

Microsoft 70-431

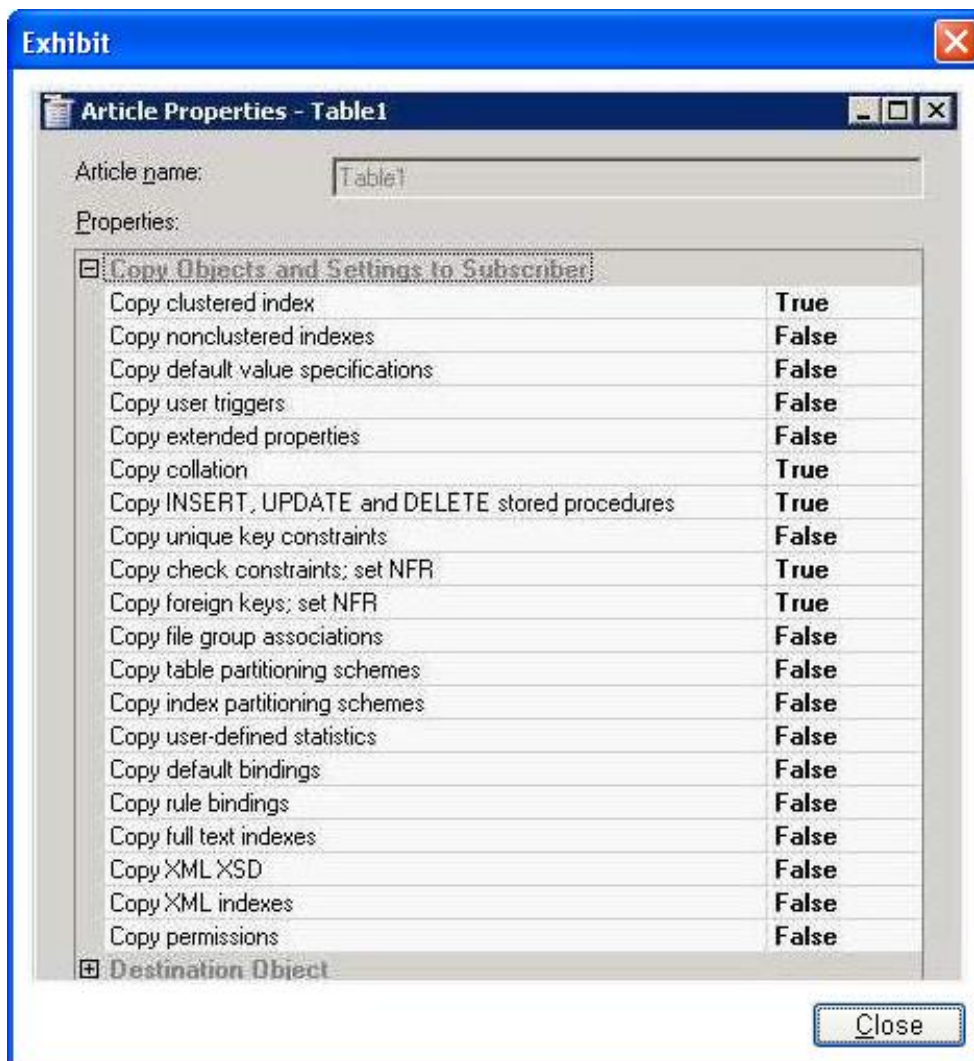
**Microsoft SQL Server 2005 Implementation &
Maintenance**
Version: 31.1

Topic 1, MAIN QUESTIONS

QUESTION NO: 1

Your company has two SQL Server 2005 computers named SQL1 and SQL2. A database named DB1 is located on SQL1. DB1 contains a table named Table1. Table1 is replicated to a database named DB1 Repl, which is located on SQL2. Full-Text Search is not being used.

Users report that the queries they run against Table1 in DB1 Repl are very slow. You investigate and discover that only the clustered index of Table1 is replicated. All other indexes in DB1 Repl are missing. You examine the Table1 article properties. The current Table1 article properties are shown in the exhibit. (Click the Exhibit button.)



You need to change the article properties so that all indexes of Table1 in DB1 are replicated when the subscription is reinitialized.

Which two article properties should you change? (Each correct answer presents part of the solution. Choose two.)

- A. Copy clustered index
- B. Copy nonclustered indexes
- C. Copy extended properties
- D. Copy unique key constraints
- E. Copy index partitioning schemes
- F. Copy XML indexes

Answer: B,F

Explanation:

QUESTION NO: 2

Your company has a server named SQL1 that runs SQL Server 2005 Enterprise Edition. SQL1 has 2 GB of RAM, 1.6 GB of which are used by the default SQL Server database engine instance. The average data growth of all databases combined is 100 MB a month. Users state that report execution times are increasing. You want to assess whether more RAM is needed. You need to use System Monitor to create a counter log that will help you decide whether to add RAM. Which performance object should you add to the counter log?

- A. MSAS200S: Cache
- B. MSAS 2005: Memory
- C. MSAS 2005: Proactive Caching
- D. SQL Server: Buffer Manager
- E. SQL Server: SQL Statistics
- F. SQL Server: General Statistics

Answer: D

Explanation:

The SQL Server:Buffer Manager object will show you

Low Buffer cache hit ration

Low Page life expectancy

High number of Checkpoint pages/sec

High number Lazy writes/sec

Insufficient memory and I/O overhead are usually related bottlenecks. SQL Server performance depends heavily on the I/O subsystem. Unless your database fits into physical memory, SQL Server constantly brings database pages in and out of the buffer pool. This generates substantial I/O traffic. Similarly, the log records need to be flushed to the disk before a transaction can be declared committed. And finally, SQL Server uses tempdb for various purposes such as to store intermediate results, to sort, to keep row versions and so on. So a good I/O subsystem is critical to the performance of SQL Server.

QUESTION NO: 3

You manage a SQL Server 2005 computer that was installed using default settings. After a power failure, the SQL Server (MSSQLSERVER) service on your database server does not start. You need to find out the cause of the problem. Which three actions should you perform? (Each correct answer presents part of the solution. Choose three.)

- A. In Event Viewer, view the system log.
- B. In Event Viewer, view the application log.
- C. In Notepad, view the C:\Program Files Microsoft SQL Server\MSSQL1\MSSQL\LOG>ErrorLog.1 file.
- D. In Notepad, view the C:\Program Files Microsoft SQL Server\MSSQL.1\MSSQL\LOG>ErrorLog.1 file.
- E. In Notepad, view the C:\Program Files Microsoft SQL Server\MSSQL1\MSSQL\LOG\SQLAgent.out file.

Answer: A,B,D

Explanation:

To resolve database errors, you should review the SQL Server error logs and the Windows Application and the Windows System Event Log for errors. Even though a variety of informational messages is logged to each of these locations, any error with a severity level of 16 or higher automatically gets logged to the SQL Server error log and the Windows Application Event log. Errors with a severity of 16 or higher are critical errors that you need to investigate immediately.

QUESTION NO: 4

A full backup of your database named DB1 is created automatically at midnight every day. Differential backups of DB1 occur twice each day at 10:00 and at 16:00. A database snapshot is created every day at noon. A developer reports that he accidentally dropped the Pricelist table in DB1 at 12:30. The last update to Pricelist occurred one week ago. You need to recover the Pricelist table. You want to achieve this goal by using the minimum amount of administrative effort. You must also minimize the amount of data that is lost. What should you do?

- A. Restore the most recent backup into a new database named DB1bak. Apply the most recent differential backup. Copy the Pricelist table from DB1bak to DB1.
- B. Delete all database snapshots except the most recent one.
- C. Restore DB1 from the most recent database snapshot.
- D. Recover DB1 from the most recent backup. Apply the most recent differential backup.
- E. Copy the Pricelist table from the most recent database snapshot into DB1.

Answer: D

Explanation:

QUESTION NO: 5

Your company uses SQL Server 2005. Users report with increasing frequency that they receive deadlock error messages in an order processing application. You need to monitor which objects and SQL Server session IDs are involved when deadlock conditions occur. You want information about each participant in the deadlock. What should you do?

- A. Trace the Lock: Timeout event by using SQL Server Profiler.
- B. Observe the SQL Server Locks - Number of Deadlocks/sec counter by using System Monitor.
- C. Trace the Lock: Deadlock event by using SQL Server Profiler.
- D. Trace the Lock Deadlock Chain event by using SQL Server Profiler.

Answer: D

Explanation:

Use SQL Server Profiler to identify the cause of a deadlock. A deadlock occurs when there is a cyclic dependency between two or more threads, or processes, for some set of resources within SQL Server. Using SQL Server Profiler, you can create a trace that records, replays, and displays deadlock events for analysis.

QUESTION NO: 6

Your company has a server named SQL1 that is dedicated to SQL Server 2005. SQL1 is configured for SQL Server and Windows Authentication mode. Using Windows Authentication, you have sysadmin privileges on SQL1. All users who work with SQL1 report that all applications that use SQL1 stop responding. Your supervisor asks you to investigate the problem. In Task Manager, you notice that CPU usage is at 100 percent. You try to connect to SQL1 by using SQL Server Management Studio. SQL1 does not respond. You need to connect to SQL1 to diagnose the problem. What are two possible ways to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. Connect to SQL1 by using the osql command-line utility.
- B. Connect to SQL1 by using the sqicmd command-line utility using a dedicated administrator connection.
- C. Connect to SQL1 by using SQL Server Management Studio. Open a query window and execute the following command. DBCC TRACEON 1224
- D. In SQL Server Management Studio, on the toolbar, click Database Engine Query to open the

Connect to Database dialog box. In the Server name text box, type ADMIN: SQL1. Select SQL Server Authentication and provide your sa credentials. Click the Connect button.

Answer: B,D

Explanation:

The DAC is implemented as a specific TCP endpoint in a SQL Server instance that is always attached to a dedicated UMS. The DAC was created to provide a connection that could always be used by a member of the sysadmin role to access a SQL Server instance, thereby guaranteeing that an administrator could not be locked out of SQL Server due to resource allocation issues. You can connect to the DAC in two ways:

Via SQLCMD or

Through the query window in SSMS.

QUESTION NO: 7

Your data warehouse includes three separate SQL Server 2005 tables for storing customer data. Each table is used to horizontally partition data that is migrated from a SQL Server 2000 installation to SQL Server 2005. Data is partitioned as shown in the following table.

Table name	Stores customer surname beginning with letter
Customer1	A - H
Customer2	I - P
Customer3	Q - Z

The database does not allow customers to opt-out of e-mail marketing offers.

You need to add a new OptOutFlag column to each of the customer tables. You must ensure that this column is added to either all three customer tables or none of them. You must also ensure that the database is always in a consistent state. Which Transact-SQL batch should you use?

A. BEGIN TRAN
 ALTER TABLE CUSTOMER1
 ADD OptOutFlag BIT NULL
 ALTER TABLE CUSTOMER2
 ADD OptOutFlag BIT NULL
 ALTER TABLE CUSTOMER3
 ADD OptOutFlag BIT NULL
 COMMIT TRAN

B. ALTER TABLE CUSTOMER1
ADD OptOutFlag BIT NULL
ALTER TABLE CUSTOMER2
ADD OptOutFlag BIT NULL
ALTER TABLE CUSTOMER3
ADD OptOutFlag BIT NULL

C. BEGIN TRAN
ALTER TABLE CUSTOMER1
ADD OptOutFlag BIT NULL
ALTER TABLE CUSTOMER2
ADD OptOutFlag BIT NULL
ALTER TABLE CUSTOMER3
ADD OptOutFlag BIT NULL
IF @@error <> 0ROLLBACK TRAN
ELSE
COMMIT TRAN

D. BEGIN TRAN
BEGIN TRY
ALTER TABLE CUSTOMER1
ADD OptOutFlag BIT NULL
ALTER TABLE CUSTOMER2
ADD OptOutFlag BIT NULL
ALTER TABLE CUSTOMER3
ADD OptOutFlag BIT NULL
END TRY
BEGIN CATCH
ROLLBACK TRAN
RETURN
END CATCH
COMMIT TRAN

Answer: D

Explanation: You should do this using the BEGIN CATCH. If there is an error reported in any of the tables when the OptOutFlag column is inserted will result that the whole transaction will rollback.

QUESTION NO: 8

You discover that the schema changes that were recently made to your SQL Server 2005 database have caused your Web site to stop functioning. It is unclear who made the changes. Your company now mandates that all changes to the database schema be tracked. You need to implement a mechanism that will track schema changes in your database. What should you do?

A. Implement a stored procedure that writes data about schema changes to a log table.

- B. Implement DDL AFTER triggers that write user and schema information to a log table.
- C. Implement a DML INSTEAD OF trigger that writes data about schema changes to a log table.
- D. Implement a DML AFTER trigger that writes data about schema changes to a log table.

Answer: B

Explanation:

A trigger is a specialized implementation of a Transact-SQL or CLR batch that automatically runs in response to an event within the database. You can create two types of triggers in TESTKING-DB04: data manipulation language (DML) triggers and data definition language (DDL) triggers. DDL triggers, which run in response to DDL events that occur on the server such as creating, altering, or dropping an object are used for database administration tasks such as auditing and controlling object access.

QUESTION NO: 9

Your company uses SQL Server 2005. Users report that report execution is slow. You investigate and discover that some queries do not use optimal execution plans. You also notice that some optimizer statistics are missing and others are out of date. You need to correct the problem so that reports execute more quickly. Which two Transact-SQL statements should you use? (Each correct answer presents part of the solution. Choose two.)

- A. DBCC CHECKTABLE
- B. ALTER INDEX REORGANIZE
- C. UPDATE STATISTICS
- D. CREATE STATISTICS
- E. DBCC SHOW_STATISTICS
- F. DBCC UPDATEUSAGE

Answer: C,D

Explanation:

QUESTION NO: 10

You are writing a query that returns data from tables in a SQL Server 2005 database. Some of the data is stored in relational columns, and some of the data is stored in XML data type columns. Your query needs to return a relational result set that contains data from relational fields and attribute values from XML data type columns. Which two methods of the XML data type should you use? (Each correct answer presents part of the solution. Choose two.)

- A. the value () method
- B. the exist () method

- C. the query () method
- D. the nodes () method
- E. the modify () method

Answer: A,D

Explanation:

QUESTION NO: 11

Your company uses SQL Server 2005. A user reports that an order processing application stopped responding in the middle of an order transaction. The user's SQL Server session ID is 54. You need to find out if session 54 is blocked by another connection. If it is, you need to identify the blocking session ID. What are two possible ways to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. In SQL Server Management Studio, open the Activity Monitor Open the Process Info page. View the BlockedBy column for session 54.
- B. In SQL Server Management Studio, open the Activity Monitor. Open the Locks by Process page. View the Request Mode column for session 54.
- C. In SQL Server Management Studio, open a new query window and execute the following statement. `SELECT * FROM sys.dm_exec_requests WHERE session_id = 54` View the blocking_session_id column.
- D. In SQL Server Management Studio, open a new query window and execute the following statement. `SELECT * FROM sys.dm_exec_sessions WHERE session_id = 54` View the status column.

Answer: A,C

Explanation:

The Process Info page shows the properties of current processes. Blocked By tells you the Process ID (SPID) of a blocking process. Blocking indicates whether this process is blocking others. 1=yes; 0=no.

The sys.dm_exec_request DMV can be used to determine whether a process is being blocked and to identify the process that is creating the blocking. If a value greater than 0 exists in the blocking_process_id column, the process is being blocked by the SPID logged in the column.

QUESTION NO: 12

You use a SQL Server 2005 database named DB1, which is located on a server named SQL1. DB1 is in use 24 hours a day, 7 days a week. A recent copy of DB1 exists on a second server named SQLtest that also runs SQL Server 2005. You detect a high number of full scans on SQL1

and conclude that additional indexes in DB1 are needed. A workload file that is suitable for Database Engine Tuning Advisor (DTA) already exists. You need to analyze the workload file by using DTA. You must ensure maximum performance on SQL1 during analysis. You must also ensure availability during the implementation of any recommendations suggested by the DTA. What should you do?

- A.** Store the workload file on SQL1. Start DTA on SQLtest and connect to SQL1. Specify all workload and tuning options as necessary. In the Advanced Tuning Options dialog box, select the Generate only online recommendations check box.
- B.** Store the workload file on SQLtest. Start DTA on SQLtest and connect to SQLtest. Specify all workload and tuning options as necessary. In the Advanced Tuning Options dialog box, select the Generate only online recommendations check box.
- C.** Store the workload file on SQL1. Start DTA on SQL1 and connect to SQL1. Specify all workload and tuning options as necessary. In the Advanced Tuning Options dialog box, select the All recommendations are online check box.
- D.** Store the workload file on SQLtest. Start DTA on SQLtest and connect to SQLtest. Specify all workload and tuning options as necessary. In the Advanced Tuning Options dialog box, select the All recommendations are online check box.

Answer: B

Explanation:

One of the requirements is to ensure maximum performance on SQL1 during analysis, therefore storing the workload file on SQL1 when running DTA is out of the question. You only want suggestions from the DTA that can be implemented while the database is online and therefore you should choose the Generate only online recommendations.

QUESTION NO: 13

You are responsible for implementing maintenance jobs on a SQL Server 2005 database server. Certain jobs run every Sunday and other jobs run at the beginning of every month. You need to schedule the jobs in the way that uses the least amount of administrative effort. What should you do?

- A.** Create a job schedule that runs every Sunday. Assign weekly tasks to this schedule. Create a second schedule that runs on the first day of every month. Assign monthly tasks to this schedule.
- B.** Create a job for each task that runs once a day. Use a Transact-SQL statement to check the date and day of the week. If the day is either a Sunday or the first day of the month, execute the code.
- C.** Create a job schedule that runs once a day. Assign jobs to this job schedule. If the day is either a Sunday or the first day of the month, execute the jobs.
- D.** Create a job for each task that runs once a week on Sunday. Add a second job schedule that runs the job on the first of the month.