

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

(19) World Intellectual Property Organization International Bureau



(43) International Publication Date
25 March 2004 (25.03.2004)

PCT

(10) International Publication Number
WO 2004/025471 A2

(51) International Patent Classification⁷: G06F 11/08 (74) Agent: DE PENA, Alain; Compagnie IBM France, Direction de la Propriété Intellectuelle, F-06610 La Gaude (FR).

(21) International Application Number: PCT/EP2003/010080 (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.

(22) International Filing Date: 13 August 2003 (13.08.2003) (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(25) Filing Language: English (26) Publication Language: English

(30) Priority Data: 10/241,213 11 September 2002 (11.09.2002) US

(71) Applicant: INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; New Orchard Road, Armonk, NY 10504 (US). (71) Applicant (for MC only): COMPAGNIE IBM FRANCE [FR/FR]; Tour Descartes-La Défense 5, 2, Avenue Gambetta, F-92400 Courbevoie (FR).

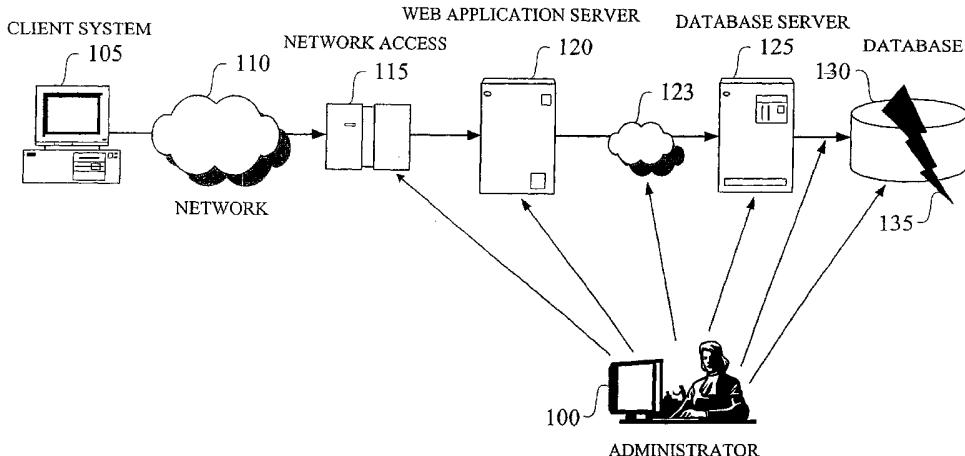
(72) Inventors: KELLER, Alexander; 301 W. 22nd Street, Apt. 48, New York, NY 10011 (US). KAR, Gautam; 211 Sara Court, Yorktown Heights, NY 10598 (US).

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: METHODS AND APPARATUS FOR ROOT CAUSE IDENTIFICATION AND PROBLEM DETERMINATION IN DISTRIBUTED SYSTEMS



(57) Abstract: A technique for determining a root cause of a condition (e.g., service outage) of at least one subject component in a computing environment comprises the following steps/operations. First, one or more components in the computing environment upon which the at least one subject component depends (e.g., antecedents) are identified. Identification comprises traversing at least a portion of a model representative of an existence of one or more relationships associated with at least a portion of components of the computing environment and which is capable of accounting for a full lifecycle (e.g., including deployment, installation and runtime) associated with at least one component of the computing environment. Then, one or more procedures are performed in accordance with the one or more identified components to determine a condition status associated with each of the one or more identified components. By way of example, the inventive techniques may be applied to a distributed computing environment. The computing environment may also be an autonomic computing environment.

WO 2004/025471 A2

METHODS AND APPARATUS FOR ROOT CAUSE IDENTIFICATION AND
PROBLEM DETERMINATION IN DISTRIBUTED SYSTEMS

Cross Reference to Related Applications

The present application is related to the 5 concurrently-filed U.S. patent applications respectively identified as: attorney docket no. YOR920020097US1 entitled: "Methods And Apparatus For Managing Dependencies in Distributed Systems;" attorney docket no. SOM920020003US1 entitled: "Methods And Apparatus For Topology Discovery and 10 Representation of Distributed Applications and Services;" attorney docket no. SOM920020004US1 entitled: "Methods And Apparatus For Impact Analysis and Problem Determination;" and attorney docket no. SOM920020005US1 entitled: "Methods And 15 Apparatus For Dependency-based Impact Simulation and Vulnerability Analysis;" the disclosures of which are incorporated by reference herein.

Field of the Invention

The present invention relates to distributed computing systems and, more particularly, to methods and apparatus for 20 analyzing and determining the root cause of a service outage based on dependencies between the various components of such distributed computing systems.

Background of the Invention

The identification and tracking of dependencies between 25 the components of distributed systems is becoming increasingly important for integrated fault management. Applications, services and their components rely on a variety of supporting services that might be outsourced to a service provider. Moreover, emerging web-based (world wide web-based) business

- 2 -

architectures allow the composition of web-based e-business (electronic business) applications at runtime.

It is to be understood that the term "runtime" generally refers to the time period when a piece of software is being 5 executed and active in a computer system's memory, as opposed to being dormant and merely sitting in storage on a computer's hard drive. Thus, being able to compose e-business applications at runtime means having the capability to do so without the need to bring down and restart the 10 system/application and without the need to recompile the application. Traditionally, the lifecycle of a computer program is: write program code -> compile (translate into machine code) -> run. Thus, with the above capability, one can assemble several pieces of software to form a new 15 application "on-the-fly," i.e., without the need to bring down/compile/restart the application.

Consequently, however, failures occurring in one service affect other services being offered to a customer, i.e., services have dependencies on other services. Dependencies 20 exist between the components of different services on a single system and also between the client and server components of a service across multiple systems and domains. Herein, services that depend on other services are referred to as dependents, while services on which other services depend are referred to 25 as antecedents.

It is important to note that a service often plays both roles (e.g., a name service is required by many applications and services but depends, itself, on the proper functioning of other services, such as the operating system and the network 30 protocols and infrastructure). Furthermore, dependency relationships are transitive, i.e., the dependent of a given component requires, in addition to the component itself, the components' antecedent(s).

Dependencies exist between various components of a 35 distributed system, such as end-user services, system services, applications and their logical and physical

- 3 -

components. However, service dependencies are not made explicit in today's systems, thus making the task of problem determination, isolation and resolution particularly difficult.

5 Existing art in the area of software development (such as U.S. Patent No. 4,751,635 and U.S. Patent No. 5,960,196), maintenance (such as U.S. Patent No. 5,493,682) and software packaging (such as U.S. Patent No. 5,835,777) deal with individual software elements and modules that form the atomic
10 parts of a program package and require the availability of program source code in order to build software and bundle it into software products. Source code is available to the software developer and not to the service user. The invention primarily focuses on software products that are already
15 packaged.

The Institute of Electrical and Electronics Engineers Standard 1387.2 (entitled "Portable Operating System Interface (POSIX) system administration, part 2: Software Administration," IEEE, 1995) addresses software distribution/deployment/installation. The IEEE standard defines a mechanism for ensuring that new software components (which are going to be installed) do not conflict with an already existing software installation. The IEEE standard identifies three kinds of relationships: prerequisite, exrequisite, corequisite, that facilitate such compatibility checks. This is done individually for every system on which new software needs to be installed. With the IEEE standard, the software inventories present on other systems are not taken into account. Furthermore, the IEEE standard does not deal with instantiated applications and services and therefore does not represent any means of determining the dependencies between components at runtime.

Open Group (Systems Management: Distributed Software Administration, CAE Specification C701, The Open Group, 35 January 1998) extends IEEE 1387.2 by defining several commands (swinstall, swlist, swmodify, etc.) that are invoked by

- 4 -

software installation tools on a specific system. Open Group also defines a software definition file format to make sure that the information required by the aforementioned commands is available from the system on which the commands are invoked. The shortcomings of IEEE 1387.2 (i.e., confined to a single isolated system, no means for determining software dependencies at runtime) also apply to the Open Group specification.

Current Operating System Inventory implementations (such as the IBM AIX Object Data Manager (ODM), the Linux Red Hat Package Manager (RPM) or the Microsoft Windows Registry) follow either the OpenGroup specification and the IEEE 1387.2 standard or describe the software inventory in a proprietary format. Thus, the aforementioned limitations also apply to such Current Operating System Inventory implementations.

Techniques for electronic software distribution of whole program packages (such as U.S. Patent No. 6,009,525 and U.S. Patent No. 5,721,824) or updates/corrections/ fixes/patches (such as U.S. Patent No. 5,999,740, U.S. Patent No. 5,805,891, and U.S. Patent No. 5,953,533) are, by definition, restricted to the distribution/deployment/ installation of (one or many at a time) physical software packages and do not take the runtime stages of applications into account. In addition, they deal with one system at a time and do not take the cross-system aspects of applications and services into account.

Techniques for determining conflicts in existing software/hardware configurations (such as U.S. Patent No. 5,867,714) are also confined to a single system and do not take runtime aspects into account.

While existing work (such as U.S. Patent No. 5,917,831), often within the scope of event correlation (see, e.g., Gruschke et al., "Integrated Event Management: Event Correlation Using Dependency Graphs, DSOM '98, 1998 and Kätke et al., "Fault Isolation and Event Correlation for Integrated Fault Management, IM '97, 1997), has focused on identifying

- 5 -

and describing service dependencies in a proprietary format, it has remained unclear how dependency information can actually be exchanged between different entities of the fault management process. Since it is unlikely that the different 5 parties involved in the fault management process of outsourced applications use the same toolset for tracking dependencies, it is of fundamental importance to define an open format for specifying and exchanging dependency information.

Also, due to the heterogeneity associated with components 10 of the distributed system with which the fault management process is involved, determining the root cause of a system failure (e.g., service outage) is extremely difficult, given the limitations of existing techniques.

To sum up, a few techniques relating to the determination 15 of relationships between software products have been described and implemented in the existing art. These existing techniques suffer from one or more of the following shortcomings:

(a) they address only the installation and deployment 20 phases of a software product; i.e., they do not attempt to capture the design and runtime aspects;

(b) they do not deal with end-to-end applications and services that span multiple systems; i.e., they address the 25 characteristics of software residing on a single, isolated system;

(c) software inventory information is described in a proprietary format that makes it extremely difficult to share this information among various heterogeneous systems; and

(d) they do not effectively identify the root cause of a 30 service outage.

Summary of the Invention

The present invention provides techniques for identifying the root cause of a component failure and performing appropriate problem determination procedures in accordance 35 with a computing environment. By way of example, the

- 6 -

inventive techniques may be applied to a distributed computing environment. The computing environment may also be an autonomic computing environment.

For example, in one aspect of the invention, a computer-based technique for determining a root cause of a condition (e.g., service outage) of at least one subject component in a computing environment comprises the following steps/operations. First, one or more components in the computing environment upon which the at least one subject component depends (e.g., antecedents) are identified. Identification comprises traversing at least a portion of a model representative of an existence of one or more relationships associated with at least a portion of components of the computing environment and which is capable of accounting for a full lifecycle (e.g., including deployment, installation and runtime) associated with at least one component of the computing environment.

Then, one or more procedures are performed in accordance with the one or more identified components to determine a condition status associated with each of the one or more identified components. The procedures may be carried out either stepwise or in combination, and may include, by way of example, process inspection, exercising, heartbeats, and status indicators.

By way of example, a component may be a service, an application, middleware, hardware, a device driver, an operating system or a system associated with the computing environment. However, the term "component" is not limited to these examples.

These and other objects, features and advantages of the present invention will become apparent from the following detailed description of illustrative embodiments thereof, which is to be read in connection with the accompanying drawings.

35 **Brief Description of the Drawings**

- 7 -

FIG. 1 is a block diagram illustrating an example of a client-server application architecture with which features of the present invention can interact to produce information;

5 FIG. 2A is a block diagram illustrating a system for providing dependency management according to an embodiment of the present invention;

10 FIG. 2B is a block diagram illustrating a generalized hardware architecture of a computer system suitable for implementing a system for providing dependency management according to an embodiment of the present invention;

FIG. 3 is a block diagram illustrating a functional dependency model of services according to an embodiment of the present invention;

15 FIG. 4 is a block diagram illustrating a structural dependency model of services according to an embodiment of the present invention;

20 FIG. 5 is a block diagram illustrating a service lifecycle addressed by functional, structural and operational dependency models according to an embodiment of the present invention;

FIG. 6 is a block diagram illustrating relationships between functional, structural and operational dependency models according to an embodiment of the present invention;

25 FIG. 7 is a block diagram illustrating components involved in analyzing and computing a root cause of a service outage according to an embodiment of the present invention;

FIG. 8 is a block diagram illustrating components of a root cause analyzer according to an embodiment of the present invention;

30 FIG. 9 is a flow diagram illustrating steps for invoking a dependency service and performing root cause analysis on an operational model according to an embodiment of the present invention;

35 FIG. 10 is a flow diagram illustrating tasks of an administrator for creating and updating a functional

- 8 -

dependency model according to an embodiment of the present invention;

FIG. 11 is a flow diagram illustrating steps for updating a structural dependency model by installing or removing hardware/software components on a computer system according to an embodiment of the present invention;

FIG. 12 is a flow diagram illustrating performance of root cause analysis on an operational model according to an embodiment of the present invention;

FIG. 13 is a flow diagram illustrating performance of root cause analysis on antecedents of a service according to an embodiment of the present invention;

FIG. 14 is a flow diagram illustrating steps for determining a status of a service according to an embodiment of the present invention; and

FIG. 15 depicts examples of root cause analyzer application programming interfaces according to an embodiment of the present invention.

Detailed Description of Preferred Embodiments

The present invention will be explained below in the context of an illustrative distributed computing environment. However, it is to be understood that the present invention is not limited to such a particular computing environment. Rather, the invention is more generally applicable to any computing environment in which it is desirable to manage (e.g., compute, query, etc.) dependencies in order to make the task of problem determination, isolation and resolution significantly easier.

As used herein, depending on the context of the discussion, the term "system" may be used to refer to a computer system, a software system and/or some combination thereof. The term "system" may also be used to refer to an application and/or a service. Thus, the phrase "multiple systems" refers to a collection of several systems. Also, the

- 9 -

term "component" may refer to a system, itself, or one or more portions of a system.

As mentioned above, service dependencies are not made explicit in today's systems, thus making the task of problem determination, isolation and resolution particularly difficult. Solving this problem requires the determination and computation of dependencies between services and applications across different systems and domains, i.e., establishing a "global" service dependency model and enabling system administrators to navigate through the resulting directed graph from the top to the bottom and in reverse order. The need for such a mechanism is best illustrated by the following two scenarios.

The first scenario deals with managing outsourced services, typically offered by Internet or Application Service Providers (ISP/ASP). Outsourcing services leads to layered service hierarchies where, e.g., the services of an ASP depend on the IP-connectivity (Internet Protocol-connectivity) offered by an ISP, which, in turn, relies on the wide area network of a telecom carrier. At every layer, a service is accessed through a Service Access Point (SAP). A SAP delimits the boundary between the different organizational domains and is the place where Service Level Agreements (SLAs) are defined and observed. Usually, this is done at every layer by monitoring a set of specific parameters that are exposed by the provider. In case of an outage or performance degradation in an upper-layer service, it is necessary to traverse the service hierarchy from the top to the bottom to identify the root cause of the problem.

The second scenario deals with the regular maintenance tasks that can not be done "on-the-fly" and therefore affect services and their customers: e.g., e-mail servers get updated with a new release of their operating system, network devices are exchanged or upgraded with a new firmware version, etc. In all cases, it is important for the network and server administrators to determine in advance how many and, more

- 10 -

specifically, which services and users are affected by the maintenance. We refer to this task as impact analysis.

The aforementioned tasks are further exacerbated by the following factors.

5 Dependency models provide a straightforward means to identify possible root causes of an observed problem. If the dependency graph for a system is known, navigating the graph from an impaired service towards its antecedents (being either co-located on the same host or on different systems) will 10 reveal which entities might have failed. Traversing the graph towards its root (i.e., in upwards direction) yields the dependents of a service, i.e., the components that might fail if this service experiences an outage. The following issues need to be addressed.

15 (a) Scale: The number of dependencies between many involved systems can be computed, but may become very large. From an engineering viewpoint, it is often undesirable (and sometimes impossible) to store a complete, instantiated dependency model at a single place. Traditional mechanisms 20 used in network management platforms such as keeping an instantiated network map in the platform database therefore can not be applied to dependencies due to the sheer number and the dynamics of the involved dependencies.

25 These two facts make it prohibitive to follow a "network-management-style" approach for the deployment of application, service and middleware dependency models. As an example, typical data centers of service outsourcees host vast amounts (several thousands) of web application and database servers. This implies a huge number of simultaneously running 30 program instances of, e.g., web application and database servers. A system capable of constructing a dependency model should provide features that allow appropriate scalability by distributing the storage and computation of dependencies across the systems involved in the management process.

35 (b) Dynamics: Hosted applications (running within web application servers) have a very short lifetime, often only a

- 11 -

few seconds. Upon reception of a request, the business logic of a web application (usually implemented as one or more Java Servlets) gets instantiated by the servlet engine of the application server, performs its tasks and is then removed by 5 the servlet engine. Consequently, a system for computing dependencies among these dynamic entities should address the trade-off between the accuracy of the data and the workload generated for retrieving this data.

10 (c) Heterogeneity: Heterogeneity comes in three different flavors. First, the services offered to customers differ to a large degree. Second, there may be various providers involved in offering a service to a customer. Finally, the products implementing a service may stem from a variety of vendors. A 15 system for computing dependencies should provide a language that is independent of specific operating systems, network protocols, software products and the services offered to a customer.

20 (d) Manual maintenance of Dependency Data: The acquisition of a service dependency model, even confined to a single host system, is a challenge on its own as today's 25 systems usually do not provide appropriate management instrumentation. It is to be understood that the term "instrumentation" refers to program code that exposes the management characteristics and capabilities of a (managed) resource through a well-defined (sometimes even standardized) 30 interface so that it can be accessed by management applications. Further, even if available from managed resources, dependency data is not exploited by today's management systems. Instead, the dependency information not 35 only has to be entered manually into a specific management component, but also in a proprietary format. The dependency information is therefore incomplete, outdated (due to the error-prone manual processing), and sometimes even inconsistent because different operators enter the rules independently and there is no way to check the rule base for consistency in an automated way.

- 12 -

(e) Taxonomy for Dependencies: The notion of dependencies is very coarse and needs to be refined in order to be useful. Examples for this are the strength of a dependency (indicating the likelihood and the degree to which a component is affected 5 if its antecedent fails), the criticality (how important this dependency is with respect to the goals and policies of an enterprise), the degree of formalization (i.e., how difficult it is to obtain the dependency) and many more. There is a need to add attributes to dependencies that allow them to be 10 qualified more appropriately; and accordingly, a need to reflect these attributes in the dependency representation.

(f) Problem Determination Features: Further facilities for combining local dependency graphs, stored on every system, into a uniform dependency model are desired. In addition, 15 these facilities should provide an API (application programming interface) allowing management applications to issue queries against the dependency model. These queries will be allowed to retrieve the entities that a specific service directly depends on, or to recursively determine a 20 whole set of nodes, including sub-antecedents. The list of nodes received by the management application enables it to perform specific problem determination routines to check whether these services are operational.

The previous discussion shows that it is important to 25 establish a mapping between three different stages of a service lifecycle:

- (a) an (abstract) service being offered to customers, e.g., "Web Hosting," "Managed Storage," "IP Connectivity," "Managed Database," etc.;
- 30 (b) the implementation of a service, i.e., the product(s) being used for providing the service, e.g., "IBM Universal Database version 7.1," "WebSphere Application Server version 3.2;" and
- 35 (c) the running instance(s) of an implementation, i.e., the process or task, e.g., "db2 daemon," "nfs daemon."

- 13 -

While the task of individually obtaining information available at every single stage is feasible, the combination of the three stages into a uniform dependency model is challenging and has not been done in previous work. In

5 addition, there is a need for establishing an efficiently computable dependency model that addresses the requirements of scale, dynamics and heterogeneity of the underlying environment while eliminating the need for human interaction and maintenance of dependency data.

10 As will be illustrated below in the context of the figures, the present invention addresses these and other needs. That is, the present invention has features which compute runtime dependencies (a "dependency model") between components of a distributed system on behalf of a management 15 application. The invention offers a generic and uniform approach for retrieving dependency information from computer systems that provide a mechanism for retrieving configuration information of individual computer systems or that provide such data in a machine-readable format.

20 One benefit of the aforementioned system is that a large amount of application/service management information can be obtained from these computer systems without having the need to instrument individual applications/services. However, if such application/service instrumentation is available, it can 25 be used by the present invention.

The execution of the system described by the present invention may be triggered either by a specific (management) application (such as: impact analyzer, root cause analyzer), a network management platform (such as IBM/Tivoli NetView, HP 30 OpenView or Aprisma Spectrum) or a management application based on traditional network management systems and platforms.

The present invention provides, *inter alia*, features for:

(a) observing performance degradations and outages of subscribed services;

35 (b) tracking down the root cause of the problem by traversing the different layers of the dependency model from

- 14 -

the top to the bottom (since various services may be outsourced to other service providers, this (recursive) traversal of a dependency model crosses domain boundaries); and

5 (c) analyzing the impact of a service outage by navigating the dependency model from the bottom to the top.

The invention combines dependency information that is available during the lifecycle of an application or service (i.e., from the design to deployment, installation and runtime 10 stages of an application/service). This information is kept within the following models:

(a) Functional Model: In a preferred implementation, the functional model defines dependencies between different generic services (database service, name service, web 15 application service, connectivity service, etc.). The functional model does not describe client/server relationships within a specific service. In addition, the functional model neither takes into account which concrete products have been chosen to implement the services nor their actual 20 configuration. The functional model establishes the principal constraints to which the other models (described below) are bound, i.e., further models may refine the dependencies defined in the functional model with respect to a concrete system infrastructure but should not introduce new 25 dependencies between service categories. The model is very compact and generic and is preferably stored on the managing system.

(b) Structural Model: In a preferred implementation, the structural model contains the detailed descriptions of 30 software components that realize the services defined in the functional model. The structural model provides details captured during the installation/deployment phase and complements the functional model by taking the software inventory of concrete systems into account. The structural 35 model provides information about which services are installed and configured on a specific system and, for every service,

- 15 -

whether the system operates in a client or a server role. The potentially high numbers of systems and services make it difficult to track these dependencies from a remote location. It is thus desirable to store this model close to or at the 5 managed resources.

(c) Operational Model: In a preferred implementation, when software packages get instantiated and bindings between services and applications are established, the operational model of dependencies is created. The high degree of dynamics 10 of this model and large number of involved systems places a limitation on the extent to which the complete model can be instantiated and stored. It is not practical to define and store such a model, rather, the model has to be computed 15 dynamically and stepwise. The operational model is therefore computed "on demand" and relies on the functional and structural models.

As is to be expected, the amount of dependencies and their dynamics are extremely high in large-scale distributed 20 systems. The features of the present invention keep their impact on the distributed system (in terms of resource and bandwidth usage) as small as possible and leave as many 25 configuration options that might affect the performance up to the user. Examples for this are: the time intervals for retrieving an updated dependency model, the range of the systems whose dependencies should be tracked, the depth of the dependency model (only services immediately affected versus 30 transitive closure for a given service versus whole service hierarchy).

The present invention preferably exploits the following 35 characteristics of dependency information:

(a) Dependencies between different services are layered. Furthermore, their dependency graph is directed and acyclic. The latter statement also reflects experiences with IP-based networked services, such as DNS (Domain Name System), NFS 35 (Network File System), DFS (Distributed File System), NIS (Network Information System), etc., but there may be cases

- 16 -

where mutual dependencies may occur in some systems. A pathological example for such a mutual dependency is a DNS server that mounts the file system in which its DNS configuration is stored via NFS from a remote system. While 5 such a configuration is technically feasible, it reflects flaws in the system design because this leads to an unstable system whose bootstrapping might be non-deterministic and thus should be avoided. A dependency-checking application that discovers cyclic dependencies should issue a warning to an 10 administrator.

(b) Every dependency is visible at a customer/provider domain boundary and made explicit by means of SLAs. It follows that the number of observable dependencies is finite.

15 (c) Dependency models allow a top-down traversal of dependency chains.

(d) Dependencies between different systems ("inter-system") are perceived as dependencies between the client and server parts of the same service. It is not possible that a client for service A issues requests to a 20 server which provides a different service B.

One goal of the present invention is to retrieve 25 information mainly from a few well-known/well-defined places (e.g., system repositories) in order to achieve a maximum degree of independence from concrete service/application instrumentation. In order to achieve this, the invention defines a minimal and sufficient amount of commonly available dependency information.

The invention comprises facilities for storing dependency 30 models persistently or leaves this to the discretion of a management application or another service that uses the invention.

The invention is enabled to have a notion of history in order to detect and determine changes in the dependency model. In this case, the invention provides a publish/subscribe 35 interface for notifying software components that have previously registered for changes within the dependency model.

- 17 -

Another possible use of the invention is to leave the detection of changes in the dependency model up to the discretion of a management application (or a change management service) to issue periodic calls to the invention in order to 5 determine whether changes in the dependency model have occurred.

Furthermore, as will be illustrated herein below, the present invention provides techniques for identifying the root cause of a service outage and performing appropriate problem 10 determination procedures. Finding the root cause of a service outage involves the traversal of the service dependency graph from the top to the bottom to identify the candidate services that may have experienced a problem, which is then propagated to the service in question. This traversal from the service 15 where the outage or degradation has been observed towards its antecedents retrieves the entities that a specific service directly depends on, or recursively selects the complete set of nodes, including sub-antecedents, that this service requires to function properly. The list of nodes received by 20 a root cause analyzer enables it to perform, in a second step, specific problem determination routines to check whether these services are operational. These problem determination procedures involve determining whether a service is functioning properly or not. The problem determination 25 procedures may be carried out either:

(i) stepwise, i.e., the function test is carried out for every single service returned, or

(ii) combined, i.e., the whole list of antecedent services (or a subset of the operational model) is obtained 30 first, and problem determination procedures are carried out for all of them simultaneously. The order in which these tests are carried out may be determined by precedence values computed from the strength of dependencies between services, which, in turn, are expressed as weights associated with such 35 dependencies.

- 18 -

As mentioned, in order to determine whether a service is functioning properly, problem determination procedures may be used. Examples of problem determination procedures for distributed services and applications may include, but are not limited to:

(i) Process Inspection: For software applications that implement a service, one way to determine whether they are functioning properly is to verify if their processes (daemons) are running. This can be done in a non-intrusive way by inspecting the process table (or task list) of an operating system and does not require any instrumentation of the application. If the process(es) of an application are running, the application is considered to be in a healthy state.

(ii) Exercising an application is an intrusive (and more precise) way of determining whether the application is fully operational, i.e., running, and performing its business function. An "exerciser" is a transaction or command invoked from outside of the application that exercises the application in some fairly complete way to determine if it is really alive and able to deliver its functionality in a timely way. What is exercised is a function of the application itself. An analogy in networking is to test the connectivity via the ICMP (Internet Control Message Protocol) "ping" command, which sends time-stamped IP (Internet Protocol) packets to a resource. These packets are returned by the resource and thus allow a system to determine whether the resource is alive and to measure the round-trip delay.

(iii) Heartbeats allow the application to demonstrate that it is alive and well. The application regularly announces that it is in a healthy operational state by generating events automatically and repeatedly. The system listening for heartbeats must understand that if the timeout period expires between the heartbeat events then the application may not be functioning correctly.

- 19 -

(iv) Status indicators are specific management variables of an application that reflect the current status of an application implementing a service. Querying one or more status variables indicates the overall health of a service.

5 However, this approach requires the application to be instrumented appropriately by exposing this information to the outside.

Further, the root cause determination methodology of the invention may be performed for one or more subject components 10 substantially concurrently (e.g., for one component at a time, or for multiple components in parallel). At least a portion of results of the root cause determination methodology of the invention may be persistently stored, however, such results may not be persistently stored. Still further, a history of 15 results associated with the root cause determination methodology of the invention may be maintained. Such history may be used to derive heuristics for subsequently determining the sequence of steps for determining the root cause. For example, a ranking may be established, from the history, of 20 the most common outages and the next time a root cause analysis is invoked, the component that has been failing the most in the past is checked first.

Given the above realizations made in accordance with the present invention and general features associated with the 25 present invention, the remainder of the detailed description will provide an illustrative explanation of techniques for implementing such realizations and features in the context of FIGs. 1 through 15.

Referring initially to FIG. 1, a block diagram 30 illustrates an example of an electronic commerce system in the form of a client-server application architecture with which the features of the present invention can interact to produce information. The architecture of FIG. 1 will be described below to illustrate how such an architecture may handle a 35 transaction in the absence of the techniques of the invention.

- 20 -

As depicted, a client system 105 is used to initiate a request, for example via keyboard. Requests, however, could be initiated by any conventional means such as by mouse click, voice command, bar code swipe, etc. Examples of the client system 105 are personal computers, kiosks, data entry terminals, scanners, telephones, pagers, handheld or wearable devices, wireless devices, personal digital assistants, network-enabled watches, etc.

The request is acted upon locally where the request is formulated and forwarded to a web application server 120 over a network 110 and by passing through one or many network access 115 devices. An example of the network 110 and communication protocol are socket-based communications riding on a TCP/IP (Transmission Control Protocol/Internet Protocol) transport across a local area network (LAN) that is connected by network access 115 devices such as routers and switches to a wide area network (WAN) containing many switching locations that create a virtual circuit to a service provider and eventually to a web application server 120. Examples of a web application server 120 are high-end personal computers, RISC-based PowerPC's, UNIX-based workstations, minicomputers or mainframe computers running software fielding requests from clients and distributing the requests to appropriate back-end database servers when appropriate.

For illustrative purposes, an electronic commerce transaction initiated within a web browser (running on the client system 105) to purchase an item using the Internet will now be described. It is to be understood that the techniques of the invention may work with any form of transaction. Examples of web application servers include, but are not limited to, those available from IBM Corporation under the trademark WEBSPHERE, from BEA Systems, Inc. under the trademark WEBLOGIC, or from Lotus under the trademark LOTUS DOMINO SERVER.

In the example transaction, the business logic of the web application server 120 processes the incoming request and

- 21 -

provides authentication and/or identification of the client system 105. Once the business logic implemented by the web application server 120 determines that the client may proceed with the purchase, it communicates another request via a 5 network 123 to a database server 125 to decrement the inventory. The database server 125 processes the request, accesses its database 130 and prepares a response to the web application server 120. Examples of database servers include, but are not limited to those sold by Microsoft under the 10 trademark SQL/SERVER or TRANSACTION SERVER and by IBM Corporation under the trademark DB2 UNIVERSAL DATABASE SERVER.

The web application server 120 receives the response from the database server 125 and returns it via the network 110 to the client system 105. The client system 105 then processes 15 the response to format it for display and presents the response for the transaction initiator to review.

An administrator 100 observes the various software and hardware components, through which the business transaction is processed, located at the site of the service provider to 20 determine if they function properly. In case of an outage 135 occurring at the database 130, such as a corrupted tablespace or a failure of the database runtime system, the task of the administrator 100 is to locate the cause of the outage, correct the problem and verify if the overall system is 25 functioning properly again. It is to be understood that the invention is intended to operate with any form of outage or performance degradation.

The administrator 100 interacts with the software and hardware components either directly or through a management 30 system that processes the administrative information (such as status and health data) exposed by software and hardware components at a well-defined management interface. In either case, it is important to note that the hardware and software components are perceived by the administrator as isolated 35 resources and not as being part of an overall system serving a specific business purpose.

- 22 -

In particular, errors occurring in one component may not be noticed because the administrator does not become aware of them due to lack of continuous monitoring. In addition, absent the techniques of the invention, no explicit 5 information regarding the interdependencies between the various components is available to an administrator in a straightforward way. Thus, it is possible that an error within a component, which is not continuously monitored may go unnoticed until the failure propagates to a monitored 10 component.

In case of the aforementioned database outage 135, an administrator may eventually only become aware of the outage if the web application server 120 is not functioning properly anymore (for example, the load on the web application server 15 increases drastically because it continuously retries to connect to the database server 125, and is unable to complete the requests sent by the client system 105). Thus, the administrator 100 would first examine the web application server 120, then determine if there is a network 123 connectivity problem and finally verify if the database server 20 125 is experiencing difficulties that might stem from an internal error in the database 130.

The client-server application architecture described above can be regarded as a precursor to an emerging computing 25 environment referred to by IBM Corporation as an "autonomic" computing environment. P. Horn, "Autonomic Computing: IBM's Perspective on the State of Information Technology," IBM Research, October 2001, the disclosure of which is incorporated by reference herein, defines autonomic computing 30 as a comprehensive and holistic approach to self-managed computing systems with a minimum of human interference. The term derives from the body's autonomic nervous system, which controls key functions without conscious awareness or involvement. More specifically, one of the goals of autonomic 35 computing is to automate some or all of the tasks an

- 23 -

administrator 100 would typically carry out. Motivation for doing so is as follows.

As computing evolves, the overlapping connections, dependencies, and interacting applications call for administrative decision-making and responses faster than any human can deliver. Pinpointing root causes of failures becomes more difficult, while finding ways of increasing system efficiency generates problems with more variables than any human can hope to solve. The problem of identifying and tracking dependencies between different systems of an autonomic computing environment may be characterized in the following way. Since a system can exist at many levels, an autonomic system needs detailed knowledge of its components, current status, ultimate capacity, and all connections with other systems to govern itself. Those skilled in the art will appreciate that the present invention may be performed in an autonomic computing environment.

Referring now to FIG. 2A, a block diagram illustrates a system for providing dependency management according to an embodiment of the present invention. More specifically, FIG. 2A depicts a dependency management system that addresses the aforementioned issues. The system comprises four layers (application layer 200, service layer 205, middleware layer 210 and resource layer 215) and an administrator graphical user interface 285 by which an administrator 100 interacts with the system.

The lowest layer is the resource layer 215. The resource layer 215 comprises managed resources 220, a resource dependency repository 225 and a repository agent 230. Examples of managed resources 220 include, but are not limited to, physical and logical hardware components (examples of the former are hard disks, random access memory, central processing units, network adapters, channel controllers, etc.; examples of the latter are disk partitions, file systems, etc.) and software components (such as operating system,

- 24 -

system services like print spoolers or name services, and end-user applications).

The resource dependency repository 225 contains the inventory of the hardware and software components of every 5 managed resource 220 and the dependency information on a per-resource basis (i.e., the dependencies between the components within a managed resource 220). The resource dependency repository 225 can be either co-located with every individual managed resource 220 or reside at a centralized 10 location. The resource dependency repository 225 can be queried, updated and modified through a repository agent 230, which makes the information of the resource dependency repository 225 available to other components of the system.

The middleware layer 210 comprises a management 15 communication infrastructure 235 such as the protocols and object request brokers by which the different components of the system exchange (management) information.

The service layer 205 comprises various generic management services 250 such as policy, event and directory, 20 which can be used by a variety of management applications. A specifically important service is the dependency service 245, which retrieves information, both from the managed resources 220 and from the repository agent 230, and processes this information to establish an end-to-end dependency model of the 25 overall resource environment. This model (or parts of it) is stored in the end-to-end dependency repository 240, according to the needs (e.g., caching for faster retrieval) of the dependency service 245. Note that the dependency service 245 is the only component in the described system that directly 30 interacts with the end-to-end dependency repository 240.

It is to be appreciated that the above-mentioned dependency model and its parts may be generated in accordance with the techniques disclosed in the above-referenced and concurrently-filed U.S. patent application identified by 35 attorney docket no. Y0R920020097US1 entitled: "Methods And Apparatus For Managing Dependencies in Distributed Systems,"

- 25 -

some illustrative details of which are provided herein below. However, other model generation techniques may be employed.

The application layer 200 comprises various management applications that use the generic management services 250 and/or the dependency service 245. Examples of such management applications include, but are not limited to, fault manager 260, topology generator 265, impact analyzer 270, impact simulator 275 and root cause analyzer 280.

The root cause analyzer 280, as will be described further herein, determines the root cause of an outage (i.e., the component that initially caused the outage) based on traversing the dependency model (provided by the dependency service 245) from a component affected by an outage towards its antecedents.

The impact analyzer 270 determines the impact of an outage (i.e., the components that are likely to be affected by the outage) based on traversing the dependency model (provided by the dependency service 245) from a component experiencing an outage towards its dependents. The impact analyzer may employ the techniques disclosed in the above-referenced and concurrently-filed U.S. patent application identified by attorney docket no. SOM920020004US1 entitled: "Methods And Apparatus For Impact Analysis and Problem Determination." However, other impact analysis techniques may be employed.

The impact simulator 275, based on the impact analyzer 270, allows an administrator 100 to carry out a "what-if" analysis by simulating the effects of an outage of a specific component on the overall system. This enables the provisioning of appropriate failover solutions. The impact simulator may employ the techniques disclosed in the above-referenced and concurrently-filed U.S. patent application identified by attorney docket no. SOM920020005US1 entitled: "Methods And Apparatus For Dependency-based Impact Simulation and Vulnerability Analysis." However, other impact simulation techniques may be employed.

- 26 -

The fault manager 260 performs appropriate "sanity checks" or tests on components that have been identified either by the root cause analyzer 280 or the impact analyzer 270 as candidates for a failure. That is, the fault manager 5 can perform such tests upon the direction of the root cause analyzer 280 or the impact analyzer 270 (i.e., serve as an interface for these modules), and report back results thereto. However, the root cause analyzer 280 or the impact analyzer 270 can perform their own tests independent of the fault 10 manager.

It is to be understood that the fault manager is preferably comprised of a collection of application-specific or resource-specific tools that allow a determination to be made whether the component that is being tested is working 15 properly or not. Thus, after testing the component with a related tool, the fault manager may return a message indicating whether the component "works" or "does not work." These tools may be automated and/or manual. By way of one automated example, a so-called "ping" program checks network 20 connectivity. If the subject remote system answers to a ping, it is online and its network protocol stack (and all the underlying hardware, e.g., network adapter, cables, intermediate network components, etc.) works. If the remote system does not answer, it is known that at least something is 25 wrong and another (set of) tool(s) may be employed to determine the problem. Thus, the fault manager may employ the ping program, as well as any number and type of other tools needed to test components of the distributed computing environment (e.g., heartbeat detection, status indication, 30 etc.).

The topology generator 265 establishes (a subset of) the overall topology of a distributed system, comprising a vast amount of highly dynamic components, such as web applications, database instances and transactions. An example for using the 35 topology generator 265 is to display the components of a distributed system that are involved in fulfilling the

- 27 -

requests of a specific client system 105. The dependency model (or parts of it) is stored in the topology database 255, according to the needs (e.g., caching for faster retrieval) of the topology generator 265. Note that the topology generator 5 265 is the only component in the described system that directly interacts with the topology database 255. The topology generator may employ the techniques disclosed in the above-referenced and concurrently-filed U.S. patent application identified by attorney docket no. SOM920020003US1 10 entitled: "Methods And Apparatus For Topology Discovery and Representation of Distributed Applications and Services." However, other topology generation techniques may be employed.

Referring now to FIG. 2B, a block diagram is shown 15 illustrating a generalized hardware architecture of a computer system suitable for implementing the various functional components/modules of a system for providing dependency management as depicted in the figures and explained in detail herein. It is to be understood that the individual components 20 of the dependency management system, namely, components associated with the graphical user interface 285, the application layer 200, the service layer 205 and the middleware layer 210 (FIG. 2A), may be implemented on one or more computer systems having an architecture as shown in FIG. 25 2B. The other components shown in FIG. 2A, e.g., the components associated with the resource layer 215, may also be implemented on similar computer systems.

As shown, the computer system may be implemented in accordance with a processor 290, a memory 292 and I/O devices 30 294. It is to be appreciated that the term "processor" as used herein is intended to include any processing device, such as, for example, one that includes a CPU (central processing unit) and/or other processing circuitry. The term "memory" as used herein is intended to include memory associated with a 35 processor or CPU, such as, for example, RAM, ROM, a fixed memory device (e.g., hard drive), a removable memory device

- 28 -

(e.g., diskette), flash memory, etc. In addition, the term "input/output devices" or "I/O devices" as used herein is intended to include, for example, one or more input devices (e.g., keyboard) for entering data to the processing unit, and/or one or more output devices (e.g., CRT display and/or printer) for presenting results associated with the processing unit.

It is also to be understood that the term "processor" may refer to more than one processing device and that various elements associated with a processing device may be shared by other processing devices.

Accordingly, software components including instructions or code for performing the methodologies of the invention, as described herein, may be stored in one or more of the associated memory devices (e.g., ROM, fixed or removable memory) and, when ready to be utilized, loaded in part or in whole (e.g., into RAM) and executed by a CPU.

Referring now to FIG. 3, a block diagram illustrates a functional dependency model of services according to an embodiment of the present invention. More specifically, FIG. 3 depicts a functional application dependency graph between various components in an electronic commerce system such as the one depicted in FIG. 1. This functional dependency model represents both the functional components of a distributed system and their dependencies. Thus, the model defines dependencies between generic services, which are considered atomic from a business perspective. This implies that the functional model is not concerned with dependencies occurring within a business service. Such a decomposition makes sense in the scope of a specific product being used to implement the service and will be discussed in more detail with reference to FIG. 4.

Dependencies between components are depicted as arrows. An arrow always points from the dependent to the antecedent. Functional components are the (sub-) services that a service provider needs to deploy for providing an end-to-end service

- 29 -

to a customer, the latter being defined in a service level agreement. The functional model focuses on the design of an end-to-end service and abstracts from the details of an end-to-end service's technical realization, such as the products being used for the service offering, their location (local or remote system), the provider domain (i.e., whether the provider itself outsources some of its services to another service provider, transparent to the customer), etc.

As shown, an e-business application 300 service depends on a web application service 305 for hosting the business logic. In order to function properly, the web application service 305 requires two further services. The static content of the electronic commerce website is provided by a web service 310, while a back-end database service 330 stores the dynamic content (such as product descriptions, user and manufacturer data, shopping carts, user profiles and preferences, payment information, etc.) of the e-business application 300 being offered to a customer. The web service 310 depends, itself, on two services, namely, the name service 315 for mapping host names to IP addresses, and the IP service 320 for network connectivity.

Recall that dependency relationships are transitive, i.e., the dependent of a given component requires, in addition to the component itself, also the components' antecedent(s). Consequently, in addition to the IP service 320 and the database service 330, all the depicted services require the presence of an operating system (OS) 325 service. For the sake of brevity, the dependency relationships of an OS 325 on hardware components is not depicted, although they are present in a functional model.

Referring now to FIG. 4, a block diagram illustrates a structural dependency model of services according to an embodiment of the present invention. More specifically, FIG. 4 depicts a structural application dependency graph between various components in an electronic commerce system such as the one depicted in FIG. 1.

- 30 -

This structural dependency model extends the functional model (FIG. 3) in the following ways. The structural dependency model deals with the implementation of a business service and is focused on concrete products and their logical 5 (modules, components) and physical (files, shared libraries) architecture. The structural dependency model captures the detailed descriptions of software components, i.e., the system inventory, which is usually recorded in the various system repositories or in well-defined places, e.g., the 10 configuration files of a managed resource 220.

Note that although the structural model deals with the components of a single system, it may maintain references to services and applications hosted by other systems, because the configuration files located on the system may contain this 15 information. Examples of system repositories include, but are not limited to, the IBM AIX Object Data Manager (ODM), the Linux Red Hat Package Manager (RPM) or the Microsoft Windows Registry. Information relating to software components is typically captured during the installation and deployment of a 20 software package. In addition, the structural model contains the dependencies between the various system components, depicted as arrows. For the sake of clarity, the names of the business services are written without quotes, while the names of the elements of the structural model are written with 25 quotes in FIG. 4.

The system with the fully qualified domain name wslab8.watson.ibm.com 400 hosts the following components: the e-business application (a business service defined in the functional model), which is implemented as storefront servlets 30 410, the latter encapsulate the business logic of the application. The web application service is implemented by IBM WebSphere version 3.5 415, while the web service is realized by the IBM HTTP Server version 1.3.6 420. The IP service is implemented by the default IP protocol stack 430, 35 the operating system (OS) is Win(dows) NT version 4 425.

- 31 -

The system with the fully qualified domain name rslab2.watson.ibm.com 405 hosts the following components: a database service implemented by (IBM) DB2 Universal Database (UDB) version 5.2 435, and an operating system, here (IBM)

5 Advanced Interactive Executive (AIX) version 4.3.3 440.

Referring now to FIG. 5, a block diagram illustrates a service lifecycle addressed by functional, structural and operational dependency models according to an embodiment of the present invention. More specifically, FIG. 5 depicts the 10 relationships between a functional model 500 and a structural model 510, described above, and introduces a third dependency model, an operational model 520. These three models enable the invention to track the services during their whole lifecycle, i.e., from the design stage to the installation and 15 deployment stage, to the operational or runtime stage.

As explained above, the functional model 500 relates to the design of the business services and is thus captured at the design time of a business system. Once the system described by the functional model 500 becomes instantiated or 20 deployed (step 505), the structural model 510 is established. The operational model 520 is created when the various components of the structural model 510 become instantiated (step 515) and when runtime bindings between them are established. The operational model represents the 25 characteristics of the previously described models at runtime. Several scenarios that illustrate the aforementioned concepts will now be described.

The web application service 305 is implemented by IBM WebSphere 415; one or more instances of the latter are 30 referred to as websphere-daemon 545. Here, the web (or WWW) service 310 is implemented by two products, namely, Apache 1.3.4 525 and Lotus Domino 530. The running instances of these products can be identified as http daemons "httpd" 550. The database service 330 is implemented by two products, 35 namely, Oracle v7 535 and DB2 UDB 435; however, no instances of Oracle v7 535 are active because no server processes are

- 32 -

visible in the operational model 520. In contrast, four instances of DB2 UDB 435 are running, as can be seen from the presence of four DB2 daemons "db2d" 555 in the operational model 520. The name service 315 is implemented by BIND version 5.6 540; the running instance of BIND can be observed as "named" 560 in the operational model 520.

Note that the dependencies are propagated from the functional to the structural and the operational model. This is necessary because it is not possible to determine from a running application instance which other application instance(s) it requires to function properly.

Due to the short-livedness of some application instances, the operational model 520 is highly dynamic and potentially very large. In contrast to the functional and structural dependency models, the operational model 520 is not stored in a repository or a database, but computed on demand and to the extent needed.

Referring now to FIG. 6, a block diagram illustrates relationships between functional, structural and operational dependency models according to an embodiment of the present invention. More specifically, FIG. 6 depicts the details of the data template used for the three dependency models and the means for tying these models together by an example. The example details the template and its associated values for describing the name service during its lifecycle.

The functional template 605, used for the functional model 500 contains the "hostName" (the unique name of the computer system that hosts the service), the "serviceName" (name of the service) and the "componentType" (the role in which this service acts, i.e., client or server). With this information, a service can be uniquely identified within a distributed environment. However, further fields containing descriptive data (such as a description of the purpose of the service, the customer that subscribes to this service, etc.) may be added without departing from the spirit of the present

- 33 -

invention. Finally, the "Antecedent" field contains the service(s) this service requires for proper functioning.

The structural template 610, used for the structural model 510, contains all the fields of the functional template 605, which allows to link the functional template 605 with the structural template 610 in order to navigate from the functional model 500 to the structural model 510, and vice-versa. In addition, the structural template 610 contains the "componentName" (the name of the product component), an "identifier" (a globally unique name for identifying the component), the "version," "release" and "modification" (e.g., maintenance or patch/fix level) numbers, the "installState" (indicating whether the component has been successfully and completely installed) and "processName" (the name of the process(es) that identify this product component at runtime). Furthermore, the "Antecedent" field lists the component(s) this component requires to be operable.

The operational template 615, used for the operational model 520, contains the fields "hostName" (the unique name of the computer system that hosts the service) and "processName" (the name of the process(es) that identify the product component at runtime). These two fields link the structural template 610 with the operational template 615 in order to navigate from the structural model 510 to the operational model 520, and vice-versa. In addition, the operational template 615 contains the fields "operState" (the operational state of the process, i.e., running, interrupted, zombie, etc.), "portNumber" (the number of the TCP/UDP port by which an application can connect to the process), and, "instanceID" (to distinguish the various application instances within the scope of a computer system).

The three dependency models are stored and computed at different places to achieve a maximum degree of efficiency. The functional model 500 is gathered and stored at the management system 620, i.e., the central point of control by which the administrator 100 interacts with the distributed

- 34 -

environment. Some of the reasons for this choice are as follows. As seen in the descriptions of FIG. 3 and FIG. 5, the functional model 500 is fairly compact because the amount of possible business services is limited. In addition, the 5 functional model is not subject to overly frequent changes. The functional model is defined at the time when a business service is provided to a customer and remains unchanged until the service offering period ends. Since the administrator 100 is responsible for setting up and updating the functional 10 model 500, it is a natural choice to keep it close to the management system 620.

As mentioned in the description of FIG. 4 and FIG. 5, the structural model 510, in contrast, captures the detailed descriptions of software components, i.e., the system 15 inventory, which is usually recorded in the various system repositories or in well-defined places, e.g., the configuration files of managed resource 220. Consequently, it is both large sized (the content of a system repository tends to be between several hundred kilobytes and up to a few 20 Megabytes) and also subject to frequent changes. Therefore, keeping the structural model 510 of a system at the managed resource 220, itself, eliminates both the communication overhead for updating the model and the need for vast amounts 25 of storage that would occur if the structural models 510 of all managed resources (220) is stored at a centralized place.

The operational model 520 has been described in FIG. 5 as very dynamic and also extremely large because it covers potentially multiple instances of every application present on the computer systems of the distributed environment and the 30 dependency relationships between them. Given the fact that current data centers of Internet/Application/Storage Service Providers and outsourcers consist of several thousands of computer systems with each hosting close to 100 applications and system services, an operational model comprising all the 35 currently instantiated applications and their dependencies may be impractical. Thus, a practical approach is to compute the

- 35 -

relevant parts of the operational model on demand (step 625). This is the purpose of the dependency service 245.

Referring now to FIG. 7, a block diagram illustrates components involved in analyzing and computing the root cause 5 of a service outage based on dynamic information technology (IT) service dependencies according to an embodiment of the present invention. More specifically, FIG. 7 depicts the data flows between the various components used for analyzing and computing such a root cause. It is assumed that the managed 10 resources 220 are able to provide XML (Extensible Markup Language) descriptions of their system inventory, configuration files and their various dependencies. However, it should be noted that any data description format may be used in accordance with the invention. The details on how 15 this information can be acquired are as follows.

One straightforward way is to provide appropriate instrumentation within the system and its applications and services. This information is described in flat XML files 740 and made available to the other components of the system 20 through a web server 725.

Alternatively, the dependency service 245 makes use of information stored in system repositories 745 for generating appropriate service dependency information. This information is made available to the other components of the system 25 through a web server 730.

Third, the managed resources 220 expose their information by means of an instrumentation agent, called CIM (Common Information Model, which is a standardized management framework) provider 750, which interacts with a CIM Object 30 Manager (CIMOM) 735, as proposed by the Distributed Management Task Force (DMTF). The CIMOM then exposes the necessary information to the interested components.

In the center of FIG. 7, various management services, being part of the service layer 205, are depicted. These are: 35 a name service 700, a trader service 710, an event service 715 and the dependency service 245. The dependency service 245,

- 36 -

triggered by queries of the administrator 100 through either the root cause analyzer 280, its management system or any management application located in the application layer 200 using a communication protocol (e.g., Java Remote Method Invocation (RMI)), processes them and sends the results back to the root cause analyzer 280 which, in turn, forwards the results for further processing to the administrator 100. The main tasks of the dependency service 245 are as follows:

(a) Interacting with the management system or any management application located in the application layer 200. The management system issues queries to the application programming interface (API) of the dependency service (245).

(b) Expose a 'drill-down' method that, upon receiving the identifier of a service, returns:

(i) either descriptions of its direct antecedents, i.e., the first level below the node representing the service, or

(ii) the whole subgraph below the node representing the service,

(iii) an arbitrary subset of the dependency graph (levels m to n below a given node).

(c) Provide a 'drill-up' method with the same facilities, targeting the dependents of the service.

(d) Additional methods for gathering and filtering information for classes and properties of managed objects are present.

(e) Obtaining the dependency information from the managed resources 220 by issuing queries over http (HyperText Transport Protocol) and applying filtering rules (as specified by the administrator 100) to it.

(f) Combining the information into a data structure that is sent back to the management system as XML document.

As mentioned above, due to its fully distributed nature, the invention aims at keeping the load on every involved system as low as possible. The invention decouples the management system from the managed resources 220 and

- 37 -

encapsulates the time consuming filter and join operations in the dependency service 245, which can be replicated on various systems. It is therefore possible to achieve a maximum level of parallelism for query operations, since the selection of an 5 instance of the dependency service 245 can be done flexibly by the management system.

Another important advantage is that the (very large and highly dynamic) operational model 520 is not stored at a specific place but computed on demand in a stepwise manner. 10 The different parts of the structural model 510 are stored at the managed resources 220. The management system therefore always receives the most recent information but is still free to store it according to elaborate caching policies.

Referring now to FIG. 8, a block diagram illustrates 15 components of a root cause analyzer according to an embodiment of the present invention. As shown, a root cause correlator 870, which functions as a flow coordinator for the entire root cause analysis process, receives service problem reports 880 from the administrator 100, providing the name of the service 20 and the host name where the problem has been observed or reported. The root cause correlator 870 interacts with the dependency service 245 in order to obtain a list of the basic services on which the problematic service depends. Examples 25 of basic services would be: domain name service, IP connectivity service, etc.

The task of the dependency service 245 is to find the antecedents of the service in question, even if the e-commerce environment spans different managed domains 800. In order to deal with multiple domains, various (cascaded) instances of 30 the dependency service 245 may work jointly together. In FIG. 8, the e-commerce environment is shown by the dotted rectangle. Typically, such environments contain one or more managed domains, eventually each with its own dependency database 810 and dependency service 245. The dependency 35 service 245 returns the names and identifiers of the antecedents back to the root cause correlator 870, which then

- 38 -

initiates further problem determination procedures to verify whether one or more antecedents experience difficulties. This can be accomplished via a number of approaches, some of which are illustrated below.

5 A first approach assumes the presence of an event monitoring and distribution function, provided by an event monitor 820. Examples of such event monitors include, but are not limited to HP OpenView Event Services and Tivoli Enterprise Console. The event monitor 820 receives
10 events/alerts related to resources within the e-commerce environment and forwards the ones, which are relevant to an event database 830. Some events are informational in nature or are related to errors that are fixed automatically by the domain management services. These are typically filtered out
15 and not forwarded to the event database 830. In practice, the event monitor 820 may comprise of a hierarchical arrangement of multiple event monitors, one for each managed domain.

20 The main function of an event-service correlator 840 is to provide, for a given service or resource, a list of outstanding alerts associated with that resource or service. It does so by interacting with event monitors 820 and stores the events in the event database 830 for further retrieval. When a service returned by the dependency service 245 is one
25 of the suspected root causes, the root cause correlator 870 uses the event-service correlator 840 to estimate the likelihood of this being the case by getting a list of events that may have been reported about the suspected service or resource. If there is an undesirable event, for example buffer overflow, the likelihood of this resource being a root
30 cause is high.

35 The root cause correlator 870 also interacts with the dependency services 245 for each of the managed domains in order get a set of possible root causes that lie within that domain. The dependency service 245 interacts with domain dependency data 810, which is constructed during deployment

- 39 -

and periodically updated during the operation phase of the e-commerce setup.

A second approach is to obtain the current status of a service or resource from a status monitor 860, which directly 5 interacts with the services and verifies whether they are functioning properly or not (e.g., failed, not failed, degraded, etc.). As mentioned above, if the system incorporates a fault manager (as does the system of FIG. 2A), the status monitor may serve as the interface between the root 10 cause analyzer and the fault manager so as to, among other advantages, make the root cause analyzer independent of any particular test procedure. The status monitor 860 may therefore serve as a single point of contact for the root cause correlator 870 to determine the health of the service in 15 question.

In general, it can not be assumed that an external fault manager is always available. Since, however, the functionality of determining the status of a service is crucial for the root cause analyzer, this functionality must 20 be present and is provided by the status monitor 860, which may extend the functionality offered by a fault manager. Thus, external fault managers may provide either:

(a) all the needed functionality (thus, the status manager will "wrap" the functions performed by the fault 25 manager and adapt the interface to what the root cause correlator 870 expects;

(b) only a part of the functionality (e.g., testing of network connectivity only - but no application and middleware status verification), thus the needed functionality must be 30 provided by the status monitor; or

(c) no functionality at all; it is either not present or does not offer its functionality to other components through a programmatic interface. So while a human user may interact with the fault manager through a GUI (graphical user 35 interface), no program (such as the correlator 870) can use it.

- 40 -

Also note that a fault manager is meant to refer to a loose collection of system administration tools having completely different interfaces. The status monitor thus serves to integrate these various tools under a uniform 5 interface. That is, the status monitor preferably operates to account for any mismatch between a root cause correlator which expects everything to be testable for status and the fault manager which may provide between 0 percent and 100 percent of this required functionality.

10 The root cause correlator 870 thus has two ways to determine whether a service is functioning properly or experiences a problem, i.e., the correlator can query the event-service correlator 840 for problem event reports relating to a service or directly look up the status of a 15 service from the status monitor 860. The root cause correlator 870 is free to choose by which of these at least two means it determines the health of a service. Further, the root cause correlator 870 may choose to obtain information both from the event-service correlator 840 and from the status 20 monitor 860.

After the root causes have been narrowed down to the smallest possible set of services and resources, the root cause correlator 870 returns this information 890 to the administrator 100.

25 Referring now to FIG. 9, a flow diagram illustrates steps for invoking a dependency service and performing root cause analysis on an operational model according to an embodiment of the present invention. More specifically, FIG. 9 depicts a methodology of invoking a dependency service (e.g., dependency 30 service 245), collecting its results, and applying root cause analysis to them. This methodology is initiated either by an administrator 100 or a management application being part of the application layer 200, as described in FIG. 2A.

35 The methodology begins at block 900 and proceeds as follows. First, a business service is selected (step 905), typically from the functional model, since an administrator is

- 41 -

interested in the business services provided by the distributed system. Upon selection of a business service, the structural model is queried to provide a selection of hosts that are involved in the provisioning of the business service.

5 This can be either done by locating the structural models present on every host of the distributed system, or (for efficiency purposes) by querying a (periodically updated) service/host lookup table stored at the management system, which contains a mapping between the services and hosts
10 present in the distributed system. The administrator then selects a host at his discretion (step 910).

In addition, the administrator composes a query (step 915). Examples of query parameters include, but are not limited to, the direction of the traversal (towards the
15 service dependents, or towards its antecedents), the depth of the traversal (e.g., only the immediate antecedents/dependents; the overall possible antecedents/dependents, i.e., the complete transitive closure of the operational model; only between the m -th and n -th
20 layers of the operational model), filtering criteria related either to the presence of attributes or to their values.

The fact that the order of the steps for selecting service (step 905), host (step 910) and the options for composing the query is prescribed here, emphasizes the
25 "service-centric" approach (versus the "host-centric" approach of existing art) of the present invention. However, one having skill in the relevant art will recognize that modifications in the order of steps (steps 905, 910 and 915) may be made without departing from the spirit and scope of the
30 present invention.

Examples of such modifications are: to offer the user (e.g., by means of a graphical user interface) the choice of performing the three steps of the selection process in an arbitrary order; allowing at first the selection of a host and
35 then looking up the services present on that host by means of

- 42 -

querying the structural model, thereby restricting the possible service candidates for selection.

After the service and host selection and the composition of the query, the dependency service is invoked (step 920) with these parameters. Note that the mode of invocation can be either synchronous (i.e., blocking the invoker until the results are returned by the dependency service) or asynchronous (thus allowing the invoker to perform additional tasks during the computation).

The dependency service computes the appropriate part of the operational model and, depending on the mode of invocation, either sends the results back to the invoker or notifies the invoker that the results are available. The invoker then collects the results and applies root cause analysis and appropriate problem determination procedures to them (step 925). The methodology ends at block 930.

Referring now to FIG. 10, a flow diagram illustrates tasks of an administrator for creating and updating a functional dependency model according to an embodiment of the present invention. This is necessary either if new (business) services are deployed and offered, or changes are applied to an existing model, or existing (business) services are withdrawn from an offering.

The methodology begins at block 1000 and proceeds as follows. An administrator or a management application evaluates whether a new business service should be added or an existing service is to be deleted (step 1005). If this is not necessary, the methodology proceeds directly to block 1025. Otherwise, in step 1010, the service and its description are entered in (or removed from) the template 605 of the functional model, which has been described in FIG. 6.

Then, in step 1015, the service dependencies, i.e., its relationships regarding its antecedents, need to be added to (or removed from) the template 605 of the functional model.

In case of a deletion, note that the dependencies from the service dependents need to be adjusted to point to the

- 43 -

antecedents of the service that is to be removed. This may involve checking for eventual duplicate descriptions within the dependencies of the antecedents. Finally, the updated functional model is stored in the repository of the management system (step 1020). The methodology ends at block 1025.

Referring now to FIG. 11, a flow diagram illustrates steps of updating a structural dependency model by installing or removing hardware/software components on a computer system according to an embodiment of the present invention. This is necessary either if new components are deployed and installed on a host, or existing components are removed from the host.

The methodology begins at block 1100 and proceeds as follows. If new hardware components are installed/removed, the verification and adjustment of their dependencies is usually performed by the operation system and is thus not described further here. Instead, the following description focuses on the task of adding/removing software components. An administrator or a management application performing software distribution and installation evaluates whether a new software component should be added or an existing software component is to be deleted (step 1105). If this is not necessary, the methodology proceeds directly to block 1125. Otherwise, in step 1110, the software components' descriptions are entered in (or removed from) the template 610 of the structural model, which has been described in FIG. 6. Then, in step 1115, the software components' dependencies, i.e., its relationships regarding its antecedents, needs to be added to (or removed from) the template 610 of the structural model.

In case of a deletion, note that the dependencies from the software components' dependents need to be adjusted to point to the antecedents of the software component that is to be removed. This may involve checking for eventual duplicate descriptions within the dependencies of the antecedents. Finally, the updated structural model is stored in the resource dependency repository of the host (step 1120). The methodology ends at block 1125.

- 44 -

Referring now to FIG. 12, a flow diagram illustrates performance of root cause analysis on an operational model according to an embodiment of the present invention. The methodology begins at block 1200 and proceeds as follows. The 5 system performing the root cause analysis on the operational dependency model is continuously listening for requests at a specific port of the host on which the system is executed, which is illustrated by the loop that connects step 1205 with itself. This is the standard behavior for server processes 10 ("daemons") that implement services, which can be invoked by applications at any time.

On reception of a request, the system extracts the input parameters from the request (step 1210). As mentioned in the description of FIG. 9, examples of input parameters include, 15 but are not limited to, the names of the service and the host in question, the direction of the traversal, the depth of the traversal, filtering criteria related either to the presence of attributes or to their values. These input parameters are then used to invoke the computation of the operational model, 20 which is invoked in step 1215.

In addition, the results of the computation, i.e., the operational model, are gathered. The root cause analysis is then performed on the operational model (step 1220). The results of the root cause analysis are returned to the 25 invoking application, in accordance with the mode of invocation specified at the time of the invocation (step 1225). After this step, any allocated resources of the host on which the system is running are released (step 1230). Examples of host resources include, but are not limited to, 30 memory, disk space or CPU registers. Finally, the system returns back to its initial stage and listens for subsequent incoming requests (return to step 1205).

Referring now to FIG. 13, a flow diagram illustrates performance of root cause analysis on antecedents of a service 35 according to an embodiment of the present invention. The methodology begins at block 1300 and proceeds as follows.

- 45 -

First, the names of the targeted service and the host are obtained (step 1305). These parameters are provided by the invoking management application, which obtains these parameters either directly from the administrator or from an 5 event message that arrives at the management console. In addition, the search criteria according to the status of the antecedents are specified. They indicate whether the system should return the services that have encountered a problem ("defective") or the services that perform well. Usually, a 10 management application is interested in the former, i.e., the defective services. Furthermore, the search path is specified. A management application can be either interested in the direct antecedents of a service (search path length = 1), the whole set of services the service directly (or 15 indirectly) depends on (search path = recursive), or the lowest level of the antecedents, i.e., the base services that a service depends on (search path = max).

Then, the computation of the operational model is performed by the dependency service, according to the 20 parameters "Service Name," "Host Name," "Antecedent status," "search path" (step 1310). Next, in step 1315, the results, i.e., the list of antecedent service elements ("Candidate List"), are obtained from the dependency service.

The following steps are performed until the list of 25 candidate service elements is empty (1320):

The first service element of the candidate list is selected (step 1325) and checked to determine whether it is functioning properly (step 1330). The steps of this status check procedure are described in detail in FIG. 14. If the 30 service element in question is functioning properly (status = "OK"), it is added to the list of working service elements, i.e., an "OK" list (step 1335). If, however, it turns out that the service element is experiencing a problem, it is added to the "defective" list (step 1340). Finally, the 35 service element in question is removed from the candidate

- 46 -

service elements list (step 1345) and the methodology proceeds back to block step 1320.

If the list of candidate service elements is empty, the methodology then proceeds directly to step 1350, where either the list of working service elements or the list of defective service elements is returned to the invoker. The results content depends on whether the invoker has been asking in step 1305 for working or defective services. The methodology ends at block 1355.

Referring now to FIG. 14, a flow diagram illustrates steps for determining a status of a service according to an embodiment of the present invention. More specifically, FIG. 14 illustrates interactions of the root cause correlator 870 with the event-service correlator 840 and the status monitor 860, depicted in FIG. 8. The methodology begins at block 1400 and proceeds as follows.

First, the name of the targeted service element is obtained (step 1405). This problem determination problem is described from the perspective of checking the status of a single service element, as may be done by the root cause correlator 870. It is to be understood that this procedure is repeated for every service in question. It is to be further understood that this procedure is carried out in step 1330 of FIG. 13.

Also, with respect to step 1405, the names of the targeted service elements are provided by the invoking management application, which obtains these parameters either directly from the administrator or from an event message that arrives at the management console.

Then, the root cause correlator 870 queries the event-service correlator 840 for the presence of events relating to the service in question (step 1410). These events indicate whether any problems with this service have been observed. If this is the case, the service element is flagged as "defective" (step 1415), and the result is returned to the invoker (step 1430).

- 47 -

If step 1410 yields that no fault events have occurred for this service, it is necessary to perform further work to determine the status of a service. This is done by querying the status monitor 860 for the status of the service (step 5 1420). The task of the status monitor 860 comprises providing a uniform interface to status information for every service in the distributed system, regardless of how the status of the service is determined.

As mentioned above, four illustrative ways of determining 10 the status of a service may comprise: (i) process inspection; (ii) exerciser; (iii) heartbeat; and (iv) status indicators. If the query (step 1420) to the status monitor 860 yields that the service in question is functioning properly (status = "OK"), the service element representing the service is flagged 15 as "OK" (step 1425). In any case, the result is returned to the invoker (step 1430). The methodology ends at block 1435.

Referring now to FIG. 15, examples are depicted of root cause analyzer application programming interfaces (APIs) according to an embodiment of the present invention. The 20 table includes base APIs that can generate, send and request receipt of appropriate operational models for a given service and host name. Those skilled in the art will appreciate that the APIs can use one or more parameters (not shown) to identify characteristics (specified in the functional 25 description) used by the APIs.

Specifically, a "getDirectAntecedents(parameters)" API retrieves direct antecedents, regardless of their current status, of a service located on a specific host. The "getAntecedentsRecursive(parameters)" API performs a recursive 30 "drill-down," i.e., it retrieves all the antecedents of a given service (regardless of their current status), which is located on a specific host. It therefore returns all the services residing "below" a given service in the dependency hierarchy. The "getLowestAntecedents(parameters)" API 35 retrieves the lowest antecedents of a service located on a specific host. This method yields the services that do not

- 48 -

have any antecedents themselves (i.e., they reside in the lowest layer of the dependency hierarchy of a given service).

The "getWorkingDirectAntecedents(parameters)" API retrieves direct antecedents with status "up" of a service located on a specific host.

5 The "getWorkingAntecedentsRecursive(parameters)" API performs a recursive "drill-down," i.e., it retrieves all the antecedents with status "up" of a given service, which is located on a specific host. It therefore returns all the working services 10 residing "below" a given service in the dependency hierarchy. The "getWorkingLowestAntecedents(parameters)" API retrieves the lowest antecedents with status "up" of a service located on a specific host. This method yields the services that do not have any antecedents themselves (i.e., they reside in the 15 lowest layer of the dependency hierarchy of a given service).

The "getDefectiveDirectAntecedents(parameters)" API retrieves direct antecedents with status "down" of a service located on a specific host.

20 The "getDefectiveAntecedentsRecursive(parameters)" API performs a recursive "drill-down," i.e., it retrieves all the antecedents with status "down" of a given service, which is located on a specific host. It therefore returns all the failed services residing "below" a given service in the dependency hierarchy. The "getDefectiveLowestAntecedents(parameters)" API retrieves 25 the lowest antecedents with status "down" of a service located on a specific host. This method yields the services that do not have any antecedents themselves (i.e., they reside in the lowest layer of the dependency hierarchy of a given service). This API is particularly useful if a user or management 30 application wants to determine quickly whether any "fundamental" service (e.g., network connectivity service) has failed without having to verify the status of "intermediate" services (such as name services).

35 Although illustrative embodiments of the present invention have been described herein with reference to the accompanying drawings, it is to be understood that the

- 49 -

invention is not limited to those precise embodiments, and that various other changes and modifications may be made by one skilled in the art without departing from the scope or spirit of the invention.

Claims

What is claimed is:

1. A computer-based method for determining a root cause of a condition of at least one subject component in a computing environment, the method comprising the steps of:

5 identifying one or more components in the computing environment upon which the at least one subject component depends, wherein identification comprises traversing at least a portion of a model representative of an existence of one or more relationships associated with at least a portion of components of the computing environment and which is capable of accounting for a full lifecycle associated with at least one component of the computing environment; and

10 performing one or more procedures in accordance with the one or more identified components to determine a condition status associated with each of the one or more identified components.

2. The method of claim 1, wherein the computing environment comprises a distributed computing environment.

20 3. The method of claim 1, wherein the computing environment comprises an autonomic computing environment.

4. The method of claim 1, wherein the identifying step further comprises identifying the one or more components upon which the at least one subject component directly depends.

25 5. The method of claim 1, wherein the identifying step further comprises identifying a complete set of components upon which the at least one subject component depends.

30 6. The method of claim 1, wherein the performing step further comprises performing the one or more procedures in a stepwise manner.

7. The method of claim 1, wherein the performing step further comprises performing the one or more procedures in a combined manner.

8. The method of claim 1, wherein the model is in a form 5 comprising a functional categorization, a structural categorization and an operational categorization.

9. The method of claim 8, wherein the identifying step further comprises computing the operational categorization of the model based on at least one input criterion and traversing 10 the operational categorization to identify the one or more components in the computing environment upon which the at least one subject component depends.

10. The method of claim 9, wherein the at least one input criterion comprises at least one of a name of the at 15 least one subject component, a host associated with the at least one subject component, a direction of the traversal, a depth of the traversal, a path of the traversal and at least one filtering criterion.

20 11. The method of claim 1, wherein the one or more procedures of the performing step comprise at least one of a component inspection procedure, a component exercising procedure, a component heartbeat detection procedure and a component status indication procedure.

25 12. The method of claim 1, wherein the condition of the at least one subject component is one of failed, not failed and degraded.

13. The method of claim 1, wherein a component is one of a service, an application, middleware, hardware, a device

- 52 -

driver, an operating system and a system associated with the computing environment.

14. The method of claim 1, wherein the root cause determination method is performed for one or more subject components substantially concurrently.

15. The method of claim 1, wherein at least a portion of results of the root cause determination method is one of persistently stored and not persistently stored.

16. The method of claim 1, further comprising the step of maintaining a history of results associated with the root cause determination method.

17. The method of claim 14, wherein at least a portion of the history is used to identify a most likely root cause.

18. Apparatus for determining a root cause of a condition of at least one subject component in a computing environment, the apparatus comprising:

at least one processor operative to: (i) identify one or more components in the computing environment upon which the at least one subject component depends, wherein identification comprises traversing at least a portion of a model representative of an existence of one or more relationships associated with at least a portion of components of the computing environment and which is capable of accounting for a full lifecycle associated with at least one component of the computing environment; and (ii) cause performance of one or more procedures in accordance with the one or more identified components to determine a condition status associated with each of the one or more identified components; and

30 memory, coupled to the at least one processor, operative to store at least a portion of results associated with the identifying and performing operations.

19. The apparatus of claim 18, wherein the computing environment comprises a distributed computing environment.

20. The apparatus of claim 18, wherein the computing environment comprises an autonomic computing environment.

5 21. The apparatus of claim 18, wherein the identifying operation further comprises identifying the one or more components upon which the at least one subject component directly depends.

10 22. The apparatus of claim 18, wherein the identifying operation further comprises identifying a complete set of components upon which the at least one subject component depends.

15 23. The apparatus of claim 18, wherein the performing operation further comprises causing the performance of the one or more procedures in a stepwise manner.

24. The apparatus of claim 18, wherein the performing operation further comprises causing the performance of the one or more procedures in a combined manner.

20 25. The apparatus of claim 18, wherein the model is in a form comprising a functional categorization, a structural categorization and an operational categorization.

25 26. The apparatus of claim 25, wherein the identifying operation further comprises computing the operational categorization of the model based on at least one input criterion and traversing the operational categorization to identify the one or more components in the computing environment upon which the at least one subject component depends.

27. The apparatus of claim 26, wherein the at least one input criterion comprises at least one of a name of the at least one subject component, a host associated with the at least one subject component, a direction of the traversal, a 5 depth of the traversal, a path of the traversal and at least one filtering criterion.

28. The apparatus of claim 18, wherein the one or more procedures of the performing operation comprise at least one 10 of a component inspection procedure, a component exercising procedure, a component heartbeat detection procedure and a component status indication procedure.

29. The apparatus of claim 18, wherein the condition of the at least one subject component is one of failed, not 15 failed and degraded.

30. The apparatus of claim 18, wherein a component is one of a service, an application, middleware, hardware, a device driver, an operating system and a system associated with the computing environment.

20 31. The apparatus of claim 18, wherein root cause determination is performed by the at least one processor for one or more subject components substantially concurrently.

25 32. The apparatus of claim 18, wherein at least a portion of results of the root cause determination performed by the at least one processor is one of persistently stored and not persistently stored.

33. The apparatus of claim 18, wherein the at least one processor is further operative to maintain a history of results associated with the root cause determination method.

- 55 -

34. The apparatus of claim 33, wherein at least a portion of the history is used to identify a most likely root cause.

35. An article of manufacture for determining a root cause of a condition of at least one subject component in a computing environment, comprising a machine readable medium containing one or more programs which when executed implement the steps of:

identifying one or more components in the computing environment upon which the at least one subject component depends, wherein identification comprises traversing at least a portion of a model representative of an existence of one or more relationships associated with at least a portion of components of the computing environment and which is capable of accounting for a full lifecycle associated with at least one component of the computing environment; and

performing one or more procedures in accordance with the one or more identified components to determine a condition status associated with each of the one or more identified components.

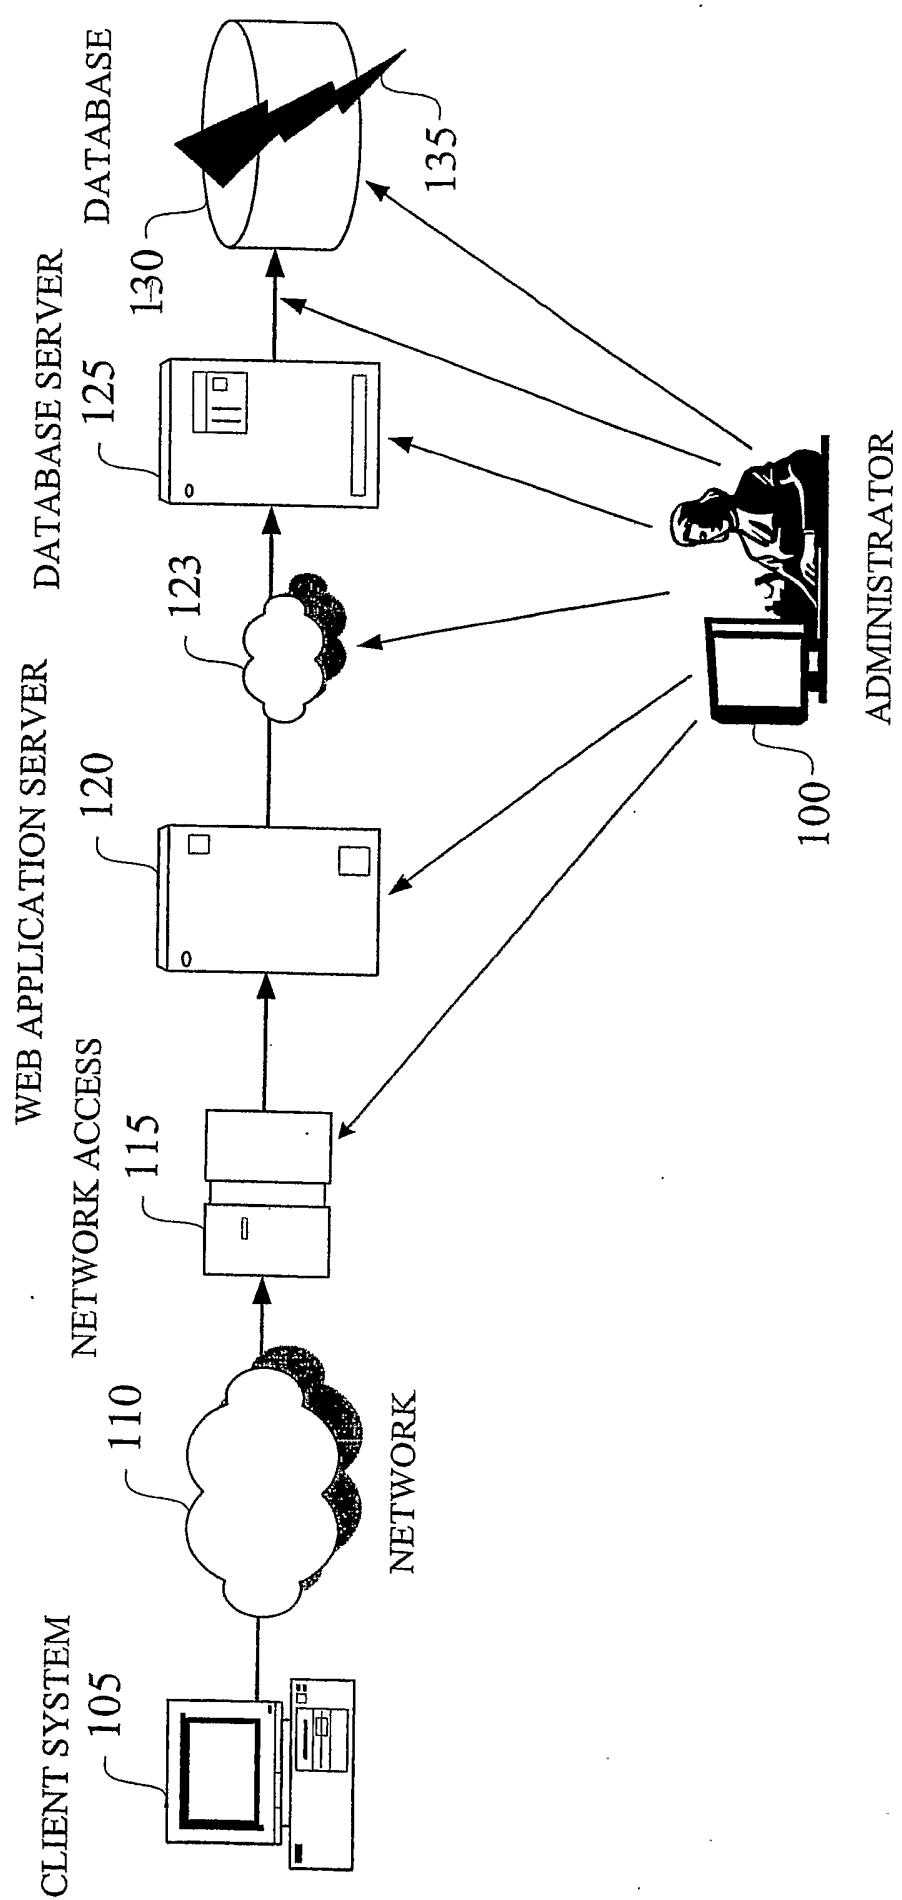
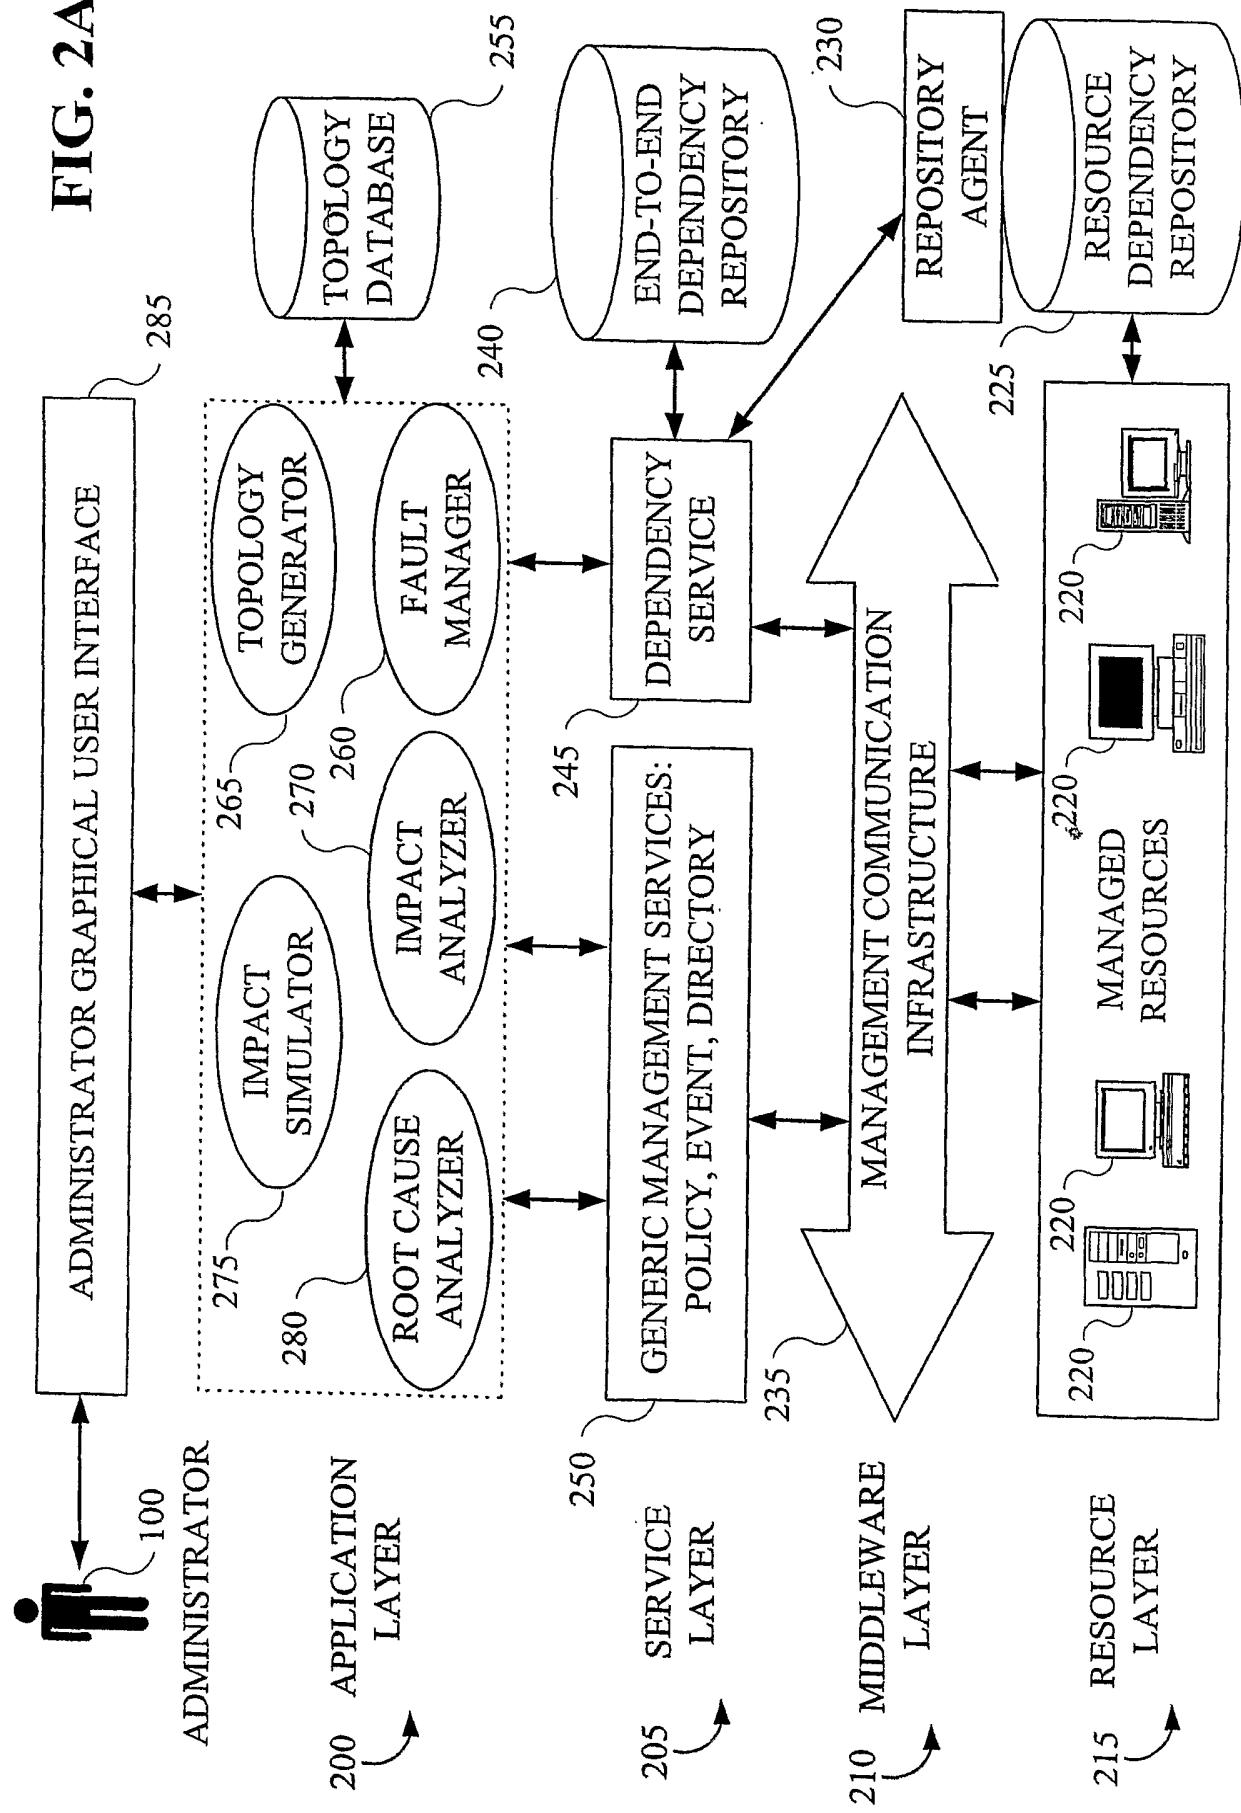
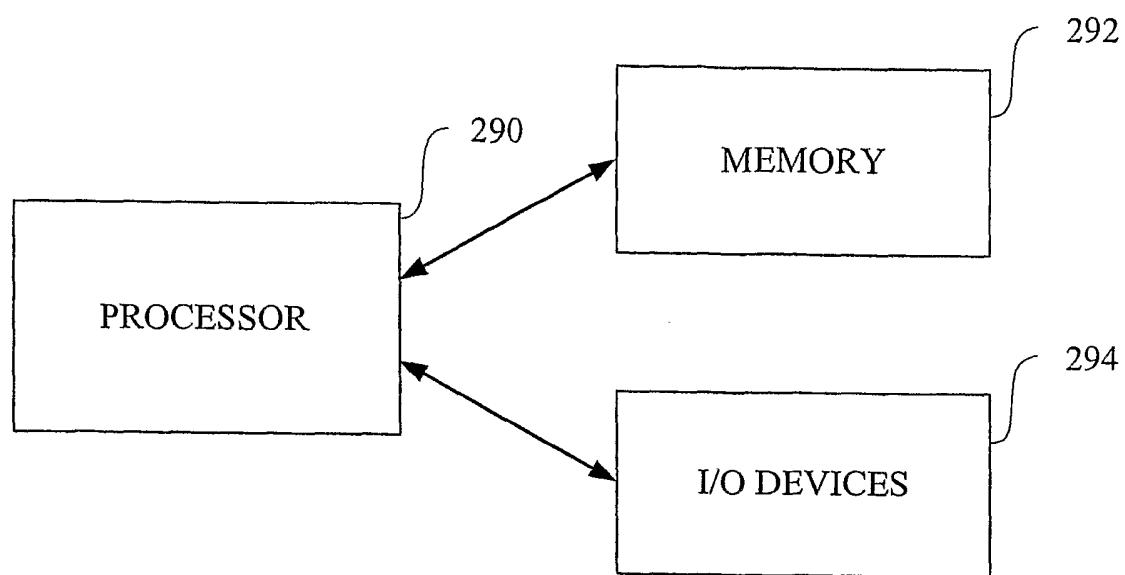
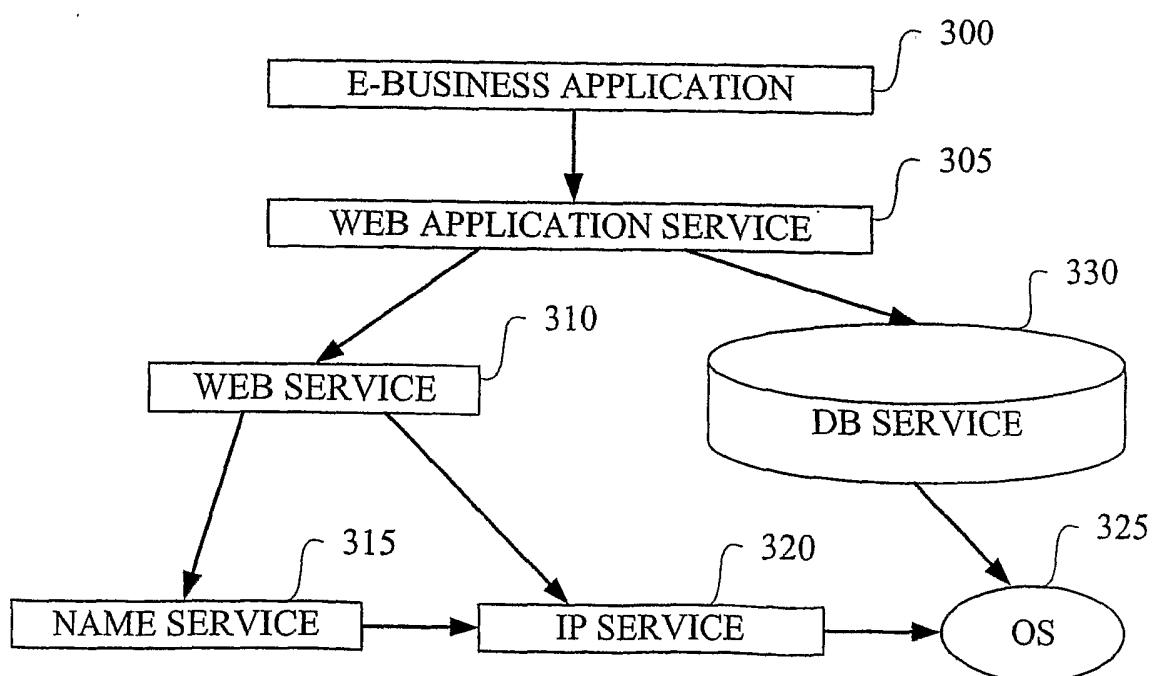
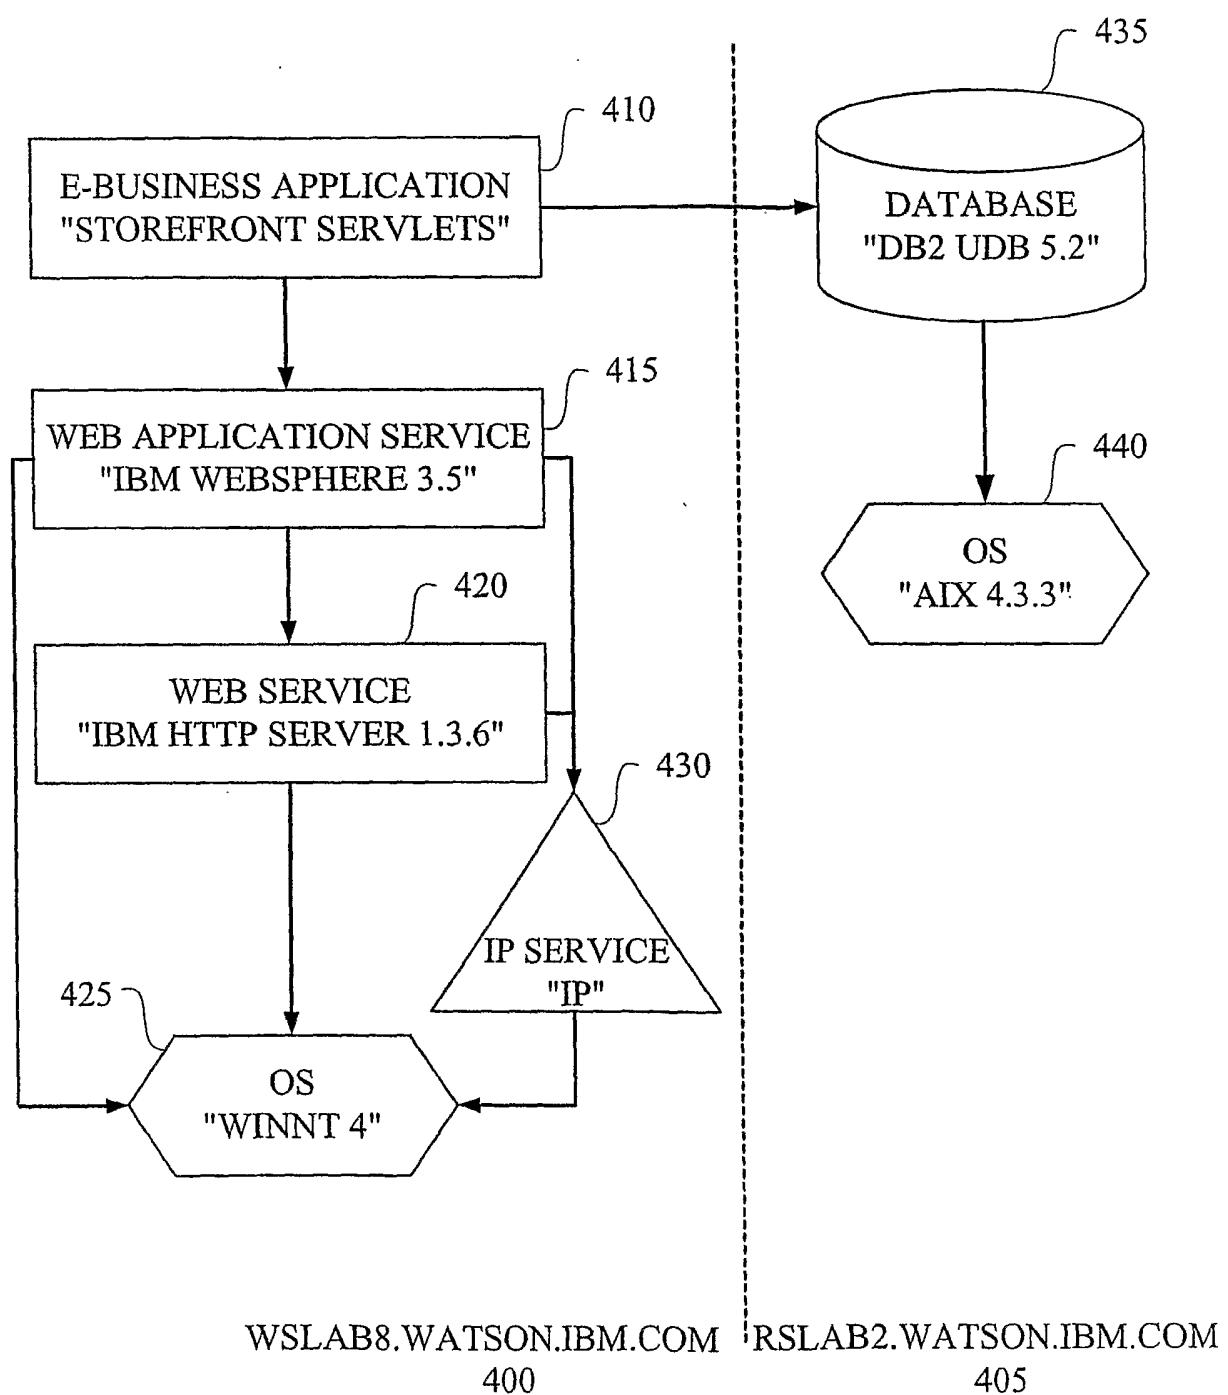


FIG. 1

FIG. 2A



**FIG. 2B****FIG. 3**

**FIG. 4**

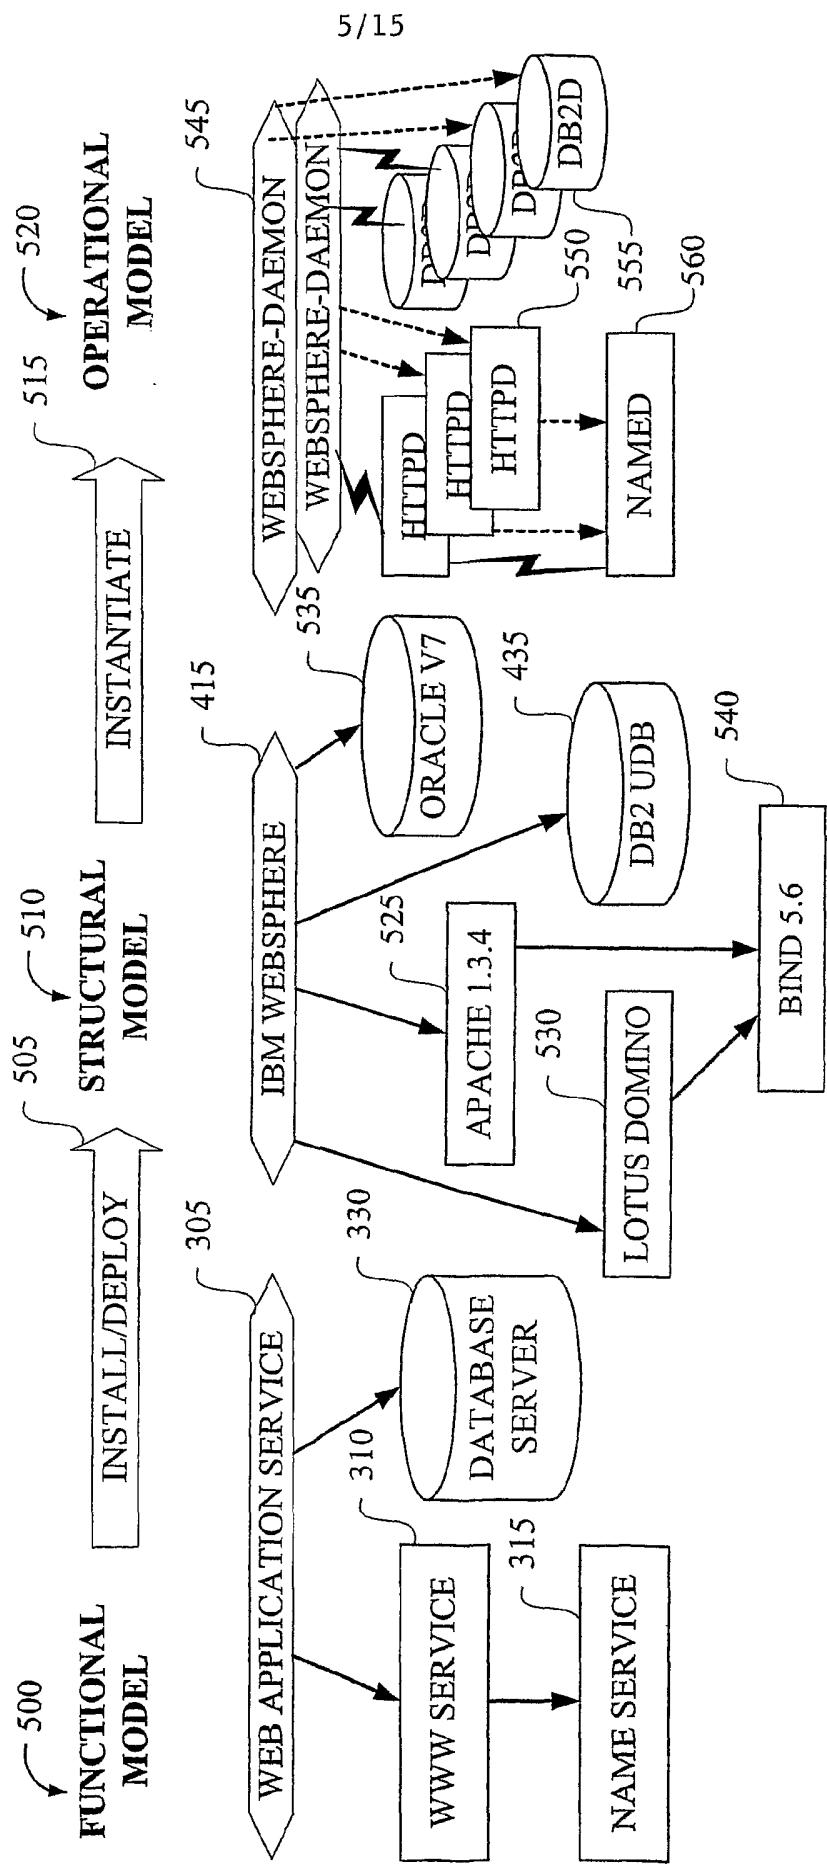


FIG. 5

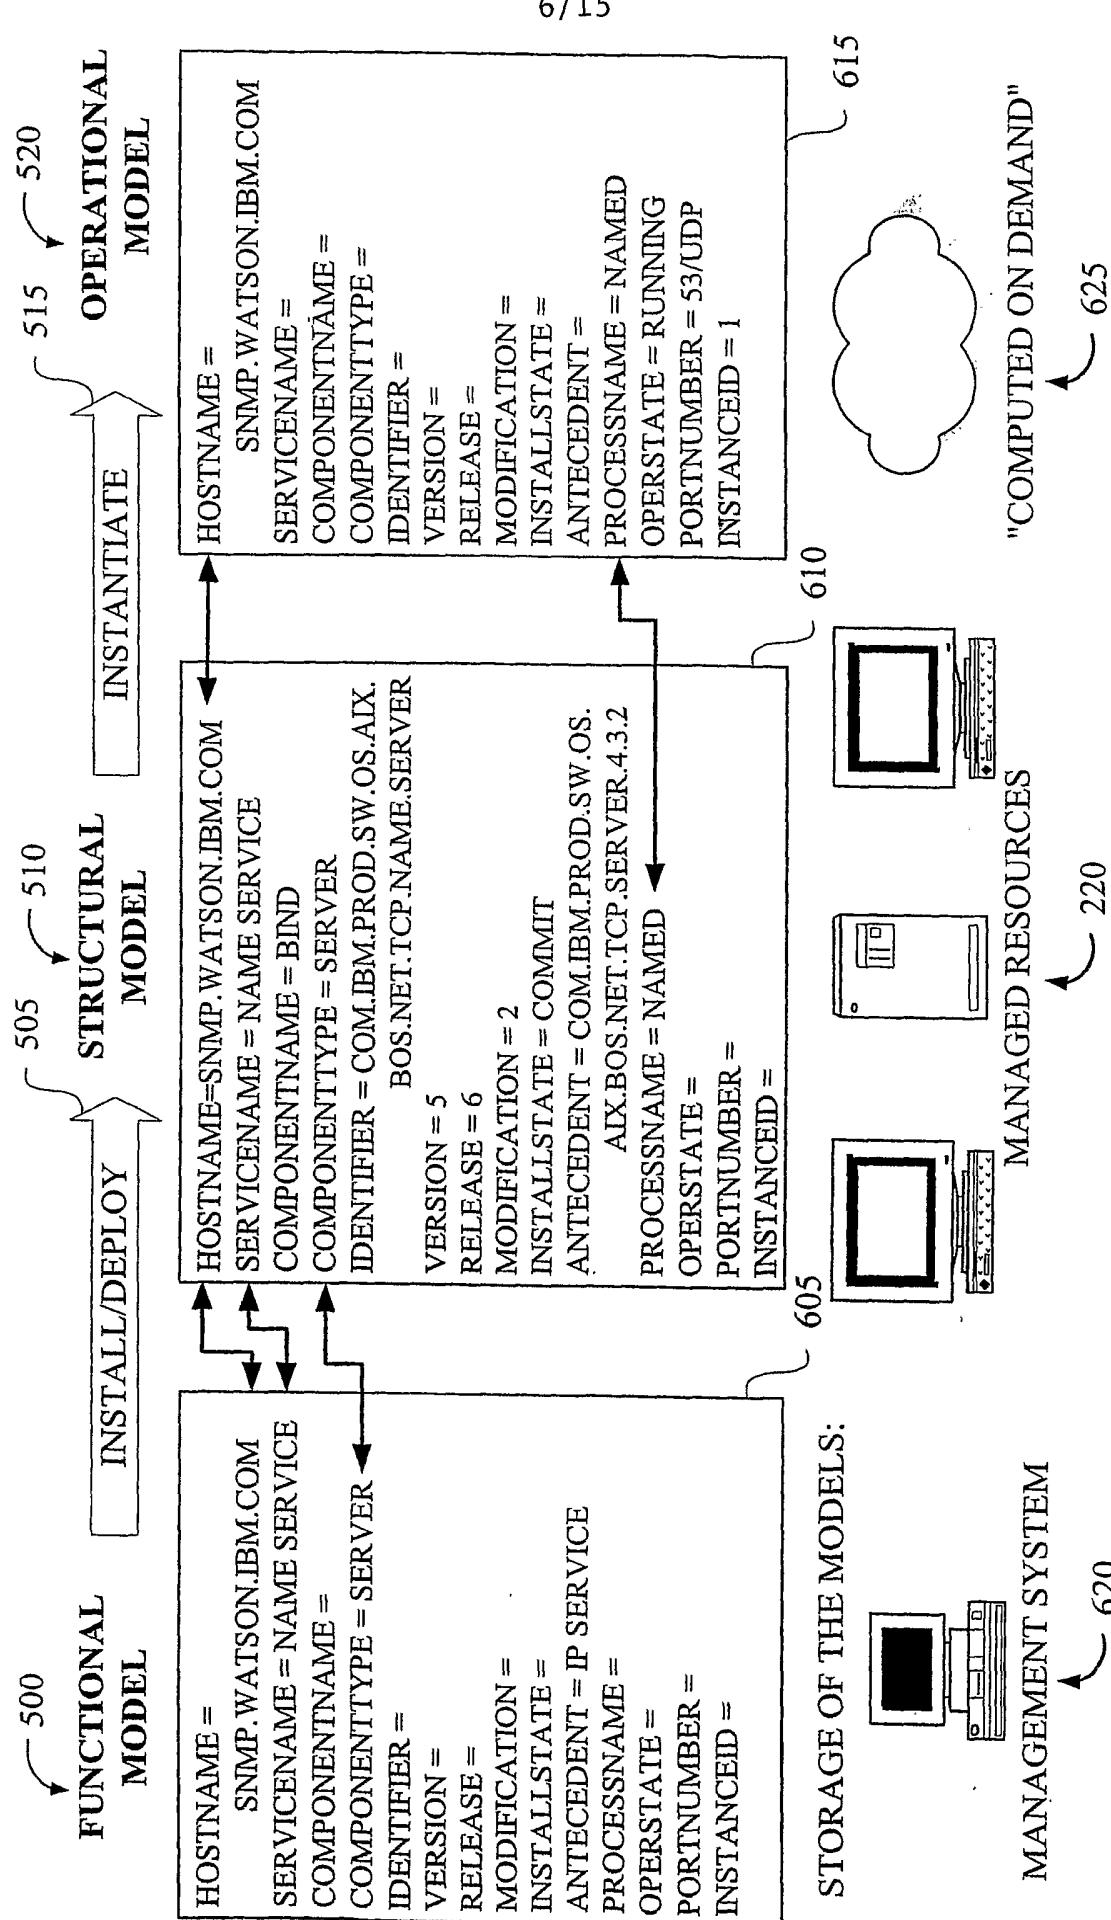


FIG. 6

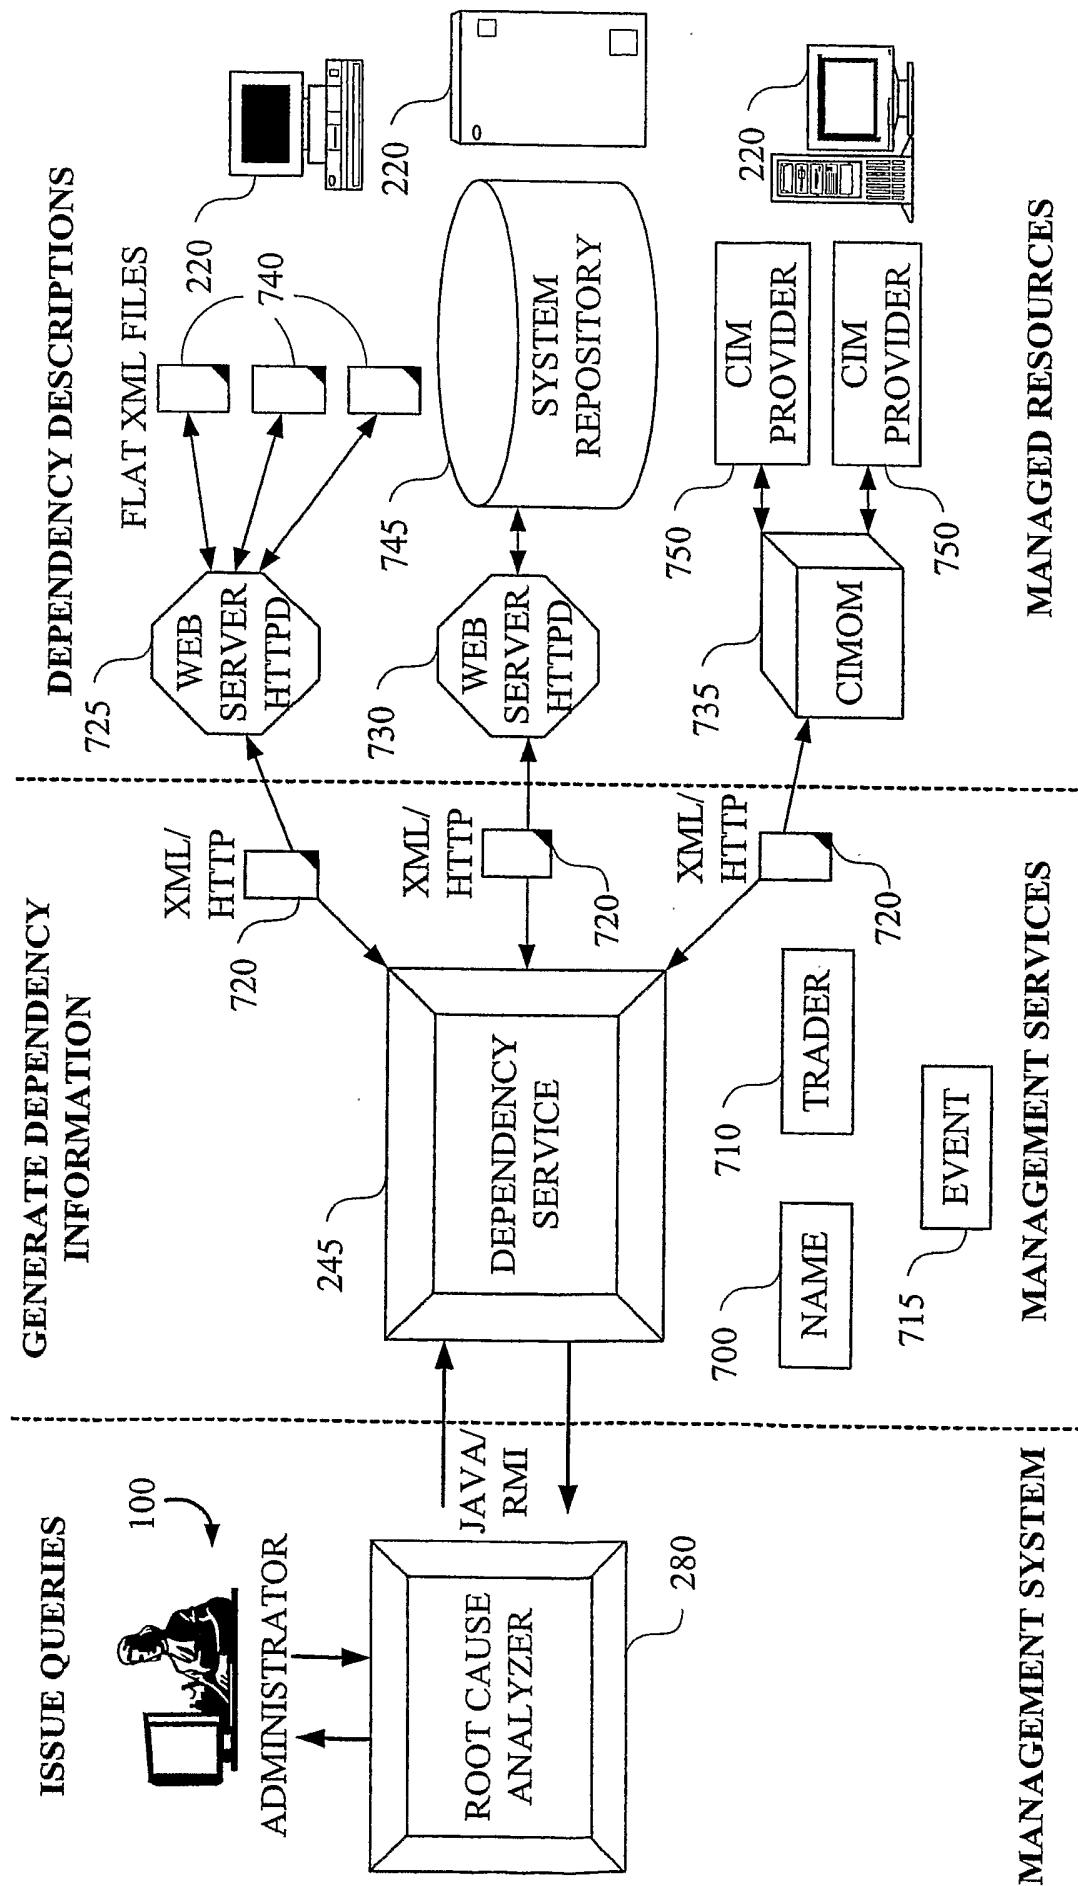


FIG. 7

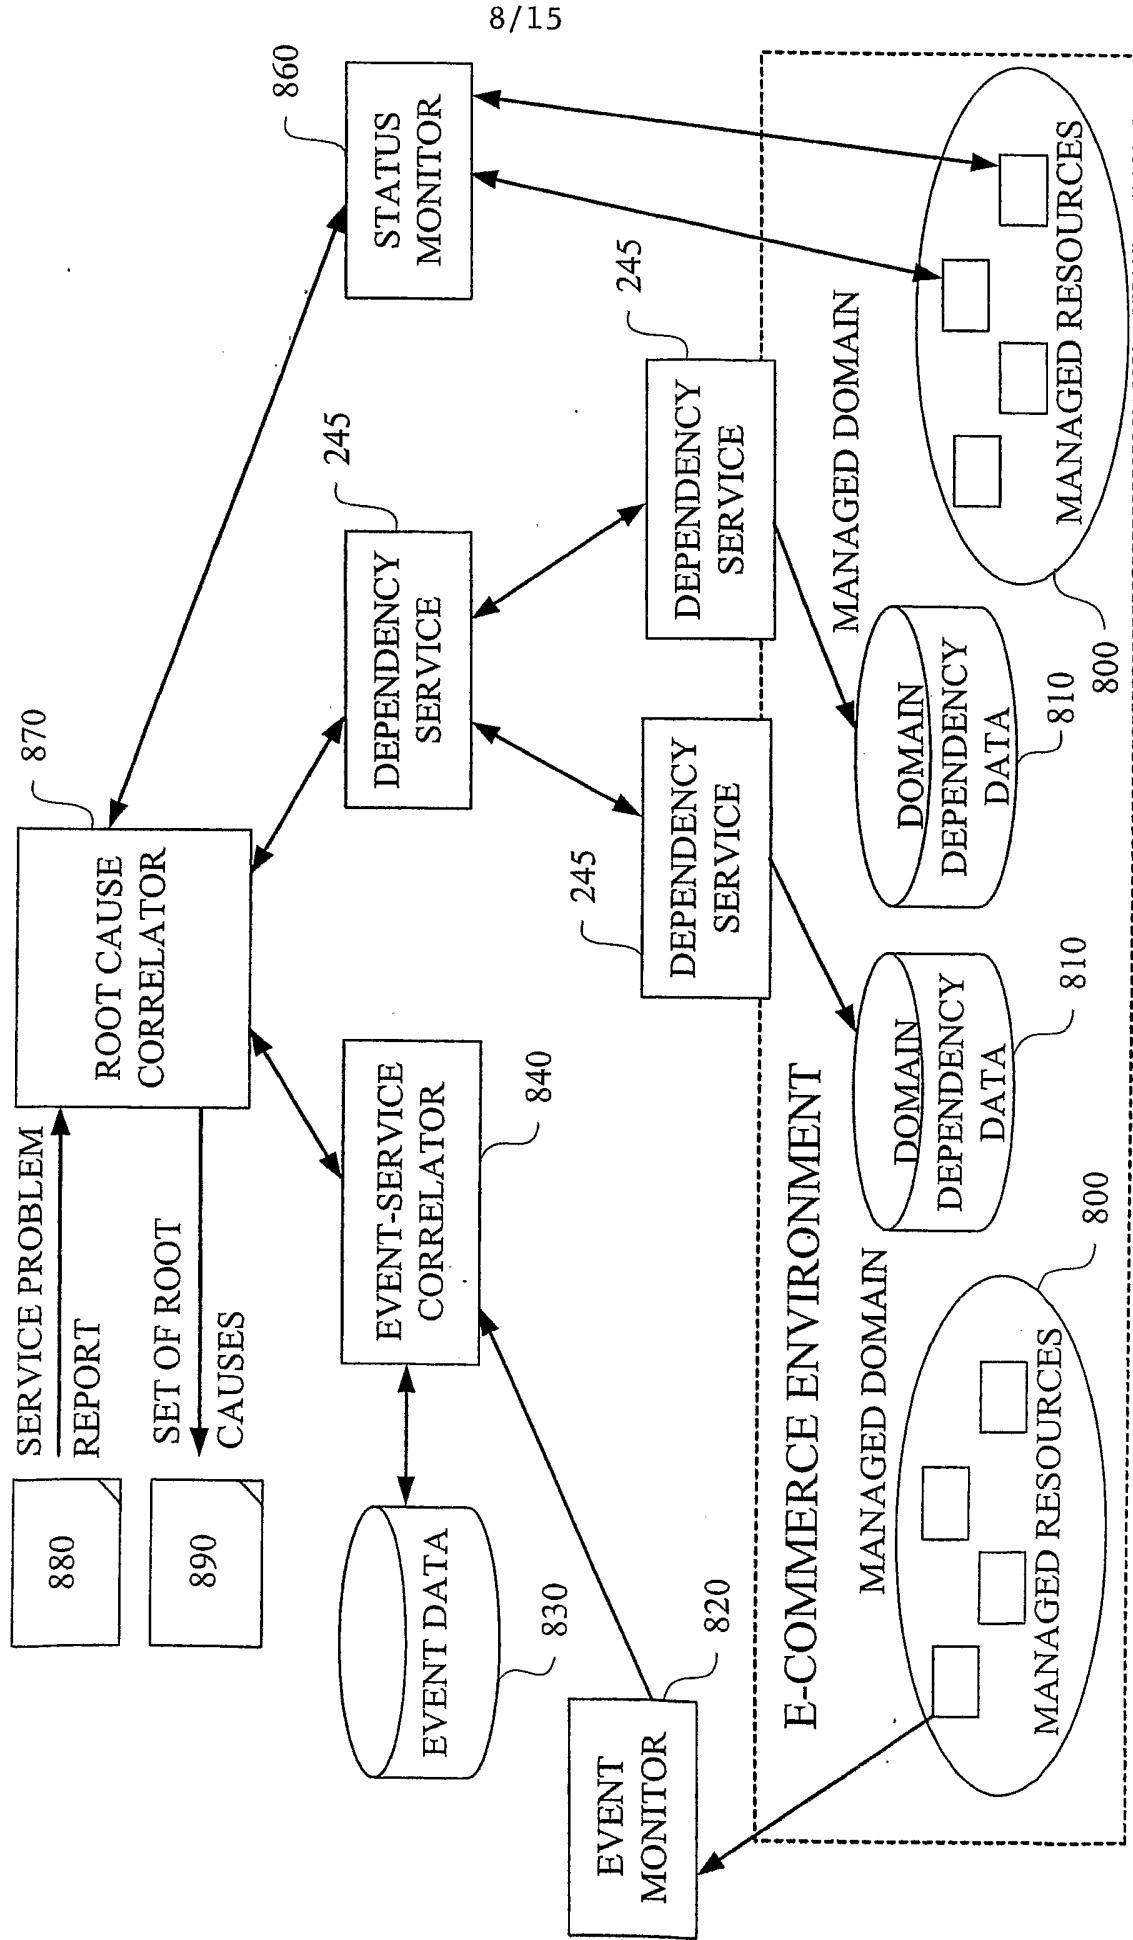
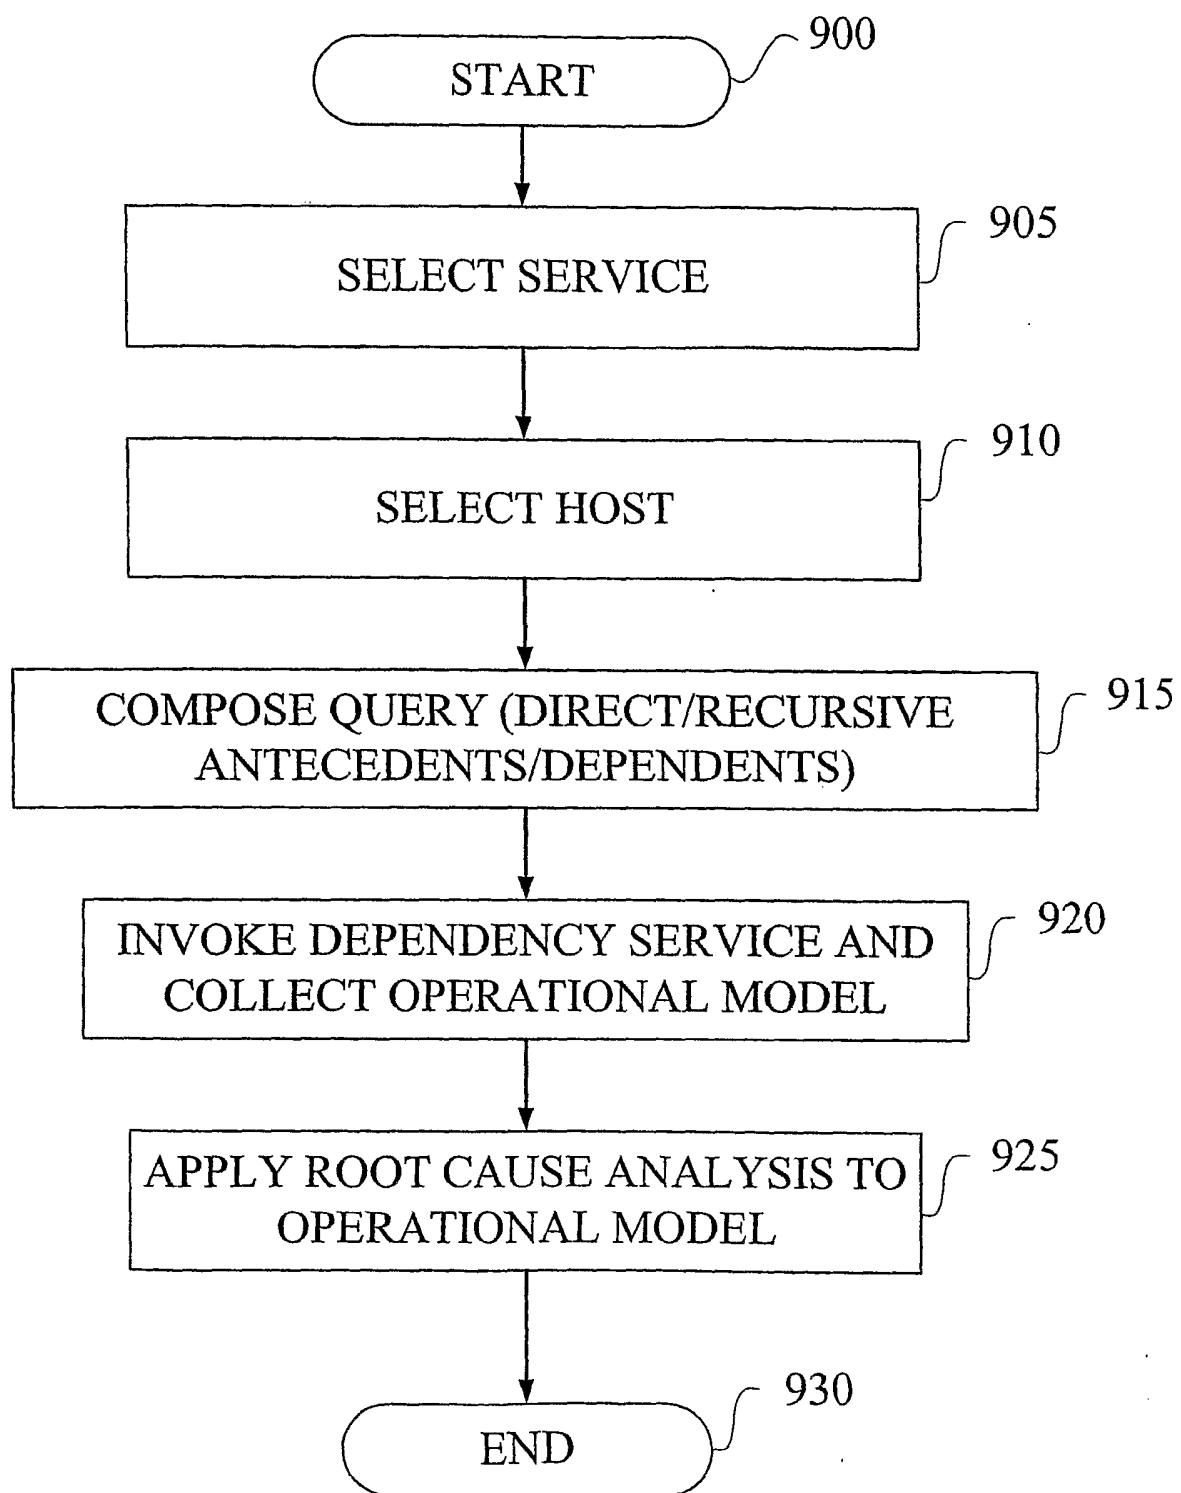
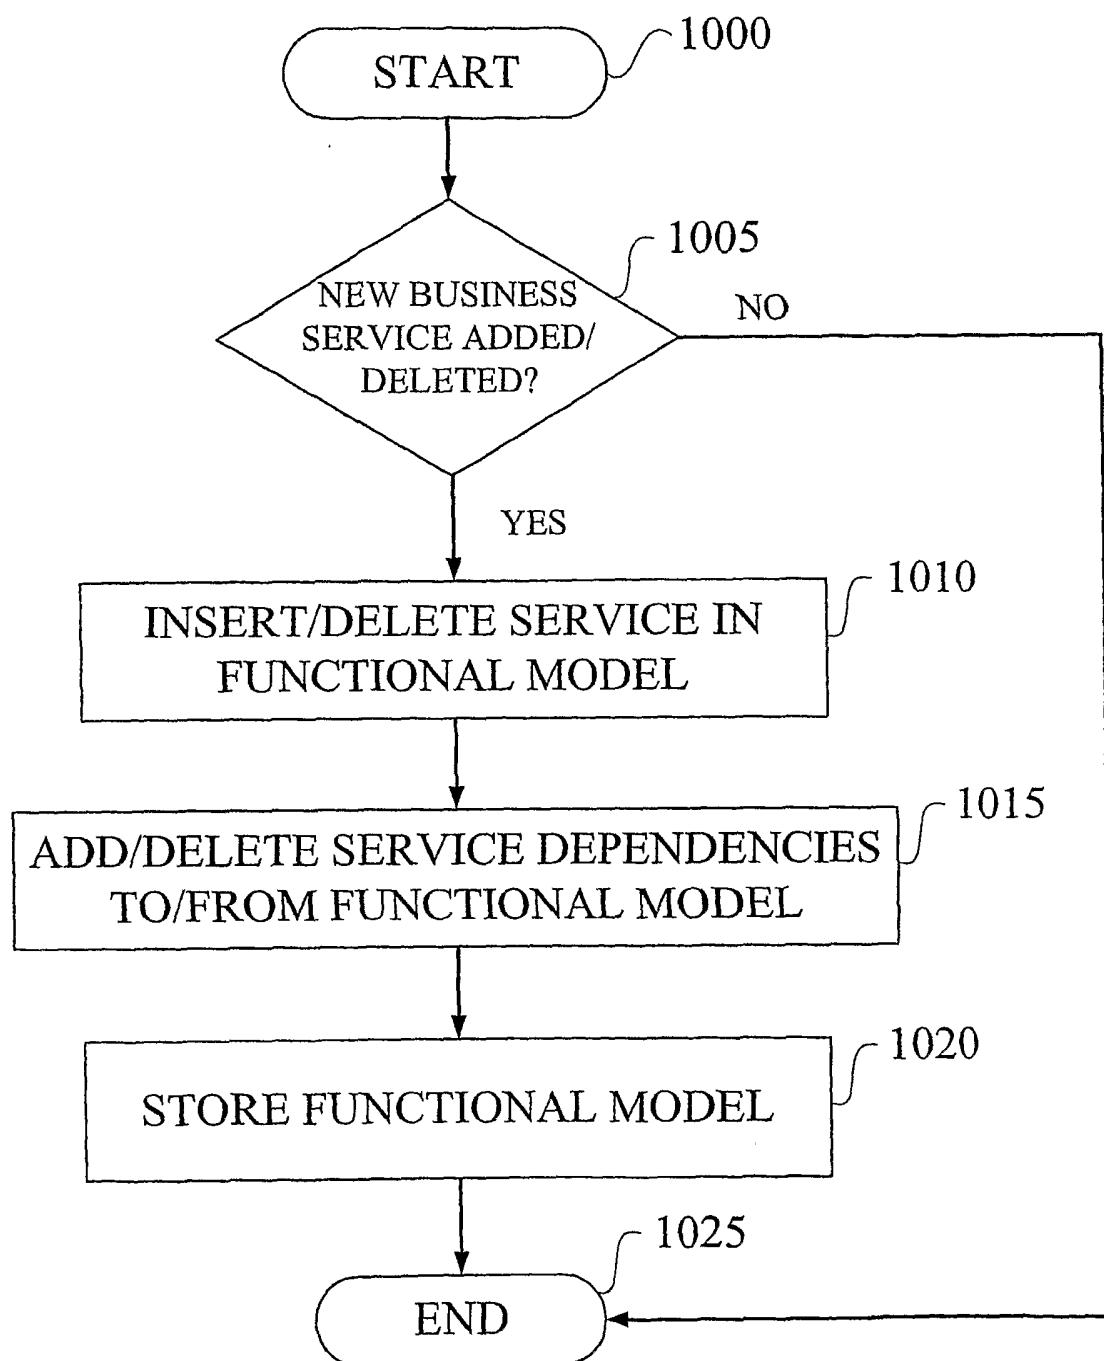
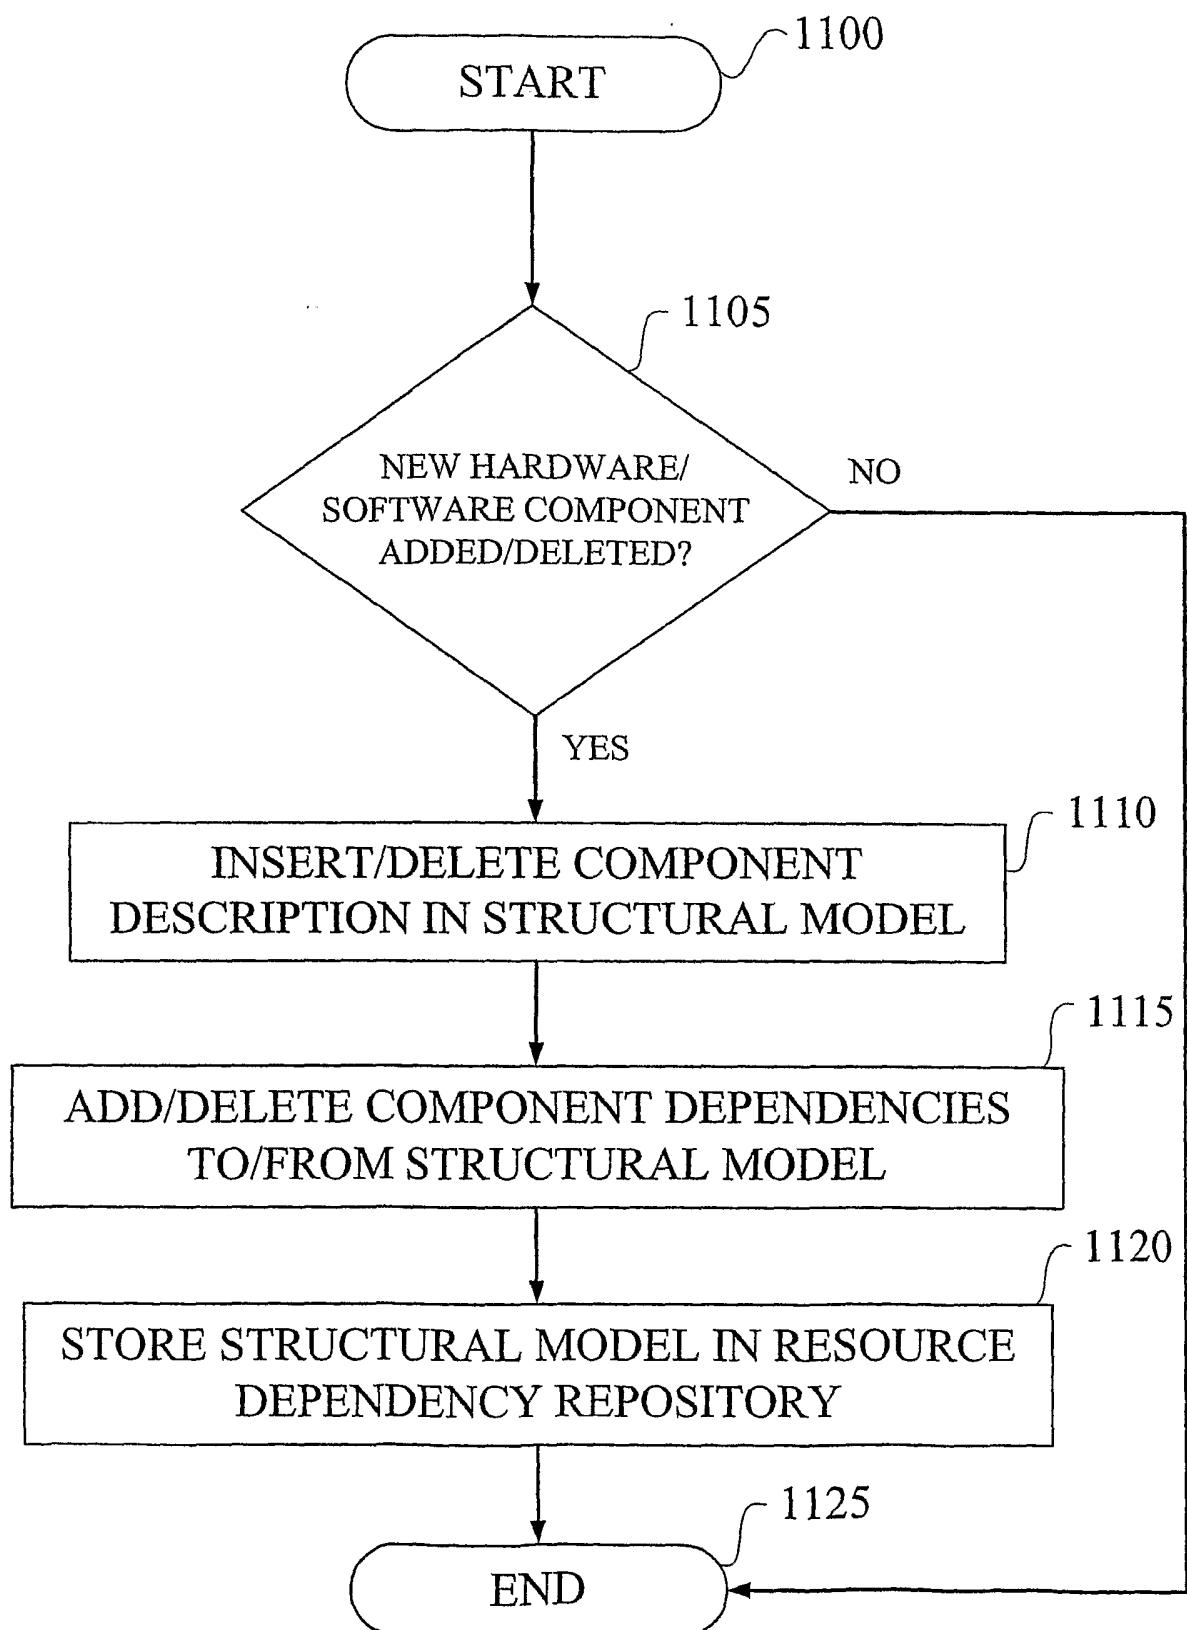


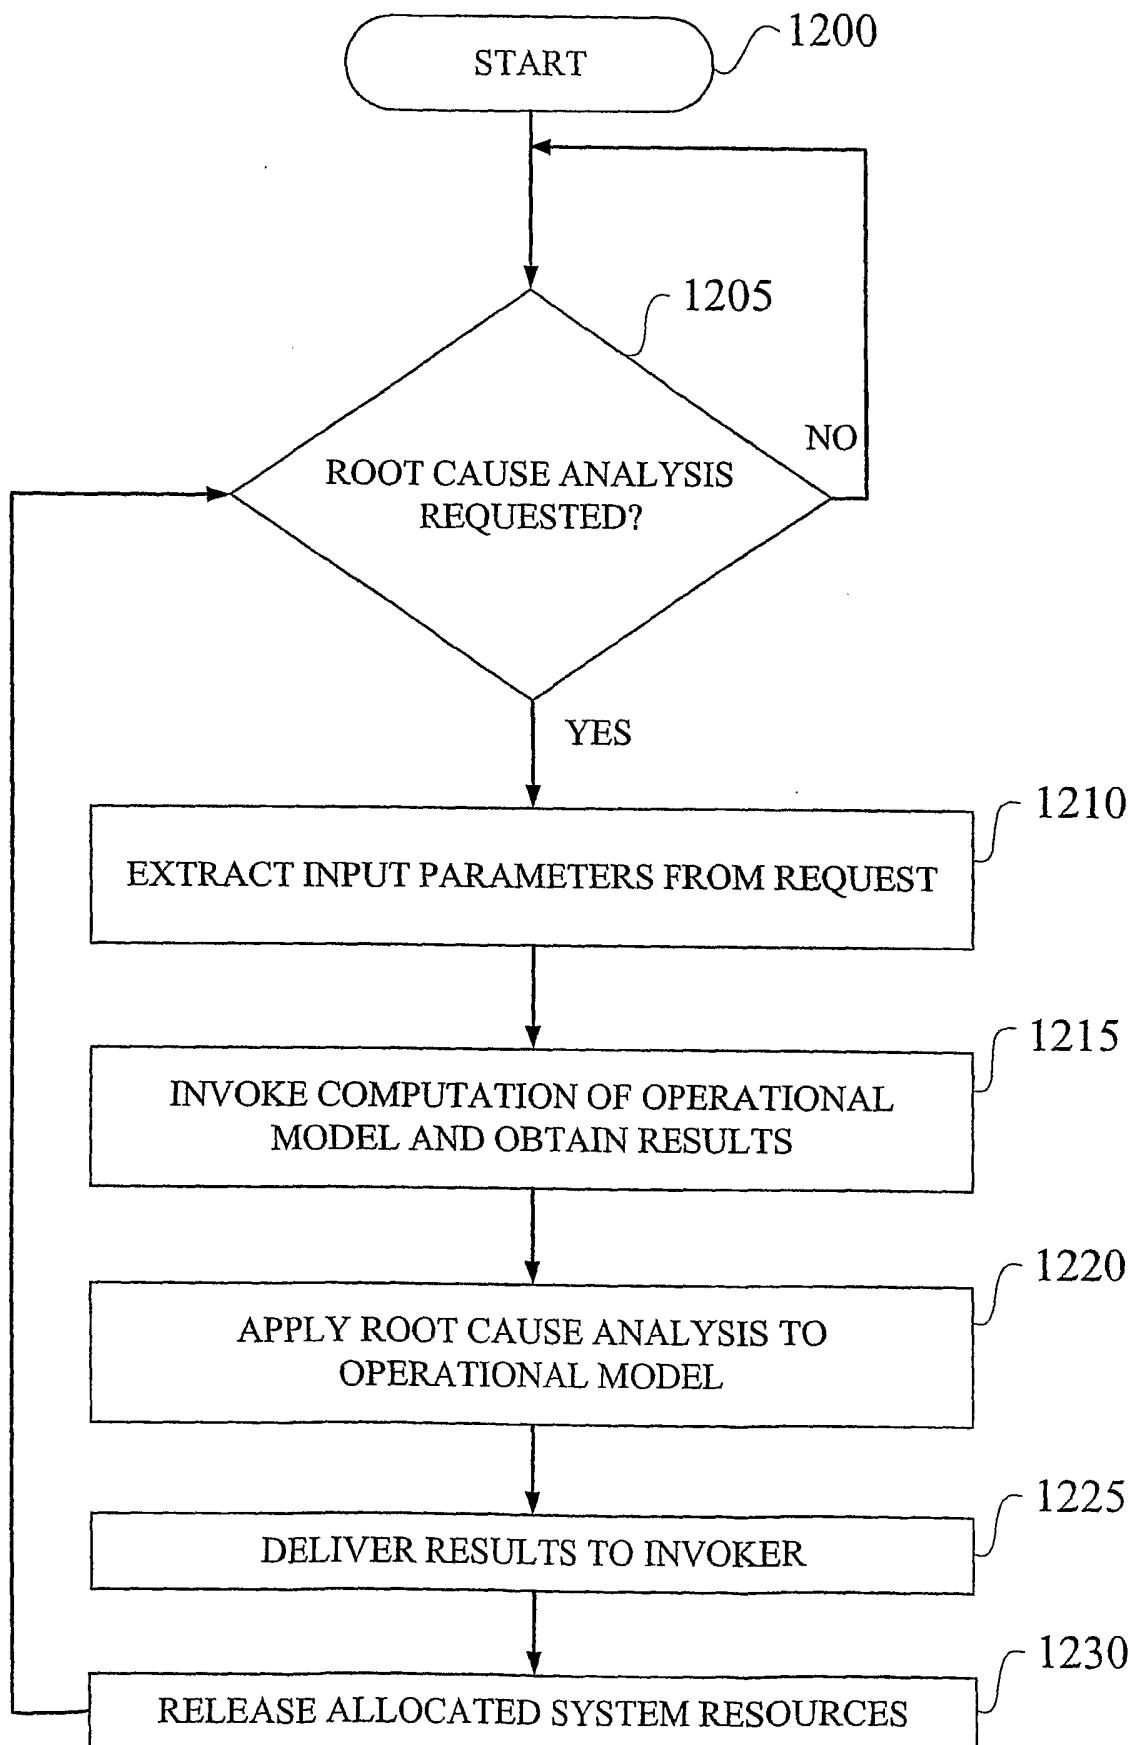
FIG. 8

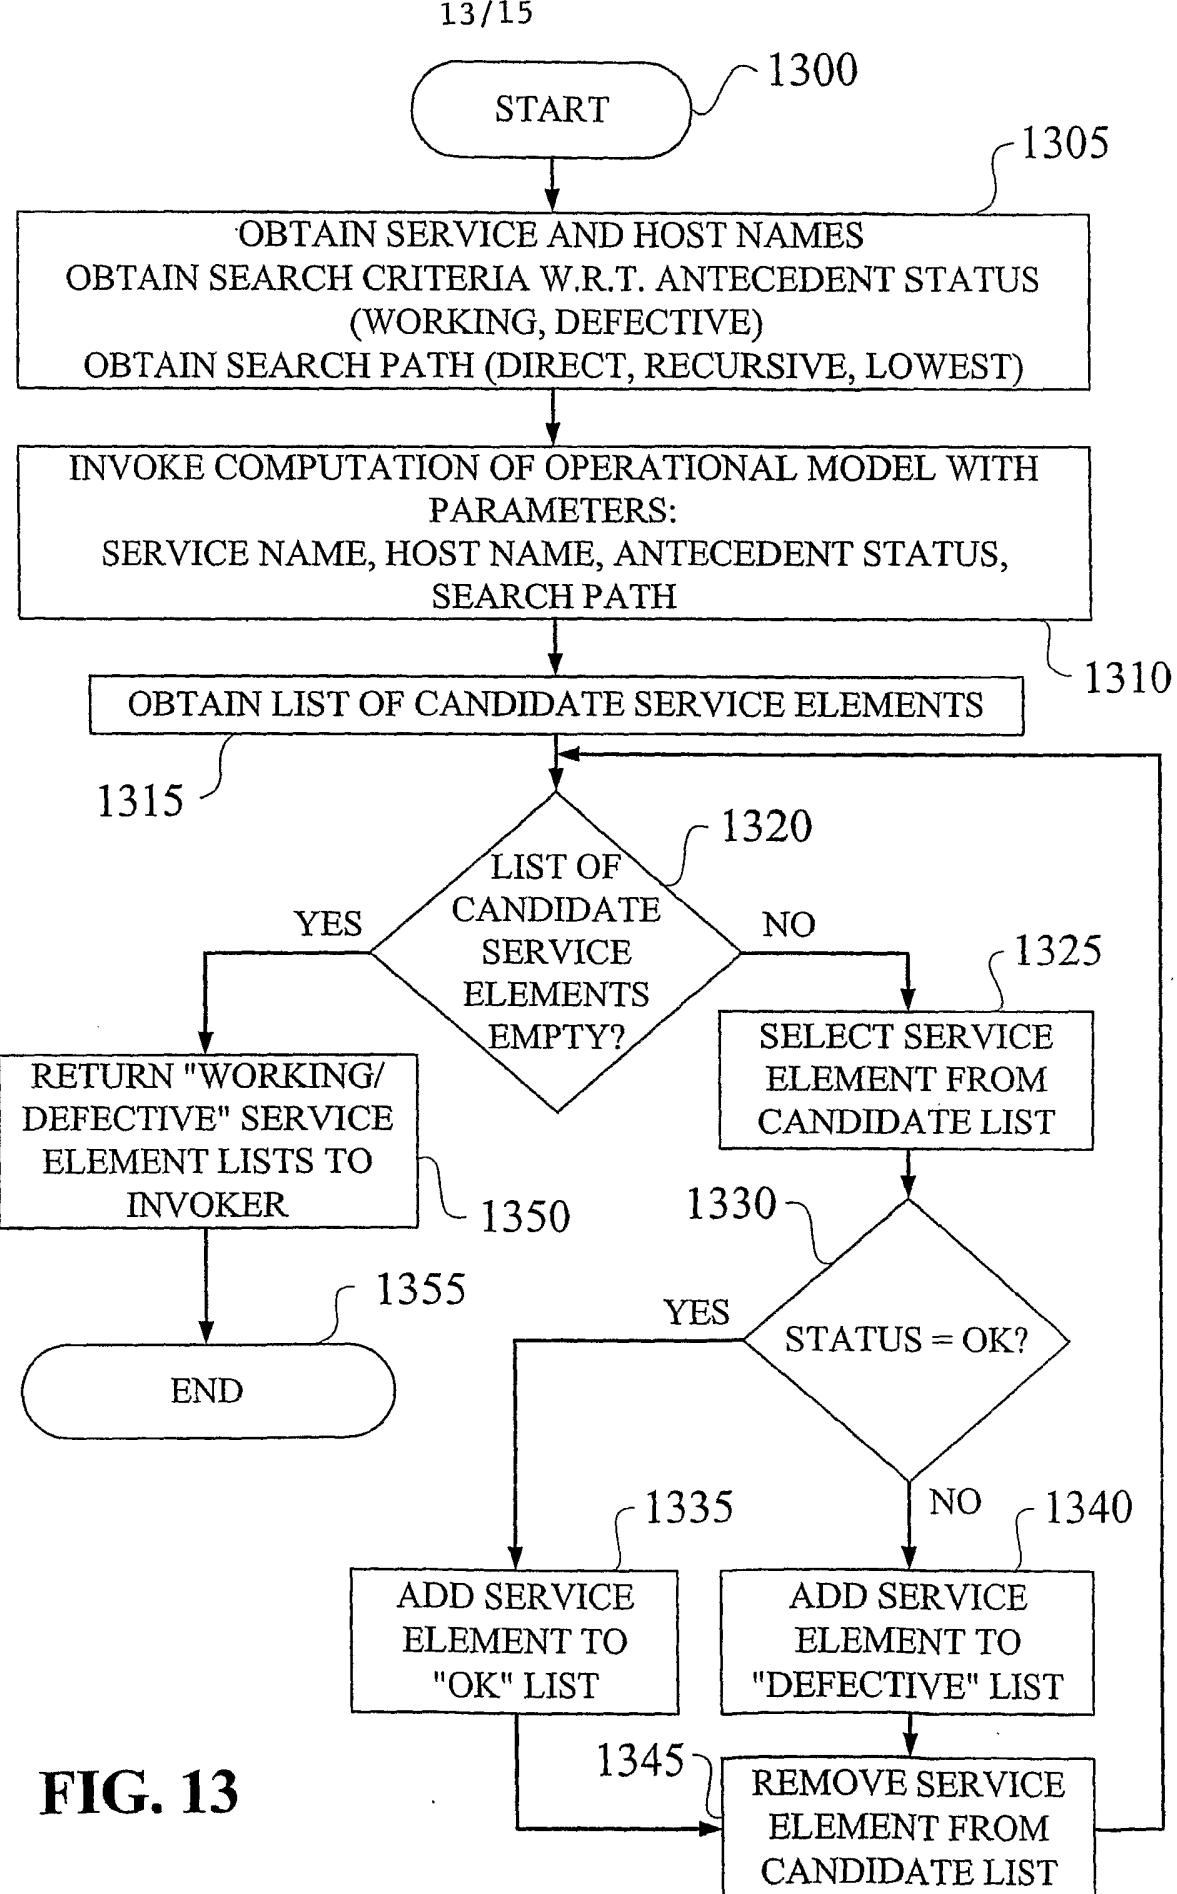
**FIG. 9**

**FIG. 10**

**FIG. 11**

12/15

**FIG. 12**

**FIG. 13**

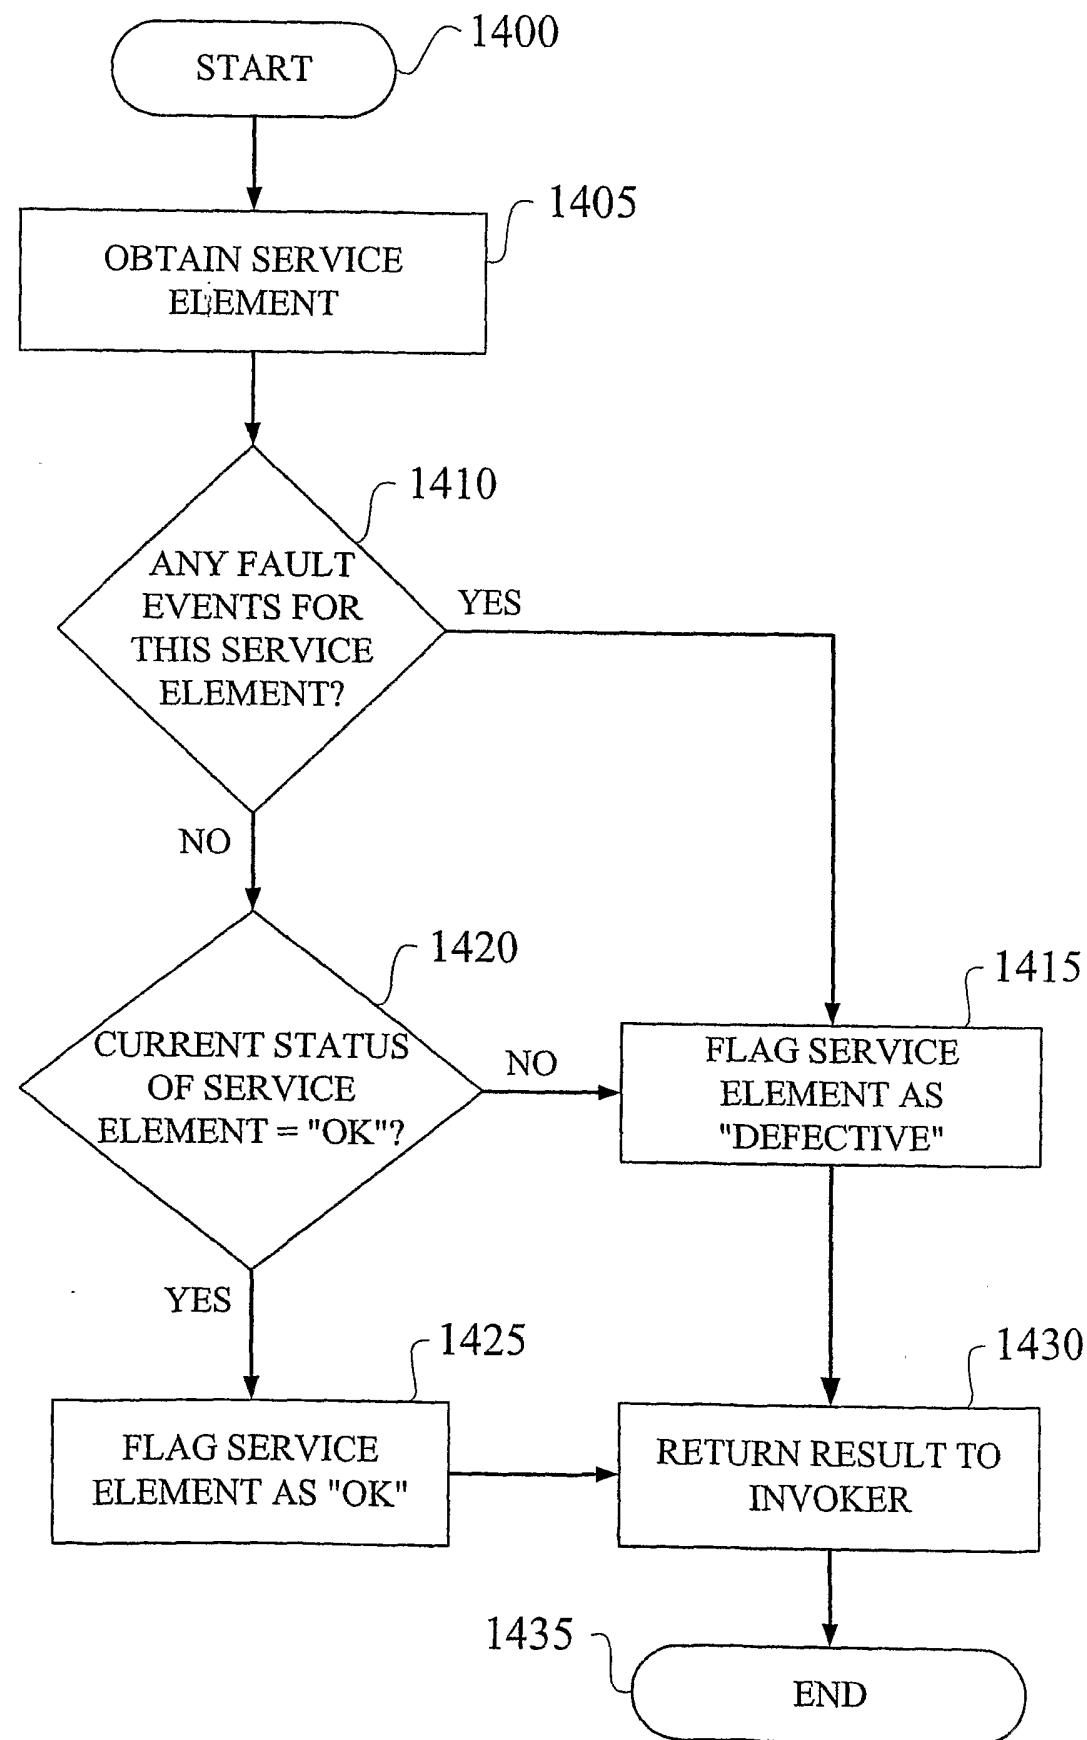


FIG. 14

APIS	FUNCTIONAL DESCRIPTION
GETDIRECTANTECEDENTS ()	RETRIEVE DIRECT ANTECEDENTS OF A SERVICE LOCATED ON A SPECIFIC HOST
GETLOWESTANTECEDENTS ()	RETRIEVE LOWEST ANTECEDENTS OF A SERVICE LOCATED ON A SPECIFIC HOST
GETANTECEDENTSRECURSIVE ()	RECURSIVE DRILL-DOWN; RETRIEVE ALL ANTECEDENTS OF A SERVICE LOCATED ON A SPECIFIC HOST
GETWORKINGDIRECTANTECEDENTS ()	RETRIEVE DIRECT ANTECEDENTS WITH STATUS "UP" OF A SERVICE LOCATED ON A SPECIFIC HOST
GETWORKINGLOWESTANTECEDENTS ()	RETRIEVE LOWEST ANTECEDENTS WITH STATUS "UP" OF A SERVICE LOCATED ON A SPECIFIC HOST
GETWORKINGANTECEDENTSRECURSIVE ()	RECURSIVE DRILL-DOWN; RETRIEVE ALL ANTECEDENTS WITH STATUS "UP" OF A SERVICE LOCATED ON A SPECIFIC HOST
GETDEFECTIVEDIRECTANTECEDENTS ()	RETRIEVE DIRECT ANTECEDENTS WITH STATUS "DOWN" OF A SERVICE LOCATED ON A SPECIFIC HOST
GETDEFECTIVELOWESTANTECEDENTS ()	RETRIEVE LOWEST ANTECEDENTS WITH STATUS "DOWN" OF A SERVICE LOCATED ON A SPECIFIC HOST
GETDEFECTIVEANTECEDENTSRECURSIVE ()	RECURSIVE DRILL-DOWN; RETRIEVE ALL ANTECEDENTS WITH STATUS "DOWN" OF A SERVICE LOCATED ON A SPECIFIC HOST

FIG. 15