

Syllabus

Course Overview

This course focuses on the data access layer of an application. You will explore application objects that interact with an assortment of databases and data repositories, and examine data modeling, connections, queries, and data manipulation for contemporary applications.

Technology Resources

Capella offers tutorials, labs, or a virtual desktop as part of this course. These resources offer software or guided practice in performing tasks related to achieving course competencies and completing assessments. If you require the use of assistive technology or alternative communication methods to participate in these activities, please contact DisabilityServices@Capella.edu to request accommodations.

Disability Services

Note: If you require the use of assistive technology or alternative communication methods to participate in any activity in this course, please contact DisabilityServices@Capella.edu to request accommodations.

Course Competencies

(Read Only)

To successfully complete this course, you will be expected to:

- 1 Apply database management system features and tools.
- 2 Manage a database administration infrastructure.
- 3 Troubleshoot database administration issues.
- 4 Implement security strategies and policies within a database environment.
- 5 Communicate effectively and professionally.

Course Prerequisites

Prerequisite(s): Completion of or concurrent registration in IT4731 or IT-FP4731.

The materials listed below are required to complete the learning activities in this course.

Library

The following required readings are provided in the Capella University Library or linked directly in this course. To find specific readings by journal or book title, use [Journal and Book Locator](#). Refer to the [Journal and Book Locator library guide](#) to learn how to use this tool.

- Cherry, D. (2015). [Securing SQL Server: Protecting your database from attackers \(3rd ed.\)](#). Rockland, MA: Syngress.
- Giesenow, H. (2015). [OWASP Top 10: SQL Server injection mitigation \[Video\]](#). Skillsoft Ireland.
- Petković, D. (2017). [Microsoft SQL Server 2016: A beginner's guide \(6th ed.\)](#). New York, NY: McGraw-Hill/Osborne.
- Sampson, A. (n.d.). [Transactions and isolation levels \[Tutorial\]](#). Skillsoft Ireland.
- Sampson, A. (n.d.). [Triggers and functions \[Tutorial\]](#). Skillsoft Ireland.
- Sampson, A. (n.d.). [Microsoft SQL Server 2016: Backing up databases \[Tutorial\]](#). Skillsoft Ireland.
- Sampson, A. (n.d.). [Microsoft SQL Server 2016: Data access and permissions \[Tutorial\]](#). Skillsoft Ireland.
- Sampson, A. (n.d.). [Performance tuning \[Tutorial\]](#). Skillsoft Ireland.
- Sampson, A. (n.d.). [Programmability objects \[Tutorial\]](#). Skillsoft Ireland.

External Resource

Please note that URLs change frequently. While the URLs were current when this course was designed, some may no longer be valid. If you cannot access a specific link, contact your instructor for an alternative URL.

Permissions for the following links have been either granted or deemed appropriate for educational use at the time of course publication.

- Microsoft. (2017). [Install SQL Server](#). Retrieved from <https://docs.microsoft.com/en-us/sql/database-engine/install-windows/install-sql-server?view=sql-server-2017>
- Microsoft. (2017). [Tutorials for SQL Server Management Studio \(SSMS\)](#). Retrieved from <https://docs.microsoft.com/en-us/sql/ssms/tutorials/tutorial-sql-server-management-studio?view=sql-server-2017>

Suggested

The following materials are recommended to provide you with a better understanding of the topics in this course. These materials are not required to complete the course, but they are aligned to course activities and assessments and are highly recommended for your use.

Optional

The following optional materials are offered to provide you with a better understanding of the topics in this course. These materials are not required to complete the course.

Library

The following optional readings may be available in the Capella University Library. To find specific readings by journal or book title, use [Journal and Book Locator](#). Refer to the [Journal and Book Locator library guide](#) to learn how to use this tool. If the full text is not available, you may be able to request a copy through the [Interlibrary Loan](#) service.

- Beard, B. (2018). [Beginning backup and restore for SQL Server: Data loss management and prevention techniques](#). New York, NY: Apress.
- Fritchey, G. (2018). [SQL Server 2017 query performance tuning: Troubleshoot and optimize query performance \(5th ed.\)](#). New York, NY: Apress.
- Nevarez, B. (2016). [High performance SQL server: The go faster book](#). New York, NY: Apress.

External Resource

Please note that URLs change frequently. While the URLs were current when this course was designed, some may no longer be valid. If you cannot access a specific link, contact your instructor for an alternative URL. Permissions for the following links have been either granted or deemed appropriate for educational use at the time of course publication.

- Erkeç, E. (2018). [T-SQL best practices](#). Retrieved from <https://codingsight.com/t-sql-best-practices/>
- Instant SQL Formatter. (n.d.). [SQL Online](#). Retrieved from <http://www.dpriver.com/pp/sqlformat.htm>
- Sheldon, R. (2009). [Transact-SQL formatting standards \(coding styles\)](#). Retrieved from <https://www.red-gate.com/simple-talk/sql/t-sql-programming/transact-sql-formatting-standards-coding-styles/>
- Sheldon, R. (2017). [The basics of good T-SQL coding style](#). Retrieved from <https://www.red-gate.com/simple-talk/sql/t-sql-programming/basics-good-t-sql-coding-style/>

Unit 1 >> Database Backup and Restore

Introduction

In this unit, you will:

- Discuss different database backup and recovery scheduled plans.
- Create a database backup and recovery plan suitable for guiding developers.

Whenever we work on a database, especially if new features are being tested on, it is a good idea to perform a backup. This way, if we accidentally modify or delete data, we are able to quickly restore it to its previous state. By default, the Microsoft SQL Server Management Studio (SSMS) creates a full backup of the database and

stores the file for the database in the backup directory, which by default will be under C:\Program Files\Microsoft SQL Server\MSSQL13.SQLEXPRESS\MSSQL\Backup. The file for the backup is the name of the database with the extension .bak.

By default, the backup is set to expire in zero days, which means that the backup file will be saved on the disk until the backup is run again. Once the backup is complete, the old backup file will be replaced by the new backup. For smaller databases, these settings are generally adequate but we can also use incremental backups or set the number of expiration days for the backup higher.

To restore a database from a backup copy, the Restore Database dialog restores the current database to the most recent backup of the database. However, if we want to restore the database to a specific point in time, you can use the Restore Database dialog box to specify a date and time. To access it, you can right click on the database and select **Tasks > Restore > Database**.

The SQL Server 2016 database engine is backwards compatible and can run older versions of SQL Server databases as if it were running an older SQL Server database engine. If you attach a database, you may want to change the compatibility level for the database so that it is appropriate for our purposes. For example, if we attached a database file that was originally created under SQL Server 2008 to the SQL Server 2016 database engine, the compatibility level will remain set to SQL Server 2008. This means that we would be able to use SQL Server 2008 features even if the features are no longer available under SQL Server 2016. However, we will not be able to use the new SQL Server 2016 features for that database. To access this feature, you can right click the database, select *Properties* and click **Option** item to select SQL Server 2016 from the Compatibility Level drop-down list.

Learning Activities

u01s1 - Studies

Skillsoft Resources

Read or view the following resources related to database backup and recovery:

- Petković, D. (2017). [Microsoft SQL Server 2016: A beginner's guide \(6th ed.\)](#). New York, NY: McGraw-Hill/Osborne.
 - Chapter 17, "Automating System Administration Tasks."
- Sampson, A. (n.d.). [Microsoft SQL Server 2016: Backing up databases \[Tutorial\]](#). Skillsoft Ireland.
 - Back Up Databases.
 - Using Recovery Models.

u01s1 - Learning Components

- Write command line code to backup and restore data.

- Create database tables in SQL Server.
- Import data into SQL Server.
- Understand how to backup and recover data using MS SQL Server using SQL Server Management Studio.
- Explain how to set up automated backups in SSMS.
- Run full backups in SSMS.
- Create incremental backups in SSMS.

u01s2 - Software Preparation and Technology Access

In this course, you will be using software and technology needed to complete designated activities and assignments. There is no additional cost for this software and technology. Some software packages will be made available to you at no additional cost through Capella's subscription with Microsoft, while other software packages are available for free download through open-source licensing.

Capella University requires learners to meet certain minimum [computer requirements](#). Please note that some software required for a course may exceed these minimum requirements. Check the requirements for the software you may need to download and install to make sure it will work on your device. Most software will require a Windows PC. If you use a Mac, refer to the [Installing a Windows Virtual Environment](#)

The software and technologies below are strongly recommended to support you in completing the course objectives. If you have access to other tools that you believe may still meet the requirements of this course, please discuss your selected alternatives with your instructor.

If you require the use of assistive technology or any alternative communication methods to access course content, please contact DisabilityServices@Capella.edu with any access-related questions or to request accommodations.

The database server you will use in this course is Microsoft SQL Server 2017. The software you will use in this course is Microsoft SSMS. Follow the instructions provided through the links below to download and install software or register for an account, as required.

Microsoft Software

1. Visit Capella's [Microsoft Software](#) page for instructions on obtaining free Microsoft software.
2. Identify the version of MS Visio / Project / Access / Visual Studio Enterprise / SQL Server / Etc. that is compatible with your operating system.
3. Download and install.

For more information on downloading and installing SQL Server, see the following resources:

- Microsoft. (2017). [Install SQL Server](https://docs.microsoft.com/en-us/sql/database-engine/install-windows/install-sql-server?view=sql-server-2017). Retrieved from <https://docs.microsoft.com/en-us/sql/database-engine/install-windows/install-sql-server?view=sql-server-2017>

- Microsoft. (2017). [Tutorials for SQL Server Management Studio \(SSMS\)](https://docs.microsoft.com/en-us/sql/ssms/tutorials/tutorial-sql-server-management-studio?view=sql-server-2017). Retrieved from <https://docs.microsoft.com/en-us/sql/ssms/tutorials/tutorial-sql-server-management-studio?view=sql-server-2017>

u01s3 - Optional Resources

Skillsoft Resources

You may read the following:

- Beard, B. (2018). [*Beginning backup and restore for SQL Server: Data loss management and prevention techniques*](#). New York, NY: Apress.
 - Chapters 1, 2, and 5.

Code Formatting and Optimization

Optimization

You may read the following:

- Sheldon, R. (2017). [The basics of good T-SQL coding style](https://www.red-gate.com/simple-talk/sql/t-sql-programming/basics-good-t-sql-coding-style/). Retrieved from <https://www.red-gate.com/simple-talk/sql/t-sql-programming/basics-good-t-sql-coding-style/>
- Erkeç, E. (2018). [T-SQL best practices](https://codingsight.com/t-sql-best-practices/). Retrieved from <https://codingsight.com/t-sql-best-practices/>

Formatting

You may read the following:

- Instant SQL Formatter. (n.d.). [SQL Online](http://www.dpriver.com/pp/sqlformat.htm). Retrieved from <http://www.dpriver.com/pp/sqlformat.htm>
- Sheldon, R. (2009). [Transact-SQL formatting standards \(coding styles\)](https://www.red-gate.com/simple-talk/sql/t-sql-programming/transact-sql-formatting-standards-coding-styles/). Retrieved from <https://www.red-gate.com/simple-talk/sql/t-sql-programming/transact-sql-formatting-standards-coding-styles/>

u01d1 - Scheduled Backup and Recovery

Having a structured database backup and recovery plan can be quite important to ensure the safety and security of the database. Because the database is the central piece to any organizational software application, it is crucial that if there are database issues, the organization be able to recover quickly.

Scenario

SmartHomes is a mid-sized organization that has a majority of their business transactions performed during the opening hours of their business (8 a.m.–9 p.m. local time). They also have an online website and business-to-business transactions that occur throughout the week. SmartHomes has service level agreements that their online services, including their database, need to be up and running 99.99 percent of the time. SmartHomes's database transactions are the lowest on Sundays at 1 a.m.–2 a.m. local time.

Discuss

Consider the scenario above (or describe one of your own) and search the Internet for database backup plans. Select two of them and summarize each in your post (also share links to them). Discuss the pros and cons of applying each to the scenario above. Explain which one you think is the most appropriate and why.

Response Guidelines

Review the posts of your peers and respond to two. Offer comments or arguments that contribute to your peers' ideas.

Course Resources

Undergraduate Discussion Participation Scoring Guide

u01d1 - Learning Components

- Perform backup and recovery process using various methods in SQL Server.

u01a1 - Database Backup and Recovery

Overview

The unit introduction discussed creating a backup using the SSMS tool. On top of that, you can also perform backup and recovery using T-SQL or automate the backup process rather than manually performing them. The DEV and TEST environments are frequently outdated and generally can require a refresh of data from the PROD environment.

In this assignment, you will backup and recover data and create a plan that documents the procedure.

Preparation

By now, you should have installed MS SQL Server 2017.

Access the following documents are given in the resources:

- Backup and Recovery Template.
 - Use this to complete this assignment.
- SmartHomesData.
 - Use data in this Excel document to populate MS SQL Server database.

Directions

Use the Backup and Recovery Template to complete and submit the following:

Part 1: Database Backup and Recovery

Use of SQL Server to do the following:

1. Create the SmartHomes database and tables and load the data.
2. Perform a proper backup and recovery using SSMS.
3. Perform a proper backup and recovery using T-SQL.
4. Set up an automated backup process to run a full backup every Sunday at midnight.
5. Set up the automated process to create incremental backups on all other days other than Sunday.

Part 2: Database Backup and Recovery Plan

Create a step-by-step procedure for use by subsequent developers to guide them to proper backup and restore the data files. Include supporting graphics to illustrate important or complex aspects of the backup.

Submit your completed template to this assignment.

Course Resources
Backup and Recovery Template [DOCX]
SmartHomesData [XLSX]

Introduction

In this unit, you will:

- Discuss the pros and cons of different database procedures.
- Create a stored procedure that takes in the customer ID, product ID, and quantity, and execute the associated transaction.

A stored procedure is a database object that contains one or more SQL statements. To create a stored procedure, we can use the CREATE PROC statement. When the procedure is executed for the first time, each SQL statement is compiled and executed to create an execution plan. Then the procedure is stored in compiled form within the database. So, for each subsequent execution, the SQL statements are executed without compilation since they are then precompiled. This makes the execution of the stored procedure faster than the execution of the SQL statements using a separate script.

To execute a stored procedure, we can use the EXEC statement. We can call a stored procedure from another stored procedure or even in itself through recursive calling, but this is rarely used in SQL programming but it is possible. One big advantage to using stored procedures is that they can help to restrict and control access to the database to help prevent accidental errors in SQL statements as well as the possibility of malicious users.

In a stored procedure, we can add in input and output parameters to accept values passed from the calling program or have values passing back to the calling program. We can also incorporate optional parameters that do not require a value to be passed in but we will need to assign it a default value so that this value will be used if it is not passed from the calling program.

A trigger is a special type of procedure executed automatically, if we have an action that is executed on a table or view. They provide a powerful way to control how action queries will modify the data in our database. We can use these triggers to help enforce business logic and help prevent data inconsistency, but they can cause some potential locking and performance issues if we overly use them.

Learning Activities

u02s1 - Studies

Skillsoft Resources

Read or view the following resources related to database backup and recovery and triggers:

- Petković, D. (2017). [*Microsoft SQL Server 2016: A beginner's guide \(6th ed.\)*](#). New York, NY: McGraw-Hill/Osborne.
 - Chapter 8, "Stored Procedures and User-Defined Functions."
 - Chapter 14, "Triggers."

- Sampson, A. (n.d.). [Triggers and functions \[Tutorial\]](#). Skillsoft Ireland.
 - Triggers.
 - Functions.
- Sampson, A. (n.d.). [Microsoft SQL Server 2016: Backing up databases \[Tutorial\]](#). Skillsoft Ireland.
 - Creating Users and Configuring Access.
- Sampson, A. (n.d.). [Programmability objects \[Tutorial\]](#). Skillsoft Ireland.
 - Stored Procedures.
- Cherry, D. (2015). [Securing SQL Server: Protecting your database from attackers \(3rd ed.\)](#). Rockland, MA: Syngress.
 - Chapter 9, "Chapter 9: SQL Injection Attacks."
- Giesenow, H. (2015). [OWASP Top 10: SQL Server injection mitigation \[Video\]](#). Skillsoft Ireland.

u02s1 - Learning Components

- Write code that implements save points.
- Write code that executes stored procedures.
- Describe exception handling.
- Explain the execution process of a stored procedure and trigger.
- Write code that creates a trigger in T-SQL.
- Write code that causes a trigger to be executed.
- Explain how to prevent SQL injections.

u02d1 - Database Versus Application Code

Some organizations have all of the database-related code created within the database as functions, stored procedures, and triggers. Other organizations only use the database for data storage.

Research online to find some pros and cons of each approach. If the database is only meant to be used by one application, provide an argument for which approach would be the best. If the database is planned to be used by multiple applications, provide an argument of which approach would be the best.

Response Guidelines

Respond to at least two of your peers' posts. Discuss similarities and differences that you find in the posts.

- Explain transactions in SQL Server.

u02a1 - Creating a Stored Procedure

Overview

Database stored procedures and triggers are frequently set up to help facilitate data entry and validation in a consistent manner. Rather than having users have direct access to insert, update, or delete data in tables, developers can make use of the stored procedures to make those calls. Triggers are used to keep data accurately updated.

In this assignment, you will implement a stored procedure to handle a complete transaction. You will also implement a trigger to validate data before a transaction is permitted to execute.

Preparation

Complete the following:

- Open the SmartHomesData Excel document in MS SQL Server.
- Access the Stored Procedure and Trigger Template given in the resources.
 - Use this template to complete this assignment.

Directions

Use the Stored Procedure and Trigger Template to complete the following:

Part 1: Database Stored Procedure and Testing

Using the SmartHomes database do the following:

1. Create and run a stored procedure (write code) in SQL Server to take in a customer ID, product ID, and quantity.
2. Test the three executions of the stored procedure (that include valid and invalid data) to help test extraneous information.
3. Explain how the transaction is executed within the stored procedure.

Part 2: SQL Trigger and Testing

To ensure data integrity, we need to validate that the existing quantity of the product table has enough items before the stored procedure is executed. This can be done using a trigger.

1. Create a trigger in SQL Server to execute when the transaction is performed in the stored procedure. The trigger should be based on the update of the product table to check if the current quantity has enough items to deduct the updated quantity.
2. Test three executions of the trigger with various inputs (that include valid and invalid data) to test for extraneous information.
3. Explain how the trigger helps prevent SQL injection.

Submit your completed template to this assignment.

Course Resources

[Stored Procedure and Trigger Template \[DOCX\]](#)

Unit 3 >> Data Concurrency and Transactions

Introduction

In this unit, you will:

- Discuss database locking conflicts.
- Resolve four concurrency problems.

In any database, we often have two or more users concurrently accessing the same set of database data. This is called concurrency and it is not an issue if the two users are just retrieving the same data at the same time. Where it can be an issue if one of the users is updating, deleting, or inserting into that dataset while the other user is trying to retrieve the data. We will learn about concurrency and see how SQL Server uses locking to prevent concurrency problems. We will also learn how to control the types of problems that are allowed.

If we have a database that has a small number of users, the chances of concurrency issues are quite low. However, as the number of users and transactions increases, we should expect to see concurrency issues. One method that we can apply is to use locking to help avoid concurrency issues. By holding a lock on the data, the transaction can prevent others from using that data. Once the transaction is complete and releases the lock, the next transaction can work with the data. SQL Server automatically enables and manages locking and will prevent some of the concurrency issues. However, as the system becomes larger, the default locking mechanism does not work as well, and we will need to override the default locking behavior.

To change the default locking behavior in SQL Server, we can use the `SET TRANSACTION ISOLATION LEVEL` statement to set the transaction isolation level for the current session.

There are four concurrency problems that the use of locks can prevent:

1. **Lost update:** If we have two transactions that select and then update the data from the same row in the same table, and if these two transactions are submitted at the same time, the one that executes first will be overwritten by the one that executes second. This means that one of the two updates is lost.
2. **Dirty reads:** This can happen if a transaction reads data not committed by another transaction. If one user changes a row, the second user can then read the changed row before the first user commits the change. If the first user then rolls back the change, the second user has selected a row that does not actually exist in the database.
3. **Inconsistent analysis:** This can occur if we have two reads of the same data result in different values because a third user has updated the data in between the first two users reading the data.
4. **Phantom reads:** This can occur when we perform an update or a delete on a set of rows when another transaction is performing an insert or a delete that affects one or more rows in the same set of rows. For example, if a first user updates all products to have the shipping price to be 0, and if the second user inserts a new product that includes a shipping cost, once the first user is complete, there are still products that has a shipping charge.

Learning Activities

u03s1 - Studies

Skillsoft Resources

Read or view the following resources related to database backup and recovery:

- Petković, D. (2017). [Microsoft SQL Server 2016: A beginner's guide \(6th ed.\)](#). New York, NY: McGraw-Hill/Osborne.
 - Chapter 7, "Modification of a Table's Contents."
 - Chapter 13, "Concurrency Control."
- Sampson, A. (n.d.) [Transactions and isolation levels \[Tutorial\]](#). Skillsoft Ireland.
 - Transaction Statements.
 - Isolation Levels.

u03s1 - Learning Components

- Modify the transaction isolation level in SQL Server Management Studio.
- Write code that modifies the transaction isolation level.
- Identify typical concurrency problems found in SQL Server.
- Write code to handle transaction isolation levels.
- Explain the various transaction isolation levels in SQL Server.

u03d1 - Lock Conditions

Imagine that users of your organization's e-commerce application have been reporting issues with pages hanging when they are trying to update content. Your organization's DBA has reported that there have been some locking conflicts that may be the cause.

Discuss a possible cause and solution for one of the following database lock origins:

- Table lock.
- Row lock.
- Dead lock.

Run a simulation that reflects the lock in SSMS. Take a screenshot of the lock and another depicting the results after applying your solution. Share the code and screenshots in your post.

Discuss:

- Problems or challenges you encountered and how you solved or were stifled by them.
- Other possible solutions for the problem.

Response Guidelines

Review the posts of your peers and respond to two. Offer comments or arguments that contribute to your peers' ideas.

Course Resources

Undergraduate Discussion Participation Scoring Guide

u03a1 - Concurrency Problems

Overview

In the unit introduction, you read about four concurrency problems that the use of locks can prevent. In this assignment, you will write code that demonstrates the lost update, dirty reads, inconsistent analysis, and phantom reads concurrency problems. You will next resolve the concurrency issues based on the transaction isolation levels and document why your selected transaction isolation level is the optimal choice for that transaction.

Preparation

Access the Concurrency Template given in the resources, review the four concurrency problems within it, and use the document to complete this assignment.

Directions

Use the Concurrency Template to complete the following **for each** problem listed in the template.

- Create and execute the code for each transaction.
- Describe briefly the main issue that is caused for each concurrency problem.
- Resolve the transaction at the isolation level by making updates to either SSMS or in the code.
- Explain briefly your approach to resolving the problem and why the selected transaction isolation level used to resolve the concurrency problem is the optimal choice.

Submit your completed template to this assignment.

Course Resources

Concurrency Template [DOCX]

Unit 4 >> User Roles and Permissions

Introduction

In this unit, you will:

- Discuss the different approaches to handling database user accounts.
- Create various scripts to permit database communication.

As you have seen already, any connection to SQL Server is done through the database user accounts. Some administrative accounts are created automatically, such as SA. This account has database administration privileges.

Each user account has a unique user name and all of the attributes belonging to the user, including the password, role, privileges, and default schema for any database objects that the user creates. When you create a new user, a schema for the user is also created with the user's user name. Since these objects belong to the user, if the user is dropped from the database, the schema objects must also be dropped first (or moved to another user) before the user can be removed.

There are two different main privileges when it comes to users: system privileges and object privileges. System privileges focus on the ability to perform an action such as creating tables or updating rows in any table. Object

privileges focus on an action done on a given object within a schema. For example, updating the rows on a table would be considered an object privilege.

Learning Activities

u04s1 - Studies

Skillsoft Resources

Read or view the following resources related to database backup and recovery:

- Petković, D. (2017). [Microsoft SQL Server 2016: A beginner's guide \(6th ed.\)](#). New York, NY: McGraw-Hill/Osborne.
 - Chapter 12, "Security System of the Database Engine."
- Sampson, A. (n.d.). [Microsoft SQL Server 2016: Data access and permissions \[Tutorial\]](#). Skillsoft Ireland.
 - Working With Roles.
 - Configuring Data Security.

u04s1 - Learning Components

- Create user accounts using SQL Server Management Studio.
- Create user accounts using T-SQL.
- Grant privileges in T-SQL to user accounts.
- Grant permissions to user accounts.
- Apply permission changes in SQL Server Management Studio.
- Write code to grant permissions to roles and accounts.

u04s2 - Optional Resources

Skillsoft Resources

You may read the following:

- Cherry, D. (2015). [Securing SQL Server: Protecting your database from attackers \(3rd ed.\)](#). Rockland, MA: Syngress.
 - Chapters 1, 5, 6, and 12.

u04d1 - User Account Permissions

Some organizations create individual user accounts in the database while others create a generic database account for certain groups of individuals.

Research online to explain the differences between both approaches. Which approach would you recommend for a mid-sized company and why? Would this approach change if it were for a smaller or larger organization?

Response Guidelines

Review the posts of your peers and respond to two. Offer comments or arguments that contribute to your peers' ideas.

Course Resources

Undergraduate Discussion Participation Scoring Guide

u04a1 - User Roles and Permissions

Overview

Imagine your organization's environment requires you to set up various database user accounts to support a new application. An administrator's script has to be written to connect to the database without having the password embedded in the script. In this assignment, you will create user accounts and scripts to control database functions such as this.

Preparation

Access the User Roles and Accounts Template given in the resources and use it to complete this assignment.

Directions

Use the User Roles and Accounts Template to complete and submit the following:

1. **OS User Account:** Create an OS user account that uses Windows authentication and grants the user database access.
2. **Developer User Account:** Create a user account using T-SQL for developers named DEVELOPER with the password TESTACCOUNT that grants the user the ability to:

- Select and modify any table.
- Connect to and have access to all resources.

3. **QA User Account:** Create a user account using T-SQL for Quality Assurance with the password QAACCOUNT, that:

- Forces the user to change the password when first logging on.
- Grants the user access to select any table and to connect and have access to all resources.
- Allows the user to connect and have access to all resources.

4. **Locked Script:** Create a script that:

- Removes the DEVELOPER account.
- Locks the QA account.
- Expires the Password for the QA account.

Submit your completed template to this assignment.

Course Resources

User Roles and Accounts Template [DOCX]

Unit 5 >> Performance Tuning

Introduction

In this unit, you will:

- Discuss two different methods that can be used to optimize query performance in SQL Server and compare these two methods to the use of execution plans in SSMS.
- Optimize several queries.

During the process of query optimization, the task of the optimizer is to find the most efficient execution plan for a given query. The first phase is to parse the query's syntax and to transform it into a tree. The validation is performed on all of the database objects that are references by the query. For example, it will check on the table names and column names to verify that they exist. The next phase for the query optimizer is to compile the query tree.

For the next phase, the query optimizer takes the compiled query tree and investigates multiple access strategies before it decides how to process the given query. To find the most efficient execution plan, the query optimizer first makes the query analysis during which it searches for search arguments and join conditions. The optimizer then selects which indices to use. Finally, if join operations exist, the optimizer selects the join order

and chosen one of the join processing techniques. Once the execution plan is generated, it is permanently stored and executed.

During the query analysis, the optimizer checks the query for search arguments, the use of the OR operator and if there are any joins in the query in that order. The search arguments are the part of the query that restricts the result set of the query.

Two examples include:

```
firstName = 'Bob'
```

```
price >= 50
```

There are several expression forms that cannot be optimized in search arguments and should be evaluated to determine if these search arguments can be changed. The first group has all of the expressions using the NOT or <> operator. The other type of expression is if the expression is on the left side of the operator. In each of these cases, the optimizer uses a full table scan.

For example:

```
firstName NOT IN ('Bob','Carl')
```

```
customerID <> 1000
```

```
price * 10 > 25
```

During this phase, the optimizer also checks each search argument to see if there are any indices in relation to the corresponding expression. If there is an index available, the optimizer decides whether or not to use it. This depends on the selectivity of the expression. The selectivity of the expression is defined as the ratio of the number of rows that satisfy the condition compared to the number of total rows in the table. If there are a small number of rows that satisfies the condition, the higher the level of selectivity exists for that filter.

Learning Activities

u05s1 - Studies

Skillsoft Resources

Read or view the following resources related to database backup and recovery:

- Petković, D. (2017). [Microsoft SQL Server 2016: A beginner's guide \(6th ed.\)](#). New York, NY: McGraw-Hill/Osborne.
 - Chapter 6, "Queries."
 - Chapter 10, "Indices."
 - Chapter 19, "Query Optimizer."
 - Chapter 20, "Performance Tuning."

- Sampson, A. (n.d.). [Performance tuning \[Tutorial\]](#). Skillsoft Ireland.
 - SQL Server Performance.
 - Monitoring Tools.
 - Extended Events.
 - SQL Trace.

u05s2 - Optional Resources

Skillsoft Resources

You may read the following:

- Fritchey, G. (2018). [SQL Server 2017 query performance tuning: Troubleshoot and optimize query performance \(5th ed.\)](#). New York, NY: Apress.
 - Chapters 1, 7, and 8.
- Nevarez, B. (2016). [High performance SQL server: The go faster book](#). New York, NY: Apress.
 - Chapter 8.

u05d1 - Execution Plans

Query optimization is crucial to implement to ensure that each query runs optimally in the database. There are many different methods in which can be implemented depending on the type of query that is used.

Research online to find two different methods that can be used to optimize query performance in SQL Server. Compare these two methods to the use of execution plans within SSMS.

Response Guidelines

Review the posts of your peers and respond to two. Offer comments or arguments that contribute to your peers' ideas.

u05d1 - Learning Components

- Describe query optimization techniques.
- Write code that uses optimizer hints.
- Write code to create indices.

u05a1 - Query Optimization

Overview

Queries written by developers are not always optimized for performance. It is crucial to ensure that all queries are fully optimized in the database whenever possible so that they can execute optimally. There may be changes in code or addition of indices to help with optimization.

In this assignment, you will review six queries that require optimization.

Designers note: Database designers should consider how the eventual growth of the data in tables affect query optimization.

Preparation

Complete the following:

- Make sure to read the guidelines in the unit introduction to help you with this assignment.
- Access the Query Optimization Template given in the resources. Review the six queries found within it and use the document to complete this assignment.

Directions

Use the Query Optimization Template to complete the following for each of the six queries found in the template:

1. Run each query in the query optimization tool in SSMS and take a screenshot.
2. Optimize query performance using appropriate strategies (that is, optimizer hints, additional indices, changing the query search arguments, or some combination thereof as appropriate). Provide a screenshot of the updated query optimization tool result.
3. Explain briefly why you choose the strategies you employed.

Submit the completed template to this assignment.

Course Resources

Query Optimization Template [DOCX]