes a network 102, to which servers 104A and 104B, collectively referred to as the servers 104, are communicatively coupled. Although FIG. 1 shows two of the servers 104, there may be more than two of the servers 104 in one embodiment. A management console 106, a management server 108, and a coordinator server 110, which together may be considered to be a management system for the servers 104, are also each communicatively coupled to the network 102. The network 102 may be or include one or more of: a local-area network (LAN), a wide-area network (WAN), an intranet, an extranet, the Internet, a wired network, a wireless network, and a telephony network, among other types of networks.
- [021] Each of the servers 104 is a server computing device. The server 104A includes a network adapter 112 to communicatively couple the server 104A to the network 102, whereas the server 104B includes a network adapter 118 to communicatively couple the server 104B to the network 102. The server 104A also includes a network adapter 114 to communicatively couple the server 104A to a separate storage area network (SAN) 116, whereas the server 104B includes a network adapter 120 to communicatively couple the server 104B to the SAN 116. The SAN 116 is a networked

collection of storage, such as hard disk drives, that are available to both of the servers 104. Furthermore, the servers 104 may each include other hardware components beside those indicated in FIG. 1, such as processors, memory, and so on.

[022] The server 104A has a virtual host computer program 122 running thereon, while the server 104B has a virtual host computer program 124 running thereon. The virtual host computer programs 122 and 124 are supervisory computer programs that enable a number of virtual machines to run on the servers 104. For instance, the server 104A has virtual machine computer programs 126A and 126B, collectively referred to as the virtual machine computer programs 126, running thereon as supported by the virtual host computer program 122. Similarly, the server 104B has virtual machine computer programs 128A and 128B, collectively referred to as the virtual machine computer programs 128, running thereon as supported by the virtual host computer program 124. Although there are two of the virtual machine computer programs 128 and 126 for each of the servers 104 in FIG. 1, in other embodiments of the invention there can be more or less virtual machine computer programs.

[023] The virtual machine computer programs 126 and 128 are separate partitions that enable the servers 104 to run independent and multiple instances of operating systems thereon, such as versions of the Microsoft Windows® operating system, versions of the UNIX® operating system, and versions of the Linux® operating system. As such, each of the virtual machine computer programs 126 and 128 can be considered a virtual computer, in that the operating systems running on the virtual machine computer programs 126 and 128 operate as if they were the only operating systems running on their respective servers 104. Therefore, a virtual machine computer program is a single system image within a computer that supports multiple system images. Each system image contains the operating system and its associated applications, and each image may have the same operating system or a different operating system. The operating systems running on the virtual machine computer programs 126 and 128 allow application computer programs to be executed thereon.

[024] The coordinator server 110 is a server computing device, and may include hardware such as processors, memory, storage, and network adapters, as well as other types of hardware, which are not particularly depicted in FIG. 1. The coordinator server 110 coordinates the virtual machine computer programs 126 and 128 supported by the virtual host computer programs 122 and 124 running on the servers 104. In particular, the coordinator server 110 monitors in conjunction with which of the virtual host computer programs 122 and 124, and thus, indirectly, which of the servers 104, each of the virtual machine computer programs 126 and 128 is running.

[025] The coordinator server 110 further is able to accomplish migration of virtual machine computer programs from one of the servers 104 to another of the servers 104.

That is, the coordinator server 110 specifically initiates, facilitates, and orchestrates such migration. For instance, the coordinator server 110 is able to cause either or both of the virtual machine computer programs 128 to migrate, or move, from being supported by the virtual host computer program 124 of the server 104B to being supported by the virtual host computer program 122 of the server 104A. In at least some embodiments of the invention, the migration of virtual host computer programs among the servers 104 is dynamic, which means that the virtual host computer programs do not have to be shut down or restarted before and after moving to a new server.

[026] As a result, the operating systems running on the virtual machine computer programs and the application computer programs running on the operating systems do not have to be shutdown or restarted during the migration process. For example, instances of VMware® virtual machine computer programs, available from VMware, Inc., may be dynamically migrated among virtual host computer programs of servers. In another embodiment, the migration of virtual host computer programs among the servers 104 is static, in that virtual machine computer programs may first have to be shut down before being moved to another server, and then restarted at the new server.

[027] The management server 108 is a server computing device, and may include hardware such as processors, memory, storage, and network adapters, as well as other types of hardware, which are not particularly depicted in FIG. 1. The management server 108 is able to cause the virtual machine computer programs 126 and 128 supported by the virtual host computer programs 122 and 124 running on the servers 104 to migrate between the virtual host computer programs 122 and 124 and thus between the servers 104. For instance, in one embodiment, the management server 108 may migrate the virtual machine computer programs 126 and 128 between the virtual host computer programs 122 and 124 and thus between the servers 104 by sending commands to the coordinator server 110, which actually performs the migration process. In an alternative embodiment, the management server 108 may itself migrate virtual machine computer programs 126 and 128 between the virtual host computer programs 122 and 124 and thus between the servers 104, especially where such migration is static migration, without involving the coordinator server 110, which may not be present in the system 100 in this embodiment.

[028] The management server 108 may control the servers 104 and 108, such as restarting and shutting them down, via the network adapters 112 and 118, which may be remote supervisor adapters in one embodiment. The management server 108 is further to monitor the occurrence or satisfaction of one or more conditions that occur relative to any of the servers 104. In response to any of the conditions being satisfied with respect to a particular one of the servers 104, the management server 108 causes the virtual

machine computer programs on this server to migrate to the other server, and then shuts down or causes the server to shut down.

[029] In one embodiment, the management server 108 performs its monitoring functionality in relation to a management agent computer program 130 running on the virtual host 122 of the server 104A, and a management agent computer program 132 running on the virtual host 124 of the server 104B. In another embodiment, the management server 108 may be able to accomplish static migration of the virtual machine computer programs 126 and 128 between the virtual host computer programs 122 and 124 between the servers 104, without involving the coordinator server 110, via the management agent computer programs 130 and 132. As such, the management agent computer programs 130 and 132 may allow the management server 108 to control their servers 104 therethrough. An agent may generally be considered a computer program that runs at the behest of another computer or computer program.

[030] The management agent computer programs 130 and 132 or the network adapters 112 or 118, thus report back to the server 108 when one of the conditions has occurred. For example, one of the conditions may be that one of the servers 104 has been detected or predicted as being failure prone. That is, the management agent computer program running on a server may predict that a failure is likely to occur on the server in the near future. As an illustrative example, the operating temperature of the server may be outside of a desired range, indicating that the cooling subsystem of the server is malfunctioning, and which may portend failure of the entire server itself. Migration of the virtual machine computer programs away from such a failure-prone server therefore provides for maximum server availability. In this example, the management server 108 monitors predictive failure of the servers 104.

[031] Another condition may be that one of the servers 104 is determined as having resource utilization less than a threshold. For instance, the processors of a server may on average have less than 20% utilization, the memory of the server may on average have less than 30% utilization, and so on. Migration of the virtual machine computer programs away from such an under-utilized server may be accomplished so that the server may be shut down to decrease power consumption within the system 100. The virtual machine computer programs utilizing such a low amount of their server's resources are likely to be able to easily be absorbed by another server.

[032] A third example of one of the conditions may be that one of the servers 104 is determined as having power consumption less than a threshold. For instance, a server may be consuming less than 30 watts of power, when during more regular operation the server is likely to consume more than 100 watts of power. Such low power consumption indirectly indicates that the server is being under utilized, such that migration of the virtual machine computer programs away from the server may be ac-

complished so that the server can be shut down to decrease power consumption. As before, the virtual machine computer programs utilizing such a low amount of their server's resources are likely to be able to easily be absorbed by another server.

[033] The management console 106 is a computing device, such as a client computing device. The console 106 may include hardware such as processors, memory, storage, and network adapters, as well as other types of hardware, which are not particularly depicted in FIG. 1. The management console 106 allows a network administrator or another user to modify parameters associated with the monitoring of the servers 104 and the migration of the virtual machine computer programs 126 and 128 among the servers 104. For instance, the management console 106 may be the mechanism by which a user is able to interact with the management server 108 and change the various thresholds associated with the conditions monitored by the management server 108. A user may be able to change how predictive faults are determined by the management server 108, the thresholds at which resource utilization of the servers 104 is considered low, and/or the thresholds at which power consumption of the servers 104 is considered low, in one embodiment of the invention.

[034] FIG. 2 shows the system 100, according to another embodiment of the invention. Like-numbered features and components of the system 100 between FIGs. 1 and 2 indicate comparable functionality between the embodiments of FIGs. 1 and 2, and description thereof is not represented in relation to FIG. 2 to avoid redundancy. In the embodiment of FIG. 2, the servers 104 are implemented as blade servers physically connected to a common blade chassis 202. The blade chassis 202 has a backplane into which the blade servers 104 are connected, where the blade servers 104 obtain power and interconnect with other system components via their insertion into the blade chassis 202.

[035] The blade servers 104 are therefore insertable into and controllable by or via the blade chassis 202. The blade servers 104 are thus single-board computers or input/output (I/O) boards. The blade chassis 202 has network adapters 114 and 120 for the blade servers 104A and 104B, through which the blade servers 104 are communicatively connected to the SAN 116. The network adapters 114 and 120 in this embodiment may thus be host bus or board adapters. The blade chassis 202 further has a management module 212, via which the blade chassis 202, and thus the blade servers 104, communicate with the network 102 and thus with the management server 108 and the coordinator server 110. The management module 212 is a hardware module that enables the management server 108 in particular to control the servers 104 of the chassis 202.

[036] The coordinator server 110 is specifically depicted in FIG. 2 as including a management agent computer program 210, through which the management server 108

communicates with the coordinator server 110. The management agent computer program 210 may also be included in the embodiment of FIG. 1, although it is not specifically depicted in FIG. 1. The management server 108 itself is depicted in FIG. 2 as including a resource consumption monitor service 208. The resource consumption monitor service 208 specifically is the mechanism in the embodiment of FIG. 2 by which the management server 108 monitors resource utilization and/or power consumption of the blade servers 104. As such, the blade servers 104 include additional agent computer programs 204 and 206 that are resource consumption monitor agent computer programs monitoring resource utilization and/or power consumption of the blade servers 104. The resource consumption monitor agent computer programs 204 and 206 monitor the resource utilization and/or power consumption of their respective blade servers 104, communicating back to the resource consumption monitor service 208, and thus to the management server 108, when resource utilization and/or power utilization drops below a given threshold.

[037] Ensuring maximum server availability

[038] FIG. 3 shows a method 300 for ensuring maximum server availability, according to an embodiment of the invention. The method 300 may be performed in relation to the system 100 of FIG. 1 or FIG. 2. The method 300 is specifically described in relation to the system 100 of FIG. 1 for examples purposes only. Furthermore, the method 300 may be implemented as a computer program that can be stored on and/or executed from a computer-readable medium. Such a medium may include a volatile or non-volatile medium, a semiconductor medium, a magnetic medium, and/or an optical medium.

[039] The management server 108 first receives user specification of failure conditions as entered by a user at the management console 106 (302). For instance, the user may specify the types of failures of the servers 104 to be predicted or detected by the management server 108, by the agent computer programs 130 and 132, or by the network adapters 112 and 118. The user may further specify the courses of action to be taken by the management server 108 when such imminent failure of any of the servers 104 is detected. Most generally, the user specifies one or more conditions corresponding to predictive failure of any of the servers 104.

[040] Thereafter, one of the agent computer programs 130 and 132 or one of the network adapters 112 and 118 detects a potential coming failure at the server 104A or 104B (304). That is, one of the agent computer programs 130 and 132 predicts that either the server 104A or 104B may fail. For descriptive simplicity, the server 104B is presumed to be the failing server, such that the agent computer program 132 or the network adapter 118 detects that the server 104B may imminently fail. The agent computer program 132 or the network adapter 118 reports back to the management server 108,

indicating that the server 104B may fail.

[041] The management server 108 communicates with the coordinator server 110 and the agent 132 to determine which of the virtual machine computer programs 126 and 128 is running on the virtual host 124 (306). That is, the management server 108 receives alerts from the agent 132 or the network adapter 118 and correlates the server 104B with virtual host 124 through previous interface with the agent. The management server then passes to the coordinator server 110 the identity of virtual host 124. In response, the coordinator server 110 correlates the virtual machine computer programs 128A and 128B supported by the virtual host computer program 124. The coordinator server 110 reports this information to the management server 108.

[042] The virtual machine computer programs 128 are dynamically migrated from the virtual host computer program 124 running on the potentially failing server 104B to the virtual host computer program 122 running on the server 104A (308). In one embodiment, the management server 108 may communicate with the coordinator server 110 to cause the coordinator server 110 to accomplish such dynamic migration. Once the migration of the virtual machine computer programs 128 to the server 104A is complete, the coordinator server 110 may report back to the management server 108 that the migration has been successfully completed.

[043] The virtual host computer program 124 is then shut down (310). In one embodiment, the management server 108 may communicate with the coordinator server 110 to stop execution of the virtual host computer program 124. The potentially failing server 104B is finally shut down (312). In one embodiment, the management server 108 may remotely shut down the server 104B via its agent computer program 132. In another embodiment, the management server 108 may instruct the user at the management console 106 to manually shut down the server 104B.

[044] Ensuring power conservation by monitoring server resource utilization

[045] FIG. 4 shows a method 400 for achieving power conservation by monitoring server resource utilization, according to an embodiment of the invention. The method 400 may be performed in relation to the system 100 of FIG. 1 or FIG. 2. The method 400 is specifically described in relation to the system 100 of FIG. 2 for examples purposes only. Furthermore, like the method 300 of FIG. 3, the method 400 of FIG. 4 may be implemented as a computer program that can be stored on and/or executed from a computer-readable medium. Such a medium may include a volatile or non-volatile medium, a semiconductor medium, a magnetic medium, and/or an optical medium.

[046] The management server 108 first receives user specification of resource utilization conditions as entered by a user at the management console 106 (402). For instance, the user may specify the resource utilization threshold at which the servers 104 are considered to be utilizing a low amount of their resources, as detected by the

management server 108 and/or by the agent computer programs 204 and 206. The user may further specify the courses of action to be taken by the management server 108 when such low resource utilization of any of the servers 104 is detected. Most generally, the user specifies one or more conditions corresponding to resource utilization of any of the servers 104.

[047] Thereafter, one of the agent computer programs 204 or 206 detects resource utilization of the server 104A or 104B as being less than a threshold (404). That is, one of the agent computer programs 204 and 206 detects that either the server 104A or 104B has less than a threshold percentage of its resources being utilized. For descriptive simplicity, the server 104B is presumed to be the server having its resources utilized less than the threshold percentage, such that the agent computer program 206 detects this condition. The agent computer program 206 reports such detection back to the resource consumption monitor service 208 of the management server 108.

[048] The management server 108 communicates with the coordinator server 110 and the agent 132 to determine which of the virtual machine computer programs 126 and 128 is running on the virtual host 124 (406). That is, the management server 108 passes to the coordinator server 110 the identity of virtual host 124. In response, the coordinator server 110 correlates the virtual machine computer programs 128A and 128B with the virtual host computer program 124. The coordinator server 110 reports this information to the management server 108.

[049] The virtual machine computer programs 128 are dynamically migrated from the virtual host computer program 124 running on the server 104B to the virtual host computer program 122 running on the server 104A (408). In one embodiment, the management server 108 may communicate with the coordinator server 110 to cause the coordinator server 110 to accomplish such dynamic migration. Once the migration of the virtual machine computer programs 128 to the server 104A is complete, the coordinator server 110 may report back to the management server 108 that the migration has been successfully completed. Communication between the coordinator server 110 and the management server 108 may be accomplished via the agent computer program 210 of the coordinator server 110.

[050] The virtual host computer program 124 is then shut down (410). In one embodiment, the management server 108 may communicate with the coordinator server 110 to stop execution of the virtual host computer program 124. The server 104B is finally shut down (412). In one embodiment, the management server 108 may remotely shut down the server 104B via its agent computer program 132 or the management module 212. In another embodiment, the management server 108 may instruct the user at the management console 106 to manually shut down the server 104B.

[051] Ensuring power conservation by monitoring server power consumption

[052] FIG. 5 shows a method 500 for achieving power conservation by monitoring server power consumption, according to an embodiment of the invention. The method 500 may be performed in relation to the system 100 of FIG. 1 or FIG. 2. The method 500 is specifically described in relation to the system 100 of FIG. 2 for examples purposes only. Furthermore, like the method 300 of FIG. 3 and the method 400 of FIG. 4, the method 500 of FIG. 5 may be implemented as a computer program that can be stored on and/or executed from a computer-readable medium. Such a medium may include a volatile or non-volatile medium, a semiconductor medium, a magnetic medium, and/or an optical medium.

[053] The management server 108 first receives user specification of power consumption conditions as entered by a user at the management console 106 (502). For instance, the user may specify the power consumption threshold at which the servers 104 are considered to be utilizing a low amount of power, as detected by the management server 108, by the agent computer programs 204 and 206, or by the management module 212. The user may further specify the courses of action to be taken by the management server 108 when such low power consumption of any of the servers 104 is detected. Most generally, the user specifies one or more conditions corresponding to power consumption of any of the servers 104.

[054] Thereafter, one of the agent computer programs 204 or 206 or the management module 212 detects power consumption of the server 104A or 104B as being less than a threshold (504). That is, one of the agent computer programs 204 and 206 detects that either the server 104A or 104B has less than a threshold amount of power being consumed. For descriptive simplicity, the server 104B is presumed to be the server consuming less than the threshold amount of power, such that the agent computer program 206 detects this condition. The agent computer program 206 reports such detection back to the resource consumption monitor service 208 of the management server 108.

[055] The management server 108 communicates with the coordinator server 110 and the agent 132 to determine which of the virtual machine computer programs 126 and 128 is running on the virtual host 124 (506). That is, the management server 108 passes to the coordinator server 110 the identity of virtual host 124. In response, the coordinator server 110 correlates the virtual machine computer programs 128A and 128B with the virtual host computer program 124. The coordinator server 110 reports this information to the management server 108.

[056] The virtual machine computer programs 128 are dynamically migrated from the virtual host computer program 124 running on the server 104B to the virtual host computer program 122 running on the server 104A (508). In one embodiment, the

management server 108 may communicate with the coordinator server 110 to cause the coordinator server 110 to accomplish such dynamic migration. Once the migration of the virtual machine computer programs 128 to the server 104A is complete, the coordinator server 110 may report back to the management server 108 that the migration has been successfully completed. Communication between the coordinator server 110 and the management server 108 may be accomplished via the agent computer program 210 of the coordinator server 110.

[057] The virtual host computer program 124 is then shut down (510). In one embodiment, the management server 108 may communicate with the coordinator server 110 to stop execution of the virtual host computer program 124. The server 104B is finally shut down (512). In one embodiment, the management server 108 may remotely shut down the server 104B via its agent computer program 132 or the management module 212. In another embodiment, the management server 108 may instruct the user at the management console 106 to manually shut down the server 104B.

[058] It is noted that, although specific embodiments have been illustrated and described herein, it will be appreciated by those of ordinary skill in the art that any arrangement calculated to achieve the same purpose may be substituted for the specific embodiments shown. This application is intended to cover any adaptations or variations of embodiments of the present invention. It is manifestly intended that this invention be limited only by the claims and equivalents thereof.

Claims

- [001] A system comprising: a plurality of server computing devices communicatively coupled to a network, each server computing device having a virtual host computer program executable thereon to support one or more virtual machine computer programs, each virtual machine computer program operable to execute an instance of an operating system on which application computer programs are executable; and a management server computing device communicatively coupled to the network to monitor the server computing devices and cause the virtual machine computer programs supported by the virtual host computer program of one of the server computing devices to migrate to the virtual host computer program of another of the server computing devices upon one or more conditions being satisfied.
- [002] The system of claim 1, wherein the management server computing device is operable to cause the virtual machine computer programs supported by the virtual host computer program of one of the server computing devices to dynamically migrate to the virtual host computer program of another server computing device, such that the virtual machine computer programs supported by the virtual host computer program are migrated without restarting.
- [003] The system of claim 1, wherein the one or more conditions comprises one of the server computing devices having the virtual host computer program from which the virtual machine computer programs are to be migrated being predicted as failure prone.
- [004] The system of claim 1, wherein the one or more conditions comprises the one of the server computing devices having the virtual host computer program from which the virtual machine computer programs are to be migrated being determined as having resource utilization less than a threshold.
- [005] The system of claim 1, wherein the one or more conditions comprises one of the server computing devices having the virtual host computer program from which the virtual machine computer programs are to be migrated being determined as having power consumption less than a threshold.
- [006] The system of claim 1, further comprising a coordinator server computing device communicatively coupled to the network to coordinate the virtual machine computer programs supported by the virtual host computer program of each server computing device, such that the management server computing device is operable to cause the virtual machine computer programs supported by the virtual host computer program of the one of the server computing devices to migrate to the virtual host computer program of the other of the server computing

- devices by sending commands to the coordinator server computing device.
- [007] The system of claim 6, further comprising a management agent executable on the management server computing device through which the management server computing device communicates with the coordinator server computing device.
- [008] The system of claim 1, further comprising a management agent executable on the virtual host computer program of each server computing device, through which the management server computing device is operable to monitor and control the server computing device.
- [009] The system of claim 8, wherein the management server computing device is operable to cause the virtual machine computer programs supported by the virtual host computer program of the one of the server computing devices to migrate to the virtual host computer program of the other of the server computing devices through the management agent computer program running on the one of the server computing devices and the management agent computer program running on the other of the server computing devices.
- [010] The system of claim 8, wherein the management agent is operable to monitor predictive failure of the server computing device for the management server computing device.
- [011] The system of claim 1, further comprising a resource consumption monitor agent executable on the virtual host computer program of each server computing device, through which the management server computing device is operable to monitor at least one of power consumption and resource utilization of the server computing device.
- [012] The system of claim 11, further comprising a resource consumption monitor service computer program executable on the management server computing device through which the management server computing device communicates with the resource consumption monitor agent computer program running on the virtual host computer program of each server computing device.
- [013] The system of claim 1, further comprising a remote supervisor adapter for each server computing device communicatively coupled to the network to enable the management server computing device to power down the server computing device as needed.
- [014] A method comprising: predicting a failure at a first server computing device; correlating the first server computing device with a virtual host computer program supporting one or more virtual machine computer programs; and dynamically migrating the virtual machine computer programs from the virtual host computer program of the first server computing device to a virtual host computer program of a second server computing device.

- [015] The method of claim 14, further comprising initially specifying one or more conditions corresponding to the failure predicted at the first server computing device.
- [016] The method of claim 14, further comprising: stopping execution of the virtual host computer program of the first server computing device; and turning off the first server computing device.
- [017] The method of claim 14, wherein dynamically migrating the virtual machine computer programs comprises sending a command from a management server computing device to a coordinator server computing device able to control the virtual host computer programs of the first and the second server computing devices.
- [018] A computer program comprising program code for carrying out the steps of any of claims 14 to 17.

FIG 1

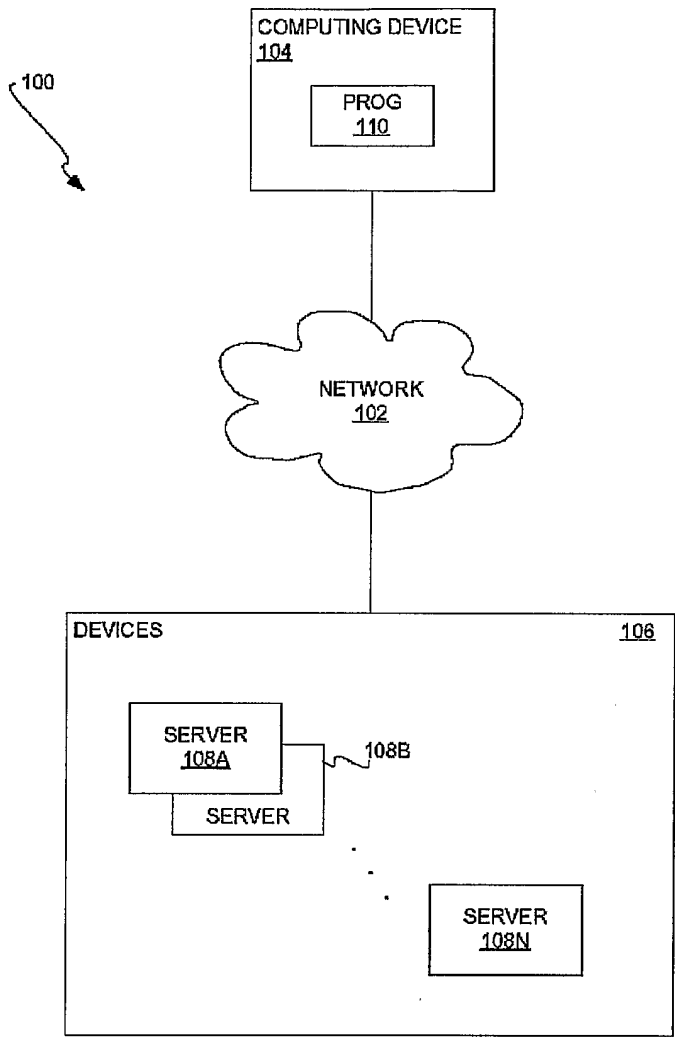



FIG 2



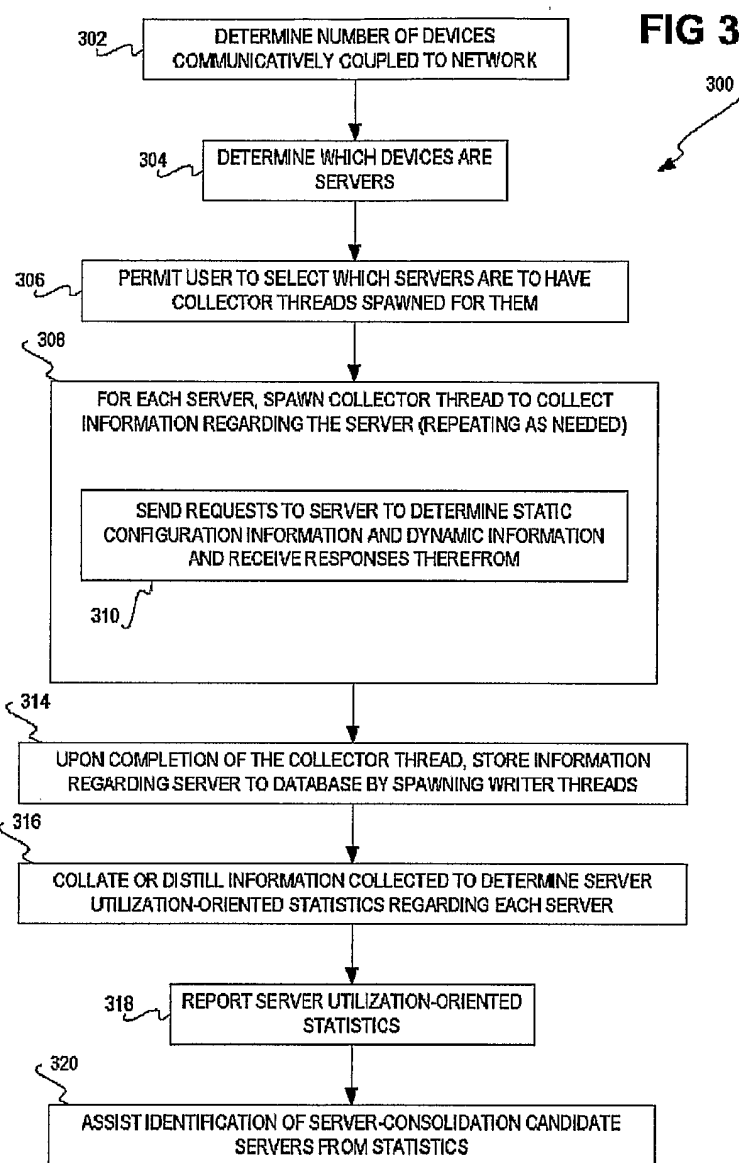


FIG 4

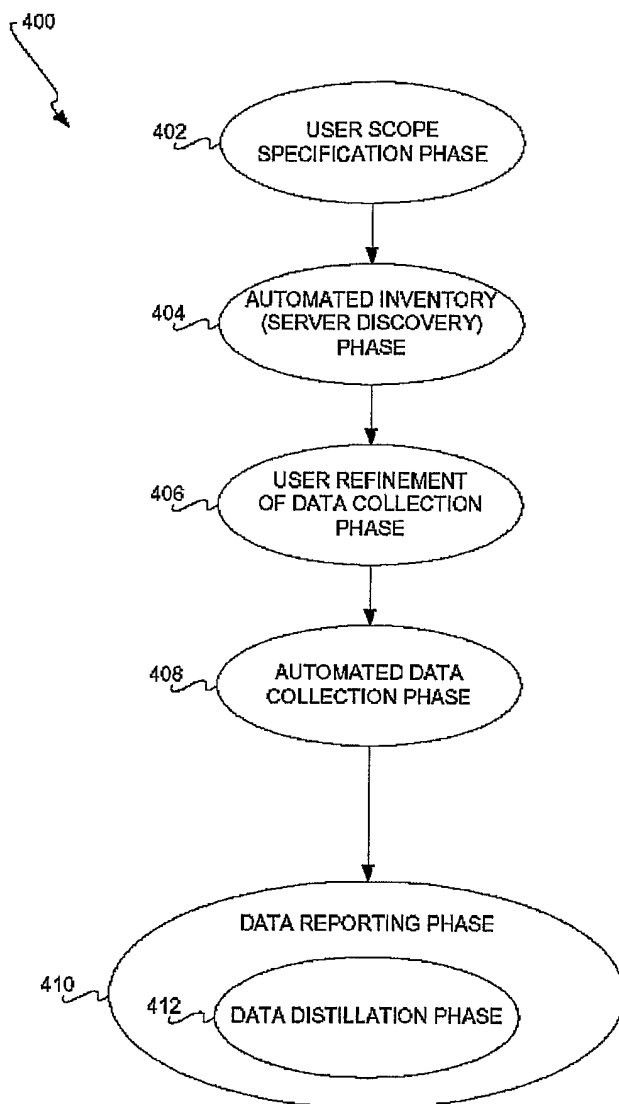


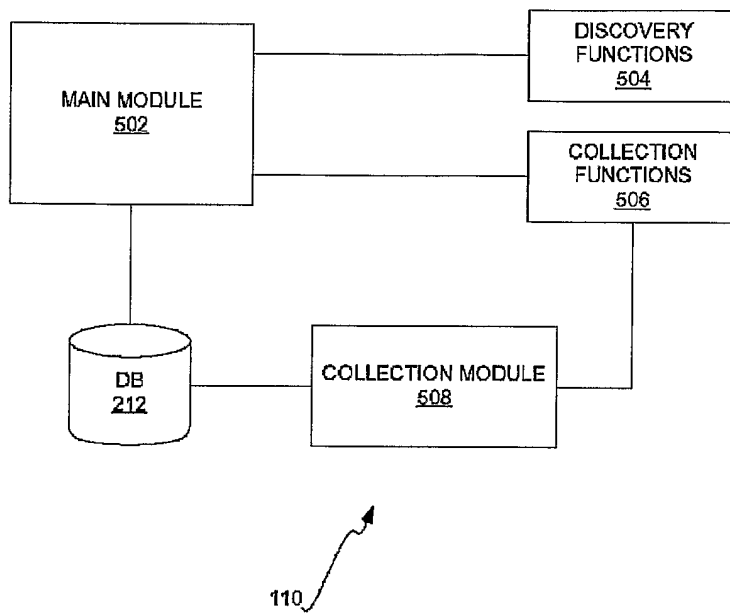
FIG 5

FIG 6

