

(19) World Intellectual Property Organization
International Bureau



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

PCT

(10) International Publication Number
WO 01/25954 A2

(51) International Patent Classification⁷: G06F 17/00

(21) International Application Number: PCT/US00/26987

(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): **SARGENT, Scott,**

R. [US/US]; 415 Essex Road, Kenilworth, IL 60043 (US). **MULLEN, Nancy, K.** [US/US]; 412 Kings Way, Mundelein, IL 60060 (US). **TAYLOR, Stanton, J.** [US/US]; 31475 N. Reigate Lane, Green Oaks, IL 60048 (US).

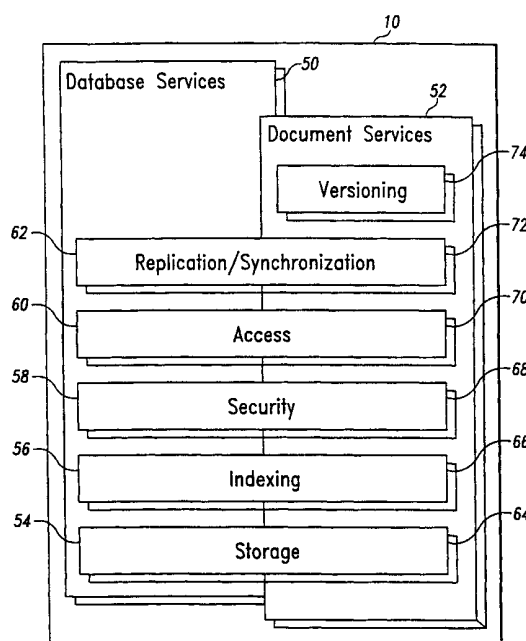
(74) Agent: **MCCONNELL, Dean, E.**; Brinks Hofer Gilson & Lione, One Indiana Square, Suite 2425, Indianapolis, IN 46204 (US).

(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.

(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

[Continued on next page]

(54) Title: INFORMATION SERVICE ARCHITECTURES FOR NETCENTRIC COMPUTING SYSTEMS



(57) Abstract: An information service architecture for a netcentric computing system is disclosed by the present invention. The information service architecture includes at least one client connected with a web server. A database service that includes a database storage service, a database indexing service, a database security service, a database access service and a database replication/synchronization service is located on said web server for interacting and exchanging information with the client. A document service that includes a document storage service, a document indexing service, a document security service, a document access service, a document replication/synchronization service and a document versioning service that are located on the web server for interacting and exchanging information with the client.



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).

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.

Published:

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

INFORMATION SERVICE ARCHITECTURES FOR NETCENTRIC COMPUTING SYSTEMS

5 This application claims the benefit under 35 U.S.C. §119(e) of U.S. provisional 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 information service architectures 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 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-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.

25

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).

5 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
10 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
15 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
20 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
25 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 they so desire. 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.

30 **Summary of the Invention**

The preferred embodiment of the present invention discloses an information service architecture for a netcentric computing system. The information service architecture for the netcentric computing system includes at least one client connected with a web server. A

database service that includes a database storage service, a database indexing service, a database security service, a database access service and a database replication/synchronization service is located on said web server for interacting and exchanging information with the client. A document service that includes a document storage service, a document indexing service, a document security service, a document access service, a document replication/synchronization service and a document versioning service that are located on the web server for interacting and exchanging information with the client.

In the preferred embodiment, the database storage service manages and physically stores data in the form of files in at least one database. The database indexing service uses indexing applications that enable the client to retrieve information from the files in the database storage service. The indexing applications may be selected from the group of applications consisting of bitmap indexing applications, context indexing applications and star indexing applications.

The database security service use at least one application to enforce access control to ensure that records and data stored in the database storage service are only visible or editable by authorized clients. The database access service enables applications on the client to retrieve data from a database in the database storage service. The database replication/synchronization service uses applications to support making multiple copies of at least one database in the database storage service. As known in the art, databases may be stored on database servers or web servers and may be stored on hard drives, tape or digital media. Those skilled in the art would recognize various methods exist to store files.

The document storage services manage physical storage of a plurality of document files within the document services. The document indexing service use index applications that support searching a plurality of document files stored in the document storage services. The document index applications may be selected from the group of applications consisting of attribute searching applications, full-text searching applications, context searching applications or Boolean searching applications.

The document security service use security applications that restrict access to users on the clients to a plurality of document files stored in the document storage service. The document access service provides access applications that support document creation, maintenance and retrieval. The access applications may be selected from the group

consisting of attribute search applications, full-text search applications, context search applications and Boolean search applications. The document replication/synchronization service uses applications to maintain multiple copies of document files stored in the document storage services. The document versioning service uses at least one application to maintain a historical record of changes to document files stored in the document storage service.

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.

Brief Description of the Drawings

Figure 1 illustrates a block diagram of a preferred information service architecture for a netcentric computing system.

Figure 2 illustrates a representative netcentric computing system.

Detailed Description of the Preferred Embodiments of the Invention

Referring to Figs. 1 and 2, the present invention discloses an information service architecture 10 for a netcentric computing system 12. The preferred information service architecture 10 is used to distribute, manage and store various types of data. In the preferred embodiment, data can take many forms, such as graphic files, image files, voice files, text files and full-motion video files. In order to gain a better understanding of the present invention, a brief discussion of the netcentric computing system 12 is set forth below, which is followed by a detailed discussion of the preferred information service architecture 10.

Referring to Fig. 2, the physical picture of an illustrative netcentric computing system 12 is illustrated. A business enterprise 18 may include at least one client 14, 20, at least one database server 22, at least one firewall 24, at least one application server 26, at least one web server 28 and a LAN connection 30 connected as illustrated in Fig. 2. The LAN connection 30 is used to interconnect various components or computing devices that are located at a first business enterprise location 32 within the business enterprise 18. Those skilled in the art would recognize that various types of LAN connections 30 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. As known in the art, firewalls isolate the web servers 28 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 on the web servers 28. All requests for other applications (e.g., FTP, Telnet) and other IP addresses are blocked by the firewall 24.

5 The web servers 28 are the primary interface to the clients 14, 20 for all interactions with the applications or services of the netcentric computing system 12. The main task of the web servers 28 is to authenticate the clients 14, 20, establish a secure connection from the clients 14, 20 to the web servers 28 using encrypted messages, and allow the applications the clients 14, 20 are using to transparently access the resources of the netcentric computing system 12. The web servers 28 are responsible for accepting
10 incoming HTTP messages and fulfilling the requests. For dynamic HTML page generation, requests are forwarded to the application servers 26. Static pages, such as help pages, are preferably generated by the web servers 28.

In the preferred embodiment, the primary function of the application servers 26 is to provide a link through which the web servers 28 can interact with the host, trigger business
15 transactions, and send back resulting data to the clients 14, 20. A fundamental role of the application servers 26 is to manage the logical flow of the transactions and keep track of the state of the sessions. The application servers 26 are also responsible for managing all sessions.

Further, in the preferred embodiment of the present invention, the main purpose of
20 the database servers 22 is to handle an application log. All requests sent to the web servers 28 and application servers 26 as well as 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 26. Those skilled in the art would recognize that any number of data items can be monitored by the application log.

25 As further illustrated in Fig. 2, a second business enterprise location 34 may be connected with the first business enterprise location 32 using a intranet connection 36. Those skilled in the art would recognize that various intranet connections 36 exist and may be used in the present invention. As those skilled in the art would recognize, the intranet connection 36 allow the computing resources of the second business enterprise location 34 to be shared
30 or connected with the computing resources available at the first business enterprise location 32. Although not fully illustrated in Fig. 2, several other enterprise locations may be connected with the netcentric computing system 12.

In the preferred embodiment the firewall 24 of the first business enterprise location 32 is connected with a dedicated Internet connection 38 to a plurality of remote clients 14. Preferentially, the remote clients 14 that are connected to the Internet connection 38 access data at the business enterprise 18 through the Internet connection 38 using a browser application. The Internet connection 38 gives the remote clients 14 the ability to gain access to information and data content contained on the database server 22, the application server 26 and the web server 28.

For a detailed discussion of the preferred architecture for the netcentric computing system 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 back to Fig. 1, the preferred information service architecture 10 manages information and enables applications to access and manipulate data stored locally or remotely from documents, databases, or external data sources. The preferred information service architecture 10 is one that minimizes a software applications dependency on physical storage and location within the netcentric computing system 12. In the preferred embodiment, the information service architecture 10 can also be accessed and used by the end user when ad-hoc data and document access are integral to the software application work tasks.

In the preferred information service architecture 10, a two-layer approach is used to keep information distinct from the processes that access and use it and include a logical layer and a physical layer. Within the netcentric computing system 12, the information services architecture 10 maintains this logical/physical distinction. The logical layer acts to isolate the physical aspects of information (e.g., location, storage format, and access language) from applications and applications developers. This layer provides all the detail services associated with information and with access to or from that information.

The physical layer can be used within a netcentric architecture to isolate the detailed technical implementations of information. This layer insulates an organization and its applications from the rapid pace of change in information management technology. This layer can also be used to position legacy information sources into the netcentric computing environment, independent from migrating applications and implementing new applications.

As illustrated in Fig. 1, the preferred information service architecture 10 includes database services 50 and document services 52. Database services 50 are responsible for providing access to a local or remote database as well as maintaining integrity of the data

within the database. Data may also be stored on either a single physical platform or in some cases across multiple platforms. The database services 50 are typically provided by database management systems (DCMS) vendors and accessed via embedded or call-level SQL variants and supersets. Depending upon the underlying storage model, non-SQL access methods may
5 be used instead.

The preferred database services 50 include database storage services 54, database indexing services 56, database security services 58, database access services 60 and database replication/synchronization services 62. Database storage services 54 manage and store the actual physical data. As known in the art, database storage services 54 provide a mechanism
10 for saving information so that data will live beyond program execution. Data is often stored in relational format (an RDBMS) but may also be stored in an object-oriented format (OODBMS) or other structures such as IMS and VSAM.

The database indexing services 56 provide a mechanism for speeding up data retrieval from the database storage services 54. In relational databases one or more fields can be used
15 to construct the index. Therefore, when a user searches for a specific record, rather than scanning the whole table sequentially, an index is used to find the location of that record faster. Revolutionary advances in indexing techniques, such as bitmapped indexing, context indexing, and star indexes, provide rich capabilities for database indexing services 56.

The preferred database security services 58 contain applications that enforce access
20 control to ensure that records and data are only visible or editable by authorized people for approved purposes. As known in the art, most database management systems (DBMS) provide access control at the database, table, or row levels to specific users and groups as well as concurrency control. They also provide execution control for such things as stored procedures and database functions. However, there may be severe limitations in DBMS's
25 ability to pass data needed for security authentication across a network, forcing the architect to build those services into the database security services 58.

The database access services 60 enable an application to retrieve data from a database as well as manipulate (insert, update, or delete) data in a database. As known in the art, SQL is the primary approach for accessing records in today's database management systems.
30 Netcentric computing systems 12 often require data access from multiple databases offered by different vendors. This is often due to integration of new database systems with existing legacy database systems. In the preferred database access service 60, the key architectural concern is in building the application where the multi-vendor data problem is transparent to

the application needing the data. This provides future portability and flexibility and also makes it easier for application developers to write to a single database access interface.

Achieving database access transparency requires the following: standard-based SQL API, SQL Gateways, Distributed Relational Data Access (DRDA) and IBI's EDA/SQL and the Sybase/MDI Open Server. Standards-Based SQL API uses a single, standards-based set of APIs to access any database and includes the following technologies: Open Database Connectivity (ODBC), Java Database Connectivity (JDBC), and Object Linking and Embedding (OLE DB). SQL Gateways provide a mechanism for clients to transparently access data in a variety of databases (e.g., Oracle, Sybase, or DB2) by translating SQL calls written using the format and protocols of the gateway server or primary server to the format and protocols of the target database. Currently there are three contending architectures for providing this gateway function.

Distributed Relational Data Access (DRDA) is a standard promoted by IBM for distributed data access between heterogeneous databases. In this case, the conversion of the format and protocols occurs only once. It supports SQL89 and a subset of the SQL92 standard and is built on top on APPC/APPN and TCP/IP transport stacks. IBI's EDA/SQL and the Sybase/MDI Open Server use SQL to access relational and non-relational database systems. They use API/SQL or T-SQL, respectively, as the standard interface language. A large number of communication protocols are supported, including NetBIOS, SNA, DecNET, and TCP/IP. The main engine translates the client requests into specific server calls to the web server. It handles security, authentication, statistics gathering, and some system management tasks.

The preferred database replication/synchronization services use applications to support an environment in which multiple copies of databases must be maintained. This is very important to enterprises that must maintain data in the event that one or more systems go down due to hardware or software failure. In addition, the database replication/synchronization services use applications that perform the transactions required to make one or more information sources that are intended to mirror each other consistent. Those skilled in the art would recognize that the method in which the preferred database replication/synchronization services performs its tasks will vary from enterprise to enterprise, depending on the particular needs of the enterprise.

Replication can be either complete or partial. During complete replication, all records are copied from one destination to another; during partial replication, only a subset

of data is copied, as specified by the user or the program. Replication can also be done either real time or on demand (i.e., initiated by a user, program, or scheduler). The following becomes possible if databases are replicated on alternate server(s): better availability or recoverability of distributed applications; better performance and reduced network cost, particularly in environments where users are widely geographically dispersed; and improved access to wider ranges of data, as data replicas may be more readily available.

Synchronization services perform the transactions required to make consistent one or more information sources that are intended to mirror each other. This function is especially valuable when implementing applications for remote users or users of mobile devices because it allows a working copy of data or documents to be available locally without a constant network attachment. The emergence of applications that allow teams to collaborate and share knowledge (e.g., the Knowledge Xchange® at Andersen Consulting) has heightened the need for synchronization services in the information service architecture.

10. The terms replication and synchronization are used interchangeably in regards to the database replication/synchronization service 62.

The database document services 52 provide similar structure and control for documents that database services 50 apply to record-oriented data. As known in the art, a document is defined as a collection of objects of potentially different types (e.g., structured data, unstructured text, images or multimedia) that a business user deals with. Regardless of the application used to create and maintain the component parts, all parts together constitute the document, which is managed as a single entity, usually in the form of a document file. An individual document might be a table created using a spreadsheet package, a report created using a word processing package, a Web page created using an HTML authoring tool, unstructured text, or a combination of these object types.

As further illustrated in Fig. 1, the preferred document services 52 include document storage services 64, document indexing services 66, document security services 68, document access services 70, document replication/synchronization services 72 and document versioning services 74. The document storage services 64 manage the physical storage of documents within the document services 52 of the information service architecture 10. In the preferred embodiment, most document management products store documents as objects that include two basic data types: attributes and content.

Document attributes are key fields used to identify the document, such as author name or created date. Document content refers to the actual unstructured information stored within the document. Generally, the documents are stored in a repository in one of the document storage services 64 using one of the following methods: proprietary database, industry standard database, or industry standard database and file system.

In proprietary databases documents (attributes and contents) are stored in a proprietary database, one that the vendor has specifically developed for use with its product. In industry standard databases documents (attributes and contents) are stored in an industry standard database such as Oracle or Sybase. Attributes are stored within traditional database data types (e.g., integer or character); contents are stored in the database's BLOB (Binary Large Objects) data type. In industry standard database and file systems documents' attributes are stored in an industry standard database, and documents' contents are usually stored in the file system of the host operating system. Most document management products use this document storage method today because this approach provides the most flexibility in terms of data distribution and also allows for greater scalability.

Document indexing services 66 are applications that allow users to locate documents stored in the database. As known in the art, locating documents and content within documents is a complex problem and involves several alternative methods. The preferred document indexing services 66 provide index applications that support searching document repositories by using attribute searching, full-text searching, context searching or Boolean searching.

The preferred document security services 68 provide applications that only allow documents to be accessed exclusively through the document management backbone. The document management backbone determines how documents are accessed and who may access the documents. Preferentially, if a document is checked in, checked out, routed, viewed, annotated, archived, or printed, it should be done only by authorized users. As known in the art, document security services 68 can be designed to control access at the user, role and group levels. Those skilled in the art would recognize that the type of document security services 68 used in the present invention will vary depending upon the needs of the particular business enterprise.

The document access services 70 are applications that support document creation, maintenance, and retrieval and allow users to capture knowledge or content through the

creation of unstructured information, i.e., documents. Document access services 70 allow users to effectively retrieve documents that were created by them and documents that were created by others. As previously set forth, document files can be comprised of many different data types, including text, charts, graphics, or even audio and video.

5 Locating document files as well as content within document files, is a more complex problem and involves several alternative methods. The Windows file manager is a simplistic implementation of a hierarchical organization of files and collections of files. If the user model of where documents should be stored and found can be represented in this way, the use of structure and naming standards can be sufficient. However, a hierarchical
10 document-filing organization is not suitable for many types of document queries (e.g., retrieving all sales order documents for over \$1,000). Therefore, most document management products that may be used in the document access service 70 provide index services that support the following methods for searching repositories of documents: attribute searches, full-text searches, context searches and Boolean searches.

15 Attribute searches scan short lists (attributes) of important words that are associated with a document and returns documents that match the search criteria. Attribute search brings the capabilities of the SQL-oriented database approach to finding documents by storing in a database the values of specially identified fields within a document and a reference to the actual document itself. To support attribute search, an index maintains
20 document attributes, which it uses to manage, find, and catalog documents. This is the least complicated approach of the searching methods.

 Full-text searches are searches that scan repository contents for exact words or phrases and returns documents that match the search criteria. To facilitate full-text search, full-text indexes are constructed by scanning documents once and recording in an index file
25 which words occur in which documents. Leading document management systems have full-text search services built-in, which can be integrated directly into applications.

 Context searching is a method of searching repository contents for exact words or phrases. It also searches for related words or phrases by using synonyms and word taxonomies. For example, if the user searches for auto, the search engine should look for
30 car, automobile, motor vehicle, etc. Boolean searches are searches that scan repository contents for words or phrases that are joined together using Boolean operators (e.g., AND,

OR or NOT). The same type of indexes used for Boolean searches are used for full-text searches.

Documents should be accessed exclusively through the document services 52. If a document is checked in, checked out, routed, viewed, annotated, archived, or printed, it should be done only by users with the correct security privileges. Those access privileges should be able to be controlled by user, role and group. Analogous to record locking to prevent two users from editing the same data, document management access control services include check-in/check-out services to limit concurrent editing.

The document replication/synchronization services 72 use applications to support an environment in which multiple copies of documents are maintained. In addition, the document replication/synchronization services 72 perform the transactions required to make one or more information sources that are intended to mirror each other consistent. As known in the art, there are many ways to provide document replication/synchronization services 72, depending upon the needs of the particular enterprise using the netcentric computing system 12.

In the preferred embodiment, documents are shareable and searchable across the entire netcentric computing system 12. Therefore, the architecture needs to logically provide a single repository, even though the documents are physically stored in different locations. In some embodiments, replicating documents on alternative server(s) may have some benefits: better availability or recoverability of a distributed application, better performance, reduced network cost, or increased information access and availability.

Document synchronization services perform the transactions required to make consistent one or more information sources that are intended to mirror each other. They support the needs of intermittently connected users or sites. As with the database services 50, the synchronization services are valuable for users of remote or mobile devices that need to be able to work locally without a constant network connection and then be able to synchronize with the web server 28 at a given point in time. The terms replication and synchronization may be used interchangeably for the purpose of the present invention.

In the preferred embodiment, the document versioning services 74 use applications to maintain a historical record of the changes to a document over time. By maintaining this record, the document versioning services 74 allow for the recreation of a document as it looked at any given point in time during its evolution. Depending on the document, versioning can be accomplished using one of several methods known and presently used in

the art. Additional document versioning service 74 features may include recording who made changes, when, and why they were made.

5 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. Provisional Application Serial No: 60/156,962, entitled NETCENTRIC AND CLIENT/SERVER COMPUTING, which is herein incorporated by reference, in its entirety.

What is claimed is:

1. An information service architecture for a netcentric computing system, comprising:

at least one client connected with a server;

5 a database service including a database storage service, a database indexing service, a database security service, a database access service and a database replication/synchronization service located on said server for interacting and exchanging information with said client; and

10 a document service including a document storage service, a document indexing service, a document security service, a document access service, a document replication/synchronization service and a document versioning service located on said server for interacting and exchanging information with said client.

2. The information service architecture of claim 1, wherein said database storage service manages and physically stores data in the form of at least one database file.

3. The information service architecture of claim 1, wherein said database indexing service uses at least one database indexing application that enables said client to retrieve information from said database storage service.

20 4. The information service architecture of claim 3, wherein said database indexing application may be selected from the group of applications consisting of bitmap indexing applications, context indexing applications and star indexing applications.

25 5. The information service architecture of claim 1, wherein said database security service uses at least one database security application to enforce access control to ensure that records and data stored in said database storage service are only visible or editable by authorized clients.

30 6. The information service architecture of claim 1, wherein said database access service enables applications on said client to retrieve data from a database file in said database storage service.

7. The information service architecture of claim 1, wherein said database replication/synchronization service uses applications to make multiple copies of at least one database file in said database storage service.

5 8. The information service architecture of claim 1, wherein said document storage service manages physical storage of a plurality of document files within said document services.

10 9. The information service architecture of claim 1, wherein said document indexing service uses at least one document index application that is capable of searching a plurality of document files stored in said document storage services.

15 10. The information service architecture of claim 9, wherein said document index applications may be selected from the group of applications consisting of attribute searching applications, full-text searching applications, context searching applications or Boolean searching applications.

20 11. The information service architecture of claim 1, wherein said document security service includes at least one document security application that restricts access to a plurality of document files stored in said document storage service.

25 12. The information service architecture of claim 1, wherein said document access service provides document access applications that support document file creation, maintenance and retrieval.

30 13. The information service architecture of claim 12, wherein said document access applications may be selected from the group of applications consisting of attribute search applications, full-text search applications, context search applications and Boolean search applications.

14. The information service architecture of claim 1, wherein said document replication/synchronization service uses applications to maintain multiple copies of a plurality of document file stored in said document storage services.

15. The information service architecture of claim 1, wherein said document versioning service includes at least one version application that maintains a historical record of changes to document files stored in said document storage service.

5 16. An information service architecture for a netcentric computing system, comprising:

 at least one client connected with a server;

 a database storage service for managing and physically storing a plurality of database files on said server;

10 a database indexing service that includes database indexing applications that enable said client to retrieve information from said database files located in said database storage service;

 a database security service that uses at least one database security application to enforce access control to ensure that said database files are only visible or
15 editable by authorized clients;

 a database access service for enabling applications on said client to retrieve data from a predetermined database file in said database storage service;

 a database replication/synchronization service for making multiple copies of at least one database file in said database storage service;

20 a document storage service for managing and physically storing a plurality of document files;

 a document indexing service including document index applications that support searching said document files stored in said document storage service;

25 a document security service that includes at least one document security application that restricts access to said document files stored in said document storage service;

 a document access service that includes document access application that support document creation, maintenance and retrieval;

30 a document replication/synchronization service for maintaining multiple copies of predetermined document files stored in said document storage services;
and

a document versioning service including at least one version application that maintains a historical record of changes to document files stored in said document storage service.

5 17. The information service architecture of claim 16, wherein said database indexing applications may be selected from the group of applications consisting of bitmap indexing applications, context indexing applications and star indexing applications.

10 18. The information service architecture of claim 16, wherein said document index applications may be selected from the group of applications consisting of attribute searching applications, full-text searching applications, context searching applications or Boolean searching applications.

15 19. The information service architecture of claim 16, wherein said document access applications may be selected from the group of applications consisting of attribute search applications, full-text search applications, context search applications and Boolean search applications.

20 20. A method of providing an information service architecture in a netcentric computing system, comprising the steps of:

- providing at least one client connected with a server;
- managing and physically storing data in the form of a plurality of files in at least one database with a database storage service;
- enabling said client to retrieve information from said database storage service
- 25 with a database indexing service that includes indexing applications;
- enforcing access control to ensure that records and data stored in said database storage service are only visible or editable by authorized clients with a database security service that uses at least one security application;
- enabling applications on said client to retrieve data from a database in said
- 30 database storage service with a database access service;
- making multiple copies of at least one database in said database storage service with a database replication/synchronization service;

managing and physically storing a plurality of document files with a document storage service;

searching said document files stored in said document storage service with a document indexing service including at least one document index application;

5 restricting access to said clients to said document files stored in said document storage service with a document security service including at least one security application;

creating, maintaining and retrieving document files with a document access service that includes at least one document access application;

10 maintaining multiple copies of predetermined document files stored in said document storage services with a document replication/synchronization service; and

maintaining a historical record of changes to said document files stored in said document storage service with a document versioning service including at least one version application.

15

21. The method of claim 20, wherein said database indexing applications may be selected from the group of applications consisting of bitmap indexing applications, context indexing applications and star indexing applications.

20

22. The method of claim 21, wherein said document index applications may be selected from the group of applications consisting of attribute searching applications, full-text searching applications, context searching applications or Boolean searching applications.

25

23. The information service architecture of claim 20, wherein said document access applications may be selected from the group of applications consisting of attribute search applications, full-text search applications, context search applications and Boolean search applications.

1/2

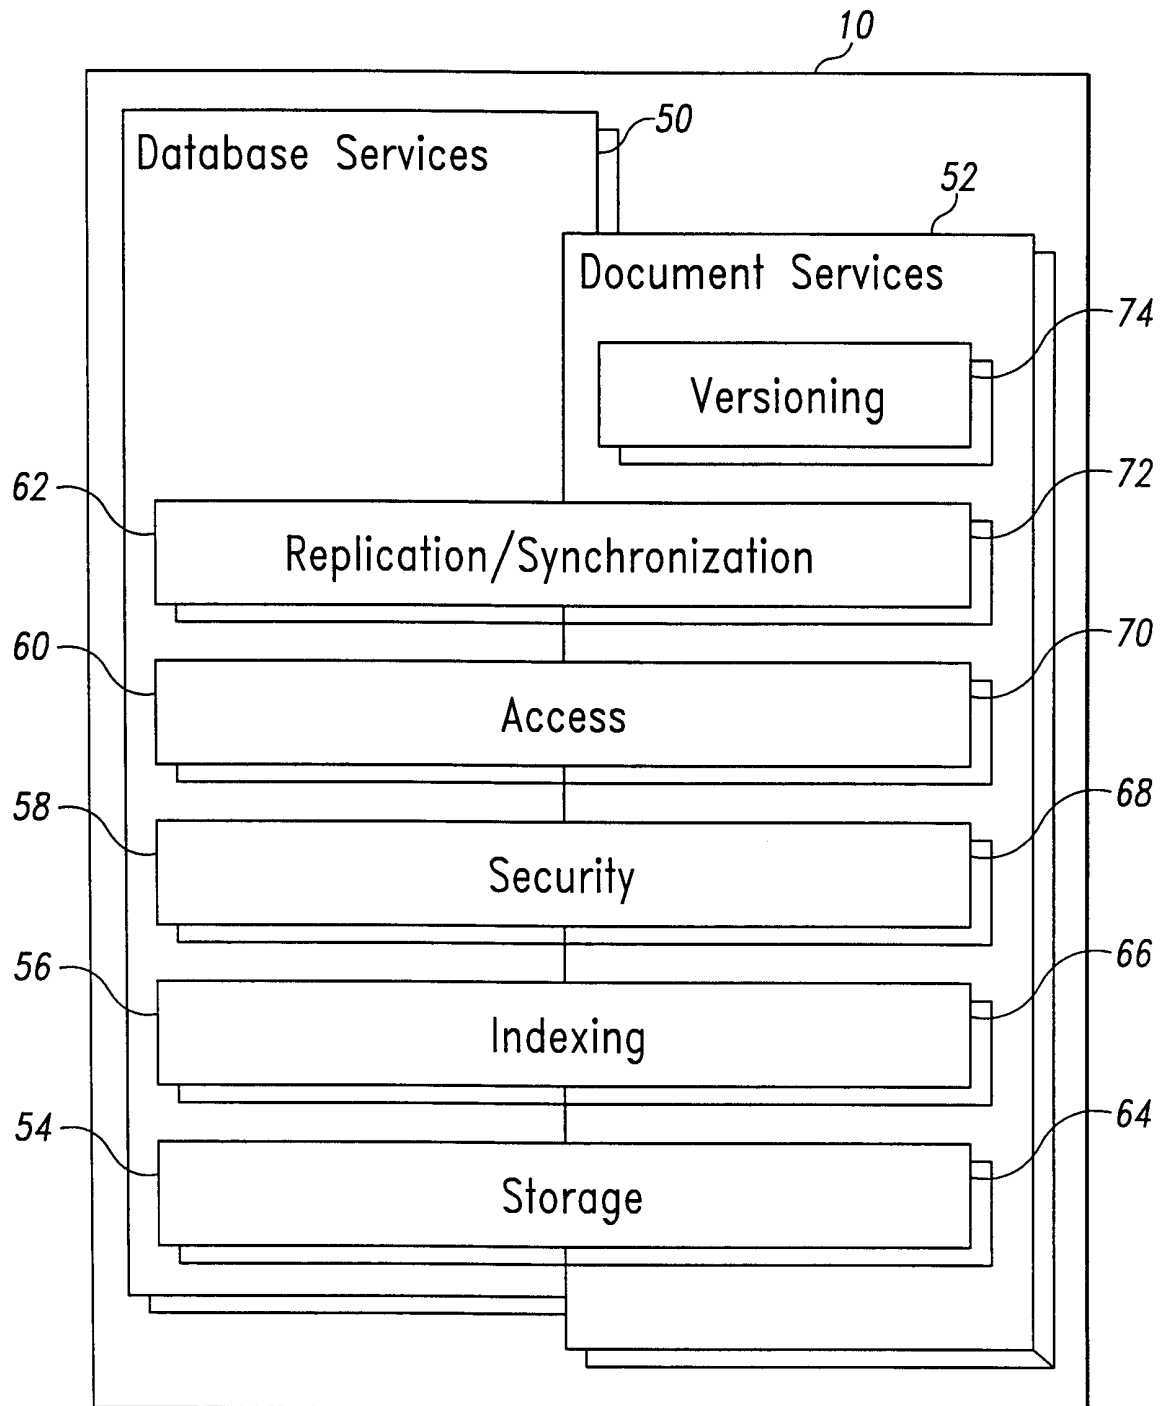


Fig. 1

2/2

Enterprise

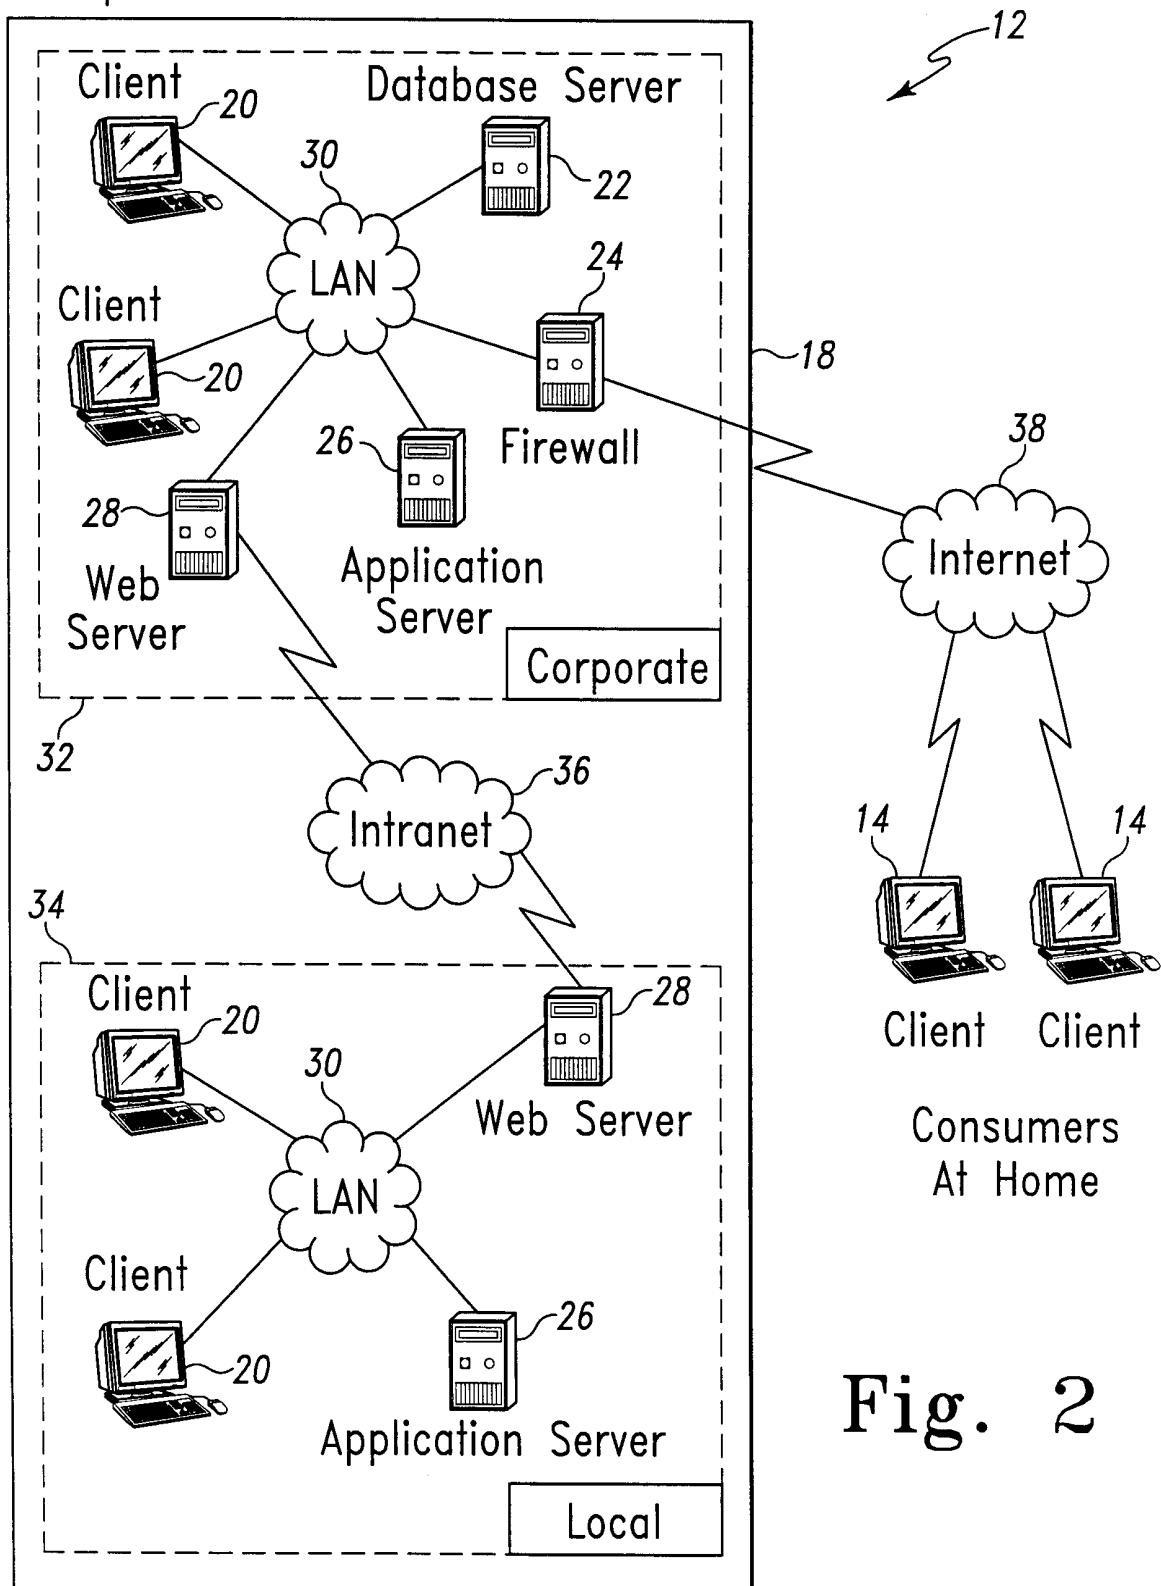


Fig. 2