

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 April 2001 (12.04.2001)

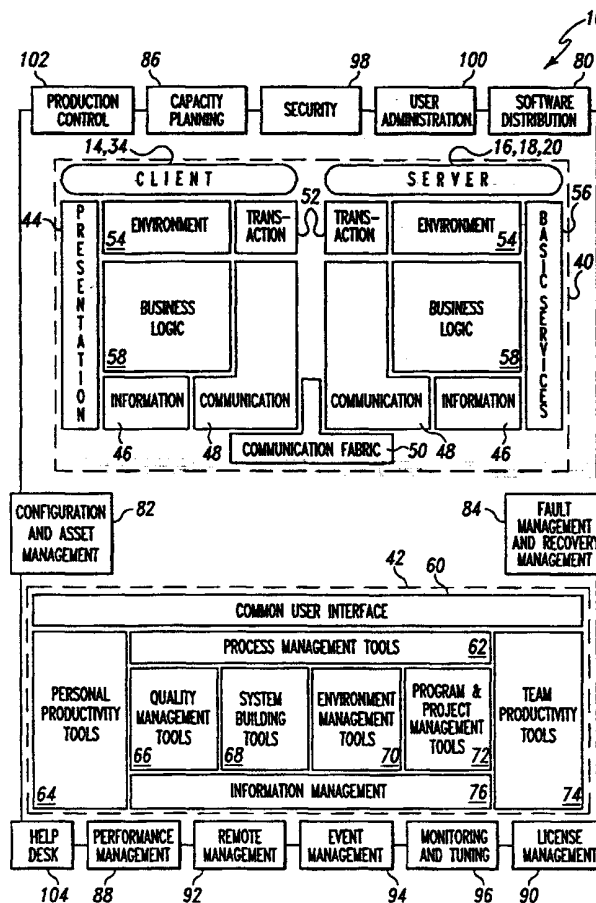
PCT

(10) International Publication Number
WO 01/25914 A2

- (51) International Patent Classification⁷: G06F 9/445
- (21) International Application Number: PCT/US00/26692
- (22) International Filing Date:
29 September 2000 (29.09.2000)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/156,962 1 October 1999 (01.10.1999) US
- (71) Applicant (for all designated States except US): ANDERSEN CONSULTING L.L.P. [US/US]; 100 South Wacker Drive, Chicago, IL 60603 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): KALTENMARK, John, K. [US/US]; 1202 Keim Trail, St. Charles, IL 60174 (US). DEGIORGIO, Christopher, M. [US/US]; 860 West Buckingham #2W, Chicago, IL 60657 (US).
- (74) Agent: MCCONNELL, Dean, E.; Brinks Hofer Gilson & Lione, One Indiana Square, Suite 2425, Indianapolis, IN 46204 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

[Continued on next page]

(54) Title: OPERATIONS ARCHITECTURES FOR NETCENTRIC COMPUTING SYSTEMS



[Continued on next page]



WO 01/25914 A2



(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

— *Without international search report and to be republished upon receipt of that report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(57) **Abstract:** An operations architecture for a netcentric computing system including a server connected with a client. The preferred operations architecture includes a software distribution tool, a configuration and asset management tool, a fault management and recovery management tool, a capacity planning tool, a performance management tool, a license management tool, a remote management tool, an event management tool, a systems monitoring and tuning tool, a security tool, a user administration tool, a production control application set and a help desk tool that support the server and the client in the netcentric computing system.

OPERATIONS ARCHITECTURES FOR NETCENTRIC COMPUTING SYSTEMS

This application claims the benefit under 35 U.S.C. §119(e) of U.S. provisional
5 application Serial No: 60/156,962 filed on October 1, 1999.

Field of the Invention

The present invention relates generally to business computing systems, and more
particularly, to an operations architecture for netcentric computing systems.

10

Background of the Invention

Computer-based business solutions have existed for various different types of
transactions since the mid-to-late 1960s. During this time period, the technology focused on the
use of batch technology. In batch processing, the business user would present a file of
15 transactions to the application. The computer system would then run through the transactions,
processing each one, essentially without user intervention. The system would provide reporting
at some point in the batch processing. Typically, the reports would be batch-printed, which, in
turn, would be used by the business user to correct the input transactions that were resubmitted
along with the next batch of transactions.

20

In the 1970s, businesses began a transition to on-line, interactive transactions. At a
conceptual level, this processing opened up the file of transactions found in batch transactions
and allowed the user to submit them one at a time, receiving either immediate confirmation of
the success of the transaction or else feedback on the nature of the transaction error. The
conceptually simple change of having the user interact with the computer on a transaction-at-a-
25 time basis caused huge changes in the nature of business computing. More important, users saw
huge changes in what they could do on a day-to-day basis. Customers were no longer forced to
wait for a batch run to process the particular application. In essence, the computer had an impact
on the entire work flow of the business user.

30

Along with the advent of on-line interactive systems, it was equally significant that the
systems provided a means for the business user to communicate with others in the business as the
day-to-day business went along. This capability was provided on the backbone of a wide area
network (WAN). The WAN was in itself a demanding technology during this time period and,

because of these demands, telecommunications groups emerged within organizations, charged with the responsibility to maintain, evolve and manage the network over a period of time.

The theme of the 1980s was database management systems (DBMSs). Organizations used and applied database technology in the 1970s, but in the 1980s, they grew more confident in the application of DBMS technology. Because of the advances in network technology, the focus was now on the sharing of data across organizational and application boundaries. Curiously, database technology did not change the fundamental way in which business processing was done. DBMS made it more convenient to access the data and to ensure that it could be updated while maintaining the integrity of the data.

In the 1990s, technology began to shift toward client/server computing. Client/server computing is a style of computing involving multiple processors, one of which is typically a workstation, and across which a single business transaction is completed. Using the workstation, the transaction entered by the user could now be processed on a keystroke-by-keystroke basis.

Furthermore, there was a change in the communications. With client/server, users could communicate with others in the work group via a local area network (LAN). The LAN permitted workstation-to-workstation communications at speeds of 100 to 1,000 times what was typically available on a WAN. The LAN was a technology that could be grown and evolved in a local office with little need for direct interaction from the telecommunications group.

During the late 1990s, the Internet began to receive widespread use by consumers and businesses. In the business world, the Internet has caused the concept of business users to expand greatly because of the way in which computers are now capable of being interconnected. In addition, the cost of computers has dropped to the point that it is affordable for almost every household to own a computer if so desired. As such, a need to expand the reach of computing both within and outside the enterprise, and that enables the sharing of data and content between individuals and applications has developed.

In the mainframe environment, operations tasks are performed by those in the data center who constantly watch monitor, and react to problems with the host or network. Keeping a mission-critical client/server application system available and under control, while providing a high level of service to the end user, is more complex and difficult than in a mainframe environment. Unfortunately, not all organizations are aware of this complexity, as they should be.

When client/server computing first emerged, organizations expected the cost and complexity of operations to be reduced because of reduced administration and because of common operating systems on workstations and servers. Time has shown that client/server environments tend instead to add rather than reduce complexity, therefore increasing operations costs.

More recently, netcentric computing has emerged as the next technology generation which will coexist with host and client/server environments. Again, while the initial hype around netcentric suggested that it would significantly simplify operations, experience is beginning to indicate that netcentric only adds an additional level of complexity through additional processes, tools, and support services, thus creating an environment even more potentially difficult and expensive to manage. The operations architecture now needs not only to keep an organization's internal production systems up and running, but also to maintain production systems that extend to business partners and customers.

The complexity and cost of operations architecture keeps increasing, which suggests a strong need for a structured and disciplined approach to implementation of tools and technologies to support eased operations.

Summary of the Invention

The present invention discloses an operations architecture, as well as a method of providing an operations architecture, for a netcentric computing system that includes a server connected with a client. The client may be a remote client or a client that is connected with the network of the netcentric computing system through a LAN connection or some other equivalent network connection. Preferentially, the client accesses all of the tools and resources of the netcentric computing system through a web browser application that interacts with the server of the netcentric computing system.

The preferred operations architecture includes a software distribution tool for providing automated delivery to, and installation of, an application on the server or the client. A configuration and asset management tool is also included in the operations architecture for managing a plurality of predetermined assets connected with said netcentric computing system. These assets could be servers, clients, printers, and various other computing devices that are connected with the netcentric computing system.

A fault management and recovery management tool is also provided in the preferred operations architecture for assisting in the diagnosis and correction of a plurality of system faults in said netcentric computing system. Those skilled in the art would recognize that several system faults may occur in the netcentric computing system and that the preferred fault management and recovery tool is able to deal with and correct these system faults.

In addition, the preferred operations architecture also includes a capacity planning tool that monitors a plurality of predetermined system usage levels in the netcentric computing system. The system usage levels may be selected from, but are not limited to, the group consisting of server processing usage, server bandwidth usage, server storage usage and client usage. A performance management tool is also included in the operations architecture for monitoring the performance of applications running on the netcentric computing system. A license management tool of the operations architecture manages and controls software license information for applications running on the netcentric computing system.

The preferred operations architecture also includes a remote management tool that allows support personnel from the netcentric computing system to take control of the client if required. This allows support personnel from to diagnose and repair problems with the client if they occur during operation. An event management tool of the operations architecture handles a plurality of predetermined events in the netcentric computing system. The predetermined events the event management tool are designed to handle include a broad category of events, including, but not limited to disk space indications, central processing unit utilization, database error indications, network error indications, application error indications and file and print service indications.

A systems monitoring and tuning tool is also provided by the preferred operations architecture for monitoring applications and computing devices connected with the netcentric computing system. The preferred monitoring and tuning tools are capable of monitoring applications, middleware, databases, networks, clients and servers and the tuning tools are capable tuning applications and dealing with network capacity issues. The preferred operations architecture also includes a security tool that includes a security application that provides security to the resources of the netcentric computing system. A user administration tool is also provided in the preferred operations architecture for administering users of the netcentric

computing system. Administering users includes such tasks as adding new users, deleting users, setting up access rights for users and removing access rights for users, to name just a few.

5 A production control application set for scheduling and handling a plurality of production processes on said netcentric computing system. In the preferred embodiment, the production control application set may be selected from the group consisting of a print management tool, a file transfer and control tool, a mass storage management tool, a backup and restore tool, an archiving tool and a system startup and recovery tool. A help desk tool is also part of the preferred operations architecture and includes a help application that provides
10 users of applications on the netcentric computing system with assistance during times of need.

Further objects and advantages of the present invention will be apparent from the following description, reference being made to the accompanying drawings wherein preferred embodiments of the present invention are clearly shown.

15 **Brief Description of the Drawings**

Figure 1 illustrates operations architecture for a netcentric computing system.

Figure 2 illustrates a representative netcentric computing system.

Figure 3 illustrates a preferred implementation of public key cryptography for the security tools of the operations architecture.

20

Detailed Description of the Presently Preferred Embodiments of the Invention

Referring to Figs. 1 and 2, the present invention discloses an operations architecture 10 for a netcentric computing system 12. The preferred netcentric computing system 12 includes at least one client 14 that is connected, via a network connection, with at least one server 16, 18,
25 20. The operations architecture 10 is located on the servers 16, 18, 20 and is used by the servers 16, 18, 20 during operation of the netcentric computing system 12, as set forth in greater detail below. Referring to Fig. 2, the physical picture of an illustrative netcentric computing system 12 is illustrated. In this example, a business enterprise 22 includes at least one client 14, at least one database server 16, at least one firewall 24, at least one application server 18, at least one web
30 server 20 and a local area network (LAN) connection 26, which are electrically connected as illustrated in Fig. 2.

As generally known in the art, LAN connections 26 are comprised of networking software applications and various computing devices (network cards, cables, hubs, routers, etc.) that are used to interconnect various computing devices (i.e. - clients 14 and servers 16, 18, 20) that are located at a first business enterprise location 28o form a network at that location. The term LAN connection 26 as used herein, should be broadly construed to include any and all hardware and software applications that allows clients 14, servers 16, 18, 20, or other computing devices to be electrically connected together to share and transfer data. Although not illustrated, other devices such as printers may be connected with the LAN connection 26 so that the resource is available to users of the network. Those skilled in the art would recognize that various types of LAN connections 26 exist and may be used in the present invention.

For the purpose of the present invention, the firewall 24 is used to isolate internal systems from unwanted intruders. In particular, firewalls 24 isolate web servers 20 from all Internet traffic that is not relevant to the netcentric computing system 12. In the preferred embodiment, the only requests allowed through the firewall 24 are for services located on the web servers 20. All requests for other applications (e.g., FTP, Telnet) and other IP addresses that the netcentric computing system 12 receives are typically blocked by the firewall 24 during operation of the netcentric computing system 12.

The web servers 20 are the primary interface to the clients 14 for all interactions with the applications or services that are provided by the netcentric computing system 12. The main task of the web servers 20 is to authenticate the clients 14, establish a secure connection from the clients 14 to the web servers 20 using encrypted messages, and allow applications the clients 14 are using to transparently access the resources of the netcentric computing system 12. The web servers 20 are responsible for accepting incoming HTTP (Hypertext Transfer Protocol) messages and fulfilling the requests. For dynamic HTML (Hypertext Markup Language) page generation, requests are forwarded to the application servers 18. During operation, static pages, such as help pages, are preferably generated entirely by the web servers 20.

In the preferred embodiment, the primary function of the application servers 18 is to provide a link through which the web servers 20 can interact with the clients 14, trigger business transactions, and send back resulting data to the clients 14. A fundamental role of the application servers 18 is to manage the logical flow of transactions and keep track of the state of sessions. The application servers 18 are also responsible for managing all sessions within the netcentric

computing system 12. A session is a period of time in which a client 14 is interacting with, and using, a resource of the netcentric computing system 12.

In the preferred embodiment of the present invention, the main purpose of the database servers 16 is to handle an application log. All requests sent to the web servers 20 and application servers 18, as well as their respective responses, are logged in the application log. The application log is preferentially used for traceability. In the preferred embodiment, requests are logged in the application log directly by the application server 18. Those skilled in the art would recognize that any number of data items can be monitored and kept track of in the application log.

As further illustrated in Fig. 2, a second business enterprise location 30 may be connected with the first business enterprise location 28 using an intranet connection 32. Those skilled in the art would recognize that various intranet connections 32 exist and may be used in the present invention. The intranet connection 32 allows the computing resources of the second business enterprise location 30 to be shared or connected with the computing resources available at the first business enterprise location 28. The term intranet connection 32, as used herein, should be broadly construed to include communication devices and software applications as well as various other connection devices used to physically interconnect two or more business networks. Although not illustrated, several other enterprise locations, each containing its own computing resources, may be connected with the netcentric computing system 12 using other intranet connections 32.

In the preferred embodiment illustrated in Fig. 2, the firewall 24 of the first business enterprise location 28 is connected with an Internet connection 32 to a plurality of remote clients 34. The remote clients 34 that are connected to the Internet connection 32 preferentially access data and communicate with the services of the netcentric computing system 12 through the Internet connection 32 using web browser applications that are located and running on the clients 34. The Internet connection 32 gives the remote clients 34 the ability to gain access to applications, information and data content that may be located on the database server 16, the application server 18 and the web server 20, preferably by means of the web server 20.

As used herein, the term Internet connection 32 should be broadly construed to include any software application and hardware device that is used to connect the clients 34 and the servers 16, 18, 20 with an Internet service provider (not illustrated) that establishes the

connection to the Internet. Those skilled in the art would recognize that the clients 34 and the servers 16, 18, 20 may establish the Internet connection 32 with the Internet service provider using modems, cable modems, ISDN connections and devices, DSL connections and devices, fiber optic connections and devices, and satellite connections and devices to name a few. For the purpose of the present invention, it is important to understand that the remote clients 34 and servers 16, 18, 20 are connected with one another through the Internet connection 32.

For a detailed discussion of the elements of the technical architecture of the preferred netcentric computing system 12, as well as netcentric computing systems 12, refer to co-pending U.S. patent application Serial Number _____ entitled Architectures For Netcentric Computing Systems, which was filed on September 29, 2000 and is hereby incorporated by reference in its entirety.

Referring to FIG. 1, the operations architecture 10 includes the tools and support services required to keep a production system up and running well in a netcentric computing system 12. The preferred operations architecture 10 supports a netcentric execution architecture 40 and a development architecture 42 that are located on the netcentric computing system 12. It differs from the execution architecture 40 and the development architecture 42 in that the operations architecture's 10 primary users are systems administrators and production support personnel. In addition, it differs from the operations infrastructure in that the operations infrastructure represents operations processes and organization as well as the technologies and tools.

The netcentric execution architecture 40 includes common run-time services that are required when an application executes on the netcentric computing system 12. The preferred netcentric execution architecture includes presentation services 44, information services 46, communication services 48, communication fabric services 50, transaction services 52, environment services 54, base services 56, and business logic services 58. For a detailed discussion of each of the services set forth above, as well as the role they play in the preferred netcentric computing system 12, refer to co-pending U.S. Patent Application Serial No. _____ entitled ARCHITECTURES FOR NETCENTRIC COMPUTING SYSTEMS.

The development architecture 42 is used to create the applications that perform business processes in the netcentric execution architecture 40 of the netcentric computing system 12. The purpose of the development architecture 34 is to support the tasks involved in the analysis, design, construction and maintenance of business systems and applications, as well as the

associated management processes of the netcentric computing system 12. The preferred development architecture 42 includes common user interface tools 60, process management tools 62, personal productivity tools 64, quality management tools 66, system building tools 68, environment management tools 70, program and project management tools 72, team productivity tools 74, and information management tools 76. For a detailed discussion of the preferred development architecture 42, refer to co-pending U.S. Patent Application Serial No. _____, entitled DEVELOPMENT ARCHITECTURES FOR NETCENTRIC COMPUTING SYSTEMS, which is incorporated herein by reference in its entirety.

As illustrated in FIG. 1, the preferred operations architecture includes a wide variety of tool categories. Tool categories cover the spectrum of functions provided by the operations architecture 10, which range from software distribution tools to help desk tools, as set forth in more detail below. The preferred operations tools in the operations architecture 10 include a software distribution tool 80, a configuration and asset management tool 82, a fault management and recovery management tool 84, a capacity planning tool 86, a performance management tool 88, a license management tool 90, a remote management tool 92, a event management tool 94, a systems monitoring and tuning tool 96, a security tool 98, a user administration tool 100, a production control application set 102 and a help desk tool 104.

The tools set forth above provide visibility and control over the events of a distributed environment, which is how netcentric computing system 12 operates. In essence, they can be viewed as the support and control mechanisms for both the netcentric execution architecture 40 and the development architectures 42. This relationship is illustrated in FIG. 1, where the major categories of operations tools are depicted as supporting the netcentric execution architectures 40 and the development architectures 42.

In the preferred embodiment, the software distribution tools 80 provide automated delivery to, and installation of, applications and systems software on the servers 16, 18, 20 and end user devices (e.g., clients 14, 34, kiosks, etc.). This can be for an organization's internal computing environment, as well as for its extended one, i.e., its business partners and customers. The architectural support required to support the operations architecture 10 software distribution is largely driven by the number of clients 14, 34, servers 16, 18, 20, and the geographic locations to be served.

When it is unrealistic to use a manual approach to software distribution, an organization should consider adding automated software distribution tools 80 to the operations architecture 10. Many products from leading vendors such as Microsoft, Tivoli, and Hewlett-Packard are on the market today that include or specialize in automated software distribution. Systems developers must look for several important features, depending on the specific support requirements of the netcentric computing system 12.

The server 16, 18, 20 component of a preferred software distribution tool 80 enables administrators to build distribution packages and to control distribution amongst the netcentric computing system 12. A distribution is a package of related software files, data, and installation scripts that form an installable unit. Few significant application installations, systems software installations, or even upgrades can be achieved simply by sending a single file. Configuration files (e.g., config.sys) and system files (e.g., autoexec.bat, login as well as multiple software files for a particular application or systems software component, often require changes. In addition, it is usually desirable to upgrade multiple applications or combinations of systems software and applications in a single distribution rather than performing multiple independent software distributions. Bundling software upgrades together also reduces the amount of release testing required.

A distribution is created by selecting the files and scripts, often through a point-and-click interface on the clients 14, 34 or servers 16, 18, 20, depending on what system is being upgraded. The components are then combined into a single file for transmission to the respective devices. Some software distribution tools 80 might provide compression capabilities to reduce the physical size of the distribution. This is particularly important in a WAN environment where line speeds are an issue.

There are multiple approaches to scheduling software distributions. Some solutions use a rigid scheduling mechanism that requires all target machines (i.e., clients 14, 34 or servers 16, 18, 20) to be powered on at a specified time when the software distribution is to occur. This mechanism could be characterized as a "push" strategy, where the server 16, 18, 20 performing the software distribution pushes the application to the clients 14, 34 or servers 16, 18, 20 at a specified time.

A more flexible approach is a pull strategy, where the clients 14, 34 or servers 16, 18, 20 check for software updates and pull the software from the designated servers 16, 18, 20 at log-in time. Thus, when the user signs on either in the morning or at some point during the day, any pending updates are downloaded to the respective clients 14, 34 or servers 16, 18, 20. When
5 combined with a forced log-off capability, which most networks support, this can effectively mimic the push strategy without the attending problem of some machines being powered off.

Neither the push nor pull scheduling approach is sufficient when large numbers of target clients 14, 34 are involved. For example, a sales office automation system developed several years ago and used by 1,400 salespeople distributed across scores of locations encountered a
10 problem with these strategies on its first major software upgrade. The sales office used the pull strategy because it was not feasible to have all workstations, locations, and dialup users connected and powered up at the same time. The distribution was scheduled to be available when the users logged in on Monday morning. This was a substantial functional upgrade to the system, so the software distribution was several megabytes in size. The problem was that 1,400
15 machines could not simultaneously download one copy of software off of a server. As a result, most users were unable to retrieve the new software or use the system for several days.

Faced with the problem of scale, two alternatives can be used when performing a software distribution on a large scale. One is simply to acquire more servers 16, 18, 20 with more copies of the software to be distributed. Of course, this is an expensive solution,
20 particularly when these machines are not needed for any other purpose. The preferred solution provided by the software distribution tools 80 involves staging software distribution. Software distribution staging works by sending a new version of the software in advance of the cut-over date. In effect, the clients 14, 34 or servers 16, 18, 20 have two versions of the application physically resident simultaneously, but only one is in use. The existing software is used until the
25 present cut-over date is reached. At that time, the client 14, 34 portion of the software distribution tool 80 automatically completes the installation and redirects the user to the new version. Using this approach, it is possible to selectively download the software update to subsets of machines well in advance of the cut-over date, thus eliminating the bottleneck.

An enhancement of staging is the ability to cut over to the new version on the receipt of a
30 small command file rather than a preset date. This gives operations more flexibility to alter the cut-over date due to unanticipated events. For example, many adopters fail to anticipate the

requirements of having multiple copies of applications stored simultaneously when determining the size of the workstation hard disks required for the users.

Remote Installation

Most software distribution tools 80 used in the operations architecture 10 include a client 5 14, 34 portion as well as a server 16, 18, 20 portion that resides on the target machine. The client 14, 34 software is responsible for installation of the software distribution onto the target machine's hard disk. The first step in the installation process is the unbundling (and uncompressing) of the software distribution into the component files, data sets, and scripts (although the better products will first check to see that the required disk space is in fact 10 available). Next, any pre-installation scripts are executed. These scripts may do such various tasks as checking for required components or adding or modifying lines in the target machine configuration or systems files that will be required by the new software (e.g., changing the number of buffers or adding a line to install a necessary driver at startup time). The directories in which the software is to reside are checked or created, and then the actual software files are 15 moved into the proper location on the hard disk. At this point a post-installation script may be invoked that could include rebooting the machine so that the changes to the system and configuration files can take effect.

In large netcentric computing system 12, where tens or even hundreds of servers 16, 18, 20 support individual groups of clients 14, 34, a cascaded software distribution approach may be 20 used. A cascaded software distribution allows for a central administrator to schedule the distribution of software updates to designated servers 16, 18, 20 within the netcentric computing system 12. These servers 16, 18, 20 in turn, distribute the software updates to their associated clients 14, 34. This approach allows the simple push and pull strategies to be used for larger numbers of clients 14, 34 without requiring staging. It also better utilizes the servers 16, 18, 20 25 and communications links in these larger environments.

Another aspect of the software distribution tool 80 is that it supports error handling reporting. When dealing with larger networks of clients 14, 34, errors inevitably occur in the software distribution process. There may be insufficient disk space or a required component may be missing. The preferred software distribution tool 80 is capable of reporting errors and 30 taking appropriate actions. Error reporting normally takes the form of a distribution log file that records success, failure, or errors encountered. In some cases a more active form of error

reporting may be required, where email messages may be automatically generated and sent to either the administrator or, in some cases, the affected clients 14, 34. If a fatal error is detected, the software distribution tool 80 will reverse any changes made to that point and restore the clients' 14, 34 machines to their previous state.

5 As illustrated in FIG. 1, the preferred operations architecture 10 includes configuration and asset management tools 82. To manage a netcentric computing system 12, one must have a solid understanding of what is located where, and one must maintain rigor in the change control procedures that govern modifications to the netcentric computing system 12. Configuration and asset management information that the configuration and asset management tools track includes
10 such details as product licensing information, warranty information, vendor names, logical and physical device information (such as total capacity and current utilization), product configuration tracking, software and data version levels, network configuration parameters, physical location, and perhaps accounting information.

In larger netcentric computing systems 12 an underlying configuration and asset
15 management database is used to keep track of configuration and asset information in the netcentric computing system 12. This database becomes a key information source for those managing, maintaining, and adding to the netcentric computing system 12. Automatic asset and configuration collection capability is included in many vendor solutions, including OpenView from Hewlett-Packard (HP), and POLYCENTER Systems Census from Digital Equipment Corp.
20 These products can be used to interrogate the netcentric computing system 12 and discover network and computing devices, and collect related information about these devices. In addition, these products can perform the needed periodic auditing to detect changes to the environment over time - for example, when a client 14, 34 moves a machine or installs a network game. Those skilled in the art would recognize that various configuration and asset information may be
25 collected and stored about the resources using the netcentric computing system 12.

Another important and related feature that is provided by the configuration and asset management tools 82 is the ability to restore a machine (i.e., clients 14, 34 or servers 16, 18, 20) to a known or initial configuration for problem resolution. The configuration and asset management tools 82 provide facilities for determining the correct initial state for a given
30 machine or network device and initiates any software distribution or configuration changes needed to bring the device back within compliance. For more dynamic netcentric computing

systems 12, where machine and network configurations are changing frequently, it is even more important to have an active configuration and asset management system. The capability to automatically change configurations of a large number of machines and network components or even to roll back to previous configuration settings for any particular device becomes
5 increasingly important and is provided by the preferred configuration and asset management tools 82.

A well-functioning configuration and asset management tool 82 becomes a vital information source for conducting impact analysis for any requested changes to the netcentric computing system 12. The frequency with which unexpected negative side effects are caused by
10 relatively minor configuration changes to the netcentric computing system 12 has been an embarrassing and frustrating surprise for many adopters of the technology.

Much of the source of these problems relates to the high number of execution architecture components and complex interdependencies between them. Another problem is the reality that most netcentric computing systems 12 involve numerous independent vendors. Changing even
15 the release level of one systems software component may have a ripple effect and may require updates to, or newer versions of, additional software components or applications.

To support this type of impact analysis, dependency information is maintained by the configuration and asset management tools 82. For example, version X of the Oracle database management system requires version Y or greater of the HP-UX operating system and version Z
20 of yet another vendor's Transmission Control Protocol/ Internet Protocol product to function properly. It is not uncommon for a user organization to wish to return to a previous operating system release to acquire an application package that does not yet support the latest operating system version. The configuration and asset management tools 82 maintain relationship
25 information so that it is not purely guesswork if in fact the proposed version change will break any required dependencies of the netcentric computing system 12.

The configuration and asset management tools 82 also enforce the appropriate degree of standardization across network environments in the netcentric computing system 12. For large netcentric computing systems 12, where thousands of clients 14, 34 are involved, it is not
feasible to effectively manage the environment if each client 14, 34 has its own unique
30 configuration and combination of software products. On the other hand, it is not typically appropriate to give thousands of users the exact same configuration if the users perform different

functions within the organization. For example, users in such diverse areas as sales, product development, and human resources are likely to require different computing capabilities. The goal is to strike the correct balance between standardization, which simplifies the required operations architecture and tasks, and accommodation to each business area's unique computing needs.

Referring to FIG. 1, the preferred operations architecture 10 includes fault management and recovery management tools 84. Failure control is important in a netcentric computing system 12. The presence of heterogeneous equipment, however, makes it difficult to determine the origins of a fault. Multiple messages may be generated within the system from a single fault, making it difficult to separate the fault's cause from its effects.

The fault management services and recovery management tools 84 of the operations architecture 10 assist in the diagnosis and correction of system faults in the netcentric computing system 12. Faults may include network-, server- 16, 18, 20, client- 14, 34, or even application-level faults. Fault diagnosis requires services for isolation; viewing of host, server 16, 18, 20, and client 14, 34 error logs; and determining the software and data versions and configurations of affected machines.

The fault management and recovery management tools 84 also include network management and diagnostic tools for monitoring and reporting on network traffic and failures on the netcentric computing system 12. Additional diagnostic tools such as protocol analyzers may also be included to determine the true source of the problem.

Another factor to consider in this selection is the choice between integrated operations environments (typified by HP's Open View or CA-Unicenter TNG), and point solutions that provide only one function. Although most integrated tool sets today do not adequately address the full breadth of fault management and diagnostic requirements, they can reduce the number of vendors and the complexity of integrating these point solutions.

Recovery capabilities are also included in the fault management and recovery management tools 84. Recovery capabilities span the range from those required to bring up a client 14, 34 or server 16, 18, 20 after it has failed to those required in the event of a major disaster. With critical business applications being rolled out on distributed technologies, the recovery of these systems must be easy, quick, and efficient. Loss of the netcentric computing system 12 for even a short period can result in significant financial losses to the business.

A wide variety of recovery tools may be required for fault recovery. These range from strictly network-oriented components (for restoring links or reconfiguring components) to more systems-level components (for restarting processes on machines or restoring databases). More involved tasks, such as the distribution of software fixes to clients 14, 34 or servers 16, 18, 20, may require the ability to remotely reboot and reinitialize machines, printers, or other network components. Those skilled in the art would recognize that the application of the preferred fault management and recovery management tools 84 will vary, depending on the needs and requirements placed on the netcentric computing system 12.

As illustrated in FIG. 1, the preferred operations architecture 10 includes capacity planning tools 86. The capacity planning tools 86 focus on individual components of an environment such as the network, physical space, and processing power to understand the need to change the capacity of those components based on organizational changes. The capacity planning tools 86 preferentially include applications that monitor a plurality of predetermined system usage levels in the netcentric computing system 12.

The system usage levels are preferentially selected from the group consisting of server processing usage, server bandwidth usage, server storage usage and client usage. The server processing usage information relates to the processing power being consumed by servers 16, 18, 20 during operation. The server bandwidth usage information relates to the amount of Internet traffic that is taking place over the Internet connection 32 with the servers 16, 18, 20 of the netcentric computing system 12. The server storage information relates to the amount of drive space available and being used on the servers 16, 18, 20 of the netcentric computing system 12. The client usage information relates to various items that can be stored about a respective client's 14, 34 session with the netcentric computing system 12.

The capacity planning tools 86 typically focus on components that are considered to be heavily sensitive to changes in computing resource usage. The preferred capacity planning tools 86 may use historical management data combined with estimates for growth or changes to configuration to simulate the ability of different system configurations to meet capacity needs. Capacity planning tools 86 can sometimes be integrated into a larger integration platform, or they can be standalone applications, depending on the needs and requirements of the netcentric computing system 12.

As previously set forth, referring to FIG. 1, the preferred operations architecture 10 includes performance management tools 88. The performance management tools 88 include applications that monitor the performance of computing resources and netcentric applications that are running on the netcentric computing system 12. Performance tuning issues are no longer confined to the network or to central processing units in netcentric computing systems 12. Performance tuning needs to be viewed in an end-to-end manner, accounting for all the factors that affect the performance of the netcentric computing system 12 relative to a user request from a respective client 14, 34. Those skilled in the art would recognize that the particular design of applications used in the performance management tools 88 will vary depending on the needs and requirements of the netcentric computing system 12.

The creation of a customer order, for instance, may involve multiple server 16, 18, 20 accesses for data and information to be exchanged between the client 14, 34 and the host server 16, 18, 20. The performance relative to the entire business event needs to be considered, not simply the performance of a single component involved. As such, the performance management tools 88 include applications that surround processes occurring on the netcentric computing system 12 that monitor the performance of devices (i.e., clients 14, 34; server 16, 18, 20) to calculate and provide end-to-end performance information.

The preferred operations architecture 10 for the netcentric computing system 12 also includes license management tools 90. The license management tools 90 include applications that focus on guaranteeing compliance with software license agreements for various vendor applications that are used on the netcentric computing system 12. Since the advent of computer networks that allow applications to be shipped and installed around the network as required, the issue of license management has become increasingly important. Application vendors have been experimenting with various licensing strategies, including unrestricted site licenses, fixed concurrent user licenses, and floating licenses that actually enforce the restriction on concurrent users.

Independent of these actions by software vendors, large organizations struggle to keep a handle on exactly what software products they own and how many copies they own. They have also been working to ensure that they are in compliance with software licensing agreements while not paying for more copies of software than they truly require. As such, the licensing

management tools 90 allow administrators to monitor and track applications that have licensing requirements to ensure compliance with the terms of each respective agreement.

In addition to guaranteeing compliance with software licensing agreements, the preferred license management tools 90 are capable of providing license report detailing which clients 14, 34 and how many clients 14, 34 are actually using a given software application. If, in fact, the license report indicates that the organization has over-purchased, it may be possible to realize some savings by reducing software licensing agreements, or vice versa. Those skilled in the art would recognize that several applications may be provided in the license management tools 90 to ensure license agreement compliance, depending on the particular applications provided in the netcentric computing system 12.

As distributed environments allow users more flexibility in terms of where they work, the ability of a centralized support group to effectively manage remote clients 34 has become important. Visibility to the configuration of a respective remote client 34 is only possible by physically sitting at the workstation and diagnosing problems or by accomplishing the same remotely.

As illustrated in FIG. 1, the preferred operations architecture 10 also includes remote management tools 92. The remote management tools 92 allow support personnel to "control" a user's desktop over the netcentric computing system 12 so that the support personnel do not need to be physically present at a particular client 14, 34 to diagnose problems. Once control of the desktop on the client 14, 34 is established by the remote management tools 92, screen updates for the controlled desktop are displayed at both locations. The support person will typically be located at another client 14, 34 that is connected with the netcentric computing system 12. The support person is then effectively sitting at the client 14, 34 he/she controls and can do necessary diagnostics.

In addition to problem diagnosis, the remote management tools 92 provide visual explanations to user questions. For example, if a user has a question about a certain application feature, the support person may remotely control the user's desktop, then walk through the solution while actions are displayed on the screen of the client 14, 34.

The preferred remote management tools 92 are also useful in organizations where 24x7 support is required. Rather than requiring support personnel to be physically present for all events, they may be able to dial in through the remote management tools 92 from home and

accomplish the same tasks. The ability to perform these tasks remotely can have positive effects on overall support costs through a reduction in the amount of time needed to resolve problems. Remote management tools may come bundled with an integration platform such as HP Open View or Tivoli TME, or they may be purchased as third-party software packages or designed specifically for the netcentric computing system 12.

The preferred operations architecture 10 also includes event management tools 94. The event management tools 94 include applications that manage a plurality of predetermined events generated by applications or devices on the netcentric computing system 12. The predetermined events may relate to disk space indications, central processing unit utilization indications, database error indications, network error indications, application error indications and file and print service indications. Those skilled in the art would recognize that other predetermined events could be monitored depending on the needs of the respective netcentric computing system 12.

In addition to hardware devices, applications and systems software generate events on the netcentric computing system 12. Common event-handling applications are used to provide information to management in a simple, consistent format and to forward on important events for management purposes. Those skilled in the art would recognize that events the event management tools 94 are design to monitor will vary. The applications that are designed for the preferred event management tools 94 preferentially monitor a plurality of predetermined events that might occur in various applications on the netcentric computing system 12.

The preferred operations architecture 10 also includes systems monitoring and tuning tools 96. The number of devices and the geographic disparity of devices used in a netcentric computing system 12 increase the effort required to monitor the system. The number of events generated in the netcentric computing system 12 rises due to the increased complexity. Devices such as clients 14, 34, network components (software and hardware), and servers 16, 18, 20 generate events on startup or failure to periodically report device status. The application used in the systems monitoring and tuning tools 96 are designed to detect and record predetermined events that occur on the clients 14, 34 or servers 16, 18, 20. The predetermined events may be from applications, databases, networks, clients 14, 34, servers 16, 18, 20. Those skilled in the art would recognize that the term event should be broadly construed to cover any event that can be monitored and recorded in the netcentric computing system 12.

The security tools 98 include applications that implement a predetermined security policy on the netcentric computing system 12. As illustrated in FIG. 1, the operations architecture 10 also include security tools 98. A security policy is the set of rules, directives, and practices that regulate how an organization manages, protects, and distributes sensitive information on the netcentric computing system 10. A security policy is translated into access control rules that are enforced by the security tools 98.

The preferred security tools 98 of the operations architecture 10 include identification tools and authentication tools. The identification tools are used to provide an identifier for users of the netcentric computing system 12. An identifier is a piece of data used to uniquely identify an entity in a transaction. The identifiers are unique and associate the entity with the identifier. The identifiers are issued to entities during part of a registration process that validates an entity's request to participate in a system, generates a unique identifier, binds that identifier to the requesting entity, and distributes the identifier to the now participant entity.

Once participating entities have been registered, the authentication tools validate the identifier during a transaction. Authentication applications validate that the entity requesting access to the resources of the netcentric computing system 12, whether that is a human or automated process, is the true owner of that identity. Authentication can be performed by three primary methods: by validating what the user/entity knows, what they have, or what they are. For instance, validating by what the user identity knows may be done by using secret passwords, PIN numbers, credit card numbers or mother's maiden name. Validating by what the user has can be done using an ATM card, credit card, smart card or a private key stored on an encrypted file on the client 14, 34. Validating by what the user is can be done using various biometric verification means such as voice prints, iris scan, signature verification and thumb scan.

The preferred security tools 98 provide access control to the netcentric computing system 12. Once the identity has been established, access control rules determine what resources the entity may use. Access control is used to permit or deny a specific type of use system resources on the netcentric computing system 12. For example, a user may be authorized to access a resource, but only for reading. Access control can be used to arbitrate access to files, processes, operating system ports, application functions, database tables,

portions of a network (such as through virtual or dedicated circuits and firewalls), and other types of resources. This is preferentially accomplished through the use of Access Control Lists (ACLs) in the netcentric computing system 12. An ACL for a resource specifies the user or group and the type of access permitted (read, write, etc.). ACLs may optionally include date and time restrictions and program restrictions.

Another way the security tools 98 can provide access to the netcentric computing system 12 may be through the use of role based access control. Role based access control associates a job function/role to a set of resources on the netcentric computing system 12, and then assigns the user to a particular role. So, for example, the role of junior bookkeeper may have read and write access to the petty case account, but read-only access to the general ledger. The advantage of role based access control is that it facilitates the management of access control and prevents users from retaining access to data that is no longer needed as they move from role to role.

Resource access control may be either restrictive or permissive in the netcentric computing system 12. Restrictive resource access control is based on the policy that whatever is not explicitly prohibited is allowed. Each of these methods has a use, depending on the requirements of the netcentric computing system 12. For network and firewalls 24, restrictive access control is commonly used. For most servers, 16, 18, 20, permissive access control is the norm. Those skilled in the art would recognize that variations exist on the exact manner in which access control is provided and are envisioned.

The preferred security tools 98 also include auditing tools. Auditing tools are used to record accesses to resources on the netcentric computing system 12, and may be implemented at a number of layers, including operating system, database, application, middleware, as well as in network devices such as firewalls 24 and routers. Auditing is typically implemented in combination of these layers to allow reconstruction of events after a security problem is detected. The logs kept by the auditing tools are preferentially searchable for known or suspected patterns of abuse, and are protected from alteration. Logs can monitor a variety of data, including access times, user Ids, locations, actions the user performed, and whether or not those actions were successfully completed.

The preferred security tools 98 may also include integrity tools. Integrity refers to the property that any system must have if it is to protect itself and enforce its security policy. During operation, the integrity tools protect the netcentric computing system 12 from buffer overflows, faulty parameters, or attacks on improperly-configured network ports have failed to meet the integrity requirement. The integrity tools also protect the netcentric computing system 12 from viruses. Viruses constitute what is probably the best known attack on integrity in a netcentric computing system 12.

The preferred security services 98 also includes cryptographic tools. Public key cryptography is one of the most important enabling technologies in the netcentric environment. The cryptographic tools ensure that messages are accessible only by those properly authorized, even when they traverse insecure networks. The term "message" broadly refers to an e-mail dispatch, or the more dynamic transactions of web sessions between clients 14, 34 and the web server 20. The cryptographic tools also ensure that a message is actually sent by the purported sender. Further, the cryptographic tools check for integrity to provide assurance that the message has not been modified in transit and also ensures that a sender cannot disavow a message.

The preferred cryptic tools use keys to encrypt communications. There are two types of keys used in the preferred netcentric computing system 12. A secret key is one type of key that is used and a key that is shared between two entities in a transaction. Because the same key is used to encrypt and decrypt data, this is referred to as symmetric key encryption. In order for the parties to communicate, they must establish the secret key in advance, using a secure channel. The most common implementation of a symmetric key algorithm is the Data Encryption Standard (DES). A public/private key pair or asymmetric key is the second type of key that is used and uses a pair of keys to encrypt and decrypt messages. Messages encrypted using one of the keys can only be decrypted with the other key. Each party possesses a pair of keys, one public key accessible to all participants in the system, and one private key accessible only to the party that owns it. The most common implementations of public key algorithms are supplied by RSA Data Security, Inc. In the most basic implementations, data is encrypted by the sender (i.e., client 14, 34 or server 16, 18, 20) with the public key of the recipient (i.e., client 14, 34 or server 16, 18, 20) and decrypted by the recipient with their private key.

Although public key cryptosystems do not require users to share a common secret key, key management is still a serious problem. Public key systems require a binding between a specific public/private key pair and an entity that is participating in the system. When using a public key to protect information destined for a specific entity, the user assumes that the public key he or she uses is really the one belonging to the entity. As such, in the preferred embodiment of the invention this binding is assured through the use of a trusted third party (TTP), called a Certificate of Authority, or CA.

Recall that the method for transmitting a message using public key cryptography is to encrypt the message with the receiver's public key. The benefit is that a user's public keys can be sent as clear text, or even published in a directory. So, if Alice wants to send a message to Bob, but is tricked into using Eve's public key, then Even will be able to intercept the message. (Eve can then, if she chooses, re-encrypt the message using Bob's actual public key, and neither Alice nor Bob will be the wiser.) In a netcentric computing system 12, which is in effect a global network lacking face-to-face contact, users must be assured they are using the right key. The CA provides this in the preferred netcentric computing system 12.

The CA serves a function analogous to that of a passport or drivers license in the netcentric computing system 12. The CA binds public keys to users and services similar to the way a passport agency issues you a passport that ties your name and relevant personal information to you with a picture. CAs deliver public keys through the use of certificates which are preferentially compliant with the X.509 standard. The CA will publish information to a directory, which contains an entry for each entity with a certificate.

Public key cryptosystems provide transaction authentication through the use of digital signatures. Digital signatures are created by the application of a has function to a piece of data (e.g., a message). This message hash is then encrypted with a sender's private key. The message recipient can use the sender's public key to decrypt the message hash, and rerun the hashing algorithm to make sure the hash has not changed. If the two hashes match, the sender has been properly authenticated. Note that for authentication, the pattern of public/private key use is the reverse of that for confidentiality. For confidentiality, the sender encrypts with the receiver's public key. To provide authenticity, the senders encrypt with their own private key.

The preferred cryptographic tools also include certification services that support activities needed to verify that the certificates are properly used, to ensure the authenticity and confidentiality of communications and stored data. Key recovery services are also provided under the cryptographic tools. Data encrypted under a public key cannot be recovered without the private key. If the private key is rendered inaccessible (through file corruption, token
5 destruction, or failure), it is essential that the cryptosystem owner/operator provide a means for recovering that data. As such, the key recovery services allow private keys to be recovered in case the private key is lost or damaged.

The preferred cryptographic tools also include revocation services. In any public key
10 cryptosystem, keys will eventually be compromised, either because they are lost or stolen. The revocation services allow users to notify an administrator if their keys are compromised, to disseminate the list of compromised keys to all participating entities, and to issue new keys to replace compromised keys. Since public key binding is typically carried out using X.509 compliant certificates, this process may also be referred to as certificate revocation.

Referring to FIG. 3, a preferred implementation of public key cryptography for the
15 netcentric computing system 12 is illustrated. The remote client 34 has a personal private key stored on an encrypted file in the client 34 or possibly on a more secure device such as a smart card. The client has access to a repository 112 used for a public key storage to obtain public keys of other entities. The web server 20 has its own private key stored in a secure file or
20 cryptographic device in the web server 20. The combination of a public and private key pair associated with the web server 20 and the remote client 34, along with cryptographic software on either end, enable the two entities to authenticate to each other, send encrypted data, and digitally sign documents over the Internet connection 32. In some cases, an authentication server 114, or directory service such as LDAP, can be used to validate the user's current
25 access rights to the system. Authentication services and directories are often used to supplement certification revocation lists to provide faster and more granular authorization.

As illustrated in FIG. 1, the preferred operations architecture 10 also includes user
administration tools 100. The netcentric computing system 12 introduces many new challenges to the task of user administration. The majority of these stem from the dramatically increased
30 number of system components. Adding a user to the netcentric computing system 12 may

require adding a user to the network, one or more server 16, 18, 20 operating systems, one or more database systems (so that the user can access data), an e-mail system, and an existing host-based system. In some cases, the addition of a user may require entries to be added to several individual system components. The preferred user administration tools 100 allow an administrator to add users to the netcentric computing system 12 in an orderly and automated fashion to eliminate the problems encountered when adding users. The preferred user administration tools 100 also allow the administrator to delete users from the netcentric computing system 12. Unless careful records are kept, it can be very difficult to determine to which machines, databases, and applications the user had been added originally so that this information can be deleted. From an administration standpoint this may seem to be only a headache, but from a security standpoint it represents a substantial risk. The preferred user administration tools 100 keep track and allow the deletion of users to be accomplished in an orderly and automated fashion.

Most user administration products on the market today focus on the operating system aspect of the problem (adding user access to the server, setting file permissions, group associations). Although these solutions are certainly helpful, they do not cover many of the more difficult user administration challenges such as database access, e-mail, and networking software. Each of these products often comes with its own administration tools which may simplify the individual administration tasks but do little to help with providing an integrated user administration approach. The preferred user administration tools 100 provide an integrated approach to handling the types of access that users are granted to the netcentric computing system 12.

An alternative approach to the user administration tools is to implement a single sign-on (SSO) application in the netcentric computing system 12. These applications are meant to eliminate the need for users to remember user names and passwords to all of their business applications. The first time they log in, users enter a user name and password into the SSO application which then automatically logs into applications through a scripting process. An advantage to this approach is that through implementing SSO, a database that maps users to the applications they access is created. This significantly simplifies user administration, and can increase security as well. A key drawback to SSO applications is failover. If a SSO server fails, users cannot access applications as they do not remember passwords to all their applications.

The preferred operations architecture 10 also includes production control application set 102. In distributed environments, processes may be taking place across the entire system on multiple platforms in either a parallel or a serial fashion. Batch dependencies may be required across platforms, and multiple time zones may be involved. In addition, many non-mainframe-based applications do not provide production scheduling capabilities included with the application. For these reasons, scheduling processes across the netcentric computing system 12 can be quite complex, requiring significant management effort to ensure that the processes run smoothly. The preferred production control application set 102 includes print management tools, file transfer and control tools, mass storage management tools, backup and restore tools, archiving tools, and system startup and shutdown tools that ensure that processes run smoothly on the netcentric computing system 12.

The print management tools include applications that handle printing documents on the netcentric computing system 12. The file transfer and control tools handle the transferring of files from one location to another location in the netcentric computing system 12. The mass storage management tools monitor and control database files and various other kinds of data files that are stored in the netcentric computing system 12.

The backup and restore tools are used by the netcentric computing system 12 to backup and restore files that are used on the netcentric computing system 12. Backup and restoration processes become more complex in a distributed environment as business-critical information becomes distributed across the netcentric computing system 12. Backup strategies used coordinate information across the netcentric computing system 12 and determine where the backup copy or copies of information will reside. As with centralized computing environments, restoration processes are directly dependent on how backup was performed. A single restore process no longer suffices. Depending on a particular fault, restoration services may only need to be performed for a portion of the netcentric computing system 12, while the rest of the netcentric computing system 12 stays up and running.

The archiving tools include applications that are responsible for archiving files on the netcentric computing system 12. The issues surrounding archiving are quite similar to those surrounding backup. The archiving tools place limitations on the amount of information that may be archived on the netcentric computing system 12 as a result of the space limitations on servers 16, 18, 20 and clients 14, 34. Additional problems are created with archiving in a

distributed environment, because users have no incentives to perform housekeeping tasks on their devices. Depending on the users' ability to store information on the clients 14, 34 or on the server 16, 18, 20, the clients 14, 34 may become cluttered with seldom-used files. Lack of space may affect other processes that need to take place on these devices, such as software and data
5 distribution. The preferred archiving tools solve these problems by providing regular archiving procedures that automatically archive a predetermined set of files.

Referring to FIG. 1, the preferred operations architecture 10 also includes a help desk tool 104. The netcentric computing system 12 puts the operations help desk tools 104 closer to the "end user" in terms of visibility and influence. The help desk tools 104 are integrated with the
10 business processes being supported through the netcentric computing system 12. If the help desk tools 104 are well integrated with the business process, there is risk, that the user may be given incorrect information, be forwarded to the wrong department, or otherwise mishandled. It is also important that the information collected by the help desk tools 104 about a user be properly shared with other stakeholders in the business process, which is also provided by the preferred
15 help desk tools 104.

The preferred help desk tools 104 turn web browsers on the clients 14, 34 into interactive clients of the help desk with the power to enter, query and modify help desk requests. The preferred help desk tools 104 allow users to directly perform most of the help services provided by the help desk tools 104 without assistance from the help desk staff. As such, the preferred
20 help desk tools 104 are capable of providing automated assistance to users on the clients 14, 34.

Another key function provided by the help desk tools 104 in the netcentric computing system 12 is for users to more effectively support themselves. In Internet environments, it is usually prohibitively expensive for a service provider to provide interactive help desk support to all interested Internet users. This is due to potential volumes of support requests as well as the
25 diversity of technical environments that could be encountered. Consequently, it is often more reasonable to provide Internet users with access to the required applications of the help desk tools 104. In the preferred embodiment, the preferred help desk tools 104 include a download site where patches, drivers, and self-help support materials are available.

The help desk tools 104 also use metrics to measure the performance of support
30 personnel that consider interactions via e-mail or video. An example metric might be the "number of e-mails answered per hour." In addition, existing metrics may need to be refined to

fairly reflect netcentric characteristics. Those skilled in the art would recognize that several metric values can be monitored and kept track of by the netcentric computing system 12.

The preferred help desk tools 104 are available continuously in the netcentric computing system 12. In addition, in netcentric computing systems 12 there may be additional complexities of help desk operations introduced by global interactions. For example, the preferred help desk tools 104 support web page generation and e-mail support using multiple languages. Those skilled in the art would recognize that the applications used in the preferred help desk tools 104 will vary depending on the needs of each particular enterprise.

As set forth above, the operations architecture 10 consists of different operations tools that focus on different functions, such as the help desk tools or fault management and recovery management tool 84. Each tool introduces a predetermined set of operations services such as core management logic and event generation. Although product selection decisions are often based on the functions that a product provides, true integration of these tools into a cohesive operations architecture requires a service-based view, rather than a functional view and many specialized applications to integrate the tools.

It is therefore important to consider the services provided by the operations architecture tools when selecting operations tools. The services provided by the operations architecture 10 are core management logic, integration platform, event/data generation, event processing, and repositories.

The core management logic services apply business roles to management data. Core management logic is typically specific to the function being served by an operations tool. For example, core management logic of a backup/restore application of the production control application set 102 would initiate a backup process based on the time of day information it receives from a system clock. Core management logic receives data from event/data generation, event processing, and repositories services and then sends data for presentation or to repositories services. In addition, core management logic often polls the event/data generators for information.

The integration platform services provide a common platform for the tools of the operations architecture 10. At the lowest level this means common standards, interfaces, message formats, and file logging forms to be used with all the tools. Though the integration

platform can be homegrown, these applications are growing extremely complex, suggesting the use of one of many available third party integration platforms.

There are two types of third party platforms available. The first group are framework type products such as HP Open View, CA-Unicenter TNG, and Tivoli Management

5 Environment. These products are modular. Each module within the suite can be run separately; however, they all conform to a common framework which allows for greater compatibility, integration and better performance. The second type of integration platform is point-solution oriented. Products like Boole and Babbage implement this approach which typically results in best-of-breed solutions for various management solutions, but a larger amount of integration
10 work between tools is required.

The event/data generation services interact with all the managed components in the execution and development environments in order to produce the required management information. The output of event/data generation services is actual raw management data that can then be processed and acted upon.

15 The event processing services manipulate the raw data obtained by event/data generation services into a form on which operations personnel can take action. This service may perform several functions such as event filtering, alert generation, event correlation, event collection and logging, and automatic trouble ticket generation. When management events are generated, event filtering mechanisms constantly compare predetermined event thresholds to current management
20 events to determine the need for a management alert. If the threshold is exceeded, the event filtering function takes a specific action based on predetermined rules. When an event filter has noted the need for an alert, the alert generation function creates the proper notification. This may take one of several forms; a page, an email, a display change (icon changes color to red), etc.

Event correlation functions use logic to tie different events together with the intention of
25 understanding potentials causes of problems. For example, nightly processing utilization shortages may be tied by event correlation functions back to a nightly batch job. It may be determined that historical analysis of management events is important. If so, the collection and logging of management events into repositories is important so that reporting and correlation activities can be performed at a future time. Automated trouble ticket generation. For certain
30 events, it may be desirable for trouble tickets to be generated automatically in an organization's help desk system so that action can be taken.

The repository services contain all the management data generated or used during the management process. This data includes historical data, capacity data, performance data, problem knowledge bases, asset databases, solution sets, and management information bases (MIBs).

5 The preferred operations architecture 10 consists of a set of tools that allow administrators to effectively manage a distributed environment. While the invention has been described in its currently best known modes of operation and embodiments, other modes and embodiments of the invention will be apparent to those skilled in the art and are contemplated. For other features, advantages and combinations of the present invention refer to U.S.
10 provisional application Serial No: 60/156,962, which is herein incorporated by reference in its entirety.

What is claimed is:

1. An operations architecture for a netcentric computing system, comprising:
a server connected with a client; and
a software distribution tool, a configuration and asset management tool, a fault
5 management and recovery management tool, a capacity planning tool, a performance
management tool, a license management tool, a remote management tool, a event
management tool, a systems monitoring and tuning tool, a security tool, a user
administration tool, a production control application set and a help desk tool supporting
said server and said client in said netcentric computing system.
10
2. The operations architecture of claim 1, wherein said software distribution tool
provides automated delivery to, and installation of, applications on said server and said client.
3. The operations architecture of claim 1, wherein said configuration and asset
15 management tool that manages a plurality of predetermined assets connected with said
netcentric computing system.
4. The operations architecture of claim 3, wherein said predetermined assets may
be selected from the group consisting of said server, said client, a product license information
20 file, a warranty information file, a vendor name file, a logical device information file and a
physical device information file.
5. The operations architecture of claim 1, wherein said fault management and
recovery management tool assists in the diagnosis and correction of a plurality of system faults
25 in said netcentric computing system.
6. The operations architecture of claim 1, wherein said capacity planning tool
monitors a plurality of predetermined system usage levels in said netcentric computing system.

7. The operations architecture of claim 6, wherein said system usage levels may be selected from the group consisting of server processing usage, server bandwidth usage, server storage usage and client usage.

5 8. The operations architecture of claim 1, wherein said performance management tool monitors the performance of applications running on said netcentric computing system.

9. The operations architecture of claim 1, wherein said license management tool manages and controls license information for applications running on said netcentric computing
10 system.

10. The operations architecture of claim 1, wherein said remote management tool allows support personnel from said netcentric computing system to take control of said client.

15 11. The operations architecture of claim 1, wherein said event management tool is responsible for handling a plurality of predetermined events in said netcentric computing system.

12. The operations architecture of claim 11, wherein said predetermined events may
20 be selected from the group consisting of disk space indications, central processing unit utilization indications, database error indications, network error indications and file and print service indications.

13. The operations architecture of claim 1, wherein said systems monitoring and
25 tuning tool monitors applications, middleware, databases, networks, clients and servers on said netcentric computing system.

14. The operations architecture of claim 1, wherein said security tool includes applications that provide security to said netcentric computing system.

30

15. The operations architecture of claim 1, wherein said user administration tool is used for administering users of said netcentric computing system.

5 16. The operations architecture of claim 1, wherein said production control application set is used for scheduling and processing a plurality of production processes on said netcentric computing system.

10 17. The operations architecture of claim 16, wherein said production control application set may be selected from the group consisting of a print management tool, a file transfer and control tool, a mass storage management tool, a backup and restore tool, a archiving tool and a system startup and recovery tool.

15 18. The operations architecture of claim 1, wherein said help desk tool provides a help application for assisting users of applications on said netcentric computing system.

19. An operations architecture for a netcentric computing system, comprising:
a server connected with a client;
a software distribution tool for providing automated delivery to, and installation of, an application on said server or said client;
20 a configuration and asset management tool for managing a plurality of predetermined assets connected with said netcentric computing system;
a fault management and recovery management tool for assisting in the diagnosis and correction of a plurality of system faults in said netcentric computing system;
a capacity planning tool for monitoring a plurality of predetermined system
25 usage levels in said netcentric computing system;
a performance management tool for monitoring the performance of applications running on said netcentric computing system;
a license management tool for managing and controlling license information for applications running on said netcentric computing system;

a remote management tool allowing support personnel from said netcentric computing system to take control of said client;

a event management tool for handling a plurality of predetermined events in said netcentric computing system;

5 a systems monitoring and tuning tool for monitoring applications, middleware, databases, networks, clients and servers;

a security tool that includes a security application that provides security to said netcentric computing system;

10 a user administration tool for administering users of said netcentric computing system;

a production control application set for scheduling and handling a plurality of production processes on said netcentric computing system; and

a help desk tool including a help application that provides users of applications on said netcentric computing system with assistance.

15

20. The operations architecture of claim 19, wherein said predetermined assets may be selected from the group consisting of said server, said client, a product license information file, a warranty information file, a vendor name file, a logical device information file and a physical device information file.

20

21. The operations architecture of claim 19, wherein said system usage levels may be selected from the group consisting of server processing usage, server bandwidth usage, server storage usage and client usage.

25

22. The operations architecture of claim 19, wherein said predetermined events that said event management tool handles may be selected from the group consisting of disk space indications, central processing unit utilization, database error indications, network error indications and file and print server indications.

23. The operations architecture of claim 19, wherein said production control application set may be selected from the group consisting of a print management tool, a file transfer and control tool, a mass storage management tool, a backup and restore tool, a archiving tool and a system startup and recovery tool.

5

24. A method of providing an operations architecture for a netcentric computing system including a client and a server, comprising the steps of:

- 10 using a software distribution tool for providing automated delivery to, and installation of, a predetermined application on said server or said client;
- managing a plurality of predetermined assets connected with said netcentric computing system with a configuration and asset management tool;
- 15 assisting in the diagnosis and correction of a plurality of system faults in said netcentric computing system with a fault management and recovery management tool;
- monitoring a plurality of predetermined system usage levels in said netcentric computing system with a capacity planning tool;
- 20 monitoring the performance of applications running on said netcentric computing system with a performance management tool;
- managing and controlling license information for applications running on said netcentric computing system with a license management tool;
- allowing support personnel to take control of said client with a remote management tool;
- 25 handling a plurality of predetermined events in said netcentric computing system with a event management tool;
- monitoring a plurality of computing devices connected with said netcentric computing system with a systems monitoring and tuning tool;
- securing said netcentric computing system with a security tool;
- administering users of said netcentric computing system with a user administration tool;

scheduling and handling a plurality of production processes on said netcentric computing system with a production control application set; and
helping users encountering problems with applications on said netcentric computing system with a help desk tool.

5

25. The method of claim 24, wherein said predetermined assets may be selected from the group consisting of said server, said client, a product license information file, a warranty information file, a vendor name file, a logical device information file and a physical device information file.

10

26. The method of claim 24, wherein said system usage levels may be selected from the group consisting of server processing usage, server bandwidth usage, server storage usage and client usage.

15

27. The method of claim 24, wherein said predetermined events that said event management tool handles may be selected from the group consisting of disk space indications, central processing unit utilization, database error indications, network error indications, application error indications and file and printer service indications.

20

28. The method of claim 24, wherein said production control application set may be selected from the group consisting of a print management tool, a file transfer and control tool, a mass storage management tool, a backup and restore tool, a archiving tool and a system startup and recovery tool.

25

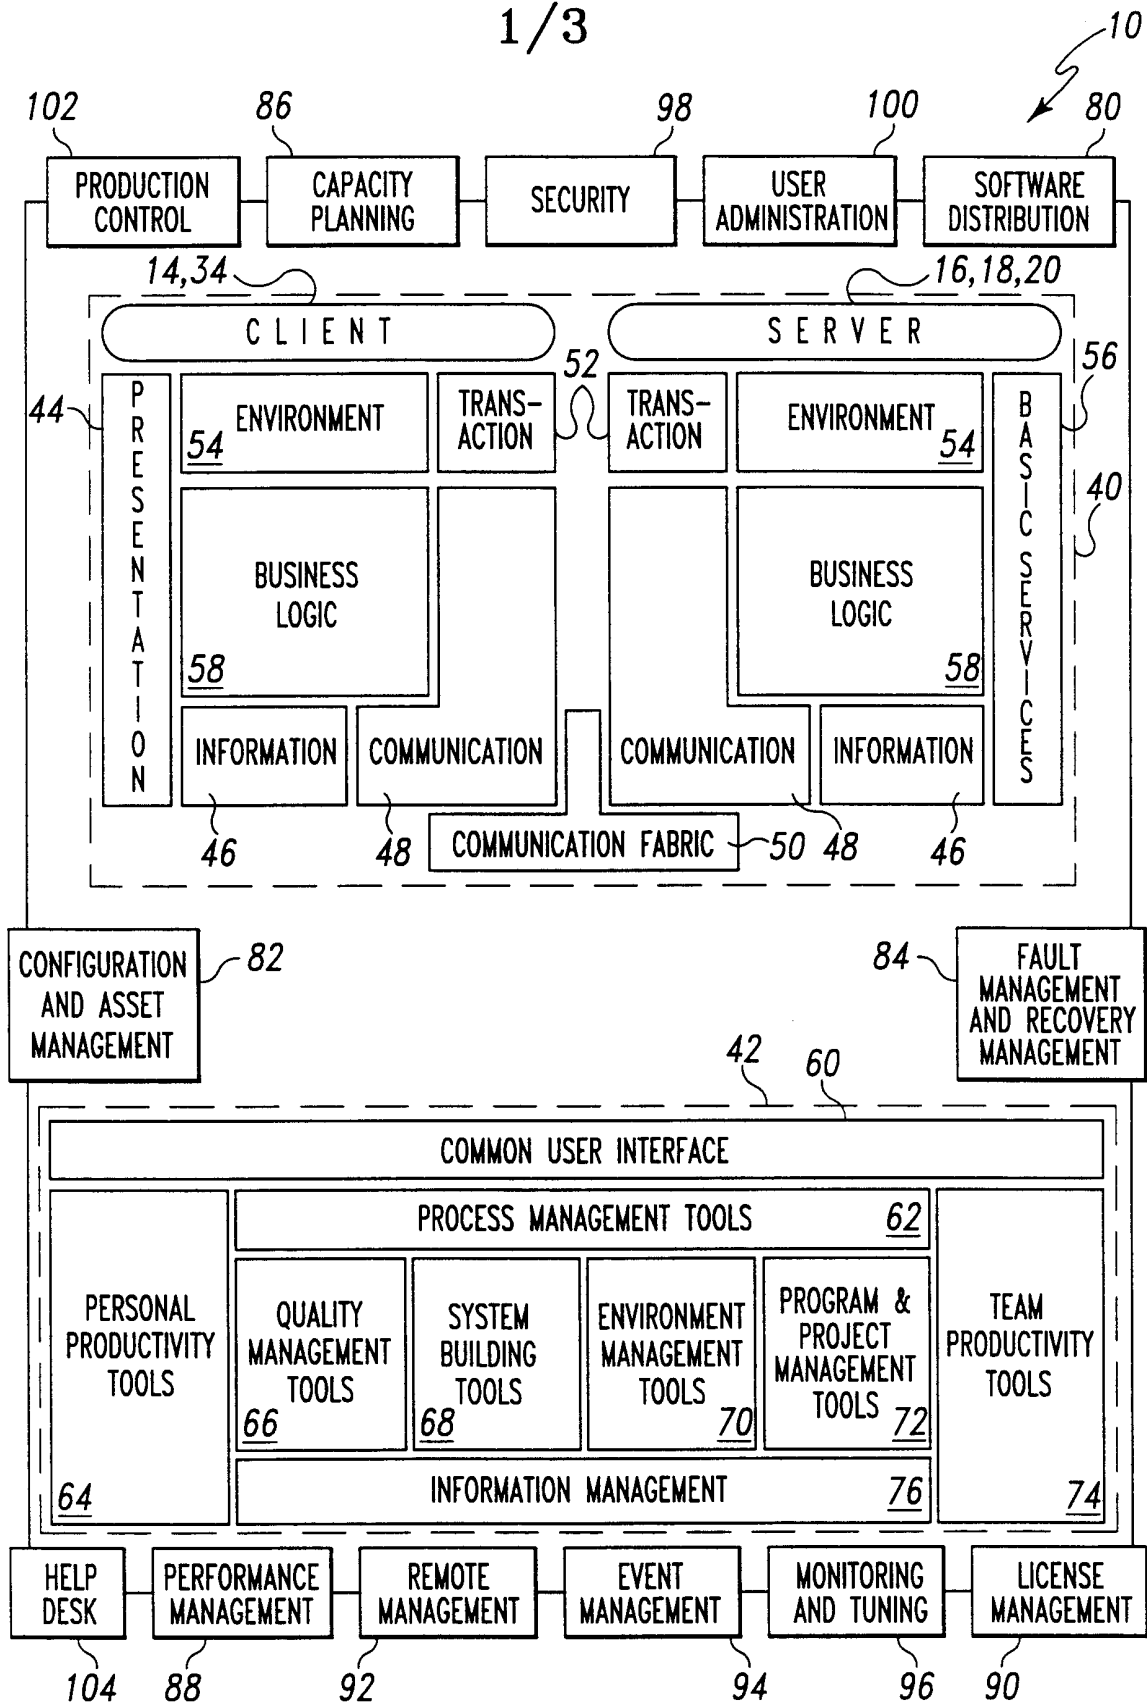


Fig. 1

Enterprise

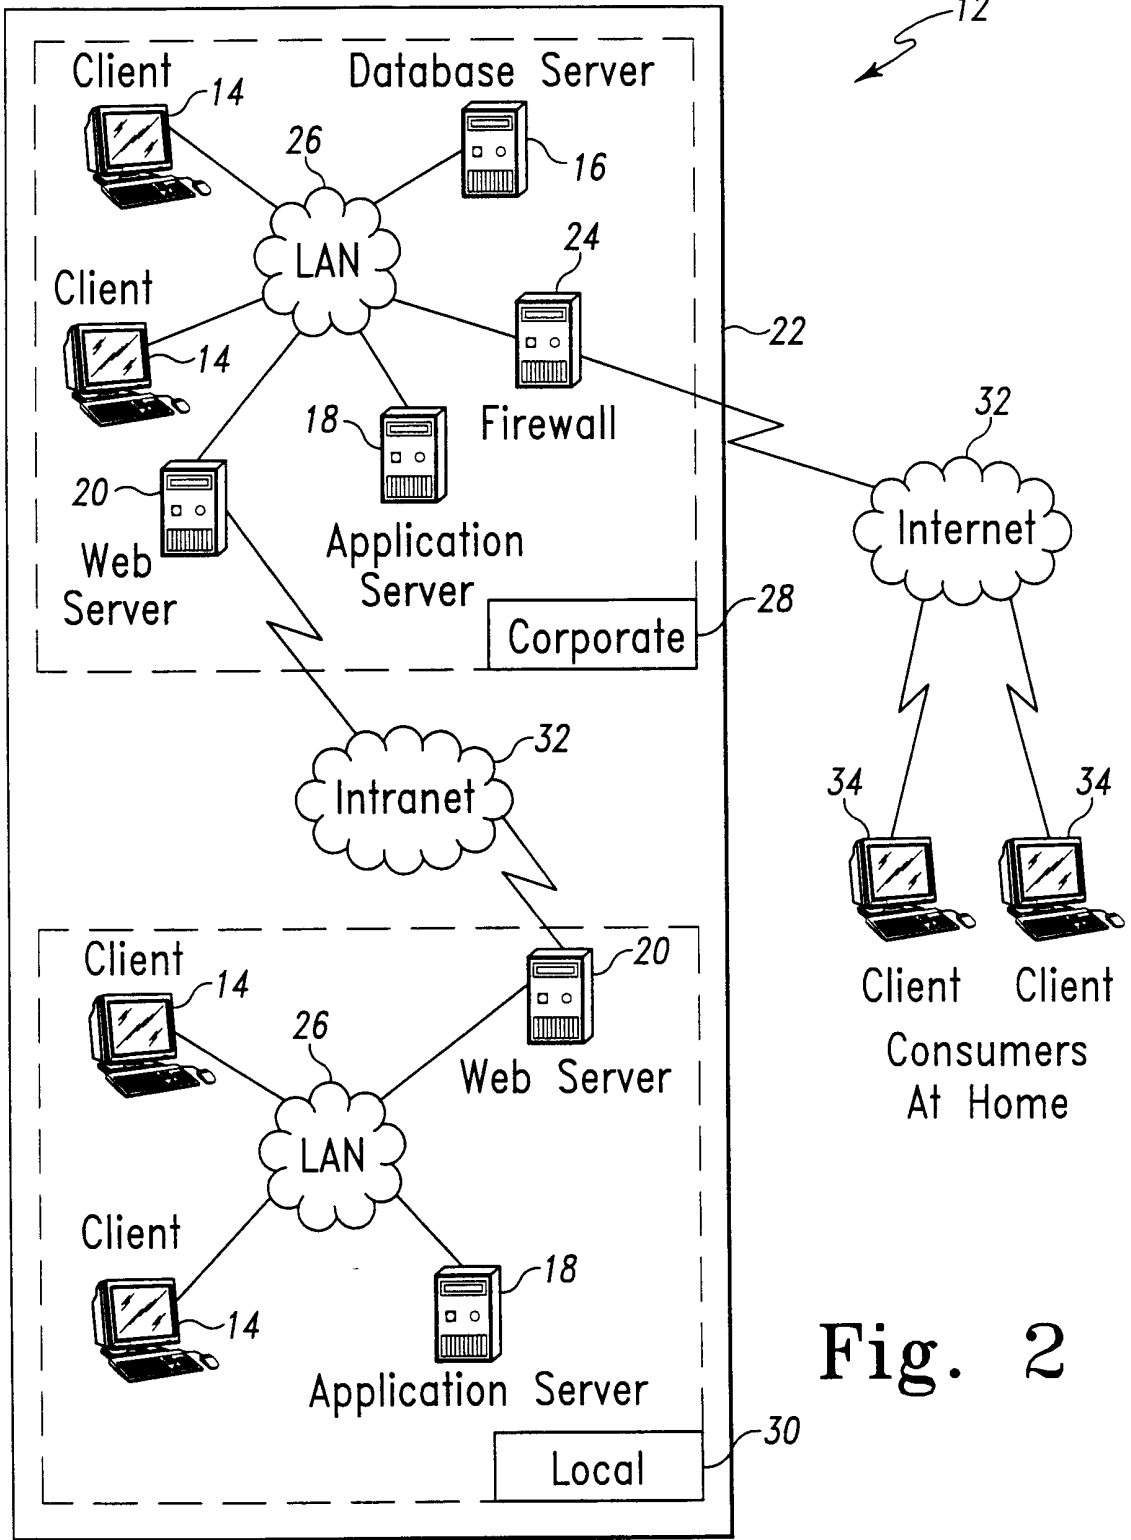


Fig. 2




Fig. 3