524 Chapter 18 Security

0√.

Table 18.1 Examples of threats.

Linear								2001
	्रमा जिम्ब	nt and Idaan	11-0 11-0	is of fidentiality	Loss of L	Loss c Integri	LOSS Of	
Using another person's means of access								lan Geol
Unauthorized amendment or copying of data	1				~			
Program alteration	1					· · ·		
Inadequate policies and procedures that allow a mix of confidential and normal output					ł.			
Wire tapping	1	£	· ·		- Witsterner -	38		
Illegal entry by hacker	1						(°	
Blackmail	1						٠	
Creating 'trapdoor' into system	1			~				
Theft of data, programs, and equipment	~~~			1				
Failure of security mechanisms, giving greater access than normal								
Staff shortages or strikes			•	~	· •			
Inadequate staff training			,					
Viewing and disclosing unauthorized data	1		×				1	
Electronic interference and radiation			•	1		а.		
Data corruption owing to power loss or surge								
Fire (electrical fault, lightning strike, arson), flood, bomb							,	
Physical damage to equipment		. °.						
Breaking cables or disconnection of cables		10						
Introduction of viruses					1		1	

The extent that an organization suffers as a result of a threat's succeeding depends upon a number of factors, such as the existence of countermeasures and contingency plans. For example, if a hardware failure occurs corrupting secondary storage, all processing activity must cease until the problem is resolved. The recovery will depend upon a number of factors, which include when the last backups were taken and the time needed to restore the system.

An organization needs to identify the types of threat it may be subjected to and initiate appropriate plans and countermeasures, bearing in mind the costs of implementing them. Obviously, it may not be cost-effective to spend considerable time, effort, and money on potential threats that may result only in minor inconvenience. The organization's business may also influence the types of threat that should be considered, some of which may be rare. However, rare events should be taken into account, particularly if their impact would be significant. A summary of the potential threats to computer systems is represented in Figure 18.1.

H-18pighed x Inipublic

18.2 Countermeasures - Computer-Based Controls



illegal entry by hacker Blackmall Introduction of viruses

procedures Staff shortages or strikes

Figure 18.1 Summary of potential threats to computer systems.

Countermeasures – Computer-Based Controls

The types of countermeasure to threats on computer systems range from physical controls to administrative procedures. Despite the range of computer-based controls that are available, it is worth noting that, generally, the security of a DBMS is only as good as that of the operating system, owing to their close association. Representation of a typical multiuser computer environment is shown in Figure 18.2. In this section we focus on the following computer-based security controls for a multi-user environment (some of which may not be available in the PC environment):

18.2

525

- ----,
- authorization
- views
- backup and recovery
- integrity

5;

T

6

1.0.

- encryption
- RAID technology.



Authorization The granting of a right or phylioge that enables

18.2.1 Authorization

asubject

System a solar

18.2 Countermeasures – Computer-Based Controls

Table 18.3 Access control matrix.

User Ider	illier: propertyNe	i ivoe	anine) Solite	Cinternio	Asiemino	Sectore and	e - Quêry kowili	
Sales	0001	0001	0001	0000	0000	0000	15	
SG37	0101	0101	0111	0101	0111	0000	100	
SG5	1111 -	1111	1111	1111	1111	1111	none	

propertyNo, type, and ownerNo attributes and Select, Update, and Insert privileges (shown as 0001 + 0010 + 0100 = 0111) for the price and staffNo attributes, with a limit of 100 rows for any query result set. Finally, user SG5 (Susan Brand) has Select, Update, Insert, and Delete privileges (shown as 0001 + 0010 + 0100 + 1000 = 1111), in other words All privileges for all attributes, with no limit set on the number of rows for any query result set.

DBMSs use similar matrices to implement access control, although the precise details of implementation vary from one system to another. On some DBMSs, a user has to tell the system under which identifier he or she is operating, especially if the user is a member of more than one group. It is essential to become familiar with the available authorization and other control mechanisms provided by the DBMS, particularly where priorities may be applied to different authorization identifiers and where privileges can be passed on. This will enable the correct types of privileges to be granted to users based on their requirements and those of the application programs that many of them will use.

Views (Subschemas)

The view mechanism provides a powerful and flexible security mechanism by hiding parts of the database from certain users. The user is not aware of the existence of any attributes or rows that are missing from the view. A view can be defined over several relations with a user being granted the appropriate privilege to use it, but not to use the base relations. In this way, using a view is more restrictive than simply having certain privileges granted to a user on the base relation(s). We discussed views in detail in Sections 3.4 and 6.4.

periodically taking a copy of the

(and possibly programs) on to offline storage media

Aview is the dynamic result of one opmore pelational operations

the base relations to produce another relation A view is coestrict actually, exist in the catabase but is produce

particular user, at the time of request

Backup and Recovery

Backup The proce

18.2.3

18.2.2

529

552 . Chapter 19 Transaction Management

Figure 19.1 Example transactions.

delete(staffNo=x) for all PropertyForRent's read(staffNo=x, salary) Degin salary = salary * 1.1 read propertyNo =pno, staffNo write staffNo = x, new_salary if (staffNo = x) then begin staffNo = newStaffNo write(propertyNo = pno; staffNo) (a) (b)

Staff PropertyForRent

(staffNo, fName, IName, position, sex, DOB, salary, branchNo) (propertyNo, street, city, postcode, type, rooms, rent, ownerNo, staffNo, branchNo)

A simple transaction against this database is to update the salary of a particular member of staff given the staff number, x. At a high level, we could write this transaction as shown in Figure 19.1(a). In this chapter we denote a database read or write operation on a data item x as read(x) or write(x). Additional qualifiers may be added as necessary; for example, in Figure 19.1(a), we have used the notation read(staffNo = x, salary) to indicate that we want to read the data item salary for the tuple with primary key value x. In this example, we have a **transaction** consisting of two database operations (read and write) and a non-database operation (salary = salary*1.1).

A more complicated transaction is to delete the member of staff with a given staff number x, as shown in Figure 19.1(b). In this case, as well as having to delete the tuple in the Staff relation, we also need to find all the PropertyForRent tuples that this member of staff managed and reassign them to a different member of staff, *newStaffNo* say. If all these updates are not made, referential integrity will be lost and the database will be in an inconsistent state: a property will be managed by a member of staff who no longer exists in the database.

A transaction should always transform the database from one consistent state to another, although we accept that consistency may be violated while the transaction is in progress. For example, during the transaction in Figure 19.1(b), there may be some moment when one tuple of PropertyForRent contains the new *newStaffNo* value and another still contains the old one, x. However, at the end of the transaction, all necessary tuples should have the new *newStaffNo* value.

A transaction can have one of two outcomes. If it completes successfully, the transaction is said to have committed and the database reaches a new consistent state. On the other hand, if the transaction does not execute successfully, the transaction is aborted, If a transaction is aborted, the database must be restored to the consistent state it was in before the transaction started. Such a transaction is rolled back or undone. A committed transaction cannot be aborted. If we decide that the committed transaction was a mistake, we must perform another compensating transaction to reverse its effects (as we discuss in Section 19.4.2). However, an aborted transaction that is rolled back can be restarted later and, depending on the cause of the failure, may successfully execute and commit at that time. Figure 19.4 The lost update problem.



Figure 19.5 The uncommitted dependency problem.







The problems described in Examples 19.1–19.3 resulted from the mismanagement of concurrency, which left the database in an inconsistent state in the first two examples and presented the user with the wrong result in the third. Serial execution prevents such problems occurring. No matter which serial schedule is chosen, serial execution never leaves the database in an inconsistent state, so every serial execution is considered correct, although different results may be produced. The objective of **serializability** is to find non-serial schedules that allow transactions to execute concurrently without interfering with one another, and thereby produce a database state that could be produced by a serial execution.

If a set of transactions executes concurrently, we say that the (nonserial) schedule is correct if it *produces the same results as some serial execution*. Such a schedule is called **serializable**. To prevent inconsistency from transactions interfering with one another, it is essential to guarantee serializability of concurrent transactions. In serializability, the ordering of read and write operations is important:

- If two transactions only read a data item, they do not conflict and order is not important.
- If two transactions either read or write completely separate data items, they do not conflict and order is not important.
- If one transaction writes a data item and another either reads or writes the same data item, the order of execution is important.

Consider the schedule S_1 shown in Figure 19.7(a) containing operations from two concurrently executing transactions T_7 and T_8 . Since the write operation on bal_x in T_8 does not conflict with the subsequent read operation on bal_y in T_7 , we can change the order of these operations to produce the equivalent schedule S_2 shown in Figure 19.7(b). If we also now change the order of the following non-conflicting operations, we produce the equivalent serial schedule S_3 shown in Figure 19.7(c):

- Change the order of the write(bal_x) of T₈ with the write(bal_y) of T₇.
- Change the order of the read(bal_x) of T₈ with the read(bal_y) of T₇.
- Change the order of the read(bal_x) of T₈ with the write(bal_y) of T₇.

Figure 19.7 Equivalent schedules: (a) nonserial schedule S_1 ; (b) nonserial schedule S_2 equivalent to S_1 ; (c) serial schedule S_3 , equivalent to S_1 and S_2 .



560 Chapter 19 Transaction Management

Schedule S_3 is a serial schedule and, since S_1 and S_2 are equivalent to S_3 , S_1 and S_2 are serializable schedules.

This type of serializability is known as conflict serializability. A conflict serializable schedule orders any conflicting operations in the same way as some serial execution. Under the constrained write rule (that is, a transaction updates a data item based on its old value, which is first read by the transaction), a precedence (or serialization) graph can be produced to test for conflict serializability. For a schedule S, a precedence graph is a directed graph G = (N, E) that consists of a set of nodes N and a set of directed edges E, which is constructed as follows:

- Create a node for each transaction.
- Create a directed edge $T_i \rightarrow T_j$, if T_j reads the value of an item written by T_i .
- Create a directed edge $T_i \rightarrow T_j$, if T_j writes a value into an item after it has been read by T_i .
- Create a directed edge $T_i \rightarrow T_j$, if T_j writes a value into an item after it has been written by T_i .

If an edge $T_i \rightarrow T_j$ exists in the precedence graph for S, then in any serial schedule S' equivalent to S, T_i must appear before T_j . If the precedence graph contains a cycle the schedule is not conflict serializable.

Example 19.4 Non-conflict serializable schedule

Consider the two transactions shown in Figure 19.8. Transaction T_9 is transferring £100 from one account with balance bal_x to another account with balance bal_y, while T_{10} is increasing the balance of these two accounts by 10%. The precedence graph for this schedule, shown in Figure 19.9, has a cycle and so is not conflict serializable.



Figure 19.8 Two concurrent update transactions

9 19.9

561



Figure 19.9 Precedence graph for Figure 19.8.

View serializability

There are several other types of serializability that offer less stringent definitions of schedule equivalence than that offered by conflict serializability. One less restrictive definition is called view serializability. Two schedules S_1 and S_2 consisting of the same operations from *n* transactions T_1, T_2, \ldots, T_n are view equivalent if the following three conditions hold:

- For each data item x, if transaction T_i reads the initial value of x in schedule S_1 , then transaction T_i must also read the initial value of x in schedule S_2 .
- For each read operation on data item x by transaction T_i in schedule S_1 , if the value read by x has been written by transaction T_j , then transaction T_i must also read the value of x produced by transaction T_j in schedule S_2 .
- For each data item x, if the last write operation on x was performed by transaction T_i in schedule S_1 , the same transaction must perform the final write on data item x in schedule S_2 .

A schedule is view serializable if it is view equivalent to a serial schedule. Every conflict serializable schedule is view serializable, although the converse is not true. For example, the schedule shown in Figure 19.10 is view serializable, although it is not conflict serializable. In this example, transactions T_{12} and T_{13} do not conform to the constrained write rule; in other words, they perform *blind writes*. It can be shown that any view serializable schedule that is not conflict serializable contains one or more blind writes.



Figure 19.10 View serializable schedule that is not conflict serializable.



If upgrading of locks is allowed, upgrading can take place only during the growing phase and may require that the transaction wait until another transaction releases a shared lock on the item. Downgrading can take place only during the shrinking phase. We now look at how two-phase locking is used to resolve the three problems identified in Section 19.2.1.

Example 19.6 Preventing the lost update problem using 2PL

A solution to the lost update problem is shown in Figure 19.11. To prevent the lost update problem occurring, T_2 first requests an exclusive lock on bal_x. It can then proceed to read the value of bal_x from the database, increment it by £100, and write the new value back to the database. When T_1 starts, it also requests an exclusive lock on bal_x. However, because the data item bal_x is currently exclusively locked by T_2 , the request is not immediately granted and T_1 has to wait until the lock is released by T_2 . This occurs only once the commit of T_2 has been completed.



Figure 19.11 Preventing the lost update problem.

Example 19.7 Preventing the uncommitted dependency problem using 2PL

A solution to the uncommitted dependency problem is shown in Figure 19.12. To prevent this problem occurring, T_4 first requests an exclusive lock on bal_x. It can then proceed to read the value of bal_x from the database, increment it by £100, and write the new value back to the database. When the rollback is executed, the updates of transaction T_4 are undone and the value of bal_x in the database is returned to its original value of £100. When T_3 starts, it also requests an exclusive lock on bal_x. However, because the data item bal_x is currently exclusively locked by T_4 , the request is not immediately granted and T_3 has to wait until the lock is released by T_4 . This occurs only once the rollback of T_4 has been completed.

566 Chapter 19 Transaction Management

Figure 19.12 Preventing the uncommitted dependency problem.

Time T3	and a start of the second s Second second	T ₄	bal
6		been impraction	100
1		write lock(bal.)	100
f g	and the second second second	read(bal,)	100
t ₄ begin_transaction	President of the second	bal_x = bal_x + 100	-100
ts write_look(ba	њ	write(bal _x)	200
t ₆ WAIT	and the second second	rollback/unlock(bal _k)	100
ty read(Dal _x)	10		1000
le write(bal.)			- 100
t ₁₀ commit/unlock(b	al.)		± 90
and the second	Standard States	A. F. S.	2019年4月1日1月1日1日

Example 19.8 Preventing the inconsistent analysis problem using 2PL

- Aller

A solution to the inconsistent analysis problem is shown in Figure 19.13. To prevent this problem occurring, T_5 must precede its reads by exclusive locks, and T_6 must precede its reads with shared locks. Therefore, when T_5 starts it requests and obtains an exclusive lock on bal_x. Now, when T_6 tries to share lock bal_x the request is not immediately granted and T_6 has to wait until the lock is released, which is when T_5 commits.

Time	τ,	T ₆	bal _x	bal	y bai _z	TE SUM
4		begin_transaction	100	50	25	
±2	. begin_transaction	sum = 0	100	- 50	25	• 0
t,	write_lock(bal _x)		100	50	25	0
4	read(bal _x)	read_lock(balx)	100	50	25 -	0
t5	bal _x =bal _x -10	WAIT	100	50	25; 🚖	0,
t6 1 1	write(bal _x)	WAIT	90	50	- 25	0
t7 .	write_lock(balz)	WAIT	. 90	50	25	0
tg	read(balz)	TIAW	90	50	25	. 0
tg i	bal_z = bal_z + 10	WAIT	90	50	.25 -	, O
t10	write(balz)	WAIT	90 -	50	35	0
t11	commit/unlock(bal _x , bal _z)	WAIT	90	50	35.	0
t12		read(baly)	. 90	50	. 35	0
t ₁₃		sum = sum + bal _x .	90	50	35	90
t ₁₄		read_lock(baly)	90	50	35	90
t15		read(baly)	90 /	50	35	90
t ₁₆		sum = sum + baly	90	50	35	140
t ₁₇		read_lock(balz)	90	50	35	140
t ₁₈	$ \frac{\partial M}{\partial x} = \sum_{i=1}^{n} \sum_{j=1}^{n} \frac{\partial M}{\partial x_{ij}} + \sum_{i=1}^{n} \frac{\partial M}{\partial x_{ij}} + \sum_$	read(bal _z)	90	50	35	140
t19 .	Sec. Sec.	sum = sum + balz	90	50	35	175
t ₂₀		commit/unlock(baly, baly, bal,)	90	50	35	175

Figure 19.13 Preventing the inconsistent analysis problem. It can be proved that if *every* transaction in a schedule follows the two-phase locking protocol, then the schedule is guaranteed to be conflict serializable (Eswaran *et al.*, 1976). However, while the two-phase locking protocol guarantees serializability, problems can occur with the interpretation of when locks can be released, as the next example shows.

Example 19.9 Cascading rollback

Consider a schedule consisting of the three transactions shown in Figure 19.14, which conforms to the two-phase locking protocol. Transaction T_{14} obtains an exclusive lock on bal, then updates it using bal, which has been obtained with a shared lock, and writes the value of bal, back to the database before releasing the lock on bal,. Transaction T_{15} then obtains an exclusive lock on bal, reads the value of bal, from the database, updates it, and writes the new value back to the database before releasing the lock. Finally, T_{16} share locks bal, and reads it from the database. By now, T_{14} has failed and has been rolled back. However, since T_{15} is dependent on T_{14} (it has read an item that has been updated by T_{14}), T_{15} must also be rolled back. Similarly, T_{16} is dependent on T_{15} , so it too must be rolled back. This situation, in which a single transaction leads to a series of rollbacks, is called **cascading rollback**.



