

Enterprise Resource Planning Systems

Until recently, most large and mid-sized organizations designed and programmed custom information systems in-house. This resulted in an array of stand-alone systems that were designed to the unique needs of the users. Although these systems dealt efficiently with their designated tasks, they did not provide strategic decision support at the enterprise level because they lacked the integration needed for information transfer across organization boundaries. Today the trend in information systems is toward implementing highly integrated, enterprise-oriented systems. These are not custom packages designed for a specific organization. Instead, they are generalized systems that incorporate the best business practices in use. Organizations mix and match precoded software components to assemble an **enterprise resource planning (ERP)** system that best meets their business requirements. This means that an organization may need to change the way it conducts business to take full advantage of the ERP.

This chapter is composed of five major sections. The first section outlines the key features of a generic ERP system by comparing the function and data-storage techniques of a traditional flat file or database system to that of an ERP. The second section describes various ERP configurations related to servers, databases, and bolt-on software. The topic of the third section is data warehousing. A data warehouse is a relational or multidimensional database that supports online analytical processing (OLAP). The fourth section examines key risks associated with ERP implementation. The fifth section reviews the internal control and auditing issues related to ERPs, and the discussion follows the COSO framework.

CHAPTER

II

■ Learning Objectives

After studying this chapter, you should:

- Understand the general functionality and key elements of ERP systems.
- Understand the various aspects of ERP configuration, including servers, databases, and the use of bolt-on software.
- Understand the purpose of data warehousing as a strategic tool and recognize the issues related to the design, maintenance, and operation of a data warehouse.
- Recognize the risks associated with ERP implementation.
- Be aware of the key considerations related to ERP implementation.
- Understand the internal control and auditing implications associated with ERPs.

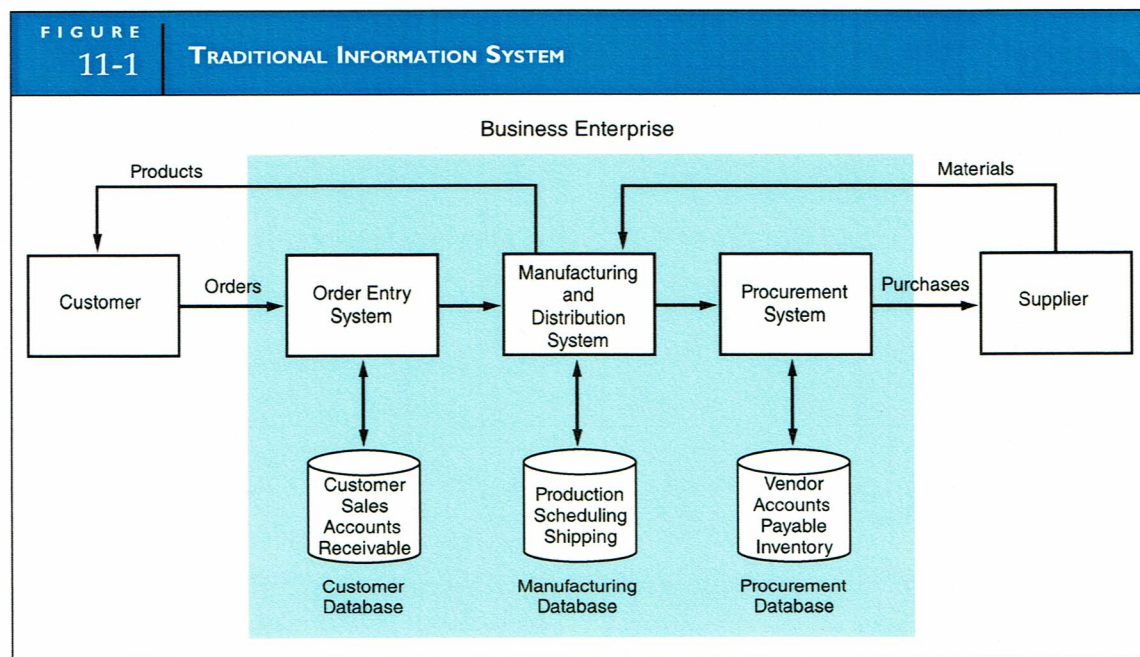
What Is an ERP?

ERP systems are multiple module software packages that evolved primarily from traditional manufacturing resource planning (MRP II) systems. The Gartner Group coined the term *ERP*, which has become widely used in recent years. The objective of ERP is to integrate key processes of the organization, such as order entry, manufacturing, procurement and accounts payable, payroll, and human resources. By doing so, a single computer system can serve the unique needs of each functional area. Designing one system that serves everyone is an undertaking of massive proportions. Under the traditional model, each functional area or department has its own computer system optimized to the way it does its daily business. ERP combines all of these into a single, integrated system that accesses a single database to facilitate the sharing of information and to improve communications across the organization.

To illustrate, consider the traditional model for a manufacturing firm illustrated in Figure 11-1. This company employs **closed database architecture**, which is similar in concept to the basic flat-file model. Under this approach, a database management system is used to provide minimal technological advantage over flat-file systems. The database management system is little more than a private but powerful file system. As with the flat-file approach, the data remain the property of the application. Thus, distinct, separate, and independent databases exist. As is true with the flat-file architecture, there is a high degree of data redundancy in a closed database environment.

When a customer places an order, the order begins a paper-based journey around the company, where it is keyed and rekeyed into the systems of several different departments. These redundant tasks cause delays and lost orders, and promote data entry errors. During transit through various systems, the status of the order may be unknown at any point in time. For example, when responding to a customer query, the marketing department may be unable to look into the production database to determine whether an order has been manufactured and shipped. Instead, the frustrated customer is told to call the manufacturing department. Similarly, the procurement of raw materials from suppliers is not linked to customer orders until they reach the manufacturing stage. This results either in delays because manufacturing awaits the arrival of needed materials, or in excessive investment in inventories to avoid stock-outs.

The lack of effective communication between systems in the traditional model is often the consequence of a fragmented systems design process. Each system tends to be designed as a solution to a specific operational problem rather than as a part of an overall strategy. Furthermore,



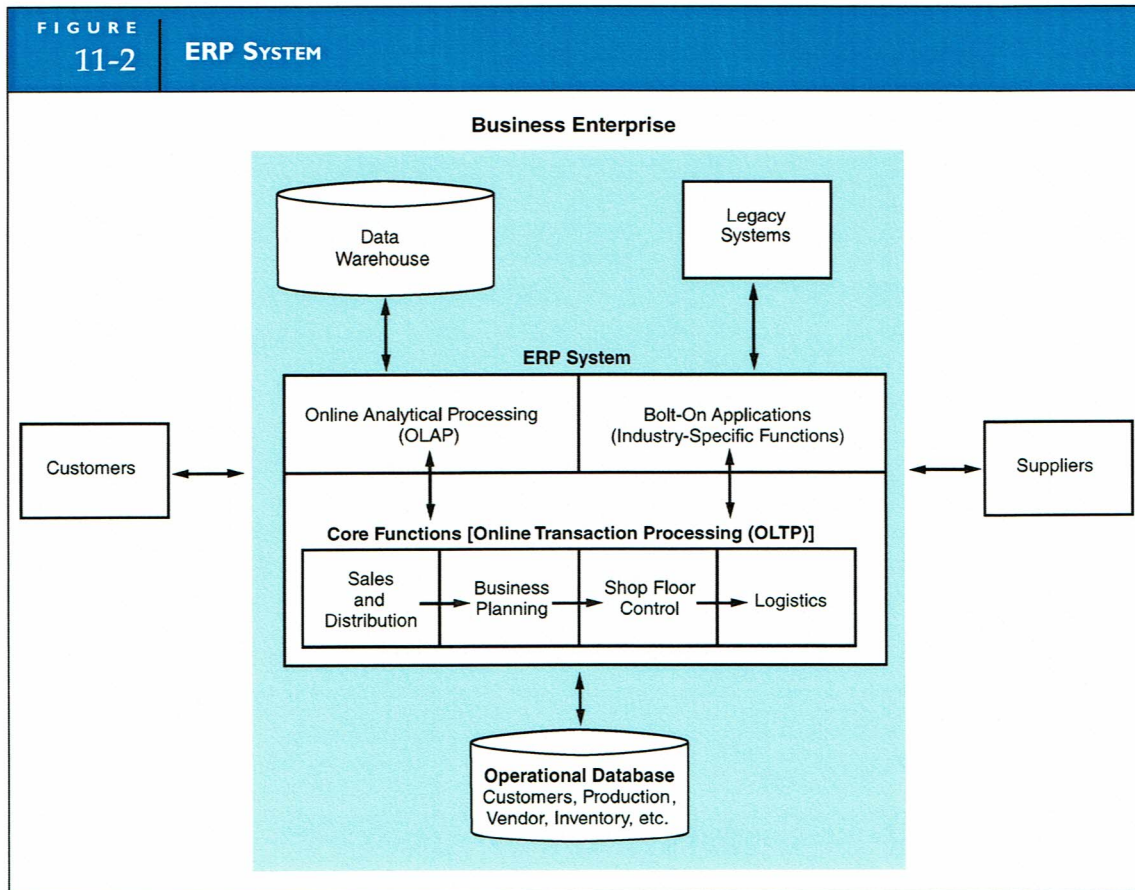
because systems designed in-house emerge independently and over time, they are often constructed on different and incompatible technology platforms. Thus, special procedures and programs need to be created so that older mainframe systems using flat files can communicate with newer distributed systems that use relational databases. Special software patches are also needed to enable commercial systems from different vendors to communicate with each other as well as with custom systems that were developed in-house. Although communications among such a hodgepodge of systems is possible, it is highly fragmented and not conducive to efficient operations.

ERP systems support a smooth and seamless flow of information across the organization by providing a standardized environment for a firm's business processes and a common operational database that supports communications. An overview of ERP is presented in Figure 11-2. Data in the operational database are modeled, structured, and stored in accordance with the internal attributes of the data. They remain independent of any specific application. Extensive data sharing among users occurs through application-sensitive views that present the data in a way that meets all user needs.

ERP CORE APPLICATIONS

ERP functionality falls into two general groups of applications: core applications and business analysis applications. **Core applications** are those applications that operationally support the day-to-day activities of the business. If these applications fail, so does the business. Typical core applications include, but are not limited to, sales and distribution, business planning, production planning, shop floor control, and logistics. Core applications are also called **online transaction processing (OLTP)** applications. Figure 11-2 illustrates these functions applied to a manufacturing firm.

Sales and distribution functions handle order entry and delivery scheduling. This includes checking on product availability to ensure timely delivery and verifying customer credit limits.



Unlike the previous example, customer orders are entered into the ERP only once. Because all users access a common database, the status of an order can be determined at any point. In fact, the customer will be able to check the order directly via an Internet connection. Such integration reduces manual activities, saves time, and decreases human error.

Business planning consists of forecasting demand, planning product production, and detailing routing information that describes the sequence and the stages of the actual production process. Capacity planning and production planning can be very complex; therefore, some ERPs provide simulation tools to help managers decide how to avoid shortages in materials, labor, or plant facilities. When the master production schedule is complete, the data are entered into the materials requirements planning (MRP) module, which provides three key pieces of information: an exception report, a materials requirements list, and inventory requisitions. The exception report identifies potential situations that will result in rescheduling production, such as late delivery of materials. The materials requirements list shows the details of vendor shipments and expected receipts of products and components needed for the order. Inventory requisitions are used to trigger material purchase orders to vendors for items not in stock.

Shop floor control involves the detailed production scheduling, dispatching, and job costing activities associated with the actual production process. Finally, the logistics application is responsible for ensuring timely delivery to the customer. This consists of inventory and warehouse management, as well as shipping. Most ERPs also include their procurement activities within the logistics function.

ONLINE ANALYTICAL PROCESSING

An ERP is more than simply an elaborate transaction processing system. It is a decision support tool that supplies management with real-time information and permits timely decisions that are needed to improve performance and achieve competitive advantage. **Online analytical processing (OLAP)** includes decision support, modeling, information retrieval, ad hoc reporting/analysis, and what-if analysis. Some ERPs support these functions with their own industry-specific modules that can be added to the core system. Other ERP vendors have designed their systems to accept and communicate with specialized bolt-on packages that third-party vendors produce. Sometimes, the user organization's decision support requirements are so unique that the company needs to integrate in-house legacy systems into the ERP. Whether business analysis applications are obtained or derived, they are central to their successful function as a data warehouse. A **data warehouse** is a database constructed for quick searching, retrieval, ad hoc queries, and ease of use. The data are normally extracted periodically from an operational database or from a public information service. An ERP system could exist without having a data warehouse; similarly, organizations that have not implemented an ERP may deploy data warehouses. The trend, however, is that organizations that are serious about competitive advantage deploy both. The recommended data architecture for an ERP implementation includes separate operational and data warehouse databases. We will examine issues related to the creation and operation of a data warehouse later in the chapter.

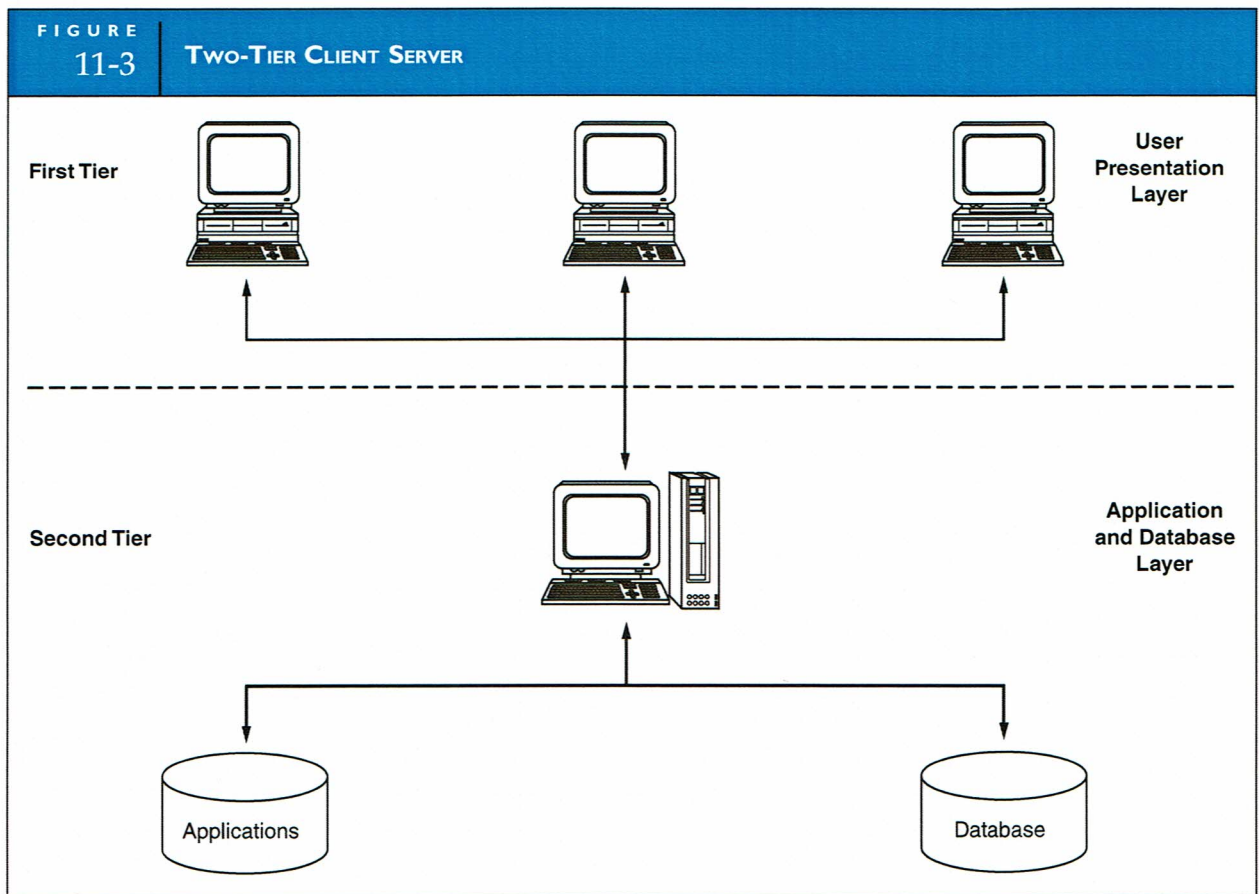
ERP System Configurations

SERVER CONFIGURATIONS

Most ERP systems are based on the **client-server model**, which will be discussed in detail in Chapter 12. Briefly, the client-server model is a form of network topology in which a user's computer or terminal (the client) accesses the ERP programs and data via a host computer called the server. The servers may be centralized, but the clients are usually located at multiple locations throughout the enterprise. Two basic architectures are the two-tier model and the three-tier model, as described in the following sections.

Two-Tier Model

In a typical **two-tier model**, the server handles both application and database duties. Client computers are responsible for presenting data to the user and passing user input back to the server. Some ERP vendors



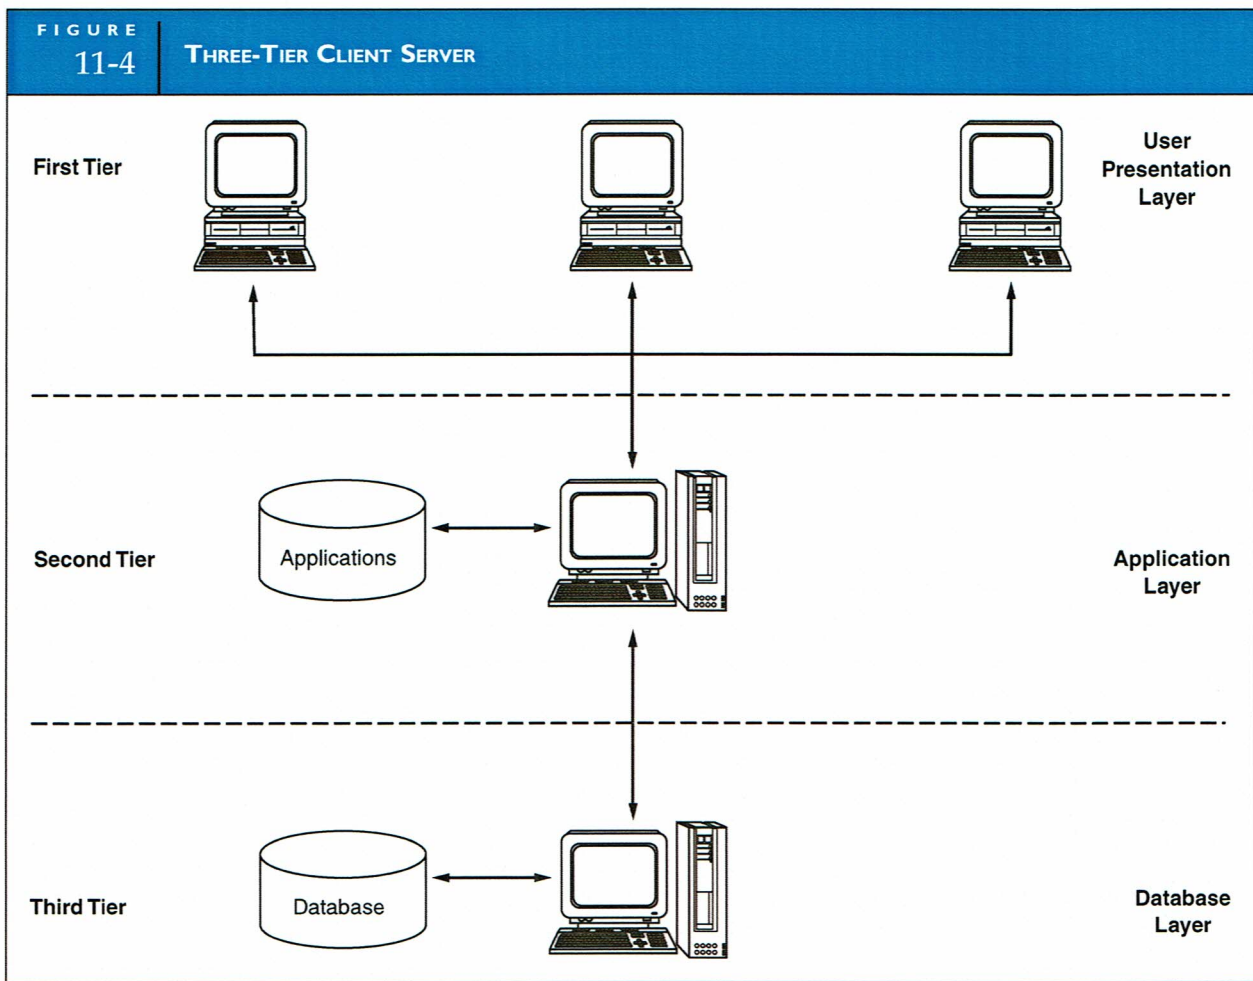
use this approach for local area network (LAN) applications for which the demand on the server is restricted to a relatively small population of users. This configuration is illustrated in Figure 11-3.

Three-Tier Model

The database and application functions are separated in the **three-tier model**. This architecture is typical of large ERP systems that use wide area networks (WANs) for connectivity among the users. Satisfying client requests requires two or more network connections. Initially, the client establishes communications with the application server. The application server then initiates a second connection to the database server. Figure 11-4 presents the three-tier model.

OLTP VERSUS OLAP SERVERS

When implementing an ERP system that will include a data warehouse, a clear distinction needs to be made between the competing types of data processing: OLTP and OLAP. OLTP events consist of large numbers of relatively simple transactions, such as updating accounting records that are stored in several related tables. For example, an order entry system retrieves all of the data relating to a specific customer to process a sales transaction. Relevant data are selected from the Customer table, Invoice table, and a detailed Line Item table. Each table contains an embedded key (i.e., CUSTOMER NUMBER), which is used to relate rows between different tables. The transaction processing activity involves updating the customer's current balance and inserting new records into the Invoice and Line Item tables. The relationships between records in such OLTP transactions are generally simple, and only a few records are actually retrieved or updated in a single transaction.

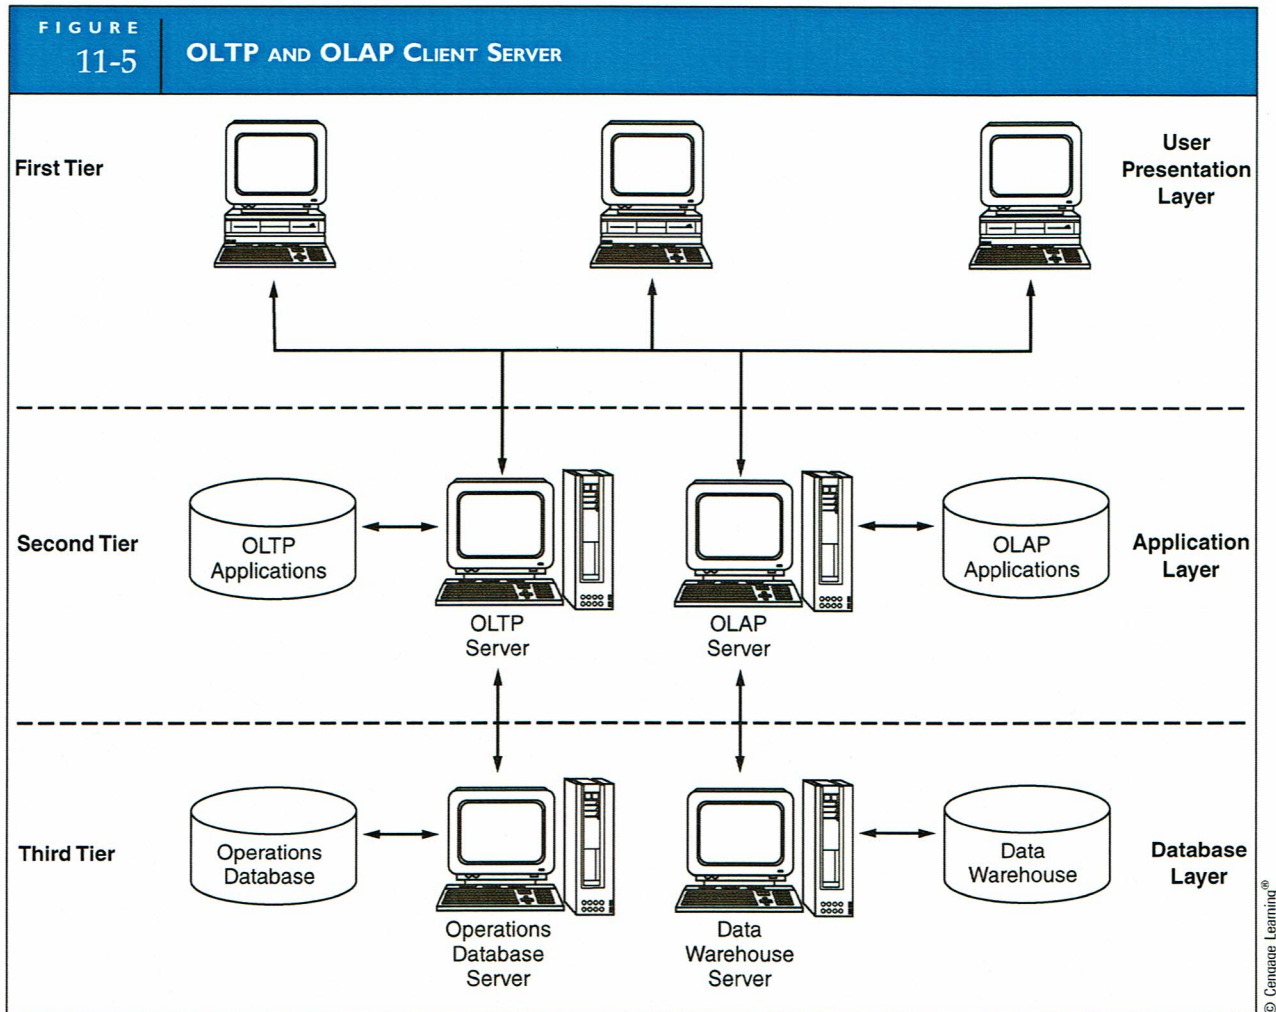


OLAP can be characterized as online transactions that:¹

- Access very large amounts of data (e.g., several years of sales data).
- Analyze the relationships among many types of business elements, such as sales, products, geographic regions, and marketing channels.
- Involve aggregated data, such as sales volumes, budgeted dollars, and dollars spent.
- Compare aggregated data over hierarchical time periods (e.g., monthly, quarterly, yearly).
- Present data in different perspectives, such as sales by region, by distribution channel, or by product.
- Involve complex calculations among data elements, such as expected profit as a function of sales revenue for each type of sales channel in a particular region.
- Respond quickly to user requests so they can pursue an analytical thought process without being stymied by system delays.

An example of an OLAP transaction is the aggregation of sales data by region, product type, and sales channel. The OLAP query may need to access vast amounts of sales data over a multi-year period to find sales for each product type within each region. The user can further refine the query to identify sales volume by product for each sales channel within a given region. Finally,

¹ © Cengage Learning 2015.



the user may decide to perform year-to-year or quarter-to-quarter comparisons for each sales channel. An OLAP application must be able to support this analysis online with rapid response.

The difference between OLAP and OLTP can be summarized as follows. OLTP applications support mission-critical tasks through simple queries of operational databases. OLAP applications support management-critical tasks through analytical investigation of complex data associations that are captured in data warehouses. OLAP and OLTP have specialized requirements that are in direct conflict. Figure 11-5 shows how the client-server architecture enables organizations to deploy separate and specialized application and database servers to resolve these conflicting data management needs. OLAP servers support common analytical operations, including consolidation, drill-down, and slicing and dicing.

Consolidation is the aggregation or roll-up of data. For example, sales offices data can be rolled up to districts, and districts can be rolled up to regions.

Drill-down permits disaggregating data to reveal the underlying details that explain certain phenomena. For example, the user can drill down from total sales returns for a period to identify the actual products returned and the reasons for their return.

Slicing and dicing enables the user to examine data from different viewpoints. One slice of data might show sales within each region. Another slice might present sales by product across regions. Slicing and dicing is often performed along a time axis to depict trends and patterns.

OLAP servers allow users to analyze complex data relationships. The physical database itself is organized in such a way that related data may be rapidly retrieved across multiple dimensions. Thus, OLAP database servers need to be efficient when storing and processing multidimensional data. Later in the chapter, data modeling and storage techniques that improve data warehouse efficiency will be examined. In contrast, relational databases for operations are modeled and optimized to handle OLTP applications. They concentrate on reliability and transaction processing speed instead of decision support need.

DATABASE CONFIGURATION

ERP systems are composed of thousands of database tables. Each table is associated with business processes that are coded into the ERP. The ERP implementation team, which includes key users and information technology (IT) professionals, selects specific database tables and processes by setting switches in the system. Determining how all the switches need to be set for a given configuration requires a deep understanding of the existing processes used in operating the business. Often, however, choosing table settings involves decisions to reengineer the company's processes so that they comply with the best business practices in use. In other words, the company typically changes its processes to accommodate the ERP rather than modifying the ERP to accommodate the company.

BOLT-ON SOFTWARE

Many organizations have ascertained that ERP software alone cannot drive all the processes of the company. These firms use a variety of **bolt-on software** that third-party vendors provide. The decision to use bolt-on software requires careful consideration. Most of the leading ERP vendors have entered into partnership arrangements with third-party vendors that provide specialized functionality. The least risky approach is to choose a bolt-on that is endorsed by the ERP vendor. Some organizations, however, take a more independent approach. Domino's Pizza is a case in point.

Domino's Pizza

Domino's US distribution delivered 338 million pizzas in 1998.² The Company manufactures an average of 4.2 million pounds of dough per week in its 18 US distribution centers. A fleet of 160 trucks carries the dough, along with other food and paper products, to the 4,500 US Domino's franchises. Domino's has no cutoff time for ordering supplies. Therefore, a franchise can call and adjust its order even after the truck has rolled away from the distribution center. To help anticipate demand, Domino's uses forecasting software from Prescient Systems, Inc., which bolts on to its PeopleSoft ERP system. In addition, it uses a system from Manugistics, Inc., to schedule and route the delivery trucks. Each truck has an onboard computer system that feeds data into a time-and-attendance system from Kronos, Inc., which connects to the PeopleSoft human resources module. Domino's also has an extensive data warehouse. To anticipate its market, Domino's performs data mining with software from Cognos, Inc., and Hyperion Solutions Corp.

Domino's had been using these and other applications before it implemented an ERP. The company did not want to retire its existing applications but discovered that the legacy system required data fields that the ERP did not provide. For instance, the routing system tells the truck drivers which stores to visit and in what order. The ERP system did not have a data field for specifying the delivery stop sequence. The warehousing system needs this information, however, to tell loaders what to put in the trucks and in what order. Having confidence in its in-house IT staff, Domino's management decided to take the relatively drastic step of modifying the ERP software to include these fields.

Supply Chain Management

Another development regarding the bolt-on software issue is the rapid convergence between ERP and bolt-on software functionality. **Supply chain management (SCM)** software is a case in point. The supply chain is the set of activities associated with moving goods from the raw materials

² Slater, D. "The Ties That Bolt," *Enterprise Resource Planning, CIO Magazine* (April 15, 1999), 4-9.

stage to the consumer. This includes procurement, production scheduling, order processing, inventory management, transportation, warehousing, customer service, and forecasting the demand for goods. SCM systems are a class of application software that supports this task. Successful SCM systems coordinate and integrate these activities into a seamless process. In addition to the key functional areas within the organization, SCM links all of the partners in the chain, including vendors, carriers, third-party logistics companies, and information systems providers. An organization can achieve competitive advantage by linking the activities in its supply chain more efficiently and effectively than its competitors.

ERP systems and SCM systems are now on converging paths. Many ERP vendors have moved decisively to add SCM functionality to their ERP products. On the other hand, SCM software vendors are expanding their functionality to appear more like ERP systems. As larger ERP vendors move into the midsize company market, the smaller SCM and ERP vendors will likely be pushed out of business.

Data Warehousing

Data warehousing is one of the fastest growing IT issues for businesses today. Not surprisingly, data warehousing functionality is being incorporated into all leading ERP systems. A data warehouse is a relational or multidimensional database that may consume hundreds of gigabytes or even terabytes of disk storage. When the data warehouse is organized for a single department or function, it is often called a **data mart**. Rather than containing hundreds of gigabytes of data for the entire enterprise, a data mart may have only tens of gigabytes of data. Other than size, we make no distinction between a data mart and a data warehouse. The issues discussed in this section apply to both.

The process of data warehousing involves extracting, converting, and standardizing an organization's operational data from ERP and legacy systems and loading them into a central archive—the data warehouse. Once loaded into the warehouse, data are accessible via various query and analysis tools that are used for data mining. Data mining, which was introduced in Chapter 8, is the process of selecting, exploring, and modeling large amounts of data to uncover relationships and global patterns that exist in large databases but are hidden among the vast number of facts. This involves sophisticated techniques that use database queries and artificial intelligence to model real-world phenomena from data collected from the warehouse.

Most organizations implement a data warehouse as part of a strategic IT initiative that involves an ERP system. Implementing a successful data warehouse involves installing a process for gathering data on an ongoing basis, organizing it into meaningful information, and delivering it for evaluation. The data warehousing process has the following essential stages:³

- Modeling data for the data warehouse
- Extracting data from operational databases
- Cleansing extracted data
- Transforming data into the warehouse model
- Loading data into the data warehouse database

MODELING DATA FOR THE DATA WAREHOUSE

Chapters 9 and 10 stressed the importance of data normalization to eliminate three serious anomalies: update, insertion, and deletion anomalies. Normalizing data in an operational database is necessary to reflect accurately the dynamic interactions among entities. Data attributes are constantly updated, new attributes are added, and obsolete attributes are deleted on a regular basis. Even though a fully normalized database will yield the flexible model needed for supporting multiple users in this dynamic operational environment, it also adds to complexity that is translated into performance inefficiency.

3 P. Fiore, "Everyone Is Talking about Data Warehousing," *Evolving Enterprise* (Spring 1998), 2.

Denormalized Data Warehouse

Because of the vast size of a data warehouse, performance inefficiency can be devastating. A three-way join between many thousands of tables in a large data warehouse may take an unacceptably long time to complete, and may be unnecessary. In the data warehouse model, the relationship among attributes does not change. Because historical data are static in nature, nothing is gained by constructing normalized tables with dynamic links.

For example, in an operational database system, Product X may be an element of work-in-process (WIP) in Department A this month and part of Department B's WIP next month. In a properly normalized data model, it would be incorrect to include Department A's WIP data as part of a Sales Order table that records an order for Product X. Only the product item number would be included in the Sales Order table as a foreign key linking it to the Product table. Relational theory would call for a join (link) between the Sales Order table and Product table to determine the production status (i.e., which department the product is currently in) and other attributes of the product. From an operational perspective, complying with relational theory is important because the relation changes as the product moves through different departments over time. Relational theory does not apply to a data warehousing system because the Sales Order/Product relation is stable.

Wherever possible, therefore, normalized tables pertaining to selected events may be consolidated into denormalized tables. Figure 11-6 illustrates how sales order data are reduced to a single denormalized Sales Order table for storage in a data warehouse system.

EXTRACTING DATA FROM OPERATIONAL DATABASES

Data extraction is the process of collecting data from operational databases, flat files, archives, and external data sources. Operational databases typically need to be out of service when data extraction occurs to avoid data inconsistencies. Because of their large size and the need for a speedy transfer to minimize the downtime, little or no conversion of data occurs at this point. A technique called **changed data capture** can dramatically reduce the extraction time by capturing only newly modified data. The extraction software compares the current operational database with an image of the data taken at the last transfer of data to the warehouse. Only data that have changed in the interim are captured.

Extracting Snapshots versus Stabilized Data

Transaction data stored in the operational database go through several stages as economic events unfold. For example, a sales transaction first undergoes credit approval, then the product is shipped, then billing occurs, and finally payment is received. Each of these events changes the state of the transaction and associated accounts, such as inventory, accounts receivable, and cash. A key feature of a data warehouse is that the data contained in it are in a nonvolatile, stable state. Typically, transaction data are loaded into the warehouse only when the activity on them has been completed. Potentially important relationships between entities may, however, be absent from data that are captured in this stable state. For example, information about canceled sales orders will probably not be reflected among the sales orders that have been shipped and paid for before they are placed in the warehouse. One way to reflect these dynamics is to extract the operations data in slices of time. These slices provide snapshots of business activity. For example, decision makers may want to observe sales transactions approved, shipped, billed, and paid at various points in time, along with snapshots of inventory levels at each state. Such data may be useful in depicting trends in the average time taken to approve credit or ship goods that might help explain lost sales.

CLEANSING EXTRACTED DATA

Data cleansing involves filtering out or repairing invalid data prior to being stored in the warehouse. Operational data are dirty for many reasons. Clerical, data entry, and computer program errors can create illogical data, such as negative inventory quantities, misspelled names, and blank fields. Data cleansing also involves transforming data into standard business terms with standard data values. Data are often combined from multiple systems that use slightly different spellings to represent common terms, such as `cust`, `cust_id`, or `cust_no`. Some operational systems may use entirely different

FIGURE 11-6 DENORMALIZED DATA

A. Normalized Representation for an Operational Database System

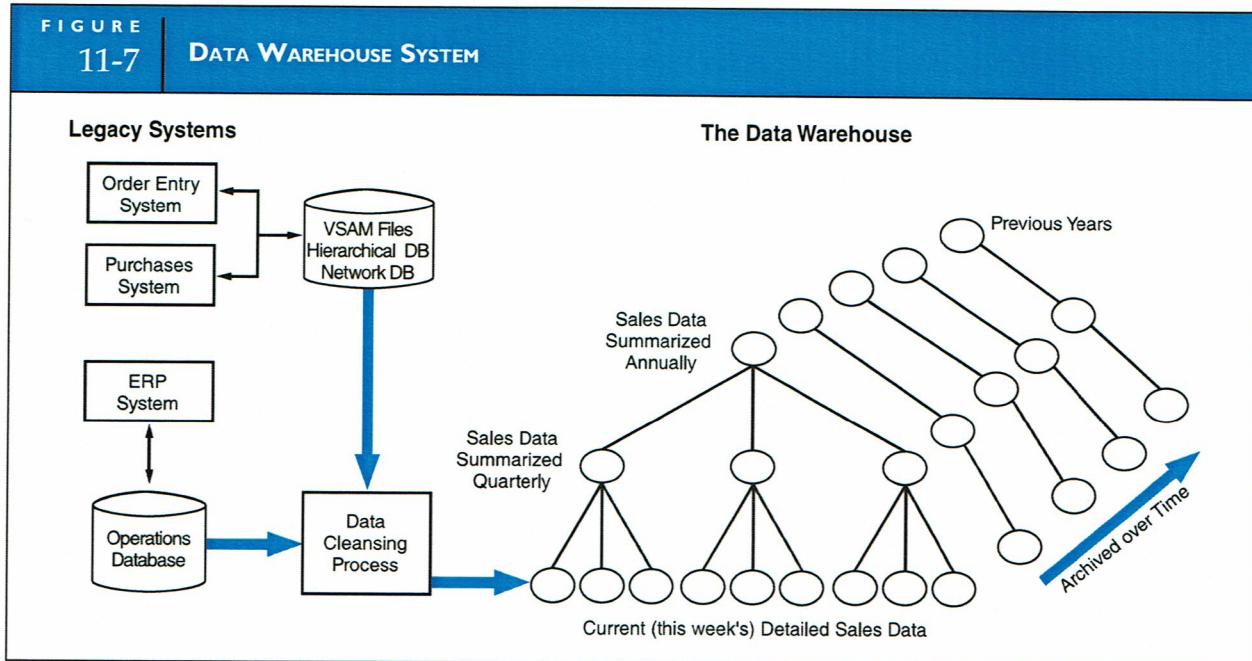
Customer Table				
Customer Number	Name	Street	City	State
34675	John Smith	10 Elm	Bath	PA

Invoice Table				
Invoice Number	Invoice Date	Shipped Date	Invoice Amount	Customer Number
8866376	06/12/09	06/23/09	600	34675

Line Item Table			
Invoice Number	Item Number	Quantity	Extended Price
8866376	j683	2	400
8866376	r223	5	200

B. Denormalized Representation for Data Warehouse System

Customer Number	Name	Street	City	State	Invoice Number	Invoice Date	Shipped Date	Invoice Amount	Item Number	Quantity	Price	Extended Price
34675	John Smith	10 Elm	Bath	PA	8866376	06/12/09	06/23/09	600	j683	2	200	400
34675	John Smith	10 Elm	Bath	PA	8866376	06/12/09	06/23/09	600	r223	5	40	200



terms to refer to the same entity. For example, a bank customer with a certificate of deposit and an outstanding loan may be called a lender by one system and a borrower by another. The source application may use cryptic or difficult-to-understand terms for a number of reasons. For example, some older legacy systems were designed at a time when programming rules placed severe restrictions on naming and formatting data attributes. Also, a commercial application may assign attribute names that are too generic for the needs of the data warehouse user. Businesses that purchase commercial data, such as competitive performance information or market surveys, need to extract data from whatever format the external source provides and reorganize them according to the conventions used in the data warehouse. During the cleansing process, therefore, the attributes taken from multiple systems need to be transformed into uniform, standard business terms. This tends to be an expensive and labor-intensive activity, but one that is critical in establishing data integrity in the warehouse. Figure 11-7 illustrates the role of data cleansing in building and maintaining a data warehouse.

TRANSFORMING DATA INTO THE WAREHOUSE MODEL

A data warehouse is composed of both detail and summary data. To improve efficiency, data can be transformed into summary views before they are loaded into the warehouse. For example, many decision makers may need to see product sales figures summarized weekly, monthly, quarterly, or annually. It may not be practical to summarize information from detail data every time the user needs it. A data warehouse that contains the most frequently requested summary views of data can reduce the amount of processing time during analysis. Referring again to Figure 11-7, we see the creation of summary views over time. These are typically created around business entities such as customers, products, and suppliers. Unlike operational views, which are virtual in nature with underlying base tables, data warehouse views are physical tables. Most OLAP software will, however, permit the user to construct virtual views from detail data when one does not already exist.

A data warehouse will often provide multiple summary views based on the same detailed data such as customers or products. For example, several different summary views may be generated from sales order detail data; these may include summaries by product, customer, and region. From such views, an analyst can drill down into the underlying detail data. Many business problems require a review of detail data to fully evaluate a trend, a pattern, or an anomaly exhibited in the summarized reports. Also, a single anomaly in detail data may manifest itself differently in different summary views.

LOADING THE DATA INTO THE DATA WAREHOUSE DATABASE

Most organizations have found that data warehousing success requires that the data warehouse be created and maintained separately from the operational (transaction processing) databases. This point is developed further in the next sections.

Internal Efficiency

One reason for a separate data warehouse is that the structural and operational requirements of transaction processing and data mining systems are fundamentally different, making it impractical to keep both operational (current) and archive data in the same database. Transaction processing systems need a data structure that supports performance, while data mining systems need data organized in a manner that permits broad examination and the detection of underlying trends.

Integration of Legacy Systems

The continued influence of legacy systems is another reason that the data warehouse needs to be independent of operations. A remarkably large number of business applications continue to run in the mainframe environment of the 1970s. By some estimates, more than 70 percent of business data for large corporations still resides in the mainframe environment. The data structures these systems employ are often incompatible with the architectures of modern data mining tools. Hence, transaction data that are stored in navigational databases and Virtual Storage Access Method systems often end up in large tape libraries that are isolated from the decision process. A separate data warehouse provides a venue for integrating the data from legacy and contemporary systems into a common structure that supports entity-wide analysis.

Consolidation of Global Data

Finally, the emergence of the global economy has brought about fundamental changes in business organizational structure and has profoundly changed the information requirements of business entities. Unique business complexities challenge decision makers in the global corporation. For example, they need to assess the profitability of products built and sold in multiple countries with volatile currencies. Such challenges add complexity to data mining. A separate centralized data warehouse is an effective means of collecting, standardizing, and assimilating data from diverse sources.

In conclusion, the creation of a data warehouse separate from operational systems is a fundamental data warehousing concept. Many organizations now consider data warehouse systems to be key components of their information systems strategy. As such, they allocate considerable resources to build data warehouses concurrently with the operational systems being implemented.

DECISIONS SUPPORTED BY THE DATA WAREHOUSE

By making the data warehouse as flexible and friendly as possible, it becomes accessible to many end users. Some decisions that a data warehouse supports are not fundamentally different from those that traditional databases support. Other information uses, such as multidimensional analysis and information visualization, are not possible with traditional systems. Some users of the data warehouse need routine reports based on traditional queries. When standard reports can be anticipated in advance, they can be provided automatically as a periodic product. Automatic generation of standard information reduces access activity against the data warehouse and will improve its efficiency in dealing with more esoteric needs.

Drill-down capability is a useful data analysis technique associated with data mining. Drill-down analysis begins with the summary views of data described previously. When anomalies or interesting trends are observed, the user drills down to lower-level views and ultimately into the underlying detail data.

Obviously, such analysis cannot be anticipated like a standard report. Drill-down capability is an OLAP feature of data mining tools available to the user. Tools for data mining are evolving rapidly to satisfy the decision maker's need to understand the business unit's behavior in relation to key entities,

TABLE 11-1 APPLICATIONS OF DATA MINING	
Business Field	Application
Banking/Investments	Detect patterns of fraudulent credit card use. Identify loyal customers and predict those likely to change their credit card affiliation. Examine historical market data to determine investors' stock trading rules. Predict credit card spending of key customer groups. Identify correlations between different financial indicators.
Health Care and Medical Insurance	Predict office visits from historical analysis of historical patient behavior. Identify successful and economic medical therapies for different illnesses. Identify which medical procedures tend to be claimed together. Predict which customers will buy new policies. Identify behavior patterns associated with high-risk customers. Identify indicators of fraudulent behavior.
Marketing	Identify buying patterns based on historical customer data. Identify relationships among customer demographic data. Predict response to various forms of marketing and promotion campaigns.

© Cengage Learning®

including customers, suppliers, employees, and products. Standard reports and queries produced from summary views can answer many *what* questions, but drill-down capability answers the *why* and *how* questions. Table 11-1 summarizes some of the applications of data mining in decision support.

SUPPORTING SUPPLY CHAIN DECISIONS FROM THE DATA WAREHOUSE

The primary reason for data warehousing is to optimize business performance. Many organizations believe that more strategic benefit can be gained by sharing data externally. By providing customers and suppliers with the information they need when they need it, the company can improve its relationships and provide better service. The potential gain to the giving organization is seen in a more responsive and efficient supply chain. Using Internet technologies and OLAP applications, an organization can share its data warehouse with its trading partners and, in effect, treat them like divisions of the firm. A few examples of this approach are outlined in the following extract.⁴

Western Digital Corporation, a leading manufacturer of hard drives, plans to grant certain suppliers access to its data warehouse so suppliers can view performance data on their parts. Because Western Digital maintains a limited engineering staff, the company relies on its suppliers to act as strategic partners in product development. Providing suppliers with performance data allows them to make improvements and participate in the engineering process. The suppliers improve their parts, which in turn improves Western Digital's products.

The company's data warehouse holds more than 600 gigabytes of raw data collected from more than 100,000 drives that it manufactures each day. Approximately 800 attributes are collected on each drive, which can be analyzed using OLAP software. The systems feeding the warehouse include ERP applications, data from trouble-call centers, data from failure-analysis systems, and field test data from customer sites and service centers. The company routinely searches the data warehouse for failure information on every drive that it manufactures. All failures and their causes can be linked back to the supplier.

⁴ Davis, "Data Warehouses Open Up," *Information Week Online News in Review* (June 28, 1999).

The General Motors (GM) supply-chain data warehouse is available via the web to more than 5,000 suppliers worldwide. Suppliers can log on to a secure website and query information on the quantities of supplies shipped, delivery times, and prices. This information will help GM suppliers optimize their product planning, ability to source materials, and shipping-fulfillment processes.

MIM Health Plans, Inc., an independent pharmacy benefits management company, lets its customers view warehouse data to promote better buying decisions. For instance, benefits managers can view reports and drill down into the warehouse to see claims costs, overall costs, the number of prescriptions ordered in a given time period, the number of brand versus generic drugs, and other decision metrics.

Risks Associated with ERP Implementation

The benefits from ERP can be significant, but they do not come risk-free to the organization. An ERP system is not a silver bullet that will, by its mere existence, solve an organization's problems. If that were the case, there would never be ERP failures, but there have been many. This section examines some of the risk issues that need to be considered.

BIG BANG VERSUS PHASED-IN IMPLEMENTATION

Implementing an ERP system has more to do with changing the way an organization does business than it does with technology. As a result, most ERP implementation failures are the result of cultural problems within the firm that stand in opposition to the objective of process reengineering. Strategies for implementing ERP systems to achieve this objective follow two general approaches: the big bang approach and the phased-in approach.

The **big bang** method is the more ambitious and risky of the two. Organizations taking this approach attempt to switch operations from their old legacy systems to the new system in a single event that implements the ERP across the entire company. Although this method has certain advantages, it has been associated with numerous system failures. Because the new ERP system means new ways of conducting business, getting the entire organization on board and in sync can be a daunting task. On day 1 of the implementation, no one within the organization will have had any experience with the new system, that is, everyone in the company is a trainee learning a new job.

The new ERP system will initially meet with opposition because using it involves compromise. The legacy systems, which everyone in the organization was familiar with, had been honed over the years to meet exact needs. In most cases, ERP systems have neither the range of functionality nor the familiarity of the legacy systems that they replace. Also, because a single system is now serving the entire organization, individuals at data input points often find themselves entering considerably more data than they did previously with the more narrowly focused legacy system.

As a result, the speed of the new system often suffers and causes disruptions to daily operations. These problems are typically experienced whenever any new system is implemented. The magnitude of the problem is the issue under the big bang approach, in which everyone in the company is affected. Once the initial adjustment period has passed and the new culture emerges, however, the ERP becomes an effective operational and strategic tool that provides competitive advantage to the firm.

Because of the disruptions associated with the big bang, the **phased-in** approach has emerged as a popular alternative. It is particularly suited to diversified organizations whose units do not share common processes and data. In these types of companies, independent ERP systems can be installed in each business unit over time to accommodate the adjustment periods needed for assimilation. Common processes and data, such as the general ledger function, can be integrated across the organization without disrupting operations throughout the firm.

Organizations that are not diversified can also employ the phased-in approach. The implementation usually begins with one or more key processes, such as order entry. The goal is to get ERP up and running concurrently with legacy systems. As more of the organization's functions are converted to ERP, legacy systems are systematically retired. In the interim, the ERP is interfaced to legacy systems. During this period, the objectives of system integration and process reengineering, which are fundamental to the ERP model, are not achievable. To take full advantage of the ERP,

process reengineering will still need to occur. Otherwise, the organization will have simply replaced its old legacy system with a very expensive new one.

OPPOSITION TO CHANGES IN THE BUSINESS'S CULTURE

To be successful, all functional areas of the organization need be involved in determining the culture of the firm and in defining the new system's requirements. The firm's willingness and ability to undertake a change of the magnitude of an ERP implementation is an important consideration. If the corporate culture is such that change is not tolerated or desired, then an ERP implementation will not be successful.

The technological culture must also be assessed. Organizations that lack technical support staff for the new system or have a user base that is unfamiliar with computer technology face a steeper learning curve and a potentially greater barrier to acceptance of the system by their employees.

CHOOSING THE WRONG ERP

Because ERP systems are prefabricated systems, users need to determine whether a particular ERP fits their organization's culture and its business processes. A common reason for system failure is when the ERP does not support one or more important business processes. In one example, a textile manufacturer in India implemented an ERP only to discover afterward that it did not accommodate a basic need.

The textile company had a policy of maintaining two prices for each item of inventory that it sold. One price was used for the domestic market, and a second price, which was four times higher, was for export sales. The ERP that the user implemented was not designed to allow two different prices for the same inventory item. The changes needed to make the ERP work were both extensive and expensive. Serious system disruptions resulted from this oversight. Furthermore, modifying an ERP program and database can introduce potential processing errors and can make updating the system to later versions difficult.

Goodness of Fit

Management needs to make sure that the ERP it chooses is right for the company. No single ERP system is capable of solving all the problems of all organizations. For example, SAP's R/3 was designed primarily for manufacturing firms with highly predictable processes that are relatively similar to those of other manufacturers. It may not be the best solution for a service-oriented organization that has a great need for customer-related activities conducted over the Internet.

Finding a good functionality fit requires a software selection process that resembles a funnel, which starts broad and systematically becomes more focused. It begins with a large number of software vendors that are potential candidates. Evaluation questions are asked of vendors in iterative rounds. Starting with a large population of vendors and a small number of high-level qualifier questions, the number of vendors is reduced to a manageable few. With proper questioning, more than half the vendors are removed from contention with as few as 10 to 20 questions. In each succeeding round, the questions asked become more detailed, and the population of vendors decreases.

When a business's processes are truly unique, the ERP system must be modified to accommodate industry-specific (bolt-on) software or to work with custom-built legacy systems. Some organizations, such as telecommunications service providers, have unique billing operations that off-the-shelf ERP systems cannot satisfy. Before embarking on the ERP journey, the organization's management needs to assess whether it can and should reengineer its business practices around a standardized model.

System Scalability Issues

If an organization's management expects business volumes to increase substantially during the life of the ERP system, then there is a scalability issue that needs to be addressed. **Scalability** is the system's ability to grow smoothly and economically as user requirements increase. The term *system* in this context refers to the technology platform, application software, network configuration, or database. Smooth and economic growth is the ability to increase system capacity at an acceptable incremental

cost per unit of capacity without encountering limits that would demand a system upgrade or replacement. User requirements pertain to volume-related activities, such as transaction processing volume, data entry volume, data output volume, data storage volume, or increases in the user population.

To illustrate scalability, four dimensions of scalability are important: size, speed, workload, and transaction cost. In assessing scalability needs for an organization, each of these dimensions in terms of the ideal of linear scaling must be considered.⁵

Size. With no other changes to the system, if database size increases by a factor of X, then query response time will increase by no more than a factor of X in a scalable system. For example, if business growth causes the database to increase from 100 to 500 gigabytes, then transactions and queries that previously took 1 second will now take no more than 5 seconds.

Speed. An increase in hardware capacity by a factor of X will decrease query response time by no less than a factor of X in a scalable system. For example, increasing the number of input terminals (nodes) from 1 to 20 will increase transaction processing time proportionately. Transactions that previously took 20 seconds will now take no more than 1 second in a system with linear scaling.

Workload. If workload in a scalable system is increased by a factor of X, then response time, or throughput, can be maintained by increasing hardware capacity by a factor of no more than x. For example, if transaction volume increased from 400 per hour to 4,000 per hour, the previous response time can be achieved by increasing the number of processors by a factor of 10 in a system that is linearly scalable.

Transaction Cost. In a scalable system, increases in workload do not increase transaction cost. Therefore, an organization should not need to increase system capacity faster than demand. For example, if the cost of processing a transaction in a system with one processor is 10 cents, then it should still cost no more than 10 cents when the number of processors is increased to handle larger volumes of transactions.

Vendors of ERP systems sometimes advertise scalability as if it were a single-dimension factor. In fact, it is a multifaceted issue. Some systems accommodate growth in user populations better than others. Some systems can be scaled to provide more efficient access to large databases when business growth demands it. All systems, however, have their scaling limits. Because infinite scalability is impossible, prospective users need to assess their needs and determine how much scalability they want to purchase up front, and what form it should take. The key is to anticipate specific scalability issues before making an ERP investment and before the issues become reality.

CHOOSING THE WRONG CONSULTANT

Implementing an ERP system is an event that most organizations will undergo only once. Success of the projects rests on skills and experience that typically do not exist in-house. Because of this, virtually all ERP implementations involve an outside consulting firm, which coordinates the project, helps the organization to identify its needs, develops a requirements specification for the ERP, selects the ERP package, and manages the cutover. ERP consulting has grown into a \$20 billion-per-year market. The fee for a typical implementation is normally between three and five times the cost of the ERP software license.

Consulting firms with large ERP practices at times have been desperately short of human resources. This was especially true in the mid- to late 1990s, when thousands of clients were rushing to implement ERP systems before the new millennium to avoid Y2K (year-2000) problems. As demand for ERP implementations grew beyond the supply of qualified consultants, more and more stories of botched projects materialized.

A frequent complaint is that consulting firms promise experienced professionals but deliver incompetent trainees. They have been accused of employing a bait-and-switch maneuver to get contracts. At the initial engagement interview, the consulting firm introduces their top consultants, who are sophisticated,

⁵ Winter, "Scalable Systems: Lexicology of Scale," *Intelligent Enterprise Magazine* (March 2000), 68–74.

talented, and persuasive. The client agrees to the deal with the firm, but incorrectly assumes that these individuals, or others with similar qualifications, will actually implement the system.

The problem has been equated to the airline industry's common practice of overbooking flights. Some suggest that consulting firms, not wanting to turn away business, are guilty of overbooking their consulting staff. The consequences, however, are far graver than the inconvenience of missing a flight—a free hotel room and meal cannot compensate for the damages done. Therefore, before engaging an outside consultant, management should:

- Interview the staff proposed for the project and draft a detailed contract specifying which members of the consulting team will be assigned to which tasks.
- Establish in writing how staff changes will be handled.
- Conduct reference checks of the proposed staff members.
- Align the consultants' interests with those of the organization by negotiating a pay-for-performance scheme based on achieving certain milestones in the project. For example, the actual amount paid to the consultant may be between 85 and 115 percent of the contracted fee, based on whether a successful project implementation comes in under or over schedule.
- Set a firm termination date for the consultant to avoid consulting arrangements becoming interminable, resulting in dependency and an endless stream of fees.

HIGH COST AND COST OVERRUNS

Total cost of ownership (TCO) for ERP systems varies greatly from company to company. For medium- to large-sized systems implementations, costs range from hundreds of thousands to hundreds of millions of dollars. TCO includes hardware, software, consulting services, internal personnel costs, installation, and upgrades and maintenance to the system for the first two years after implementation. The risk comes in the form of underestimated and unanticipated costs. Some of the more commonly experienced problems occur in the following areas.

Training. Training costs are invariably higher than estimated because management focuses primarily on the cost of teaching employees the new software. This is only part of the needed training. Employees also need to learn new procedures, which is often overlooked during the budgeting process.

System Testing and Integration. In theory, ERP is a holistic model in which one system drives the entire organization. The reality, however, is that many organizations use their ERP as a backbone system that is attached to legacy systems and other bolt-on systems, which support the unique needs of the firm. Integrating these disparate systems with the ERP may involve writing special conversion programs or even modifying the internal code of the ERP. Integration and testing are done on a case-by-case basis; thus, the cost is extremely difficult to estimate in advance.

Database Conversion. A new ERP system usually means a new database. Data conversion is the process of transferring data from the legacy system's flat files to the ERP's relational database. When the legacy system's data are reliable, the conversion process may be accomplished through automated procedures. Even under ideal circumstances, a high degree of testing and manual reconciliation is necessary to ensure that the transfer was complete and accurate. More often, the data in the legacy system are not reliable (sometimes called dirty). Empty fields and corrupted data values cause conversion problems that demand human intervention and data rekeying. Also, and more importantly, the structure of the legacy data is likely to be incompatible with the reengineered processes of the new system. Depending on the extent of the process reengineering involved, the entire database may need to be converted through manual data entry procedures.

Develop Performance Measures

Because ERPs are extremely expensive to implement, many managers are often dismayed at the apparent lack of cost savings that they achieve in the short term. In fact, a great deal of criticism about the relative success of ERPs relates to whether they provide benefits that outweigh their cost.

To assess benefits, management first needs to know what it wants and needs from the ERP. It should then establish key performance measures such as reductions in inventory levels, inventory turnover, stock-outs, and average order fulfillment time that reflect its expectations. To monitor performance in such key areas, some organizations establish an independent value assessment group that reports to top management. Although financial breakeven on an ERP will take years, by developing focused and measurable performance indicators, an operational perspective on its success can be developed.

DISRUPTIONS TO OPERATIONS

ERP systems can wreak havoc in the companies that install them. In a Deloitte Consulting survey of 64 Fortune 500 companies, 25 percent of the firms surveyed admitted that they experienced a drop in performance in the period immediately following implementation. The reengineering of business processes that often accompanies ERP implementation is the most commonly attributed cause of performance problems. Operationally speaking, when business begins under the ERP system, everything looks and works differently from the way it did with the legacy system. An adjustment period is needed for everyone to reach a comfortable point on the learning curve. Depending on the culture of the organization and attitudes toward change within the firm, adjustment may take longer in some firms than in others. The list of major organizations that have experienced serious disruptions includes Dow Chemical, Boeing, Dell Computer, Apple Computer, Whirlpool Corporation, and Waste Management.

The most notorious case in the press was Hershey Foods Corporation, which had trouble processing orders through its new ERP system and was unable to ship products. As a result of these disruptions, Hershey's 1999 third-quarter sales dropped by 12.4 percent compared to the previous year's sales, and earnings were down by 18.6 percent. Hershey's problem has been attributed to two strategic errors related to system implementation. First, because of schedule overruns, it decided to switch to the new system during its busy season. The inevitable snags that arise from implementations of complex systems like SAP's R/3 are easier to deal with during slack business periods. Secondly, many experts believe that Hershey attempted to do too much in a single implementation. In addition to the R/3 system, it implemented a customer relations management system and logistics software from two different vendors, which had to interface with R/3. The ERP and these bolt-on components were all implemented using the big bang approach.

Implications for Internal Control and Auditing

As with any system, the internal control and audit of ERP systems are issues. Key concerns are examined next within the COSO framework.

TRANSACTION AUTHORIZATION

A key benefit of an ERP system is its tightly integrated architecture of modules. This structure, however, also poses potential problems for transaction authorization. For example, the bill of materials drives many manufacturing systems. If the procedures for the creation of the bill of materials are not configured correctly, every component that uses the bill of materials could be affected. Controls need to be built into the system to validate transactions before other modules accept and act upon them. Because of ERP's real-time orientation, systems are more dependent on programmed controls than on human intervention, as was the case with legacy systems. The challenge for auditors in verifying transaction authorization is to gain a detailed knowledge of the ERP system configuration, as well as a thorough understanding of the business processes and the flow of information between system components.

SEGREGATION OF DUTIES

Operational decisions in ERP-based organizations are pushed down to a point as close as possible to the source of the event. Manual processes that normally require segregation of duties are, therefore, often eliminated in an ERP environment. For example, shop supervisors may order inventories from

suppliers, and receiving dock personnel may post inventory receipts to the inventory records in real time. Furthermore, ERP forces together many different business functions, such as order entry, billing, and accounts payable, under a single integrated system. Organizations using ERP systems must establish new security, audit, and control tools to ensure duties are properly segregated. An important aspect of such control is the assignment of *roles*, which is discussed in a later section.

SUPERVISION

An often-cited pitfall of an ERP implementation is that management does not fully understand its impact on business. Too often, after the ERP is up and running, only the implementation team understands how it works. Because their traditional responsibilities will be changed, supervisors need to acquire an extensive technical and operational understanding of the new system. Typically, when an organization implements an ERP, many decision-making responsibilities are pushed down to the shop floor level. The employee-empowered philosophy of ERP should not eliminate supervision as an internal control. Instead, it should provide substantial efficiency benefits. Supervisors should have more time to manage the shop floor and, through improved monitoring capability, increase their span of control.

ACCOUNTING RECORDS

ERP systems have the ability to streamline the entire financial reporting process. In fact, many organizations can and do close their books daily. OLTP data can be manipulated quickly to produce ledger entries, accounts receivable and payable summaries, and financial consolidation for both internal and external users. Traditional batch controls and audit trails are no longer needed in many cases. The risk of financial statement misrepresentations is mitigated by improved data entry accuracy through the use of default values, cross checking, and specified user views of data. In spite of ERP technology, some risk to accounting record accuracy may still exist. Because of the close interfaces with customers and suppliers, some organizations run the risk that corrupted or inaccurate data may be passed from these external sources and corrupt the ERP accounting database. Additionally, many organizations need to import data from legacy systems into their ERP systems. These data may be laden with problems, such as duplicate records, inaccurate values, or incomplete fields. Consequently, strict data cleansing is an important control. Special scrubber programs are used as interfaces between the ERP and the exporting systems to reduce these risks and ensure that the most accurate and current data are being received.

INDEPENDENT VERIFICATION

Because ERP systems employ OLTP, traditional, independent verification controls such as reconciling batch control numbers (discussed in Chapter 3) serve little purpose. Similarly, process reengineering to improve efficiency also changes the nature of independent verification. For example, the traditional three-way match of the purchase order, receiving report, and invoice, and the subsequent writing of a check, may be completely automated in an ERP environment. The focus of independent verification thus needs to be redirected from the individual transaction level to one that views overall performance. ERP systems come with canned controls and can be configured to produce performance reports that should be used as assessment tools. Internal auditors also play an important role in this environment and need to acquire a thorough technical background and comprehensive understanding of the ERP system. Ongoing independent verification efforts can be conducted only by a team well versed in ERP technology.

ACCESS CONTROLS

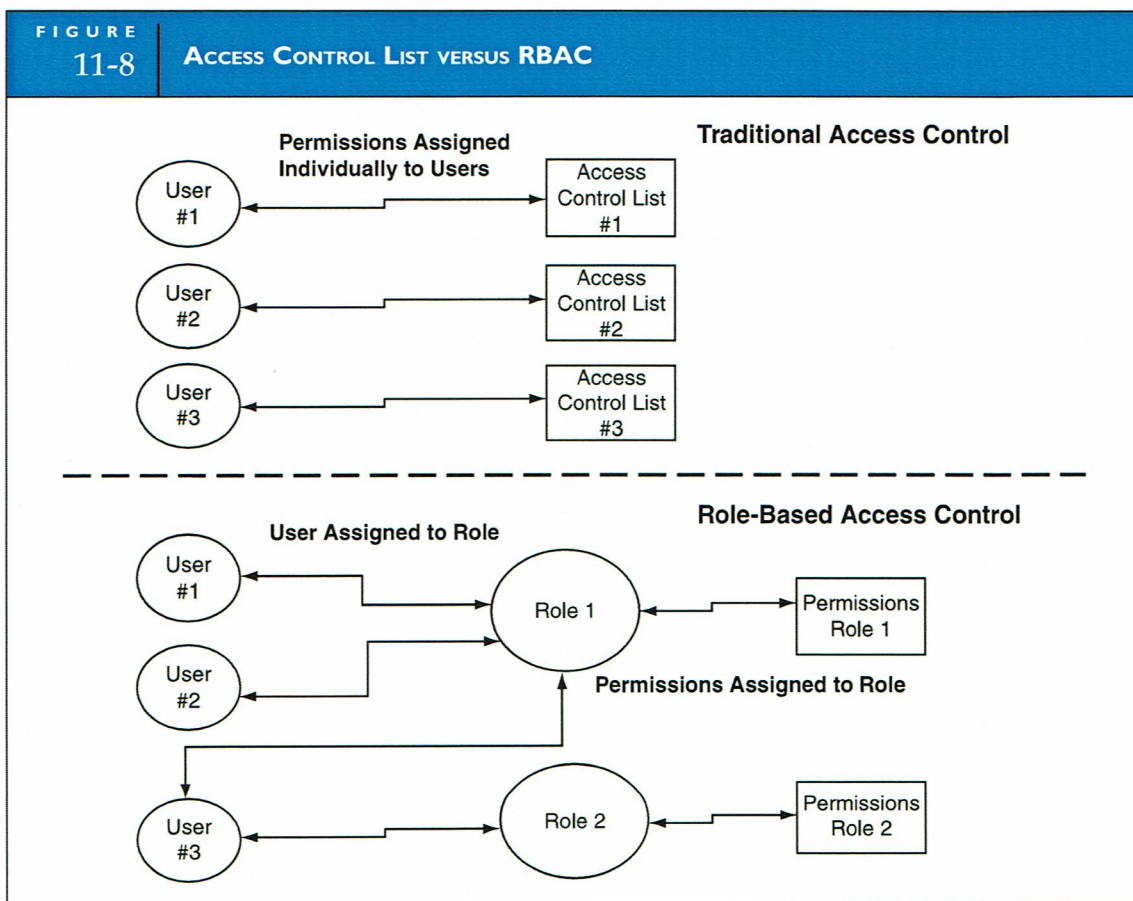
Access security is one of the most critical control issues in an ERP environment. The goal of ERP access control is to maintain data confidentiality, integrity, and availability. Security weaknesses can result in transaction errors, irregularities, data corruption, and financial statement misrepresentations. Also, uncontrolled access exposes organizations to cyber criminals who steal and subsequently sell critical data to competitors. Security administrators therefore need to control access to the tasks and operations that process or otherwise manipulate sensitive corporate data.

Traditional Access Control Models

Traditionally, the owner of system resources (data, functions, and processes) grants access privileges individually to users based on the individual's trust level and job description. Access control is typically achieved via an **access control list** (or access token) within the user's application.⁶ The access control list specifies the user ID, the resources available to the user, and the level of permission granted, such as read only, edit, or create. Although this model allows for the assignment of specific access privileges to individuals, it is quite inflexible. The sheer volume and variety of access privilege needs in modern ERP environments present a significant administrative burden. Any access-granting model must efficiently keep up with new hires, changes to existing privileges brought about by promotions and individuals transferring from one department to another, and personnel terminations. To meet these demands, modern ERP systems employ **role-based access control (RBAC)**, which is discussed next.

Role-Based Access Control (RBAC)

A **role** is a formal technique for grouping together users according to the system resources they need to perform their assigned tasks. For example, a system administrator could create a Sales Role for the sales department personnel that permits access only to the ERP's Sales Module and certain documents such as customer orders, sales orders, and customer records. When an employee joins the sales department (a new hire or transfer from another department), he or she will be assigned to the Sales Role and, through it, can access its prespecified resources. Figure 11-8 illustrates the difference between RBAC and the traditional access control list approach.



⁶ Access control list and access tokens are discussed in more detail in Chapter 16 within the broader context of general controls.

FIGURE 11-9		SAP ROLE DEFINITION	
ERP Role Name: AP Processing		Location: AP Department	
Role Goal: Data Entry of Vendor Invoices and Check Requests			
Knowledge, Skills: Knowledge of invoice processing from receipt to payment and relevant finance and procurement SAP modules, Vendor Master Data, P-card processes. Understanding of IRS 1099 reporting & processing requirements.			
Responsibilities/Tasks:			
--Review and verify parked invoices and submit to Accounts Payable --Park non-P.O. invoices (check request) --Enter invoices for goods received into SAP; code invoices for payment (park)			
SAP Security Considerations:			
AP	Transaction Type		
FB03	Display Document	FB04	Display Changes
FBD3	Display Recurring Entry	FBL1	
FBL1N	Display Vendor Line Item	FBL3N	Account Line Item Display
FBV2	Change Parked Document	FBV3	Display Parked Documents
FCH1	Display Check Information	FCH2	Display Payment Document checks
FCHN	Check Register	FK10N	Vendor Balance Display
FV60	Park/Edit Invoice	PV65	Park/Edit Credit Memo
MIR4	Display Invoice Document	MIR5	Display List of Invoice

© Cengage Learning®

Notice from the figure how this technique assigns access permissions to the role an individual plays in the organization rather than directly to the individual. Therefore, more than one individual can be assigned to a role and a predefined set of access permissions. Also, an individual may be assigned more than one role but can log into the system only under one role at a time. Thus, RBAC conveniently handles many-to-many relationships between users and permissions and facilitates dealing efficiently with vast numbers of employees.

ERPs come with predefined roles with preassigned permissions. Administrators and line managers may also create new roles, modify existing roles, and delete roles that are no longer needed. Creating a role involves defining the following role attributes:

1. A stated set of business responsibilities to be performed within the role
2. The technical competencies needed to perform the role
3. The specific transactions (permissions) required to carry out the stated responsibilities

Figure 11-9 presents a SAP role definition for an accounts payable role in a hypothetical organization. The individual(s) assigned this role is governed in his or her access to specific SAP program modules (in this case AP) and to specific activities within the module by the transaction code list in the *SAP Security Considerations* section of the role definition.

INTERNAL CONTROL ISSUES RELATED TO ERP ROLES

Although RBAC is an excellent mechanism for efficiently managing access control, the process of creating, modifying, and deleting roles is an internal control issue of concern for management and auditors alike. The following points highlight the key concerns:

1. The creation of unnecessary roles
2. The rule of least access should apply to permission assignments
3. Monitor role creation and permission-granting activities

The Creation of Unnecessary Roles

The fundamental objective of RBAC is to provide access in accordance with an organization's needs, which derive from defined tasks rather than an individual's wants. Managers in ERP environments, however, have significant discretion in creating new roles for individuals. This may be done for employees who need access to resources for special and/or one-time projects. Such access-granting authority needs to be tempered with judgment to prevent the number of roles from multiplying to the point of becoming dysfunctional and thus creating a control risk. Indeed, an oft-cited problem in ERP environments is that roles tend to proliferate to a point at which their numbers actually exceed the number of employees in the organization. Policies need to be in place to prevent the creation of unnecessary new roles and to ensure that temporary role assignments are deleted when the reason for them terminates.

The Rule of Least Access

Access privileges (permissions) should be granted on a need-to-know basis only. Nevertheless, ERP users tend to accumulate unneeded permissions over time. This is often due to two problems:

1. Managers fail to exercise adequate care in assigning permissions as part of their role-granting authority. Because managers are not always experts in internal controls, they may not recognize when excessive permissions are awarded to an individual.
2. Managers tend to be better at issuing privileges than removing them. As a result, an individual may retain unneeded access privileges from a previous job assignment that creates a segregation-of-duties violation when combined with a newly assigned role.

Policies should be in place to require managers to apply due diligence in assigning permissions to roles to avoid the granting of excessive access. They should assign privileges based on the task at hand and be aware of the individual's existing permissions that may pose a control violation.

Monitor Role Creation and Permission-Granting Activities

Effective RBAC management demands procedures that monitor role creation and permission granting to ensure compliance with internal control objectives. Verifying role compliance across all applications and users in an ERP environment, however, poses a highly complex and technical problem that does not lend itself to manual techniques. **Role-based governance** systems are available for this purpose. These systems allow managers to:

- View the current and historical inventory of roles, permissions granted, and the individuals assigned to roles.
- Identify unnecessary or inappropriate access entitlements and segregation-of-duties violations.
- Verify that changes to roles and entitlements have been successfully implemented.

These systems can continually monitor for risk and issue alerts when violations are detected so that remedial action can be taken. In addition, role-based governance can maintain an audit trail to provide a record of violations and evidence of compliance.

CONTINGENCY PLANNING

The implementation of an ERP creates an environment with a single point of failure, which places the organization at risk from equipment failure, sabotage, or natural disaster. To control this risk, an organization needs an effective contingency plan that can be invoked quickly in the event of a disaster. Two general approaches are outlined next, but a more extensive discussion is presented in Chapter 15 within the broader context of IT governance.

Centralized organizations with highly integrated business units may need a single global ERP system that is accessed via the Internet or private lines from around the world to consolidate data from subsidiary systems. A server failure under this model could leave the entire organization

unable to process transactions. To control this, two linked servers may be connected in redundant backup mode. All production processing is done on one server. If it fails, processing is automatically transferred to the other server. Organizations that want more security and resilience may arrange servers in a cluster of three or more that dynamically share the workload. Processing can be redistributed if one or more of the servers in the cluster fail.

Companies whose organizational units are autonomous and do not share common customers, suppliers, or product lines often choose to install regional servers. This approach permits independent processing and spreads the risk associated with server failure. For example, BP Amoco implemented SAP's R/3 into 17 separate business groups.

Summary

This chapter opened by comparing the function and data storage techniques of a traditional flat file or database system with that of an ERP. An important distinction was drawn between OLTP and OLAP applications. Similarly, the differences between the ERP's operational database and the data warehouse were discussed. Next, ERP configurations were examined relating to servers, databases, and bolt-on software. We discussed SCM as an area of contention. ERP vendors are moving quickly to provide SCM functionality. Simultaneously, SCM vendors are encroaching on traditional ERP territory. Data warehousing was the topic of the third section. A data warehouse is a relational or multidimensional database that supports OLAP. A number of data warehouse issues were discussed, including data modeling, data extraction from operational databases, data cleansing, data transformation, and data loading into the warehouse.

The fourth section examined common risks associated with ERP implementation. Among these are the risks associated with the big bang approach, internal opposition to changing the way a company does its business, choosing the wrong ERP model, choosing the wrong consultant, cost overrun issues, and disruptions to operations. Also presented were a number of issues to consider when implementing an ERP. These include selecting a system that is a good fit for the organization, understanding that the term *scalability* can mean different things to different people, potential problems associated with customizing the software, the need for assigning performance measures, and the need to control outside consultants. The chapter concluded with a review of the internal control and auditing issues related to ERPs.

Key Terms

access control list (491)
 big bang (485)
 bolt-on software (478)
 changed data capture (480)
 client-server model (474)
 closed database architecture (472)
 consolidation (477)
 core applications (473)
 data mart (479)
 data warehouse (474)
 drill-down (477)
 enterprise resource planning (ERP) (471)

online analytical processing (OLAP) (474)
 online transaction processing (OLTP) (473)
 phased-in (485)
 role (491)
 role-based access control (RBAC) (491)
 role-based governance (493)
 scalability (486)
 slicing and dicing (477)
 supply chain management (SCM) (478)
 three-tier model (475)
 two-tier model (474)

Review Questions

1. Define ERP.
2. What is closed database architecture?
3. Define *core applications*, and give some examples.
4. Define OLAP, and give some examples.