



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) **EP 0 735 469 B1**

(12) **EUROPEAN PATENT SPECIFICATION**

(45) Date of publication and mention  
of the grant of the patent:  
**08.01.2003 Bulletin 2003/02**

(51) Int Cl.7: **G06F 9/46**, G06F 11/30

(21) Application number: **96104077.1**

(22) Date of filing: **14.03.1996**

(54) **System and method to control and administer distributed object servers using first class distributed objects**

System und Verfahren zum Steuern und Verwalten von verteilten Objektservern unter Verwendung von erstklassigen verteilten Objekten

Système et méthode pour contrôler et administrer des serveurs d'objets distribués utilisant des objets distribués de première classe

(84) Designated Contracting States:  
**DE FR GB IT SE**

• **McChesney, Roderick J.**  
**Redwood City, CA 94063 (US)**

(30) Priority: **30.03.1995 US 413402**

(74) Representative: **Liesegang, Roland, Dr.-Ing.**  
**FORRESTER & BOEHMERT**  
**Pettenkoflerstrasse 20-22**  
**80336 München (DE)**

(43) Date of publication of application:  
**02.10.1996 Bulletin 1996/40**

(73) Proprietor: **SUN MICROSYSTEMS, INC.**  
**Mountain View, CA 94043 (US)**

(56) References cited:  
**EP-A- 0 636 985**

(72) Inventors:  
• **Balick, Maurice**  
**Fairhaven, MD 20754 (US)**  
• **van Hoff, Arthur A.**  
**Palo Alto, CA 94306 (US)**

• **"SOMobjects Developers Toolkit Users Guide" ,**  
**IBM , ARMONK, USA XP002019177 \* Chap. 6,**  
**Distributed SOM (DSOM) \*\* page XI - page 6-16**  
**\*\* page 6-27 - page 6-37 \*\* page 6-44 - page 6-45**  
**\*\* page 6-50 - page 6-51 \*\* page 6-59 - page 6-62**  
**\*\* page 6-85 \***

**EP 0 735 469 B1**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

**Description**COPYRIGHT NOTICE

5 **[0001]** A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by any one of the patent disclosure, as it appears in the Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

RELATED APPLICATION

10 **[0002]** This application is related to the European patent application EP 0732834 entitled SYSTEM AND METHOD FOR DETERMINING AND MANIPULATING CONFIGURATION INFORMATION OF SERVERS IN A DISTRIBUTED OBJECT ENVIRONMENT, our reference FB 5839, filed simultaneously with this application.

FIELD OF INVENTION

15 **[0003]** This invention relates to the field of object oriented application and operating systems development, and more particularly, to methods and systems for performing administrative operations on object oriented applications operating in a distributed object programming environment.

20

BACKGROUND OF INVENTION

25 **[0004]** Software applications are typically designed as self contained programs that have all the necessary procedural code and data included in the application or its immediate operating environment. Such programs generally have only limited access to the code and data of other applications, and particularly, are unable to execute procedures or access data that are not in the address space of the application. With the rise of object oriented programming, however, an application need not be designed as an atomic entity, but rather, applications can be comprised of objects that can access each other's operational code and data.

30 **[0005]** In object oriented programming, a first class object is an entity that can be directly manipulated by other entities in the program, instead of merely providing information or structures that define the programming environment. For example, in C++, a class definition is not a first class object because a class definition is only accessible to a compiler, and cannot be manipulated by other objects during actual execution of the program. In some object oriented languages, such as Smalltalk, a class definition is a first class object because it can be manipulated during execution of the application.

35 **[0006]** The objects can be designed to work in a distributed environment, typically a network of computers, where objects on various computers can provide various functionality to a given user. This style of application development is called distributed object programming.

40 **[0007]** Distributed object programming may be used to implement a distributed computing system. One predominant model of distributed computing is the client-server model. Here, servers are applications that provide operations and data to client applications. The clients and servers may be distributed in various computers on a network, or may reside in a single computer. In a client-server system based on a distributed object programming environment, a client is code that invokes or manipulates an object, and the entity that is invoked is the server. A server, commonly known as an object server, may contain multiple objects for performing various operations. The client and server can reside in the same process, a different process, in the same computer, or even in different remotely situated computers. A client invokes a server through an object reference, without having information as to where the invoked object resides in the system, that is, without specifically knowing the computer or process where the invoked object is stored or where it's running. Rather, the client passes the object reference through an object request broker which locates the computer executing the server as a process, and passes back any output from the server to the client. To the user, the client and server, and any group of distributed objects, behaves transparently as an integrated application, even though the objects may be distributed in various computers. One standard for the distributed object programming is the *Common Object Request Broker Architecture* (CORBA) specified by the Object Management Group (OMG).

45 **[0008]** Information about a server can be classified into two categories. Dynamic information arises out of a particular process that is executing a server, such as its process identification, its current memory requirements or location, how many objects are active within the server, and the like. This information is dynamic because it depends on a current executing process which is transient. Static information does not devolve from a particular execution or process. This includes external information such as the location of where the server is installed, its pathname, its command line arguments, its configuration setup data, and internal information, such as the identity and description of the interfaces and implementations the server provides.

55

**[0009]** Most operating systems, such as UNIX®, provide system calls useful for determining dynamic information about the applications that are currently executing or ready to execute, such as how much primary memory they have been allocated, the address offset for the address space, which applications are currently running, on which machines, and the like. Such system facilities are used by application developers to debug applications during development, and by system administrators to diagnose and solve problems occurring during actual use. In conventional systems where the processes created by an application all reside on a single computer, information about a process executing an application can be obtained using conventional operating system routines.

**[0010]** However, in a distributed object environment, an object server does not exist as an object, but rather as a collection of discrete, and possibly remotely situated, objects providing related functionality to various clients. In this environment, there is typically no means for determining the state of remote server processes. This is for two reasons. First, there is no first class server object per se, but only a set of distributed objects, and thus a client cannot identify a specific server object for obtaining process information. Second, in a distributed object environment, the only remote constructs on which clients can operate at run-time are first class objects; since servers are embodied by processes (e.g. UNIX® processes), which in most computing platforms are not themselves objects, clients can not easily administrate servers. Without such information, generic debugging, error tracing, and logging are not possible, thereby impeding the development, debugging, and administration of applications created from distributed objects.

**[0011]** Accordingly, it is desirable to provide a computer implemented system and method for obtaining and modifying administrative information and behavior of server processes on remote computers in a distributed object programming environment.

SUMMARY OF THE INVENTION

**[0012]** This invention enables the implementation of generic administration tools for distributed object environments by providing an instance of a first class distributed object (hereafter referred to as the server spy) in each server on a host machine, and by making the interface and implementation of this server spy well suited to administrate the state and behavior of that server by a client holding an object reference to the server spy.

**[0013]** Each server spy object provides structures and methods for obtaining and modifying runtime information for processes created to embody the activated server. The server administrator object is used by client objects to access the server spy object for this information. The server spy object is embedded in each server during application development, by including it in a library of runtime code that is linked into server code developed by an applications developer.

**[0014]** The server spy object allows clients, both within the host computer maintaining a given server, and other remotely situated clients, to determine and control the operation of the server associated with the particular server spy object. Through the server spy object, the client can determine the process identifier of the server, even where the server is executing on a remote computer on the network. The client can control whether the server is subject to automatic deallocation of idle resource. The client may also manipulate, through the server spy, the output of tracing commands that may be included in the server by the applications developer, and where this information is output as log files.

BRIEF DESCRIPTION OF THE DRAWINGS

**[0015]**

- Figure 1 is an illustration of the system of distributed computers for using server spy objects;
- Figure 2 is an illustration of the organization of a host computer supporting various servers;
- Figure 3 is a flowchart of the process of automatically deallocating idle objects in a server;
- Figure 4 is a flowchart of the process for manipulating log files;
- Figure 5 is a flowchart of the process for swapping a set of log files;
- Figure 6 is a dataflow diagram of the method of obtaining state information about a server.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

**[0016]** Referring now to Figure 1 there is shown a hardware view of a system 100 for providing runtime information for servers on remotely distributed computers in a distributed object programming environment. The system 100 includes a number of host computers 101 connected along a network 103. Each host computer 101 includes secondary storage 107 for long term storage of implementation code and data for servers and clients, and the like, an input device 109 and an output device 115 for receiving and outputting commands and data into the system 100, and an addressable memory 113 for storing server (not shown) and client implementation code during execution by a processor 111. During execution by the processor 111, servers exist as processes in the addressable memory 113. Also coupled to the network

103 are a number of clients 105. Each client 105 is an object executing as a process in a remotely situated computer similar in structure to a host computer 101, or alternatively, existing as separate processes in any of the host computers 101. Each client 105 requests services or data from servers in host computers 101 on the network 103. The host computers 101 may be realized by most general purposes computers, such as a SPARCstation™ computer manufactured by Sun Microsystems of Mountain View, Ca. Any other general purpose computer may also be adapted for use with the invention. Each host computer 101 executes a general purpose operating system, such as Sun Microsystems' Solaris® operating system. In addition, each host computer 101 is part of an object request broker environment, satisfying the CORBA standard set by the Object Management Group, in *The Common Object Request Broker: Architecture and Specification*, Rev. 1.2, OMG TC Document Number 93-12-43, available by anonymous ftp to omg.org. Other equivalent environments for distributed object systems may also be used. In the preferred embodiment, the object request broker environment is Sun Microsystems' Project DOE (Distributed Objects Everywhere).

**[0017]** Referring now to Figure 2, there is shown a logical view of system 100, particularly illustrating the configuration of addressable memory 113 in any of the host computers 101 during runtime operation of various servers. The addressable memory 113 of each host computer 101 includes a number of executable objects for performing various functions, including user application functions, such as word processing, database management, graphics design, CAD engineering, software development, financial management, and the like, and operating system functions, such as memory management, resource allocation, storage management, and system accounting information, and network interfacing with remote clients and hosts.

**[0018]** More particularly, a host computer 101 includes a number of servers 201, each providing a particular type of functionality to clients 105 distributed on the network 103 (or locally executing in the host computer 101). Each server 201 is comprised of various object 209. For each server 201 on the host computer 101 there is a persistent server administrator 203 object instantiated from a server administrator class. Each server administrator 203 maintains execution configuration information for the particular server 201 associated with the server administrator 203, and all information that can be known whether or not the server 201 is executing. A server administrator 203 is created by a server administrator factory 205 when a server 201 is installed on the host computer 101. The server administrator 203 is more fully described the related application identified above.

**[0019]** Facilitating communication between the clients 105 and the host computers 101 is a basic object adapter 225, or BOA. The BOA 225 conforms to the requirements for a basic object adapter specified in CORBA. The BOA 225 provides an interface for clients 105 to access object implementations of objects 209 included in servers 201 on various host computers 101 distributed on network 103. Object implementations are actual code which implements an object 209. In order to efficiently manage the addressable memory 113 of the host computer 101, the BOA 225 maintains records indicating which servers 201 are active and inactive, and in each server 201, which objects 209 are active or inactive, and which implementations of object methods are also active or inactive. In this way the BOA 225 can readily manage memory resources in the addressable memory 113 on an as-needed basis, allocating memory to active servers 201, objects 209 and implementations, and deallocating memory from inactive entities. The BOA 225 receives a request from a client 105 for an object 209 in a server 201, the client 105 passing the object reference to the BOA 225. The BOA 225 determines whether there is an active implementation of the object 209 or method and if so, the BOA 225 passes the request to the object 209, otherwise, the BOA 225 retrieves the implementation code for the object 209 from secondary storage 107, starts a new process for the invoked object, executing the implementation with the processor 111. The BOA 225 will also manage deactivation of objects 209 and implementations that are no longer being used, or that are explicitly terminated by a client 105 or server 201, reclaiming resources allocated to the object 209 in the processes or databases maintained by the BOA 225.

**[0020]** Each server 201 in the host computer 101 includes a number of objects 209, each of which encapsulates data associated with the object 209 and its state, and executable code for implementing functional operations, or methods of the object 209 on the data or other objects 209. An object 209 is manipulated by clients 105 or other objects 209 through an interface definition that specifies the operations, attributes, and exceptions that an object 209 provides. The objects 209 are distributed objects, distinguished from conventional objects in object oriented programming languages in that the interface of an object 209 is defined in an interface definition language, which is then mapped to an implementation language, such as C++ or C. This allows objects 209 to have multiple implementations in various languages, while still maintaining a consistently defined interface. A client 105 does not have to have any information about the underlying implementation of the object 209, but merely has to have the interface for the object 209. The objects 209 are distributed in that a client 105 may invoke an object 209 existing anywhere on the network 103, including in the address space of a client 105, or any number of host computers 101, by using an object reference for object 209 (such object reference might be obtained, for example, through the use of a naming server 227) and the various BOAs 225 on the host computers 101 to access the object 209 from the object reference.

**[0021]** Coupled to the host computer 101 by the network 103 is a naming server 227. The naming server 227 maintains a name space directory 229 including number of naming contexts, each of which includes a set of name bindings. A name binding is an association between an arbitrary object name, and an object reference uniquely identifying an

object 209 within a server 201 and further identifying the server 201, and host computer 101 including the object 209. The naming server 227 provides methods for binding a name to an object reference so that the object 209 associated with the object reference can be accessed from the name, and for resolving names provided by clients 105 in order to provide an object reference to a client 105. With an object reference, a client 105 can directly invoke a server 201 or any object 209 within a server 201. The naming server 227 provides clients 105 with a service that allows access to distributed servers 201 before the client 105 has an object reference. Multiple naming servers 227 may be supported, with clients 105 transparently accessing various naming servers 27 on the network 103. In the preferred embodiment, a naming server 227 is included in each host computer 101, and supports objects 209 within the host computer 101; multiple naming servers 227 then work in conjunction to provides clients 105 with object references for local objects 209.

**[0022]** A factory object 211 may be associated with each object 209, for performing a single method that is called by a client 105 to create a new instance of the type of object 209. The factory object 211 is registered in the name space directory 229 of the naming server 227, which is accessed by clients 105 for obtaining object references. Each factory object 211 includes a static create member function that creates a new object 209 of the type specific to the factory object 211, and returns an object reference to the new object 209. An initialize function is provided in created object 209 to perform any needed initialization on the object's state. Typically there is only one instance of a factory object 211 implementation for each object 209 in the server 201. The instance of each factory object 211 is created when the server 201 is installed in the host computer 201 and registered in the naming server 227. An object 209 need not have an associated factory object 211, for example, where the object 209 is not public and available only to the server 201, or where the server 201 is designed to have a single instance of a particular type of object 209.

**[0023]** Each server 201 also includes a server spy 213 object for determining runtime information about the server 201, so that remotely situated clients 105, for example during application development and debugging, can obtain current process and related information. The server spy 213 is a distributed first class object embedded in the server 201 during application development. The server spy 213 differs from other objects 209 in the server 201 in that the methods and data of the server spy 213 operate on, and reflect, the current runtime state of the server 201, whereas the other objects maintained by the server 201 will generally provide application methods and data. Throughout the remainder of this disclosure, references to "the server" and "the server spy" refer to one server 201 and its embedded server spy 213, and it is understood that the functionality described for these respective entities applies to each respective server 201 and server spy 213 in the host computer 101.

**[0024]** When a server 201 is executed by the processor 111 it is an active process in the host computer 101, and the server spy 213 is part of that server 201 process. Accordingly, the operations and attributes of each the server spy 213 are with regard to the current state of server 201 process in which the server spy 213 is embedded. Thus, with multiple server 201 processes executing on host computer 101 at one time, there is a server spy 213 in each server 201 process, and each such server spy 213 reports on the particular server 201 process in which it is embedded.

**[0025]** The server spy 213 does not persistently maintain any information, again, because the server spy 213 exists to provide runtime information about a server 201, and this information is transient and limited to the process created for the particular execution of the server 201 by the processor 111. In an alternative embodiment, a server 201 may contain multiple server spy objects 213, providing for example, different levels of access to the server 201 to different user classes, domains, or the like.

**[0026]** A server spy 213 is included in each server 201 by making it a part of a runtime library which is linked with each server 201 developed and compiled in the distributed object programming environment. The server spy 213 is included in the runtime library by creating an interface definition language file, commonly known as an IDL file, defining the interface of the server spy 213. This interface may be defined by either the provider of the development environment, the operating system provider, or the systems programmer, to best accommodate the needs of users of the system 100. Separate implementation definition and implementation code files are then written to implement the methods of the server spy 213; these files are compiled and then linked into the runtime library, which is used by applications developers to create new servers 201. Thus, the server spy 213 has an implementation like any other object 209, but the applications developer does not have to include the server spy 213 explicitly. Rather, the server spy 213 appears automatically in the servers 213 during compilation and linking. This embedding process makes all servers 201 in the distributed object programming environment capable of responding to a known set of messages, defined in the interface of the server spy 213, without any work by the applications developers. This allows the system administrator to develop generic administration tools that are guaranteed to be able to communicate with any server 201 via the server spy 213. Once such administration tool is the server administrator 203, which is more completely described in the related application referenced above.

**[0027]** The server spy 213 implements an IDL interface, and thus can be called by clients 105 which exist in processes external to the server 201 process including the server spy 213. This distinguishes the server spy 213 from conventional language specific objects that are normally linked into a server 201 or application program by the development environment compiler and linker, since these objects are only accessible to other objects that are part of the server 201 process.

**[0028]** Because the state of server spy 213 reflects the state of the server 201 at any given time, the server spy 213 effectively instantiates the server 201 as a distributed object in the system 100, providing clients 105 with the same type of access to the internals of the server 201 that such clients 105 have to other distributed objects 209 in system 100. In the absence of the server spy 213, clients 105 would be unable to determine and manipulate the server 201 in the same manner as any other distributed object 209, since conventionally, a server 201 or application program, when executing, is merely a transient process that cannot be accessed separately by clients 105 in the distributed environment. Accordingly, server spy 213 provides for increased flexibility and efficiency in operation and development of servers 201.

**[0029]** The server spy 213 is created when a server 201 is registered with the object request broker environment. Registration notifies the BOA 225 of the existence of the server 201, particularly the path name, arguments, and environment for the server 201, and which interfaces the server 201 implements. Registration also creates a factory object 211 for those objects 209 for which a factory 211 is publicly available, binding a name specified for the factory 211 object by the applications programmer with an object reference in the naming server 227. The server spy 213 is created by an invocation of a constructor function for the server spy 213. In the preferred embodiment, the constructor function determines the server administrator 203 associated with the server 201 that the server spy 213 is to be embedded in, and then assigns the server spy 213 to that server administrator 203 by providing the object reference of the server spy 213 to the server administrator 203. This allows the server administrator to invoke the methods of the server spy 213 as needed. In embodiments where there is no server administrator 203, other methods may be employed to make the server spy 213 publically available. During registration, the server spy 213 will execute its facilities attribute (as described below) to determine what tracing facilities are available in the server 201, providing this information to the server administrator 203. The server spy 213 may also be registered with the BOA 225 so that the server administrator 203, or remote clients 105 may access the server spy 213 through its object reference.

**[0030]** Figure 6 illustrates the call architecture for using the server spy 213 to obtain current state information about a server 201. In the preferred embodiment, clients 105 do not have public access to the server spy 213. Rather, access is obtained indirectly through the server administrator 203. A client 105 invokes 601 an appropriate method on the server administrator 203 for obtaining current state information about the server 201 associated with the invoked server administrator 203. (In an alternative embodiment, the server spy 213 is publicly accessible to clients 105 on the network 103, and may be invoked 601 directly). The server administrator 203 in turn calls 603 the server spy 213 for the desired information. The server spy 213 executes a particular method, such as those described below, either invoking 605 further methods provided in the BOA 225, or the operating system 609, or other lower level object or process. The server spy 213 then returns 607 the resulting information to the server administrator 203. The server administrator 203 may further return 611 this information to the client 105, if desired by the client 105. The current state information obtained by the server spy 213 can include both dynamic information dependent on the server 201 process, such as the process identification of the server 201, non-runtime information that alters the runtime operation of the server 201, such as the operation of predefined API calls in the server 201, or even configuration information, such as the execution definition of the server 201.

**[0031]** Each server spy 213 includes a number of attributes used to transiently store and manipulate information about the server 201 process that the server spy 213 is part of, but the specific functionality of the server spy 213 can vary according to the needs of the system administrator and other users. A pseudo-interface definition file of one embodiment of the interface of the server spy 213 is shown in Appendix A; the interface disclosed in Appendix A is useful to define the functionality of the server spy 213 as described herein. It is understood that other interface definition files may also be suitably used to define a server spy 213 object. Similarly, various implementation files can be used to implement the interface definition of the server spy 213.

**[0032]** Because one of the functions of the server spy 213 is to make the server 201 available to remote client 105 as an observable process in the host computer 101, the server spy 213 includes an attribute for identifying and storing the process identifier provided by the operating system for the server 201 process including the server spy 213. An example of the interface for this attribute is the pid attribute. This attribute allows a client 105 that is invoking the server 201 to access any operating system functions that examine a process, including its memory allocation, execution time, and the like, for determining such information about the server 201 containing the server spy 213. In the absence of the server spy 213, a client 105 could not obtain this information about the server 201 because the client 105 does not have a name or object reference for the server 201. In the preferred embodiment, the pid attribute of the server spy 213 provides an interface to a method that calls an existing operating system function that returns the process identifier of the server 201, for example, using the getpid() command in the UNIX® operating system. With the pid attribute, the server spy 213 will return the process id of the server 201 process to the server administrator 203 associated with the server 201, which would then pass the process id to client 105.

**[0033]** The server spy 213 also includes an attribute for determining and storing a count of the active objects 209 within the server 201, that is those objects 209 which are currently running and allocated resources, particularly portions of the addressable memory 113. An example of the interface for this attribute is the active\_object attribute. Where the

implementation language of the server 201 is C++, for example, the active\_object attribute indicates the number of C++ objects 209 that exist in the server 201 to service requests originating from clients 105 in the system 100. The active\_object attribute provides a measure of how heavily loaded the server 201 is at any one time, and supports resource management and utilization by the client 105 and the operating system.

5 **[0034]** Each server 201 contains one or more implementation, and each object 209 in server 201 belongs to exactly one of these implementations. The server spy 213 maintains separate counts for the number of active objects created by each implementation when the implementation is executing. An example is the ActiveObjects struct.

10 **[0035]** In the preferred embodiment, the count of active objects 209 is maintained as follows. When a server 201 is initialized, a linked list of all implementations is created; this is a list of the individual classes for the object 209 provided in the server 201. Each implementation in the list is a header to a separate linked list of active objects 209 for that implementation; each such linked list is updated as new instances of objects 209 of the given class are created or destroyed. The code for creating these linked lists is provided by a runtime library that is linked into the code of the server 201 when the server 201 is compiled by the applications developer. The server spy 213 has direct access to these linked lists. When a count of active objects is requested by a client 105, by invocation of the active\_objects attribute, the server spy 213 traverses the linked list of implementations, and for each implementation, counts the current number of object instances in the linked list associated with the implementation. This count is stored in the ActiveObject struct with the name of each implementation. The server spy 213 totals the count of active objects across all implementations, and returns this to the client 105. Alternatively, a count of object instances for each implementation may be maintained in the implementation linked list, with the server spy 213 traversing only this list. In the preferred embodiment, the server spy 213 has exclusive access to the linked list structures, thereby ensuring an accurate count of active objects. Exclusive access prevents automatic deallocation by the reaper 217 because the reaper 217 cannot obtain access to the list of active objects, which is necessary to perform the deallocation.

15 **[0036]** The server spy 213 also includes a flag for indicating whether automatic deallocation is to be performed by the reaper 217. An example of the interface for this attribute is the boolean attribute reaping. A client 105 may wish to disable the reaper 217 in order to debug the server 201, thereby preventing the server 201 from being terminated by the reaper 217 during the debugging process, which may make the server 201 appear inactive for long periods of time. Similarly, a system administrator may want to disable the reaper 217 to inspect the server 201, alter its configuration information, and the like.

20 **[0037]** The reaper 217 is a separate thread executing in each server 201 process. When the reaping flag is set, the reaper 217 will operate normally, automatically deallocating resources from objects 209 that have been idle for a specified period of time, as preferably established by the applications developer. To determine which objects 209 are active, the reaper 217 obtains exclusive access to the linked lists of active objects, as described above. Figure 3 illustrates a flowchart for one embodiment of the method implemented by the reaper 217. After waking 301, the reaper 217 loop 303 over all implementations in the server 201 in which the reaper 217 is embedded. Looping 305 over the linked list of active objects for each implementation for each active object 209, the reaper 217 determines 307 whether the object 209 has not been invoked within the time period specified, that is, if the difference between current time and a time stamp value for the last invocation of the object 209 is greater than the specified idle time. If so, the reaper 217 deallocates 309 resources from the idle object 209 by indirectly calling its destructor functions in conventional manner. Otherwise, if the reaping flag is not set, the reaper 217 will itself be inactive, and active objects in the server 201 will be deallocated only by explicit termination of the server 201. Alternatively, instead of a single attribute functioning a flag to the reaper 217, the server spy 203 may include distinct methods for directly starting and stopping the reaper 217; this would allow direct manipulation of the reaping function.

25 **[0038]** The server spy 213 includes an attribute for determining and setting the cycle time of the reaper 217, that is, how frequently the reaper 217 of each server 201 is activated by the host computer 101 to perform deallocation. An example of the interface for this attribute is the reaper\_cycle\_time attribute.

30 **[0039]** The server spy 213 also has an attribute for determining and setting an idle time specific to the server 201 containing the server spy 213, that is, the number of seconds during which no object 209 within the server 201 is invoked. An example of the interface for this attribute is the server\_timeout attribute. Each time an object within the server 201 is invoked, current time is stored in the server 201. Referring to Figure 3, the reaper 217 compares 311 this time value with the current time, and if the difference exceeds the specified idle time in the server\_timeout attribute, then the reaper 217 deactivates 313 the server 201 by calling the server administrator 203 associated with the server 201, indicating server 201 is being terminated, and then invoking a shutdown method provided by the shutdown API 223.

35 **[0040]** The distributed object programming environment used in system 100 supports a number of application programming interfaces (APIs) for enabling each server 201 to access various operating system calls and predefined application routines, such as input and output functions, user interface management, and the like. In this fashion an API alters the operational characteristics of a server 201 by providing functionality within the server 201 that is not directly coded by the applications developer. Rather, an API provides the applications developer with set of calls that can be included in the implementation code of an object 209 or server 201 for invoking API functions during execution.

There is provided in each server 201 a tracing API 219, a logging API 221, a reaper API 217, and a shutdown API 223.

**[0041]** The tracing API 219 provides for control of the conditional output of diagnostic information about a server process, useful in the determination of errors generated at runtime, or for identifying the causes of improper results from application execution not involving an error. The logging API 221 provides for directing where such diagnostic information is output. Separating these two functions increases the flexibility of diagnostic tools, and allows for various remote clients 105 to individually control the logging and tracing behavior of a given server 201. The reaper API 217, provides a method for automatic deallocation of inactive or idle objects 209 and resources. The shutdown API 223 provides a method for the graceful shutdown and cleanup of the server 201.

**[0042]** Conventionally, an applications developer is limited to including various API calls in the server code, and permanently configuring their operation, so that an application user has no means for dynamically configuring which API calls are executed once the server 201 is compiled into an executable application. For example, to trace execution of a server 201, an applications developer would include specific API calls, in the desired tracing locations; this code could not be changed once the server 201 is installed as an executable application, so that the API call is always operable, thereby limiting the flexibility of such tracing features. The applications developer may provide for conditional control of the trace statements, for example to output all data stored in a server 201, but only when trying to diagnose a particular problem. However, such routines add complexity to the server 201. Additionally, the applications developer would identify in the server 201 code specific log files for outputting the trace information, again, limiting the user's or system administrator's ability to alter such output.

**[0043]** In order to provide for uniform tracing functions, and simplify the process for the applications developer, the distributed object programming environment provides specific APIs for tracing and logging functions. The server spy 213 object can enable and configure these various APIs within a server 201 during execution, even where the user is remotely situated from the host computer 201 maintaining the server 201.

**[0044]** The tracing API 217 provides for conditional tracing control through facilities. A facility is a set of tracing flags that are available as a predefined set. The applications developer can define various sets of facilities to include in the server 201 code. For example, facilities can be defined to track loop entry or exit, loop execution time, disk reads or writes, and the like. Each trace flag in a given facility is associated with a specific method that implements a trace function included in runtime library of the distributed object programming environment, or separately provided by the applications developer. The applications developer specifies in the server 201 code, at various desired locations, a trace message specifying a given facility, one or more flags for controlling various trace functions, and an output value. A function prototype (in C++) is:

`DOT_TRACE(char* facility_name, char* trace_flags, char* format, ...);` where DOT\_TRACE is an API call provided by the tracing API 219; any number of output values can be specified following the format specification.

**[0045]** In a distributed object programming environment, these tracing and logging functions cannot be directly accessed for remote servers 201 by the clients 105. However, the server spy 213 provides remote clients 105 the ability to control this information for any server 201 on the network 103. The server spy 213, in conjunction with the trace 219 and log 221 APIs, enables such users to directly manipulate tracing and logging conditions for a server 201 during execution, by providing access to API calls that control the operation of the API functions. The applications developer incorporates various desired API calls in the server 201 application. During execution of the server 201, a remotely situated client 105 can invoke the server spy 213 to configure the operation of these calls at runtime.

**[0046]** For controlling tracing, the server spy 213 can invoke (enable/disable) a given facility and specific trace flags available in the facility, resulting in the generation of the output value when the trace message is encountered during execution of the server 201. The output of trace messages for a given server 201 are sent to the tracing API 217 which separately maintains a list of log files for storing all trace outputs from the server 201. The log 219 is used to determine the appropriate log file or files for the server 201, and outputs the trace message there. The separation of tracing control and log control enables remote clients 105 to individually manipulate both the tracing operation and the output location for a server 201.

**[0047]** The distributed object programming environment of the preferred embodiment provides a "default" trace facility with the following trace flags:

admin	provides administrator control
config	configuration information
heartbeats	track heartbeats sent by server
install	installation details
object	object lifecycle internal details
reaper	removal of inactive objects
run_verbose	extensive server startup and shutdown information



(continued)

data	persistent data operations
------	----------------------------

5 Other trace facilities may be defined by the applications developer as desired in an API call to the tracing API 219.

**[0048]** Accordingly, to access the tracing facilities included in the server 201, the server spy 213 includes the following functionality. First, the server spy 213 is able to inform a client 105 of what facilities are available within a server 201. A facilities attribute outputs to the client 105 a list of facilities, including for each facility a facility name, a description of the function or purpose of the facility, the set of trace flags included in the facility, and boolean settings of the trace flags. The server spy 213 then can enable or disable any of the listed facilities, with the client 105 invoking an enable or disable method of the server spy 213, specifying a facility name, and a set of trace flags to be set. An example of the interface for these methods are the trace\_enable() and trace\_disable() methods.

**[0049]** For example, a client 105 may enable a default facility called "default" with the flags "run," and "method," by passing the following command to the server spy 213:

```
15     trace_enable("default","run:method");
```

**[0050]** In addition, the server spy 213 allows the client 105 to enable all facilities in a server 201, or all trace flags in a given facility, or both:

```
20     trace_enable("all", "flag_name");
     trace_enable("facility_name", "all");
     trace_enable("all", "all");
```

**[0051]** Similarly, the client 105 can disable all facilities previously enabled with:

```
25     trace_disable("all","all").
```

**[0052]** As noted above, separation of logging and tracing functionality provides for increased flexibility in the control of the server 201. Conventionally, the operating system will support a limited number of predefined log files. These log files will be used to record all system information, and not just information related to a specific server 201 or host computer 101. However, a system administrator may be interested only in a specific host computer 101 or server 201. Accordingly, the server spy 213 includes methods for controlling the printing for log files specific to any individual server 201 or host computer 101, for example, with the file\_start () method starting a new log file, specified by the user, and a complementary file\_stop () method for closing and saving a log file specified by the user. In the preferred embodiment of the server spy 213, a log file structure is maintained for each set of log files, storing the name of the log file, its open mode, the maximum file size in megabytes, and the number of swap files, if any. When the file\_start() method is invoked, the log file structure is passed to the method, the a new log file of the specified parameters in the log file structure is created. When the file\_stop() method is invoked, only the name of the log file need by passed, after which no further output is written to the log file.

**[0053]** To create a new log file, the server spy 213 calls the log API 221, passing a log file path name, which in turn creates a new log file, and places the log file on a global list stored within the server 201 so that all trace messages and error messages are output by the trace API 219 to the new log file. The designation of log files is read by the trace API 201 of the server 201 when an enabled trace facility in encountered during execution of the server 201, the trace API 219 automatically storing output value of the trace message in the new log file. The server spy 213 supports conventional exception handling for file manipulation.

**[0054]** In addition, because the user, such as a system administrator, may desire to log information in the standard system log files, the server spy 213 includes a flag, the boolean attribute use\_syslog, to control this option. When the flag is set, the trace API 219 and log API 221 will mirror trace output value, error information log information for the server 201 to the operating system's log files, otherwise this information will be stored only in the specified log files for the server 201.

**[0055]** A server 201 that has been operating for significant periods of time may create very large log files, whereas a server 201 that is infrequently used may have a very small log file. In order to accommodate this varying behavior while efficiently conserving secondary storage space, the server spy 213 provides a method, file\_set\_max\_bytes(), for changing the maximum size of log files for the server 201. The server spy 213 will call the log API 221 which then stores a maximum file size for a specified log file of the server 201. This limit value is used by the trace API 219 and log API 221 when information is output to the log files, as further described with respect to Figure 4.

**[0056]** Over time, most servers 201 will exceed the size of their log files. The server spy 213 provides a system

administrator the ability to specify a number of log files that can be swapped on a rotating basis. An example of the interface for this method is the file\_set\_swap\_count() method. Swapping allows a given log file to be examined, and either deleted or archived as needed, while still providing an active log file for logging the server 201. Figure 4 shows a flowchart for one embodiment of a method of the log API 221, executed by the processor 111 for setting and manipulating swappable log files.

**[0057]** To print to swap files the processor 111 prints 401 the output value of the trace message, or other log information to the designated log file, previously specified in the server 201. The processor 111 calls 403 an operating system function to obtain the current size of the log file. This size is compared 405 to the maximum file size as specified by a client 105 via the server spy 213. If the file size equals or exceeds the maximum, the processor 111 checks 407 whether the swap count, number of swap files specified for use in swapping, is greater than zero. If so, the processor 111 invokes 411 a swap files routine to physically swap the log files. Otherwise, the processor 111 simply resets 409 an end of file pointer, so that the log file is written over by new log or trace information.

**[0058]** Figure 5 illustrates one embodiment of swap file method for swapping log files. In this method, the oldest of a set of swap files is overwritten by successive renaming of newer swap files. Each log file is given a name in the form of filename.log.i, where *i* is an iterated value; the main, or active log file is simply filename.log. When invoked, the trace facility or logging function will loop 501 over the set of log files specified by the user, beginning with the value (swap\_count-2) and iterating to zero. From the iterated value of *i*, a new file name of the form filename.log.i is created 503. If a file of this name already exists 505, and it will where there is more than one log file, then the file is renamed from filename.log.i to filename.log.(i+1). This will overwrite each of the existing log files, beginning with the oldest one. Once all the old log files have been renamed in this fashion, the main log file is renamed 507 from filename.log to filename.log. 1.

**[0059]** The server spy 203 includes a shutdown method for invoking the shutdown API 223. This method allows a client 105 to gracefully terminate the server 201. In a multithreaded environment, there is considerable difficulty in handling signals received from the operating system, and determining the proper target for such signals. Conventionally, a UNIX® kill command would be used to shutdown a given server 201. However, because of the multithreaded environment, there is no way to ensure that the proper thread in the server 201 will receive the kill command, and thus the wrong thread may be terminated. The shutdown method the server spy 203 bypasses this problem by directly invoking the shutdown API 223. This allows greater control over the operation of the server 201 by system administrator during debugging and other administrative functions. In addition, this method allows a server 201 to hold itself down before the shutdown. Without such a hold down operation, a client 105 cannot be sure that a call will not come between a shutdown() invocation and a hold\_down() invocation, since a hold\_down() cannot be invoked first as the shutdown() invocation would be rejected, as would any client call.

APPENDIX A

**[0060]**

```

// File: ServerSpy.idl
//
// © 1994 Sun Microsystems, Inc.
//
5
#ifndef _SERVER_SPY_IDL
#define _SERVER_SPY_IDL

#pragma ident "@(#)ServerSpy.idl 1.19 95/01/16 Sun Microsystems"
10
#include <admin/DomfServerAdmin.idl>
#include <admin/DomfAdmin.idl>

// A single ServerSpy instance is embedded in every ODF server.
// It provides runtime information about the server as it executes
// and supports controlling tracing/logging configuration.
// It is provided to the domf as the Server Administrator for graceful
// shutdowns when the domf daemon stops.
//
15

module DatInt {

    //*****
    //
    // Structure and typedef declarations
    //
    //*****
25

    typedef unsigned long Seconds;
    typedef unsigned long ObjectCount;

    // Interface and server version structure
30
    struct Version {
        unsigned long major;
        unsigned long minor;
        string date;
    };

    // TraceFlag: used to control conditional output.
35
    struct TraceFlag {
        string name;
        string description;
        boolean enabled;
    };
    typedef sequence<TraceFlag> TraceFlags;

    // Facility: used to group sets of flags.
40
    struct Facility {
        string name;
        string description;
        TraceFlags flags;
    };
    typedef sequence<Facility> Facilities;

    exception UnknownFacility {};
    exception UnknownTraceFlag {
50
        string trace_flag; // Unknown flag
    };

    // LogFile: a single file being logged to
55
    struct LogFile {
        string name;

```

```

    string open_mode;
    float max_mbytes;
    long swap_count;
};
5 typedef sequence<LogFile> LogFiles;

exception FileError { // UNIX file error
    string details; // Explanation of error
};
10 exception AlreadyOpened {}; // Log file has already been opened.
exception SpecialFile {}; // Log file is not a regular UNIX file,
// and max_mbytes or swap_count > 0.
exception NotOpened {}; // Not currently logging to this file.

//*****
//
15 // ServerSpy interface
//
//*****

interface ServerSpy: DOMF_ServerAdmin, ObserverAdmin {
20
// Shutting down
void shut_down_hold(
    in BoaDbAdmin::HoldDownDuration hold_down
);

25 //
// Information about the running server
//

// Total number of active (invoked and unreaped) objects.
30 readonly attribute ObjectCount active_object_count;

// Number of active objects for the implementation.
struct ActiveObjectCount {
    string implementation_name;
    ObjectCount active_object_count;
35 };
typedef sequence<ActiveObjectCount> ActiveObjectCounts;
readonly attribute ActiveObjectCounts active_object_counts;

// List of all facilities provided by the server.
40 readonly attribute Facilities facilities;

//
// Inspecting and changing the behavior of the reaper
//

// Get/set seconds server must be idle before reaper will shut it down.
45 attribute Seconds server_timeout;

// Get/set how often the reaper wakes up and checks for inactivity.
attribute Seconds reaper_cycle_time;

// Get/set whether to preform automatic deactivation of objects and
50 // server itself.
attribute boolean reaping;

//
55 // Tracing control
//

```

EP 0 735 469 B1

```

// Enable one or more trace flags. Flags are space delimited
// keywords. Either facility or flag can be "all".
// Adds to the currently enabled set, does not
// override any current settings.
5 void trace_enable(
    in string facility_name,
    in string flags
) raises(UnknownFacility, UnknownTraceFlag);

10 // Disable one or more flags. To turn all tracing off, use "all"
// for both facility and flags.
void trace_disable(
    in string facility_name,
    in string flags
) raises(UnknownFacility, UnknownTraceFlag);

15 // Check whether flags for a facility are enabled.
boolean trace_is_enabled(
    in string facility_name,
    in string flags
) raises(UnknownFacility, UnknownTraceFlag);

20 //
// Logging output control
//

// Flag whether to send output to syslog daemon.
25 attribute boolean use_syslog;

// Information about log files being printed to.
readonly attribute LogFiles log_files;

30 // Start printing to a new log file.
void file_start(
    in LogFile new_file
) raises(FileError, AlreadyOpened, SpecialFile);

// Stop printing and close a log file.
35 void file_stop(
    in string name
) raises(FileError, NotOpened);

// Change the maximum size in Mb of a log file. Raises
// NotOpened if not currently logging to this file.
40 void file_set_max_mbytes(
    in string file_name,
    in float max_mbytes
) raises (FileError, NotOpened, SpecialFile);

// Change the maximum number of swapped out log files
// (file_name.1, .2, etc.). Raises NotOpened if not
// currently logging to this file.
45 void file_set_swap_count(
    in string file_name,
    in long swap_count
) raises (FileError, NotOpened, SpecialFile);

50 // This interface's version
readonly attribute Version spy_interface_version;

}; // ServerSpy interface
55 };

#endif // _SERVER_SPY_IDL

```

## Claims

1. A computer system (100) for administering servers and for providing a distributed object interface to a server executing as a server process, the system comprising:

a server (201) including at least one distributed object (209) as machine readable and executable structures, the server (201) executing as a server process and not separately existing as a first class object; at least one distributed first class object (213) within the server (201) including at least one machine executable structure adapted to execute as part of the server process and to determine current administrative state information about the server without interrupting execution of the server process, the first class object (213) having an interface available to objects external to the server (201) to enable manipulation of the server (201) as a distributed first class object for determining current administrative state information about the server.

2. The computer system (100) of claim 1, further comprising:

a machine executable structure adapted to determine a process identification code for the server process.

3. The computer system (100) of claim 1, wherein each distributed object (209) within the server (201) is an instance of an object implementation, and is either active, idle, or inactive, further comprising:

a first machine executable structure adapted to determine a number of distinct object implementations within the server (201); and,  
a second machine executable structure adapted to determine for each object implementation a number of active objects (209) of the object implementation.

4. The computer system (100) of claim 3 further comprising:

primary storage (107) for storing the server process;  
a processing unit (111) for executing the server (201) as the server process in the primary storage (107), the server process in the primary storage (107) including at least one object (209) executable by the processing unit (111) as part of the server process;  
a storage structure stored in the primary storage (107) and having a stored representation of the object implementations and object instances contained within the server (201);

wherein:

the processing unit (111) executing the server process generates the process identification code;  
the first machine executable structure reads the storage structure (107) to determine the number of distinct object implementations; and  
the second machine executable structure reads the storage structure (107) to determine the number of active objects (209).

5. The computer system (100) of one of claims 1 to 4, in which the least one first class object (213) further comprises:

a machine executable structure adapted to enable and disable automatic deallocation of resources held by idle objects (209) within the server (201);  
a machine executable structure adapted to establish a cycle time for the automatic deallocation of idle objects (209); and  
a machine executable structure adapted to establish a maximum time interval for which the server (201) containing the first class object (213) can be idle, such that the server (201) is automatically shutdown after a time interval equal to or exceeding the maximum time interval.

6. The computer system (100) of one of claims 1 to 4, in which the least one first class object (213) further comprises:

a machine readable structure adapted to identify at least one facility contained within the server (201), the facility including a set of trace flags;  
a trace enable machine executable structure adapted to selectively enable a user specifiable facility and set of trace flags within the facility; and,

a trace disable machine executable structure adapted to selectively disable a user specifiable facility and a set of trace flags within the facility.

7. The computer system (100) of claim 6, further comprising:

a machine executable structure adapted to indicate whether trace information output by the facilities within the server (201) is incorporated into an operating system message log; and  
a machine executable structure adapted to specify at least one log file for storing trace information output by facilities within the server (201).

8. A computer implemented method for administering servers and of embedding a distributed object (209) within a server (201) to determine current administrative state information about the server executing as a server process, comprising the steps of:

defining an object interface definition language file for a distributed first class object (213), the file specifying:

an interface to a first attribute or operation for obtaining a process identification number for a server process not separately existing as a first class object;  
an interface to a second attribute or operation for obtaining a host name for a computer (101) executing the server process;  
defining an implementation file specifying source code of the distributed first class object (213) for implementing the first and second attributes or operations, respectively;  
compiling at least one distributed object (209) for performing a function of the server (201), the object (209) incorporated into the server (201), to produce a server file;  
compiling the interface definition language file of the distributed first class object (213) to produce a set of source code files for interfacing the interface definition language file of the distributed first class object (213) to the implementation file of the distributed first class object (213);  
compiling the source code files and implementation file of the distributed first class object (213) to produce a first class object file;  
linking the first class object file to the compiled server file to produce an executable server file; and  
executing the server file to create a server process containing an instance of a first class object.

9. A computer implemented method for administering servers and for determining current administrative state information about a server (201) executing as a server process in a distributed object environment, the server (201) not separately existing as a first class object, the method comprising the steps of:

receiving in a first object (203) external to the server (201) a request from a client object to determine current administrative state information about the server (201);  
invoking from the first object (203) via an interface available to objects external to the server (201) to enable manipulation of the server as a distributed first class object for determining current administrative state information about the server an operation on a second object (213) being a first class object having said interface and being embedded in the server (201), the operation determining the current administrative state information about the server (201);  
executing the operation of the second object (213) as part of the server process without interrupting the server process to determine the current administrative state information; and  
providing the current administrative state information to the first object (203).

10. The method of claim 9, further comprising, before the step of receiving, the step of:

embedding in the server (201) at least one second object (213), each second object (213) having at least one operation for determining current administrative state information about the server.

11. The method of claim 9 or 10, wherein the operation of the second object (213) provides a process identifier for the server (201).

12. The method of one of claims 9 to 11, wherein the operation of the second object (213) provides a count of active objects (209) within the server (201).

13. The method of one of claims 9 to 12, wherein the operation of the second object (213) provides a count of object implementations within the server (201).
- 5 14. The method of one of claims 9 to 13, wherein the operation of the second object (213) provides a list of tracing facilities within the server (201).
- 15 15. The method of claim 14, wherein a second operation of the second object (213) enables at least one tracing facility within the server (201).
- 10 16. The method of claim 14, wherein the second operation of the second object (213) selectively enables at least one trace flag in a tracing facility.
17. The method of one of claims 9 to 16, wherein the operation of the second object (213) designates a log file for outputting trace and log information for conditions resulting from execution of the server (201).
- 15 18. The method of claim 17, wherein a second operation of the second object (213) establishes a number of log files to be used on a rotating basis.
- 20 19. The method of one of claims 9 to 18, wherein the operation of the second object (213) establishes a maximum time interval for which the server (201) can be idle, further comprising the step of:
- shutting down the server (201) after a time interval equal to or exceeding the maximum time interval.
- 25 20. The method of one of claims 9 to 19, wherein the operation of the second object (213) establishes a cycle time for automatic deallocation of resources held by idle objects (209) within the server (201).
21. The method of claim 20, wherein a second operation of the second object (213) enables automatic deallocation of resources held by idle objects (209) within the server (201).
- 30 22. A computer readable memory for use in a computer system (100) and for administering servers and for providing a distributed object interface to a server (201) executing as a server process, the memory comprising:
- a server (201) including at least one distributed object (209) as machine readable and executable structures, the server (201) executable by the computer system (100) as a server process and not separately instantiated as a first class object (213) in the system (100) upon execution;
- 35 style="padding-left: 40px;">at least one distributed first class object (213) within the server (201) including at least one method adapted to execute as part of the server process upon execution of the server (201), and further adapted to determine current administrative state information about the server without interrupting execution of the server process, the distributed first class object (213) providing an interface to objects (209) external to the server (201) to
- 40 style="padding-left: 40px;">enable manipulation of the server (201) as a distributed first class object (213) by such external objects (209) for determining current administrative state information about the server.

### Patentansprüche

- 45 1. Computersystem (100) zur Verwaltung von Servern und zur Bereitstellung einer verteilten Objektschnittstelle für einen Server, der Serverprozesse durchführt, wobei das System umfaßt:
- einen Server (201), der wenigstens ein verteiltes Objekt (209) als maschinenlesbare und ausführbare Struktur aufweist, wobei der Server (201) als Serverprozeß ausgeführt wird und nicht als getrenntes Basisklassenobjekt besteht;
- 50 style="padding-left: 40px;">zumindest ein verteiltes Basisklassenobjekt (213) innerhalb des Servers (201), das wenigstens eine maschinenausführbare Struktur aufweist, die eingerichtet ist, als Teil des Serverprozesses abzulaufen, um die aktuelle den Server betreffende administrative Statusinformation festzustellen, ohne die Ausführung des Serverprozesses zu unterbrechen, wobei das Basisklassenobjekt (213) eine Schnittstelle hat, die Objekten außerhalb
- 55 style="padding-left: 40px;">des Servers (201) zugänglich ist, um Manipulationen am Server als ein verteiltes Basisklassenobjekt zu ermöglichen, um die aktuelle administrative Statusinformation des Servers festzustellen.



2. Computersystem (100) nach Anspruch 1, weiterhin umfassend:

eine maschinenausführbare Struktur, die eingerichtet ist, einen Prozeßkennzeichnungscode für den Serverprozeß festzustellen.

3. Computersystem (100) nach Anspruch 1, wobei jedes verteilte Objekt (209) innerhalb des Servers (201) eine Instanz einer Objektimplementierung ist und entweder aktiv, im Leerlauf oder inaktiv ist und weiterhin umfaßt:

eine erste maschinenausführbare Struktur, die eingerichtet ist, eine Anzahl von verschiedenen Objektimplementierungen innerhalb des Servers (201) festzustellen; und  
eine zweite maschinenausführbare Struktur, die eingerichtet ist, für jede Objektimplementierung eine Anzahl von aktiven Objekten (209) der Objektimplementierung festzustellen.

4. Computersystem (100) nach Anspruch 3, weiterhin umfassend:

einen Primärspeicher (107) zum Speichern des Serverprozesses;  
eine Prozeßeinheit (111) zur Ausführung des Servers als der Serverprozeß in dem Primärspeicher (107), wobei der Serverprozeß in dem Primärspeicher (107) zumindest ein Objekt (209) umfaßt, das durch die Prozeßeinheit (111) als Teil des Serverprozesses ausführbar ist;  
eine in dem Primärspeicher (107) gespeicherte Speicherstruktur mit einer im Server (201) enthaltenen gespeicherten Darstellung der Objektimplementierungen und der Objektinstanzen,

worin:

die Prozeßeinheit (111), die den Serverprozeß ausführt, einen Prozeßkennzeichnungscode erzeugt;  
eine erste maschinenausführbare Struktur die Speicherstruktur (107) liest, um die Anzahl der verschiedenen Objektimplementierungen festzustellen; und  
die zweite maschinenausführbare Struktur die Speicherstruktur (107) liest, um die Anzahl der aktiven Objekte (209) festzustellen.

5. Computersystem (100) nach einem der Ansprüche 1 bis 4, in dem zumindest ein Basisklassenobjekt (213) im weiteren umfaßt:

eine maschinenausführbare Struktur, die eingerichtet ist, die automatische Freigabe von Ressourcen, die durch im Leerlauf befindliche Objekte innerhalb des Servers (201) belegt sind, ein- und auszuschalten;  
eine maschinenausführbare Struktur, die eingerichtet ist, eine Durchlaufzeit für die automatische Freigabe von im Leerlauf befindlichen Objekten (209) zu bilden; und  
eine maschinenausführbare Struktur, die eingerichtet ist, eine maximale Zeitdauer, während der sich der Basisklassenobjekte (213) enthaltende Server (201) im Leerlauf befinden kann, zu bilden, so daß der Server (201) nach einer Zeitdauer gleich oder größer der maximalen Zeitdauer automatisch heruntergefahren wird.

6. Computersystem nach einem der Ansprüche 1 bis 4, in dem zumindest ein Basisklassen-Objekt (213) im weiteren umfaßt:

eine maschinenlesbare Struktur, die eingerichtet ist, mindestens eine im Server (201) enthaltene Einrichtung zu bestimmen, wobei die Einrichtung ein Satz von Ablauf-Flags enthält;  
eine maschinenausführbare Struktur zur Ermöglichung der Ablaufverfolgung, die eingerichtet ist, wahlweise eine durch den Bediener bezeichnbare Einrichtung und einen Satz von Ablauf-Flags innerhalb der Einrichtung einzuschalten; und  
eine maschinenausführbare Struktur zur Deaktivierung der Ablaufverfolgung, die eingerichtet ist, wahlweise eine durch den Bediener bezeichnbare Einrichtung und einen Satz von Ablauf-Flags innerhalb der Einrichtung auszuschalten.

7. Computersystem (100) nach Anspruch 6, im weiteren umfassend:

eine maschinenausführbare Struktur, die eingerichtet ist, anzuzeigen, ob die Ausgabe der Ablauf-Flags durch Einrichtungen innerhalb des Servers (201) in ein Mitteilungsprotokoll eines Betriebssystems aufgenommen wird; und

eine maschinenausführbare Struktur, die eingerichtet ist, zumindest ein Mitteilungsprotokoll zum Speichern der Ausgabe der Ablauf-Flags durch die Einrichtung innerhalb des Servers (201) zu bestimmen.

- 5 8. In einem Computer implementiertes Verfahren zur Verwaltung von Servern und zur Einbettung eines verteilten Objekts (209) innerhalb eines Servers (201) zur Feststellung der aktuellen administrativen Statusinformation des Servers, der als Serverprozeß ausgeführt wird, mit folgenden Schritten:

10 Definieren einer Sprachdefinitionsdatei der Objektschnittstelle für ein verteiltes Basisklassenobjekt (213), wobei die Datei festlegt:

eine Schnittstelle zu einem ersten Attribut oder zu einer ersten Operation, um eine Prozeßkennzeichnungsnummer für ein Serverprozeß zu erhalten, der nicht getrennt als Basisklassenobjekt besteht; eine Schnittstelle zu einem zweiten Attribut oder einer zweiten Operation, um einen Hostnamen für ein Computer (101) zu erhalten, der den Serverprozeß ausführt;

15 Definieren einer Implementierungsdatei, die den Quellcode des verteilten Basisklassenobjektes (213) bestimmt, um jeweils die ersten und zweiten Attribute oder Operationen zu implementieren;

Kompilieren von zumindest einem verteilten Objekt (209), um eine Funktion des Servers (201) durchzuführen, wobei das Objekt (209) in dem Server (201) enthalten ist, um eine Serverdatei zu erstellen;

20 Kompilieren einer Sprachdefinitionsdatei der Objektschnittstelle des verteilten Basisklassenobjektes (213), um einen Satz von Quellcodedateien zu erzeugen, die die Sprachdefinitionsdatei der Objektschnittstelle des verteilten Basisklassenobjektes (213) mit der Implementierungsdatei des verteilten Basisklassenobjektes (213) koppelt;

Kompilieren der Quellcodedateien und Implementierungsdateien der verteilten Basisklassenobjekte (213), um eine Basisklassenobjektdatei zu erzeugen;

25 Linken der Basisklassenobjektdatei der kompilierten Serverdatei, um eine ausführbare Serverdatei zu erzeugen; und

Ausführen der Serverdatei, um einen Serverprozeß zu erstellen, der eine Instanz des Basisklassenobjektes enthält.

- 30 9. In einem Computer implementiertes Verfahren zur Verwaltung von Servern zur Feststellung der aktuellen administrativen Statusinformation eines Servers (201), der als Serverprozeß in einer Umgebung mit verteilten Objekten ausgeführt wird, wobei der Server (201) nicht getrennt als Basisklassenobjekt vorhanden ist, wobei das Verfahren folgende Schritte umfaßt:

35 Empfangen einer Anfrage innerhalb eines ersten Objekts (203) außerhalb des Servers (201), die von einem Klientenobjekt stammt, um die aktuelle administrative Statusinformation des Servers (201) festzustellen;

40 Aufrufen einer Operation von einem ersten Objekt aus (203), das ein zweites Objekt (213) betrifft, über eine Schnittstelle, die Objekten außerhalb des Servers (201) zur Verfügung steht, um die Manipulation der Servers als verteiltes Basisklassenobjekt zu gestatten, um die aktuellen administrativen Statusinformation des Servers festzustellen, wobei das zweite Objekt (213) ein Basisklassenobjekt mit der Schnittstelle und in den Server eingebettet ist, wobei die Operation die aktuelle administrative Statusinformation des Servers feststellt;

Ausführen der Operation des zweiten Objektes (213) als Teil des Serverprozesses, um die aktuelle administrative Statusinformation festzustellen, ohne den Serverprozeß zu unterbrechen; und  
Liefen der aktuellen administrativen Statusinformation an das erste Objekt (203).

- 45 10. Verfahren nach Anspruch 9 mit folgendem weiteren Schritt vor dem Schritt des Empfangens:

Einbetten von zumindest einem zweiten Objekt (213) in dem Server (201), wobei jedes zweite Objekt (213) zumindest eine Operation zur Feststellung der aktuellen administrativen Statusinformation des Servers hat.

- 50 11. Verfahren nach Anspruch 9 oder 10, wobei die Operation des zweiten Objekts (213) dem Server (201) eine Prozeßkennung liefert.

- 55 12. Verfahren nach einem der Ansprüche 9 bis 11, wobei die Operation des zweiten Objektes (213) eine Anzahl der aktiven Objekte (209) innerhalb des Servers (201) liefert.

13. Verfahren nach einem der Ansprüche 9 bis 12, wobei die Operation des zweiten Objektes (213) die Anzahl der Objektimplementierungen innerhalb des Servers (201) liefert.

## EP 0 735 469 B1

14. Verfahren nach einem der Ansprüche 9 bis 13, wobei die Operation des zweiten Objektes (213) dem Server (201) eine Liste von Einrichtungen zur Ablaufverfolgung zur Verfügung stellt.

5 15. Verfahren nach Anspruch 14, wobei eine zweite Operation des zweiten Objektes (213) zumindest eine Einrichtung zur Ablaufverfolgung innerhalb des Servers (201) ermöglicht.

16. Verfahren nach Anspruch 14, wobei die zweite Operation des zweiten Objektes (213) mindestens eine Kennzeichnung zur Ablaufverfolgung selektiv einschaltet.

10 17. Verfahren nach einem der Ansprüche 9 bis 16, wobei die Operation des zweiten Objektes (213) eine Protokolldatei festlegt, um Ablaufverfolgungs- und Protokollinformationen auszugeben, die Bedingungen betreffen, welche das Ergebnis des Ausführens des Servers (201) sind.

15 18. Verfahren nach Anspruch 17, worin eine zweite Operation eines zweiten Objektes (213) eine Anzahl von Protokolldateien erstellt, die in rotierender Weise zu verwenden sind.

19. Verfahren nach einem der Ansprüche 9 bis 18, wobei die Operation des zweiten Objektes (213) eine maximale Zeitdauer erstellt, während der der Server (201) sich im Leerlauf befinden kann, im weiteren den folgenden Schritt umfassend:

20 Herunterfahren des Servers (201) nach einer Zeitdauer, die gleich oder über der maximalen Zeitdauer ist.

25 20. Verfahren nach einem der Ansprüche 9 bis 19, wobei die Operation des zweiten Objektes (213) eine Durchlaufzeit zum automatischen Freigeben von Ressourcen erstellt, die durch sich im Leerlauf befindende Objekte (209) innerhalb des Servers (201) gebunden sind.

30 21. Methode nach Anspruch 20, wobei eine zweite Operation des zweiten Objektes (213) das automatische Freigeben von Ressourcen ermöglicht, die durch sich im Leerlauf befindende Objekte (209) innerhalb des Servers (201) gebunden sind.

22. Computerlesbarer Speicher für den Betrieb in einem Computersystem (100) und zum Verwalten von Servern zur Bereitstellung einer verteilten Objektschnittstelle für einen Server (201), der als Serverprozeß ausgeführt wird, wobei der Speicher umfaßt:

35 einen Server (201) einschließlich mindestens ein verteiltes Objekt (209) als maschinenlesbare und ausführbare Struktur, wobei der Server (201) durch das Computersystem (100) als Serverprozeß ausführbar ist und nicht getrennt als Basisklassenobjekt (213) in dem System (100) bei einer Ausführung instantiiert wurde; zumindest ein verteiltes Basisklassenobjekt (213) innerhalb des Servers (201) einschließlich zumindest ein Verfahren, das eingerichtet ist, bei Ausführung des Servers (201) als Teil des Serverprozesses abzulaufen, und im weiteren eingerichtet ist, die aktuelle administrative Statusinformation des Servers festzustellen, ohne den Ablauf des Serverprozesses zu unterbrechen, wobei das Basisklassenobjekt (213) eine Schnittstelle für Objekte (209), die außerhalb des Servers liegen, bereitstellt, um Manipulationen des Servers (201) als verteiltes Basisklassenobjekt (213) durch solche externe Objekte (209) zur Feststellung der aktuellen administrativen Statusinformationen zu ermöglichen.

### Revendications

50 1. Système informatique (100) pour administrer des serveurs et pour fournir une interface d'objet répartie à un serveur fonctionnant en tant que processus serveur, le système comprenant :

un serveur (201) comprenant au moins un objet réparti (209) en tant que structure lisible et exécutable par la machine, le serveur (201) fonctionnant comme un processus serveur et n'existant pas séparément en tant qu'objet de première classe ;

55 au moins un objet de première classe réparti (213) dans le serveur (201) comprenant au moins une structure exécutable par la machine adaptée à s'exécuter en tant que partie du processus serveur et à déterminer des informations d'état administratif courant sur le serveur sans interrompre l'exécution du processus serveur, l'objet de première classe (213) ayant une interface accessible à des objets externes au serveur (201) pour

permettre une manipulation du serveur (201) en tant qu'objet de première classe réparti pour déterminer des informations d'état administratif courant sur le serveur.

5 2. Système informatique (100) selon la revendication 1, comprenant en outre une structure exécutable par la machine adaptée à déterminer un code d'identification de processus pour le processus serveur.

3. Système informatique (100) selon la revendication 1, dans lequel chaque objet réparti (209) dans le serveur (201) est une instance d'une mise en oeuvre d'objet et est actif, au repos ou inactif, comprenant en outre :

10 une première structure exécutable par la machine adaptée à déterminer un nombre d'implémentations d'objets distinctes dans le serveur (201) ; et  
une seconde structure exécutable par la machine adaptée à déterminer pour chaque implémentation d'objets un nombre d'objets actifs (209) de l'implémentation d'objets.

15 4. système informatique (100) selon la revendication 3, comprenant en outre :

20 une mémoire primaire (107) pour mémoriser le processus serveur ;  
un module de traitement (111) pour exécuter le serveur (201) en tant que processus serveur dans la mémoire primaire (107), le processus serveur dans la mémoire primaire (107) comprenant au moins un objet (209) exécutable par le module de traitement (111) en tant que partie du processus serveur ;  
une structure de mémorisation mémorisée dans la mémoire primaire (107) et comportant une représentation mémorisée des implémentations d'objets et des instances d'objets contenues dans le serveur (201) ;

25 dans lequel :

30 le module de traitement (111) exécutant le processus serveur produit le code d'identification de processus ;  
la première structure exécutable par la machine lit la structure de mémoire (107) pour déterminer le nombre d'implémentations d'objets distinctes ; et  
la seconde structure exécutable par la machine lit la structure de mémorisation (107) pour déterminer le nombre d'objets actifs (209).

5. système informatique (100) selon l'une des revendications 1 à 4, dans lequel ledit au moins un objet de première classe (213) comprend en outre :

35 une structure exécutable par la machine adaptée à valider et invalider une désallocation automatique de ressources maintenue par les objets au repos (209) dans le serveur (201) ;  
une structure exécutable par la machine adaptée à établir un temps de cycle de la désallocation automatique des objets au repos (209) ; et  
40 une structure exécutable par la machine adaptée à établir un intervalle de temps maximum pendant lequel le serveur (201) contenant l'objet de première classe (213) peut être au repos de sorte que le serveur (201) est automatiquement coupé après un intervalle de temps égal ou supérieur à l'intervalle de temps maximum.

6. Système informatique (100) selon l'une des revendications 1 à 4, dans lequel ledit au moins un objet de première classe (213) comprend en outre :

45 une structure lisible par la machine adaptée à identifier au moins une fonction contenue dans le serveur (201), la fonction comprenant un ensemble de drapeaux de trace ;  
une structure exécutable par la machine de validation de trace adaptée à valider sélectivement une fonction spécifiable par l'utilisateur et un ensemble de drapeaux de trace dans la fonction ; et  
50 une structure exécutable par la machine d'invalidation de trace adaptée à invalider sélectivement une fonction spécifiable par l'utilisateur et un ensemble de drapeaux de trace dans la fonction.

7. Système informatique (100) selon la revendication 6, comprenant en outre :

55 une structure exécutable par la machine adaptée à indiquer si la sortie d'information de trace par les fonctions dans le serveur (201) est incorporée dans le journal de message de système d'exécution ; et  
une structure exécutable par la machine adaptée à spécifier au moins un fichier de journal pour mémoriser une sortie d'informations de trace par les fonctions dans le serveur (201).

8. Procédé mis en oeuvre par ordinateur pour administrer des serveurs et contenir un objet réparti (209) dans un serveur (201) pour déterminer des informations d'état administratif courant sur le serveur fonctionnant en tant que processus serveur, comprenant les étapes suivantes :

5 définir un fichier de langage de définition d'interface objet pour un objet de première classe réparti (213), le fichier spécifiant :

une interface vers un premier attribut ou opération pour obtenir un numéro d'identification de processus pour une porte de sélection n'existant pas séparément en tant qu'objet de première classe ;

10 une interface vers un second attribut ou opération pour obtenir un nom d'hôte pour un ordinateur (101) exécutant le processus serveur ;

définir un code source spécifiant un fichier d'implémentation de l'objet de première classe réparti (213) pour mettre en oeuvre les premier et second attributs ou opérations, respectivement ;

15 compiler au moins un objet réparti (209) pour réaliser une fonction du serveur (201), l'objet (209) étant incorporé dans le serveur (201), pour produire un fichier de serveur ;

compiler le fichier de langage de définition d'interface de l'objet de première classe (213) pour produire un ensemble de fichiers de code source pour servir d'interface avec des fichiers de langage de définition d'interface de l'objet de première classe réparti (213) vers le fichier d'implémentation de l'objet de première classe réparti (213) ;

20 compiler les fichiers de code source et le fichier d'implémentations de l'objet de première classe réparti (213) pour produire un fichier d'objet de première classe ;

lier le fichier d'objet de première classe au fichier de serveur compilé pour produire un fichier de serveur exécutable ; et

25 exécuter le fichier de serveur pour créer un processus serveur contenant une instance d'un objet de première classe.

9. Procédé mis en oeuvre par ordinateur pour administrer des serveurs pour déterminer des informations d'état administratif courant sur un serveur (201) fonctionnant en tant que processus serveur dans un environnement d'objet réparti, le serveur (201) n'existant pas en tant qu'objet de première classe, le procédé comprenant les étapes suivantes :

recevoir dans un premier objet (203) externe au serveur (201) une requête d'un objet client pour déterminer des informations d'état administratif courant sur le serveur (201) ;

35 invoquer à partir du premier objet (203), par l'intermédiaire d'une interface disponible aux objets externes au serveur (201) pour valider la manipulation du serveur en tant qu'objet de première classe réparti pour déterminer des informations d'état administratif courant sur le serveur, une opération sur un second objet (213) qui est un objet de premier classe ayant ladite interface et qui est logé dans le serveur (201), l'opération déterminant les informations d'état administratif courant sur le serveur (201) ;

40 exécuter l'opération du second objet (213) en tant que partie du processus serveur sans interrompre le processus serveur pour déterminer les informations d'état administratif courant ; et

fournir les informations d'état administratif courant au premier objet (203).

10. Procédé selon la revendication 9, comprenant en outre, avant l'étape de réception, l'étape consistant à loger dans le serveur (201) au moins un second objet (213), chaque second objet (213) ayant au moins une opération pour déterminer des informations d'état administratif courant sur le serveur.

11. Procédé selon la revendication 9 ou 10, dans lequel l'opération du second objet (213) fournit un identificateur de processus pour le serveur (201).

- 50 12. Procédé selon l'une des revendications 9 à 11, dans lequel l'opération du second objet (213) fournit un comptage d'objets actifs (209) dans le serveur (201).

13. Procédé selon l'une des revendications 9 à 12, dans lequel l'opération du second objet (213) fournit un comptage d'implémentations d'objets dans le serveur (201).

- 55 14. Procédé selon l'une des revendications 9 à 13, dans lequel l'opération du second objet (213) fournit une liste de fonctions de traçage dans le serveur (201).

15. Procédé selon la revendication 14, dans lequel une seconde opération du second objet (213) valide au moins une fonction de traçage dans le serveur (201).

5 16. Procédé selon la revendication 14, dans lequel la seconde opération du second objet (213) valide sélectivement au moins un drapeau de trace dans une fonction de traçage.

10 17. Procédé selon l'une des revendications 9 à 16, dans lequel l'opération du second objet (213) désigne un fichier de journal pour fournir des informations de trace et de journal sur les conditions résultant de l'exécution du serveur (201).

18. Procédé selon la revendication 17, dans lequel une seconde opération du second objet (213) établit plusieurs fichiers de journal destinés à être utilisés sur une base tournante.

15 19. Procédé selon l'une des revendications 9 à 18, dans lequel l'opération du second objet (213) établit un intervalle de temps maximum pendant lequel le serveur (201) peut être au repos, comprenant en outre l'étape consistant à couper le serveur (201) après un intervalle de temps égal ou supérieur à l'intervalle de temps maximum.

20 20. Procédé selon l'une des revendications 9 à 19, dans lequel l'opération du second objet (213) établit un temps de cycle pour la désallocation automatique des ressources maintenues par des objets au repos (209) dans le serveur (201).

21. Procédé selon la revendication 20, dans lequel une seconde opération du second objet (213) valide une désallocation automatique des ressources maintenues par des objets au repos (209) dans le serveur (201).

25 22. Mémoire lisible par ordinateur destinée à être utilisée dans un système informatique (100) et pour administrer des serveurs et pour fournir une interface d'objet répartie à un serveur (201) fonctionnant en tant que processus serveur, la mémoire comprenant :

30 un serveur (201) comprenant au moins un objet réparti (209) en tant que structures lisibles et exécutables par la machine, le serveur (201) étant exécutable par le système informatique (100) en tant que processus serveur et pas séparément instancié en tant qu'objet de première classe (213) dans le système (100) à la suite d'une exécution ;

35 au moins un objet de première classe réparti (213) dans le serveur (201) comprenant au moins un procédé adapté à s'exécuter en tant que partie du processus serveur par suite de l'exécution du serveur (201) et en outre adapté à déterminer des informations d'état administratif courant sur le serveur sans interrompre l'exécution du processus serveur, l'objet de première classe réparti (213) assurant une interface aux objets (209) externes au serveur (201) pour valider une manipulation du serveur (201) en tant qu'objet de première classe réparti (213) par de tels objets externes (209) pour déterminer des informations d'état administratif courant sur le serveur.




FIGURE 1




FIGURE 2



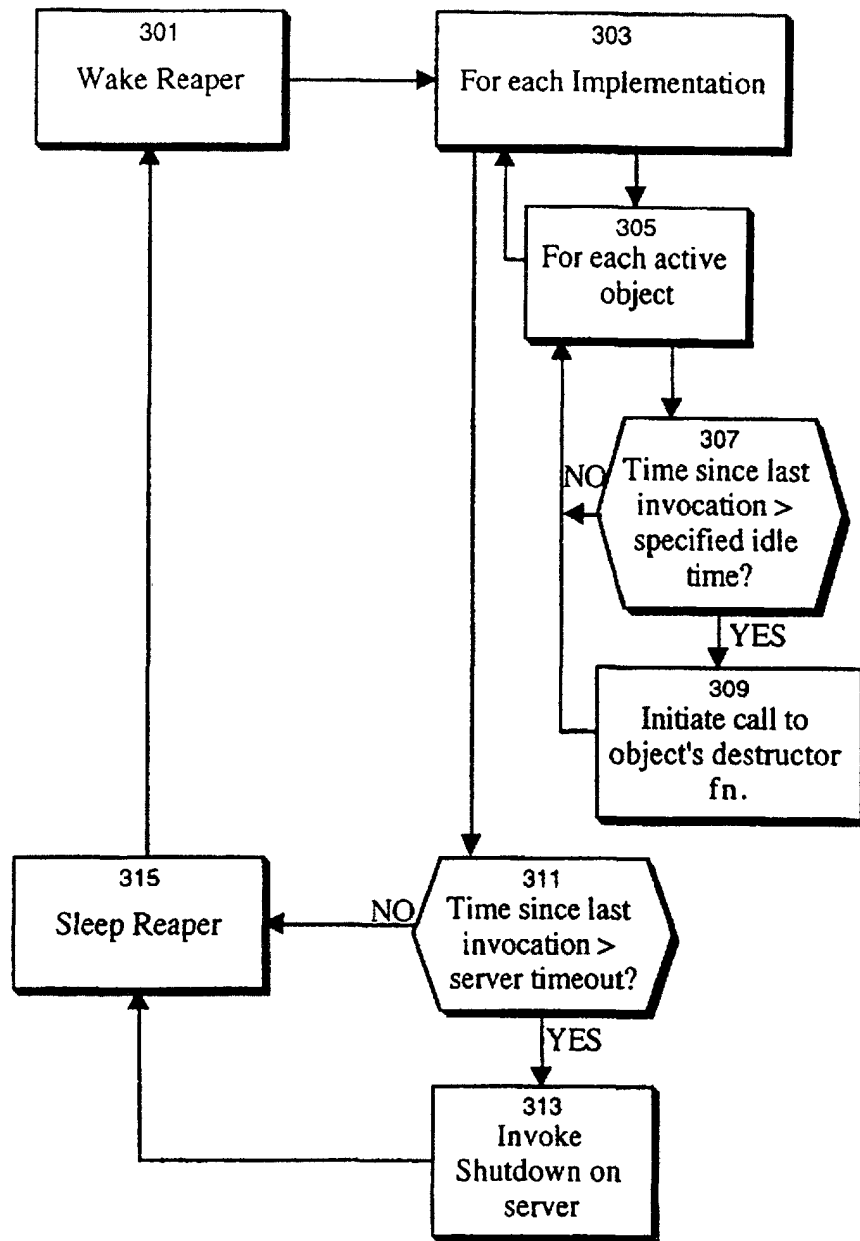


FIGURE 3

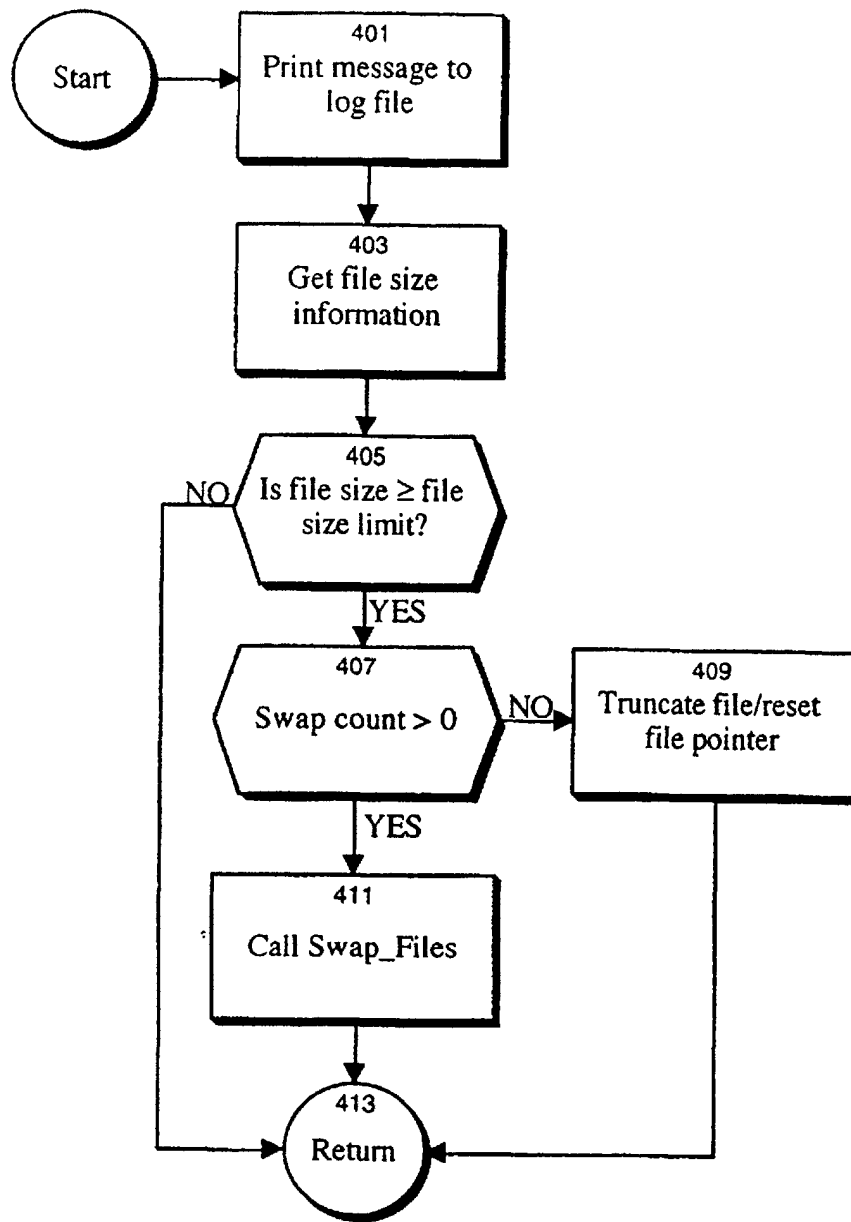


FIGURE 4

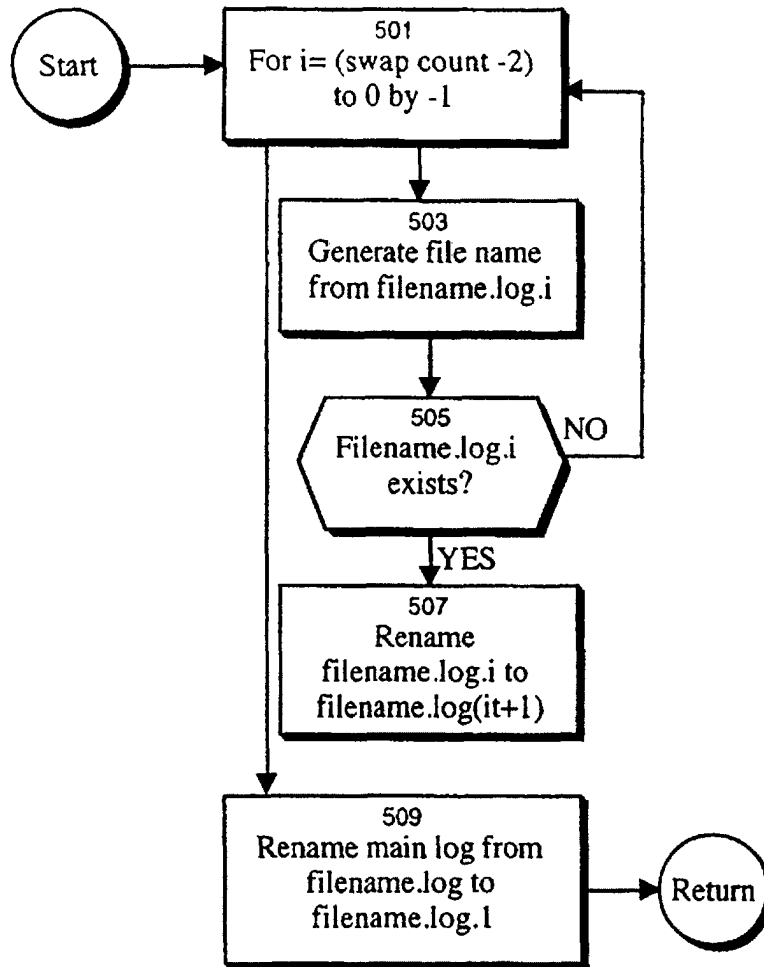
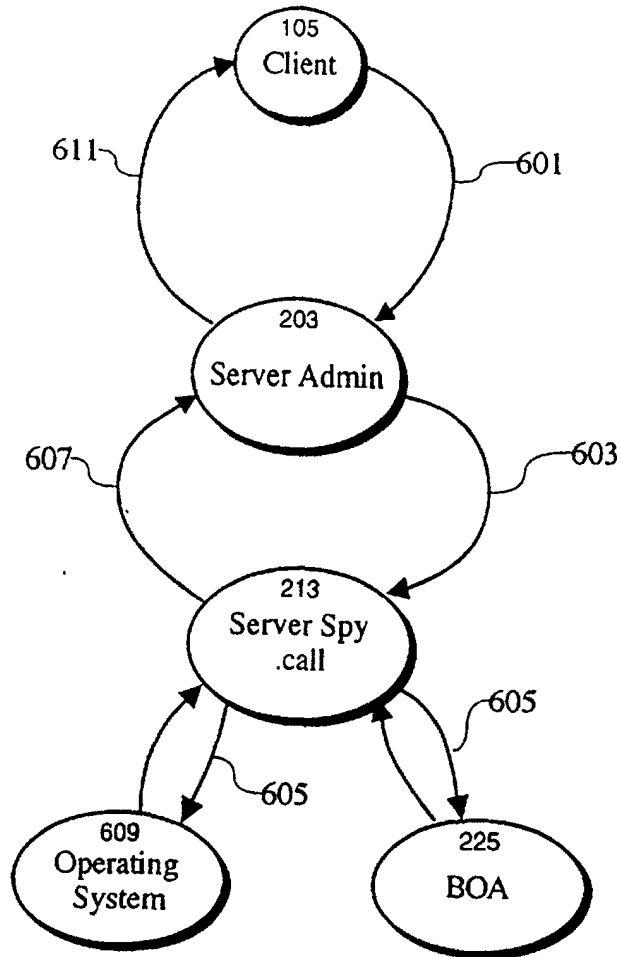


FIGURE 5



**FIGURE 6**