Chapter 12 – Database, Security, and Controls

Solutions to End-of-Chapter Problems

Review Questions

1. List the components of a DBMS and describe the function of each.

Application program interface – An interface engine or library of precompiled subroutines that enable application programs (such as those written in C or Java) to interact with the database.

End-user query processor – A program or utility that allows end users to retrieve data and generate reports without writing application programs.

Data definition interface – A program or utility that allows a database administrator to define or modify the content and structure of the database (for example, add new fields or redefine data types or relationships).

Data access and control logic – The system software that controls access to the physical database and maintains various internal data structures (for example, indexes and pointers).

Database – The physical data store (or stores) combined with the schema.

Schema – A store of data that describes various aspects of the “real” data, including data types, relationships, indexes, content restrictions, and access controls.

Physical data store – The “real” data as stored on a physical storage medium (for example, a magnetic disk).

2. What is a database schema? What information does it contain?

A database schema is a store of data that describes the content and structure of the physical data store (sometimes called metadata—data about data). It contains a variety of information about data types, relationships, indexes, content restrictions, and access controls.

3. Why are databases the preferred method of storing data used by an information system?

Databases are a common point of access, management, and control. They allow data to be managed as an enterprise-wide resource while providing simultaneous access to many different users and application programs. They solve many of the problems associated with separately maintained data stores, including redundancy, inconsistent security, and inconsistent data access methods.

4. With respect to relational databases, briefly define the terms row and attribute value.

Row – The portion of a table containing data that describes one entity, relationship, or object.

Attribute value – The portion of a table (a column) containing data that describes the same fact about all entities, relationships, or objects in the table.

5. What is a primary key? Are duplicate primary key values allowed? Why or why not?

A primary key is a field or set of fields, the values of which uniquely identify a row of a table. Because primary keys must uniquely identify a row, duplicate key values aren't allowed.

6. What is the difference between a natural key and an invented key? Which type is most commonly used in business information processing?

A natural key is a naturally occurring attribute of or fact about something represented in a database (for example, a human fingerprint or the atomic weight of an element). An invented key is one that is assigned by a system (for example, a social security or credit card number). Most keys used in business information processing are invented.

7. What is a foreign key? Why are foreign keys used or required in a relational database? Are duplicate foreign key values allowed? Why or why not?

A foreign key is a field value (or set of values) stored in one table that also exists as a primary key value in another table. Foreign keys are used to represent relationships among entities that are represented as tables. Duplicate foreign keys are not allowed within the same table because they would redundantly represent the same relationship. Duplicate foreign keys may exist in different tables because they would represent different relationships.

8. Describe the steps used to transform a domain class diagram into a relational database schema.

1. Create a table for each class.
2. Choose a primary key for each table.
3. Add foreign keys to represent one-to-many relationships.
4. Create new tables to represent many-to-many relationships.
5. Define referential integrity constraints.
6. Evaluate schema quality and make necessary improvements.
7. Choose appropriate data types and value restrictions for each field.

9. What is referential integrity? Describe how it is enforced when a new foreign key value is created, when a row containing a primary key is deleted, and when a primary key value is changed.

Referential integrity is content constraint between the values of a foreign key and the values of the corresponding primary key in another table. The constraint is that values of the foreign key field(s) must either exist as values of a primary key or must be NULL. A valid value must exist in the foreign key field(s) before the row can be added. When a row containing the primary key is deleted, the row with the foreign key must also be deleted for the data to maintain referential integrity. A primary key should never be changed; but in the event that it is, the value of the foreign key must also be changed.

10. What types of data (or attributes) should never be stored more than once in a relational database? What types of data (or attributes) usually must be stored more than once in a relational database?

Non-key fields should never be stored more than once.

If a table represents a class, the primary key values of each class represented in the table are redundantly stored (as foreign keys) for every relationship in which the class participates.

11. What is relational database normalization? Why is a database schema in third normal form considered to be of higher quality than an unnormalized database schema?

Relational database normalization is a process that increases schema quality by minimizing data redundancy. A schema with tables in third normal form has less non-key data redundancy than a schema with unnormalized tables. Less redundancy makes the schema and database contents easier to maintain over the long term.

12. Describe the process of relational database normalization. Which normal forms rely on the definition of functional dependency?

The process of normalization modifies the schema and table definitions by successively applying higher order rules of table construction. The rules each define a normal form, and the normal forms are numbered one through three. First normal form eliminates repeating groups that are embedded in tables.

Second and third normal forms are based on a concept called *functional dependency*—a one-to-one correspondence between two field values. Second normal form ensures that every field in a table is functionally dependent on the primary key. Third normal form ensures that no non-key field is functionally dependent on any other non-key field.

13. What is the difference between a primitive data type and a complex data type?

A primitive data type (for example, integer, real, or character) is directly supported (represented) by the CPU or a programming language. A complex data type (for example, record, linked list, or object) contains one or more data elements constructed using the primitive data types as building blocks.

14. Briefly describe these distributed database architectures: replicated database servers, partitioned database servers, and cloud-based database servers. What are the comparative advantages of each?

Replicated database servers – An entire database is replicated on multiple servers, and each server is located near a group of clients. Best performance and fault tolerance for clients because all data is available from a “nearby” server.

Partitioned database servers – A database is partitioned so that each partition is a database subset used by a single group of clients. Each partition is located on a separate server, and each server is located close to the clients that access it. Better performance and less replication traffic than replicated servers if similar collocated clients use only a subset of database content.

Cloud-based database servers – A cloud based database server is really one of the previously defined architectures, but implemented using the cloud-based services of a cloud computing vendor. The cloud provider hosts the database and provides the services across defined geographical areas. The cloud provider manages the database including synchronization and backup.

15. What additional database management complexities are introduced when database contents are replicated in multiple locations?

Replicated copies are redundant data stores. Thus, any changes to data content must be redundantly implemented on each copy. Implementing redundant maintenance of data content requires all servers to periodically exchange database updates.

16. Describe the risk factors associated with database design.

Since the database is an integral part of any information system, the performance and operation of the database is critical. Hence design decisions about what DBMS to use, how to configure it, and how to optimize it are usually quite complex and critically important.

Another risk factor is how to integrate a new database with existing databases. New systems normally are not completely independent of existing systems and database, and usually must interface with existing architectures. It is important that the new database not only integrate well, but that it not cause problems with existing configurations.

Finally, good database design depends on having a complete, or mostly complete problem domain model. While enhancing and adding tables and attributes is possible to already constructed databases, doing so can sometimes cause sub-optimization of data structures.

17. When should database design be performed? Can the database be designed iteratively or must the entire database be designed at once?

Database design is usually done as early as possible in the project. For an iterative project, it is usually designed and implemented in the earlier iterations. The database does not have to be designed completely all at once, however, it should be designed and refined as much as possible in the early iterations.

18. Explain four types of integrity controls for input forms. Which have you seen most frequently? Why are they important?

- Field combination controls verify that the data in one field is based on the data in another field or fields.
- Value limit controls identify when a value in a field is too large or too small.
- Completeness controls ensure that all necessary fields on an input form have been entered.
- Data validation controls validate the input data for correctness.

Answers will vary on importance.

19. What are the objectives of integrity controls in information systems? In your own words, explain what each of the three objectives means. Give an example of each.

- Ensure that only appropriate and correct business transactions occur. This objective ensures that no erroneous or fraudulent transactions are entered. Example: A control to ensure that a clerk does not request a check for a service that was never provided.

- Ensure that the transactions are recorded and processed correctly. This objective ensures that the system processes and stores the data completely. Example: a control to ensure that a double-entry bookkeeping entry always processes both entries.

- Protect and safeguard the assets (including information) of the organization. This objective ensures that information is not lost due to theft, fire, or some other mishap. Example: Storing backup data periodically off site.

20. What are the four types of input controls used to reduce input errors? Describe how each works.

Field combination control: An integrity control that verifies that the data in one field is based on the data in another field or fields.

Value limit control: An integrity control that identifies when a value in a field is too large or too small.

Completeness control: An integrity control that ensures that all necessary fields on an input form have been entered.

Data validation control: An integrity control that validates the input data for correctness and appropriateness.

21. What is the basic purpose of transaction logging?

Transaction logging takes every update to the database and logs exactly how it happened (sometimes with an image of the transaction). It is extremely important for audit trails and for recovery in case something goes wrong.

22. What are the two primary objectives of security controls?

Maintain a stable, functioning operating environment for users and application systems (usually 24 hours a day, seven days a week).

Protect information and transactions during transmission outside the organization (public carriers).

23. Briefly define or describe authentication, access control lists, and authorization.

Authentication is the process used to identify who a user is – to authenticate that this is the right person. This process is often done with user ID and password. Other methods of identifying a person might be with smart card, biometric devices, and questions and answers.

Authorization is done after authentication. Once a user has been verified, i.e. the system knows who it is, then that user will have rights or privileges to access particular data and and system functions. He/she is authorized to access particular parts of the system. What parts of the system a user has access to is often determined by an Access Control list. An access control lists contains a list of all the users and what rights or privileges he/she has.

24. How does single-key (symmetric) encryption work? What are its strengths? What are its weaknesses?

A single key is used to encrypt and decrypt a message. Both parties must have the key. Its strength is that it is simple and fast. Its weaknesses are that it might be easy to break the encryption and that it is difficult to distribute the key in a secret fashion to all the authorized participants.

25. How does public key (asymmetric) encryption work? What are its strengths? What are its weaknesses?

A public-key encryption has two keys, a public one that is widely distributed and a private one that is secret. To send data to the owner of the keys, someone uses the public key. The data can then only be decrypted with the private key. So, the owner is the only one who can decrypt the data. After the message is encrypted, it can only be decrypted with the private key.

Its strengths are that it is very secure, and the public key can be widely distributed. So if an entity wants to received secret messages, it can distribute its public key, and it will be the only one able to decrypt the message.

One weakness is that it tends to be quite slow in decrypting. So it is not suitable for high volume, rapidly changing data.

26. What is a digital certificate? What role do certifying authorities play in security systems?

A digital certificate is an institution's name and public key (plus other information such as address, Web site URL, and validity date of the certificate) that is encrypted and certified by a third party.

Certifying authorities are companies that are very well known so that everybody knows for sure what their public keys are. These certifying authorities sell digital certificates to other companies (that are not as well known) so that these companies can convince their customers that they are legitimate.

27. What is a digital signature? What does it tell a user?

A digital signature is a technique in which a document is encrypted using a private key to verify who wrote the document. If you have the public key of an entity, and that entity sends you a message with its private key, you can decode it with the public key. You know that the party is the one you want to communicate with because that entity is the only one who can encode a message with that private key.

Problems and Exercises

1. The Universal Product Code (UPC) is a bar-coded number that uniquely identifies many products sold in the United States. For example, all printed copies of this textbook sold in the United States have the same UPC bar code on the back cover. Now consider how the design of the RMO database might change if all items sold by RMO were required by law to carry a permanently attached UPC (e.g., on a label sewn into a garment or on a radio frequency ID tag attached to a product). How might the RMO relational database schema change under this requirement?

The change to the schema is relatively simple. ProductID is replaced with the UPC bar code both as the primary key of ProductItem and all corresponding foreign keys. The change might be more complex if the DBMS were previously responsible for generating values of ProductItem.Number. That function would now be removed from the DBMS because the key values would be externally assigned. This would potentially add more complexity to the system in order to determine what the UPC values were and to get them entered into the system.

2. Assume that RMO will begin asking a random sample of customers who order by telephone about purchases made from competitors. RMO will give customers a 15 percent discount on their current order in exchange for answering a few questions. To store and use this information, RMO will add two new classes and three new associations to the class diagram. The new classes are Competitor and ProductCategory. Competitor has a one-to-many association with ProductCategory, and the existing Customer class also has a one-to-many association with ProductCategory. Competitor has a single attribute called Name. ProductCategory has four attributes: Description, DollarAmountPurchased, MonthPurchased, and YearPurchased. Revise the relational database schema shown in Figure 12-10 to include the new classes and associations. All tables must be in 3NF.

The following tables must be added to the relational database schema:

Competitor = **Name**
ProductCategory = ***CompetitorName, CustomerAccountNo,* MonthPurchased,**
YearPurchased, **Description,** DollarAmountPurchased

Primary keys are shown in bold, and foreign keys are shown in italics. Note that the primary key of ProductCategory is guaranteed to be unique only if multiple customer purchases from a competitor in the same month and for the same product category (description) are combined in a single row.

3. Assume that RMO will use a relational database, as shown in Figure 12-10. Assume further that a new catalog group located in Milan, Italy, will now create and maintain the product catalog. To minimize networking costs, the catalog group will have a dedicated database server attached to its LAN. Develop a plan to partition the RMO database. Which tables should be replicated on the catalog group's local database server? Update Figure 12-18 to show the new distributed database architecture.

Assumptions:
Milan will have responsibility for describing and maintaining product items and accessory packages.
Milan will also have responsibility for supporting and maintaining Promotions.

The following tables will need to be replicated on the local LAN. Access requirements (C,R,U, and D) are shown in parentheses.

Promotion (CRUD)
PromoOffering (CRUD)
ProductItem (CRUD)
AccessoryPackage (CRUD)
InventoryItem (R)

Updates to all of these tables are assumed to be relatively infrequent and, thus, the performance cost of complete replication with immediate or frequent updates is minimal. Milan can be represented in Figure 10-32 in the same manner as the warehouse LAN or retail store LAN.

4. Visit the Web site of an online catalog vendor similar to RMO (such as www.llbean.com) or an online vendor of computers and related merchandise (such as www.cdw.com). Browse the online catalog and note the various types of information contained there. Construct a list of complex data types that would be needed to store all the online catalog information.

Answers will vary.
Some examples of typical complex data types include:

- Graphic images in formats such as GIF and JPEG.
- Motion video in formats such as MPEG and AVI.
- Sound in formats such as WAV and MP3.
- Executable programs in formats such as Java and VBScript.
- Browser-formatted documents in HTML and XML.
- Hard-copy documents in formats such as Postscript or Acrobat.

5. This chapter described various situations that emphasized the need for controls. In the first scenario presented, a furniture store sells merchandise on credit. Based on the descriptions of controls given in this chapter, identify the various controls that should be implemented in the system to ensure that corrections to customer balances are made only by someone with the correct authorization.

Answers will vary but should include at least the following:

- Transaction logging to note all changes (especially financial) made to the database. Log records should include the login ID of the person making the transaction.
- Financial transaction screens should be available (and visible) only via authorization of the correct level of registered user.
- Possibly a notification report of any changes (other than standard payments) made to correct account balances.

In the second scenario illustrating the need for controls, an accounts payable clerk uses the system to write checks to suppliers. Based on the information in this chapter, what kinds of controls would you implement to ensure that checks are written only to valid suppliers, that checks are written for the correct amount, and that all payouts have the required authorization? How would you design the controls if different payment amounts required different levels of authorization?

Answers will vary but should include at least the following:

- Both manual and automated controls might be needed for this process. The manual control will require authorization by a supervisor on paper documents for payment. Also, a paper audit trail (numbered invoice) might be required.
- Payments made only to valid suppliers can be controlled by having pre-defined PayTo fields that come from a supplier file. The supplier file should be maintained by different people to ensure separation of duties.
- Ensuring that checks are written for the correct amount can be accomplished by making sure a payment amount corresponds with the invoice amount in the system.
- A supervisor can also verify payments for correct amounts and viable suppliers. This can be done either with paper documents or with electronic forms. Before a check is written, a payment transaction can be approved by an electronic signature of a different person.
- Output reports detailing payments should be provided and reviewed.
- Internal edits can be developed to note whether payments are customary and normal. Out-of-range payments can be flagged as exceptions and verified by a manager.
- Different levels of payments will require the same types of controls; however, they may require different electronic signatures by higher-level registered users.

6. Look on the Web for an e-commerce site (such as Amazon.com or eBay). What kinds of security and controls are integrated into the system?

Answers will vary. The following explanation uses Amazon.com as an example.

Data Site

Fairly easy to navigate. Lots of tabs.
Has plenty of search capability. Can narrow focus by entering different portals.

Pages

Some pages seem fairly busy; however, most pages are laid out fairly well.
Hotlinks appear to be easy to find.

Data Entry

Mostly done with clicks, which reduces data errors.
Gives lots of opportunities to double check and correct choices.
Implements one-click method to expedite data entry and reduce errors.

Security and Controls

Use user ID and password. Do not give option of system remembering password.
Make user sign in to secure server.
Use secure socket layer.

Potential Security Problems

Will let you work with standard, non-secure servers.
If you forget your password, system allows you (or someone else) to set up a new one.
Remembers all credit card information for various credit cards.

7. Examine the information system of a local business, such as a fast-food restaurant, doctor's office, video store, grocery store, etc. Evaluate the screens (and reports, if possible). What kind of integrity controls are in place? What kinds of improvements would you make?

Answers will vary.

8. Search the Web and find out what you can about Pretty Good Privacy. What is it? How does it work? Find what you can about a pass phrase. What does it mean? Start your research at www.pgpi.org.

Answers will vary.

PGP is a program that makes your e-mail messages private. It does this by encrypting your e-mail so that nobody but the intended person can read it. When encrypted, the message looks like a meaningless jumble of random characters. PGP has proven itself quite capable of resisting even the most sophisticated forms of analysis aimed at reading the encrypted text.

PGP can also be used to apply a digital signature to a message without encrypting it. This is normally used in public postings where you don't want to hide what you are saying, but rather want to allow others to confirm that the message actually came from you. After a digital signature is created, it is impossible for anyone to modify either the message or the signature without the modification being detected by PGP.

PGP uses public key encryption to encrypt and decrypt e-mail messages.

A pass phrase is a longer version of a password. It can be an entire phrase. It is used to generate the public key/private key combination for a PGP user.

Solutions to End-of-Chapter Cases

Case Study: Computer Publishing, Inc.

1. Consider the contents of this textbook as a template for CPI's database content. Draw a class diagram that represents the book and its key content elements. Expand your diagram to include related product content, such as a set of PowerPoint slides, an electronic book formatted as a Web site or PDF file, and a Web based test bank.

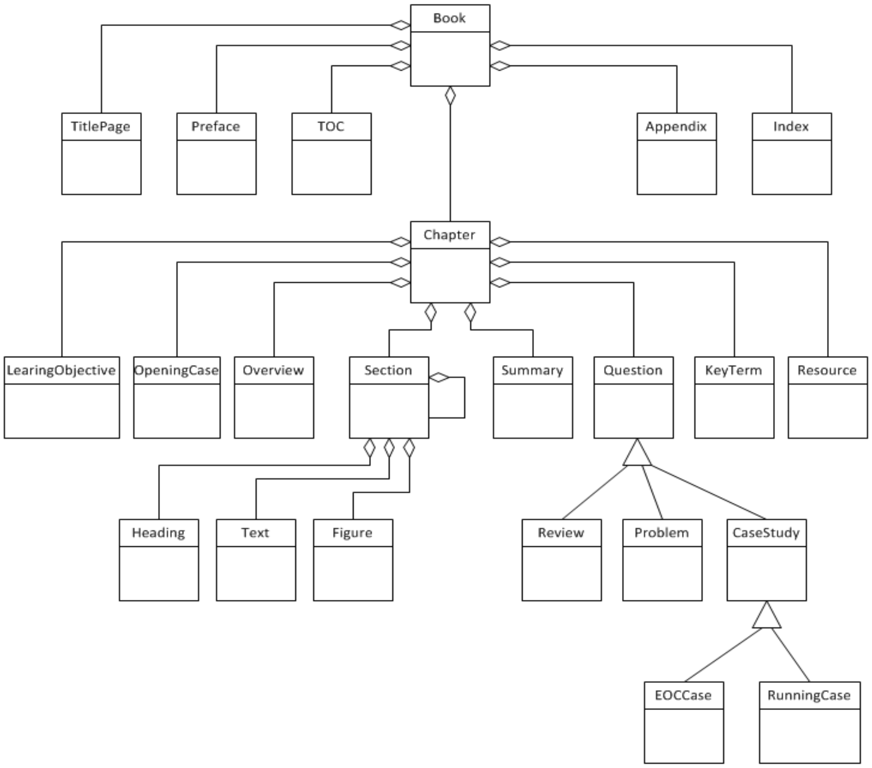

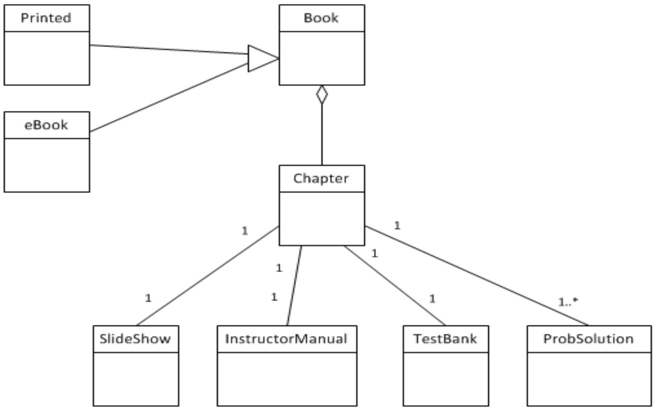

2. Develop a list of data types required to store the content of the book, slides, and Web sites. Are the relational DBMS data types listed in Figure 12-15 sufficient?

Relational databases have the option of storing the actual data or of only storing a pointer to the data. Where the data is complex such as image files, video files, slide shows, and sound files, most often those are stored outside of the database with the address or locater stored in the database. If the data itself is stored in the database, then a data type of "blob" is often used.

3. Authors and editors are often independent contractors, not publishing company employees. Consider the implications of this fact for controls and security. How would you enable authors and editors to interact with the database? How would you protect database content from hackers and other unauthorized accesses?

Interacting with an SQL relational database is not easy to do at the low level – at the SQL level. Therefore an entire system would need to be designed to allow authors to write text and create the figures and upload them into the appropriate places in the textbook structure. Writing and editing a textbook would probably require a different approach than an author simply sitting down and writing from front to back.

The new system would need to have login capability to allow only authenticated authors to access the textbook materials. Authorization would also be required to control which parts of the textbook authors and editors could access and update. One potential problem is how to prevent the database for a particular textbook from becoming corrupted, or out of order, or jumbled if the authors entered the information incorrectly.

The problem with hackers is the same as with any proprietary data that is available over the Internet. Access to the database must be through secure measures and encrypted passwords, or perhaps even more secure login procedures.

Running Cases: Community Board of Realtors

In Chapter 4, you developed a domain model class diagram. Using your previous solution or one provided to you by your instructor, update your domain model class diagram with any additional problem domain classes, new associations, or additional attributes that you have discovered as you worked through the previous chapters. Finalize this comprehensive domain model and then turn it in as part of your solution.

Using this comprehensive domain model class diagram, develop a relational database schema. In the schema, identify the foreign keys that are required. Also, identify a key attribute for each table. You may need to add a key field if there isn't an attribute that could logically serve as the key. Remember that a candidate key for an association class is the combination of the keys of the connected classes. However, it may make sense to define a shorter, more concise key field.

Verify that each table is in first, second, and third normal form. Discuss any discrepancies you had to fix from your first solution. Discuss any tables that may not be in third normal form and why you are leaving it as not-normalized.

Note: We will use the following class diagram from Chapter 4 problem 3 for this problem. The following changes/additions were made:
1. Commissioner was dropped. It is not a logical piece of the MLS system.
2. A few new attributes were added.

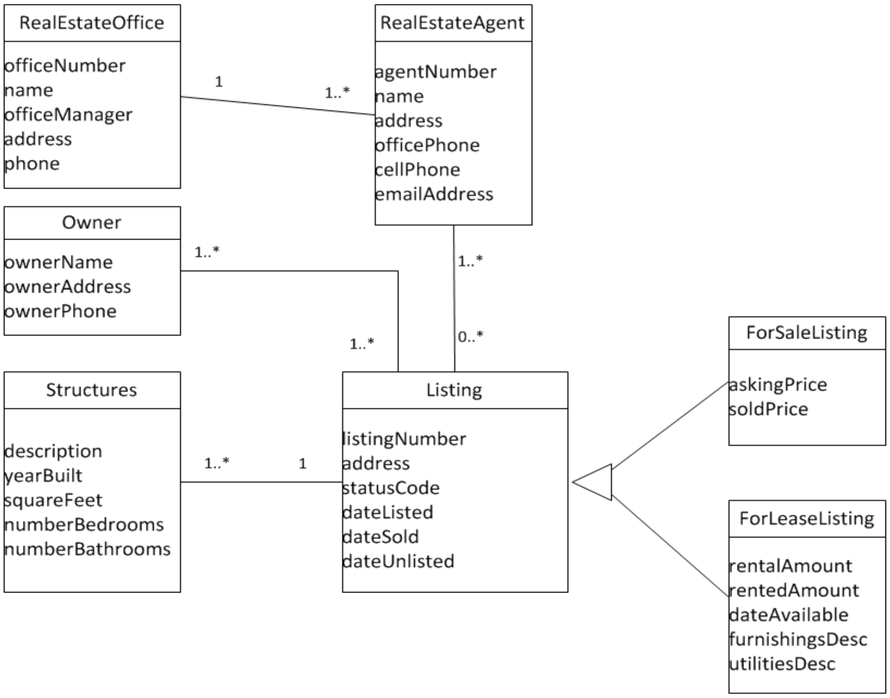

Table	Fields (columns)
REOffice	**office_number**, office_name, manager, street, city, state_province, postal_code, telephone
REAgent	**agent_number,** *office_number*, agent_lastname, agent_firstname, street, city, state_province, postal_code, office_phone, mobile_phone, email_address
Listing	**listing_number**, listing_type, property_street, property_city, property_state, property_postal, status_code, date_listed, date_sold, date_unlisted, asking_price, sold_price
ForLeaseListing	**listing_number**, rental_amount, rented_amount, date_available, furnishing_desc, utilities_desc
Owner	**owner_number**, owner_lastname, owner_firstname, street, city, state_province, postal_code
Structures	**structure_number**, *listing_number,* description, year_built, square_feet, number_bedrooms, number_bathrooms
AgentOnListing	***agent_number, listing_number***
OwnerOnListing	***owner_number, listing_number***

Note:

Primary key is bold. Foreign key is italicized

The ForSaleListing was combined with the listing table, since most listings are For Sale Listings. ForRentListing is a separate table because the information is unique and there are only a few of those types of records. A new field, listing_type, was added to denote rental listings.

The type codes (string, integer, number, etc.) nor the length have been included this information will have to be added before the tables can be entered into a database.

The tables are in 3NF with one exception. State is functionally dependent on Postal_code, i. e. a state can be determined by postal code. But due to common usage of always having state and postal-code included, they are maintained together. (Alternative is to have a separate postal_code to state translation table.)

Running Cases: The Spring Breaks 'R' Us Travel Service

As with other social networking sites and systems, users of the Spring Breaks ‘R’ Us social networking system face such risks as identity theft, phishing attacks, and viruses. Review the following information related to social networking risks and security published by the United States Computer Emergency Readiness Team, including:

■ Socializing Securely: Using Social Networking Services (www.us-cert.gov/reading_room/safe_social_networking.pdf)
■ Cyber Security Tip ST06-003: Staying Safe on Social Network Sites (www.us-cert.gov/cas/tips/ST06-003.html)
■ Cyber Security Tip ST05-013: Guidelines for Publishing Information Online (www.us-cert.gov/cas/tips/ST05-013.html)

After reviewing this information, revisit the questions for this case in Chapter 6 for the Social Networking subsystem. Based on the contents of this chapter and the information contained in the readings, what specific controls and security measures should be incorporated into the Social Networking subsystem?

Answers will vary.

There is always an issue about how easy to make it for friends to “find” and “interact” with each other, when contrasted with how secure (i.e. how difficult) the system to be. Students should address such things as:

- Login Protection:
  - Strong passwords
  - Remember MAC address
  - Personal questions
- Connection Protection:
  - What is allowed to be searched to find people
  - What is publicly viewable
  - Approval process for connected friends
- Posting Protection:
  - What can be posted on one's own account
  - What can be posted on friend's account
- Viewing Protection:
  - Different categories of friends
  - Privacy settings for posted materials
- Log off protection
  - Automatic log off based on time or activity

Running Cases: On the Spot Courier Services

In Chapter 6, you discussed hardware requirements, and in Chapter 10, you developed component and deployment diagrams. Based on your work in those chapters, take these steps:

1. For each user and each type of device, discuss what security precautions and techniques should be used to protect access to the device itself.

2. For each user and each type of device, discuss what security precautions and techniques should be used to protect access to the application programs, connect to the home system, and protect the data being transmitted to the foreign devices.

3. Discuss any security precautions and techniques you would recommend for the home office and the network servers.

Answers will vary

User	Device	Protect Device	Protect data, connection, and software
Customer	Customer computer	Customer responsibility	Customer login required Strong passwords required Encrypt sent data (HTTPS) Validate all data entered to guard against XSS attacks, etc.
Delivery Person	Mobile device	Install tracking software (GPS) Have secure place in truck Serial number on device	User login required Encrypt all data transmitted Limited inputs and outputs Validate all inputs Routes start/stop times to validate inputs
Warehouse person	Scanning device	Keep in warehouse Have secure storage location Serial number on device Check out/check in device	Limited data transmittal Keep software up to date Validate device and expected inputs
Warehouse person	Warehouse computer	Keep in secure office Automatic turn off	Login required Automatic turn off
Bill	Home laptop	Bill keep in secure location	Keep software up to date Login required MAC address remembered VPN Encrypted connections
Bill	Warehouse servers	Keep in secure office	Server security software Validate all inputs Encrypt all data transmissions Keep software up to date

Running Cases: Sandia Medical Devices

Part 1.

Review the original system description in Chapter 2, the additional project information in Chapters 3, 4, and 8, and the domain class diagram shown in Figure 12-26 to refamiliarize yourself with the proposed system. Assume that the type attribute of the AlertCondition class identifies one of three alert types:

1. Glucose levels that fall outside the specified range for 15 minutes (three consecutive readings)

2. Glucose levels that fall outside the specified range for 60 minutes (12 consecutive readings)

3. An average of glucose levels over a eight-hour period that falls outside a specified range

The specified range for an AlertCondition object is the set of values between and including lowerBound and upperBound. AlertCondition objects also include an effective time period specified by the attributes startHour and endHour, which enables physicians to set different alert parameters for sleeping and waking hours.

When an alert is triggered, an object of type Alert is created and associated with an alertCondition object. The dateTime attribute records when the Alert object was created, and the value(s) attribute record(s) the glucose levels (alert types 1 and 2) or average level (alert type 3) that fell outside the specified range. Each Alert object is indirectly related to a Patient object via the association between Alert and AlertCondition and the association between AlertCondition and Patient. Develop a set of relational database tables based on the domain class diagram. Identify all primary and foreign keys, and ensure that the tables are in 3NF.

Table	Fields (Columns)
Patient	**patientID**, *physicianID*, *deviceID*, medical_rec_number, last_name, first_name, birthdate, gender, race, height, weight
Physician	**physicianID**, last_name, first_name
MonitorDevice	**deviceID**, serial_number, manufacturer, manufacture_date, firmware_version
AlertCondition	**alert_condID**, *patientID*, type, start_hour, end_hour, upper_bound, lower_bound
Alert	**alert_number**, *alert_condID*, date_time, value
CellPhone	**phoneID**, *patientID*, phone_number, operating_system, os_version, application_version
GlucoseObservation	**observation_number**, *patientID*, date_time, level

Primary keys
Foreign Keys
All tables are in 3NF. (Correctly built class diagrams always result in tables in 3NF.)

Part 2.
Based on what you learned in this chapter about databases, controls, and system security, review your answers to the questions for this case in Chapter 6. Assume that the patient's cell phone and the centralized servers are different nodes in a replicated database architecture and are regularly synchronized. What changes, if any, should be made to your answers now that you have a deeper understanding of databases, controls, security, and related design issues?

Answers will vary by student.