METHOD AND APPARATUS FOR PROCESSING BUSINESS INFORMATION FROM MULTIPLE ENTERPRISES

BACKGROUND OF THE INVENTION

Field of the Invention

This invention generally relates to database systems for computers and more particularly to a method and apparatus for processing business information from multiple enterprises.

Description of the Related Art

In recent years, businesses have focused on creating information systems to analyze business operations and increase profitability. These information systems access business related information stored in a collection of databases called a data warehouse. The business information in the data warehouse includes data on a wide range of business activities such as inventory, sales records, shipping costs, account receivables, and employee salaries. Data warehousing was a first attempt to make the vast amount of information associated with running a business available for analysis and study.

Even though the data warehouse makes information available, many managers and executives have difficulty applying this vast amount of information to everyday business needs. Often, the volume of data in most data warehouses simply overwhelms the average business user. Consequently, managers, executives, and others must rely on specially trained analysts to sort through information in the data warehouse. Even the analysts may rely on special software tools and techniques to sort through the data warehouse and produce information useful to a business.

Data marts were created to help deal with the vast amount of information found in data warehouses. The data mart improved data warehousing by grouping portions of a data warehouse into logical subsets of information. This reduced the amount of information presented to managers and executives and made the information easier to manage. Using the data mart, a manager or executive could use focus on information related to a particular product or service being analyzed rather than the information related to all products or services provided by a company.

Despite these advances, many business people continue to have difficulties accessing and using information stored using the data warehouse and data mart technologies. First, data warehouses and data marts do not analyze information and provide users with results for making business decisions. Even with the data mart, the users must analyze information contained in these storage areas and generate their own conclusions. Consequently, many managers and executives either train themselves quickly or rely on trained analysts to analyze the information.

Second, existing data warehouse and data mart technologies are limited to information from a single enterprise rather than a combination of multiple enterprises. In most cases, business information collected by a single enterprise (i.e. an independent business unit) cannot be shared with other enterprises because their information systems and software collecting the information is incompatible. For example, enterprises may use different enterprise resource planning (ERP) systems developed by different vendors such as Bann, Oracle, SAP, and PeopleSoft to collect business information. These incompatibilities restricts the business information to a single enterprise and limit the type of analysis that can be done. For example, without access to the information systems for a sales organization of a company, the manufacturing division of the company would not be able to analyze order requirements based on recent sales.

Moreover even when compatible information systems and software are used, there may be sensitive or proprietary information an enterprise may not want to share with other enterprises. For example, a supplier of a raw materials may not want to indicate to a buyer of the raw material the markup on the material they are charging the buyer. Current information systems are not designed to limit sharing of business information in this manner.

Accordingly, there is a need to share business information generated from multiple enterprises and provide analysis that can be used in businesses.

SUMMARY OF THE INVENTION

In accordance with the present invention, as embodied and broadly described herein, a method for processing business information generated by multiple enterprises. Initially, a data warehouse is provided that has the capability of holding business information. The data warehouse is populated with business information received from a first enterprise and a

second enterprise. The business information is associated with the first and second enterprise based upon a set of standardized categories used in the data warehouse. Metastore information describing the organization of business information in the data warehouse is used to develop rules for extracting a portion of the business information from the data warehouse. This portion of business information extracted from the data warehouse is then stored in a data mart. Using at least one such data mart, a multiple dimension database is created wherein each dimension of the multiple dimension database corresponds to variables derived from business rules established in an industry. Using the multiple dimension database, metrics for measuring business performance from the multiple dimension databases are then generated.

Another aspect of the present invention includes a method performed on a data processing system for combining business information generated by multiple enterprises. Initially, a data warehouse is provided that has the capability of holding business information received from a first enterprise and a second enterprise. Business information associated with the first and second enterprise is categorized using a set of standardized categories used in the data warehouse. The set of standardized categories are then used to store the business information received from the first and second enterprises in the data warehouse.

Yet another aspect of the present invention includes a method for creating data marts from the business information associated with multiple enterprises. Initially, a data warehouse includes business information from a first and a second enterprise stored in a data warehouse in accordance with a set of standardized categories. Metastore information, describing the organization of business information in the data warehouse, is then used to develop rules for extracting business information from the data warehouse. These rules are then used to extract a portion of the business information from the data warehouse and store it in a data mart.

In another aspect of the present invention, a method for generating metric values for evaluating business performance from the business information of multiple enterprises is provided. Initially, business information from a first and a second enterprise is received into a data warehouse in accordance with a set of standardized categories. Next, a portion of the business information is extracted from the data warehouse and stored in at least one data

mart. Using one or more data marts, a multiple dimension database is created wherein each dimension of the multiple dimension database corresponds to at least one variable derived from business rules established in an industry. These multiple dimension databases are then used to generate metrics for measuring business performance.

BRIEF DESCRIPTION OF THE DRAWINGS

The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate an implementation of the invention and, together with the description, serve to explain principles of the invention.

IN THE DRAWINGS:

FIG. 1 is a block diagram representation of multiple enterprises sharing information in accordance with one implementation of the present invention;

FIG. 2 is a block diagram of metastorage server implemented consistent with the present invention and used to facilitate the exchange of business information;

FIG. 3 is a block diagram representation of a transform module designed in accordance with one implementation of the present invention;

FIG. 4 is a block diagram representation of metadata tools designed in accordance with one implementation of the present invention;

FIG. 5 is a block diagram of data mart modules designed in accordance with one implementation of the present invention;

FIG. 6 is a block diagram of metric modules used to generate business performance metrics in accordance with one implementation of the present invention;

FIG. 7 is a flowchart of the method steps associated with generating business performance metrics from multiple enterprises in accordance with one implementation of the present invention;

FIG. 8 is a block diagram representation of the user interface generated in accordance with one implementation of the present invention;

FIG. 9 is a block diagram representation of the organization selection portion of the user interface as depicted in FIG. 8;

FIG. 10 is a block diagram representation of the location selection portion of the user interface as depicted in FIG. 8;

FIG. 11 is a block diagram representation of the supply chain and the demand chain of the user interface as depicted in FIG. 8; and

FIG. 12 is a flowchart of the method steps associated with creating the supply and demand chain from multiple enterprises in accordance with one implementation of the present invention.

DETAILED DESCRIPTION

Reference will now be made in detail to an implementation of the present invention as illustrated in the accompanying drawings. The same reference numbers will be used throughout the drawings in the following description to refer to the same or like parts.

MULTIPLE ENTERPRISE SYSTEM

FIG. 1 is a block diagram representation of multiple enterprises sharing information in accordance with one implementation of the present invention. The enterprise, in this context, ca be any independent business entity that supplies or uses a product or service from another business entity. These enterprises may be part of a larger corporation or may be separate business organizations. By sharing business information, each enterprise can measure their business performance in light of the business information from other enterprises. This can help make more accurate business forecasts in markets where the supply and demand for products or services varies rapidly.

Multiple enterprise configuration 100 includes enterprise 102 and enterprise 104. Enterprises 102 and 104 may exchange products and services in the normal course of business. For example, enterprise 102 may provide products to enterprise 104 and utilize services provided by enterprise 104. Enterprise 104 may modify or improve products generated by enterprise 102 and send them to other enterprises for further processing. Enterprise 102 is considered to be part of a supply chain of enterprises or simply the supply chain because enterprise 102 is supplying products and services to enterprise 104. Similarly, enterprise 104 is considered to be part of the demand chain of enterprises or the demand chain because enterprise 104 demands products or services from enterprise 102. The supply chain and demand chain are discussed in further detail below in conjunction with operation of the user interface. Further, although only two enterprises are illustrated in FIG. 1, implementations of the present invention can be adapted to work with any number of enterprises.

In FIG. 1, enterprise 102 is operatively connected to local transaction system 108, local transaction system 110, local transaction system 112, metastorage server 114, network 116, and client system 118. Similarly, enterprise 104 is operatively connected to a local transaction system 120, a local transaction system 122, a local transaction system 124, a metastorage server 126, a network 128, and a client system 130. Communication between metastorage server 126 and metastorage server 114 facilitates the sharing of business information generated by enterprise 104 and enterprise 102, respectively.

Enterprise 102 generates business information collected by local transaction system 108, local transaction system 110, and local transaction system 112. These local transaction systems use computer systems running software such as enterprise resource planning (ERP) software designed to collect business information associated with the respective enterprise. For example, the ERP software may collect information such as manufacturing data, information on the suppl of raw materials, sales information, and inventory control from the enterprise.

Information systems associated with enterprise 104 operate in a similar manner to the information systems associated with enterprise 102. That is, local transaction systems are used to collect business information about the enterprise. For example, local transaction system 120 may be used to collect business information related to manufacturing while local transaction system 122 may be used to collect business information associated with inventory. Unfortunately, the local transaction systems associated with enterprise 104 may execute different software than the local transaction systems associated with enterprise 102 and consequently may produce information that is incompatible across the software systems. Even within the same enterprise, such as enterprise 104, local transaction system 120 may execute software that is incompatible with software executing on local transaction system 122. This incompatibility between software systems could arise if some local transaction systems executed ERP software developed by BANN/SAP while other local transaction systems executed software developed by Oracle or People Soft.

Metastorage server 126 combines the business information from these aforementioned local transaction systems with business information from enterprise 102 using metastorage server 114. For example, sharing business information allows managers and executives associated with enterprise 104 to carefully monitor information on product

and services supplied by enterprise 102. Similarly, managers and executives at enterprise 102 can track the demand for goods and services from enterprise 104 to assist in managing their manufacturing and production processes as well.

Metastorage server 114 and metastorage server 126 use meta-information to describe the organization of data in each of the local transaction systems and solve the incompatibilities between the systems. The meta-information maps data categories in each local transaction system to standard categories describing products, services, and other business information. These standard categories then enable metastorage server 114 and metastorage server 126 to access and extract information from local transaction systems executing different software applications. The combined information is used to generate databases, data warehouses, and dat marts in accordance with the present invention. In addition, the metastore servers also use meta-information to generate metrics for measuring business performance.

FIG. 2 is a block diagram of metastorage server 114 designed in accordance with one implementation of the present invention. Metastorage server 114 includes a memory 202, a processor 204, a secondary storage 206, an input/output device 208, a display 210, an input device 212, and a network interface 214. Processor 204 can be a central processor unit (CPU) capable of executing the instructions associated with software modules stored in memory 202. Processor 204 may also be implemented as an application specific integrated circuit (ASIC) designed to execute processes implemented in accordance with one embodiment of the present invention such as those stored in memory 202.

Secondary storage 206 is a storage area such as a hard disk or cache memory having a file system 232 and a database 234. File system 232 can be used to store files and other data usually associated with operation of metastorage server 114. Likewise, databases 234 includes libraries and application programming interfaces (API) useful in developing databases and extracting data from the databases. Input/output device 208 may include scanners, printers, and mass storage devices such as tape drives. Display 210 includes any display device capable of displaying information such as a monitor, a television, or a liquid crystal display (LCD). Input device 212 can be used to control operation of metastorage server 114 and may include a keyboard, a mouse a keyboard and mouse combination, and voice recognition input systems. Network interface 214 connects metastorage server 114 to

a network such as the Internet and provides a communication protocol suitable for transmitting information.

Memory 202 includes an operating system 230, a database engine 228, a transform module 226, metadata tools 224, a data warehouse module 222, a metadata module 221, a data mart module 220, a business performance metrics module 218, an application module 216, a network module 215, and a user interface 213. Operating system 230 schedules execution of software modules in memory 202 on processor 204 and manages other computer resources associated with metastorage server 114. Database engine 228 processes database requests using database languages such as structured query language (SQL) and accesses data stored in database 234. In one implementation, database engine 228 communicates with database 234 using industry standard relational database management system (RDBMS) software. Alternatively, database engine 228 and database 234 may utilize object-oriented protocols.

TRANSFORM MODULE

Transform module 226 acquires business information from local transactions systems running databases and software applications such as ERP software. For example, metastorage server 114 initially uses transform module 226 to extract information from ERP systems according to the particular ERP system design. SQL statements and templates are used to describe the data structures in the different ERP systems and control the extraction of business information. A template contains the names, descriptions, and data types of every data element in the abstract file. A template for an inventory system, such as an inventory system designed by BAAN or SAP, may map each of the tables and fields used in the inventory system to a neutral format data file. The neutral format data file contains standardized row-and-column database tables compatible with databases 234 and database engine 228.

FIG. 3 is a block diagram illustrating software modules associated with transform modul 226. Transform module 226 includes native interface programs 302, industry metadata dictionaries 304, industry data warehouse models 306, and industry and subject component models 308.

Native interface programs 302 are compatible with the application programming interfaces (APIs) for the different software applications executing on the local transaction

systems within an enterprise. These native interface programs 302 can extract information from the databases and software applications executing on the local transaction systems. For example, native interface program 302 may use database macros and SQL routines proprietary to a specific database or ERP system to extract the name, birth date, address, social security number and other information stored on the system.

Data associated with native interface programs 302 are then transformed into a neutral data format for further processing. The neutral data format defines proprietary database structures, data types, and data fields from different transaction systems into a single set of database structures, data types, and data fields. Industry data dictionaries 304 provide a set of standard industry definitions to map between the data in a software application and this neutral data format. By adopting a standard way of describing business information, enterprises can exchange information about their business in a uniform manner. For example, these dictionaries are used to map proprietary field entries of an ERP system into standard definitions and terms. Advanced Data Resource Management (ADRM) is one data dictionary developed for a wide range of industries to define products, services, and terms used in each industry and may include up to as many as 8,000 definitions. ADRM is available from Harmony, Inc., 107 South B Street, San Mateo, California 94401 and is herein incorporated by reference in its entirety.

Industry data warehouse model 306 describes the database tables and structure for storing business information in a data warehouse once it has been extracted and transformed to a neutral format. These industry data warehouse models 306 use specific entries from industry data dictionaries 304 to create the database tables, indices, and other database structures for the data warehouse.

Industry and subject component model 308 describe those portions of the database structure that are shared between different industries and those portions of a database that are shared within an industry. Specifically, a subject component model describes the common database structures used in a single industry. In contrast, industry component models are database structures that are common between different industries. For example, portions of a database designed to hold sales and marketing information would use industry component models since many industries have sales and marketing organizations. However, portions of a database designed to hold regulatory procedures for the

pharmaceutical industry would use subject component models since other industries, such as the electronics industry, do not require such information.

MAINTENANCE TOOLS

Referring to FIG. 2, maintenance tools 224 are used to manage user accounts and business information extracted from various local transaction systems. Maintenance tools 224, represented by the block diagrams in FIG. 4 include a data quality analyzer 402, and a metadata manager 404, a reference data editor 406, a personalization manager 408, and a systems and security administration module 410.

Data quality analyzer 402 analyzes business information extracted from the local transaction system in accordance with validation rules and generates a list of nonconforming records. A system administrator is responsible for correcting errors that may exist in these nonconforming records. For example, validation rules in data quality analyzer 402 may search for records with null entries or empty values and flag them to the system administrator as record requiring modification. Because null values may not be allowed in a given record, the system administrator may need to enter valid data to correct the records in the database.

Metadata manager 404 allows the system administrator to manage information extracted from the local transaction systems in the multiple enterprises. For example, metadata manager 404 can be used to modify how information stored in a local transaction system for an enterprise is converted to corresponding entries in databases 243. This can be used to make customized mappings between databases or adjust weights and measurements for transactions made between different countries.

To modify data extracted from the local transaction systems, the system administrator uses a reference data editor 406. Reference data editor 406 provides access and editing capabilities to information in a structured format. For example, reference data editor 406 allows the system administrator to directly edit information stored in records as well as modify relationships between records in the database.

Personalization manager 408 customizes the display of information made to a manager or executive accessing metastorage server 114. Personalization can be done based on the manager or executive's role in a company or the persons specific preferences. For each role in a company certain business information may be displayed or restricted from

display as determined by the system administrator. This enables people working in the different roles in the company to focus on business information that most directly affects their underlying responsibilities. It also can b used to limit sensitive information such as financial summaries to those people entrusted to manage financial aspects of a business. For example, a person serving as the vice president in a company may receive revenue forecast metrics, inventory in transit metrics, and total inventory metrics on a display screen. Other information not directly related to a vice president's responsibilities may be restricted or eliminated from being displayed on the vice president's display screen. By controlling information in this manner, metric values, business reports, and other information can uniformly be distributed throughout the company to various managers based upon their position and responsibilities in the organization.

in addition to their role, each person in the company may also have information displaye based on their individual requirements. This type of personalization allows the user to specify personal preferences such as language, currency, and time zone when information is being generated on a display screen. For example, a vice president of finance can further personalize his information requesting display of information in English, using United States dollar currency and in the United States Central Standard Time zone.

Systems and security administration 410 is used to control access to business information generated by metastore server 114. Like personalization, security administration is based upon a user's role in a company. Security includes an authentication and an authorization process. Initially, the user must authenticate that they are a valid user by entering a valid user ID and password combination. An authenticated user is then authorized to access business information based on the role established for him by the system administrator.

A user's authorization is established as part of the configuration process when setting up metastorage server 114. The authorization process allows certain users to access certain types of information depending upon their role in the company. For example, a vice president of worldwide manufacturing may have access to all revenue models generated worldwide for the company. In contrast, a vice president of manufacturing responsible for only a particular country may only be authorized to access business information related to manufacturing in that corresponding country.

DATA WAREHOUSE MODULE

Referring to FIG. 2, data warehouse module 222 represents the database used to store the business information collected from local transaction systems used in the multiple enterprises. Although the information in data warehouse module 222 may be too detailed for average users, is quite useful for deriving smaller databases such as data marts or generating metrics for measuring business performance. Each of these derivatives, which are discussed in further detail below, are more useful to high-level users such as business managers and executives. In accordance with the present invention, information in data warehouse module 222 is used to create more refined information such as data marts or metrics that users can use more easily.

METADATA MODULE

Metadata module 221 includes information used to describe the format and structure of information stored in data warehouse module 222, data mart module 220, and business performance metrics module 218. Any information used to describe other information can be classified as metadata and may be included in metadata module 221. This includes information that defines units of measure, security information, authentication and authorization methods for different user roles, and the design of different databases. Metadata module 221 can also include information on generating data mart module 220 from data warehouse module 222 and metrics module 218 from information in data mart module 220. Further, information in metadata module 221 may also depend on the particular industry and the way information is organized within the particular industry. For example, metadata module 221 may describe information in the pharmaceutical industry using metric weights and measurements and describe information in the United States automotive industry using English weights and measurements. Details on data marts and metrics are described in further detail below.

DATA MART MODULE

Data mart module 220 provides a subset of information in data warehouse module 222 that relates to a particular product or service provided by enterprise 102. This subset of information from data warehouse module 222 used to generate data mart module 220 is more easily searched and provides faster results. Metadata information in metadata module 221 includes rules for generating the desired data mart module 220 from information stored

in a data warehouse for a given industry. For example, metadata information can be used to create a data mart that lists the capital equipment and corresponding costs for creating a semiconductor chip. The metadata information in metadata module 221 would assist in identifying all capital equipment used to create the semiconductor chip using the business information stored in data warehouse module 222 from multiple enterprises. By consolidating and summarizing information stored in data warehouse module 222 into data mart module 220, users can focus on specific information, analyze the relevant variables, and ignore information unrelated to the particular analysis.

FIG. 5 provides a block diagram of the elements within data mart module 220. Data mart module 220 includes industry metastore information for data marts 502, aggregate subject datamart models 504, and an industry based data mart 506. To create the data mart, industry-based metastore information for data marts 502 indicates which fields in the data warehouse should be extracted for consolidation. Although described as part of data mart module 220, industry metastore information for data marts 502 could also be associated with metadata module 221 since it contains information used to describe other information.

Aggregate subject datamart models 504 indicates which field selected from the data warehouse can be combined together and presented as a single value. For example, aggregate subject datamart models 504 may be used to summarize sales revenue by region, product line, or division. By reducing worldwide sales to region, a business manager could quickly determine which regions in a business were producing more revenue than other regions. Similar analysis could be applied to a specific product line in a company or division within a company. The result of this process, creates industry-based data marts 506 containing summarized information originally contained in the data warehouse.

BUSINESS PERFORMANCE METRICS MODULE

Business performance metric module 218 provides benchmark measurements for the performance of a business. These benchmarks are tailored to the particular industry and are generated from the consolidated business information stored in data mart module 220 and, in some cases, from data warehouse module 222. Metrics within business performance metrics module 218 are tailored to meet the individual needs of managers and executives and correspond to their role in the organization. For example, one metric useful in the high volume low profit margin electronics business relates to inventory turnover of materials. By

tracking the inventory turnover metric, executives or business managers can readily manage the costs for keeping inventory, reduce the costs of goods sold, and improve profits.

FIG. 6 is a block diagram of components used in the business performance metric modul 218. Business performance metric module 218 includes industry metric components 602, metric metastore content definitions 604, and industry metric data 606. Industry metric components 602 are standards set up in the industry for measuring performance of businesses in the industry. For example, industry metric components 602 in the pharmaceutical industry includes standards for measuring tested purity of drugs and overall quality as set up by the pharmaceutical standards group or regulatory agencies such as the Federal Drug Administration (FDA). For auditing and accounting purposes, industry metric components 602 may include metrics derived from well-known accounting and auditing standards such as generally acceptable accounting principles (GAAP). Other metrics may be derived from standards or rules developed by consortiums, regulatory industry organizations, or governmental bodies such as the Institute of Electrical and Electronic Engineers (IEEE), American Production and Inventory Control Society (APICS), the Securities Exchange Commission (SEC), and the Internal Revenue Services.

Metric metastore content definitions 604 describes which information in the data mart should be used to generate appropriate metric information. Essentially, metric meta store content definition 604 assists in extracting information from the data mart. For example, metric meta store content definition 604 could tell the system to extract business information on a product and all the different equivalents of the product being manufactured at different locations in the world. In the case of a product like aspirin, metric meta store content definition 604 would be used to identify the different names for aspirin in different countries, collect information on each named product, and generate appropriate metrics for managing the manufacture and other aspect of the product. Once information is gathered, the necessary metric values are generated for the industry and collected in industry metric data 606 for further analysis. Typical industry metric data 606 may include information such as revenue forecasts, total inventory on hand, and inventory in transit.

Although described as part of data mart module 220, industry metastore information for data marts 502 could also be associated with metadata module 221 since it contains information used to describe other information.

APPLICATION MODULE

Referring to FIG. 2, application module 216 organizes information for display and presentation on display 210. Application module 216 acts as a layer between the actual user interface and various portions of the database and displays information tailored to a person's role in a business and personal preferences. As previously discussed, a person's role in a business dictates how much or little business information they are allowed to view. With respect to personal preferences, application module 216 may include a preferred language, currency, and time zone.

NETWORK MODULE

Network module 215 transports business performance metrics, graphics and other information over a network to a manager or executive's desktop. This information is secured an encrypted to keep information confidential when transmitted over the network. Network module 215 may use a conventional Internet server such as the Apache server, an extensible markup language (XML) based server, an E-commerce server such as the Commerce Interchange Pipeline (CIP), Electronic Data Interchange (EDI) server or other similar network based servers.

DATAMARTS AND METRICS FROM MULTIPLE ENTERPRISES

FIG. 7 is a flowchart diagram indicating the steps associated with creating data marts and metric values from the data warehouse information associated with one or more enterprises. Initially, a logical data model is created for the particular industry being studied (step 702). The logical data model includes all the names, descriptions, and data types of data elements required to describe business information in the particular industry. Logical data models may also describe how information is organized in a data file, the file naming convention used to store the data file on the computer, and the frequency for collecting information. Although each industry has a different logical model, certain portions of the logical data model may be shared or reused if several industries track similar types of information. For example, portions of the logical module related to marketing and sales may be shared between a wide range of industries.

For example, the categories information described in logical data models can be defined using industry data dictionaries 304 in FIG. 3. These definitions categorize and standardize products and services in a particular industry to facilitate uniform

implementation of databases and information systems. Enterprises adopting appropriate industry data dictionaries 304 can later exchange business information with each other. ADRM dictionaries, discussed above, are one type of industry data dictionaries 304 useful for describing the logical data models for an industry.

The logical data model is then converted into physical data models for different types of databases (step 704). For example, the logical data model can be used to generate different database schema programs for different relational database management systems (RDBMS) such as Oracle, Informix, Sybase, or Ingres. The database schema program uses database commands based on structured query language (SQL) and various programming languages to describe the organization and structure of a database. Different database schemas are generated because each RDBMS system may use different commands and syntax to create related database structures. In accordance with the present invention, industry data warehouse models 306 in FIG. 3 are used to create the structures within a database that hold a data warehouse and includes both the logical and physical data models discussed above.

A derivative of the physical data models are then used to create the structures within the database for a data mart (step 706). The data mart may consolidate information stored in the data warehouse or may contain a subset of information contained in the data warehouse. Accordingly, database programs using SQL and other programming languages use portions of the physical data model to generate the structure for data marts in the database.

For example, the initial physical model is created by invoking a database application from a normalized schema. In accordance with the present invention, the normalized schema creates tables, fields, and data types corresponding to the data warehouse. A star schema is used to derive the data mart structure from the data warehouse. A fact table generated by the star schema includes selected dimensions from the data warehouse and partially consolidated information useful in the multi-dimensional data mart. Surrounding the fact table, thus forming a star, are selected tables from the data warehouse having the information used in the data mart. The star schema facilitates viewing multiple dimensions of a database from different perspectives. For example, a star schema can be used to derive a fact table having primary keys corresponding to a product table, a supplier table, and a purchase order table. The fact table thus facilitates viewing permutations of information

along any of these dimensions or fields within each dimension. A product based view could display each product and the corresponding manufacture dates, geographic information or corporate organizations associated with the product.

Once the data warehouse and data mart structures are created, business information from the local transaction systems associated with multiple enterprises is extracted and stored in a data warehouse (step 708). In this context, an enterprise can be any independent business unit collecting business information that cannot be shared directly with another enterprise. Local transaction systems can be associated with enterprises from different companies or enterprises corresponding to the subdivisions with a single corporate entity. In accordance with the present invention, native interface programs 302 in FIG. 3 can be used to access the local transaction systems associated with these different enterprises. These native interface programs 302 include the appropriate database commands and applications necessary to extract information from different databases and ERP systems developed by vendors such as Oracle, Informix, BAAN, SAP, and PeopleSoft.

Business information is then extracted from the data warehouse and stored within the structure of data mart 506 (step 710). Extracting the information may include consolidating information, aggregating information, and converting between different units of measurement or monetary currencies. Industry metastore information for data marts 502 includes rules to help extract the appropriate information from the data warehouse. Further, aggregate subject data mart module 504 helps determine which information in the data warehouse can be aggregated. For example, daily sales figures can be aggregated by products, customers, and time to reduce the amount of raw data and make the information easier to analyze.

Next, multi-dimensional data tables of information are created from the data marts in accordance with established business rules (step 712). These multi-dimensional data tables are sometimes referred to as "cubes" of infonnation because the information can be analyzed from multiple different dimensions. As previously mentioned, the business rules correspond to business rules established by industry standards groups, governmental bodies, or consortiums. Variables used in these multi-dimensional tables may also depend on the person's role in the company and their business responsibilities. For example, a 3-

dimensional data table for a vice president of operations in a company may include variables such as units produced, marginal cost of units produced, and profit margin.

Business rules used to create multi-dimensional data tables are described in industry metric components 602. Industry metric components 602 includes information used to derive industry metric values such as "in transit inventory". The multi-dimensional data table or cube holding the "in transit inventory" metric value would describe the passing of inventory ownership (from either the seller or buyer) as well as information on how the goods are delivered (i.e. ship, train, or plane).

Metric values are then generated from the multi-dimensional data tables that measure business performance (step 714). These metric values summarize large amounts of information from multiple enterprises that an executive or manager can readily measure and track. Using metric values rather than data, managers and executive can quickly assess business performance and take appropriate corrective actions. For example, the manufacturing and distribution industry can measure business performance using metric values such as total inventory value, purchase order stability ratio, supplier order file rate, warehouse inventory accuracy, and customer satisfaction index (i.e. compares various factors in one industry against several other industries). The purchase order stability ratio, for example, indicates over time that a product is purchased consistently and used as planned. These metric values are displayed in accordance with the user interface discussed later herein.

USER INTERFACE TO METASTORAGE SERVER

In accordance with the present invention, user interface 203 controls the processing and display of business information on a user's desktop. The user interface can be implemented within a web browser utilizing hypertext markup language (HTML), JavaScript, Java, or other programming languages. FIG. 8 is a block diagram representation of a user interface 800 used to control the display of information associated with metastorage server 114. In accordance with one implementation of the present invention, user interface 800 includes an alert reports indicator 802, a flash report selection 804, an organization selection 806, a location selection 808, a product selection 810, a score card selection 812, a supply chain selection 814, a demand chain selection 816.

Alert reports indicator 802 monitors when a metric value associated with a product or service exceeds or drops below specified limit values. These specified limit values can be set by default or set according to a user's requirements. The metric values are good indicators of overall business activities because they are derived from the business information of multiple enterprises. Alert reports indicator 802 can monitor any one of the many metrics provided in industry metric data 606. For example, metrics used in a manufacturing industry may include monitoring missed shipments, a cash balance alert indicating that a company's cash balance is low, product shortages, and product returns. All these metrics are monitored using alert report 802 because they need to be addressed immediately.

Flash reports selection 804 provides continuous monitoring for a set of metric values. The metric values monitored using flash reports selection 804 are displayed continuously to the user for easier monitoring. The metric values associated with flash reports selection 804 can default to certain values based upon a users role in an enterprise or may be customized to each individual person. For example, a vice president of sales may review metric values relating net profit while a line manager manufacturing products in a factory may review metric values relating to total manufacturing capacity.

Organization selection 806, location selection 808, and product selection 810 provide techniques to isolate portions of an enterprise for analysis. Metric values and other information displayed thereby correspond to the portion of the enterprise selected using these selection items.

Product selection 810 identifies a portion of the enterprise associated with the production of a particular product. For example, selecting production selection 810 may generate a list of products associated with the enterprise. These products can be produced in different geographic areas or from different business units within the same enterprise. The user can select a product from the list and then analyze business metrics within the enterprise associated with the particular product.

Score card selection 812 displays metric values associated with the above selection methods on user interface 203. For example, score card selection 812 will display metric values measured for a particular region or a particular product in the region if the selection methods on user interface 203 have been used to select a product or particular region.

Supply chain selection 814 displays information associated with enterprises "supplying" products and services to create a selected product. Products displayed in the supply chain can also be selected and will display business information and metric values associated with the particular product. Similarly, demand chain selection 816 displays the demand chain of enterprises or products "demanding" a particular product. Enterprises in the demand chain can also be selected to reveal metric values and business information associated with the particular enterprise. These metric values and business information can be used to gauge demand for a product and determine which enterprises are using the product.

FIG. 9 is a block diagram representation of the selectable elements used by organization selection 806 in FIG. 8 to identify a portion of an enterprise for analysis. These selectable elements on the user interface correspond to a corporate selection 902, a division selection 904, and a business unit selection 906. Corporate selection 902 is used to select different corporate entities within an enterprise; division selection 902 is used to select a division within an enterprise; and business unit selection 906 is used to select a business unit 906 within an enterprise. The terms used to distinguish business divisions within an organization are exemplary and may differ depending on the organization. Consequently, in an alternative embodiment the selection categories corresponding to organization selection 806 may differ fron those depicted in FIG. 9 and typically will correspond to the structure of the particular organization.

FIG. 10 is a block diagram representation of the selectable elements used to isolate a geographic area for analysis. These selectable elements correspond to location selection 808 on user interface 800 and include a region selection 1002, a country selection 1004, a county selection 1006, and a site selection 1008. Using these location based selections, managers and executives can compare business performance metrics of an enterprise in different geographic areas.

FIG. 11 is a block diagram representation of a value chain displayed in accordance with one implementation of the present invention. The value chain includes a supply chain and a demand chain. Depending on whether a user selects supply chain selection 814 or demand selection 816, corresponding information from either the supply chain or the demand chain portion of the value chain is displayed.

The exemplary supply chain in FIG. 11 shows parts and subassemblies supplied by different enterprises to make a particular product. In this example, selecting a part or subassembly in the supply chain produces additional business information and metric values associated with supplying that particular part or subassembly. Additional business information could include shipping dates, inventory in stock, volume purchase costs, and physical specifications.

The demand chain in FIG. 11 shows distributors, retailers and a wholesaler that demand use of the product. The demand chain operates in a similar manner to the supply chain. For example, selecting a distributor in the supply chain would produce additional business information and metric values associated with the distributor. This information could include volume requirements, desired price, desired shipping arrangements, existing volume purchase arrangements, and physical specifications required for the identified product.

FIG. 12 is a flowchart diagram of the steps used to display a supply chain and demand chain on a user interface in accordance with one implementation of the present invention. Initially, a product is identified in an enterprise for analysis (step 1202). Typically, a user will select a product for analysis. Business information and metric values associated with the identified product can be displayed by selecting the identified product on the user interface.

To display the supply chain of products, metastorage server 114 detennines supply products used to create the identified product (step 1204). Often, the identified product is associated with one enterprise and the supply products are associated with another enterprise. Consequently, metastorage server 114 may communicate with metastorage server 126 in accordance with the present invention to receive business information and metric values associated with the supply products. The business information and metric values related to the supply products are then displayed from the supply chain of information (step 1206).

In addition to the supply information, a user may also produce information associated with the demand for the identified product. To identify demand chain information, metastorage server 114 first determines the products and enterprises using or demanding the identified product (step 1208). Metastorage server 114 may communicate

with metastorage server 126 to receive business information and metric values associated with the enterprises and products in the demand chain. Once the demand products are identified, metastorage server 114 displays the business information and metric values related to the demand products and enterprises on the user interface.

While specific implementations have been described herein for purposes of illustration, various modifications may be made without departing from the spirit and scope of the invention. For example, implementations of the present invention discusses metric values and business information related products however alternate implementations of the present invention could readily be adapted to services or a combination of services and products. Furthermore, although aspects of the present invention are described as being stored in memory and other storage mediums, one skilled in the art will appreciate that these aspects can also be stored on or read from other types of computer-readable media, such as secondary storage devices, like hard disks floppy disks, or CD-ROM; a carrier wave from the Internet; or other forms of RAM or ROM. Accordingly, the invention is not limited to the above described embodiments, but instead is defined by the appended claims in light of their full scope of equivalents.

WHAT IS CLAIMED IS:

1. A method performed on a data processing system for processing business information generated by multiple enterprises, comprising:

providing a data warehouse having the capability of holding business information;

receiving business information from a first enterprise;

receiving business information from a second enterprise;

associating business information from the first and second enterprise with a set of standardized categories used in the data warehouse;

storing the business information received by the first and second enterprises in the data warehouse in accordance with the association of business information and the set of standardized categories;

using metastore information describing the organization of business information in the data warehouse to develop rules for extracting business information from the data warehouse;

extracting a portion of the business information from the data warehouse in accordance with the rules;

storing the portion of business information extracted from the data warehouse in a data mart;

creating a multiple dimension database from at least one data mart wherein each dimension of the multiple dimension database corresponds to variables derived from business rules established in an industry; and

generating metrics for measuring business performance from the multiple dimension database.

2. The method of claim 1 wherein the step of providing a data warehouse further comprises:

standardizing categories in the data warehouse for holding business information according to business information associated with an enterprise.

3. The method of claim 2 wherein the standardizing step further comprises:

generating standard categories for business information associated with an industry;

defining fields in tables of a database that correspond to the standard categories; and utilizing at least one database to store business information in the data warehouse.

4. The method of claim 2 wherein the business information relates to products.

5. The method of claim 1 wherein the step of receiving business information from a first enterprise occurs over a network.

6. The method of claim 1 wherein the step of receiving business information further comprises extracting the business information from local transaction systems associated with each enterprise.

7. The method of claim 1 wherein the set of standardized categories correspond to Advanced Data Resource Management (ADRM) dictionaries that describe business information.

8. A method performed on a data processing system for processing business information generated by multiple enterprises, comprising:

providing a data warehouse having the capability of holding business information;

receiving business information from a second enterprise;

associating business information from the first and second enterprise with a set of standardized categories used in the data warehouse; and

storing the business information received by the first and second enterprises in the data warehouse in accordance with the association of business information and the set of standardized categories.

9. The method of claim 8 wherein the step of providing a data warehouse further comprises:

standardizing categories in the data warehouse for holding business information according to business information associated with an enterprise.

10. The method of claim 9 wherein the standardizing step further comprises:

generating standard categories for business information associated with an industry;

defining fields in tables of a database that correspond to the standard categories; and

utilizing at least one database to store business information in the data warehouse.

11. The method of claim 9 wherein the business information relates to products.

12. The method of claim 8 wherein the step of receiving business information from a first enterprise occurs over a network.

13. The method of claim 8 wherein the step of receiving business information further comprises extracting the business information from local transaction systems associated with each enterprise.

14. The method of claim 8 wherein the set of standardized categories correspond to Advanced Data Resource Management (ADRM) dictionaries that describe business information.

15. A method performed on a data processing system for processing business information generated by multiple enterprises, comprising:

receiving business information from a first and a second enterprise into a data warehouse in accordance with a set of standardized categories;

using metastore information describing the organization of business information in the data warehouse to develop rules for extracting business information from the data warehouse;

extracting a portion of the business information from the data warehouse in accordance with the rules; and

storing the portion of business information extracted from the data warehouse in a data mart.

16. The method of claim 15 wherein the set of standardized categorizes are created using steps comprising:

generating standard categories for business information associated with an industry; and

defining fields in tables of a database that correspond to the standard categories.

17. The method of claim 16 wherein the business information relates to products.

18. The method of claim 15 wherein the step of receiving business information occurs over a network.

19. The method of claim 15 wherein the set of standardized categories correspond to Advanced Data Resource Management (ADRM) dictionaries that describe business information.

20. A method performed on a data processing system for processing business information generated by multiple enterprises, comprising:

receiving business information from a first and a second enterprise into a data warehouse in accordance with a set of standardized categories;

extracting a portion of the business information from the data warehouse;

storing a portion of business information extracted from the data warehouse in at least one data mart;

creating a multiple dimension database from the at least one data mart wherein each dimension of the multiple dimension database corresponds to at least one variable derived from business rules established in an industry; and

generating metrics for measuring business performance from the multiple dimension databases.

21. The method of claim 20 wherein the set of standardized categorizes are created using steps comprising:

generating standard categories for business information associated with an industry; and defining fields in tables of a database that correspond to the standard categories.

22. The method of claim 21 wherein the business information relates to products.

23. The method of claim 20 wherein the step of receiving business information occurs over network.

24. The method of claim 20 wherein the set of standardized categories correspond to Advanced Data Resource Management (ADRM) dictionaries that describe business information

25. A method performed in a data processing system for displaying business information from multiple enterprise systems, comprising:

receiving business information from multiple enterprise systems associated with the creation of products;

determining a supply of products that are supplied in association with an identified product; and

displaying business information related to the supply of products on a display device.

26. A method performed in a data processing system for displaying business information from multiple enterprise systems, comprising:

combining business information from multiple enterprise systems associated with the creation of products;

determining a set of businesses and products that demand an identified product; and

displaying business information related to the set of businesses and products.

27. A data processing system for analyzing business information from multiple enterprise systems, comprising:

a processor:

a memory having instructions capable of execution on the processor that provide a data warehouse having the capability of holding business information, receive business information from a first enterprise, receive business information from a second enterprise, associate business information from the first and second enterprise with a set of standardized categories used in the data warehouse, store the business information received by the first and second enterprises in the data warehouse in accordance with the association of business information and the set of standardized categories, use metastore information describing the organization of business information in the data warehouse to develop rules for extracting business information from the data warehouse, extract a portion of the business information from the data warehouse in accordance with the rules, store the portion of business information extracted from the data warehouse in the data mart, create a multiple dimension database from at least one data mart wherein each dimension of the multiple dimension database corresponds to variables derived from business rules established in an industry, and generate metrics for measuring business performance from the multiple dimension databases.

28. A data processing system for analyzing business information from multiple enterprise systems, comprising:

a processor:

a memory having instructions capable of execution on the processor that receive business information from a first enterprise, receive business information from a second enterprise, associate business information from the first and second enterprise with a set of standardized categories used in the data warehouse, and store the business information received by the first and second enterprises in the data warehouse in accordance with the association of business information and the set of standardized categories.

29. A data processing system for analyzing business information from multiple enterprise systems, comprising:

a processor:

a memory having instructions capable of execution on the processor that receive business information from a first and a second enterprise into a data warehouse in

accordance with a set of standardized categories, extract a portion of the business information from the data warehouse, store a portion of business information extracted from the data warehouse in at least one data mart, create a multiple dimension database from the at least one data mart wherein each dimension of the multiple dimension database corresponds to at least one variable derived from business rules established in an industry, and generate metrics for measuring business performance from the multiple dimension databases.