Database Principles: Fundamentals of Design, Implementation, and Management

Tenth Edition

Chapter 10

Distributed Databases

Objectives

In this chapter, you will learn:

- About distributed database management systems (DDBMSs) and their components
- How database implementation is affected by different levels of data and process distribution
- How transactions are managed in a distributed database environment

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Objectives (cont’d.)

- How distributed database design draws on data partitioning and replication to balance performance, scalability, and availability
- About the trade-offs of implementing a distributed data system

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

The Evolution of Distributed Database Management Systems

- Distributed database management system (DDBMS)
  - Governs storage and processing of logically related data
  - Interconnected computer systems
  - Both data and processing functions are distributed among several sites
- Centralized database required that corporate data be stored in a single central site

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.1 Centralized database management system

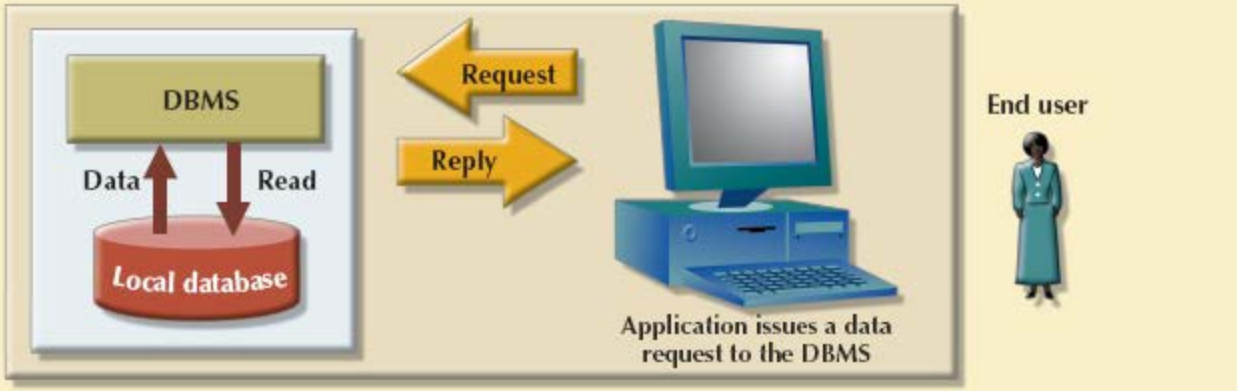

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

DDBMS Advantages and Disadvantages

- Advantages:
  - Data are located near “greatest demand” site
  - Faster data access
  - Faster data processing
  - Growth facilitation
  - Improved communications
  - Reduced operating costs
  - User-friendly interface
  - Less danger of a single-point failure
  - Processor independence

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

DDBMS Advantages and Disadvantages (cont'd.)

- Disadvantages:
  - Complexity of management and control
  - Security
  - Lack of standards
  - Increased storage requirements
  - Increased training cost
  - Costs (duplicate hardware, licensing, etc.)

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

TABLE 10.1 Distributed DBMS Advantages and Disadvantages

ADVANTAGES	DISADVANTAGES
• *Data are located near the site of greatest demand.* The data in a distributed database system are dispersed to match business requirements.	• *Complexity of management and control.* Applications must recognize data location, and they must be able to stitch together data from various sites. Database administrators must have the ability to coordinate database activities to prevent database degradation due to data anomalies.
• *Faster data access.* End users often work with only the nearest stored subset of the data.	• *Technological difficulty.* Data integrity, transaction management, concurrency control, security, backup, recovery, and query optimization must all be addressed and resolved.
• *Faster data processing.* A distributed database system spreads out the system's workload by processing data at several sites.	• *Security.* The probability of security lapses increases when data are located at multiple sites. The responsibility of data management will be shared by different people at several sites.
• *Growth facilitation.* New sites can be added to the network without affecting the operations of other sites.	• *Lack of standards.* There are no standard communication protocols at the database level. For example, different database vendors employ different and often incompatible techniques to manage the distribution of data and processing in a DDBMS environment.
• *Improved communications.* Because local sites are smaller and located closer to customers, local sites foster better communication among departments and between customers and company staff.	• *Increased storage and infrastructure requirements.* Multiple copies of data are required at different sites, thus requiring additional storage space.
• *Reduced operating costs.* It is more cost-effective to add nodes to a network than to update a mainframe system. Development work is done more cheaply and quickly on low-cost PCs than on mainframes.	• *Increased training cost.* Training costs are generally higher in a distributed model than they would be in a centralized model, sometimes even to the extent of offsetting operational and hardware savings.
• *User-friendly interface.* PCs and workstations are usually equipped with an easy-to-use graphical user interface (GUI). The GUI simplifies training and use for end users.	• *Costs.* Distributed databases require duplicated infrastructure to operate, such as physical location, environment, personnel, software, and licensing.
• *Less danger of a single-point failure.* When one of the computers fails, the workload is picked up by other workstations. Data are also distributed at multiple sites.	
• *Processor independence.* The end user can access any available copy of the data, and an end user's request is processed by any processor at the data location.	

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Distributed Processing and Distributed Databases

- Distributed processing
  - Database's logical processing is shared among two or more physically independent sites
  - Connected through a network
- Distributed database
  - Stores logically related database over two or more physically independent sites
  - Database composed of database fragments

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.2 Distributed processing environment

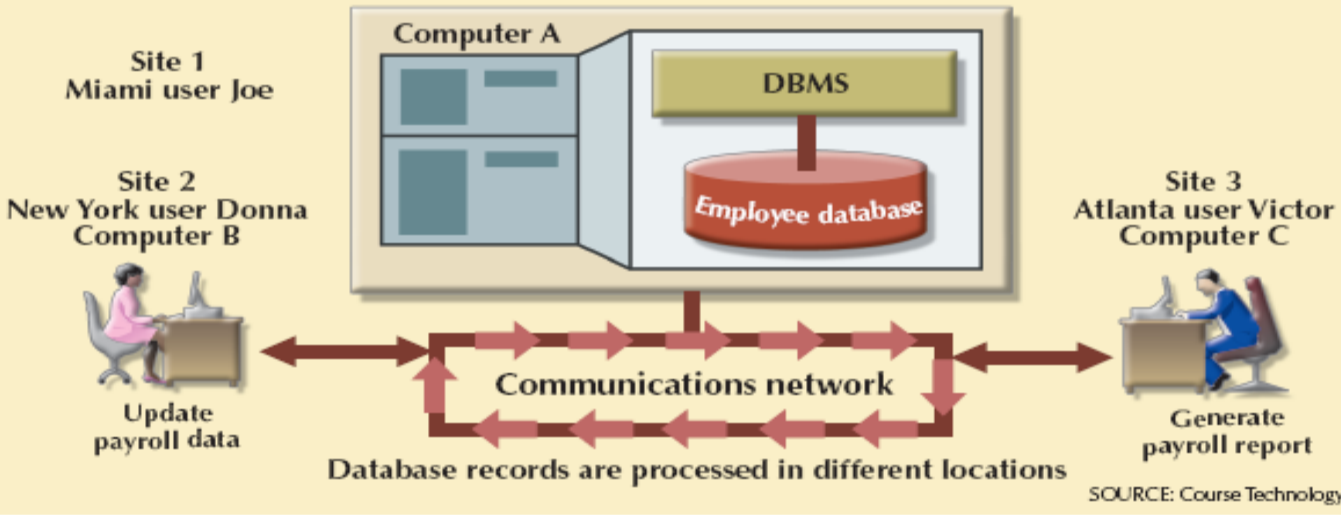

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.3 Distributed database environment

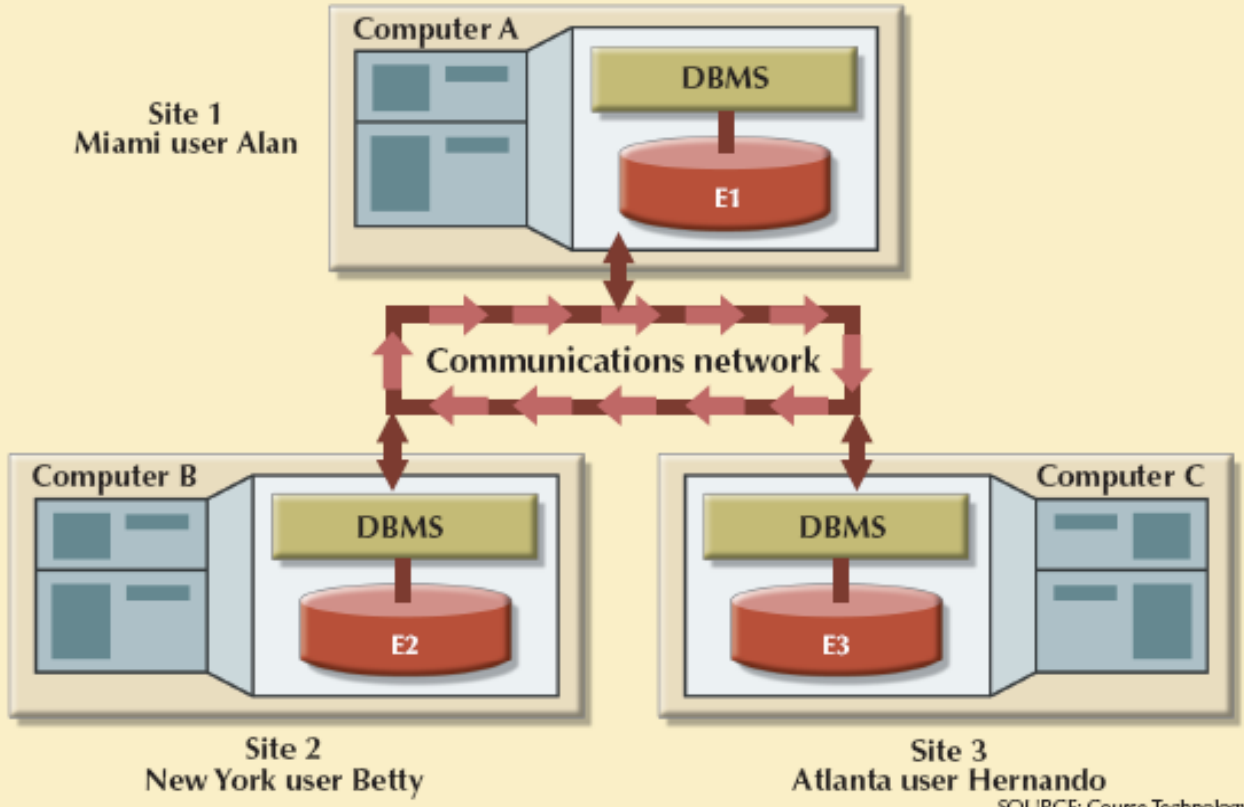

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Characteristics of Distributed Management Systems

- Application interface
- Validation
- Transformation
- Query optimization
- Mapping
- I/O interface

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Characteristics of Distributed Management Systems (cont'd.)

- Formatting
- Security
- Backup and recovery
- DB administration
- Concurrency control
- Transaction management

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Characteristics of Distributed Management Systems (cont'd.)

- Must perform all the functions of centralized DBMS
- Must handle all necessary functions imposed by distribution of data and processing
  - Must perform these additional functions transparently to the end user

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.4 A fully distributed database management system

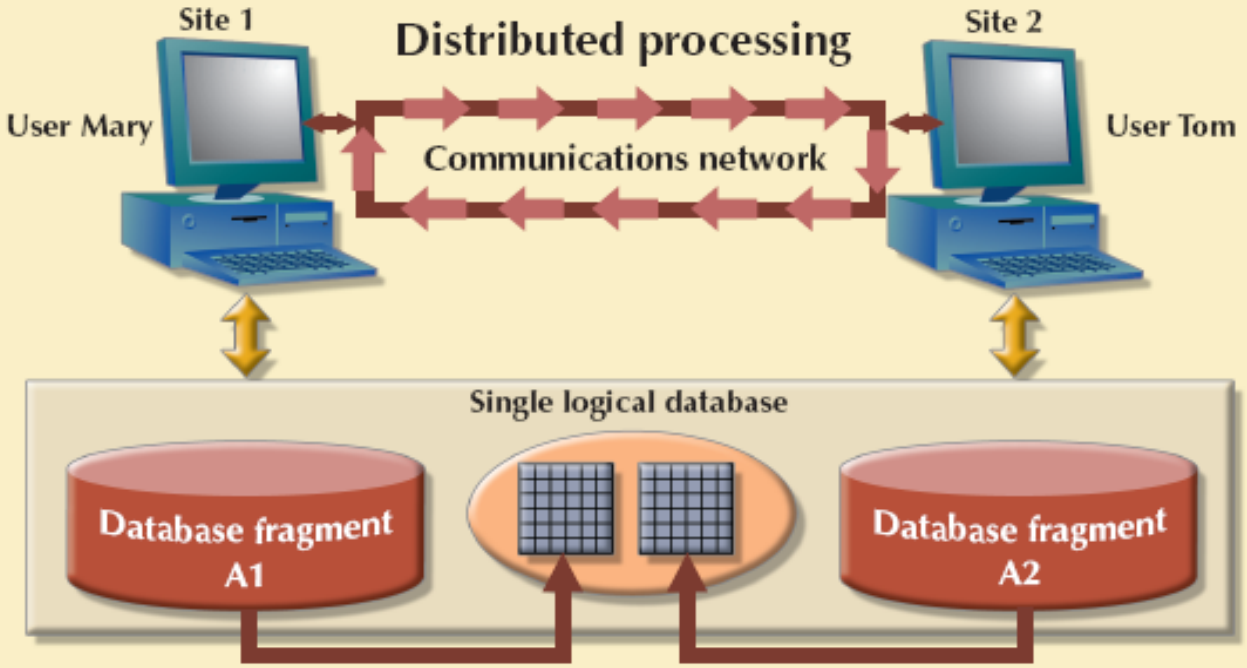

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

DDBMS Components

- Must include (at least) the following components:
  - Computer workstations
  - Network hardware and software
  - Communications media
  - Transaction processor (application processor, transaction manager)
    - Software component found in each computer that requests data

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

DDBMS Components (cont'd.)

- Data processor or data manager
  - Software component residing on each computer that stores and retrieves data located at the site
  - May be a centralized DBMS

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.5 Distributed database system management components

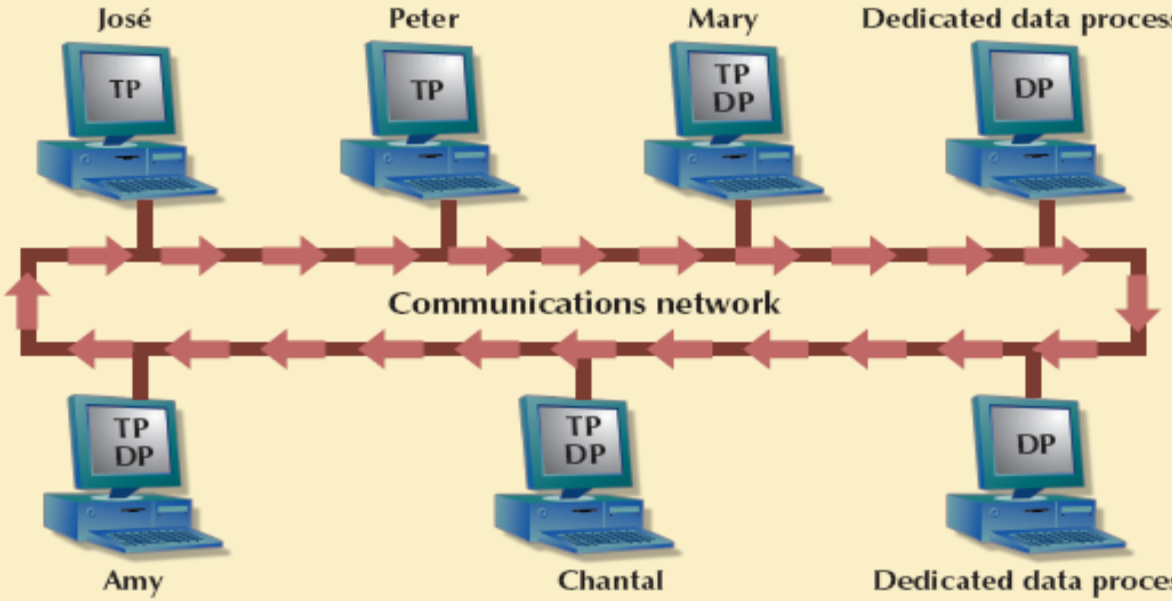

Note: Each TP can access data on any DP, and each DP handles all requests for local data from any TP.

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Levels of Data and Process Distribution

- Current systems classified by how process distribution and data distribution are supported

TABLE 10.2 Database Systems: Levels of Data and Process Distribution

	SINGLE-SITE DATA	MULTIPLE-SITE DATA
Single-site process	Host DBMS	Not applicable (Requires multiple processes)
Multiple-site process	File server Client/server DBMS (LAN DBMS)	Fully distributed Client/server DDBMS

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Single-Site Processing, Single-Site Data

- All processing is done on single CPU or host computer (mainframe, midrange, or PC)
- All data are stored on host computer's local disk
- Processing cannot be done on end user's side of system
- Typical of most mainframe and midrange computer DBMSs
- DBMS is located on host computer, which is accessed by dumb terminals connected to it

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.6 Single-site processing, single-site data (centralized)

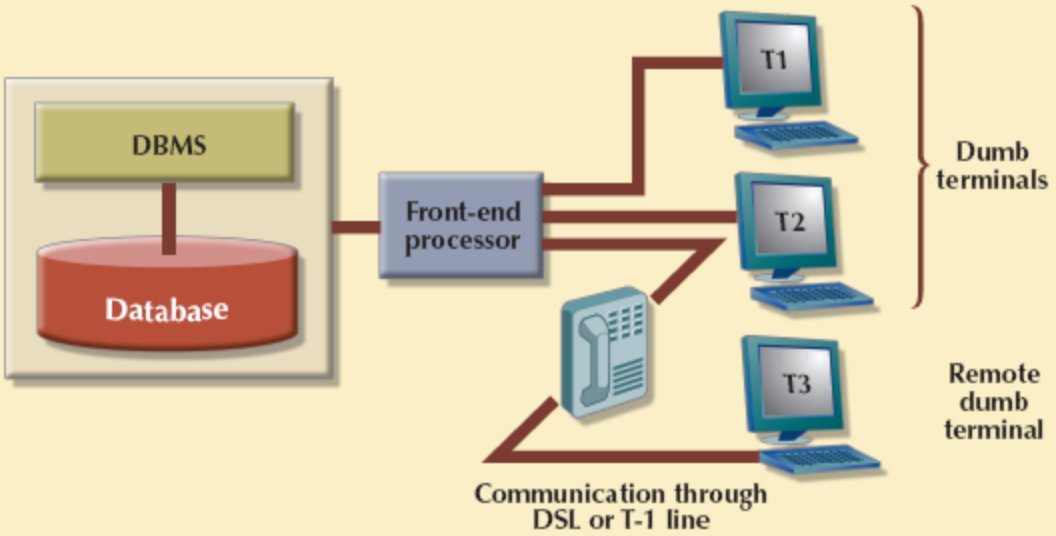

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Multiple-Site Processing, Single-Site Data

- Multiple processes run on different computers sharing single data repository
- MPSD scenario requires network file server running conventional applications
  - Accessed through LAN
- Many multiuser accounting applications, running under personal computer network

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.7 Multiple-site processing, single-site data

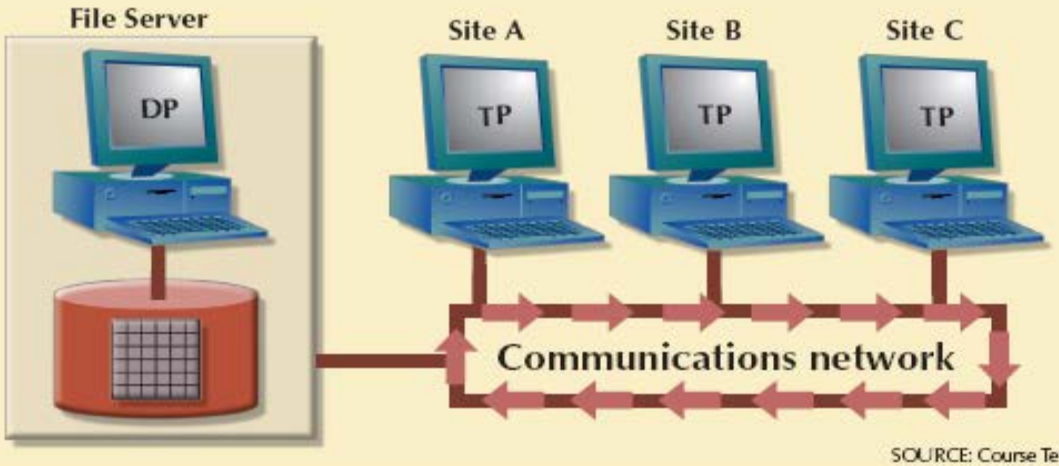

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Multiple-Site Processing, Multiple-Site Data

- Fully distributed database management system
- Support for multiple data processors and transaction processors at multiple sites
- Classified as either homogeneous or heterogeneous
- Homogeneous DDBMSs
  - Integrate only one type of centralized DBMS over a network

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Multiple-Site Processing, Multiple-Site Data (cont'd.)

- Heterogeneous DDBMSs
  - Integrate different types of centralized DBMSs over a network
- Fully heterogeneous DDBMSs
  - Support different DBMSs
  - Support different data models (relational, hierarchical, or network)
  - Different computer systems, such as mainframes and microcomputers

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

TABLE 10.3 **Heterogeneous Distributed Database Scenario**

PLATFORM	DBMS	OPERATING SYSTEM	NETWORK COMMUNICATIONS PROTOCOL
IBM 3090	DB2	MVS	APPC LU 6.2
IBM AS/400	SQL/400	OS/400	3270
RISC computer	Informix	UNIX	TCP/IP
Intel Xeon CPU	Oracle	Windows Server 2008	TCP/IP

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Distributed Database Transparency Features

- Allow end user to feel like database's only user
- Features include:
  - Distribution transparency
  - Transaction transparency
  - Failure transparency
  - Performance transparency
  - Heterogeneity transparency

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Distribution Transparency

- Allows management of physically dispersed database as if centralized
- Three levels of distribution transparency:
  - Fragmentation transparency
  - Location transparency
  - Local mapping transparency

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

TABLE 10.4 **Summary of Transparency Features**

IF THE SQL STATEMENT REQUIRES:			
FRAGMENT NAME?	LOCATION NAME?	THEN THE DBMS SUPPORTS	LEVEL OF DISTRIBUTON TRANSPARENCY
Yes	Yes	Local mapping transparency	Low
Yes	No	Location transparency	Medium
No	No	Fragmentation transparency	High

FIGURE 10.8 **Fragment locations**

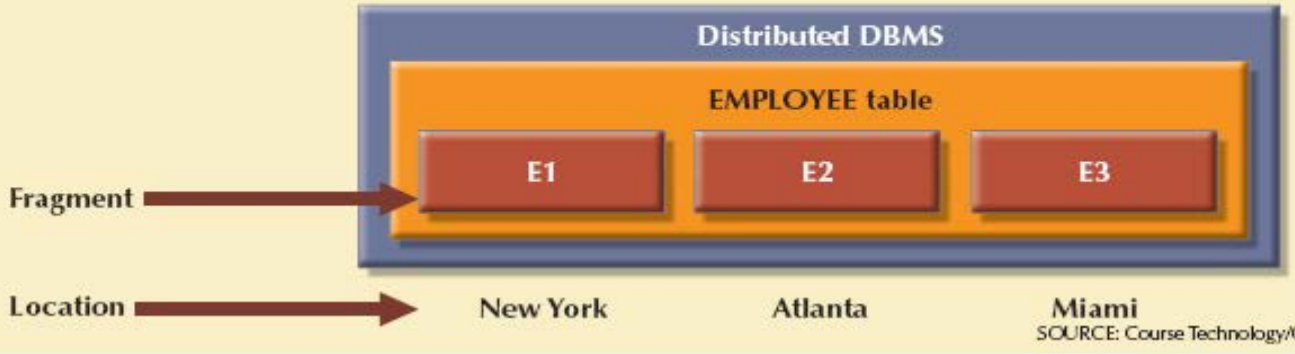

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Transaction Transparency

- Ensures database transactions will maintain distributed database's integrity and consistency
- Ensures transaction completed only when all database sites involved complete their part
- Distributed database systems require complex mechanisms to manage transactions
  - To ensure consistency and integrity

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Distributed Requests and Distributed Transactions

- Remote request: single SQL statement accesses data from single remote database
- Remote transaction: accesses data at single remote site
- Distributed transaction: requests data from several different remote sites on network
- Distributed request: single SQL statement references data at several DP sites

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Distributed Concurrency Control

- Concurrency control is important in distributed environment
  - Multisite multiple-process operations create inconsistencies and deadlocked transactions

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

FIGURE 10.14 The effect of a premature COMMIT

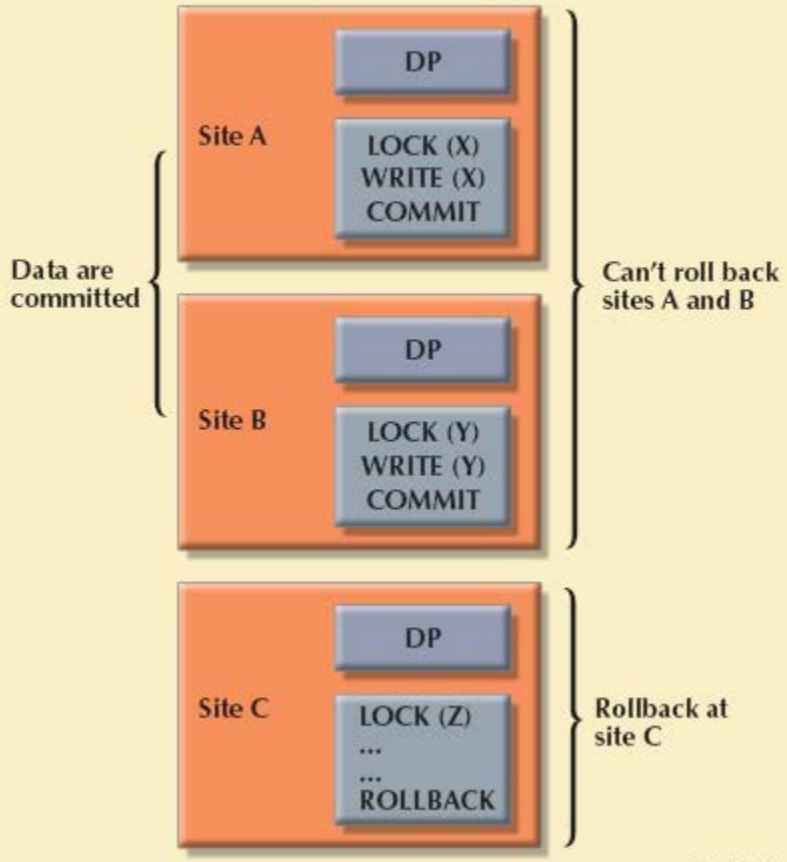

SOURCE: Course Technology/Cengage Learning

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Two-Phase Commit Protocol

- Distributed databases make it possible for transaction to access data at several sites
- Final COMMIT is issued after all sites have committed their parts of transaction
- Requires that each DP's transaction log entry be written before database fragment updated
- DO-UNDO-REDO protocol with write-ahead protocol
- Defines operations between coordinator and subordinates

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Performance and Failure Transparency

- Performance transparency
  - Allows a DDBMS to perform as if it were a centralized database
- Query optimization
  - Minimize the total cost associated with the execution of a request
- Replica transparency
  - DDBMS's ability to hide multiple copies of data from the user

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Performance and Failure Transparency (cont'd.)

- Network latency
  - Delay imposed by the amount of time required for a data packet to make a round trip from point A to point B
- Network partitioning
  - Delay imposed when nodes become suddenly unavailable due to a network failure

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Distributed Database Design

- Data fragmentation
  - How to partition database into fragments
- Data replication
  - Which fragments to replicate
- Data allocation
  - Where to locate those fragments and replicas

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Data Fragmentation

- Breaks single object into two or more segments or fragments
- Each fragment can be stored at any site over computer network
- Information stored in distributed data catalog (DDC)
  - Accessed by TP to process user requests

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Data Fragmentation (cont'd.)

- Strategies
  - Horizontal fragmentation
    - Division of a relation into subsets (fragments) of tuples (rows)
  - Vertical fragmentation
    - Division of a relation into attribute (column) subsets
  - Mixed fragmentation
    - Combination of horizontal and vertical strategies

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Data Replication

- Data copies stored at multiple sites served by computer network
- Fragment copies stored at several sites to serve specific information requirements
  - Enhance data availability and response time
  - Reduce communication and total query costs
- Mutual consistency rule: all copies of data fragments must be identical

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Data Replication (cont'd.)

- Fully replicated database
  - Stores multiple copies of each database fragment at multiple sites
  - Can be impractical due to amount of overhead
- Partially replicated database
  - Stores multiple copies of some database fragments at multiple sites
- Unreplicated database
  - Stores each database fragment at single site
  - No duplicate database fragments

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Data Allocation

- Deciding where to locate data
  - Centralized data allocation
    - Entire database is stored at one site
  - Partitioned data allocation
    - Database is divided into several disjointed parts (fragments) and stored at several sites
  - Replicated data allocation
    - Copies of one or more database fragments are stored at several sites

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

The CAP Theorem

- Initials CAP stand for three desirable properties
  - Consistency
  - Availability
  - Partition tolerance
- Basically available, soft state, eventually consistent (BASE)
  - Data changes are not immediate but propagate slowly through the system until all replicas are eventually consistent

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

TABLE 10.8 **Distributed Database Spectrum**

DBMS TYPE	CONSISTENCY	AVAILABILITY	PARTITION TOLERANCE	TRANSACTION MODEL	TRADE-OFF
Centralized DBMS	High	High	N/A	ACID	No distributed data processing
Relational DDBMS	High	Relaxed	High	ACID (2PC)	Sacrifices availability to ensure consistency and isolation
NoSQL DDBMS	Relaxed	High	High	BASE	Sacrifices consistency to ensure availability

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

C. J. Date's Twelve Commandments for Distributed Databases

- Local site independence
- Central site independence
- Failure independence
- Location transparency
- Fragmentation transparency
- Replication transparency

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

C. J. Date's Twelve Commandments for Distributed Databases (cont'd.)

- Distributed query processing
- Distributed transaction processing
- Hardware independence
- Operating system independence
- Network independence
- Database independence

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Summary

- Distributed database: logically related data in two or more physically independent sites
  - Connected via computer network
- Distributed processing: division of logical database processing among network nodes
- Distributed databases require distributed processing
- Main components of DDBMS are transaction processor and data processor

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Summary (cont'd.)

- Current distributed database systems
  - SPSD, MPSD, MPMD
- Homogeneous distributed database system
  - Integrates one type of DBMS over computer network
- Heterogeneous distributed database system
  - Integrates several types of DBMS over computer network

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Summary (cont'd.)

- DDBMS characteristics are a set of transparencies
- Transaction is formed by one or more database requests
- Distributed concurrency control is required in network of distributed databases
- Distributed DBMS evaluates every data request
  - Finds optimum access path in distributed database

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.

Summary (cont'd.)

- The design of distributed database must consider fragmentation and replication of data
- Database can be replicated over several different sites on computer network

© 2013 Cengage Learning. All Rights Reserved. This edition is intended for use outside of the U.S. only, with content that may be different from the U.S. Edition. May not be scanned, copied, duplicated, or posted to a publicly accessible website, in whole or in part.