

# ***Troy Technologies USA***



## **MCSE STUDY GUIDE**

***SQL Admin 7.0***

***Exam 70-28***

**Edition 2**

## Congratulations!!

You have purchased a ***Troy Technologies USA*** Study Guide.

This study guide is a selection of questions and answers similar to the ones you will find on the official SQL Admin 7.0 MCSE exam. Study and memorize the following concepts, questions and answers for approximately 10 to 15 hours and you will be prepared to take the exams. We guarantee it!

Remember, average study time is 10 to 15 hours and then you are ready!!!

GOOD LUCK!

### **Guarantee**

If you use this study guide correctly and still fail the exam, send your official score notice and mailing address to:

Troy Technologies USA  
8200 Pat Booker Rd. #368  
San Antonio, TX 78233

We will gladly refund the cost of this study guide. However, you will not need this guarantee if you follow the above instructions.

Ó Copyright 1999 Troy Technologies USA. All Rights Reserved.

# SQL 7.0 Admin Concepts

## Installation

The SQL 7.0 Server software can be installed on the following operating systems:

- Windows NT Server 4.0 (with SP4 or later)
- Windows NT Server Enterprise Edition (with SP4 or later)
- Windows NT Workstation 4.0 (with SP4 or later)
- Windows 95
- Windows 98

## Hardware Requirements

Processor	Pentium 166 or better
Memory	32MB 64MB for Enterprise Edition
Disk space	80MB minimum
File System	NTFS FAT
Browser	IE 4.0 (with SP1 or later)

## Network Protocols

Microsoft SQL Server uses network libraries to pass network packets back and forth between clients and a server running SQL Server. The network libraries, implemented as dynamic-link libraries (DLLs), perform the network operations required to communicate by using specific interprocess communication (IPC) mechanisms.

### Named Pipes

Named Pipes support is required on Microsoft Windows NT installations of SQL Server. Server-side Named Pipes is not supported on Microsoft Windows 95/98. You can also drop named pipe support and set SQL Server to listen only on other Net-Libraries.

### TCP/IP Sockets

This Net-Library allows SQL Server to communicate by using standard Windows Sockets as the IPC method across the TCP/IP protocol. By default, Windows 95/98 uses the TCP/IP Sockets Net-Library.

If you set SQL Server to listen on TCP/IP, type the TCP/IP port number in the **Port number** box. This is the port that SQL Server listens on when accepting connections from TCP/IP Sockets clients. The default is 1433.

### Multiprotocol

The Multiprotocol Net-Library uses the Windows remote procedure call (RPC) facility. With the Multiprotocol Net-Library, unlike other Net-Libraries, configuration parameters are not required.

The Multiprotocol Net-Library:

- Communicates over most IPC mechanisms supported by Windows NT. (Only TCP/IP Sockets, NWLink IPX/SPX, and Named Pipes are considered tested and supported.)
- Allows the use of Windows NT Authentication over all protocols that RPC supports.

- Supports encryption for user password authentication as well as data.
- Offers performance comparable to native IPC Net-Libraries for most applications.

The Multiprotocol encryption is not supported on Windows 95/98.

## NWLink IPX/SPX

This Net-Library allows SQL Server to communicate using the NWLink IPX/SPX protocol.

If you set up SQL Server to listen on NWLink IPX/SPX, the Setup program prompts you for the **Novell Bindery service name** in which to register SQL Server on the Novell network. The default service name is the computer name of the server computer. The Net-Library allows Novell SPX clients to connect.

## AppleTalk ADSP

The server AppleTalk (ADSP) Net-Library allows Apple Macintosh clients to connect to SQL Server by using native AppleTalk (as opposed to TCP/IP Sockets).

It is not necessary to enter an AppleTalk zone because the local zone is used when registering the service.

The AppleTalk Net-Library is not supported on Windows 95/98.

## Banyan VINES

SQL Server supports Banyan VINES Sequenced Packet Protocol (SPP) as the IPC method across the Banyan VINES IP network protocol. Banyan VINES support for clients and servers running Windows NT is available for SQL Server on the Intel platform only; it is not currently available on Windows 95/98 or the Alpha AXP platform.

## SQL Server Components

When installing Microsoft SQL Server, you can choose which components to install from these categories:

- SQL Server program files
- Management tools
- Client connectivity
- Online documentation
- Development tools
- Code samples

You can run SQL Server Setup again later to install components that you did not install in the initial setup, or reinstall components. When run after the initial installation, the Setup program proceeds to the **SQL Server Components** dialog box. The typical installation options are selected by default. Select the components you want to install or reinstall and clear all others.

## SQL Server Program Files

These components can be installed from the Microsoft SQL Server program files category in SQL Server Setup.

## SQL Server

Installs the SQL Server relational database engine and other core tools. If any SQL Server program files are installed, the SQL Server component must be installed.

**Note** When installing the SQL Server component, the Setup program also installs **bcp**, **isql**, **osql**, ODBC, OLE DB, and DB-Library.

## Upgrade Tools

Installs the SQL Server Upgrade Wizard, which is used to upgrade SQL Server version 6.x databases to SQL Server 7.0. The SQL Server Upgrade Wizard can be installed only on a computer running Microsoft Windows NT.

## Replication Support

Installs the scripts and binary files used for replication.

## Full-Text Search

Installs the Microsoft full-text search engine, which extends the ability to search on character columns beyond the basic equality and LIKE operators. Full-text search is not available in the Desktop Edition.

## Management Tools

SQL 7.0 has several management tools that make server administration fairly easy

### SQL Server Agent

SQL Server Agent runs on the server running Microsoft SQL Server. It is responsible for:

- Running SQL Server tasks scheduled to occur at specific times or intervals.
- Detecting specific conditions for which administrators have defined an action, such as alerting someone through pages or e-mail, or a task that will address the conditions.
- Running replication tasks defined by administrators.

SQL Server Agent is like an auxiliary operator that is responsible for handling the repetitive tasks and exception handling conditions defined through the other SQL-DMF components.

### SQL Server Enterprise Manager

Microsoft Management Console (MMC) is a tool that presents a common interface for managing different server applications in a Microsoft Windows network. Server applications provide a component called an MMC snap-in that presents MMC users with a user interface for managing the server application. SQL Server Enterprise Manager is the Microsoft SQL Server MMC snap-in.

SQL Server Enterprise Manager is the primary administrative tool for SQL Server and provides an MMC-compliant user interface that allows users to:

- Define groups of servers running SQL Server.
- Register individual servers in a group.
- Configure all SQL Server options for each registered server.
- Create and administer all SQL Server databases, objects, logins, users, and permissions in each registered server.

- Define and execute all SQL Server administrative tasks on each registered server.
- Design and test SQL statements, batches, and scripts interactively by invoking SQL Server Query Analyzer.
- Invoke the various wizards defined for SQL Server.

## **SQL Server Query Analyzer**

SQL Server Query Analyzer is a graphical user interface for designing and testing Transact-SQL statements, batches, and scripts interactively. It can be called from SQL Server Enterprise Manager.

SQL Server Query Analyzer offers the following features:

- Free-form text editor for keying in Transact-SQL statements.
- Color coding of Transact-SQL syntax to improve the readability of complex statements.
- Results presented in either a grid or a free-form text window.
- Graphical diagram of the showplan information showing the logical steps built into the execution plan of a Transact-SQL statement. This allows programmers to determine what specific part of a poorly-performing query is using a lot of resources. They can then explore changing the query in ways that minimize the resource usage while still returning the desired data.
- Index Tuning Wizard to analyze a Transact-SQL statement and the tables it references to see if adding additional indexes will improve the performance of the query.

## **SQL Server Profiler**

SQL Server Profiler is a tool that captures Microsoft SQL Server events from a server. The events are saved in a trace file that can later be analyzed or used to replay a specific series of steps when trying to diagnose a problem. SQL Server Profiler is used for activities such as:

- Stepping through problem queries to find the cause of the problem.
- Finding and diagnosing slow-running queries.
- Capturing the series of SQL statements that lead to a problem. The saved trace can then be used to replicate the problem on a test server where the problem can be diagnosed.
- Monitoring the performance of SQL Server to tune workloads.

## **SQL Server Service Manager**

SQL Server Service Manager is used to start, stop, and pause the Microsoft SQL Server components on the server. These components run as services on Microsoft Windows NT, and as separate executable programs on Microsoft Windows 95/98:

- MSSQLServer Service
- Database server for SQL Server.
- SQLServerAgent service
- Agent that runs scheduled administrative tasks.
- Microsoft Search service (Windows NT only)

- Full-text search engine.
- MSDTC service (Windows NT only)
- Manager for distributed transactions.

## Databases

There are 2 basic types of database used by SQL Server. These are the SYSTEM databases and the USER databases.

## System Databases

Microsoft SQL Server systems have four system databases:

Table	Function
<b>master</b>	The <b>master</b> database records all of the system level information for a SQL Server system. It records all login accounts and all system configuration settings. <b>master</b> is the database that records the existence of all other databases and the location of the primary files that contain the initialization information for the user databases. <b>master</b> records the initialization information for SQL Server; always have a recent backup of <b>master</b> available.
<b>tempdb</b>	<b>tempdb</b> holds all temporary tables and temporary stored procedures. It also fills any other temporary storage needs such as work tables generated by SQL Server. <b>tempdb</b> is a global resource; the temporary tables and stored procedures for all users connected to the system are stored there. <b>tempdb</b> is re-created every time SQL Server is started so the system starts with a clean copy of the database. Because temporary tables and stored procedures are dropped automatically on disconnect, and no connections are active when the system is shut down, there is never anything in <b>tempdb</b> to be saved from one session of SQL Server to another.  <b>tempdb</b> autogrows as needed. Each time the system is started, <b>tempdb</b> is reset to its default size. You can avoid the overhead of having <b>tempdb</b> autogrow by using ALTER DATABASE to increase the size of <b>tempdb</b> .
<b>model</b>	The <b>model</b> database is used as the template for all databases created on a system. When a CREATE DATABASE statement is issued, the first part of the database is created by copying in the contents of the <b>model</b> database, then the remainder of the new database is filled with empty pages. Because <b>tempdb</b> is created every time SQL Server is started, the <b>model</b> database must always exist on a SQL Server system.
<b>msdb</b>	The <b>msdb</b> database is used by SQL Server Agent for scheduling alerts and jobs, and recording operators.

In SQL Server version 7.0, every database, including the system databases, has its own set of files and does not share those files with other databases. The default location for these files is the C:\Mssql7\Data directory.

## User Databases

Microsoft SQL Server version 7.0 maps a database over a set of operating system files. Data and log information are never mixed on the same file, and individual files are used only by one database.

SQL Server 7.0 databases have three types of files:

- Primary data files

The primary data file is the starting point of the database and points to the rest of the files in the database. Every database has one primary data file. The recommended file extension for primary data files is .mdf.

- Secondary data files

Secondary data files comprise all of the data files other than the primary data file. Some databases may not have any secondary data files, while others have multiple secondary data files. The recommended file extension for secondary data files is .ndf.

- Log files

Log files hold all of the log information used to recover the database. There must be at least one log file for each database, although there can be more than one. The recommended file extension for log files is .ldf.

SQL Server 7.0 does not enforce the .mdf, .ndf, and .ldf file extensions, but these extensions are recommended to help identify the use of the file.

Creating and Modifying Databases

## Creating Databases

Creation of a database can be accomplished graphically via the Enterprise Manager, or via Transact-SQL statements. The following is the basic syntax for creating a database using Transact-SQL:

```
CREATE DATABASE database_name
```

A database can be easily deleted using the DROP DATABASE command. The syntax is simply:

```
DROP DATABASE databasename
```

## Altering a Database

Once a database has been created, some items can be changed. The ALTER DATABASE command provides the ability to add files to the database, drop files, modify file properties, and manage file groups.

```
ALTER DATABASE database
```

Note that to DECREASE the size of database or database file you must use the DBCC SHRINKDATABASE command.

```
DBCC SHRINKDATABASE (databasename, target percent free)
```

The command DBCC SHRINKDATABASE (my\_db, 20), would attempt to shrink all database files such that the database had 20 percent free space left after the operation was complete. For example. If a database file is currently 200 MB with 80 MB of data, a setting of 20 would result in a database file of 100MB with 80 MB of data (80 is 80% of 100). If the percentage is not specified then the database will be made as small as possible. If the percentage specified is greater than space currently available then the operation will simply not do anything.

If you wanted to shrink the data file instead of the whole database, you could use the DBCC SHRINKFILE command:

```
DBCC SHRINKFILE (filename, size)
```

In this case, the size is the final desired size in Megabytes.



## Inserting Data

The INSERT Statement is a T-SQL command that allows you to, appropriately, insert new rows of data into a table. This command is only used to APPEND data. It does not update existing rows, nor would it delete existing rows. If the goal is to replace the data in a table, you would need to first delete the table before populating with the INSERT command.

```
INSERT my_data VALUES ('COLUMN_1_DATA', 'COLUMN_2_DATA')
```

This basic command assumes a table name of "my\_data" and that it has two columns, each supporting the type of data entered. It also assumes that the values being inserted are in the same order as the columns within the tables. If the values were not being entered in precise order it would be necessary to specify to which columns the values should be applied. For instance:

The BULK INSERT command is used with T-SQL batches to move data from a file into a table or view. It is the fastest method of quickly inserting large amounts of data, provided the source text file also resides on the server.

```
BULK INSERT my_data FROM 'c:\my_data\my_file.txt'
```

The SELECT INTO statement allows you to create a brand new table and populate it with the results of the query being performed by the statement. You would use this statement to transfer data into a table that did not already previously exist.

## SQL Views

A view can be thought of as either a virtual table or a stored query. The data accessible through a view is not stored in the database as a distinct object. What is stored in the database is a SELECT statement. The result set of the SELECT statement forms the virtual table returned by the view. A user can use this virtual table by referencing the view name in Transact-SQL statements the same way a table is referenced. A view is used to do any or all of these functions:

- Restrict a user to specific rows in a table.
- Restrict a user to specific columns.
- Join columns from multiple tables so that they look like a single table.
- Aggregate information instead of supplying details.

## Importing/Exporting Data

Microsoft SQL Server has several components that support importing and exporting data:

### Data Transformation Services (DTS)

DTS can be used to import and export data between heterogeneous OLE DB and ODBC data sources. A DTS package is defined that specifies the source and target OLE DB data sources; the package can then be executed on an ad hoc basis or at scheduled times or intervals. A single DTS package can cover multiple tables. DTS packages are also not limited to transferring data straight from one table to another, as the package can specify a query as the source of the data. This allows packages to transform data, such as running a query that returns aggregate summary values instead of the raw data.

## Replication

Replication is used to create copies of data in separate databases and keep these copies synchronized by replicating modifications in one copy to all the others. If it is acceptable for each site to have data that may be a minute or so out of date, replication allows the distribution of data without the overhead of requiring distributed transactions to ensure all sites have an exact copy of the current data. Replication can therefore support the distribution of data for a relatively low cost in network and computing resources.

## Bulk copying

The bulk copy feature of SQL Server allows for the efficient transfer of large amounts of data. Bulk copying transfers data into or out of one table at a time. Bulk copying supports the following bulk copy transfers:

- From one SQL Server table or view to another table or view.
- From a SQL Server table or view into a data file, such as a text file or tab-delimited file.
- The result set of a query into a table, view, or data file.
- The contents of a data file into a table or view.

## Distributed Queries

Distributed queries allow Transact-SQL statements to reference data in an data source. The data sources can be another copy of SQL Server, or a heterogeneous data source such as Microsoft Access or Oracle. SELECT INTO and INSERT statements can be used to:

- Export data from a SQL Server data base to an OLE DB data source.
- Import data from an OLE DB data source into SQL Server.

## Full Text Indexing

SQL 7.0 has the ability to do full-text indexing. To make use of this feature, you must install the MS Search Service, an optional component during setup of SQL Server 7.0.

Full-text indexes are considerable different than regular table indexes. For one, they are not stored within the database. Rather, they are stored within the file system of the NT. They also depend on a service, so they only are available when SQL 7.0 is installed on NT. They are not an option on the desktop version of SQL 7.0, even if running on NT.

Each table that you wish to enable for full-text indexing must already have some other unique index created. Only one full-text index may exist per table. The easiest way to create your full-text index is by using a wizard.

Full-text indexes are not self-maintaining. You must actually populate the indexes and update them from time to time. You can schedule jobs to automate this for you in the Enterprise Manager.

## File Groups

Database files can be grouped together in filegroups for allocation and administration purposes. Some systems can improve their performance by controlling the placement of data and indexes onto specific disk drives. File groups can aid this process. The system administrator can create filegroups for each disk drive, then assign specific tables, indexes, or the **text**, **ntext**, or **image** data from a table, to specific filegroups. No

file can be a member of more than one filegroup. Log files are never a part of a filegroup. Log space is managed separately from data space.

There are three types of filegroups: Primary, Secondary and Default.

## Primary

The primary filegroup contains the primary data file and any other files not put into another filegroup. All pages for the system tables are allocated in the primary filegroup.

## User-defined

User-defined filegroups are any filegroups specified using the FILEGROUP keyword in a CREATE DATABASE or ALTER DATABASE statement.

## Default

The default filegroup contains the pages for all tables and indexes that do not have a filegroup specified when they are created. In each database, only one filegroup at a time can be the default filegroup. Members of the **db\_owner** fixed database role can switch the default filegroup from one filegroup to another. If no default filegroup was specified, it defaults to the primary filegroup.

SQL Server 7.0 can work quite effectively without filegroups, so many systems will not need to specify user-defined filegroups. In this case, all files are included in the primary filegroup and SQL Server 7.0 can effectively allocate data within the database.

## Accessing Remote Data

In SQL 7.0 you can access data remotely by using Linked Servers. They are essentially OLE DB data sources that can be referenced by T-SQL statements executing on the local server. Since they are OLE DB data sources, linked servers, are not limited to just SQL servers. They could be Oracle, Access, Excel, or other OLE DB sources. They enable you to run remote stored procedures, just like Remote Servers. You can also reference data on the linked server within a SELECT statement. You can even create a join that links the data between the local and remote data source, thus allowing for distributed queries.

When SQL Server attempts to connect to the linked server some sort of authentication information will need to be provided. By default, when a linked server is created, SQL Server sets it up so that the security credentials of the current user requesting the distributed query will be emulated. However, this may not work if the client does not allow account delegation or if the linked server does not support NT authentication. If this is the case you must create a local login mapping. This is a pre-configured account and password that SQL Server will use when it attempts to make the remote connection.

## Security

Logins, permissions, and roles are the foundation for the security mechanisms of SQL Server. Users that connect to SQL Server must identify themselves using a specific login ID. Users can then only see the tables and views they are authorized to see, and can only execute the stored procedures and administrative functions they are authorized to execute. This system of security is based on the IDs used to identify users.

## Logins

In most cases, you have to give an application only two pieces of information to connect to a Microsoft SQL Server installation:

- The network name of the server running SQL Server

- Your login ID

Login IDs are the account identifiers that control access to any SQL Server. SQL Server will not complete a connection unless it has first verified that the login ID you specified is valid. This verification process is known as authentication.

There are two types of SQL Server Authentication, each of which has a different class of login ID:

- SQL Server Authentication

A member of the **sysadmin** fixed server role first specifies to SQL Server all the valid SQL Server login accounts and passwords. These are not related to your Microsoft Windows account or network account. You must supply both the SQL Server login and password when you connect to SQL Server. You are identified in SQL Server by your SQL Server login.

- Windows NT Authentication

A member of the SQL Server **sysadmin** fixed server role must first specify to SQL Server all the Microsoft Windows NT accounts or groups that can connect to SQL Server. When using Windows NT Authentication, you do not have to specify a login ID or password when you connect to SQL Server. Your access to SQL Server is controlled by your Windows NT account or group, which is authenticated when you log on to the Windows operating system on the client. When connecting, the SQL Server client software requests a Windows *trusted connection* to SQL Server. Windows NT will not open a trusted connection unless the client has logged on successfully using a valid Windows NT account. The properties of a trusted connection include the Windows NT group and user accounts of the client that opened the connection. Because it is impossible for a trusted connection to be completed without Windows NT having first authenticated the user, SQL Server does not have to do anything to authenticate your accounts. SQL Server gets the user account information from the trusted connection properties and matches them against the Windows NT accounts defined as valid SQL Server logins. If SQL Server finds a match, it accepts the connection. You are identified in SQL Server by your Windows NT group or user account.

When SQL Server is running on Windows NT, members of the **sysadmin** fixed server role can specify one of two authentication modes:

- Windows NT Authentication Mode - Only Windows NT Authentication is allowed. Users cannot specify a SQL Server login ID.
- Mixed Mode - If a user supplies a SQL Server login ID when they connect, they are authenticated using SQL Server Authentication. If they do not supply a SQL Server login ID, they are authenticated using Windows NT Authentication.

These modes are specified during setup or with SQL Server Enterprise Manager.

The Windows 95/98 operating system does not support the server side of the trusted connection. When SQL Server is running on Windows 95/98, it does not support Windows NT Authentication. Every user must supply a SQL Server login when they connect. When SQL Server is running on Windows NT, Windows 95/98 clients can connect to it using Windows NT Authentication.

## Permissions

Permissions in databases are cumulative, except where a permission has been explicitly denied (analogous to no access NTFS permission). Implicit user permissions, such as those that are acquired through role membership or those that are implicitly given to Database Object Owners, can not be directly viewed. Database Object Owners have all permissions on objects they create and can grant, revoke or deny permissions to all users, including the Database Owner, on these objects.

## Types of permission

Permissions apply to statements and objects. Statement permissions give users the ability to execute Transact-SQL commands, such as CREATE DATABASE. Object permissions give users the ability to do something, such as viewing or updating information in a table or executing a stored procedure.

- Statement Permissions: create database, create table, create view, create rule, create default, backup database, backup log.
- Object Permissions: Select, insert, update, delete, references, execute. Select, insert, update, delete and references can be applied to tables and views, select, update, references to columns, and execute to stored procedures. Predefined permissions apply to fixed roles or object owners.

## Assigning Permissions

When assigning permissions, you can use the following commands:

- Grant (can perform action)
- Deny (Cannot perform action and cannot overridden)
- Revoke (Cannot perform action but can be overridden)

Permissions can be granted to views without having to grant permissions to the underlying tables that comprise the views, provided the ownership chain is not broken. Users who have Execute permissions on stored procedures do not need to be granted permissions to modify or view the data that the stored procedure needs access to.

## Ownership Chains

Objects, such as views, have owners. When a single owner creates a series of dependent objects, such as view that in itself is created from another view or views, and owns all the objects in the chain, there is a single ownership chain. When the dbo creates View1 and then creates View2 that is based on View1, there is a single chain. If the dbo, however, grants the permission to create a view to another user and that user creates a third view based on View2, the ownership chain is broken: the user does not own the object that his or her view depends on. SQL server will check permissions only once if there is a single ownership chain--on the view itself and not on the objects it may depend on. However, if there is a broken ownership chain, SQL will check permissions on all the objects in the chain where there is a change in ownership.

## Guest User Account

The **guest** user account allows a user without a user account to access a database. A login assumes the identity of the **guest** user when all of the following conditions are met:

- The user has access to SQL Server, but does not have access to the database through his or her own user account.
- The **guest** user account has been added.

Permissions can be applied to the **guest** user as if it were any other user account. The **guest** user can be deleted and added to all databases except **master** and **tempdb**, where it must always exist. By default, a **guest** user account does not exist in newly created databases. However, if **guest** is added to the model database, every subsequently created database will have this account.

## Roles

Roles are similar to the concept of groups. By assigning permissions to a role, and then adding users to the role, all users inherit the permission of that role. This greatly eases administration. The PUBLIC role exists in all SQL 7.0 databases. Users are always a member of the PUBLIC role. Any permissions assigned to the PUBLIC role would apply to all users. In SQL 7.0, a user can be a member of multiple roles, in addition to the PUBLIC role.

In addition to user defined roles that you create to ease administration of permissions, there are also two categories of fixed roles. Their are Database Roles and Server Roles.

<b>Fixed Server Roles</b>	
Role	Permissions
dbcreator	Create & alter databases
diskadmin	Manage disk files
processadmin	Manage SQL Server processes, such as the ability to kill errant queries.
securityadmin	Manage and audit server logins, permission to create databases and read error logs.
serveradmin	Configure server wise settings, such as those available with the sp_configure command. Shut down server using the SHUTDOWN command.
setupadmin	Has the ability to manage linked servers and designate stored procedures to run at startup.
sysadmin	Perform any activity. They essentially override all security mechanisms.
<b>Fixed Database Roles</b>	
public	Maintains all default permissions
db_owner	Perform any activity within the database.
db_accessadmin	Add/remove database users, groups, and roles
db_ddladmin	Add/modify/drop database objects
db_securityadmin	Assign statement and object permissions
db_backupoperator	Backup & restore databases
db_datareader	Read data from any table
db_datawriter	Add/change/delete data from all tables
db_denydatareader	Cannot read data from any table
db_denydatawriter	Cannot change data in any table
db_dumpoperator	Can issue DBCC commands, checkpoint, and backup.

There is one other special type of role called the Application Role. They are similar to user defined roles in that they allow you to assign permissions to the role instead of individual users. Application roles are "invoked" by a client application. Once invoked, the users permissions are suspended and they then assume all permissions of the application role.

## Stored Procedures

Transact-SQL is the language used for all commands sent to Microsoft SQL Server from all applications and contains statements that support all administrative work done in SQL Server. These statements fall into two main categories.

## Data Definition Language

The SQL language has two main divisions: Data Definition Language (DDL) that is used to define all the objects in an SQL database, and Data Manipulation Language (DML) that is used to select, insert, update, and delete data in the objects defined using DDL. Permissions are controlled using the GRANT, REVOKE and the DENY statements.

## System Stored Procedures

Administrative tasks not covered by the DDL are typically done using system stored procedures. These are stored procedures whose names start with `sp_` or `xp_` and are installed when SQL Server is installed. A stored procedure is a group of Transact-SQL statements compiled into a single execution plan.

Microsoft SQL Server stored procedures return data in four ways:

- Output parameters, which can return either data (such as an integer or character value) or a cursor variable (cursors are result sets that can be retrieved one row at a time).
- Return codes, which are always an integer value.
- A result set for each `SELECT` statement contained in the stored procedure or any other stored procedures called by the stored procedure.
- A global cursor that can be referenced outside the stored procedure.

Stored procedures assist in achieving a consistent implementation of logic across applications. The SQL statements and logic needed to perform a commonly performed task can be designed, coded, and tested once in a stored procedure. Each application needing to perform that task can then simply execute the stored procedure. Coding business logic into a single stored procedure also offers a single point of control for ensuring that business rules are correctly enforced.

Stored procedures can also improve performance. Many tasks are implemented as a series of SQL statements. Conditional logic applied to the results of the first SQL statements determines which subsequent SQL statements are executed. If these SQL statements and conditional logic are written into a stored procedure, they become part of a single execution plan on the server. The results do not have to be returned to the client to have the conditional logic applied; all of the work is done on the server.

Applications do not need to transmit all of the SQL statements in the procedure: they only have to transmit an `EXECUTE` or `CALL` statement containing the name of the procedure and the values of the parameters.

Stored procedures can also shield users from needing to know the details of the tables in the database. If a set of stored procedures supports all of the business functions users need to perform, users never need to access the tables directly; they can just execute the stored procedures that model the business processes with which they are familiar.

An illustration of this use of stored procedures is the SQL Server system stored procedures used to insulate users from the system tables. SQL Server includes a set of system stored procedures whose names usually start with `sp_`. These system stored procedures support all of the administrative tasks required to run a SQL Server system. You can administer a SQL Server system using the Transact-SQL administration-related statements (such as `CREATE TABLE`) or the system stored procedures, and never need to directly update the system tables.

SQL Server also supports temporary stored procedures that, like temporary tables, are dropped automatically when you disconnect. Temporary stored procedures are stored in **tempdb** and are useful when connected to earlier versions of SQL Server. Temporary stored procedures can be used in the case where an application builds dynamic Transact-SQL statements that are executed several times. Rather than have the Transact-SQL statements recompiled each time, you can create a temporary stored procedure that is compiled on the first execution, then execute the precompiled plan multiple times.

## Replication

Replication makes possible the transfer of data from a source to a destination or destinations. Allows for site autonomy and scalability. Can be used to ensure transactional integrity without the overhead of distributed transactions.

Distributed Transactions occur at the same time to all copies of data on all servers involved in the transaction. Using the 2-phase commit protocol (2PC) ensures transactions are committed on all servers or not at all on any. Requires good connectivity between servers. Useful when data needs to be same across all servers at the same time.

## Replication Types

**Snapshot Replication** - Takes a picture of the data at a point in time. Not as CPU intensive as Transactional Replication, which has to monitor publications for updates. Simplest type of replication. Guarantees latent transactional integrity between source and destination. Good for read-only subscribers who do not need most recent copy of data.

**Snapshot Replication with Immediate-Updating Subscribers** - An optional configuration of Snapshot Replication that allows subscribers to make changes at the subscriber and the publisher using 2PC. Transactional integrity is maintained between publisher and subscriber. This method of updating the publisher requires that only the subscriber and the publisher involved in the transaction be enlisted for the distributed transaction, not all the servers subscribing to the publication. Good for situations where subscribers have to make occasional updates to data.

**Transactional Replication** - Used for replicating tables (all or part of a table) and stored procedures. The Log Reader Agent monitors the logs of publications for INSERT, UPDATE, DELETE statements and other modifications and then stores these modifications in a queue, the distribution database, for replication to subscribers. Changes are made at the publication server, so transactional integrity is guaranteed. Given good network connections, there can be low latency between publisher and subscriber (less than a minute for push subscriptions). Can also be used for pull subscriptions where subscribers are not always connected and require read-only data, eg., salesperson who needs to get inventory and price lists.

**Transactional Replication with Immediate-Updating Subscribers** - Allows subscribers to make updates to their local data and the data on the publisher using a distributed transaction. Transactional integrity guaranteed using 2PC. All subscribers eventually have transactions replicated to them from the publisher.

**Merge Replication** - In merge replication, both the publisher and the subscriber update data. The data contained in the replica copies held by the publisher and the subscribers are the result of synchronization (convergence). With merge replication, there is no guarantee of transactional integrity and conflicts between updates can arise. SQL Server resolves conflicts based on 'generation numbers' and configured priorities--some server will "win" in the case of a conflict. Merge replication guarantees that eventually all servers will converge to the same resultant data, but the converged data may be different from the data resulting from other forms of replication that guarantee transactional integrity.

## Replication Components

The Microsoft SQL Server version 7.0 replication model is based on the publish and subscribe metaphor popularized by SQL Server 6.0. The replication model consists of:

- Publishers
- Distributors
- Subscribers



- Publications
- Articles
- Push subscriptions
- Pull subscriptions

The ***Publisher*** is a server that makes data available for replication to other servers. In addition to identifying which data is to be replicated, the Publisher detects which data has changed and maintains information about all publications at that site. Any given data element that is replicated has a single Publisher, even if it may be updated by any number of Subscribers or published again by a Subscriber.

The ***Distributor*** is the server that contains the distribution database. The exact role of the Distributor is different in each type of SQL Server 7.0 replication.

***Subscribers*** are servers that store replicas and receive updates. In earlier versions of SQL Server, updates could typically be performed only at the Publisher. However, SQL Server 7.0 allows Subscribers to make updates to data (but a Subscriber making updates is not the same as a Publisher). A Subscriber can, in turn, become a Publisher to other Subscribers.

A ***publication*** is a collection of articles, and an article is a grouping of data to be replicated. An ***article*** can be an entire table, only certain columns (using a vertical filter), only certain rows (using a horizontal filter), or even a stored procedure (in some types of replication). A publication often has multiple articles. This grouping of multiple articles makes it simpler to subscribe to a unit (the publication), which has all the relevant and required data. Subscribers subscribe only to a publication, not to individual articles within a publication.

With a ***push subscription***, the Publisher propagates the changes to a Subscriber without a request from the Subscriber to do so. Typically, push subscriptions are used in applications that are required to send changes to Subscribers whenever and as soon as they occur. Push subscriptions are best for publications that require near-real-time movement of data without polling and where the higher processor overhead at the Publisher does not affect performance. Changes can also be pushed to Subscribers on a scheduled basis.

With a ***pull subscription***, the Subscriber asks for periodic updates of all changes at the Publisher. Pull subscriptions are best for publications having a large number of Subscribers (for example, Subscribers using the Internet). Pull subscriptions are also best for autonomous mobile users because they allow the user to determine when the data changes are synchronized. A single publication can support a mixture of push and pull subscriptions.

## Replication Agents

In addition to the basic components, your replication design may have two or more replication agents:

- Snapshot Agent
- Log Reader Agent
- Distribution Agent
- Merge Agent

The ***Snapshot Agent*** prepares schema and initial data files of published tables and stored procedures, stores the snapshot on the Distributor, and records information about the synchronization status in the distribution database. Each publication has its own Snapshot Agent that runs on the Distributor and connects to the Publisher. The Snapshot Agent is typically run under SQL Server Agent and can be administered directly using SQL Server Enterprise Manager.

The **Log Reader Agent** moves transactions marked for replication from the transaction log on the Publisher to the distribution database. Each database published using transactional replication has its own Log Reader Agent that runs on the Distributor and connects to the Publisher.

The **Distribution Agent** moves the transactions and snapshot jobs held in distribution database tables to Subscribers. Transactional and snapshot publications that are set up for immediate synchronization when a new push subscription is created each have their own Distribution Agent that runs on the Distributor and connects to the Subscriber. Transactional and snapshot publications not set up for immediate synchronization share a Distribution Agent across the Publisher/Subscriber pair that runs on the Distributor and connects to the Subscriber. Pull subscriptions to either snapshot or transactional publications have Distribution Agents that run on the Subscriber instead of the Distributor. Merge publications do not have a Distribution Agent. The Distribution Agent typically runs under SQL Server Agent and can be administered directly using SQL Server Enterprise Manager.

The **Merge Agent** moves and reconciles incremental data changes that occurred after the initial snapshot was created. Each merge publication has its own Merge Agent that connects to both the Publisher and the Subscriber and updates both. In a full merge, the agent first uploads all changes from the Subscriber where the generation is 0 or is greater than the last generation sent to the Publisher. The agent gathers the rows, and those rows without conflicts are applied to the publishing database. Those rows with conflicts are handled by the conflict resolver associated with the article in the publication definition. All changes are applied using stored procedures derived from the Publisher tables at the time snapshot is generated or first applied. Finally, the agent reverses the process by downloading any changes from the Publisher to the Subscriber and applying the changes to the subscribing database. Push subscriptions to merge publications have Merge Agents that run on the Publisher, while pull subscriptions to merge publications have Merge Agents that run on the Subscriber. Snapshot and transactional publications do not have Merge Agents.

## **Replication Models**

### **Central Publisher/Distributor**

Both the Publisher and the Distributor are on the same machine with Subscribers on separate servers.

### **Central Publisher with Remote Distributor**

Like above, except Publisher and Distributor are on different machines. In heavy OLTP environments, this scenario is useful in that it reduces the load on the Publisher. Requires good network connectivity between Publisher and Distributor.

### **Publishing Subscriber**

In this scenario, the Subscriber is also responsible for republishing the received data. Useful in situations where there is low available bandwidth between locations.

### **Central Subscribers/Multiple Publishers**

A number of publishers replicate data to a common destination table on the subscriber. The data has to be partitioned and a primary key used to identify the source region/server. Useful for rolling up information.

### **Multiple Publishers/Multiple Subscribers**

Each replicates information to and receives replicated information from a common table. Useful for situations where sites have to be able to view information updated in other sites.

Any replication type can be used with any model. The model is simply the physical topology of your replication.

## **Backup and Restore**

The backup and restore components of Microsoft SQL Server provide the capability of creating a copy of a database. This copy is stored in a location that is protected from failures of the server running SQL Server. If the server running SQL Server fails, or if the database is somehow damaged, the backup copy can be used to re-create, or restore, the database.

SQL Server has sophisticated backup and restore capabilities:

- A full database backup is a full copy of the database.
- A transaction log backup copies only the transaction log.
- A differential backup copies only the database pages that have been modified after the last full database backup.
- A file or filegroup restore allows the recovery of just the portion of a database that was on the failed disk.

This allows backup and restore processes to be tailored to how critical the data in the database is. Noncritical databases that can be easily re-created from some other source may have no backups, other databases may have simple backups that can re-create the database to the night before a failure, and critical databases may have sophisticated backups that will restore the database right up to the point of failure.

## **Backup Methods**

### **Full Backup**

The full database backup does a complete backup of the entire contents of your database. It will even capture changes that occur to the database during the backup operation. It does this by backing up a portion of the transaction log as well. Note, however, that you must still backup your transaction logs separately. Full Database backups do not provide for transaction log management (truncation).

### **Differential Backup**

Each page in a database contains a Log Sequence Number (LSN). Each transaction entered into the transaction log gets a LSN. When a differential backup starts, SQL Server looks at the highest LSN number of the last full backup. It then proceeds to go through the database. When it encounters a page with an LSN higher than that from the last full backup it backs up the entire extent that contains the modified page. In this way, the differential backup captures just the changes to the database since the last time a full backup was completed.

If you think about how this works, you will realize that as successive differential backups are performed they continue to get larger and larger, until the next full backup is completed. If you were to do a full backup on Sunday and a differential every day, Mondays differential will contain only changed data from Monday. However, by the time you get to Saturday, the backup will contain data from every day except Sunday.

The benefit of Differential backups is that the length of time necessary to both backup and restore your data is reduced. In comparison to full backups, differentials take less time since you are only backing up changes. When it comes to restores, differential reduce time by allowing you to eliminate many log backups; it is only necessary to restore the most recent differential backup and an logs since then.

## Transaction Log Backup

A transaction log backup backs up just the transaction log and also truncates the inactive portion of the log. Because the transaction log gets truncated each time, a transaction log backup only contains changes since the last transaction log backup, effectively making it operate like an incremental backup. So when restoring transaction logs you must restore all transaction logs and they must be restored in order.

Another important note is an option to backup the transaction log without actually truncating the inactive portion of the log (NO\_TRUNCATE). This should be considered a part of any recovery procedure. When a failure occurs to the primary database file, you can possibly recover up to the minute of failure, but only if you have a backup of the current, non truncated, transaction log.

Conversely, there may be a time when a transaction log becomes full. When this occurs, your database can not accept updates. You therefore need to clear the transaction log using the TRUNCATE\_ONLY option.

## Database File/File Group Backup

With SQL 7.0 it is possible to backup individual database files or database file groups separately from the rest of the database. This is an extremely beneficial feature for very large databases.

## Data Recovery

Every Microsoft SQL Server database has a transaction log that records all transactions and the modifications made by the transactions in the database. This record of transactions and their modifications serves three functions:

*Recovery of individual transactions* - If an application issues a ROLLBACK statement, or if SQL Server detects an error such as the loss of communication with a client, the log records are used to roll back the modifications made by an incomplete transaction.

*Recovery of all incomplete transactions when SQL Server is started* - If a server running SQL Server fails, the databases may be left in a state where some modifications were never written from the buffer cache to the data files, and there may be some modifications from incomplete transactions in the data files. When a copy of SQL Server is started, it runs a recovery of each database. Every modification recorded in the log which may not have been written to the data files is rolled forward. Every incomplete transaction found in the transaction log is then rolled back to ensure the integrity of the database is preserved.

*Rolling a restored database forward to the point of failure* - Each time a full or differential database backup is restored, recovery is run to roll back any incomplete transactions. After a database backup has been restored, transaction log backups can be used to roll forward all completed transactions on the log backup. This allows a database to be restored to the point at which the server failed.

## Write-Ahead Transaction Log

Microsoft SQL Server, like many relational databases, uses a write-ahead log. A write-ahead log ensures that no data modifications are written to disk before the associated log record.

SQL Server maintains a buffer cache into which it reads data pages when data must be retrieved. Data modifications are not made directly to disk, but are instead made to the copy of the page in the buffer cache. The modification is not written to disk until the lazywriter process schedules a write for the page. Writing a modified data page from the buffer cache to disk is called flushing the page. A page which has been modified in the cache but is not yet written to disk is called a dirty page.

At the time a modification is made to a page in the buffer, a log record is built in the log cache recording the modification. This log record must be written to disk before the associated dirty page is flushed from the buffer cache to disk. If the dirty page were flushed before the log record, it would create a modification on disk that could not be rolled back if the server failed before the log record were written to disk. SQL Server has logic that prevents a dirty page from being flushed before the associated log record. Because log records are always written ahead of the associated data pages, the log is called a write-ahead log.

## Truncating the Transaction Log

If log records were never deleted from the transaction log, the log would keep growing until it filled all the available space on the disks holding the log. At some point in time, old log records no longer necessary for recovering or restoring a database must be deleted to make way for new log records. The process of deleting these log records is called truncating the log.

The active portion of the transaction log can never be truncated. The active portion of the log is the part of the log needed to recover the database at any time. It must always be present in the database in case the server fails because it will be required to recover the database when the server is restarted. The record at the start of the active portion of the log is identified by the minimum recovery log sequence number (MinLSN).

The backup and restore process chosen for a database determines how much of the transaction log in front of the active portion must be retained in the database. While the log records in front of the MinLSN play no role in recovery, they are required to roll forward updates when using log backups to restore a database to the point of failure.

The log records before the MinLSN are only needed to maintain a sequence of log backups. If a log backup sequence is not being maintained, all log records before the MinLSN can be truncated at any time. If a log backup sequence is being maintained, the part of the log before the MinLSN cannot be truncated until those log records have been copied to a log backup.

## Rebuilding the Registry

A utility is available called *regrebl.exe*. This new utility can be found in the \mssql7\binn\ directory and provides the capability to both back and restore the SQL Server registry items. It accepts one of two parameters, -backup or -restore. The backup option creates five files (mssql7\*.rbk) within the \mssql7\binn\ directory. Each of these files relate to various keys that have been backed up. The restore option looks for these files within the \mssql7\binn\ directory and restores them to the registry.

## Rebuilding the MASTER DATABASE

A utility exists called *rebuildm.exe*. It can be found in the \mssql7\binn\ directory. As with previous versions of SQL, the MODEL and MSDB databases will also be rebuilt. After the master database has been rebuilt you can attach your databases using the `sp_attach_db` commands.

## SQL 7.0 Exam Questions

**1: You want to repair allocation errors and delete corrupted text objects in your database. What command would you enter?**

*A: DBCC CHECKDB('my\_db', REPAIR\_ALLOW\_DATA\_LOSS)*

**2: Every night the Sales backup fails. Which DBCC command could repair the Sales database and check the data integrity and also check the indices?**

*A: DBCC CHECKDB*

**3: You restore the database from tape backup. Upon 50% completion, you have a tape drive failure. You replace the drive and want to finish your backup. What now?**

*A: EXEC RESTORE WITH RESTART*

**4: Some front end applications receive error 1205. The developers invoked error handling. The queries appear to be very slow and sometimes appear to stop responding. You want to find out which objects are causing the 1205 error message and whether the objects are degrading user response time. What would you do?**

*A: In SQL Server Profiler, monitor the Lock:DeadLock, Lock:Deadlock Chain, and Lock:Timeout events.*

**5: The server's response time has recently slowed as the result of a small number of very long-running queries. These are ad hoc queries that are poorly written. You need to prevent the queries from running for long periods of time. What must you do?**

*A: Set the query governor to stop long-running queries.*

**6: You suspect that most performance problems are caused by a small number of process intensive queries. To reduce the overhead caused by these queries, you first need to identify the specific queries that are causing the problem. What should you do to identify these queries?**

*A: Filter in SQL Server Profiler on the minimum milliseconds of CPU time to capture the appropriate queries.*

**7: You want to examine the master database file settings in SQL Server Enterprise Manager, but the master database is not listed in the Database folder. You want to set the option to make the master database visible. Where can you set this option?**

*A: In the Enterprise Manager registration*

**8: You load a lot of data to the Parts table in the Inventory database. What is the quickest way to ensure statistics reflect the changes made?**

*A: Run UPDATE STATISTICS on the Parts table*

**9: SQL1 and SQL2 are linked Servers. They have standard logins on both. You can connect to both and SELECT on both. However, when you issue a distributed query for SQL2 on SQL1, it fails. What is the reason?**

*A: Login in SQL1 not mapped in SQL2*

**10: The Trace you currently use is saved. You want to use similar traces for a quantity. How must you edit the saved Trace you currently use to include additional events?**

*A: Open the Trace definition and edit the event classes.*

**11: Your SQL server database is currently running on a single computer. To improve the availability of the SQL server database to an accounting application you are implementing a second computer running Microsoft NT Server. The accounting application must be available every business day between 9 AM and 5 PM. In the event of disk drive failure, the data must be recoverable to the end of the previous business day. How should you configure the second computer ?**

*A: Install it as a SQL Server standby server*

**12: You have a web based SQL server and Internet users need to access an order entry application to enter new orders and inquiries. You must develop a strategy to ensure availability of the application and integrity of the database. You decide to use a second computer. How would you configure it?**

*A: Failover Support*

**13: You have the following setup: Publisher (SQL6.5;TOR01), Distributor (SQL6.5;TOR02), Subscriber(SQL6.5;TOR03). In which order should you upgrade the servers to SQL 7.0?**

*A: TOR02, TOR01, TOR03*

**14: A user forgot his SQL Server password. You are a member of the SYSADMIN role. What is the easiest way for you to change his password to a new one?**

*A: Change the password in the SQL Server login properties page.*

**15: Your company acquired another company which has an Oracle database for its employees. Your company uses SQL Server for its employee database. You want to merge the data from the new company into your SQL Server Employee database. Below is the structure of both databases:**

SQL Server	Oracle
Employee ID	Employee ID
Name	Name
Addr1	Address
Addr2	City
City	State

State	Zip
Zip	

**What method should be used to accomplish the merge?**

*A: Use INSERT INTO command*

**16: The Development department uses a Oracle database with one table. You want to convert it to SQL 7 and no target table exists in SQL Server. What method would you use?**

*A: SELECT..INTO*

**17: Your company's server running SQL Server is on the same BDC as the server running Microsoft Exchange Server. Exchange Server is performing poorly. The SQL Server set working set size option is at the settings you configured for the min server memory and max server memory options.**

**What should you do to improve the performance of the server running Exchange Server while not adversely affecting SQL Server?**

*A: Increase the min server memory setting.*

**18: You plan to add 100,000 rows of new data to your main database table over the next few weeks. The table currently contains 50,000 rows, it contains three columns of int data type and one column of the nchar(20) data type, and no columns allow null values. A maximum of 144 rows can fit on a data page.**

**You want the data update processes to run as fast as possible. You decide to re-create the indexes on the table and to use a different fill factor. To what value should you set the fill factor?**

*A: 30 percent*

**19: Maria preceded you as administrator for Microsoft Windows NT Server and SQL Server. Maria created SQL Server logins and passwords matching users Windows NT accounts. When users logon to the Windows NT Server domain they obtain access to the SQL Server. Which permissions do users receive to access SQL Server by using a trusted connection?**

*A: Only the same permissions as those that are granted to their Windows NT account.*

**20: You define full-text indexing on the ProductName column in the Products table. You then execute a full-text query on the column. You specify a word that you know is present in the column, but the result set is empty. What is the most likely cause?**

*A: The catalog is not populated.*

**21: You are relocating a computer running SQL Server to a different region and you want to change the UNICODE collation. What must you do?**

*A: Script your database objects, and export your data to files; Run Rebuildm.exe; Re-create your database objects, and reload the data.*

**22: You plan to give users in the Windows NT Sales group limited access to the Sales and Inventory databases. All access for data modifications and reporting will be through a Microsoft Excel application. How should you set up access for the Sales group in both databases?**

*A: Create an application role in each database and grant permissions on each role.*



**23: To increase security, your company decides to prohibit Windows NT administrators from accessing the SQL Server administrative functions. What should you do to prohibit access?**

*A: Remove the Windows NT Administrator group from the sysadmin role on SQL Server.*

**24: The transaction log in a published database is full. You attempt to truncate the log, but you notice that the log remains full. What is the most likely cause?**

*A: The Log Reader Agent has stopped.*

**25: You have filegroup1 (fg1) on disk1 and filegroup2 (fg2) on disk2. You have a customer table on fg1 and its index on disk2. fg1 on disk1 fails. What is the quickest way to restore?**

*A: Apply fg1 backup and fg2 backup. Restore all the transaction logs since the backup.*

**26: You have 3 disks (C, D & E). On C you have program files, system database files, msdb, etc. On D you have the transaction log. On E you have the user database files. There are no frequent updates. However, there are heavy queries. Users report slow performance. What do you do?**

*A: Create new filegroups on D and place heavily used indexes in the new filegroups.*

**27: You have SQL 7.0 on Windows 98 machines in your Branch offices. In the corporate office you have SQL 7.0 on a NT Server. Branch offices report slow performance. Database Admin is in the corporate office. How should the profiler be configured in the branch offices so that they can be centrally collected in the Corporate office by the Database Admin?**

*A: Use SQL Profiler extended stored procedure and send them to the table in the corporate office.*

**28: You install transactional replication on the Manufacturing server and make several publications available. The Accounting department subscribes to one of the publications, but Accounting users need data from only the WorkSchedule article. What is the easiest way to provide only this data for them?**

*A: Create a new publication containing only the WorkSchedule table.*

**29: You start receiving continuous error messages by email indicating that the transaction log is full. After 2 days the messages suddenly stop. What could be the reason?**

*A: Windows NT Application log is full.*

**30: You have a SQL Server 6.5 application. You want developers, working on Win98, to start working on 7.0, but you do not want to upgrade the database until the new application is ready. What is the quickest way to do this?**

*A: Install SQL 7.0 desktop on the developers machine and use DTS to import the data.*

**31: The computer which was running SQL server crashes. Windows NT is reinstalled. What is the quickest way to get SQL server up and running?**

*A: Run regrebl.exe*

**32: You are a member of the fixed server role serveradmin. You create a job which backs up the database and then copies the backup file onto a network drive. When you do the individual steps, they run okay. However, when executed as a job, it backs up the database but fails to copy to the network drive. Why?**

*A: SQLAgentCmdExe does not have permissions to write to the network drive.*

**33: You want to increase security on the Finance Database so that only members of the Finance group can run and execute applications. You create a Testlogin that is not a member of the Finance group, but it is able to query and get results. What should you do?**

*A: Deny permission for the Finance database to all NT groups except the Finance users.*

**34: You are installing SQL on a network where there are Windows 95, NT Workstations and NetWare clients. All must be able to connect. What would you do? (Choose all that apply)**

*A: Create NT logons for Windows NT clients  
Create NT logons for Windows 95 clients  
Create SQL logons for Windows NT clients  
Create SQL logons for Windows 95 clients  
Create SQL logons for Netware clients*

**35: The default option was used when Windows NT Server was installed onto 10 computers. You're planning unattended installation of SQL Server onto this 10 servers and you want to avoid unnecessary modifications to the server. The SQL Server database will be accessed by Windows 95 Client computers and by NetWare Client computers. What must you first install onto these 10 Windows NT Server computers. (choose all that apply)**

*A: Install Windows NT service pack 4  
Install Internet Explorer 4.01 with service pack 1  
Install NWLink IPX/SPX*

**36: Eric needs to manage user accounts on the HumanResources database. Which minimum role can you apply to him?**

*A: db\_accessadmin*

**37: You install SQL server on an IIS machine. You need to optimize IIS. What do you configure on the SQL Server?**

*A: Max server memory*

**38: If an application generates an error you want to receive notifications. How can the Developers of the application implement this?**

*A: RAISERROR Statement*

**39: You move a server from the Finance domain to the Money domain. SQL Server is not able to start. What is the problem?**

*A: SQL Server NT account is inaccessible.*

**40: SQL server is configured with a Net-library with TCP/IP socket 1450 which is a non-standard port setting. How would you configure the client computers?**

*A: At the Client, configure the SQL Server Client Network Utility to change the TCP/IP port to 1450.*

**41: Browsing of other servers from an SQL server running on Windows98 is not possible. The Browse list is not available. When users are trying to register they are not able to do it. However, they are able to connect for file sharing. What is the problem?**

*A: Browsing is not supported in Windows98.*

**42: You have the following Network Diagram:**

**NY ( NY1----T1-----NY2) -----56kbps----->(Seattle-----T1----San Francisco----T1---Los Angeles)**

**How would you set up replication?**

*A: NY1 =pubs/dist SF=subs/dist/pubs, Seattle, NY2, LA as Subs.*

**43: Paolo is removed from a the MKTG group where he has created a table in the Sales database. He is put in a new group called Manufacturing which uses the Products database. How can you let Paolo use the table that he created in the MKTG group?**

*A: Paolo cannot be removed from the MKTG group because he is the owner of the Table.*

**44: Three users Anita, Carmen and Philippe all have db\_owner roles currently in the Finance database. You want Philippe to continue to manage the Database as before. Anita should ONLY be able to create users accounts and grant permissions for users. Carmen should be ONLY able to create tables, views, indexes and other objects. What action would you take?**

*A: Remove Carmen from db\_owner and add to db\_ddladmin role.  
Remove Anita from db\_owner role and add to db\_securityadmin and db\_accessadmin roles.*

**45: You want to change the sort order of the company's Customer database. How?**

*A: Run Rebuildm.exe*

**46: You want to export information from your database into a comma-separated format text file named output.csv. What command line do you choose?**

*A: bcp db..developer out c:\data\txt\output.csv -c -t, -Usa -Ppasswd*

**47: You are attempting to use the bcp utility to copy data from the #tmpcust temporary table to the tmpCust.txt file. The name of the server is Svr and the password is Pss0. You type the following instruction at the command line:**

**bcp tempdb..#tmpcust out tmpcust.txt -n -SSvr -Usa -Pss0**

**Your command fails. What is the most likely cause of the failure?**

*A: Temporary table must be global to be copied using bcp.*

**48: Arbor Shoes has 100 stores. Their Pricing database resides on the corporate server. The pricing database server has limited space. The pricing database contains a price table which has a high frequency of updates. The in-store checkout application requires an up-to-date local price table in order to ensure accurate pricing. Which type of SQL Server replication should you use to replication the price table to the in-store computer?**

*A: Transactional replication with a central Publisher, remote Distributor, and Multiple Subscribers*

**49: Your company has five call centers that are located worldwide. These call centers manage the reservations for an airline. The reservation application has a customer table on the corporate server. The customer table is read locally and can be updated locally but the updates must maintain consistency across all five locations. Which type of SQL Server replication should you implement?**

*A: Transaction replication with immediate-updating subscribers.*

**50: Your company uses SQL 4.2 but plans to upgrade to SQL7.0. How must you accomplish this migration?**

*A: Upgrade SQL4.2 to SQL 6.5 and then upgrade to SQL7.0*

**51: You are developing security strategy for a new installation of SQL. You want to minimize the administration required to login. What should you do?**

*A: Use Windows NT Authentication  
Add Windows NT groups as SQL Server logins.*

**52: You install SQL Server onto a Windows NT Server computer that is a member of the Workgroup SQLGROUP. Your user computers are also members of the SQLGROUP. None of the users can connect to the SQL Server by using SQL Server Query Analyzer even though they are using valid SQL Server login ID's and passwords. You use the SQL Server Client Network Utility to confirm that the default Net-Library is Named Pipes. What is the most likely cause of the problem?**

*A: The Server's Windows NT Guest account is disabled.*

**53: Your company's current applications use MSAccess. New applications will use SQL Server. The new application must be able to access data from the Access database. What must you do, so that SQL Server queries can access the Access database?**

*A: Register connection information for the Access database in SQL Server.*

**54: You install merge replication. You are using the default conflict resolution. The ContactName (CN) field is Carolyn Sweeney for customer 10 in a merge article on the Publisher and Subscriber. The value of CN for customer 10 on Publisher is changed to Amy Jones. The value of CN for customer 10 on Subscriber is changed to Jenny Sax. The merge agent runs. What is the value of CN for customer 10 on Publisher and Subscriber after conflict resolution?**

*A: Amy Jones on both Publisher and Subscriber.*

**55: In order to tighten Security on a SQL Server, you remove NT Admin group from the default membership in the sysadmin. You then create a Local group called SQLadmin and add to sysadmin role. You add your NT account to SQLadmins. The SQLServerAgent service will not start and scheduled jobs will not run. Why?**

*A: SQLServerAgent service does not have the required NT administrative privileges.*

**56: You want to administer the 6.5 and 7.0 SQL Servers from your NT Workstation. How?**

*A: Install SQL Server 7.0 and 6.5 administrative tools and client connectivity components onto your workstation.*

**57: You change the memory settings for SQL Server. Now SQL Server will not start. You want to restore SQL Server's previous functionality, and you want to minimize downtime and the potential for loss of other server settings or objects. What should you do?**

*A: Start SQL Server at the command prompt with the appropriate parameters, and reconfigure the memory settings.*

**58: As the database owner you grant Franz permission to create views and to create stored procedures in the Finance database. Franz created a stored procedure that does an update on the dbo.prices table. He creates a view that selects price information for a report. He grants Susanne select permission on the view and execute permission on the stored procedure. What must be done so that Susanne can obtain the results by using the view and the stored procedure?**

*A: You must grant Susanne select and update permissions on the Prices table.*

**59: The total size of the column definitions for one row of one table in your database is 5000 bytes. You expect to have approximately 100,000 rows in the table when the data is fully loaded. The data portion of this table will consume approximately how much disk space?**

*A: 800 MB*

**60: The History database consists of 2 files. History\_data.mdf and history\_log.ldf. The History database has grown from its original size of 500 MB to 1 GB. The database is not configured to shrink automatically. You delete data so you have free space of 800 MB in the History database. Now you want to release 600 MB of the free database space. Which statement should you execute to reduce the physical size of the History\_data file?**

*A: DBCC SHRINKFILE (History\_data, 400)*

**61: Your company's network contains UNIX and NetWare systems in a NT Domain. Many users log in to these in order to access email and proprietary software. Currently only NT accounts are able to access the SQL Server. What do you do to give UNIX and NetWare users the same amount of access to SQL Server? (choose all that apply)**

*A: Create logins in SQL Server for UNIX and NetWare users.  
Add these users to the appropriate SQL Server roles.  
Ensure mixed mode security.*

**62: SQL 6.5 is upgraded to SQL 7.0 and some 6.5 stored procedures are showing constraints by collecting data from sysobjects, syscolumns, and sysconstraints. You can no longer run them. What should you do?**

*A: Rewrite the procedure to use the appropriate information schema view.*

**63: You recently set up transactional replication and you are now testing. You update 100 rows in the publisher table. You verify corresponding changes were made on the Subscriber table. However, the MSrepl\_commands table is empty. Why?**

*A: Replication distribution clean up job ran and purged it.*

**64: You have a database of 25GB in size and want to install it on a new PC with five disks (C, D, E, F & G) of 18.2GB each. You want high disk performance and you want fault tolerance. How would you set up this PC?**

*A: Create a mirrored set that uses drives F and G.  
 Create the transaction log file on the mirrored set.  
 Create a hardware RAID 5 set that uses drives C, D, and E.  
 Create the primary data file on the RAID 5 set.*

**65: You are responsible for backups on the Sales database. The Order and Customer tables are in the Sales database. The following events occur:**

**8:00 P.M.** Full database backup of the Sales database starts  
**8:10 P.M.** User begins a transaction  
 User inserts Order 53 for Customer 10 into the Order table.  
 The name for Customer 10 is Laura Jennings.  
**8:11 P.M.** Backup of the Sales database is completed.  
**8:12 P.M.** User updates the name for Customer 10 in the Customer table to Amy Jones.  
 User commits transaction.  
**8:30 P.M.** Sales database becomes corrupted due to media failure.  
**8:40 P.M.** The backup of the Sales database from 8:00 P.M. is restored to the Sales database.

**You query the Customer and Order tables in the restored Sales database. What is the state of the data?**

*A: Order 53 does not exist in the Order table  
 Customer 10 is Laura Jennings in the Customer table.*

**66: You want to install 10 Remote SQL Servers. The remote Sysadmin has no experience with SQL Server. You want to ensure that the installation of the remote SQL Server is compatible with existing SQL. What would you do?**

*A: Install the SQL Server on a similarly configured system, use the local system account and use the resulting sql70.iss file for automation of remote installation.*

**67: Stephen tries to connect to the Sales database by using Enterprise Manager. He receives the message “Unable to open default database<ID>”. You notice that neither Stephen nor the guest account was granted access to the Sales database. What must you do to give only Stephen access with the minimum permissions?**

*A: Make Stephen a user for the database.*

**68: The database for a check-processing application contains a record for each check that is processed. Each check has an associated image column that is accessed only when the user requests a check display. You plan to create the following files in the CHECK database.**

Drive	Files
C	check0.mdf
D	check1.ndf check2.ndf
E	check3.ndf
F	check0.ldf

**Which steps should you use as a strategy for placing table indexes?**

*A: Add the filegroup Image to the Check database.  
 Add check3.ndf to the filegroup Image.*

*Place the image column in the filegroup Image.  
Specify the filegroup Default for placement of all the tables and indexes.*

**69: You create a full backup and three log backups every day.**

**Full backup at 7am  
Log backup at 10am  
Log backup at 1pm  
Log backup at 3pm**

**Your database gets corrupted after one of the employees by mistake updates one of the columns in the database at 2:30 PM. How can you solve this problem?**

*A: Restore the full backup of 7am with NORECOVERY option  
Restore the Log backup of 10am with NORECOVERY option  
Restore the Log backup of 1pm with NORECOVERY option  
Restore the Log backup of 3pm with RECOVERY option, STOPAT 2:29pm*

**70: Laura is the new sales coordinator for your company and will process all orders. You want Laura to be able to modify the Order database only through the order applications. How should you control access to the order database?**

*A: Use a User-defined application role that has UPDATE permission on the Order database.*

**71: You need to provide access to the Sales database for Internet users who do not have Windows NT accounts. How can you provide them access? (Choose all that apply)**

*A: Add a guest user account to the Sales database, and grant the guest user account the appropriate permissions in the Sales database.  
Create a SQL Server login for internet guest access.*

**72: Finance will upgrade from SQL 6.5 to SQL7.0. SQL 6.5 is locally installed on several Windows98 clients with local copies of the database. To test a SQL7.0 application you want to build up a test system with the minimum administration. How would you proceed?**

*A: Upgrade to NT and install SQL7.0. Use the same install option as the current Finance Server. Use SQL Wizard to update the database.*

**73: Your server is running SQL Server 4.21 on NT Server 3.51. You plan to upgrade this server to SQL Server7.0. Which two steps must you take prior to the upgrade? (choose two)**

*A: Upgrade SQL Server to version 6.5  
Upgrade NT Server to version 4.0*

**74: You have a Domain with two SQL Server (SQL1/SQL2) with one database on each and standard logins. Both servers are able to connect to the each others databases and are able to execute Queries on both servers. You run a query that joins Data from SQL2 to Data on SQL1. SQL1 adds SQL2 as a linked server. However, during execute the login fails. Why?**

*A: You didn't map your login on SQL1 to your login on SQL2.*

**75: You are running a web site and you want an accounting of who is accessing specific tables. What would you monitor with the Performance Monitor?**

*A: Object:opened SQL user name NT user name.*

**76:** The departmental server has Full Text search implemented on it's SQL Server database. Currently both the Full Text database and the database files reside on the same logical drive. You do not want the Full Text indexing to use more than 25% of the available drive space because you want to allow space for dynamic expansion of the data files. You want to be able to remotely monitor the amount of space in use by the catalogs. How should you remotely monitor the space usage with a minimum number of additional administrative steps?

*A: Use NT Performance Monitor to connect to the server and monitor the size of the Microsoft Search catalog.*

**77:** You load a large amount of data into your database. You notice that the select into/bulk copy option is still enabled and that your transaction log is almost full. You turn off the select into/bulk copy option. Which 2 additional steps should you take to ensure the viability of future backups? (Choose 2)

*A: Make a full backup  
Backup the transaction log with truncate\_only*

**78:** The engine manufacturing application records data about all the engines that are manufactured at your plant. The database contains approximately 500 million records and is approximately 500GB in size. The database increases by 1,000 records per day. You want to accomplish the following goals:

- Minimize the time required to recover the database
- Provide the ability to recover the database to a specific time.
- Minimize the number of transaction logs that need to be used during recovery.
- Minimize the time required to back up the database.

You take the following actions:

- Schedule a full database back up of the entire database to occur every Sunday at 1:00 AM.
- Schedule transaction log back up to occur every day at 5:00 P.M.

Which result or results do these action produce? (Choose all the apply)

*A: The database can be recovered to a specific time.  
The time required to backup the database is minimized.*

**79:** You check the free space on a log file and run a BACKUP LOG WITH NO\_LOG on a database that is a publisher for transactional replication. You check the free space after running this command and note that free space has not changed. What is the most likely reason free space has not changed?

*A: Log reader agent has stopped*

**80:** You want to transfer a database copy from Server A to Server B using a full backup and restore. A complete backup uses 5 tapes. Server A has two tape drives and Server B has one tape drive. What is the easiest way to accomplish the database transfer?

*A: Backup using two tape drives on server A and restore sequentially using the one tape drive on server B.*

**81:** Katrin is the system administrator of SQL Server and she is in the role of db\_owner of the corporate database. Carlos is in the roles of setupadmin and db\_accessadmin. Julia is in the roles securityadmin and db\_securityadmin. Katrin must ensure that only she can grant permissions in the database and that only



**Julia can grant users and groups access to SQL Server and the corporate database. How must role membership be changed to meet these requirements? (Chose two)**

*A: Remove Carlos from db\_accessadmin role.  
Add Julia from db\_accessadmin role.*

**82: You have an order-entry application running under SQL Server. The application is used over the web to enter orders 24-hours a day. What kind of setup is best for this situation?**

*A: SQL Server with cluster support*

**83: You have 2 tables, T1 & T2, and 3 groups, G1, G2, & G3. G1 should be able to update T1, G2 should be only able to select from T1 and G3 should be able to select and update on T2. How many application roles do you need?**

*A: 3*

**84: Your Sales representatives use laptops. You need to implement a replication scenario so that the sales reps can enter orders. You want to minimize network traffic. What type of replication would you use?**

*A: Snapshot with pull*

**85: A developer has mistakenly issued a large bulk-insert query without any predicate. The log file has 3GB available. About 500MB of data has been loaded and the developer estimates that about 50% of data still is remaining to be loaded. Other users are complaining they cannot access the application because of the locks held by the query. You want to get the system to function as quickly as possible. What do you do?**

*A: Let the query run to completion.*

**86: You have the following backups for the Sales database on the Salesbak device:**

Name	Date	File Seq.	Type
Monday Data	7/7/99 8:00AM	1	Complete Database
Monday Log	7/7/99 6:00PM	2	Transaction Log
Tuesday Data	8/7/99 8:00AM	3	Complete Database
Tuesday Log	8/7/99 6:00PM	4	Transaction Log

**You are asked to restore the latest backup. Which command would you use?**

*A: RESTORE DATABASE sales FROM salesbak WITH FILE=3, NORECOVERY  
RESTORE LOG sales FROM salesbak WITH FILE=4, RECOVERY*

**87: To support a new transaction, you create an index on the largest table in the Sales database. It takes you 2 hours to create the index. You want to ensure the fastest possible reconstruction of the: index in the event of a database failure. What should you do immediately after Creating the index?**

*A: Full Backup of Sales database*

**88: You are about to run a large batch job. What should you do to allow for the quickest recovery?**

*A: Perform a full backup.*

**89: You are upgrading your SQL Server 6.5 to version 7 using the Version Upgrade Wizard. One stored procedure will not convert. This stored procedure references the sysobjects, syscolumns, and sysconstraints tables. What can you do to convert the stored procedure?**

*A: Compile the stored procedure in version 7*

**90: You have a server at headquarters and 2 servers in adjacent buildings. The servers are connected. You want to set up a trace to report SQL Server starts and stops on the remote servers and you want to view the trace on the headquarter server. How do you do it?**

*A: Setup a trace on each remote server and use an extended procedure to email the trace file to headquarters. Set up the extended procedure as a startup extended procedure.*

**91: You have a SQL Server that is also an IIS Server that processes online orders. It is imperative that the server be accessible 24 hours a day. You want to add a second server to guarantee uptime. What is the best strategy for adding the second server?**

*A: Set the second server up as failover server.*

**92: You need to setup SQL Server security for two applications. The applications are named Payroll and Sales. Two types of users in the Payroll application. Those who can modify data and those who can only query on existing data. There is only one type of users in the Sales application. Those who can both modify and query on existing data. What should you do to setup security access to the table?**

*A: Create two application roles for the Payroll application and grant permissions on each role for each type of access.*

*Create one application role for the Sales application and grant permissions to the role for the single type of access.*

**93: You are relocating a SQL server to a new region. You want to change the UNICODE SQL Setup on the SQL Database to use TCP/IP Socket 1450. How?**

*A: At the Client, use Client utilities.*

**94: Your company has 5 regional offices. From your main office you require a centralized view of data every two weeks. What replication scenario should you use?**

*A: Snapshot replication*

**95: Your company has many stores in the same city. The stores only need to update their own customer profiles. The customer profiles have to be update locally for performance reasons. The central office needs the customer profiles only occasionally. What replication scenario should you use?**

*A: Snapshot replication*

**96: What is used to enforce referential integrity? (Choose two)**

*A: Triggers  
Foreign Keys*

**97: You are using replication for 10 servers. How would you set up the administrative user?**

*A: Domain-level account which is a member of the 'Domain Admins' global group.*

**98: On a distribution server, which system file holds replication tasks?**

*A: MS\_JOB*

**99: Which system table on the destination database will have the ID of the last transaction replicated from 7 columns of the publisher database?**

*A: Mslast\_job\_info*

**100: You want to quickly determine how many page reads SQL Server is doing between two points in time. What commands could you use?**

*A: sp\_monitor  
DBCC SQLPERF[IOSTATS]*

**101: What DBCC CHECKTABLE statements will report on clustered indexes ONLY?**

*A: DBCC CHECKTABLE(products, NOINDEX)*

**102: Which replication tasks are scheduled automatically when setting up replication through Enterprise Manager?**

*A: Distribution, Publishing, & Sync*

**103: The default options were used when Microsoft Windows NT Server was installed onto 10 computers. You are planning unattended installations of SQL Server onto these 10 Servers and you want to avoid unnecessary modifications to the servers. The SQL Server databases will be accessed by Microsoft Windows 95 client computers. What must you first install on the 10 windows NT Server computers? (Choose all that apply)**

*A: Windows NT 4.0 Service Pack 4 or later  
Microsoft Internet Explorer 4.01 with Service Pack 1 or later*

**104: You implement full-text indexing on 10 tables in your Production database. You want to maintain the accuracy and performance of the full-text indexes, and you want to minimize overhead associated with their maintenance. What should you do?**

*A: Schedule regular re-population of the full-text indexes.*

**105: The size of your database has increased significantly and you want to examine performance statistics on specific database objects. The original SQL Server trace file was saved in a table. How should you replay the trace one step at a time?**

*A: In SQL Server Profiler, open the trace table, and replay the trace one step at a time.*

**106: You have a 30 GB Sales database that is used for inquiry and analysis during the day (9 AM to 4 PM). Each weekday night at 2 AM a batch process updates the Sales database with the previous day's orders. A full backup or restore of this database takes 4 hours. At a minimum, you want to have the capability to restore the database to show yesterday's business. What backup strategy will produce the easiest way to accomplish your restore objectives.**

*A: Differential backup each night after the batch process.*

**107: You configured the SQL Server service load parameter with a -f option in Enterprise Manager. Now, after restarting the SQL Server service, the Enterprise Manager cannot connect so you cannot change the load parameters. How do you fix this condition?**

*A: Manually delete the appropriate MSSQLServer key value in the registry.*

**108: The Windows NT servername under which the SQL Server was running is changed. Users are no longer able to connect to SQL Server. Why?**

*A: The SQL Server name does not match with the NT Server name.*

**109: The Microsoft Windows NT Server operating system running SQL Server fails with a blue screen error after a new tape drive adapter is installed. You suspect a hardware resource conflict, and you reinstall Windows NT Server. Now SQL Server will not start. You want to restore SQL Server functionality as quickly as possible, and you want to minimize the possibility of data loss. What should you do?**

*A: Run regrebld.exe.*

**110: You have a database which is updated frequently. You have set full-text search and want it to be always up to date. How should you update your full-text search?**

*A: With a job*

**111: You have setup snapshot replication between a publisher and multiple remote subscribers. You set up a distributor on a separate computer on your LAN. You suspect that the snapshot data has become corrupted and you want to clear it out manually. From where should you clear the data?**

*A: From the snapshot folder on the distributor.*

**112: The periodicals database is 10GB in size. You specify multiple tape-backup devices, and you schedule a full database backup to occur nightly. One night, the backup fails while writing to the third tape. You resolve the problem with the tape. What is the fastest way to complete the backup?**

*A: Manually back up the database and use the RESTART option.*

**113: You have 3 Tables (Table1, Table2, Table3). You also have 4 groups (group1a, group1b, group2a, group2b). All groups have select permissions on all tables. Group1b should be the only group that can change and delete data in table1. Group2b should be the only group that can change and delete data in table2 and table3. How would you set this up? (Choose two)**

*A: Grant group1b update and delete permissions on table1  
Grant group2b update and delete permissions on table2 and table3*

**114: A user is a member of a group that has had its permissions on a database denied. The user wants to have access to the database what should you do?**

*A: Remove him from the group.*

**115: The disk drive containing your Master database failed. Your User databases are unaffected. You cannot located a recent Master database backup, so you rebuild the Master database. What should you do to ensure that your databases are accessible in the shortest amount of time?**

*A: Attach the existing user databases to the new master database.*

**116: You enable merge replication. What must you do to enable conflict resolution on the articles in your publications?**

*A: Do nothing. SQL Server will add a column to the UNIQUEIDENTIFIER data type to each article.*

**117: Most of your database queries are executed as stored procedures. You need to find out whether most of the stored procedures are precompiled when they are executed. You also need to identify stored procedures that are not precompiled but should be. How should you monitor the performance of the stored procedures?**

*A: In Windows NT Performance monitor, monitor the Ad hoc SQL Plans and the Prepared SQL Plans instances of the SQLSvr:Cache Manager Object.*

**118: You are using DTS to load 500 MB of data from a flat file to a SQL Server table. Each row contains 6000 bytes. You have 100 MB of free disk space for the transaction log. Which is the important options to set on the target database to minimize possibility that the transaction log will not have any disk space available?**

*A: Truncate log on checkpoint*

**119: You are reviewing a new snapshot replication environment for 5 large tables. In order to free up disk space, you want to delete the data objects that are generated by snapshot replication. Where can you find these objects?**

*A: In the Mssql7\Repldata directory on the distributor.*

**120: You place the distribution database on the publication server for replication. You then decide to use the distribution database that is located on a remote server. What must you do to enable your publication to use the distribution database that is on the remote server?**

*A: Disable the current publisher and enable a new publisher to use the distributor database on the remote server. Recreate all publications and subscribers.*

**121: A franchise operation that has many stores in the same city wants every store to have access to customer profiles from all the other stores. A store can update only its own customer profiles. The central office will occasionally update the customer profiles. The customer profile table needs to be local to the application for performance reasons. Which SQL Server replication scenario should you implement?**

*A: Merge replication consisting of Multiple Publishers with Multiple Subscribers*

**122: Your database is replicated among four servers. You want to schedule database consistency checks and full database backups of all four copies of the database. You also want to minimize your administrative workload. Server configurations are shown in the following table.**

Server Name	Operation System	SQL Server
SQL7NT1	NT Server	7.0
SQL7NT2	NT Server	7.0
SQL65NT1	NT Server	6.5
SQL65NT1	NT Server	6.5

**How should you implement scheduled checks and backups of all four copies of the database so that administrative workload will be minimized?**

*A: Make SQL7NT1 a master server.  
Enlist SQLTNT2 as a target server.  
Create a job on the master server to check and back up the database on both SQL Server 7.0 servers.  
Separately implement scheduled database checks and backups on the other two servers.*

**123: You are running SQL-Server on Windows98. You cannot establish trusted connection. Why?**

*A: Windows98 doesn't support trusted connections.*

**124: You created procedures on your test system. Now you want to transfer these procedures to the Productserver. How?**

*A: Make a script creating these procedures and run this script on Productserver.*

**125: The Productgroup has access to a Database. Carmen is member of the Productgroup. You want to deny productgroup access to the Database, but Carmen should receive more than just guest access to this Database. How?**

*A: Remove Carmen from Productgroup and grant her user account access to the Database.*

**126: The Claims database contains more than 50 million rows and requires 100 GB of disk space. It uses five database files. Users in the Usage Analysis department use a variety of graphical query tools to query the Claims database every business day, between 8 A.M. and 4 P.M. There is a scheduled nightly batch job that runs at 2 A.M. to update the database with the previous day's claims. A full database backup or restore takes approximately six hours. You need to develop a backup strategy that can support recovery to the end of the previous business day and that Can recover from a failure in less than 12 hours. Which set of backups should you choose?**

*A: Full back up on Sunday  
differential back up every day after batch process runs.*

**127: Maria is a member of the Windows NT HumanResource group You need to deny database access to the HumanResourCe group, but you need to ensure that Maria has more than guest access. You deny database access to the HumanResource group by adding the group to the db-denyd~areader and db-denydatawriter roles. What must you do to give Maria the access she had before the HumanResource group was denied access?**

*A: Remove Maria from the HumanResource group, and add her as a user.*

**128: Your insurance company is purchased by Duluth Mutual Life The new management wants a copy of your 120 GB Policy database. Your current procedure is to back up the Policy database nightly to a media set of two device families, named PolicyBackup1 and PolicyBackup2. Each device family contains three tapes.**

**The server at Duluth Mutual Life is named Server1, There is adequate disk space on Served to store the database, but there is only one local tape drive. What should you do to restore the Policy backup to Server1 with minimal administrative workload?**

*A: Restore the tapes from the PolicyBackup1 and PolicyBackup2 media families by using the single tape drive on Server1. Load the tapes in order of sequence number.*

**129: You implement 10 scheduled jobs on your development server, and you verify that they function**

correctly. You now want to transfer the jobs to your production server. How should you transfer the jobs .with the fewest administrative steps?

*A: Script the jobs and execute the resulting script on the production server.*

**130: You administer a 30-GB database that experiences moderate update activity. Full database backups occur every Sunday at 1 A.M., and differential database backups occur at 1 A.M. every day except Sunday. In addition, transaction log backups occur hourly at 45 minutes past the hour. Your database fails on a Saturday afternoon. Which set of steps must you take to restore your data?**

*A: Restore the full database backup from Sunday.  
Restore the differential database backup from Saturday.  
Restore the transaction logs from 1:45 A.M. Saturday to the time of failure.*

**131: Stephen tries to connect to the Sales database by using SQL Server Enterprise Manager. He receives the message "Unable to open default database '<ID>'." You notice that neither Stephen nor the guest account was granted access to the Sales database. What must you do to give only Stephen access with the minimum permissions?**

*A: Add Stephen's login as a new user in the Sales database.*

**132: You want to configure SQL Server to notify you by email when an alert occurs. Which series of steps must you perform?**

*A: Configure the SQLServerAgent service to log in as a domain user account.  
Log on to the server running SQL as the SQL Server Agent account and configure email connectivity.  
Configure the SQLServerAgent to use the resulting email profile.*

**133: You want to schedule the Finances database so that backups occur according to the following three requirements:**

- 1) every 2 hours, Monday through Friday, 8 A.M. to 6 P.M.
- 2) every 4 hours, Monday through Friday, 8 P.M. to 4 A.M. the following day
- 3) every 8 hours, 6 A.M. Saturday to 10 P.M. Sunday

**You want to minimize the number of scheduled jobs. How would you schedule backups?**

*A: Create one single-step job to backup the database.  
Schedule the job to run according to the 3 requirements.*

**134: The backup schedule for the Accounting database is shown in the following table.**

1:00 AM	Full backup
7:00 AM	Transaction log backup
8:00 AM	Transaction log backup
9:00 AM	Transaction log backup

**At 8:45 AM on October 10, a user runs a new transaction that incorrectly updates data in several tables. The manager of the Accounting department tells you that all updates after that time can be reconstructed from the paper forms. You restore the 1:00 AM full backup. Which set of steps is the best method to use to restore the database to a consistent state closest to what the state was at 8:45 AM?**

*A: Restore the 7:00 AM transaction log backup by using the WITH NORECOVERY option.*

*Restore the 8:00 AM transaction log backup by using the WITH NORECOVERY option.  
Restore the 7:00 AM transaction log backup by using the WITH RECOVERY option and the  
STOPAT='OCT 10, 8:44 AM' option.*

**135: Katrin and Stefan have NT accounts. They are not members of the NT Administrators group. You need to give each of them database owner access to the Marketing database, but you want to ensure that you can identify the actions of each of them within the database.**

*A: Add their NT accounts to the db\_owner role in the database.*

**136: Your company has three servers running SQL Server in adjoining buildings. You are responsible for monitoring performance of stored procedures and for monitoring the connections to the servers. In order to do this, you want to forward all application information to the server in your building. You also want to automatically start monitoring in the event that SQL server shutdown and then restarts on the three servers. What must you do to meet these goals?**

*A: Use SQL Server Profiler extended stored procedures to forward a trace to the server in your building and make the stored procedures startup stored procedures.*

**137: You are planning the installation of SQL Server into servers at various international offices in your company. You want to accomplish the following goals:**

- 1) Minimize the time required to complete the installations.**
- 2) Minimize network utilization during the installations.**
- 3) Minimize the possibility of errors during the installations.**
- 4) Maximize the compatibility of the SQL Servers.**
- 5) Allow each international office to stored characters that are specific to its region.**

**Which course of action will best accomplish your goals?**

*A: Instruct the server administrators at each international office to install each SQL Server by using the installation option.*

**138: Joe uses a Windows 95 computer and sometimes a UNIX computer to connect to SQL Server 7.0. You have created a login for Joe's NT user account and also a SQL Server login that has the same name and password as his NT username and password. If Joe uses a trusted connection to access the SQL server what are his effective permissions.**

*A: NT overrides SQL Server*

**139: Mark configured replication on the SQL server in the accounting department to replicate information from the Payroll database to an SQL server in the Human Resources department. The transaction log of the Payroll database is filling up more frequently than it had before he configured replication. What is the MOST likely cause of this problem?**

*A: The Log Reader Agent.*



