


MS-SQL SERVER



Microsoft[®]
SQL Server[®]

BEGINNER'S GUIDE

1.

Contents

1. Introduction to SQL Server	4
Overview of database management systems	4
Relational DBMS (RDBMS):	4
Object-oriented DBMS (OODBMS):	4
NoSQL DBMS:	4
DBMSs provide many benefits, including:	4
Improved data sharing:	4
Improved data security:	4
Improved data integrity:	5
Improved data backup and recovery:	5
• History and evolution of SQL Server	5
Enterprise Edition:	7
Standard Edition:	7
Express Edition:	7
Developer Edition:	7
Web Edition:	7
SQL Server on Azure Virtual Machines:	7
Per Core Licensing:	7
Server + CAL Licensing:	7
Azure Hybrid Benefit:	7
2. Relational Database Concepts	8
Tables:	8
Primary key:	8
Foreign key:	8
Normalization:	8
Indexes:	8
SQL: Structured Query Language (SQL)	8
Transactions:	8
ACID:	8
Tables:	9
Columns:	9
Rows:	9
Primary key:	9

Foreign key:.....	9
Normalization:.....	10
Denormalization:.....	10
3. Basic SQL Syntax	11
Creating a database:.....	11
Creating a table:	11
Inserting data:	11
Updating data:.....	11
Deleting data:.....	11
Selecting data:.....	11
4. Data Manipulation Language (DML).....	14
INSERT:.....	14
UPDATE:.....	14
DELETE:	14
5. Data Definition Language (DDL)	15
CREATE:.....	15
ALTER:	15
DROP:.....	15
6. Querying Multiple Tables.....	16
INNER JOIN:.....	16
LEFT JOIN:.....	16
RIGHT JOIN:	16
FULL OUTER JOIN:.....	17
7. Aggregate Functions	18
COUNT():.....	18
SUM():.....	18
AVG():.....	18
MAX():.....	18
MIN():.....	18
8. Subqueries.....	19
Subquery as a derived table in a FROM clause:	19
Subquery as a scalar expression.....	19
9. Indexing.....	20
Types of indexes:.....	20
Creating indexes:.....	20
Maintaining indexes:.....	20

Using indexes:.....	20
Monitoring index performance:	20
10. Stored Procedures	21
Creating stored procedures:	21
Advantages of stored procedures:	21
Executing stored procedures:	21
Modifying and dropping stored procedures:	21
Debugging stored procedures:	21
11. Views.....	22
Creating views:	22
Modifying and dropping views:.....	22
Advantages of views:	22
Using views:.....	22
Limitations of views:	22
12. Transactions and Locking	23
Transaction states:.....	23
ACID properties:	23
Transaction isolation levels:.....	23
Locking:	23
Deadlocks:	23
13. Security.....	24
Authentication:.....	24
Authorization:.....	24
Auditing:.....	24
Encryption:	24
Compliance:.....	24

1. Introduction to SQL Server

Overview of database management systems

A database management system (DBMS) is a software application that enables users to store, organize, retrieve, and manage data in a database. A database is an organized collection of data that can be accessed, managed, and updated easily. A DBMS provides a way to manage databases efficiently and effectively, and it is a critical component in modern information systems.

A DBMS consists of several components, including a database engine, a data definition language (DDL), a data manipulation language (DML), and a query language. The database engine is responsible for managing the data and providing access to it. The DDL is used to define the structure of the database, including tables, fields, and relationships. The DML is used to manipulate the data in the database, including adding, updating, and deleting records. The query language is used to retrieve data from the database.

There are several types of DBMS, including:

Relational DBMS (RDBMS): This type of DBMS uses a table-based data model and organizes data into one or more tables. The RDBMS is the most widely used type of DBMS and is the basis for SQL-based DBMSs.

Object-oriented DBMS (OODBMS): This type of DBMS stores data in objects rather than in tables. The OODBMS is designed to handle complex data types, such as multimedia objects, and is used in applications such as computer-aided design (CAD) and multimedia systems.

NoSQL DBMS: This type of DBMS is designed for handling large volumes of unstructured or semi-structured data. NoSQL DBMSs do not use the traditional table-based structure of RDBMSs and are used in applications such as social media, gaming, and e-commerce.

DBMSs provide many benefits, including:

Improved data sharing: A DBMS enables multiple users to access and share data in a controlled manner.

Improved data security: A DBMS provides security features, such as access controls and encryption, to protect sensitive data.

Improved data integrity: A DBMS enforces data integrity rules, such as referential integrity, to ensure that data is accurate and consistent.

Improved data backup and recovery: A DBMS provides backup and recovery features to protect data from loss or corruption.

In summary, a DBMS is a software application that enables users to store, organize, retrieve, and manage data in a database. DBMSs come in different types, including RDBMS, OODBMS, and NoSQL DBMS. They provide many benefits, including improved data sharing, security, integrity, and backup and recovery.

- ## History and evolution of SQL Server

SQL Server is a popular relational database management system (RDBMS) developed by Microsoft. It was first released in 1989 under the name "SQL Server for OS/2" and has since evolved into a robust, enterprise-class database platform used by organizations of all sizes.

Here is a brief overview of the history and evolution of SQL Server:

SQL Server 1.0: The first version of SQL Server was released in 1989 for the OS/2 operating system. It was based on the Sybase SQL Server engine and was designed to compete with Oracle and IBM's DB2 database platforms.

SQL Server 4.2: In 1992, Microsoft released SQL Server 4.2 for Windows NT. This version included several new features, such as stored procedures, triggers, and views.

SQL Server 6.0: In 1995, Microsoft released SQL Server 6.0, which was a major upgrade from the previous version. It included support for distributed transactions, online backup and restore, and full-text search.

SQL Server 7.0: In 1998, Microsoft released SQL Server 7.0, which introduced several new features, including support for data warehousing, OLAP, and data mining. It was also the first version of SQL Server to include the Enterprise Manager, a graphical management tool.

SQL Server 2000: In 2000, Microsoft released SQL Server 2000, which was a significant upgrade from the previous version. It included support for XML and HTTP, as well as improved scalability and performance.

SQL Server 2005: In 2005, Microsoft released SQL Server 2005, which introduced several new features, including support for native XML data, SQL Server Integration Services (SSIS), and SQL Server Reporting Services (SSRS).

SQL Server 2008: In 2008, Microsoft released SQL Server 2008, which included several new features, such as support for spatial data, policy-based management, and encryption.

SQL Server 2012: In 2012, Microsoft released SQL Server 2012, which included several new features, such as support for columnstore indexes, AlwaysOn Availability Groups, and Power View.

SQL Server 2014: In 2014, Microsoft released SQL Server 2014, which included several new features, such as support for in-memory OLTP, buffer pool extensions, and enhanced AlwaysOn Availability Groups.

SQL Server 2016: In 2016, Microsoft released SQL Server 2016, which included several new features, such as support for JSON data, enhanced security features, and support for R language analytics.

SQL Server 2017: In 2017, Microsoft released SQL Server 2017, which included several new features, such as support for graph data, machine learning services, and enhanced Linux support.

SQL Server 2019: In 2019, Microsoft released SQL Server 2019, which included several new features, such as support for big data clusters, enhanced security features, and improved performance.

SQL Server has evolved over the years to become a robust, enterprise-class database platform that supports a wide range of applications and workloads. With each new release, Microsoft has introduced new features and enhancements to improve performance, scalability, security, and manageability.

- [Editions and licensing options](#)

Microsoft SQL Server offers several editions and licensing options to meet the needs of various types and sizes of businesses. Here are some of the most common editions and licensing options available:

Enterprise Edition: This is the most comprehensive edition of SQL Server and offers advanced features such as data warehousing, advanced analytics, and online transaction processing (OLTP). It is licensed per core, with a minimum of four cores required.

Standard Edition: This edition is suitable for small to medium-sized businesses that require basic database functionality. It includes features such as basic data management, security, and analytics. It is licensed per core, with a minimum of four cores required.

Express Edition: This is a free, lightweight edition of SQL Server that is ideal for small-scale applications and development purposes. It has some limitations on database size and resource usage.

Developer Edition: This edition is designed for developers to use for building and testing applications. It includes all the features of Enterprise Edition and is licensed per user.

Web Edition: This edition is designed for hosting web applications and is licensed per core, with a minimum of four cores required.

SQL Server on Azure Virtual Machines: This option allows businesses to run SQL Server on virtual machines hosted on the Microsoft Azure cloud. It is licensed per hour and includes all the features of SQL Server Enterprise Edition.

In addition to these editions, Microsoft also offers various licensing models, including:

Per Core Licensing: This is the most common licensing model for SQL Server and is based on the number of cores in the server that SQL Server is installed on.

Server + CAL Licensing: This model is based on the number of servers and the number of client access licenses (CALs) required to access those servers.

Azure Hybrid Benefit: This licensing model allows businesses to use existing licenses for SQL Server on Azure virtual machines and save up to 40% on licensing costs.

Overall, the choice of edition and licensing option depends on the specific needs and budget of the business.

2. Relational Database Concepts

Relational databases are a type of database management system (DBMS) that store and manage data in a tabular format, organized in rows and columns. Here are some key concepts related to relational databases:

Tables: Relational databases store data in tables, which consist of rows and columns. Each row represents a single record or instance of data, and each column represents a specific attribute or field of that data.

Primary key: Each table in a relational database has one or more columns that uniquely identify each row. This column or columns are known as the primary key. The primary key is used to link data between tables.

Foreign key: A foreign key is a column or group of columns in one table that refers to the primary key of another table. This allows for the creation of relationships between tables.

Normalization: Normalization is the process of organizing data in a database so that it is efficient and reduces redundancy. This involves breaking down larger tables into smaller ones and establishing relationships between them.

Indexes: Indexes are used to improve the performance of queries on a database. They provide a way to quickly locate data based on specific criteria.

SQL: Structured Query Language (SQL) is a standard language used to manage and manipulate relational databases. It is used to create, modify, and query databases.

Transactions: Transactions are used to ensure that database operations are performed reliably and consistently. A transaction groups together a series of operations and ensures that they are either all completed successfully or all rolled back if an error occurs.

ACID: ACID is an acronym for Atomicity, Consistency, Isolation, and Durability. These are the properties that ensure that database transactions are reliable and consistent.

Relational databases are widely used in many industries and are essential for managing large amounts of data efficiently and effectively.

- [Understanding tables, columns, and rows](#)

In SQL Server, a database consists of one or more tables, which are made up of columns and rows. Here's a breakdown of each of these components:

Tables: A table is a collection of data that is organized into rows and columns. Tables are created using the CREATE TABLE statement, and they have a name that uniquely identifies them within the database. Tables are the primary objects that hold data in a database.

Columns: A column represents a single piece of data within a table. Each column has a name, a data type, and other attributes that define how the data is stored and used. Common data types include integers, strings, dates, and booleans. Columns are defined when a table is created, and they can be added, removed, or modified using the ALTER TABLE statement.

Rows: A row, also known as a record, represents a single instance of data within a table. Each row contains values for each column in the table, and the values must conform to the data types and constraints defined for the columns. Rows are inserted into a table using the INSERT statement, and they can be retrieved using the SELECT statement.

Together, tables, columns, and rows provide a powerful and flexible way to store and retrieve data in SQL Server. They form the basic building blocks for databases and allow for the creation of complex data structures that can be queried and manipulated using SQL.

- [Primary keys and foreign keys](#)

Primary keys and foreign keys are two key concepts in SQL Server and relational databases. Here's what you need to know about each:

Primary key: A primary key is a column or group of columns in a table that uniquely identifies each row in the table. The primary key constraint ensures that each value in the column or columns is unique and that the column or columns cannot contain null values. In SQL Server, primary keys are often created using the IDENTITY property, which generates a unique value for each row. Primary keys are important because they provide a way to link data between tables and ensure data integrity.

Foreign key: A foreign key is a column or group of columns in a table that refers to the primary key of another table. A foreign key constraint ensures that the values in the column or columns match the values in the primary key of the related table. Foreign keys are used to establish relationships

between tables and ensure referential integrity. In SQL Server, foreign keys are created using the FOREIGN KEY constraint.

Here's an example of how primary keys and foreign keys work together. Let's say you have two tables: Customers and Orders. The Customers table has a primary key of CustomerID, and the Orders table has a foreign key of CustomerID that refers to the Customers table. This means that each row in the Orders table is associated with a specific customer in the Customers table. When you insert a new row into the Orders table, you must specify a valid value for the CustomerID column, which ensures that the order is associated with an existing customer. If you try to insert a value that does not exist in the Customers table, you will receive an error. This ensures that the data is consistent and accurate.

Overall, primary keys and foreign keys are essential concepts in SQL Server and relational databases. They ensure that data is organized, linked, and maintained properly, and they help to ensure data integrity and consistency.

- [Normalization and denormalization](#)

Normalization and denormalization are two strategies for organizing data in SQL Server and other relational databases.

Normalization: Normalization is the process of organizing data in a database to reduce redundancy and improve data integrity. Normalization involves breaking down a large table into smaller tables and creating relationships between them. This helps to eliminate data duplication and ensures that data is consistent and accurate. There are several levels of normalization, each with its own set of rules and guidelines. The most commonly used levels are first normal form (1NF), second normal form (2NF), and third normal form (3NF).

Denormalization: Denormalization is the process of intentionally adding redundancy to a database to improve performance. Denormalization involves combining two or more tables into a single table, duplicating data across tables, or creating calculated columns that summarize data. This can improve performance by reducing the number of joins required to retrieve data and by reducing the amount of data that needs to be read from disk. However, denormalization can also lead to data inconsistency if not done carefully.

When deciding whether to normalize or denormalize a database, it's important to consider the trade-offs between data integrity and performance. Normalization can improve data integrity but may result in slower query performance, while denormalization can improve query performance but may make it harder to maintain data integrity. It's often a balancing act between these two goals, and the optimal solution will depend on the specific requirements of the application and the data being stored.

3. Basic SQL Syntax

SQL Server uses the SQL (Structured Query Language) syntax for creating, modifying, and querying databases. Here's a brief overview of some basic SQL syntax in SQL Server:

Creating a database: To create a new database, use the CREATE DATABASE statement, followed by the name of the database. For example:

```
CREATE DATABASE mydatabase;
```

Creating a table: To create a new table, use the CREATE TABLE statement, followed by the name of the table and a list of columns and their data types. For example:

```
CREATE TABLE mytable (  
  id INT PRIMARY KEY,  
  name VARCHAR(50),  
  age INT  
);
```

Inserting data: To insert data into a table, use the INSERT INTO statement, followed by the name of the table and the values to be inserted. For example:

```
INSERT INTO mytable (id, name, age) VALUES (1, 'John', 30);
```

Updating data: To update data in a table, use the UPDATE statement, followed by the name of the table and the new values to be set. For example:

```
UPDATE mytable SET name = 'Jane' WHERE id = 1;
```

Deleting data: To delete data from a table, use the DELETE FROM statement, followed by the name of the table and the conditions that specify which rows to delete. For example:

```
DELETE FROM mytable WHERE age > 40;
```

Selecting data: To retrieve data from a table, use the SELECT statement, followed by the columns to retrieve and the name of the table. For example:

```
SELECT id, name, age FROM mytable;
```

These are just a few examples of the basic SQL syntax used in SQL Server. SQL is a powerful language with many more features and capabilities, but these concepts are a good starting point for beginners.

- SELECT statement
- WHERE clause
- ORDER BY clause
- GROUP BY clause

The SELECT statement is one of the most important and commonly used SQL statements in SQL Server. It is used to retrieve data from one or more tables and can be customized to specify which columns to retrieve, filter the results based on conditions, sort the results, and more. Here's an overview of the basic syntax for the SELECT statement:

```
SELECT column1, column2, ... FROM table1
```

In this example, column1, column2, etc. are the names of the columns you want to retrieve, and table1 is the name of the table you want to retrieve data from. You can also use the wildcard * to retrieve all columns from a table:

```
SELECT * FROM table1
```

To filter the results based on conditions, use the WHERE clause:

```
SELECT column1, column2, ... FROM table1 WHERE condition
```

In this example, condition is a logical expression that evaluates to true or false, based on the values in the table. You can use operators like =, <>, <, >, <=, >=, LIKE, and BETWEEN to create conditions.

To sort the results, use the ORDER BY clause:

```
SELECT column1, column2, ... FROM table1 ORDER BY column1 ASC/DESC, column2 ASC/DESC, ...
```

In this example, ASC and DESC specify whether to sort the results in ascending or descending order. You can sort by one or more columns, in any order.

The **GROUP BY** clause is used in SQL Server to group rows that have the same values in one or more columns. This is often used in combination with aggregate functions, such as SUM, COUNT, AVG, MIN, or MAX, to calculate summary information for each group of rows.

Here's the basic syntax for using the GROUP BY clause in a SELECT statement:

```
SELECT column1, column2, ..., aggregate_function(columnN)
```

```
FROM table
```

```
WHERE condition
```

```
GROUP BY column1, column2, ...
```

In this example, column1, column2, etc. are the names of the columns you want to group by, and aggregate_function(columnN) is the aggregate function you want to use to calculate summary information for each group of rows. The WHERE clause is optional, but it can be used to filter the rows before grouping.

For example, let's say you have a table called "orders" with columns "order_id", "customer_id", "order_date", and "amount". You can use the GROUP BY clause to calculate the total amount of orders for each customer:

```
SELECT customer_id, SUM(amount) as total_amount
```

```
FROM orders
```

```
GROUP BY customer_id
```

In this example, the SUM function is used to calculate the total amount of orders for each customer, and the GROUP BY clause is used to group the rows by customer_id. The result will be a list of customer_id and their respective total_amount.

It's important to note that any column used in the SELECT statement must either be in the GROUP BY clause or used with an aggregate function. This is because when using the GROUP BY clause, SQL Server creates groups of rows based on the unique combination of values in the specified columns, and the aggregate function calculates the summary information for each group.

4. Data Manipulation Language (DML)

Data Manipulation Language (DML) is used to manipulate data stored in the database. In SQL Server, the primary DML statements are INSERT, UPDATE, and DELETE.

INSERT: The INSERT statement is used to add new rows to a table. Here's an example:

```
INSERT INTO table_name (column1, column2, column3, ...)
```

```
VALUES (value1, value2, value3, ...);
```

In this example, table_name is the name of the table you want to insert data into, and column1, column2, etc. are the names of the columns you want to insert data into. value1, value2, etc. are the values you want to insert into the corresponding columns.

UPDATE: The UPDATE statement is used to modify existing rows in a table. Here's an example:

```
UPDATE table_name
```

```
SET column1 = value1, column2 = value2, ...
```

```
WHERE condition;
```

In this example, table_name is the name of the table you want to update, and column1, column2, etc. are the names of the columns you want to update. value1, value2, etc. are the new values you want to set for the corresponding columns. condition is a logical expression that specifies which rows to update.

DELETE: The DELETE statement is used to remove rows from a table. Here's an example:

```
DELETE FROM table_name
```

```
WHERE condition;
```

In this example, table_name is the name of the table you want to delete rows from, and condition is a logical expression that specifies which rows to delete.

It's important to note that these DML statements can be combined with other SQL statements, such as SELECT, JOIN, and GROUP BY, to perform more complex data manipulations. Additionally, transactions can be used to ensure that DML statements are executed in an all-or-nothing fashion, meaning that if one statement fails, the entire transaction is rolled back.

5. Data Definition Language (DDL)

Data Definition Language (DDL) is used to define the structure of the database, including tables, columns, constraints, and indexes. In SQL Server, the primary DDL statements are CREATE, ALTER, and DROP.

CREATE: The CREATE statement is used to create new database objects, such as tables, views, and indexes. Here's an example of creating a new table:

```
CREATE TABLE table_name (  
    column1 datatype [constraint],  
    column2 datatype [constraint],  
    ...  
);
```

In this example, table_name is the name of the table you want to create, and column1, column2, etc. are the names and datatypes of the columns you want to create. Constraints, such as primary keys, foreign keys, and check constraints, can also be added to the columns.

ALTER: The ALTER statement is used to modify existing database objects, such as tables, views, and indexes. Here's an example of adding a new column to an existing table:

```
ALTER TABLE table_name  
ADD column_name datatype [constraint];
```

In this example, table_name is the name of the table you want to modify, and column_name and datatype are the names and datatypes of the new column you want to add. Constraints can also be added to the new column.

DROP: The DROP statement is used to remove database objects, such as tables, views, and indexes. Here's an example of dropping a table:

```
DROP TABLE table_name;
```

In this example, table_name is the name of the table you want to drop.

It's important to note that these DDL statements can have a significant impact on the database and should be used with caution. Additionally, transactions can be used to ensure that DDL statements are executed in an all-or-nothing fashion, meaning that if one statement fails, the entire transaction is rolled back.

6. Querying Multiple Tables

Querying multiple tables in SQL Server is accomplished through JOIN operations. JOIN operations allow you to combine rows from two or more tables based on a related column between them.

There are several types of JOIN operations in SQL Server, including:

INNER JOIN: Returns only the rows that have matching values in both tables. Here's an example:

```
SELECT column1, column2, ...
```

```
FROM table1
```

```
INNER JOIN table2
```

```
ON table1.column = table2.column;
```

In this example, table1 and table2 are the names of the two tables you want to join. column1, column2, etc. are the columns you want to select from the joined tables. table1.column and table2.column are the columns that the two tables have in common.

LEFT JOIN: Returns all the rows from the left table and the matching rows from the right table. If there is no matching row in the right table, NULL values are returned. Here's an example:

```
SELECT column1, column2, ...
```

```
FROM table1
```

```
LEFT JOIN table2
```

```
ON table1.column = table2.column;
```

In this example, table1 and table2 are the names of the two tables you want to join. column1, column2, etc. are the columns you want to select from the joined tables. table1.column and table2.column are the columns that the two tables have in common.

RIGHT JOIN: Returns all the rows from the right table and the matching rows from the left table. If there is no matching row in the left table, NULL values are returned. Here's an example:

```
SELECT column1, column2, ...
```

```
FROM table1
```

```
RIGHT JOIN table2
```

```
ON table1.column = table2.column;
```

In this example, table1 and table2 are the names of the two tables you want to join. column1, column2, etc. are the columns you want to select from the joined tables. table1.column and table2.column are the columns that the two tables have in common.

FULL OUTER JOIN: Returns all the rows from both tables, including the unmatched rows. If there is no matching row in one of the tables, NULL values are returned. Here's an example:

```
SELECT column1, column2, ...
```

```
FROM table1
```

```
FULL OUTER JOIN table2
```

```
ON table1.column = table2.column;
```

In this example, table1 and table2 are the names of the two tables you want to join. column1, column2, etc. are the columns you want to select from the joined tables. table1.column and table2.column are the columns that the two tables have in common.

It's important to note that when joining tables, it's recommended to use aliases to make the query more readable and to avoid column name conflicts. Additionally, it's important to use appropriate indexes on the join columns to improve query performance.

7. Aggregate Functions

Aggregate functions in SQL Server are used to perform calculations on a set of values and return a single value. These functions operate on a group of rows and return a single result for the entire group. Here are some of the commonly used aggregate functions in SQL Server:

COUNT(): This function returns the number of rows that match the specified condition. If no condition is specified, it returns the total number of rows in the table. Here's an example:

```
SELECT COUNT(column_name)
FROM table_name;
```

SUM(): This function returns the sum of the values in the specified column. Here's an example:

```
SELECT SUM(column_name)
FROM table_name;
```

AVG(): This function returns the average of the values in the specified column. Here's an example:

```
SELECT AVG(column_name)
FROM table_name;
```

MAX(): This function returns the maximum value in the specified column. Here's an example:

```
SELECT MAX(column_name)
FROM table_name;
```

MIN(): This function returns the minimum value in the specified column. Here's an example:

```
SELECT MIN(column_name)
FROM table_name;
```

Aggregate functions can also be used with the GROUP BY clause to group the results by one or more columns. For example:

```
SELECT column_name, COUNT(*)
FROM table_name
GROUP BY column_name;
```

This query will return the number of rows for each unique value in column_name.

8. Subqueries

In SQL Server, a subquery is a query that is nested inside another query. A subquery can be used to retrieve data that will be used in the main query as a condition to filter the results or to perform calculations.

Here are some examples of how subqueries can be used in SQL Server:

Subquery as a condition in a WHERE clause:

```
SELECT column_name1, column_name2
FROM table_name
WHERE column_name1 IN (SELECT column_name1 FROM table_name2 WHERE column_name2 = 'value');
```

This query will return the column_name1 and column_name2 values from table_name where the column_name1 values are also present in the result of the subquery, which returns all column_name1 values from table_name2 where column_name2 is equal to 'value'.

Subquery as a derived table in a FROM clause:

```
SELECT derived_table.column_name1, SUM(derived_table.column_name2) AS total
FROM (SELECT column_name1, column_name2 FROM table_name) AS derived_table
GROUP BY derived_table.column_name1;
```

This query uses a subquery as a derived table to select column_name1 and column_name2 from table_name. The derived table is then used to perform a GROUP BY and calculate the SUM of column_name2 for each unique value of column_name1.

Subquery as a scalar expression:

```
SELECT column_name1, column_name2, (SELECT MAX(column_name3) FROM table_name2 WHERE table_name2.column_name1 = table_name.column_name1) AS max_value
FROM table_name;
```

This query uses a subquery as a scalar expression to return the maximum value of column_name3 from table_name2 for each unique value of column_name1 in table_name.

Subqueries can be used in various other ways in SQL Server to achieve more complex querying requirements. It is important to use subqueries judiciously as they can impact query performance if not used properly.

9. Indexing

In SQL Server, indexing is a technique used to improve the performance of queries by creating indexes on one or more columns in a table. An index is a data structure that contains a copy of the data from a table or view, organized in a way that makes it more efficient to search and retrieve the data.

Here are some key points about indexing in SQL Server:

Types of indexes: SQL Server supports several types of indexes, including clustered indexes, nonclustered indexes, and full-text indexes. Clustered indexes define the physical order of the data in a table, while nonclustered indexes provide a separate structure for faster searching of data. Full-text indexes are used to perform fast text-based searches on large columns of text data.

Creating indexes: Indexes can be created using the CREATE INDEX statement. It is important to carefully choose the columns to be indexed as too many indexes can adversely affect the performance of write operations on the table.

Maintaining indexes: Indexes need to be maintained to ensure they remain effective over time. This involves rebuilding or reorganizing indexes periodically to optimize their performance.

Using indexes: Indexes are used automatically by SQL Server when executing queries. The query optimizer evaluates the query and determines the most efficient way to retrieve the required data using the available indexes.

Monitoring index performance: SQL Server provides tools for monitoring index usage and performance, including the Database Engine Tuning Advisor and the Dynamic Management Views.

Properly indexing tables can significantly improve the performance of queries in SQL Server. However, it is important to carefully choose the columns to be indexed, and to monitor and maintain the indexes over time to ensure they remain effective.

10. Stored Procedures

In SQL Server, a stored procedure is a precompiled collection of SQL statements and procedural logic that is stored in the database as a named object. Stored procedures can be executed by users or applications to perform specific tasks or operations on the database.

Here are some key features of stored procedures in SQL Server:

Creating stored procedures: Stored procedures are created using the CREATE PROCEDURE statement. A stored procedure can have input and output parameters, which allow the procedure to accept and return data to the caller.

Advantages of stored procedures: Stored procedures provide several advantages over ad-hoc SQL statements, including improved performance, security, and code reusability. Stored procedures can also be used to encapsulate complex business logic or data processing operations.

Executing stored procedures: Stored procedures are executed using the EXECUTE or EXEC statement. The procedure can be called with input parameters, which are passed as arguments to the procedure, and output parameters, which are used to return data to the caller.

Modifying and dropping stored procedures: Stored procedures can be modified using the ALTER PROCEDURE statement, and dropped using the DROP PROCEDURE statement.

Debugging stored procedures: SQL Server provides tools for debugging stored procedures, including the Transact-SQL Debugger and SQL Server Management Studio's Debug menu.

Stored procedures are a powerful tool for managing and manipulating data in SQL Server. They can improve performance, security, and code reusability, and are widely used in enterprise-level database applications.

11. Views

In SQL Server, a view is a virtual table that is based on the result of a SELECT statement. Views provide a way to simplify complex queries and to provide a consistent interface to data stored in multiple tables.

Here are some key features of views in SQL Server:

Creating views: Views are created using the CREATE VIEW statement. Views can include joins, subqueries, and other SQL constructs.

Modifying and dropping views: Views can be modified using the ALTER VIEW statement, and dropped using the DROP VIEW statement.

Advantages of views: Views provide several advantages over direct access to tables, including simplified querying, improved security, and abstraction of database schema changes.

Using views: Views can be used in the same way as tables in SQL queries, and can be used as the source for other views or for stored procedures.

Limitations of views: Views have some limitations in SQL Server, including restrictions on the types of queries that can be used to create them, and performance issues when using complex queries or large amounts of data.

Views are a useful tool for simplifying complex queries and abstracting database schema changes. They can provide a consistent interface to data stored in multiple tables, and can be used to enforce security restrictions on data access. However, it is important to be aware of the limitations of views and to use them appropriately in SQL Server.

12. Transactions and Locking

In SQL Server, a transaction is a sequence of one or more database operations that are treated as a single unit of work. Transactions are used to ensure the consistency and integrity of data in the database, and to maintain data accuracy and correctness.

Here are some key features of transactions and locking in SQL Server:

Transaction states: Transactions in SQL Server have four states: active, committed, rolled back, and partially committed.

ACID properties: Transactions in SQL Server follow the ACID properties: Atomicity, Consistency, Isolation, and Durability.

Transaction isolation levels: SQL Server provides several transaction isolation levels, which control the level of locking and concurrency in the database. These include Read Uncommitted, Read Committed, Repeatable Read, and Serializable.

Locking: SQL Server uses locking to prevent multiple users from modifying the same data at the same time. Locks can be acquired at various levels, such as row-level, page-level, or table-level.

Deadlocks: Deadlocks occur when two or more transactions are waiting for each other to release locks. SQL Server provides a mechanism for detecting and resolving deadlocks.

Transactions and locking are essential components of database management in SQL Server. They ensure the consistency and integrity of data, and provide a mechanism for controlling concurrency and preventing data corruption. Understanding transactions and locking is important for designing and implementing robust and scalable database applications in SQL Server.

13. Security

Security is a critical aspect of database management in SQL Server. It involves protecting the confidentiality, integrity, and availability of data in the database, and controlling access to the database by users and applications.

Here are some key features of security in SQL Server:

Authentication: SQL Server supports several authentication modes, including Windows authentication, SQL Server authentication, and Azure Active Directory authentication. Authentication is used to verify the identity of users who access the database.

Authorization: SQL Server provides a robust authorization system that controls access to database objects and data. This includes granting and revoking permissions on objects, and creating database roles to group users with similar permissions.

Auditing: SQL Server provides auditing features that allow database administrators to track and monitor database activity, including user logins, database modifications, and security-related events.

Encryption: SQL Server supports several encryption options, including Transparent Data Encryption (TDE), Cell-level Encryption, and Always Encrypted. Encryption is used to protect sensitive data in the database from unauthorized access.

Compliance: SQL Server supports several compliance standards, including the General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry Data Security Standard (PCI DSS). Compliance features help ensure that the database is secure and meets regulatory requirements.

Security is a critical component of database management in SQL Server. It is important to design and implement robust security measures to protect sensitive data from unauthorized access, and to ensure that the database meets regulatory requirements. Understanding the security features and best practices in SQL Server is essential for developing and deploying secure and compliant database applications.