

SQL Interview Questions with Answers

<http://www.sqlauthority.com>

What is RDBMS?

Relational Data Base Management Systems (RDBMS) are database management systems that maintain data records and indices in tables. Relationships may be created and maintained across and among the data and tables. In a relational database, relationships between data items are expressed by means of tables. Interdependencies among these tables are expressed by data values rather than by pointers. This allows a high degree of data independence. An RDBMS has the capability to recombine the data items from different files, providing powerful tools for data usage.

What is normalization?

Database normalization is a data design and organization process applied to data structures based on rules that help build relational databases. In relational database design, the process of organizing data to minimize redundancy. Normalization usually involves dividing a database into two or more tables and defining relationships between the tables. The objective is to isolate data so that additions, deletions, and modifications of a field can be made in just one table and then propagated through the rest of the database via the defined relationships.

What are different normalization forms?

1NF: Eliminate Repeating Groups

Make a separate table for each set of related attributes, and give each table a primary key. Each field contains at most one value from its attribute domain.

2NF: Eliminate Redundant Data

If an attribute depends on only part of a multi-valued key, remove it to a separate table.

3NF: Eliminate Columns Not Dependent On Key

If attributes do not contribute to a description of the key, remove them to a separate table. All attributes must be directly dependent on the primary key

BCNF: Boyce-Codd Normal Form

If there are non-trivial dependencies between candidate key attributes, separate them out into distinct tables.

4NF: Isolate Independent Multiple Relationships

No table may contain two or more 1:n or n:m relationships that are not directly related.

5NF: Isolate Semantically Related Multiple Relationships

There may be practical constraints on information that justify separating logically related many-to-many relationships.

ONF: Optimal Normal Form

A model limited to only simple (elemental) facts, as expressed in Object Role Model notation.

DKNF: Domain-Key Normal Form

A model free from all modification anomalies.

Remember, these normalization guidelines are cumulative. For a database to be in 3NF, it must first fulfill all the criteria of a 2NF and 1NF database.

What is Stored Procedure?

A stored procedure is a named group of SQL statements that have been previously created and stored in the server database. Stored procedures accept input parameters so that a single procedure can be used over the network by several clients using different input data. And when the procedure is modified, all clients automatically get the new version. Stored procedures reduce network traffic and improve performance. Stored procedures can be used to help ensure the integrity of the database. e.g. sp_helpdb, sp_renamedb, sp_depends etc.

What is Trigger?

A trigger is a SQL procedure that initiates an action when an event (INSERT, DELETE or UPDATE) occurs. Triggers are stored in and managed by the DBMS. Triggers are used to maintain the referential

©<http://www.sqlauthority.com>

integrity of data by changing the data in a systematic fashion. A trigger cannot be called or executed; the DBMS automatically fires the trigger as a result of a data modification to the associated table. Triggers can be viewed as similar to stored procedures in that both consist of procedural logic that is stored at the database level. Stored procedures, however, are not event-driven and are not attached to a specific table as triggers are. Stored procedures are explicitly executed by invoking a CALL to the procedure while triggers are implicitly executed. In addition, triggers can also execute stored procedures.

Nested Trigger: A trigger can also contain INSERT, UPDATE and DELETE logic within itself, so when the trigger is fired because of data modification it can also cause another data modification, thereby firing another trigger. A trigger that contains data modification logic within itself is called a nested trigger.

What is View?

A simple view can be thought of as a subset of a table. It can be used for retrieving data, as well as updating or deleting rows. Rows updated or deleted in the view are updated or deleted in the table the view was created with. It should also be noted that as data in the original table changes, so does data in the view, as views are the way to look at part of the original table. The results of using a view are not permanently stored in the database. The data accessed through a view is actually constructed using standard T-SQL select command and can come from one to many different base tables or even other views.

What is Index?

An index is a physical structure containing pointers to the data. Indices are created in an existing table to locate rows more quickly and efficiently. It is possible to create an index on one or more columns of a table, and each index is given a name. The users cannot see the indexes, they are just used to speed up queries. Effective indexes are one of the best ways to improve performance in a database application. A table scan happens when there is no index available to help a query. In a table scan SQL Server examines every row in the table to satisfy the query results. Table scans are sometimes unavoidable, but on large tables, scans have a terrific impact on performance.

Clustered indexes define the physical sorting of a database table's rows in the storage media. For this reason, each database table may have only one clustered index.

Non-clustered indexes are created outside of the database table and contain a sorted list of references to the table itself.

What is the difference between clustered and a non-clustered index?

A *clustered index* is a special type of index that reorders the way records in the table are physically stored. Therefore table can have only one clustered index. The leaf nodes of a clustered index contain the data pages.

A *nonclustered index* is a special type of index in which the logical order of the index does not match the physical stored order of the rows on disk. The leaf node of a nonclustered index does not consist of the data pages. Instead, the leaf nodes contain index rows.

What are the different index configurations a table can have?

A table can have one of the following index configurations:

- No indexes
- A clustered index
- A clustered index and many nonclustered indexes
- A nonclustered index
- Many nonclustered indexes

What is cursors?

Cursor is a database object used by applications to manipulate data in a set on a row-by-row basis, instead of the typical SQL commands that operate on all the rows in the set at one time.

©<http://www.sqlauthority.com>

©<http://www.sqlauthority.com>

In order to work with a cursor we need to perform some steps in the following order:

- Declare cursor
- Open cursor
- Fetch row from the cursor
- Process fetched row
- Close cursor
- Deallocate cursor

What is the use of DBCC commands?

DBCC stands for database consistency checker. We use these commands to check the consistency of the databases, i.e., maintenance, validation task and status checks.

E.g. DBCC CHECKDB - Ensures that tables in the db and the indexes are correctly linked.

DBCC CHECKALLOC - To check that all pages in a db are correctly allocated.

DBCC CHECKFILEGROUP - Checks all tables file group for any damage.

What is a Linked Server?

Linked Servers is a concept in SQL Server by which we can add other SQL Server to a Group and query both the SQL Server dbs using T-SQL Statements. With a linked server, you can create very clean, easy to follow, SQL statements that allow remote data to be retrieved, joined and combined with local data.

Storped Procedure `sp_addlinkedserver`, `sp_addlinkedsrvlogin` will be used add new Linked Server.

What is Collation?

Collation refers to a set of rules that determine how data is sorted and compared. Character data is sorted using rules that define the correct character sequence, with options for specifying case-sensitivity, accent marks, kana character types and character width.

What are different type of Collation Sensitivity?

Case sensitivity

A and a, B and b, etc.

Accent sensitivity

a and á, o and ó, etc.

Kana Sensitivity

When Japanese kana characters Hiragana and Katakana are treated differently, it is called Kana sensitive.

Width sensitivity

When a single-byte character (half-width) and the same character when represented as a double-byte character (full-width) are treated differently then it is width sensitive.

What's the difference between a primary key and a unique key?

Both primary key and unique enforce uniqueness of the column on which they are defined. But by default primary key creates a clustered index on the column, where are unique creates a nonclustered index by default. Another major difference is that, primary key doesn't allow NULLs, but unique key allows one NULL only.

How to implement one-to-one, one-to-many and many-to-many relationships while designing tables?

One-to-One relationship can be implemented as a single table and rarely as two tables with primary and foreign key relationships.

One-to-Many relationships are implemented by splitting the data into two tables with primary key and foreign key relationships.

Many-to-Many relationships are implemented using a junction table with the keys from both the tables forming the composite primary key of the junction table.

What is a NOLOCK?

©<http://www.sqlauthority.com>

©<http://www.sqlauthority.com>

Using the NOLOCK query optimiser hint is generally considered good practice in order to improve concurrency on a busy system. When the NOLOCK hint is included in a SELECT statement, no locks are taken when data is read. The result is a Dirty Read, which means that another process could be updating the data at the exact time you are reading it. There are no guarantees that your query will retrieve the most recent data. The advantage to performance is that your reading of data will not block updates from taking place, and updates will not block your reading of data. SELECT statements take Shared (Read) locks. This means that multiple SELECT statements are allowed simultaneous access, but other processes are blocked from modifying the data. The updates will queue until all the reads have completed, and reads requested after the update will wait for the updates to complete. The result to your system is delay(blocking).

What is difference between DELETE & TRUNCATE commands?

Delete command removes the rows from a table based on the condition that we provide with a WHERE clause. Truncate will actually remove all the rows from a table and there will be no data in the table after we run the truncate command.

TRUNCATE

TRUNCATE is faster and uses fewer system and transaction log resources than DELETE. TRUNCATE removes the data by deallocating the data pages used to store the table's data, and only the page deallocations are recorded in the transaction log. TRUNCATE removes all rows from a table, but the table structure and its columns, constraints, indexes and so on remain. The counter used by an identity for new rows is reset to the seed for the column. You cannot use TRUNCATE TABLE on a table referenced by a FOREIGN KEY constraint. Because TRUNCATE TABLE is not logged, it cannot activate a trigger. TRUNCATE can not be Rolled back. TRUNCATE is DDL Command. TRUNCATE Resets identity of the table.

DELETE

DELETE removes rows one at a time and records an entry in the transaction log for each deleted row. If you want to retain the identity counter, use DELETE instead. If you want to remove table definition and its data, use the DROP TABLE statement. DELETE Can be used with or without a WHERE clause. DELETE Activates Triggers. DELETE Can be Rolled back. DELETE is DML Command. DELETE does not reset identity of the table.

Difference between Function and Stored Procedure?

UDF can be used in the SQL statements anywhere in the WHERE/HAVING/SELECT section where as Stored procedures cannot be. UDFs that return tables can be treated as another rowset. This can be used in JOINS with other tables. Inline UDF's can be thought of as views that take parameters and can be used in JOINS and other Rowset operations.

When is the use of UPDATE_STATISTICS command?

This command is basically used when a large processing of data has occurred. If a large amount of deletions any modification or Bulk Copy into the tables has occurred, it has to update the indexes to take these changes into account. UPDATE_STATISTICS updates the indexes on these tables accordingly.

What types of Joins are possible with Sql Server?

Joins are used in queries to explain how different tables are related. Joins also let you select data from a table depending upon data from another table. Types of joins: INNER JOINS, OUTER JOINS, CROSS JOINS. OUTER JOINS are further classified as LEFT OUTER JOINS, RIGHT OUTER JOINS and FULL OUTER JOINS.

©<http://www.sqlauthority.com>

What is the difference between a HAVING CLAUSE and a WHERE CLAUSE?

Specifies a search condition for a group or an aggregate. HAVING can be used only with the SELECT statement. HAVING is typically used in a GROUP BY clause. When GROUP BY is not used, HAVING behaves like a WHERE clause. Having Clause is basically used only with the GROUP BY function in a query. WHERE Clause is applied to each row before they are part of the GROUP BY function in a query.

What is sub-query? Explain properties of sub-query.

Sub-queries are often referred to as sub-selects, as they allow a SELECT statement to be executed arbitrarily within the body of another SQL statement. A sub-query is executed by enclosing it in a set of parentheses. Sub-queries are generally used to return a single row as an atomic value, though they may be used to compare values against multiple rows with the IN keyword.

A subquery is a SELECT statement that is nested within another T-SQL statement. A subquery SELECT statement if executed independently of the T-SQL statement, in which it is nested, will return a result set. Meaning a subquery SELECT statement can standalone and is not depended on the statement in which it is nested. A subquery SELECT statement can return any number of values, and can be found in, the column list of a SELECT statement, a FROM, GROUP BY, HAVING, and/or ORDER BY clauses of a T-SQL statement. A Subquery can also be used as a parameter to a function call. Basically a subquery can be used anywhere an expression can be used.

Properties of Sub-Query

A subquery must be enclosed in the parenthesis.

A subquery must be put in the right hand of the comparison operator, and

A subquery cannot contain a ORDER-BY clause.

A query can contain more than one sub-queries.

What are types of sub-queries?

Single-row subquery, where the subquery returns only one row.

Multiple-row subquery, where the subquery returns multiple rows, and

Multiple column subquery, where the subquery returns multiple columns.

What is SQL Profiler?

SQL Profiler is a graphical tool that allows system administrators to monitor events in an instance of Microsoft SQL Server. You can capture and save data about each event to a file or SQL Server table to analyze later. For example, you can monitor a production environment to see which stored procedures are hampering performance by executing too slowly.

Use SQL Profiler to monitor only the events in which you are interested. If traces are becoming too large, you can filter them based on the information you want, so that only a subset of the event data is collected. Monitoring too many events adds overhead to the server and the monitoring process and can cause the trace file or trace table to grow very large, especially when the monitoring process takes place over a long period of time.

What is User Defined Functions?

User-Defined Functions allow to define its own T-SQL functions that can accept 0 or more parameters and return a single scalar data value or a table data type.

What kind of User-Defined Functions can be created?

There are three types of User-Defined functions in SQL Server 2000 and they are Scalar, Inline Table-Valued and Multi-statement Table-valued.

Scalar User-Defined Function

A Scalar user-defined function returns one of the scalar data types. Text, ntext, image and timestamp data types are not supported. These are the type of user-defined functions that most developers are used to in other programming languages. You pass in 0 to many parameters and you get a return value.

Inline Table-Value User-Defined Function

An Inline Table-Value user-defined function returns a table data type and is an exceptional alternative to a view as the user-defined function can pass parameters into a T-SQL select command and in essence provide us with a parameterized, non-updateable view of the underlying tables.

Multi-statement Table-Value User-Defined Function

A Multi-Statement Table-Value user-defined function returns a table and is also an exceptional alternative to a view as the function can support multiple T-SQL statements to build the final result where the view is limited to a single SELECT statement. Also, the ability to pass parameters into a T-SQL select command or a group of them gives us the capability to in essence create a parameterized, non-updateable view of the data in the underlying tables. Within the create function command you must define the table structure that is being returned. After creating this type of user-defined function, it can be used in the FROM clause of a T-SQL command unlike the behavior found when using a stored procedure which can also return record sets.

Which TCP/IP port does SQL Server run on? How can it be changed?

SQL Server runs on port 1433. It can be changed from the Network Utility TCP/IP properties -> Port number.both on client and the server.

What are the authentication modes in SQL Server? How can it be changed?

Windows mode and mixed mode (SQL & Windows).

To change authentication mode in SQL Server click Start, Programs, Microsoft SQL Server and click SQL Enterprise Manager to run SQL Enterprise Manager from the Microsoft SQL Server program group. Select the server then from the Tools menu select SQL Server Configuration Properties, and choose the Security page.

Where are SQL server users names and passwords are stored in sql server?

They get stored in master db in the sysxlogins table.

Which command using Query Analyzer will give you the version of SQL server and operating system?

```
SELECT SERVERPROPERTY('productversion'), SERVERPROPERTY ('productlevel'), SERVERPROPERTY ('edition')
```

What is SQL server agent?

SQL Server agent plays an important role in the day-to-day tasks of a database administrator (DBA). It is often overlooked as one of the main tools for SQL Server management. Its purpose is to ease the implementation of tasks for the DBA, with its full-function scheduling engine, which allows you to schedule your own jobs and scripts.

Can a stored procedure call itself or recursive stored procedure? How many level SP nesting possible?

Yes. Because Transact-SQL supports recursion, you can write stored procedures that call themselves. Recursion can be defined as a method of problem solving wherein the solution is arrived at by repetitively applying it to subsets of the problem. A common application of recursive logic is to perform numeric computations that lend themselves to repetitive evaluation by the same processing steps. Stored procedures are nested when one stored procedure calls another or executes managed code by referencing a CLR routine, type, or aggregate. You can nest stored procedures and managed code references up to 32 levels.

What is @@ERROR?

The @@ERROR automatic variable returns the error code of the last Transact-SQL statement. If there was no error, @@ERROR returns zero. Because @@ERROR is reset after each Transact-SQL statement, it must be saved to a variable if it is needed to process it further after checking it.

©<http://www.sqlauthority.com>

What is Raiseerror?

Stored procedures report errors to client applications via the RAISERROR command. RAISERROR doesn't change the flow of a procedure; it merely displays an error message, sets the @@ERROR automatic variable, and optionally writes the message to the SQL Server error log and the NT application event log.

What is log shipping?

Log shipping is the process of automating the backup of database and transaction log files on a production SQL server, and then restoring them onto a standby server. Enterprise Editions only supports log shipping. In log shipping the transactional log file from one server is automatically updated into the backup database on the other server. If one server fails, the other server will have the same db can be used this as the Disaster Recovery plan. The key feature of log shipping is that it will automatically backup transaction logs throughout the day and automatically restore them on the standby server at defined interval.

What is the difference between a local and a global variable?

A *local temporary* table exists only for the duration of a connection or, if defined inside a compound statement, for the duration of the compound statement.

A *global temporary* table remains in the database permanently, but the rows exist only within a given connection. When connection are closed, the data in the global temporary table disappears. However, the table definition remains with the database for access when database is opened next time.

What command do we use to rename a db?

`sp_renamedb 'oldname', 'newname'`

If someone is using db it will not accept `sp_renamedb`. In that case first bring db to single user using `sp_dboptions`. Use `sp_renamedb` to rename database. Use `sp_dboptions` to bring database to multi user mode.

What is sp_configure commands and set commands?

Use `sp_configure` to display or change server-level settings. To change database-level settings, use `ALTER DATABASE`. To change settings that affect only the current user session, use the `SET` statement.

What are the different types of replication? Explain.

The SQL Server 2000-supported replication types are as follows:

- Transactional
- Snapshot
- Merge

Snapshot replication distributes data exactly as it appears at a specific moment in time and does not monitor for updates to the data. Snapshot replication is best used as a method for replicating data that changes infrequently or where the most up-to-date values (low latency) are not a requirement. When synchronization occurs, the entire snapshot is generated and sent to Subscribers.

Transactional replication, an initial snapshot of data is applied at Subscribers, and then when data modifications are made at the Publisher, the individual transactions are captured and propagated to Subscribers.

Merge replication is the process of distributing data from Publisher to Subscribers, allowing the Publisher and Subscribers to make updates while connected or disconnected, and then merging the updates between sites when they are connected.

What are the OS services that the SQL Server installation adds?

MS SQL SERVER SERVICE, SQL AGENT SERVICE, DTC (Distribution transac co-ordinator)

What are three SQL keywords used to change or set someone's permissions?

©<http://www.sqlauthority.com>

What does it mean to have quoted_identifier on? What are the implications of having it off?

When SET QUOTED_IDENTIFIER is ON, identifiers can be delimited by double quotation marks, and literals must be delimited by single quotation marks. When SET QUOTED_IDENTIFIER is OFF, identifiers cannot be quoted and must follow all Transact-SQL rules for identifiers.

What is the STUFF function and how does it differ from the REPLACE function?

STUFF function to overwrite existing characters. Using this syntax, STUFF(string_expression, start, length, replacement_characters), string_expression is the string that will have characters substituted, start is the starting position, length is the number of characters in the string that are substituted, and replacement_characters are the new characters interjected into the string.

REPLACE function to replace existing characters of all occurrence. Using this syntax REPLACE(string_expression, search_string, replacement_string), where every incidence of search_string found in the string_expression will be replaced with replacement_string.

Using query analyzer, name 3 ways to get an accurate count of the number of records in a table?

```
SELECT * FROM table1
```

```
SELECT COUNT(*) FROM table1
```

```
SELECT rows FROM sysindexes WHERE id = OBJECT_ID(table1) AND indid < 2
```

How to rebuild Master Database?

Shutdown Microsoft SQL Server 2000, and then run Rebuildm.exe. This is located in the Program Files\Microsoft SQL Server\80\Tools\Binn directory.

In the Rebuild Master dialog box, click Browse.

In the Browse for Folder dialog box, select the \Data folder on the SQL Server 2000 compact disc or in the shared network directory from which SQL Server 2000 was installed, and then click OK.

Click Settings. In the Collation Settings dialog box, verify or change settings used for the master database and all other databases.

Initially, the default collation settings are shown, but these may not match the collation selected during setup. You can select the same settings used during setup or select new collation settings. When done, click OK.

In the Rebuild Master dialog box, click Rebuild to start the process.

The Rebuild Master utility reinstalls the master database.

To continue, you may need to stop a server that is running.

Source: [http://msdn2.microsoft.com/en-us/library/aa197950\(SQL.80\).aspx](http://msdn2.microsoft.com/en-us/library/aa197950(SQL.80).aspx)

What is the basic functions for master, msdb, model, tempdb databases?

The **Master** database holds information for all databases located on the SQL Server instance and is the glue that holds the engine together. Because SQL Server cannot start without a functioning master database, you must administer this database with care.

The **msdb** database stores information regarding database backups, SQL Agent information, DTS packages, SQL Server jobs, and some replication information such as for log shipping.

The **tempdb** holds temporary objects such as global and local temporary tables and stored procedures. The **model** is essentially a template database used in the creation of any new user database created in the instance.

What are primary keys and foreign keys?

Primary keys are the unique identifiers for each row. They must contain unique values and cannot be null. Due to their importance in relational databases, Primary keys are the most fundamental of all keys and constraints. A table can have only one Primary key.

Foreign keys are both a method of ensuring data integrity and a manifestation of the relationship between tables.

What is data integrity? Explain constraints?

Data integrity is an important feature in SQL Server. When used properly, it ensures that data is

accurate, correct, and valid. It also acts as a trap for otherwise undetectable bugs within applications.

A **PRIMARY KEY** constraint is a unique identifier for a row within a database table. Every table should have a primary key constraint to uniquely identify each row and only one primary key constraint can be created for each table. The primary key constraints are used to enforce entity integrity.

A **UNIQUE** constraint enforces the uniqueness of the values in a set of columns, so no duplicate values are entered. The unique key constraints are used to enforce entity integrity as the primary key constraints.

A **FOREIGN KEY** constraint prevents any actions that would destroy links between tables with the corresponding data values. A foreign key in one table points to a primary key in another table. Foreign keys prevent actions that would leave rows with foreign key values when there are no primary keys with that value. The foreign key constraints are used to enforce referential integrity.

A **CHECK** constraint is used to limit the values that can be placed in a column. The check constraints are used to enforce domain integrity.

A **NOT NULL** constraint enforces that the column will not accept null values. The not null constraints are used to enforce domain integrity, as the check constraints.

What are the properties of the Relational tables?

Relational tables have six properties:

- Values are atomic.
- Column values are of the same kind.
- Each row is unique.
- The sequence of columns is insignificant.
- The sequence of rows is insignificant.
- Each column must have a unique name.

What is De-normalization?

De-normalization is the process of attempting to optimize the performance of a database by adding redundant data. It is sometimes necessary because current DBMSs implement the relational model poorly. A true relational DBMS would allow for a fully normalized database at the logical level, while providing physical storage of data that is tuned for high performance. De-normalization is a technique to move from higher to lower normal forms of database modeling in order to speed up database access.

How to get @@error and @@rowcount at the same time?

If @@Rowcount is checked after Error checking statement then it will have 0 as the value of @@Recordcount as it would have been reset.

And if @@Recordcount is checked before the error-checking statement then @@Error would get reset. To get @@error and @@rowcount at the same time do both in same statement and store them in local variable. `SELECT @RC = @@ROWCOUNT, @ER = @@ERROR`

What is Identity?

Identity (or AutoNumber) is a column that automatically generates numeric values. A start and increment value can be set, but most DBAs leave these at 1. A GUID column also generates numbers, the value of this cannot be controlled. Identity/GUID columns do not need to be indexed.

What is a Scheduled Jobs or What is a Scheduled Tasks?

Scheduled tasks let user automate processes that run on regular or predictable cycles. User can schedule administrative tasks, such as cube processing, to run during times of slow business activity. User can also determine the order in which tasks run by creating job steps within a SQL Server Agent job. E.g. Back up database, Update Stats of Tables. Job steps give user control over flow of execution.

©<http://www.sqlauthority.com>

If one job fails, user can configure SQL Server Agent to continue to run the remaining tasks or to stop execution.

What is a table called, if it does not have neither Cluster nor Non-cluster Index? What is it used for?

Unindexed table or **Heap**. Microsoft Press Books and Book On Line (BOL) refers it as Heap.

A heap is a table that does not have a clustered index and, therefore, the pages are not linked by pointers. The IAM pages are the only structures that link the pages in a table together.

Unindexed tables are good for fast storing of data. Many times it is better to drop all indexes from table and then do bulk of inserts and to restore those indexes after that.

What is BCP? When does it used?

BulkCopy is a tool used to copy huge amount of data from tables and views. BCP does not copy the structures same as source to destination.

How do you load large data to the SQL server database?

BulkCopy is a tool used to copy huge amount of data from tables. BULK INSERT command helps to Imports a data file into a database table or view in a user-specified format.

Can we rewrite subqueries into simple select statements or with joins?

Subqueries can often be re-written to use a standard outer join, resulting in faster performance. As we may know, an outer join uses the plus sign (+) operator to tell the database to return all non-matching rows with NULL values. Hence we combine the outer join with a NULL test in the WHERE clause to reproduce the result set without using a sub-query.

Can SQL Servers linked to other servers like Oracle?

SQL Server can be linked to any server provided it has OLE-DB provider from Microsoft to allow a link. E.g. Oracle has a OLE-DB provider for oracle that Microsoft provides to add it as linked server to SQL Server group.

How to know which index a table is using?

```
SELECT table_name,index_name FROM user_constraints
```

How to copy the tables, schema and views from one SQL server to another?

Microsoft SQL Server 2000 Data Transformation Services (DTS) is a set of graphical tools and programmable objects that lets user extract, transform, and consolidate data from disparate sources into single or multiple destinations.

What is Self Join?

This is a particular case when one table joins to itself, with one or two aliases to avoid confusion. A self join can be of any type, as long as the joined tables are the same. A self join is rather unique in that it involves a relationship with only one table. The common example is when company have a hierarchal reporting structure whereby one member of staff reports to another.

What is Cross Join?

A cross join that does not have a WHERE clause produces the Cartesian product of the tables involved in the join. The size of a Cartesian product result set is the number of rows in the first table multiplied by the number of rows in the second table. The common example is when company wants to combine each product with a pricing table to analyze each product at each price.

Which virtual table does a trigger use?

Inserted and Deleted.

List few advantages of Stored Procedure.

- Stored procedure can reduced network traffic and latency, boosting application performance.
- Stored procedure execution plans can be reused, staying cached in SQL Server's memory,

©<http://www.sqlauthority.com>

reducing server overhead.

- Stored procedures help promote code reuse.
- Stored procedures can encapsulate logic. You can change stored procedure code without affecting clients.
- Stored procedures provide better security to your data.

What is DataWarehousing?

- **Subject-oriented**, meaning that the data in the database is organized so that all the data elements relating to the same real-world event or object are linked together;
- **Time-variant**, meaning that the changes to the data in the database are tracked and recorded so that reports can be produced showing changes over time;
- **Non-volatile**, meaning that data in the database is never over-written or deleted, once committed, the data is static, read-only, but retained for future reporting;
- **Integrated**, meaning that the database contains data from most or all of an organization's operational applications, and that this data is made consistent.

What is OLTP(OnLine Transaction Processing)?

In OLTP - online transaction processing systems relational database design use the discipline of data modeling and generally follow the Codd rules of data normalization in order to ensure absolute data integrity. Using these rules complex information is broken down into its most simple structures (a table) where all of the individual atomic level elements relate to each other and satisfy the normalization rules.

How do SQL server 2000 and XML linked? Can XML be used to access data? FOR XML (ROW, AUTO, EXPLICIT)

You can execute SQL queries against existing relational databases to return results as XML rather than standard rowsets. These queries can be executed directly or from within stored procedures. To retrieve XML results, use the FOR XML clause of the SELECT statement and specify an XML mode of RAW, AUTO, or EXPLICIT.

OPENXML

OPENXML is a Transact-SQL keyword that provides a relational/rowset view over an in-memory XML document. OPENXML is a rowset provider similar to a table or a view. OPENXML provides a way to access XML data within the Transact-SQL context by transferring data from an XML document into the relational tables. Thus, OPENXML allows you to manage an XML document and its interaction with the relational environment.

What is an execution plan? When would you use it? How would you view the execution plan?

An execution plan is basically a road map that graphically or textually shows the data retrieval methods chosen by the SQL Server query optimizer for a stored procedure or ad-hoc query and is a very useful tool for a developer to understand the performance characteristics of a query or stored procedure since the plan is the one that SQL Server will place in its cache and use to execute the stored procedure or query. From within Query Analyzer is an option called "Show Execution Plan" (located on the Query drop-down menu). If this option is turned on it will display query execution plan in separate window when query is ran again.