

Database Management

DBMS Functions

Objectives

- Introduce the functions, or services, provided by a DBMS
- Describe how a DBMS handles updating and retrieving data
- Examine the catalog feature of a DBMS
- Illustrate the concurrent update problem and describe how a DBMS handles this problem
- Explain the data recovery process in a database environment

Objectives (continued)

- Describe the security services provided by a DBMS
- Examine the data integrity features provided by a DBMS
- Discuss the extent to which a DBMS achieves data independence
- Define and describe data replication
- Present the utility services provided by a DBMS

Introduction

- Functions of a DBMS
 - Update and retrieve data
 - Provide catalog services
 - Support concurrent update
 - Recover data
 - Provide security services
 - Provide data integrity features
 - Support data independence
 - Support data replication
 - Provide utility services

Update and Retrieve Data

- Fundamental capability of a DBMS
- Users don't need to know how data is stored or manipulated
- Users add, change, and delete records during updates

Update and Retrieve Data (continued)

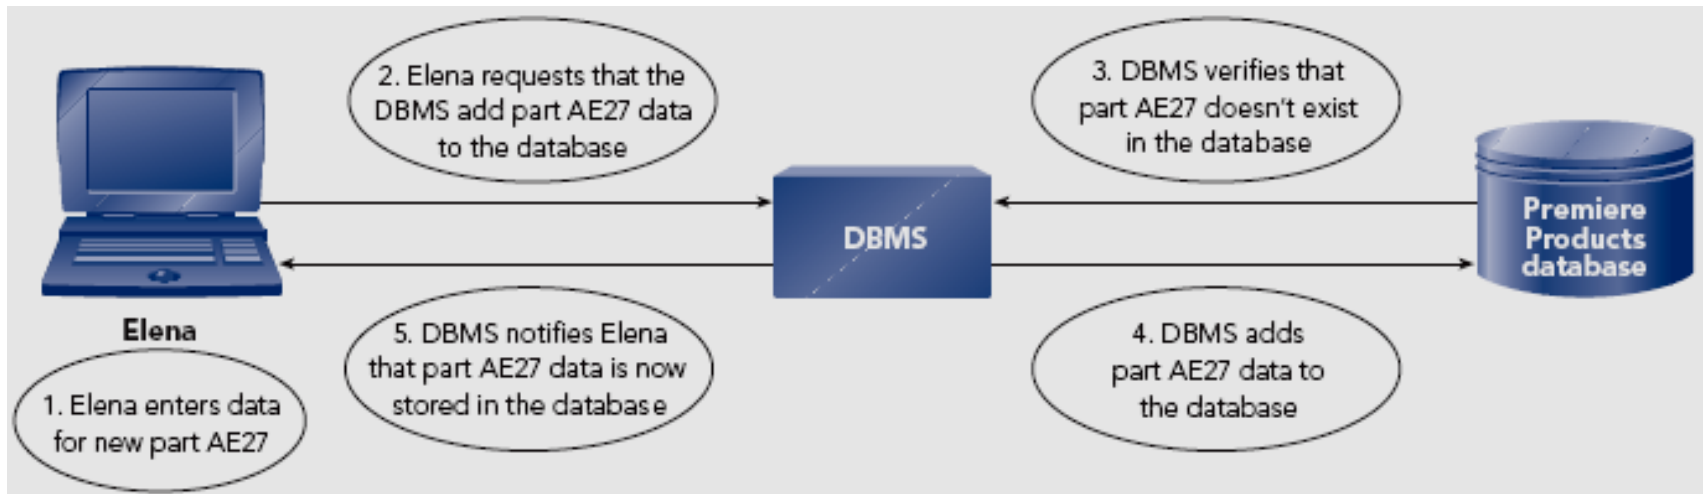


FIGURE 7-1: Adding a new part to the Premiere Products database

Update and Retrieve Data (continued)

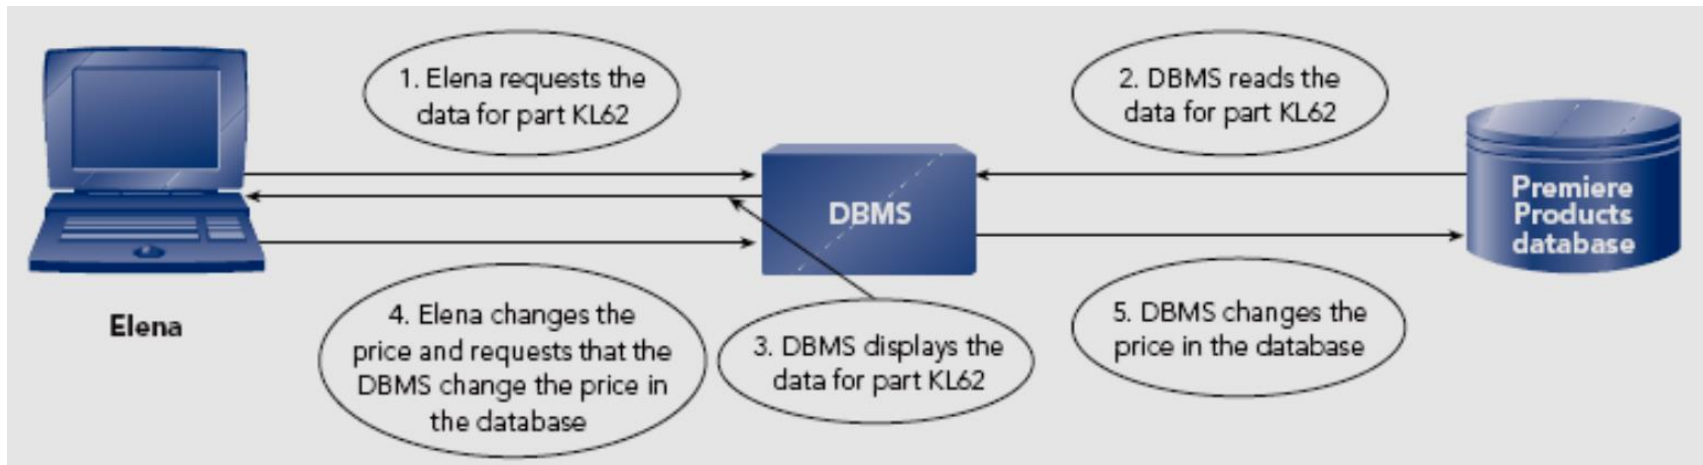


FIGURE 7-2: Changing the price of a part in the Premiere Products database

Update and Retrieve Data (continued)

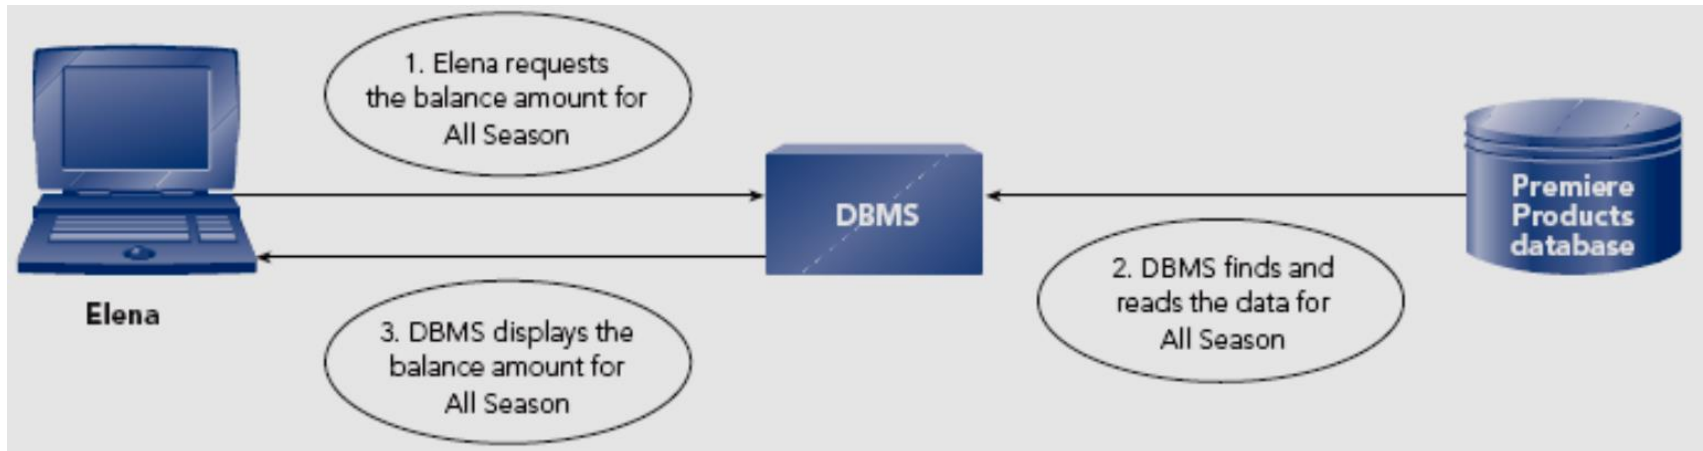


FIGURE 7-3: Retrieving a balance amount from the Premiere Products database

Provide Catalog Services

- **Metadata:** data about data
- Stores metadata and makes it accessible to users
- Enterprise DBMSs often have a **data dictionary** (a super catalog)

Support Concurrent Update

- Ensures accuracy when several users update database at the same time
- Manages complex scenarios for updates
- **Concurrent update**: multiple users make updates to the same database at the same time

The Concurrent Update Problem

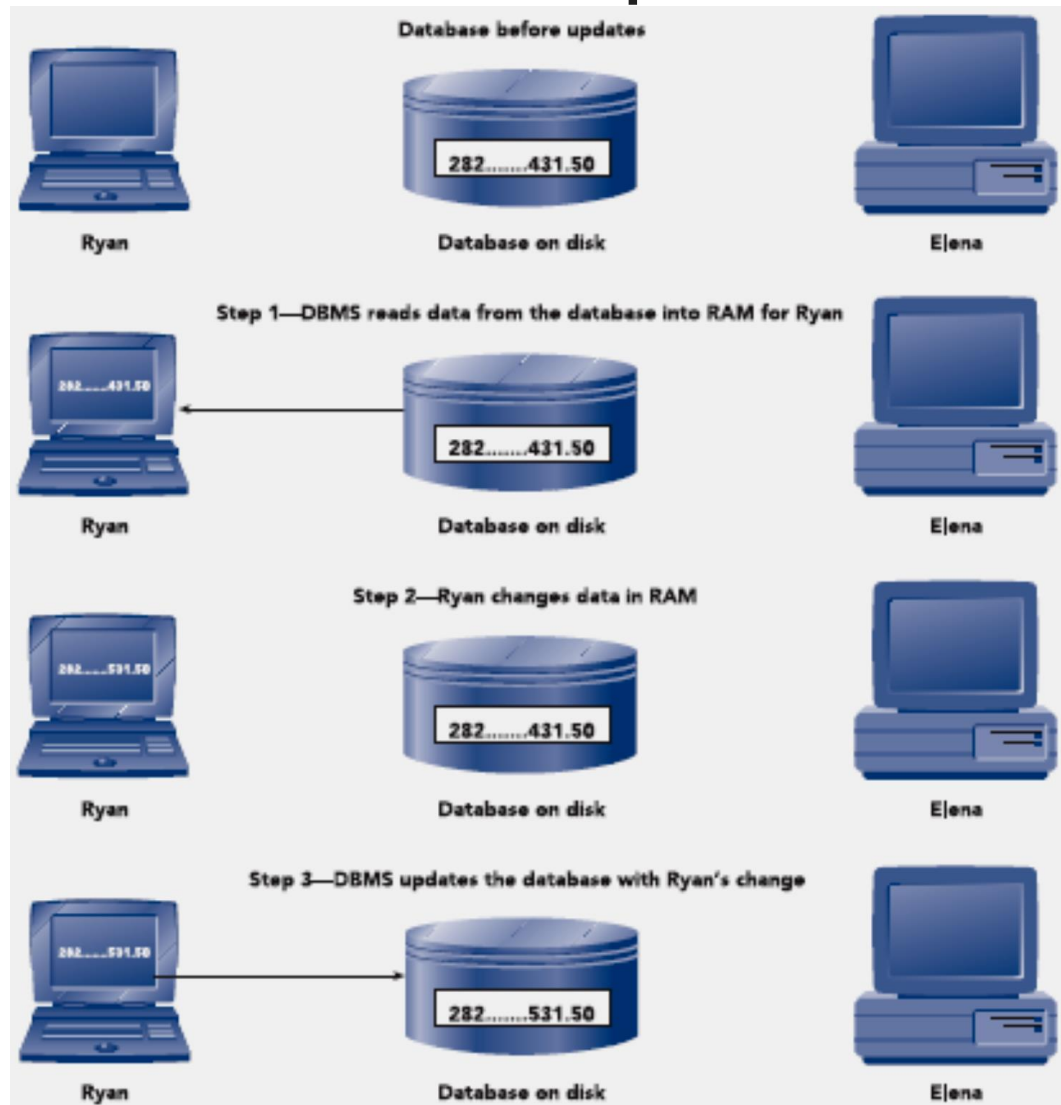


FIGURE 7-4: Ryan updates the database

The Concurrent Update Problem (continued)

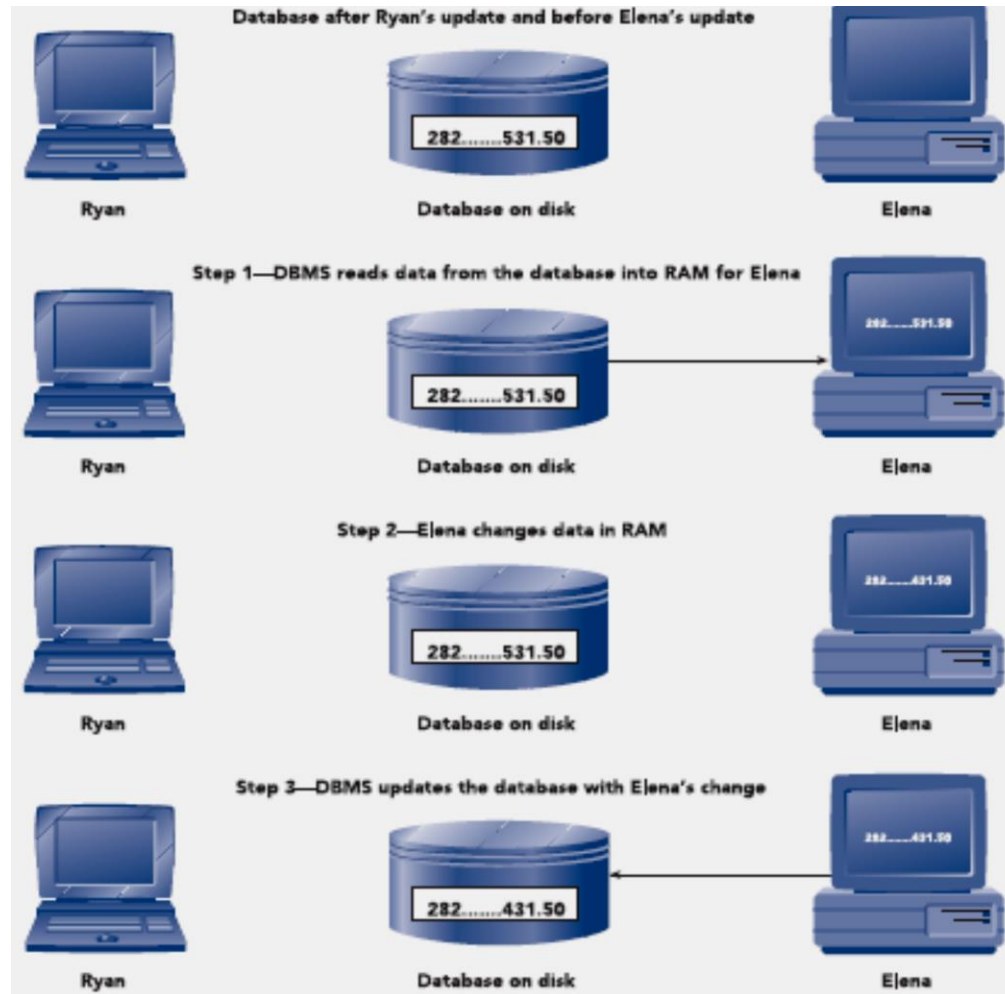


FIGURE 7-5: Elena updates the database

The Concurrent Update Problem (continued)

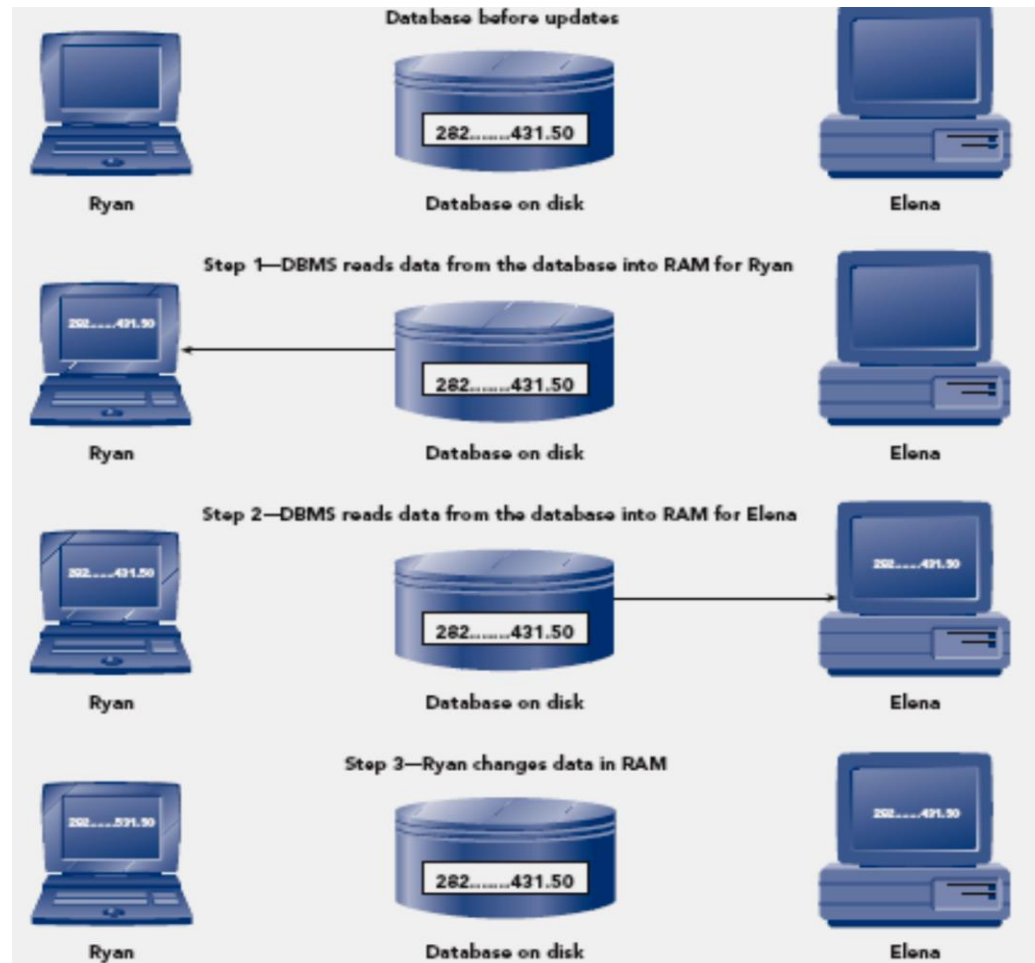


FIGURE 7-6: Ryan's and Elena's updates to the database result in a lost update
Concepts of Database Management

The Concurrent Update Problem (continued)

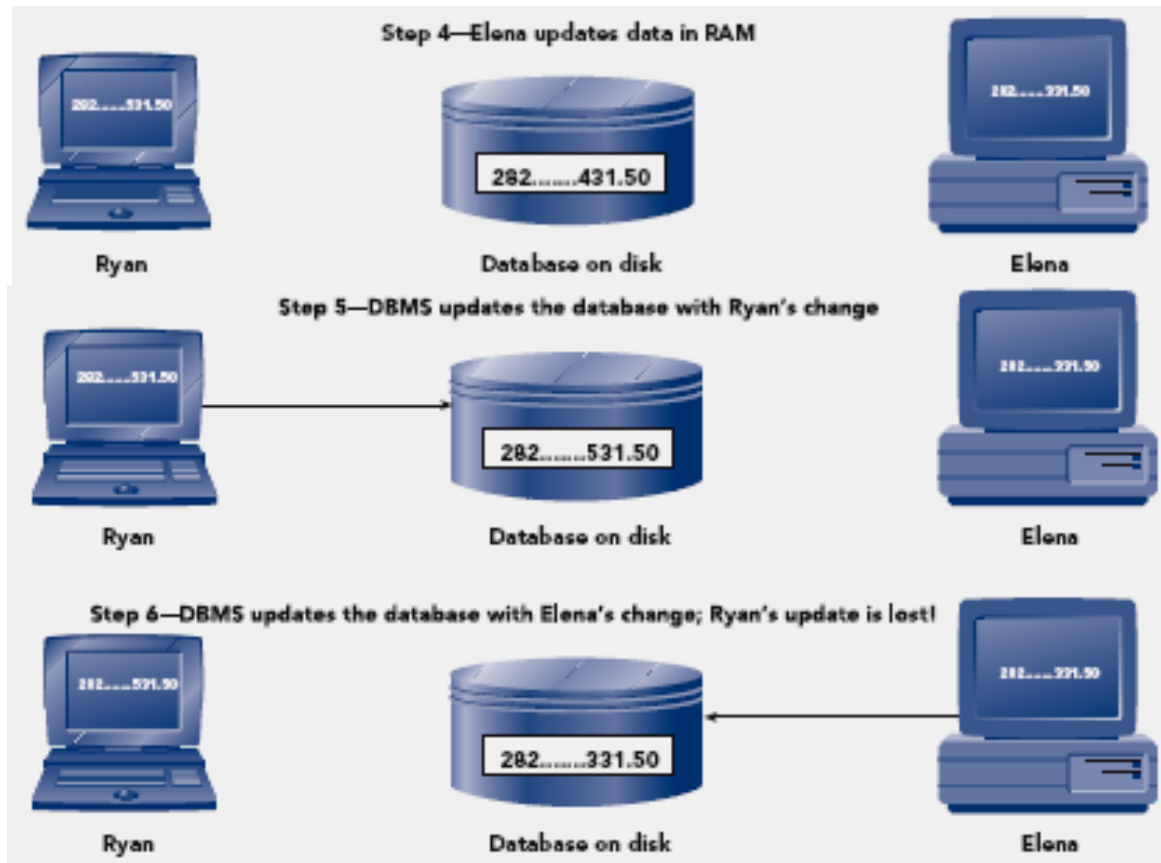


FIGURE 7-6: Ryan's and Elena's updates to the database result in a lost update (continued)

Avoiding the Lost Update Problem

- **Batch processing**
 - All updates done through a special program
 - Problem: data becomes out of date
 - Does not work in situations that require data to be current

Avoiding the Lost Update Problem (continued)

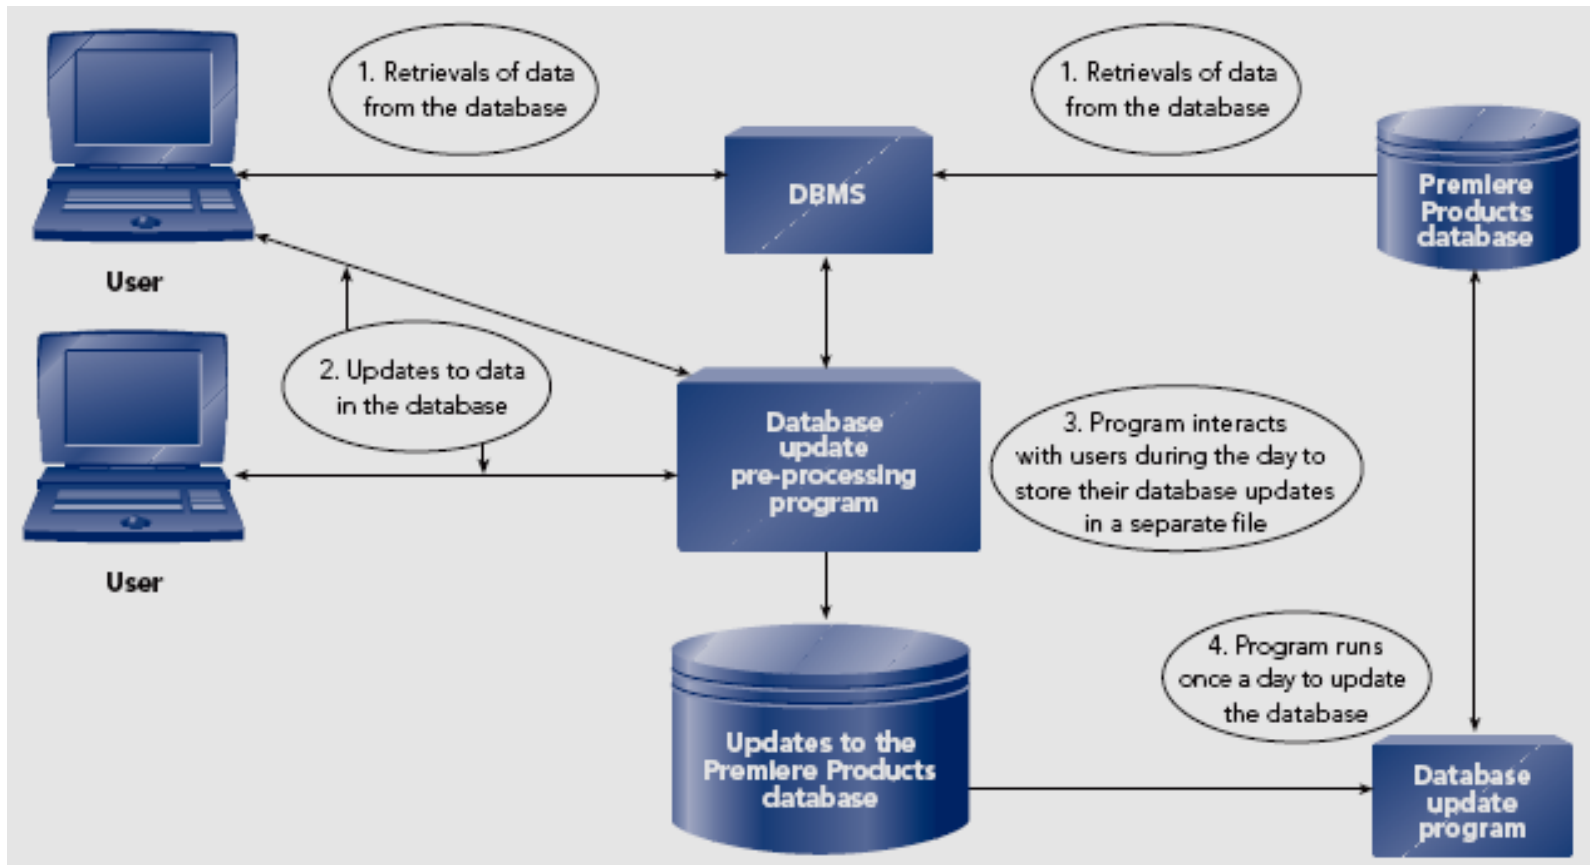


FIGURE 7-7: Delaying updates to the Premiere Products database to avoid the lost update problem

Two-Phase Locking

- **Locking:** deny other users access to data while one user's updates are being processed
- **Transaction:** set of steps completed by a DBMS to accomplish a single user task
- **Two-phase locking** solves lost update problem
 - **Growing phase:** DBMS locks more rows and releases none of the locks
 - **Shrinking phase:** DBMS releases all the locks and acquires no new locks

Two-Phase Locking (continued)

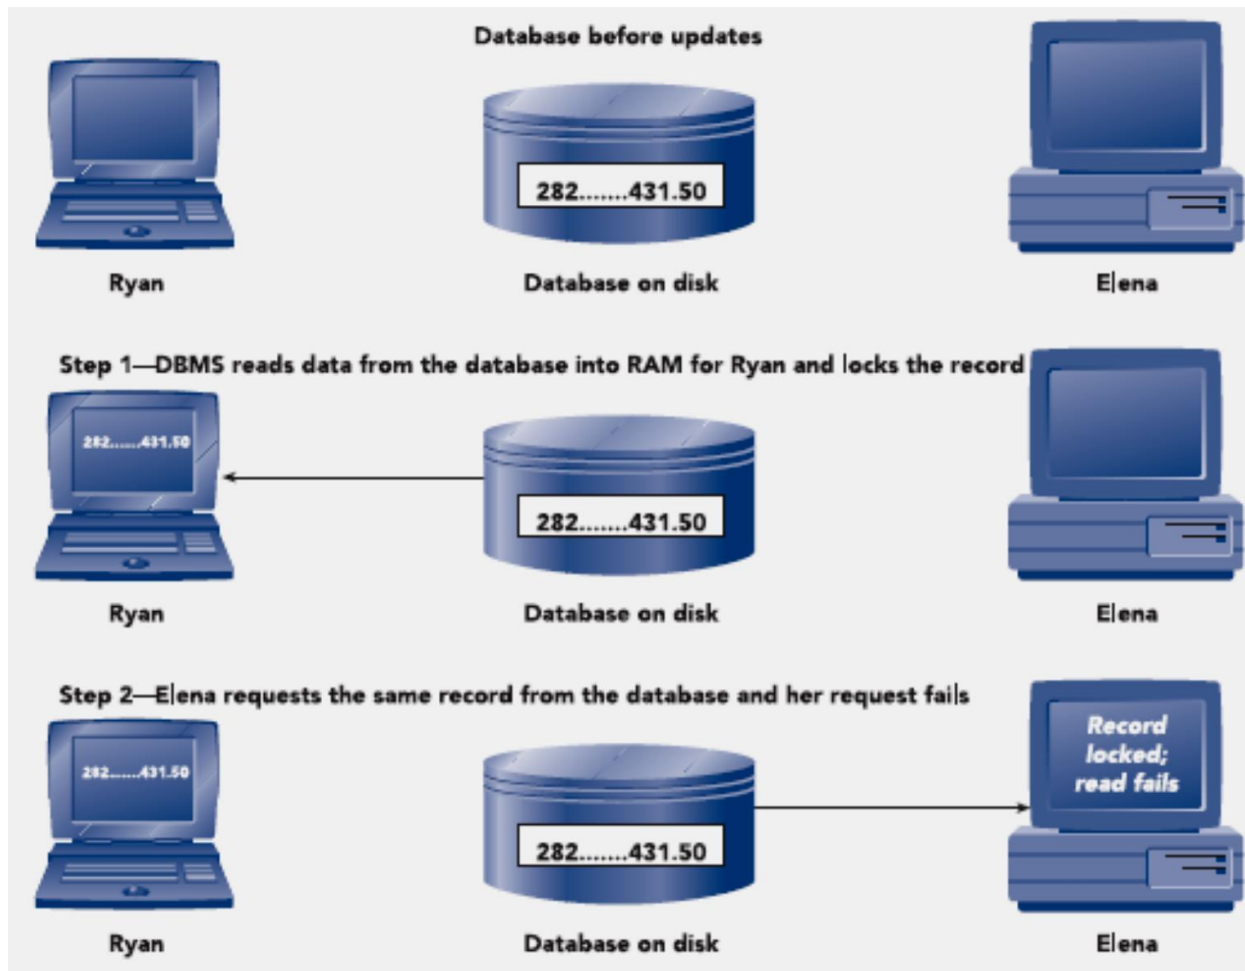


FIGURE 7-8: The DBMS uses a locking scheme to apply Ryan's and Elena's updates to the database

Two-Phase Locking (continued)

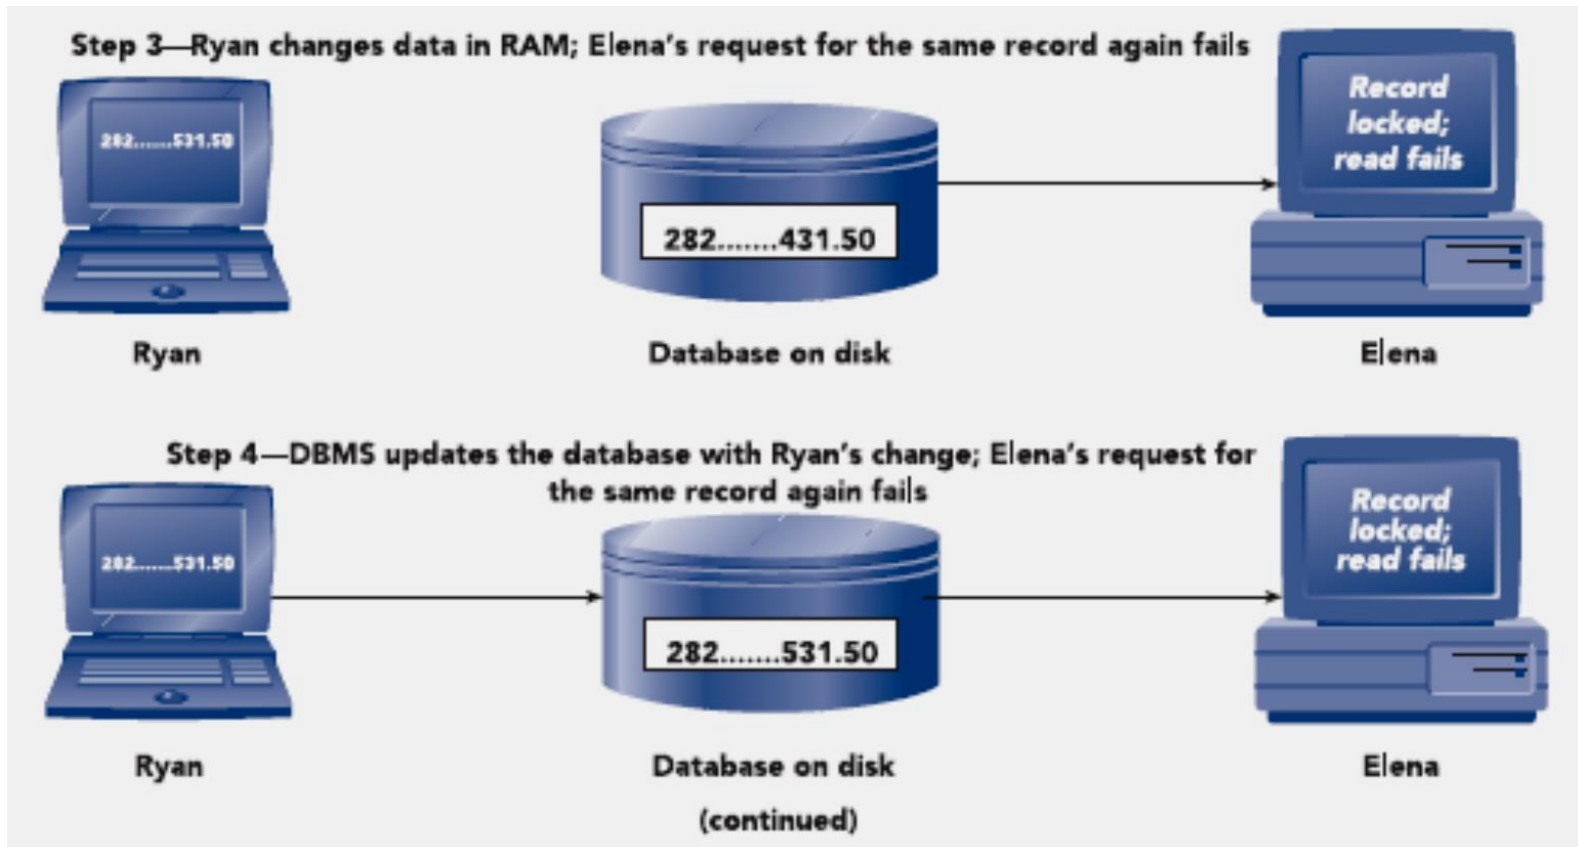


FIGURE 7-8: The DBMS uses a locking scheme to apply Ryan's and Elena's updates to the database (continued)

Two-Phase Locking (continued)

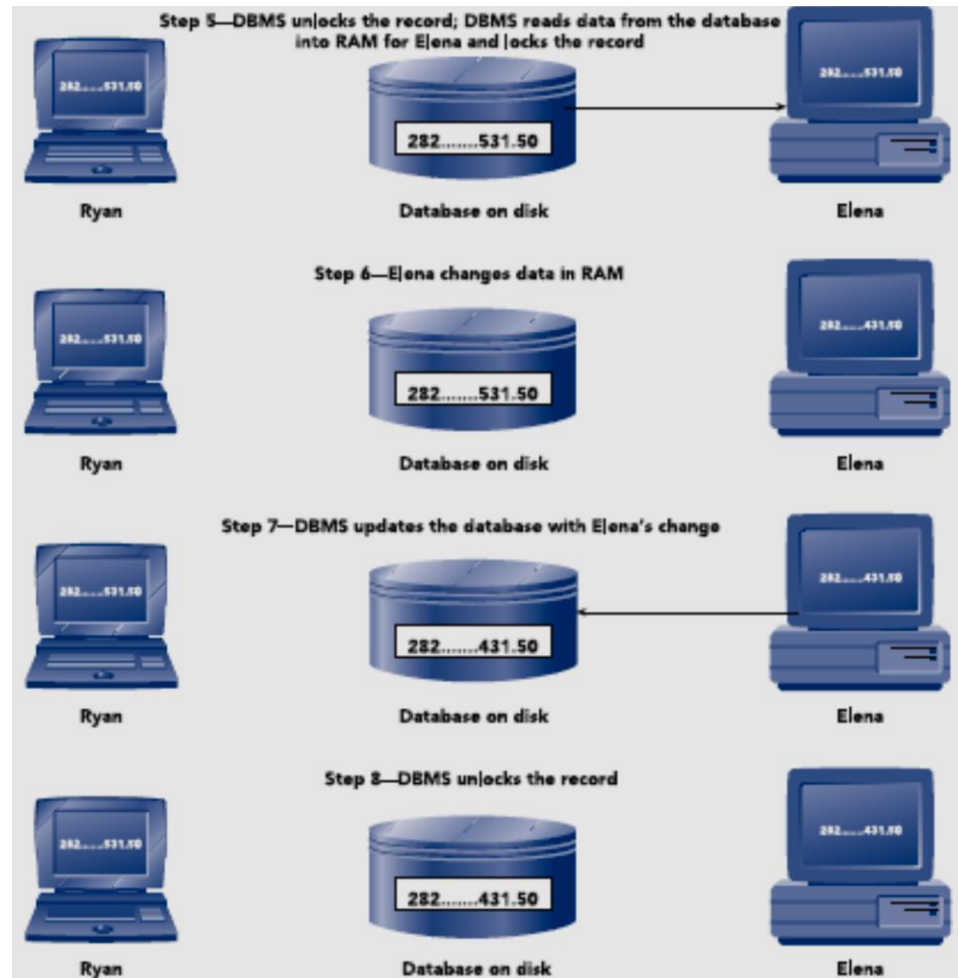


FIGURE 7-8: The DBMS uses a locking scheme to apply Ryan's and Elena's updates to the database (continued)

Deadlock

- **Deadlock or deadly embrace**
 - Two users hold a lock and require a lock on the resource that the other already has
 - To minimize occurrence, make sure all programs lock records in the same order whenever possible
- **Managing deadlocks**
 - DBMS detects and breaks any deadlock
 - DBMS chooses one user to be the **victim**

Deadlock (continued)

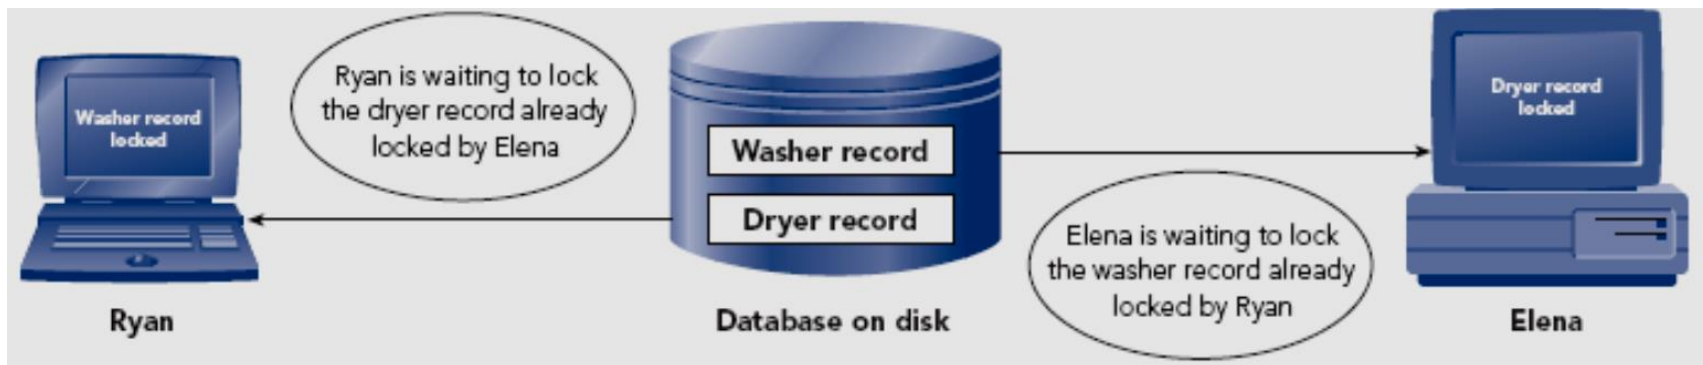


FIGURE 7-9: Two users experiencing deadlock

Locking on PC-Based DBMSs

- Usually more limited than locking facilities on enterprise DBMSs
- Programs can lock an entire table or an individual row within a table, but only one or the other
- Programs can release any or all of the locks they currently hold
- Programs can inquire whether a given row or table is locked

Timestamping

- DBMS assigns each database update a unique time (**timestamp**) when the update started
- Advantages
 - Avoids need to lock rows
 - Eliminates processing time needed to apply and release locks and to detect and resolve deadlocks
- Disadvantages
 - Additional disk and memory space
 - Extra processing time

Recover Data

- **Recovery:** returning database to a correct state from an incorrect state
- Simplest recovery involves using backups
 - **Backup** or **save:** copy of database

Journaling

- **Journaling**: maintaining a **journal** or **log** of all updates
 - Log is available even if database is destroyed
- Information kept in log for each transaction:
 - Transaction ID
 - Date and time of each update
 - **Before image**
 - **After image**
 - Start of a transaction
 - Successful completion (**commit**) of a transaction

Journaling (continued)

Transaction ID	Transaction Description
1	1. Change the Price value for part number DW11 to \$389.99
2	1. Add a record to the Orders table: OrderNum of 21700, OrderDate of 10/24/2010, CustomerNum of 282 2. Add a record to the OrderLine table: OrderNum of 21700, PartNum of DW11, NumOrdered of 3, QuotedPrice of \$389.00 3. Add a record to the OrderLine table: OrderNum of 21700, PartNum of KL62, NumOrdered of 2, QuotedPrice of \$346.50 4. Change the OnHand value for part number DW11 to 9 5. Change the OnHand value for part number KL62 to 10 6. Change the Balance value for CustomerNum 282 to \$2,321.50 7. Change the Commission value for RepNum 35 to \$39,346.20
3	1. Add customer 510
4	1. Delete part AT94

FIGURE 7-10: Four sample transactions

Forward Recovery

- DBA executes a DBMS recovery program
- Recovery program applies after images of committed transactions from log to database
- Improving performance of the recovery program
 - Apply the last after image of a record

Forward Recovery (continued)

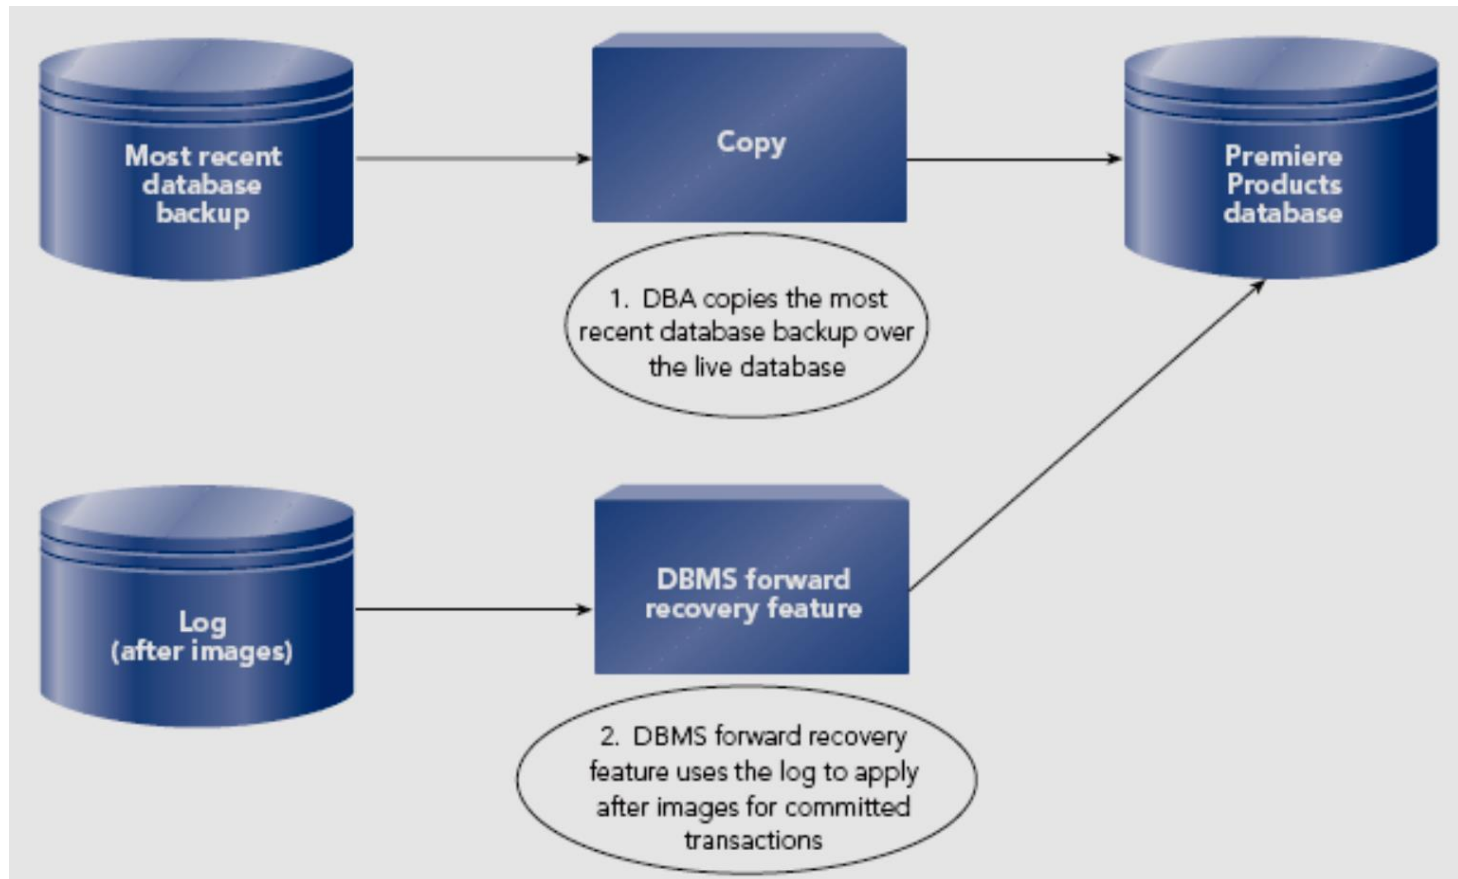


FIGURE 7-12: Forward recovery

Backward Recovery

- Database not in a valid state
 - Transactions stopped in midstream
 - Incorrect transactions
- **Backward recovery or rollback**
 - Undo problem transactions
 - Apply before images from log to undo their updates

Backward Recovery (continued)

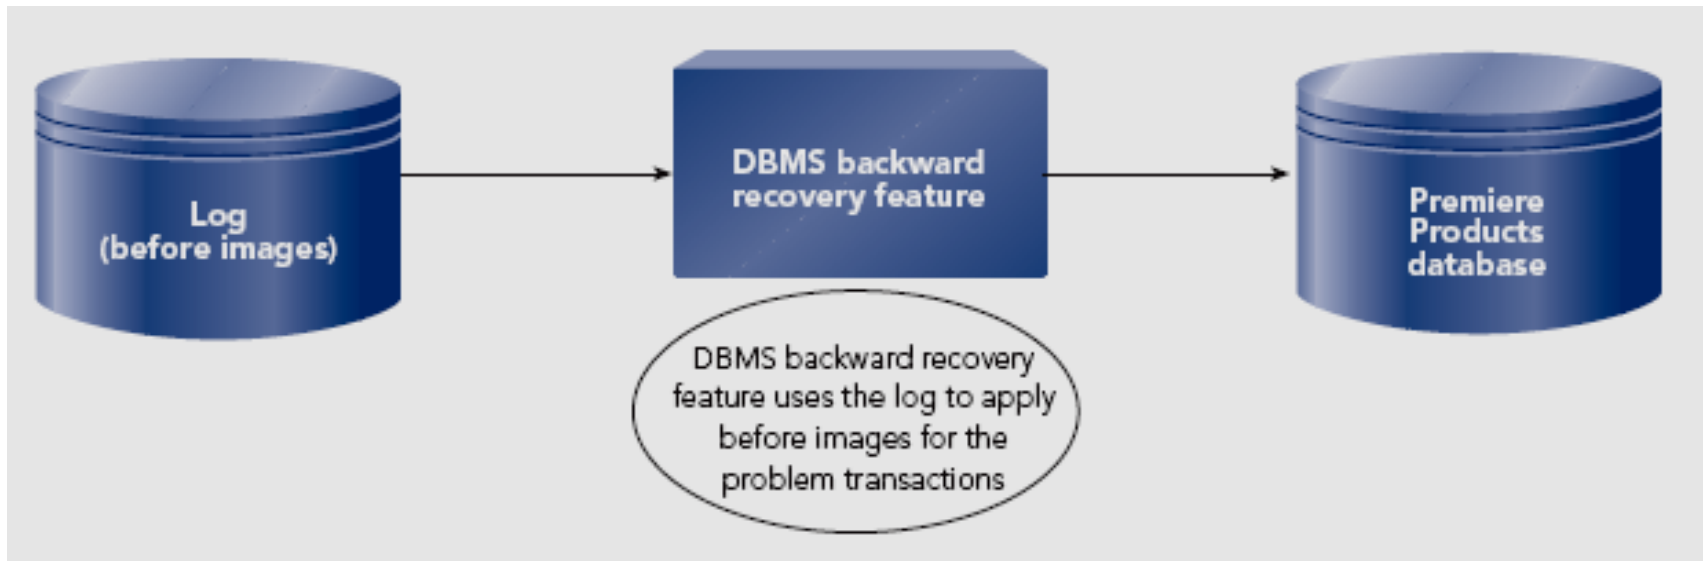


FIGURE 7-13: Backward recovery

Recovery on PC-Based DBMSs

- Sophisticated recovery features not available on PC-based DBMSs
- Regularly make backup copies using DBMS
 - Use most recent backup for recovery
- Systems with large number of updates between backups
 - Recovery features not supplied by DBMS need to be included in application programs

Provide Security Services

- Security: prevention of unauthorized access, either intentional or accidental, to a database
- Most common security features used by DBMSs:
 - Encryption
 - Authentication
 - Authorizations
 - Views

Encryption

- **Encryption:** converts data to a format indecipherable to another program and stores it in an encrypted format
- Encryption process is transparent to a legitimate user
- **Decrypting:** reversing the encryption
- In Access, encrypt a database with a password

Authentication

- **Authentication:** techniques for identifying the person attempting to access the DBMS
- **Password:** string of characters assigned by DBA to a user that must be entered for access
- **Biometrics:** identify users by physical characteristics such as fingerprints, voiceprints, handwritten signatures, and facial characteristics
- **Smart cards:** small plastic cards with built-in circuits containing processing logic to identify the cardholder

Authentication (continued)

- **Database password:** string of characters assigned to database that users must enter for accessing the database



FIGURE 7-14: Assigning a database password to the Premiere Products database

Authorizations

- DBA can use **authorization rules** to specify which users have what type of access to which data
- **Permissions**: specify what kind of access the user has to objects in the database
- **Workgroups**: groups of users

Views

- View: snapshot of certain data in the database at a given moment in time
- Can be used for security purposes

Privacy

- **Privacy:** right of individuals to have certain information about them kept confidential
- Laws and regulations dictate some privacy rules
- Companies institute additional privacy rules

Provide Data Integrity Features

- Rules followed to ensure data is accurately and consistently updated
- Key integrity
 - Foreign key and primary key constraints
- Data integrity
 - Data type
 - Legal values
 - Format

Provide Data Integrity Features (continued)

- Four ways of handling integrity constraints:
 1. Constraint is ignored
 2. Responsibility for constraint enforcement placed on users
 3. Responsibility for constraint enforcement placed on programmers
 4. Responsibility for constraint enforcement placed on DBMS

Provide Data Integrity Features (continued)

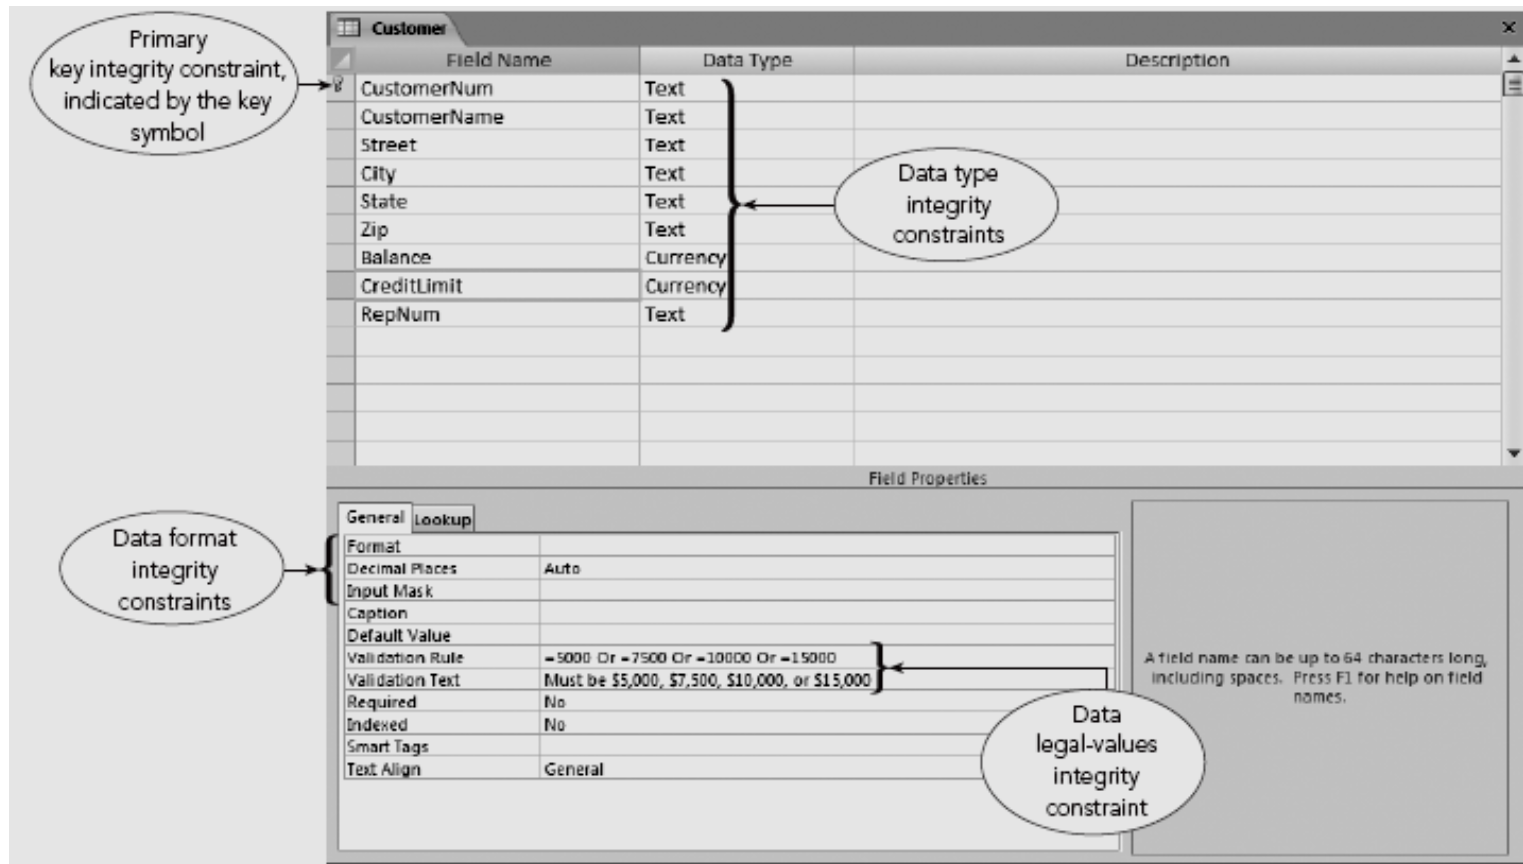


FIGURE 7-16: Example of integrity constraints in Access

Support Data Independence

- **Data independence:** can change database structure without needing to change programs that access the database
- Types of changes:
 - Adding a field
 - Changing a field property (such as length)
 - Creating an index
 - Adding or changing a relationship

Adding a Field

- Don't need to change any program except those programs using the new field
- SQL `SELECT * FROM` command will present an extra field
 - Solution: list the required fields in an SQL `SELECT` command instead of using `*`

Changing the Length of a Field

- Generally, don't need to change programs
- Need to change the program if:
 - Certain portion of screen or report is set aside for the field and the space cannot fit the new length

Creating an Index

- To create an index, enter a simple SQL command or select a few options
- Most DBMSs use the new index automatically
- For some DBMSs, need to make minor changes in already existing programs

Adding or Changing a Relationship

- Trickiest of all
- May need to restructure database

Support Data Replication

- **Replicated:** duplicated
- Manage multiple copies of same data in multiple locations
- Maintained for performance or other reasons
- Ease of access and portability
- **Replicas:** copies
- **Synchronization:** DBMS exchanges all updated data between master database and a replica

Support Data Replication (continued)

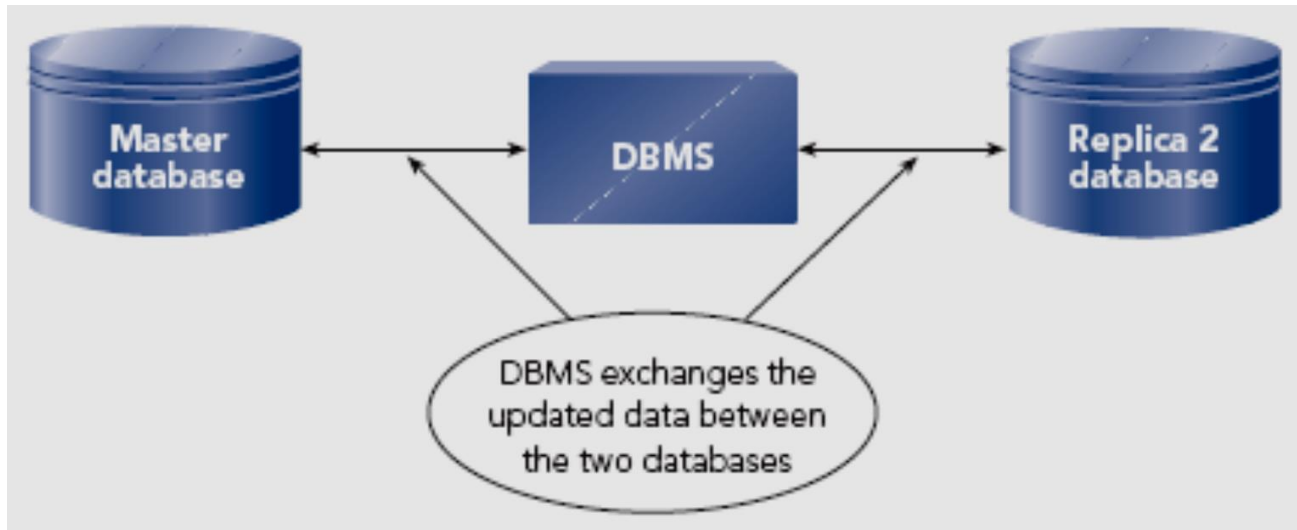


FIGURE 7-18: DBMS synchronizes two databases in a replica set

Provide Utility Services

- **Utility services** assist in general database maintenance
- Change database structure
- Add new indexes and delete indexes
- Use services available from operating system
- Export and import data
- Support for easy-to-use edit and query capabilities, screen generators, report generators, etc.

Provide Utility Services (continued)

- Support for procedural and nonprocedural languages
 - **Procedural language:** must tell computer precisely how a given task is to be accomplished
 - **Nonprocedural language:** describe task you want computer to accomplish
- Easy-to-use menu-driven or switchboard-driven interface

Summary

- DBMS allows users to update and retrieve data in a database without needing to know how data is structured on disk or manipulated
- DBMS must store metadata (data about the data) and make this data accessible to users
- DBMS must support concurrent update
- Locking denies access by other users to data while DBMS processes one user's updates
- During *deadlock* and *deadly embrace*, two or more users are waiting for the other user to release a lock before they can proceed

Summary (continued)

- In timestamping, DBMS processes updates to a database in timestamp order
- DBMS must provide methods to recover a database in the event the database is damaged
- DBMSs provide facilities for periodically making a backup copy of the database
- Enterprise DBMSs maintain a log or journal of all database updates since the last backup; log is used in recovery process

Summary (continued)

- DBMSs provide security features (encryption, authentication, authorizations, and views) to prevent unauthorized access to a database
- DBMS must follow rules or integrity constraints (key integrity constraints and data integrity constraints) so that it updates data accurately and consistently
- DBMS must support data independence
- DBMS must have facility to handle data replication
- DBMS must provide utility services that assist in general maintenance of a database