



Microsoft 70-431

Microsoft SQL Server 2005 Implementation &  
Maintenance

Q&A with explanations

Version 7.1

**Leading The Way**  
in IT Testing And Certification Tools

**[www.testking.com](http://www.testking.com)**

## **Important Note, Please Read Carefully**

### **Other TestKing products**

A) Offline Testing engine

Use the offline Testing engine product to practice the questions in an exam environment.

B) Study Guide (not available for all exams)

Build a foundation of knowledge which will be useful also after passing the exam.

### **Latest Version**

We are constantly reviewing our products. New material is added and old material is revised. Free updates are available for 90 days after the purchase. You should check your member zone at TestKing and update 3-4 days before the scheduled exam date.

Here is the procedure to get the latest version:

1. Go to [www.testking.com](http://www.testking.com)
2. Click on **Member zone/Log in**
3. The latest versions of all purchased products are downloadable from here. Just click the links.

For most updates, it is enough just to print the new questions at the end of the new version, not the whole document.

### **Feedback**

If you spot a possible improvement then please let us know. We are always interested in improving product quality.

Feedback should be sent to [feedback@testking.com](mailto:feedback@testking.com). You should include the following: Exam number, version, page number, question number, and your login ID.

Our experts will answer your mail promptly.

### **Copyright**

Each iPad file contains a unique serial number associated with your particular name and contact information for security purposes. So if we find out that a particular iPad file is being distributed by you, TestKing reserves the right to take legal action against you according to the International Copyright Laws.

## **Table of contents**

Topic 1, Main Questions (72 Questions)	3
Topic 2, Simulations (12 Questions)	55

**Total number of questions: 85 Questions**

## Topic 1, MainQuestions (73 Questions)

### QUESTION NO: 1

Your application must access data that is located on two SQL Server 2005 computers. One of these servers is named SQL1 and the other is SQL2. You have permissions to create a stored procedure on SQL1 to support your application. However, on SQL2 you only have permissions to select data.

You write the stored procedure on SQL1. The stored procedure accesses SQL2 by using the OPENQUERY Transact-SQL statement. However, the query fails when executed. You need to troubleshoot the cause of the error. What should you do?

- A. Join the two servers by using the four-part syntax of server.database.schema.table.
- B. Reference SQL2 by using an alias.
- C. Add SQL2 as a remote server to SQL1.
- D. Add SQL2 as a linked server to SQL1.

**Answer: D**

**Explanation: You need to define a linked server for each external data source you want to access and the configure the security context under which your distributed queries will run. After you create a linked server, you can use the Transact-SQL OPENQUERY function to execute your distributed queries.**

### QUESTION NO: 2

You are preparing for a new installation of SQL Server 2005. You need to select the protocols that client computers might use to connect to the server.

Which two protocols can you use to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. Named Pipes
- B. TCP/IP
- C. Shared Memory
- D. Virtual Interface Adapter (VIA)
- E. Multiprotocol

**Answer: A,B**

**Explanation:** By default, clients have TCP and Named Pipes as available protocols. You can manipulate the protocol ordering by using the SQL Server Client utility. The client application uses the protocols in the order specified on the client computer. If you are using SQL Server 2005, the protocol order is stored in the ProtocolOrder registry entry under the following registry subkey:  
**HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\MSSQLServer\Client\SNI9.0**

### **QUESTION NO: 3**

You configure a new SQL Server 2005 computer to use TCP/IP with all default settings. Your corporate policy requires that each server use a firewall. You find that you can connect to the SQL Server instance from the local computer. However, client computers cannot connect to the SQL Server instance.

You need to identify the most likely cause of the connection issues. What should you do first?

- A. Ensure that port 1433 is open in your firewall.
- B. Ensure that port 443 is open in your firewall.
- C. Ensure that client computers connect by using Shared Memory protocol.
- D. Ensure that the server is not paused.

**Answer: A**

**Explanation:** The default instance of the Database Engine uses port 1433, but that can be changed. Instances of SQL Server 2005 Express Edition, SQL Server Mobile, and named instances of the Database Engine use dynamic ports. The SQL Server Browser service lets users connect to instances of the Database Engine that are not listening on port 1433, without knowing the port number. To use SQL Server Browser, you must open UDP port 1434.

### **QUESTION NO: 4**

TestKing.com has multiple servers in a distributed environment. You work with two SQL Server 2005 computers named SQL1 and SQL2. Each server uses SQL Server Authentication and they use different logins.

You need to write a distributed query that joins the data on SQL1 with the data on SQL2. What should you do?

- A. Ensure that both SQL1 and SQL2 use the same login name as the security context for each server.
- B. Configure SQL2 as a remote server. Write the query on SQL1.
- C. Configure SQL2 as a linked server to impersonate the remote login.
- D. Configure SQL2 as a distributed server. Use pass-through authentication.

**Answer: C**

**Explanation: When you use linked servers to access external data sources, you should pay special attention to the security context for the external connection. You can configure the linked server to use one of the following three security modes:**

Self-mapping - When a linked server is created, this mode is added for all local logins, so SQL Server tries to connect to the external data source using the current user's login credentials. The same login and password must exist on the remote server. This is the default behaviour.

Delegation - This mode impersonates the Windows local credentials; the connection forwards the credentials of an authenticated Windows user to the linked server. The Windows user account and password must exist on the linked server.

Remote Credentials - This mode lets you map local logins to remote logins on the external data source.

Delegation of operating system logins is the securest mechanism.

#### **QUESTION NO: 5**

TestKing.com 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: When a DBA creates an index, the query optimizer stores statistical information about the indexed columns. Additionally, CREATE STATISTICS, will create statistics on columns that are not contained in indexes but that are used in query predicates. UPDATE STATISTICS will optimize query performance.**

#### **QUESTION NO: 6**

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.

**Answer: A**

**Explanation: You need 2 schedules or both tasks will be executed on Sundays and on the first day of the month. Creating a job schedule for each task that runs once a week instead of creating one schedule containing all the tasks would use more administrative effort than needed.**

#### **QUESTION NO: 7**

You discover that the msdb database on a SQL Server 2005 computer is corrupt and must be restored. Databases are backed up daily. The database backup files are written to a network share, but the file names do not clearly indicate which databases are in each file. You need to locate the correct backup file as quickly as possible. The first file in the list is named DB\_Backup.bak. Which Transact-SQL statement should you use?

- A. RESTORE LABELONLYFROM DISK = N\\Server1\Backup\DB\_Backup.bak
- B. RESTORE HEADERONLYFROM DISK = N\\Server1\Backup\DB\_Backup.bak
- C. RESTORE VERIFYONLYFROM DISK = N\\Server1\Backup\DB\_Backup.bak
- D. RESTORE DATABASE MSDBFROM DISK = N\\Server1\Backup\DB\_Backup.bak

**Answer: B**

**Explanation: VERIFYONLY verifies the backup but does not restore it, and checks to see that the backup set is complete and the entire backup is readable. However, RESTORE VERIFYONLY does not attempt to verify the structure of the data contained in the backup volumes.**

Syntax:

```
RESTORE VERIFYONLY
```

```

FROM <backup_device> [ ,...n ]
[ WITH
[ { CHECKSUM | NO_CHECKSUM } ]
[ [ , ] { CONTINUE_AFTER_ERROR | STOP_ON_ERROR } ]
[ [ , ] FILE =backup_set_file_number ]
[ [ , ] LOADHISTORY ]
[ [ , ] MEDIANAME = { media_name | @media_name_variable } ]
[ [ , ] MEDIAPASSWORD = { mediapassword |
@mediapassword_variable } ]
[ [ , ] MOVE 'logical_file_name' TO 'operating_system_file_name' ]
[ ,...n ]
[ [ , ] PASSWORD = { password | @password_variable } ]
[ [ , ] { REWIND | NOREWIND } ]
[ [ , ] STATS [ = percentage ] ]
[ [ , ] { UNLOAD | NOUNLOAD } ]
]
[;]

```

```

<backup_device> ::=
{
{ logical_backup_device_name |
@logical_backup_device_name_var }
| { DISK | TAPE } = { 'physical_backup_device_name' |
@physical_backup_device_name_var }
}

```

### QUESTION NO: 8

A support engineer reports that inserting new sales transactions in a SQL Server 2005 database results in an error. You investigate the error. You discover that in one of the databases, a developer has accidentally deleted some data in a table that is critical for transaction processing. The database uses the full recovery model. You need to restore the table.

You need to achieve this goal without affecting the availability of other data in the database. What should you do?

- A. Back up the current transaction log. Restore the database with a different name and stop at the point just before the data loss. Copy the table back into the original database.
- B. Back up the current transaction log. Restore the database to the point just before the data loss.
- C. Restore the database from the existing backup files to a time just before the data loss.
- D. Restore the database to the point of the last full backup.



**Answer: A**

**Explanation:** You only want to restore the table that has been accidentally altered. All other suggestions to solution will restore the complete database and therefore you will lose data that may have been written to other tables after the point in time where you stop the recovery.

**QUESTION NO: 9**

A power failure occurs on the storage area network (SAN) where your SQL Server 2005 database server is located.

You need to check the allocation as well as the structural and logical integrity of all databases, including their system catalogs. What should you do?

- A. Execute DBCC CHECKFILEGROUP for each filegroup.
- B. Execute DBCC CHECKCATALOG.
- C. Execute DBCC CHECKDB.
- D. Execute DBCC CHECKTABLE for each table.

**Answer: C**

**Explanation:** DBCC CHECKDB will check the logical and physical integrity of all the objects in the specified database by:

- Running DBCC CHECKALLOC on the database.
- Running DBCC CHECKTABLE on every table and view in the database.
- Running DBCC CHECKCATALOG on the database.
- Validating the contents of every indexed view in the database.
- Validating the Service Broker data in the database

**QUESTION NO: 10**

You are responsible for importing data into SQL Server 2005 databases. Your department is starting to receive text files that contain sales transactions from stores across the country. Columns in the data are separated by semicolons.

You need to import the files into the sales database. What should you do?

- A. Create a custom format file, specifying a semicolon as the row terminator.
- B. Use the bcp command, specifying a semicolon as the field terminator.
- C. Use the bcp command with the default arguments.
- D. Use the BULK INSERT statement with the default arguments.

**Answer: B**

**Explanation:** The bcp utility bulk copies data between an instance of Microsoft SQL Server 2005 and a data file in a user-specified format. The bcp utility can be used to import large numbers of new rows into SQL Server tables or to export data out of tables into data files. To import data into a table, you must either use a format file created for that table or understand the structure of the table and the types of data that are valid for its columns.

**QUESTION NO: 11**

You are creating a Web-based application to manage data aggregation for reports. The application connects to a SQL Server 2005 database named DataManager. One page in the application has controls that execute stored procedures in a database named ReportingDatabase. There is an existing Service Broker connection between the DataManager database and ReportingDatabase.

You want to add two new message types to the existing service. In each database, you create message types named ProcessReport and SendResult. You need to add the two new message types to the existing service. What should you do first?

- A. Create a queue on each database with the ACTIVATION argument set to DataManager.dbo.ProcessReport.
- B. Create a conversation between the databases by using the following statement.  
BEGIN DIALOG FROM SERVICE 'ProcessReport' TO SERVICE 'SendResult'
- C. Create a contract between the services by using the following statement.  
CREATE CONTRACT ProcessData (ProcessReport SENT BY INITIATOR, SendResult SENT BY TARGET)
- D. Create services for each database by using the following statement.  
CREATE SERVICE DataManager ON QUEUE ProcessReport

**Answer: C**

**Explanation: A contract contains a list of message types and the services that are allowed to send them.**

Syntax: CREATE CONTRACT contract [AUTHORIZATION owner ] ( {message\_type\_name | [DEFAULT ] } SENT BY { INITIATOR | TARGET | ANY } } [ ,...n] ) [;]  
Key: contract Name of the contract owner Owner of the contract(database user or role) message\_type\_name Name of a message type to be included as part of the contract.

**QUESTION NO: 12**

You work at the regional sales office. You are responsible for importing and exporting data in SQL Server 2005 databases. The main office asks you to send them a text file that contains updated contact information for the customers in your region. The database administrator in the main office asks that the data be sorted by the StateProvince, Surname, and FirstName columns.

You need to satisfy these requirements by using the least amount of effort. What should you do?

- A. Specify StateProvince, Surname, and FirstName in the ORDER hint in the bcp out command.
- B. Create a format file for the export operation.
- C. Specify StateProvince, Surname, and FirstName in the ORDER BY clause in the bcp queryout command.
- D. Copy the data into a new table that has a clustered index on StateProvince, Surname, and FirstName. Export the data.

**Answer: C**

**Explanation: The bcp utility bulk copies data between an instance of Microsoft SQL Server 2005 and a data file in a user-specified format. The bcp utility can be used to import large numbers of new rows into SQL Server tables or to export data out of tables into data files. Except when used with the queryout option, the utility requires no knowledge of Transact-SQL. Queryout copies from a query and must be specified only when bulk copying data from a query.**

Syntax:

```
bcp
{[[database_name.][owner].]{table_name | view_name} | "query"}
{in | out | queryout | format} data_file
[-mmax_errors] [-ffformat_file] [-x] [-eerr_file]
[-Ffirst_row] [-Llast_row] [-bbatch_size]
[-n] [-c] [-N] [-w] [-V (60 | 65 | 70 | 80)] [-6]
[-q] [-C { ACP | OEM | RAW | code_page } ] [-tfield_term]
[-rrow_term] [-iinput_file] [-ooutput_file] [-apacket_size]
[-Sserver_name[\instance_name]] [-Ulogin_id] [-Ppassword]
[-T] [-v] [-R] [-k] [-E] [-h"hint [,...n]"
```

**QUESTION NO: 13**

TestKing.com has two SQL Server 2005 computers named SQL1 and SQL2. Both servers take part in replication. SQL1 is both the Publisher and its own Distributor of a publication named Pub1. Pub1 is the only publication on SQL1, and SQL2 is the only Subscriber. Your supervisor requests a status report about the replication latencies. Using Replication Monitor on SQL1, you need to find out the current latencies between the Publisher and Distributor as well as between the Distributor and Subscriber. What should you do?

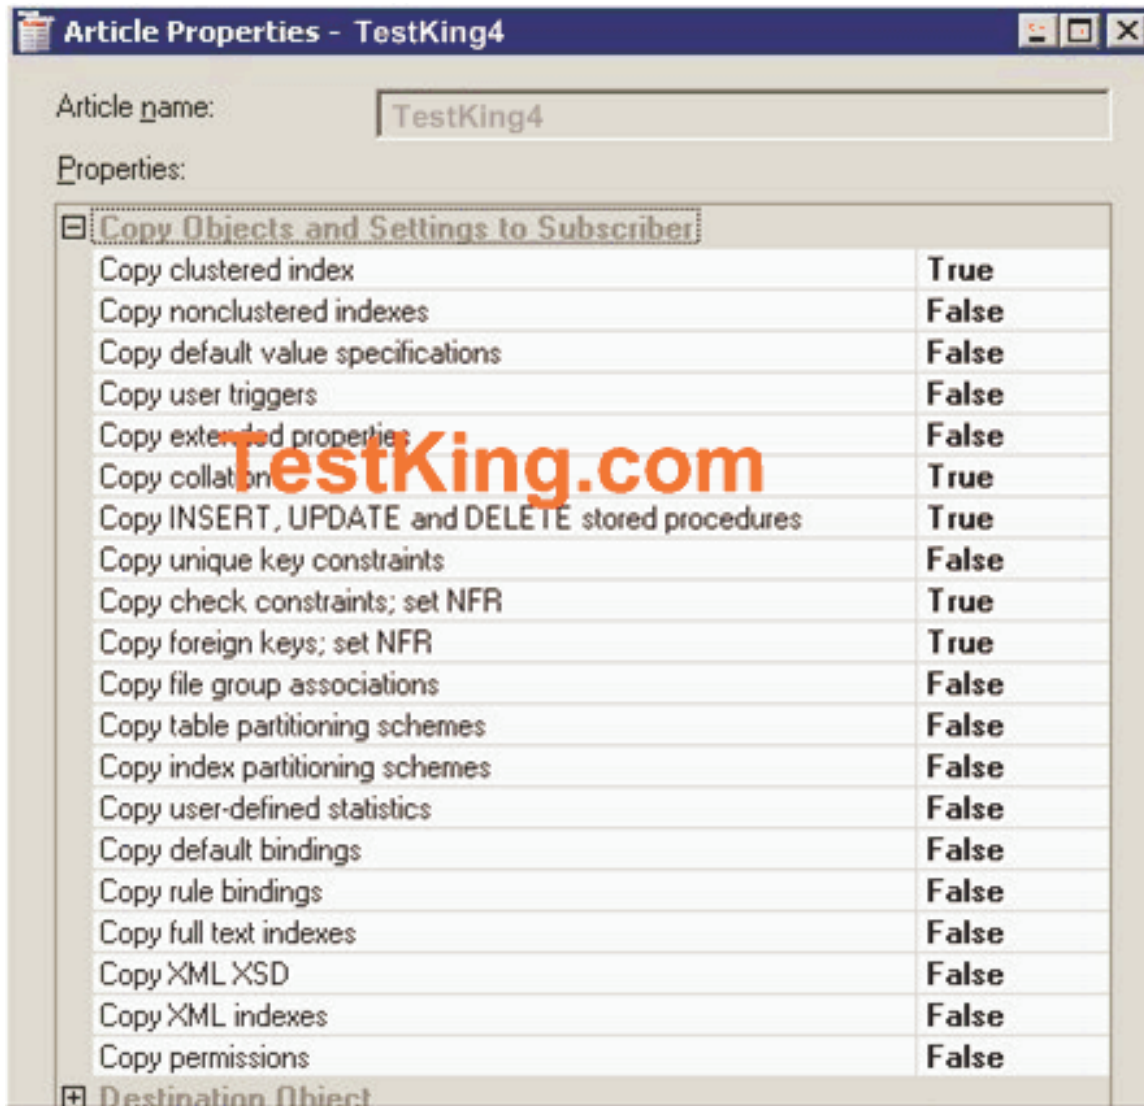
- A. Select the Subscription Watch List tab for SQL1. View the Latency column for the SQL2 subscription.
- B. Select the All Subscriptions tab for the Pub1 publication. View the Latency column for the SQL2 subscription.
- C. Select the Tracer Tokens tab for the Pub1 publication. Select the Insert Tracer option and wait for the requested latency values for the SQL2 subscription to appear.
- D. Select the Subscription Watch List tab for SQL1. Double-click the SQL2 subscription. View the duration details on the Publisher to Distributor History tab as well as on the Distributor to Subscriber History tab.

**Answer: C**

**Explanation: Tracer tokens are a new feature of the transactional replication engine in SQL Server 2005. They are a specialized transactions that is used to gain timing information. Since tracer tokens are picked up by the replication engine and transit with every other transaction, they can be used to determine the amount of time it takes a transaction to move from the publisher to distributor and then from the distributor to each subscriber. With this piece of information, you can now answer the two most predominant questions for replication: How far behind am I and how long will it take to catch up?**

**QUESTION NO: 14**

Exhibit:



TestKing.com has two SQL Server 2005 computers named SQL1 and SQL2. A database named DB1 is located on SQL1. DB1 contains a table named TestKing4. TestKing4 is replicated to a database named DB1Repl, which is located on SQL2. Full-Text Search is not being used. Users report that the queries they run against TestKing4 in DB1Repl are very slow. You investigate and discover that only the clustered index of TestKing4 is replicated. All other indexes in DB1Repl are missing. You examine the TestKing4 article properties. The current TestKing4 article properties are shown in the exhibit.

You need to change the article properties so that all indexes of TestKing4 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: 15**

You are creating an HTTP endpoint that will be used to provide customer data to external applications. Your SQL Server 2005 computer is named SQL1. You create a stored procedure named dbo.usp\_GetPersonData to retrieve the data in the TestKing database. You create the endpoint by using the following code.

```
CREATE ENDPOINT SQLEP_AWPersons AS HTTP (PATH = '/AWpersons',  
AUTHENTICATION = (INTEGRATED), PORTS = (CLEAR), SITE = 'SQL1') FOR  
SOAP (WEBSERVICE 'PersonData' (NAME='TestKing.dbo.usp_GetPersonData'),  
BATCHES = DISABLED, WSDL = DEFAULT, DATABASE = 'TestKing',  
NAMESPACE = 'http://Adventure-Works/Persons')
```

The first users to connect to the endpoint tell you that they do not get any data. You connect to the endpoint and discover that it is not responding. You need to modify the endpoint so that data is returned as expected. What should you do?

- A. Change the AUTHENTICATION property to KERBEROS.
- B. Specify BATCHES = ENABLED.
- C. Specify STATE = Started.
- D. Specify WSDL = 'pr\_GetPersonData'.

**Answer: C**

**Explanation: The possible states for an endpoint are STARTED, STOPPED, and DISABLED. For an endpoint to respond to requests, the state must be set to STARTED. To comply with the SQL Server 2005 "off by default" approach to security, the default state is STOPPED.**

#### **QUESTION NO: 16**

You work in Dublin at the main office of TestKing.com. You are responsible for managing a SQL Server 2005 database. The sales department wants a report that compares customer activity in the previous quarter between the main office in Dublin and the branch office in Buenos Aires. They want the data sorted by surname and first name. You restore a recent backup of the Buenos Aires database onto your server. You write queries to build the report, ordering the data by the Surname and FirstName columns. You review the data and notice that the customer list from the Buenos Aires database is sorted differently. The sales department needs the revised data within 15 minutes for a presentation.

You need to implement the fastest possible solution that ensures that the data from both databases is sorted identically. What should you do?

- A. Use the Copy Database Wizard to copy the data in the Buenos Aires database to a new database with the same collation as the Dublin database.
- B. Use the SQL Server Import and Export Wizard to copy the data from the Buenos Aires database into new tables, specifying the same collation as the Dublin database.
- C. Modify the format file to specify the same collation as the Dublin database. Import the table again.
- D. Modify the query on the Buenos Aires database to use the COLLATE setting in the ORDER BY clause. In the query, specify the same collation as the Dublin database.

**Answer: D**

**Explanation:**

COLLATE {collation\_name}

Specifies that the ORDER BY operation should be performed according to the collation specified in collation\_name, and not according to the collation of the column as defined in the table or view. collation\_name can be either a Windows collation name or a SQL collation name. COLLATE is applicable only for columns of the char, varchar, nchar, and nvarchar data types.

**QUESTION NO: 17**

You work for a company that sells books. You are creating a report for a SQL Server 2005 database. The report will list sales representatives and their total sales for the current month. The report must include only those sales representatives who met their sales quota for the current month. The monthly sales quota is \$2,000. The date parameters are passed in variables named @FromDate and @ToDate.

You need to create the report so that it meets these requirements. Which SQL query should you use?

A. SELECT s.AgentName, SUM(ISNULL(o.OrderTotal,0.00))AS SumOrderTotalFROM SalesAgent s JOIN OrderHeader o ON s.AgentID = o.AgentIDWHERE o.OrderDate BETWEEN @FromDate AND @ToDateGROUP BY s.AgentName

B. SELECT s.AgentName, SUM(ISNULL (o.OrderTotal,0.00))AS SumOrderTotalFROM SalesAgent s JOIN OrderHeader o ON s.AgentID = o.AgentIDWHERE o.OrderDate BETWEEN @FromDate AND @ToDate AND o.OrderTotal >= 2000GROUP BY s.AgentName

C. SELECT s.AgentName, SUM(ISNULL (o.OrderTotal,0.00)) AS SumOrderTotalFROM SalesAgent s JOIN OrderHeader o ON s.AgentID = o.AgentID WHERE o.OrderDate BETWEEN @FromDate AND @ToDateGROUP BY s.AgentNameHAVING SUM(o.OrderTotal) >= 2000

D. SELECT s.AgentName, SUM(ISNULL(o.OrderTotal,0.00)) AS SumOrderTotalFROM SalesAgent s JOIN OrderHeader o ON s.AgentID = o.AgentIDWHERE o.ordertotal = 2000 AND o.OrderDate BETWEEN @FromDate AND @ToDateGROUP BY s.AgentNameHAVING SUM(o.OrderTotal) >= 2000

**Answer: C**

**QUESTION NO: 18**

You are creating a stored procedure that will delete data from the Contact table in a SQL Server 2005 database. The stored procedure includes the following Transact-SQL statement to handle any errors that occur.

```
BEGIN TRY
BEGIN TRANSACTION
DELETE FROM Person.Contact
WHERE ContactID = @ContactID
COMMIT TRANSACTION
END TRY
BEGIN CATCH
DECLARE @ErrorMessage nvarchar(2000) DECLARE @ErrorSeverity int DECLARE
@ErrorState int SELECT @ErrorMessage = ERROR_MESSAGE(), @ErrorSeverity =
ERROR_SEVERITY(), @ErrorState = ERROR_STATE() RAISERROR
(@ErrorMessage, @ErrorSeverity, @ErrorState) END CATCH;
```

You test the stored procedure and discover that it leaves open transactions. You need to modify the stored procedure so that it properly handles the open transactions. What should you do?

- A. Add a COMMIT TRANSACTION command to the CATCH block.
- B. Remove the COMMIT TRANSACTION command from the TRY block.



- C. Add a ROLLBACK TRANSACTION command to the CATCH block.
- D. Add a ROLLBACK TRANSACTION command to the TRY block.

**Answer: C**

**Explanation: If an exception occurs from one of the statements in the TRY block, control is branched to the CATCH block, where the exception can be handled, logged, and so on. In this case you want to ROLLBACK the transaction if there is an error.**

#### **QUESTION NO: 19**

You are creating an online catalog application that will display product information on the company Web site. The product data is stored in a SQL Server 2005 database. The data is stored as relational data but must be passed to the application as an XML document by using FOR XML. You test your application and notice that not all of the items matching your query appear in the XML document. Only those products that have values for all elements in the schema appear.

You need to modify your Transact-SQL statement so that all products matching your query appear in the XML document. What should you do?

- A. Add an XML index to the table that contains the product data.
- B. Add the XSINIL argument to the ELEMENTS directive in the query.
- C. Add a HAVING clause to the query.
- D. Add the replace value of clause to the query.

**Answer: B**

**Explanation: If you add the XSINIL instruction to the ELEMENTS clause in the FOR XML construction, SQL Server 2005 generates an empty XML element for NULL values.**

#### **QUESTION NO: 20**

TestKing.com has two SQL Server 2005 computers named SQL1 and SQL2. Transaction log shipping occurs from SQL1 to SQL2 by using default SQL Server Agent schedule settings.

You need to reconfigure transaction log shipping to provide minimum latency on SQL2. What should you do?

- A. On SQL1, reschedule the transaction log backup job so that it occurs every minute. On SQL2, maintain default schedule settings for both the log shipping copy and the restore jobs.

- B. On SQL1, change the schedule type for the transaction log backup to Start automatically when SQL Server Agent starts. On SQL2, change the schedule types for both the log shipping copy and the restore jobs to Start automatically when SQL Server Agent starts.
- C. On SQL1, maintain default schedule settings for the transaction log backup job. On SQL2, change the schedule types for both the log shipping copy and the restore jobs to Start automatically when SQL Server Agent starts.
- D. On SQL1, reschedule the transaction log backup job so that it occurs every minute. On SQL2, reschedule both the log shipping copy and the restore jobs so that they occur every minute.

**Answer: D**

**Explanation: In order to minimize latency on SQL2 you need to get a copy from SQL1 as soon as possible after every write. The shortest possible interval for the transaction log backup is one minute. Once the backup is done you want to copy it and restore it on SQL2 as soon as possible.**

**QUESTION NO: 21**

You are implementing transaction log shipping for a database named DB1 from a server named SQL1 to a server named SQL2. Because DB1 is 100 GB in size, it is too big to transfer over the network in a reasonable amount of time.

You need to minimize the impact on the network while you initialize the secondary database. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Specify the simple recovery model for DB1.
- B. Specify either the full or the bulk-logged recovery model for DB1.
- C. Perform a complete backup of DB1 to portable media. Restore the secondary database from that backup; specify the RECOVERY option.
- D. Perform a complete backup of DB1 to portable media. Restore the secondary database from that backup; specify the STANDBY option.
- E. Before you activate transaction log shipping to the secondary database, execute the following statement on the primary server. `BACKUP LOG DB1 WITH TRUNCATE_ONLY`

**Answer: B,D**

**Explanation: Standby mode is the RESTORE command option you use to configure standby servers. In this mode, the secondary database is available for read-only access by users and applications. Log shipping requires that the primary database is configured with the Full or Bulk-Logged recovery model.**

**QUESTION NO: 22**

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. Restore DB1 from the most recent database snapshot.
- C. Recover DB1 from the most recent backup. Apply the most recent differential backup.
- D. Copy the Pricelist table from the most recent database snapshot into DB1.

**Answer: D**

**Explanation: SQL Server 2005 introduces the capacity for administrators to generate and use a read-only, stable view of a database. The database snapshot provides this capability without the overhead of creating a complete copy of the database or storage requirement. As the primary database diverges from the snapshot, the snapshot gets its own copy of original pages when they are modified. The snapshot may be used to recover an accidental change to a database by simply reapplying the pages from the snapshot to the primary database. As the last snapshot was created at noon the same day that the Pricelist table was dropped and the fact that there has been no changes to the table makes it possible to just copy the table from the snapshot.**

**QUESTION NO: 23**

You manage a database named DB1, which is located on a SQL Server 2005 computer. You receive a warning that the drive on which the DB1 log file is located is near capacity. Although the transaction log is backed up every five minutes, you observe that it is steadily growing. You think that an uncommitted transaction might be the cause and you want to investigate.

You need to identify both the server process ID and the start time of the oldest active transaction in DB1. What should you do?

- A. Connect to the DB1 database. Execute DBCC OPENTRAN. View the SPID and Start time rows.
- B. Connect to the master database. Execute DBCC OPENTRAN. View the SPID and Start time rows.

C. In SQL Server Management Studio, open the Activity Monitor. Select the Process Info page and apply the following filter settings. Database = DB1 Open Transactions = Yes View the Process ID and Last Batch columns.

D. Open a query window. Connect to the master database. Execute the following statement. `SELECT TOP 1 spid, last_batch FROM sys.sysprocesses WHERE dbid = db_id('DB1') AND open_tran > 0 ORDER BY last_batch`

**Answer: A**

**Explanation: The DBCC OPENTRAN command gives you the information you need by displaying information about the oldest active transaction and the oldest distributed and nondistributed replicated transactions.**

Syntax DBCC OPENTRAN

```
[ ( [ 'database' | database_id | 0 ] ) ]  
[ WITH TABLERESULTS ] [, [ NO_INFOMSGS ] ]  
]
```

#### **QUESTION NO: 24**

TestKing.com 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. MSAS 2005:Cache
- B. MSAS 2005:Memory
- C. MSAS 2005:Proactive Caching
- D. SQLServer:Buffer Manager
- E. SQLServer:SQL Statistics
- F. SQLServer: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. SQLServer performance depends heavily on the I/O subsystem. Unless your database fits into physical memory, SQLServer 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, SQLServer 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 SQLServer.

#### QUESTION NO: 25

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\MSSQL.1\MSSQL\LOG>ErrorLog.1 file.
- D. In Notepad, view the C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG>ErrorLog file.
- E. In Notepad, view the C:\Program Files\Microsoft SQL Server\MSSQL.1\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. Although 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: 26

You manage a SQL Server 2005 database that contains a table with many indexes. You notice that data modification performance has degraded over time. You suspect that some of the indexes are unused.

You need to identify which indexes were not used by any queries since the last time SQL Server 2005 started. Which dynamic management view should you use?

- A. sys.dm\_fts\_index\_population
- B. sys.dm\_exec\_query\_stats
- C. sys.dm\_db\_index\_usage\_stats
- D. sys.dm\_db\_index\_physical\_stats

**Answer: C**

**Explanation: The sys.dm\_db\_index\_usage\_stats dynamic management view contains core statistics about each index within a database. Use this view when you need to find the number of seeks, scans, lookups, or updates that have occurred with an index. You can determine the last time an index was used by examining the last\_user\_seek, last\_user\_scan and last\_user\_lookup columns.**

#### **QUESTION NO: 27**

TestKing.com uses SQL Server 2005. A user reports that an order processing application stopped responding in the middle of an order transaction. The users 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.**

To determine whether a process is being blocked and to identify the process that is creating the blocking, you could also use the sys.dm\_exec\_request DMV. 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: 28**

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 offline 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 offline 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: 29**

TestKing.com 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?

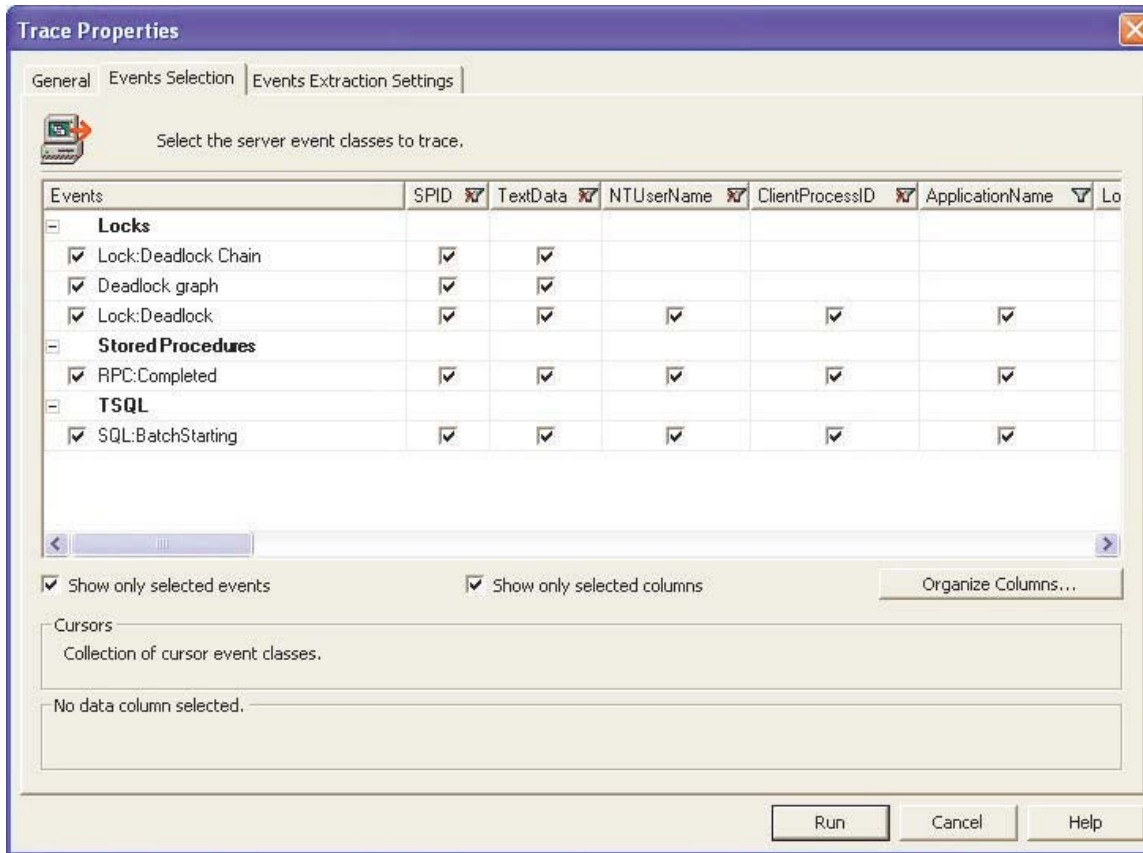
- A. Trace the Lock:Timeout event by using SQL Server Profiler.
- B. Observe the SQLServer:Locks - Number of Deadlocks/sec counter by using System Monitor.

- 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: 30 HOTSPOT

You are working as a DBA at the Cape Town office of TestKing.com. TestKing.com use a SQL Server 2005 database that does not contain any views.

You use Database Engine Tuning Advisor (DTA) to tune this database. A workload file that is suitable for DTA already exists.

You are required to locate only missing nonclustered indexes. During this process, you need to insure that existing structures remain intact, and that newly recommend structures are partitioned for best performance.

You want to accomplish this goal by configuring the tuning options in DTA.

Which tuning options should you use?

Physical Design Structures (PDS) to use in database

Indexes and indexed views

Indexes

Evaluate utilization of existing PDS only

Indexed views

Nonclustered indexes

Partitioning strategy to employ

No partitioning

Aligned partitioning

Full partitioning

Physical Design Structures (PDS) to keep in database

Do not keep any existing PDS

Keep all existing PDS

Keep aligned partitioning

Keep indexes only

Keep clustered indexes only

**Answer:**

**Explanation:**

Physical Design Structures (PDS) to use in database

Indexes and indexed views

Indexes

Evaluate utilization of existing PDS only

Indexed views

Nonclustered indexes

Partitioning strategy to employ

No partitioning

Aligned partitioning

Full partitioning

Physical Design Structure (PDS) to keep in database

Do not keep any existing PDS

Keep all existing PDS

Keep aligned partitioning

Keep indexes only

Keep clustered indexes only

For us the correct answer should be :

Nonclustered Indexes  
FullPartitioning  
Keep all existing PDS

#### On BOL explanations about DTA options:

<b>Indexes</b>	Causes Database Engine Tuning Advisor to consider adding only indexes.
<b>Indexed views</b>	Causes Database Engine Tuning Advisor to consider adding only indexed views.
<b>Nonclustered indexes</b>	Causes Database Engine Tuning Advisor to consider adding only nonclustered indexes.
<b>Evaluate utilization of existing PDS only</b>	Causes Database Engine Tuning Advisor to evaluate the existing physical design structures only, but not to consider adding any new structures.

#### Partitioning Strategy to Employ

Tuning Option	Description
No partitioning	Causes Database Engine Tuning Advisor not to use partitioning in its recommendation.
Full partitioning	Causes Database Engine Tuning Advisor to consider using partitioning in its recommendation.
Aligned partitioning	Causes Database Engine Tuning Advisor to only recommend partitions that are aligned with existing partitions of underlying table or view (enhances manageability).

### Physical Design Structures to Keep in the Database

Tuning Option	Description
Do not keep existing PDS	Causes Database Engine Tuning Advisor to consider dropping any existing physical design structure. Select the Evaluate utilization of existing PDS option with this option to set Database Engine Tuning Advisor to the drop-only mode, which only evaluates current physical design structures to determine whether they can be dropped.
Keep indexes only	Causes Database Engine Tuning Advisor to consider dropping all structures except indexes.
Keep all existing PDS	Causes Database Engine Tuning Advisor to keep all existing structures.
Keep clustered indexes only	Causes Database Engine Tuning Advisor to consider dropping all structures except clustered indexes.
Keep aligned partitioning	Causes Database Engine Tuning Advisor to consider dropping all structures except aligned partitions.

#### QUESTION NO: 31

You work for a bank that uses a SQL Server 2005 database to store line items from customer banking transactions. The bank processes 50,000 transactions every day. The application requires a clustered index on the TransactionID column. You need to create a table that supports an efficient reporting solution that queries the transactions by date. What are the two ways to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. Place a nonclustered index on the date column.
- B. Add a unique clustered index on the date column.

- C. Map each partition to a filegroup, with each filegroup accessing a different physical drive.
- D. Create a partitioning scheme that partitions the data by date.

**Answer: A,D**

**Explanation: After you build a clustered index, you can create nonclustered indexes on the table. In contrast with a clustered index, a nonclustered index does not force a sort order on the data in a table. In addition, you can create multiple nonclustered indexes to most efficiently return results based on the most common queries you execute against the table. By using partitions, you can place a subset of a table or index on a designated filegroup. This capability lets you separate specific pieces of a table or index onto individual filegroups and effectively manage file input/output (I/O) for volatile tables. Additionally, as organizations collect more and more data and keep it longer and longer for analysis purposes, tables continue to grow larger and larger. Managing such massive tables can be difficult. With partitioning, however, you can segregate data within a table based on age.**

#### **QUESTION NO: 32**

TestKing.com uses a SQL Server 2005 database. This database contains a trigger named trg\_InsertOrders, which fires when order data is inserted into the Orders table. The trigger is responsible for ensuring that a customer exists in the Customers table before data is inserted into the Orders table. You need to configure the trigger to prevent it from firing during the data import process. You must accomplish this goal while using the least amount of administrative effort. Which two Transact-SQL statements can you use to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. ALTER TABLE Orders DISABLE TRIGGER trg\_InsertOrders
- B. DROP TRIGGER trg\_InsertOrders
- C. DISABLE TRIGGER trg\_InsertOrders ON Orders
- D. ALTER TRIGGER trg\_InsertOrders ON Orders NOT FOR REPLICATION
- E. sp\_settriggerorder@triggername= 'trg\_InsertOrders', @order='None'

**Answer: A,C**

**Explanation: You do not want to alter or delete the trigger, so that you can use the trigger again with the least administrative effort. You should use the DISABLE TRIGGER Transact-SQL command.**

Syntax:

```
DISABLE TRIGGER { [ schema . ] trigger_name [ ,...n ] | ALL } ON {  
object_name | DATABASE | ALL SERVER } [ ; ]
```

**QUESTION NO: 33**

You are creating a view to join the Customers and Orders tables in a SQL Server 2005 database. You need to ensure that the view cannot be affected by modifications to underlying table schemas. You want to accomplish this goal by using the least possible amount of overhead. What should you do?

- A. Create CHECK constraints on the tables.
- B. Create a DDL trigger to roll back any changes to the tables if the changes affect the columns in the view.
- C. Create the view, specifying the WITH SCHEMABINDING option.
- D. Create the view, specifying the WITH CHECK option.

**Answer: C**

**Explanation: When you create a view and specify the SCHEMABINDING option, you cannot drop any tables, views, or functions referenced by the view without first dropping the view.**

**QUESTION NO: 34**

You are creating a SQL Server 2005 database for a mortgage company. The database will support a new Web-based application that will handle up to 1,000 simultaneous users. This application must quickly display the results of calculation-intensive operations, such as calculating mortgage payments and amortization schedules. You need to ensure that the database processes calculations as quickly and efficiently as possible. What should you do?

- A. Implement parameterized Transact-SQL queries in the application.
- B. Implement Transact-SQL stored procedures in the database.
- C. Implement CLR stored procedures in the database.
- D. Implement distributed Web services.

**Answer: C**

**Explanation: A Common Language Runtime (CLR) will generally perform better than T-SQL and CRUD (Create, Read, Update, Delete) operations for complex math, string manipulation and other tasks that go beyond data access.**

**QUESTION NO: 35**

You are modifying a table named Product in a SQL Server 2005 database. You want to add a new column named FriendlyName to the Product table. A friendly name for each product will be stored in this column. The table currently contains data. The sales department has not yet created a friendly name for each product. FriendlyName is a required value for each product. You want to add this new column by using the least amount of effort. What should you do?

- A. Define the new column as NULL. Update the FriendlyName column to the same value as the ProductName column. Alter the FriendlyName column to be NOT NULL.
- B. Define the new column as NOT NULL with a default value of 'Undefined.'
- C. Define the new column as NULL. Use application logic to enforce the data constraint.
- D. Define the new column as NULL with a default value of 'Undefined.'

**Answer: B**

**Explanation: You don't want to create the new column to accept NULL values and therefore you should create it as NOT NULL and by default enter 'Undefined' as the value for each product until the Sales Department changes the value to the correct friendly name.**

#### **QUESTION NO: 36**

TestKing.com Web site includes a page that customers use to send feedback about the company and its products. You use a SQL Server 2005 database to store the comments in the Comments column of a table named Feedback. You need to implement full-text searching so that you can run reports on the comments. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Create a nonclustered index on the Comments column.
- B. Execute the USE Master Transact-SQL statement.
- C. Create a full-text catalog.
- D. Create a full-text index on the Comments column.

**Answer: C,D**

**Explanation: Although you create full-text indexes on columns within tables in SQL databases, the full-text indexes are maintained in a structure outside of SQL Server called a full-text catalog. A full-text catalog can contain one or more full-text indexes.**

#### **QUESTION NO: 37**

You are responsible for managing a SQL Server 2005 database that stores sales information. Many values in nchar columns in the database tables contain preceding or trailing spaces. You need to implement a mechanism that selects the data from the tables without leading and trailing spaces. Your solution must be available for reuse in Transact-SQL statements and views. What should you do?

- A. Create DML triggers that query the inserted and deleted tables.
- B. Create a stored procedure that calls the LTRIM and RTRIM built-in functions.
- C. Create a Transact-SQL function that calls the LTRIM and RTRIM built-in functions.
- D. Call the TRIM built-in function.

**Answer: C**

**Explanation: RTRIM is a Transact-SQL command that returns a character string after truncating all trailing blanks. LTRIM is a Transact-SQL command that returns a character expression after it removes leading blanks.**

#### **QUESTION NO: 38**

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. TestKing.com 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 SQL Server 2005: 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: 39**



You are responsible for maintaining a SQL Server 2005 database. Business analysts in the company routinely use a view named v\_CustomerSales to join the Customers and Sales tables in the database. They use the view to aggregate total sales by customer by month. You need to increase the performance of the view. What should you do?

- A. Update the view to use an outer join between the Customers and Sales tables.
- B. Create a clustered index on the v\_CustomerSales view.
- C. Create two separate views that do not contain any joins one view named v\_Customers for the Customers table and another one named v\_Sales for the Sales table.
- D. Create a stored procedure for the business analysts that uses the v\_CustomerSales view.

**Answer: B**

**Explanation: As a view is nothing more than a virtual table containing data from one or more tables you have the option to create a clustered index on the view in order to obtain an indexed view and obtain performance benefits.**

#### **QUESTION NO: 40**

TestKing.com uses SQL Server 2005. You are implementing a series of views that are used in ad hoc queries. The views are used to enforce TestKing.com's security policy of abstracting data. Some of these views perform slowly. You create indexes on those views to increase performance, while still maintaining the company's security policy. One of the views returns the current date as one of the columns. The view returns the current date by using the GETDATE() function. This view does not allow you to create an index. You need to create an index on the view. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Remove all deterministic function calls from within the view.
- B. Remove all nondeterministic function calls from within the view.
- C. Schema-bind all functions that are called from within the view.
- D. Create the view and specify the WITH CHECK OPTION clause.

**Answer: B,C**

**Explanation: The GETDATE() is a nondeterministic function, meaning that it does NOT have to return the same value every times it is called, and therefore the results of the function can not be indexed. When you specify the SCHEMABINDING option, you cannot drop any tables, views, or functions referenced by the view without first dropping the view.**

#### **QUESTION NO: 41**

TestKing.com has multiple servers in a distributed environment. You work with two SQL Server 2005 computers named SQL1 and SQL2. Each server uses SQL Server Authentication and they use different logins. You need to write a distributed query that joins the data on SQL1 with the data on SQL2. What should you do?

- A. Ensure that both SQL1 and SQL2 use the same login name as the security context for each server.
- B. Configure SQL2 as a remote server. Write the query on SQL1.
- C. Configure SQL2 as a linked server to impersonate the remote login.
- D. Configure SQL2 as a distributed server. Use pass-through authentication.

**Answer: C**

**Explanation: You need to define a linked server for each external data source you want to access and then configure the security context under which your distributed queries will run. You can configure the linked server to use one of the following security modes:**

Self-mapping - When a linked server is created, this mode is added for all local logins, so SQL Server tries to connect to the external data source using the current user's login credentials. The same login and password must exist on the remote server. This is the default behaviour.

Delegation - This mode impersonates the Windows local credentials; the connection forwards the credentials of an authenticated Windows user to the linked server. The Windows user account and password must exist on the linked server.

Remote Credentials - This mode lets you map local logins to remote logins on the external data source.

#### **QUESTION NO: 42**

A new employee needs access to a SQL Server 2005 database that is located on a server named SQL1. You create a login named ajones by using the following Transact-SQL statement. `CREATE LOGIN ajones WITH PASSWORD = 'SQLServer$1'` The new employee reports that when he logs in, he receives the following error message: "Login failed. The user is not associated with a trusted SQL Server connection." You need to resolve the error and allow the new employee to gain access to SQL1. What should you do?

- A. Change the SQL Server security mode from Windows Authentication mode to SQL Server and Windows Authentication mode.
- B. Change the SQL Server security mode from SQL Server and Windows Authentication mode to Windows Authentication mode.
- C. Ensure that the login name is created with square brackets ([]).

D. Give the login access to a specific database by using the CREATE USER Transact-SQL statement.

**Answer: A**

**Explanation: The syntax to create a Windows login is:**

```
CREATE LOGIN [Domain\User] FROM WINDOWS
```

The Syntax to create a SQL Server login is

```
CREATE LOGIN login_name WITH PASSWORD='password'
```

You have created a SQL Server login on a SQL Server that is configured to only accept Windows Authentication mode.

### QUESTION NO: 43

TestKing.com has two SQL Server 2005 computers named SQL1 and SQL2. A user in TestKing.com named Eric writes many ad hoc queries against the company databases. Eric has access to the Customer database on SQL1. He does not have access to the Sales database on SQL2. You need to ensure that Eric can write queries that join information from both servers. What should you do first?

- A. Create a linked server on SQL1 to SQL2. Configure the linked server to use impersonation.
- B. Create a linked server on SQL1 to SQL2. Configure the linked server to use mapped logins.
- C. Instruct Eric to write the queries on SQL2 by using the OPENQUERY statement and specifying SQL1 as the server name.
- D. Instruct Eric to specify SQL Server object names by using four-part notation.

**Answer: B**

**Explanation: You need to define a linked server for each external data source you want to access and then configure the security context under which your distributed queries will run. You can configure the linked server to use one of the following security modes:**

Self-mapping - When a linked server is created, this mode is added for all local logins, so SQL Server tries to connect to the external data source using the current user's login credentials. The same login and password must exist on the remote server. This is the default behaviour.

Delegation - This mode impersonates the Windows local credentials; the connection forwards the credentials of an authenticated Windows user to the linked server. The Windows user account and password must exist on the linked server.

Remote Credentials - This mode lets you map local logins to remote logins on the external data source.

**QUESTION NO: 44**

Application developers in TestKing.com create an assembly that contains a CLR function. This CLR function reads data from a spreadsheet, performs some calculations, and returns the data to a SQL Server 2005 computer. You need to register the assembly with SQL Server 2005 by using the CREATE ASSEMBLY statement and the least privileged security permission set. Which permission set should you use?

- A. Default
- B. SAFE
- C. EXTERNAL\_ACCESS
- D. UNSAFE

**Answer: C**

**Explanation:** To keep CLR code from running amok, Microsoft created a three-tier security model for how CLR code can be invoked: **SAFE, EXTERNAL\_ACCESS and UNSAFE**. The **SAFE** permissions set is essentially the same as what a conventional stored procedure would be able to do. It can't modify anything outside of SQL Server itself. **EXTERNAL\_ACCESS** allows access through .NET to the Registry and the file system. **UNSAFE** is aptly named. Code marked as **UNSAFE** can't do anything, and it should really not be used outside of a debugging or experimental context. Most programmers should never need to use anything higher than **EXTERNAL\_ACCESS**.

**QUESTION NO: 45**

Your SQL Server 2005 database contains a table that has 500 million rows of data. Some of the data is historical and some is current. You need to partition the data on a single server to increase performance and optimize maintenance. What should you do?

- A. Implement horizontal partitioning.
- B. Implement vertical partitioning.
- C. Implement a raw partition.
- D. Implement distributed partitioning.

**Answer: A**

**Explanation: Table and index partitioning - intended to improve performance of operations performed on large tables. The basic concept is straightforward and involves splitting a table into several units (called partitions), which can be accessed independently of each other, limiting impact of I/O intensive activities performed on the table's data (queries, data loads, backups and restores, maintenance tasks - such as index rebuilds and defragmentations, as well as operations that would result in lock escalation to the table level). The most common method of splitting data is horizontal partitioning, in which rows of a table matching mutually exclusive criteria (such as, range of dates or letters in alphabet, for datetime and character data, respectively) are placed in designated partitions.**

**QUESTION NO: 46**

You work with a database named DB1, which is located on a SQL Server 2005 computer. You discover that a table scan on Table1 in DB1 causes a slow query. Table1 is a very large table that is used frequently. You ascertain that the query contains the following statement. `SELECT col1, col2 FROM table1 WHERE col3 = <value>` You need to provide maximum query performance. Table1 must remain available to users. What should you do?

- A. Update all statistics on Table1 in DB1.
- B. Use the `CREATE STATISTICS` statement in DB1 to create missing statistics on col3 of Table1.
- C. Set the priority boost server option to 1.
- D. Execute the following statement.`USE DB1GOCREATE INDEX index1 ON table1(col3) INCLUDE(col1, col2) WITH (ONLINE = ON)GO`
- E. Execute the following statement.`USE DB1GOCREATE INDEX index1 ON table1(col3, col2, col1)`

**Answer: D**

**Explanation: You want to create an index on the columns 1 to 3 in order to have quicker response when querying data in columns 1 and 2 based on a value in column 3.**

Syntax-- Relational Index`CREATE [UNIQUE] [CLUSTERED | NONCLUSTERED] INDEX index ON object (column [ASC | DESC] [,...n] ) [INCLUDE (column [ ,...n] ) ] [WITH (option [ ,...n] ) ] [ON { partition_scheme ( column ) | filegroup | default } ] [;]`

**QUESTION NO: 47**

Exhibit:

Table name	Stores customer surname beginning with letter
TestKing1	A - H
TestKing2	I - P
TestKing3	Q - Z

TestKing.com is storing its customer data in a data warehouse. It 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 table displayed in the exhibit.

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 TESTKING1
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING2
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING3
ADD OptOutFlag BIT NULL
IF @@error <> 0 ROLLBACK TRAN ELSE COMMIT TRAN
```

B. BEGIN TRAN

```
BEGIN TRY
ALTER TABLE TESTKING1
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING2
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING3
ADD OptOutFlag BIT NULL
END TRY
BEGIN CATCH
ROLLBACK TRAN
```

```
RETURN
END CATCH
COMMIT TRAN
```

```
C. ALTER TABLE TESTKING1
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING2
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING3
ADD OptOutFlag BIT NULL
```

```
D. BEGIN TRAN
ALTER TABLE TESTKING1
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING2
ADD OptOutFlag BIT NULL
ALTER TABLE TESTKING3
ADD OptOutFlag BIT NULL
COMMIT TRAN
```

**Answer: B**

**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, the whole transaction will rollback.

#### **QUESTION NO: 48**

You are responsible for managing the database schema for an order entry application in a SQL Server 2005 database. After a design review, a member of the development staff asks you to add a new column named **CommissionRate** to the **Product** table. Because the actual commission rates for all products are not known at this time, each product has a default commission rate of 12 percent. The rate can be modified as necessary. You add the new column. You need to configure the table to assign the default value as efficiently as possible. What should you do?

- A. Create an **INSERT** trigger to assign the default value to each item in the table.
- B. Create a **CHECK** constraint to validate the data and to assign the default value to each item in the table.
- C. Create an **UPDATE** trigger to update the default value for each new item in the table.

D. Create a DEFAULT constraint to assign the default value specifying the WITH VALUES argument.

**Answer: D**

**Explanation:**

[ CONSTRAINT constraint\_name] - Specifies the start of a DEFAULT definition. To maintain compatibility with earlier versions of SQL Server, a constraint name can be assigned to a DEFAULT. constraint\_name must follow the rules for identifiers, except that the name cannot start with a number sign (#). If constraint\_name is not specified, a system-generated name is assigned to the DEFAULT definition.

WITH VALUES -Specifies that the value given in DEFAULT constant\_expression is stored in a new column added to existing rows. If the added column allows null values and WITH VALUES is specified, the default value is stored in the new column, added to existing rows. If WITH VALUES is not specified for columns that allow nulls, the value NULL is stored in the new column in existing rows. If the new column does not allow nulls, the default value is stored in new rows regardless of whether WITH VALUES is specified.

**QUESTION NO: 49**

TestKing.com has a SQL Server 2005 computer. You have been assigned the task of retrieving information about a user who is currently logged in. You need to create a function that returns scalar information about the activity time for a particular user. What are two possible ways to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. Create a function that returns a list of values that represent the login times for the given user.
- B. Create a function that returns a list of values that represent the people who have logged more hours than the current user has logged.
- C. Create a function that returns a numeric value that represents the number of hours that a user has logged for the current day.
- D. Create a function that returns a numeric value that represents the number of hours that a user has logged for the current month.

**Answer: C,D**

**Explanation:** Values representing login times will only show at what time a user logs in not for how long and values representing users who have logged more hours than the current user has logged does not have to show activity time for all users.



**QUESTION NO: 50**

You work with a SQL Server 2005 database that provides banking information for customers. You want customers to see banking reports that combine data that is retrieved from the database with real-time investment information that comes from a Web service provided by a third party. The investment information must be current when the reports are executed. You need to create the appropriate objects that support the reports. What should you do?

- A. Publish the data in the database as an XML Web service by using the FOR XML AUTO clause.
- B. Create a table to store the banking information for each customer. Create a trigger that fires when data is inserted into the table that joins with the data coming from the Web service.
- C. Create a Transact-SQL stored procedure that uses a temporary table to store the banking information for each customer. Update the table with the values from the Web service.
- D. Have a developer in TestKing.com create an assembly that calls the remote Web service. Create a CLR function by using the assembly. Call the CLR function and combine the results with banking information in the database.

**Answer: D**

**QUESTION NO: 51**

You are planning to import customer data from the Customer table and the Country table in your SQL Server 2000 database into a SQL Server 2005 database. You want to ensure that during the import process, each value in the CountryCode column of the Customer table has a corresponding record in the CountryCode column in the Country table of the SQL Server 2005 database. You define a foreign key between these two tables. Thus, referential integrity will guarantee that the import routine fails if there are any records where a CountryCode value does not exist in the Country table but does exist in the Customer table. You need to ensure that the import process does not fail if records are encountered that do not exist in the Country table. What should you do?

- A. Drop the foreign key. Import the data by using the script. Re-create the foreign key.
- B. Create a CHECK constraint.
- C. Create a DML INSTEAD OF trigger that writes the failed records to a file or table.
- D. Create a DML AFTER trigger that writes the failed records to a file or table.

**Answer: C**

**Explanation:** A DML INSTEAD OF trigger causes SQL Server to execute the code in the trigger instead of the operation that caused the trigger to fire, meaning that correctly used this will write a record missing the CountryCode to an alternative destination and not stop the import.

**QUESTION NO: 52**

Exhibit:

Column name	Type
CustomerID	int
CompanyName	varchar(50)
FirstName	varchar(20)
LastName	varchar (20)
Picture	varbinary(max)
Address1	varchar(30)
Address2	varchar(30)
City	varchar(30)
StateProvince	char(2)
PostalCode	varchar(10)
CountryCode	char(2)

You use a SQL Server 2005 database that contains a table named Customers. The Customers table contains the columns displayed in the exhibit.

You create a view named v\_Customer by using the following Transact-SQL statement.  
CREATE VIEW v\_Customer AS SELECT CustomerID, CompanyName, FirstName, LastName, Picture FROM Customers  
You want to create an index on the v\_Customer view. Prior to creating the index, you need to modify the Transact-SQL statement. Which two actions should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Change the Picture column to an image data type.
- B. Include the schema name with the table name.
- C. Include the WITH SCHEMABINDING clause.
- D. Select all columns by using the \* notation.

**Answer: B,C**

**Explanation: The Syntax for the CREATE VIEW is:**

```
CREATE VIEW [schema_name . ] view_name [(column [ ,...n ] ) ]
```

```
[WITH <view_attribute> [ ,...n ] ]  
AS selected_statement [ ; ]  
[ WITH CHECK OPTION]
```

```
<view_attribute ::=  
{  
[ ENCRYPTION ]  
[ SCHEMABINDING ]  
[ VIEW_METADATA ] }
```

When you specify the SCHEMABINDING option, you cannot drop any tables, views, or functions referenced by the view without first dropping the view. Using schemas and including the schema name with the table is good practice since schemas and users now are separate entities.

#### QUESTION NO: 53

TestKing.com uses a SQL Server 2005 database. Customer data from your trading partners is imported into the Customers table every night. You need to ensure that the customer record is updated if it already exists in the Customers table. If the record does not exist, the data needs to be inserted into the Customers table. What should you do?

- A. Create a FOR trigger.
- B. Create an INSTEAD OF trigger.
- C. Create an AFTER trigger.
- D. Create a DDL trigger.

**Answer: B**

**Explanation:** By using an INSTEAD OF trigger you can write a job that inserts the records into the table and if the record already exists the INSTEAD of trigger will fire and execute an update of the record instead of an insert.

#### QUESTION NO: 54

You are creating a view to join the Customers and Orders tables in a SQL Server 2005 database. You need to ensure that the view cannot be affected by modifications to underlying table schemas. You want to accomplish this goal by using the least possible amount of overhead. What should you do?

- A. Create CHECK constraints on the tables.
- B. Create a DDL trigger to roll back any changes to the tables if the changes affect the columns in the view.

- C. Create the view, specifying the WITH SCHEMABINDING option.
- D. Create the view, specifying the WITH CHECK option.

**Answer: C**

**Explanation:** When you specify the SCHEMABINDING option, you cannot drop any tables, views, or functions referenced by the view without first dropping the view.

**QUESTION NO: 55**

Your department is responsible for creating a list of customer e-mail addresses from data in a SQL Server 2005 database. The list must include the date on which each customer was last contacted. The data in the result set must have column names. The list must be ordered by date of last contact, with the most recent contact first. The LastContact column is stored as a datetime. The date should be displayed as MM/DD/YYYY. A coworker develops the following query. `SELECT email_address AS EmailAddress, CONVERT(nvarchar, lastcontact, 101) AS LastContact FROM Contact ORDER BY LastContact DESC` You test this query and find that the data is displayed in the following order. `EmailAddress LastContact andrew@contoso.com 01/24/2003 marc@contoso.com 06/12/2005 stefan@contoso.com 07/13/2004` You need to modify the query so that the data is listed in the proper order. You want to achieve this goal without negatively affecting performance. What should you do?

- A. Change the ORDER BY clause as follows:`ORDER BY CONVERT(nvarchar, lastcontact, 101) DESC`
- B. Remove the column alias from the LastContact column.
- C. Change the ORDER BY clause as follows:`ORDER BY CAST(lastcontact AS nvarchar(20)) DESC`
- D. Change the alias on the LastContact column.

**Answer: D**

**QUESTION NO: 56**

You are responsible for managing a SQL Server 2005 database. This database is updated with data that is exported from another site. On a monthly basis, data is loaded into the database from a flat file by using the bcp utility. This operation should be minimally logged. However, users report that they receive error messages indicating that the transaction log is full. You investigate and discover that the errors occur when the data is being imported. You need to prevent this problem from occurring. What should you do?

- A. Disable the clustered index and all nonclustered indexes on the target table.

- B. Set the database recovery model to full.
- C. Set the database recovery model to bulk-logged.
- D. Use the BULK INSERT statement to load the data.

**Answer: C**

**Explanation: The bulk-logged model differs from the Full recovery model in many ways, one of which is that bulk-load operations are minimally logged so the transactions will not be rolled back by the bulk-load operation.**

#### **QUESTION NO: 57**

The sales department and the finance department are located in different buildings. The two departments are connected by a 2 megabit network. Both departments use SQL Server 2005. The sales department records all sales activities in a database named DB1. DB1 contains 380 tables and has a size of 10 GB. The finance department has read-only access to four tables in DB1 for reporting purposes. The finance department runs queries throughout the day and needs the data to be as up to date as possible. Users in the finance department tell you that report execution is slow. To optimize network capacity, your department decides to provide copies of these four tables to the finance department. You need to provide copies of these tables. You want to achieve this goal by using the copy method that uses the least amount of network bandwidth. What should you do?

- A. Mirror DB1 to the finance department. Configure mirroring for high-availability mode.
- B. Establish daily transaction log shipping of DB1 to the finance department. Use the STANDBY option to make the shipped database accessible to reporting.
- C. Establish daily snapshot replication to the finance department. Define a publication that contains the four tables as the only articles.
- D. Establish transactional replication to the finance department. Define a publication that contains the four tables as the only articles.

**Answer: D**

**Explanation: A mirror is in a recovering state, not allowing any connections, so it is not an option in this case. Daily transaction log shipping or snapshot replications are too slow to apply to the "as up to date as possible" part of the scenario.**

#### **QUESTION NO: 58**

You are responsible for maintaining scheduled jobs for the SQL Server 2005 databases in your department. One of the jobs aggregates data from multiple sources for reports. This job runs daily. The job consists of multiple steps. Each step aggregates data for a specific report. A user reports that the data for some of the reports has not been updated recently. You need to ensure that every step of the job that aggregates data executes even when errors occur. What should you do?

- A. Combine all of the steps into a single step that runs once a day.
- B. Create a notification that alerts you when an error occurs so that you can correct the error and restart the job.
- C. Change the On Failure action to go to the next step.
- D. Configure the job to retry the step.

**Answer: C**

**Explanation: When creating Microsoft SQL Server Agent jobs, you can specify what action SQL Server should take if a failure occurs during job execution. Determine the action that SQL Server should take upon the success or failure of each job step. Then use the following procedure to configure the job step action flow logic by using SQL Server Agent.**

1. In Object Explorer, expand SQL Server Agent, and then expand Jobs.
2. Right-click the job you want to edit, and then click Properties.
3. Select the Steps page, click a step, and then click Edit.
4. In the Job Step Properties dialog box, select the Advanced page.
5. In the On success action list, click the action to perform if the job step completes successfully.
6. In the Retry attempts box, enter the number of times from 0 through 9999 that the job step should be repeated before it is considered to have failed. If you entered a value greater than 0 in the Retry attempts box, enter in the Retry interval (minutes) box the number of minutes from 1 through 9999 that must pass before the job step is retried.
7. In the On failure action list, click the action to perform if the job step fails.
8. If the job is a Transact-SQL script, you can choose from the following options:
  - \*In the Output file box, enter the name of an output file to which the script output will be written. By default the file is overwritten each time the job step executes. If you do not want the output file overwritten, check Append output to existing file.
  - \*Check Log to table if you want to log the job step to a database table. By default the table contents are overwritten each time the job step executes. If you do not want the table contents overwritten, check Append output to existing entry in table. After the job step executes, you can view the contents of this table by clicking View.
  - \*Check Include step output in history if you want the output included in the step's history. Output will only be shown if there were no errors. Also, output may be truncated.
9. If the

Run as user list is available, select the proxy account with the credentials that the job will use.

**QUESTION NO: 59**

You are creating a SQL Server 2005 application that will store and manage legal documents. The original documents are stored as XML documents on a file server. The application will be used to insert the documents into the database. Later, the documents must be retrieved from the database. These documents must be identical to the originals. You need to design a table to store the document data. What should you do?

- A. Store the XML in a column that has a data type of nvarchar(max).
- B. Shred the XML and store it in a relational structure.
- C. Store the XML in a column that has a data type of XML.
- D. Store the XML in a column that has a data type of varchar(8000).

**Answer: A**

**Explanation: The possible Character Data Types are: char(n), nchar(N), varchar(n), varchar(max), nvarchar (max), text and ntext. Nvarchar(max) will store up to 2GB in maximum 536,870,912 characters and is used for Unicode data type that is variable length. The XML datatype is for well-formed documents and as you have to store the documents exactly as the original there is no guarantees that the documents are well formed.**

**QUESTION NO: 60**

You use Service Broker to manage data requests between SQL Server 2005 databases. You need to make changes to objects for an existing service. You want to ensure that no messages can be received by the service until the changes are complete. What should you do?

- A. Create the queue with ACTIVATION STATUS OFF. Set it to ON when the service is ready.
- B. Create the queue with STATUS OFF. Set it to ON when the service is ready.
- C. Create the queue with only the queue\_name argument. Set the other parameters when the service is ready.
- D. Create the queue with max\_queue\_readers = 0. Set it to an optimum value when the service is ready.

**Answer: B**

**Explanation:** The STATUS clause determines whether the queue is enabled. When a queue is set to STATUS = OFF, it does not allow any messages to be added to or removed from the queue.

**QUESTION NO: 61**

You are a database developer for a company that manufactures bicycles. You are creating a Web application that will retrieve data from a database named TestKing. The application connects to a database server that is running SQL Server 2005. One of the functions of your application will allow shift supervisors to see how many bicycles were manufactured during the current shift. A user named Tom creates a stored procedure named usp\_GetPersonData in the dbo schema. You are creating an HTTP endpoint that exposes a Web method named PersonData. This Web method allows client applications to access the stored procedure. You use the following statement to create the endpoint.

```
CREATE ENDPOINT SQLEP_AWPpersons
AS HTTP (PATH = '/AWpersons', AUTHENTICATION = (INTEGRATED), PORTS =
(CLEAR), SITE = 'SQL1')
FOR SOAP (WEBMETHOD " (NAME=)", BATCHES = DISABLED, WSDL =
DEFAULT, DATABASE = 'TestKing', NAMESPACE =
'http://Adventure-Works/Persons')
```

You need to specify the correct WEBMETHOD and NAME to create the endpoint. Which command should you use?

- A. WEBMETHOD 'http://www.adventure-works.com /PersonData' NAME = 'TestKing.dbo.usp\_GetPersonData'
- B. WEBMETHOD 'PersonData'NAME = 'TestKing.dbo.usp\_GetPersonData'
- C. WEBMETHOD 'PersonData'NAME = 'TestKing.Tom.usp\_GetPersonData'
- D. WEBMETHOD 'TestKing.PersonData'NAME = 'TestKing.Tom.usp\_GetPersonData'

**Answer: B**

**Explanation:** The general format of the WEBMETHOD portion of the SOAP payload is as follows:

```
[ { WEBMETHOD [ 'namespace' .] 'method_alias'
( NAME = 'database.owner.name'
[ , SCHEMA = { NONE | STANDARD | DEFAULT } ]
[ , FORMAT = { ALL_RESULTS | ROWSETS_ONLY | NONE } ]
)
```

**QUESTION NO: 62**



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:**

query() - Provides the ability to execute an XPATH or XQUERY expression and returns the resulting XML fragment.

value() - Provides the ability to execute an XPATH or XQUERY expression and returns a single scalar value that is converted into a SQL type.

exist() - Provides the ability to execute an XPATH or XQUERY expression to check for the existence of nodes. If the query returns a node collection, the exist() methods returns true; otherwise, the exist() method returns false.

modify() - Provides XML data-manipulation capabilities.

nodes() - Provides the ability to execute an XPATH or XQUERY expression and returns the resulting XML fragment shredded into row set.

**QUESTION NO: 63**

You are moving a SQL Server 2005 database named Timesheet to a new database server. The Timesheet database is used by the TimeEntry application. You need to minimize the amount of time that the TimeEntry application is unavailable while you move the database. What should you do?

- A. Move the database to the new server by using the SQL Management Object method in the Copy Database Wizard.
- B. Detach the current database. Copy the data files to the new server. Attach the files.
- C. Back up the database. Copy the backup file to the new server. Restore the database.
- D. Move the data files. Specify the new location by using ALTER DATABASE.

**Answer: B**

**Explanation:** You can unmount databases from a SQL Server by detaching them. This process removes the entries in the system tables for this database, causing it to no longer be accessible on the SQL Server instance. Although the database is inaccessible, the files that contain all the objects and data still exist on the operating system in the location in which you created them. After they are detached, you can copy these files to any location on your network because they are no longer being accessed by SQL Server. The detach operation requires SQL Server only to close the files and remove an entry in the system tables. And an attach requires SQL Server to simply open the files and make an entry in the system tables. Each operation requires only 1-2 seconds at most to complete.

**QUESTION NO: 64**

You use an order application named App1, which connects to a SQL Server 2005 database named DB1. Because some actions take more than three seconds, users report slow performance when they use App1. You want to identify those actions by using the Tuning template in SQL Server Profiler to trace the application. You need to capture only those App1 actions in DB1 that have a response time of two seconds or more. What should you do?

- A. Add the ApplicationName column and create the following column filters. ColumnOperator Value Duration Greater than or equal 2000 ApplicationName Like App1 DatabaseName Like DB1
- B. Add the ApplicationName column and create the following column filters. ColumnOperator Value Duration Greater than or equal 2 ApplicationName Like App1 DatabaseName Like DB1
- C. Add the CPU column and create the following column filters. ColumnOperator Value CPU Greater than or equal 2000 TextData Like App1 DatabaseName Like DB1
- D. Add the ApplicationName and CPU columns and create the following column filters. ColumnOperator Value CPU Greater than or equal 2 ApplicationName Like App1 TextData Like DB1

**Answer: A**

**Explanation:** You want to trace that Application App1 so you should use the ApplicationName column and you want to trace the actions that are greater than or equal to 2000 mS.

**QUESTION NO: 65**

TestKing.com has a network that consists of two domains named DOM1 and DOM2. There is no trust relationship between DOM1 and DOM2. DOM1 contains a server named SQL1. DOM2 contains a server named SQL2. Both servers run SQL Server 2005. The SQL Server instance on SQL1 uses the DOM1\SQLService1 service account. The SQL Server instance on SQL2 uses the DOM2\SQLService2 service account. There are no mirroring endpoints on either server. SQL1 contains a database named DB1. You want to mirror DB1 to SQL2 by using high-protection mode without a witness. You need to create the database mirroring endpoints and configure security on all participating server instances. What should you do?

- A. Create endpoints and configure security by using the Configure Database Mirroring Security Wizard for DB1. Set up the endpoints to use SQL Server Authentication.
- B. Configure database mirroring endpoints and security for each server instance by using Transact-SQL. Allow database mirroring to use encrypted certificates for outbound and inbound connections.
- C. Create logins and add them to the sysadmin group. Configure the logins to use SQL Server Authentication on each of the server instances. Create endpoints by using the Configure Database Mirroring Security Wizard on the Database Properties Mirroring page for DB1.
- D. Connect to each server instance and execute the following statement.  
`CREATE ENDPOINT mirroring STATE = STARTED AS TCP ( LISTENER_PORT = 5022 ) FOR DATABASE_MIRRORING ( AUTHENTICATION = WINDOWS NTLM, ROLE=PARTNER);`

**Answer: B**

**Explanation: The TCP Endpoint for mirroring has to be created with the T-SQL command:**

```
CREATE ENDPOINT [Mirroring]
AS TCP (LISTENER_PORT = 5022)
FOR DATA_MIRRORING (ROLE = PARTNER, ENCRYPTION = REQUIRED);
ALTER ENDPOINT [Mirroring] STATE = STARTED;
```

All communications between endpoints can be encrypted, and you specify which encryption algorithm to use for the communications. The default algorithm is RC4.

#### **QUESTION NO: 66**

TestKing.com uses a mission-critical database named DB1, which is located on a server named SQL1. You want to mirror DB1 on a second server named SQL2. Before you can establish database mirroring, you need to prepare DB1 and the servers for mirroring. Which three actions should you perform? (Each correct answer presents part of the solution. Choose three.)

- A. Set the recovery model of DB1 on SQL1 to bulk-logged recovery.
- B. Set the recovery model of DB1 on SQL1 to full recovery.
- C. Back up DB1 on SQL1. Restore the backup on SQL2; specify the NORECOVERY option.
- D. Back up DB1 on SQL1. Restore the backup on SQL2; specify the STANDBY option.
- E. Back up DB1 on SQL1. Restore the backup on SQL2; specify the RECOVERY option.
- F. Create endpoints on all participating servers.
- G. Set the AutoShrink property of DB1 on SQL1 to false.

**Answer: B,C,F**

**Explanation: In order to configure mirroring you have to ensure that the databases are set to use the Full recovery model. Because database mirroring maintains both the primary and mirror databases as exact duplicates, including synchronizing all internal structures such as Log Sequence Numbers (LSNs), the Simple and Bulk-Logged recovery models are incompatible with database mirroring.**

When restoring the database to the mirror, it is essential that you specify the NORECOVERY option for the RESTORE command, which guarantees that the starting state of the mirror reflects the state of the principal database, including the LSNs.

Finally you have to produce a communication channel for the database mirroring to use and this is done by creating a endpoint for data mirroring on all participating servers.

#### **QUESTION NO: 67**

You are responsible for configuring and maintaining the backup jobs for a SQL Server 2005 computer. Backup files are written to a tape drive that is connected to the database server. In the future, your department will install a new backup server. Then backup files will be written to disk on the new server. To simplify maintenance, you need to configure the backup jobs so that minimal changes will be necessary to back up to the new server. What should you do?

- A. Configure the backup jobs to write to an internal disk. When the new backup server is available, alter the backup jobs.
- B. Configure the backup jobs to write to backup devices that reference the tape drive. Alter the backup devices to point to the new backup server after the new server is installed.
- C. Configure the backup jobs to write directly to the tape drive. Move the tape drive to the new backup server.
- D. Configure the backup jobs to write to an internal disk. Alter the backup jobs to use backup devices after the new server is in place.

**Answer: B**

**Explanation: Backup Device - A predefined pointer to a backup location. Instead of dynamically specifying a tape or disk location within the backup command, a backup device assigns a name to the physical location, and this name can then be reused.**

**QUESTION NO: 68**

You discover that all indexes of the Sales table on your SQL Server 2005 database server are heavily fragmented. You need to decrease the fragmentation of all indexes of the Sales table to a minimum, while keeping the Sales table available to all users. What should you do?

- A. Defragment the disk that contains the Sales table.
- B. Execute the following statement.

```
USE master;GODBCC CLEAN TABLE ('DB1','sales');GO
```
- C. Execute the following statement.

```
USE DB1;GOALTER INDEX ALL ON sales REORGANIZE WITH (LOB_COMPACT = ON);GO
```
- D. Execute the following statement.

```
USE DB1;GOALTER INDEX ALL ON sales REBUILD;GO
```

**Answer: C**

**Explanation: Defragmenting the disk will not help, this might help if the files containing the database were fragmented but not if the index itself is fragmented. Executing a statement with USE master will not help as the sales table is contained in the DB1 database. You have to use ALTER INDEX on DB1 and use the REORGANIZE switch to minimize internal fragmentation while the table is online. The REBUILD switch is used for external fragmentation and will drop and rebuild the index to reduce both external and internal fragmentation and therefore the table will not be available to users during rebuild.**

**QUESTION NO: 69**

TestKing.com uses SQL Server 2005. A user reports that an order processing application stopped responding during a transaction. You investigate and discover that the user's session, which is session ID 85, is being blocked by an errant process. The session ID of the errant process is 101. You need to solve this blocking situation quickly so that the user can continue using the application without losing any work. 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 a new query window. Execute the following statement.

```
KILL 85
```

B. In SQL Server Management Studio, open a new query window. Execute the following statement.KILL 101

C. In SQL Server Management Studio, open the Activity Monitor. Open the Process Info page. Right-click session ID 85 and select Kill Process.

D. In SQL Server Management Studio, open the Activity Monitor. Open the Process Info page. Double-click session ID 101 to show the process details. Click the Kill Process button.

E. In SQL Server Configuration Manager, stop and then restart the SQL Server (MSSQLSERVER) service.

F. At a Microsoft Windows command prompt, execute the following statements.NET STOP MSSQLSERVERNET START MSSQLSERVER

**Answer: B,D**

**Explanation: You want to KILL the errant process, not the one being blocked. You can kill the process either with a query or from the Process Info page in the SQL Server Management Studio. The Kill command can only be executed by a member of the sysadmin fixed server role.**

#### **QUESTION NO: 70**

TestKing.com 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 sqlcmd 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: either via SQLCMD or through the query window in SSMS.

**QUESTION NO: 71**

You have a SQL Server 2005 computer named SQL1. Users report that the time required to run their queries on SQL1 is steadily increasing. You want to investigate this issue by using SQL Server Profiler to trace all user queries. You need to decide on which computer to run SQL Server Profiler, and where to store the trace. You must minimize the impact of this operation on SQL1 performance. What should you do?

- A. Run SQL Server Profiler from the console of SQL1. Store the trace in a table on SQL1.
- B. Run SQL Server Profiler over a Remote Desktop connection to SQL1. Store the trace in a file on a separate computer.
- C. Run SQL Server Profiler on a separate computer that is connected to SQL1. Store the trace in a table on SQL1.
- D. Run SQL Server Profiler on a separate computer that is connected to SQL1. Store the trace in a file on the computer that is running SQL Server Profiler.

**Answer: D**

**Explanation:** In order to minimize impact on SQL1 when running the SQL Server Profiler you should run this on a separate computer connected to SQL1 (and a remote desktop session is NOT a separate computer) and store the trace on the computer running the SQL Server Profiler.

**QUESTION NO: 72**

You use a SQL Server 2005 computer that was installed with the default settings. A database named DB1 is located on the server. Due to the recent failure of a hard disk controller, you decide to check the integrity of DB1. You create a maintenance task named CheckDB1 that uses default settings. Because this check must not interfere with user transactions, you schedule it to run over the weekend. The following Monday, you need to retrieve the outcome details of the CheckDB1 maintenance plan. What are two possible ways to achieve this goal? (Each correct answer presents a complete solution. Choose two.)

- A. In SQL Management Studio, right-click the CheckDB1 maintenance plan. On the context menu, click View History. See the record for the plan named CheckDB1.
- B. In SQL Management Studio, right-click the Activity Monitor. On the context menu, click View Processes. Look for a record that has "CheckDB1" in the Application column.
- C. In Windows Explorer, navigate to the C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG folder. View the SQLAgent.out text file.
- D. In Windows Explorer, navigate to the C:\Program Files\Microsoft SQL Server\MSSQL.1\MSSQL\LOG folder. View the text file that has a name that starts with "CheckDB1\_" followed by the execution year, date, and time.

**Answer: A,D**

### **QUESTION NO: 73**

Your database named DB1 is one of 15 databases located on a SQL Server 2005 computer named SQL1. Recently, users who work with DB1 report that they experience long query execution times. You want to reproduce and analyze this problem in a test environment that holds all copies of the production databases. You need to record the workload for an entire day on SQL1 for later execution on the test server. You must ensure that all events are captured without skipping any, even if SQL1 is under stress. What should you do?

- A. Start SQL Server Profiler on SQL1 and connect locally. Use the Tuning template. Create a filter for the DatabaseID column that is equal to the database ID of DB1. Save the trace to a table on SQL1.
- B. Start SQL Server Profiler on your computer and connect to SQL1. Use the Tuning template. Save the trace to a table on the test computer. Ensure that the Server processes trace data option is selected.
- C. Start SQL Server Profiler on the test computer and connect to SQL1. Use the Replay template. Save the trace to a file on the test computer. Ensure that the Server processes trace data option is selected.
- D. Start SQL Server Profiler on your computer and connect to SQL1. Use the Replay template. Save the trace to a file on your computer. Ensure that the Server processes trace data option is not selected.

**Answer: C**

**Explanation: You want to run the SQL Server Profiler from the test computer connected to the SQL1 server.**

## **Topic 2, Simulations (12 Questions)**



**QUESTION NO: 1 SIMULATION**

**You work as a database administrator at TestKing.com.**

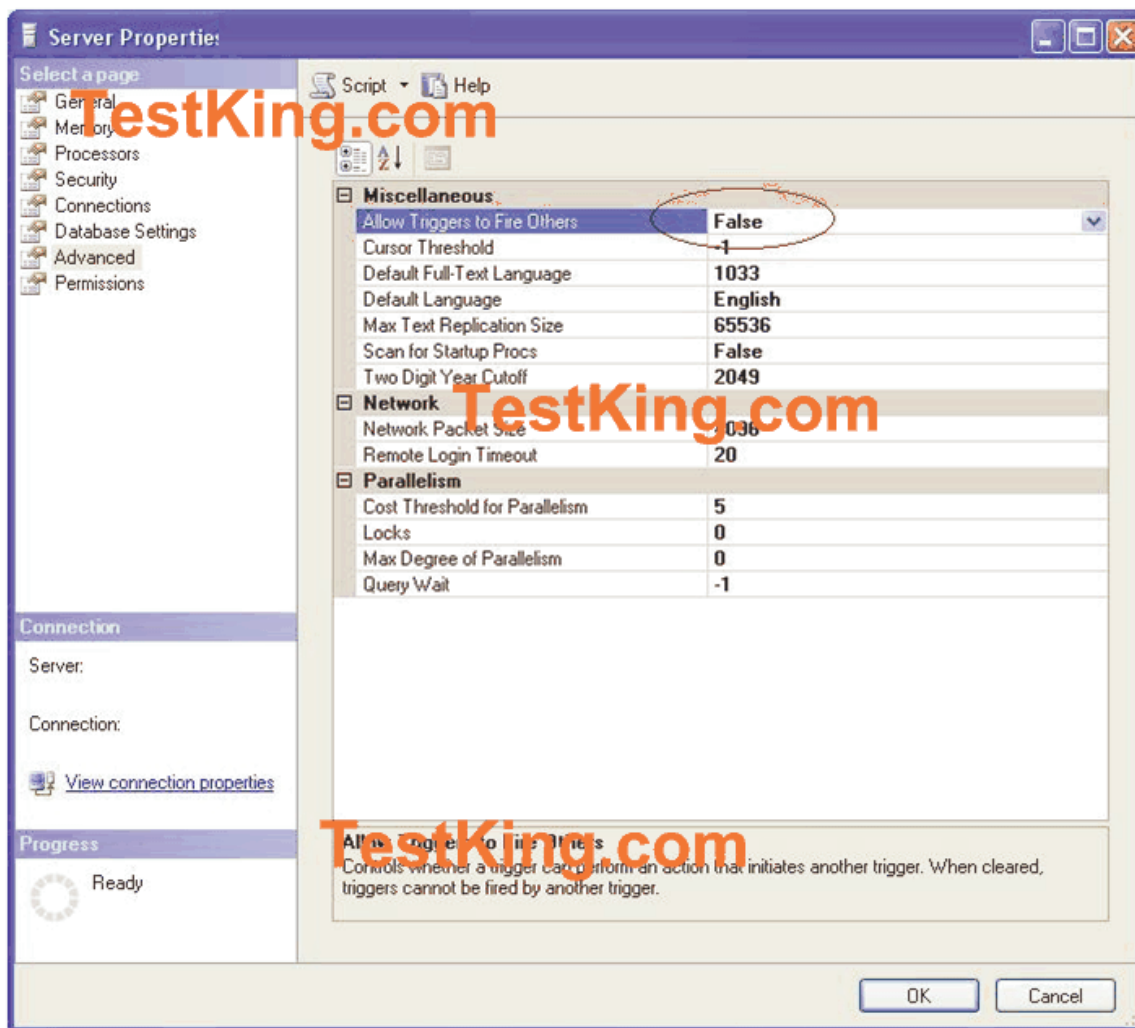
**The only non-system database on the server is TestKingDB.**

**A new written company policy states that nested trigger execution is not allowed.**

**You need to ensure that this policy is implemented by modifying the appropriate server and database settings in the Database Properties dialog box of the Server Properties dialog box, or both of these dialog boxes, in SQL Server Management Studio (SSMS).**

**Answer:**

**Explanation:**



**QUESTION NO: 2 SIMULATION**

**You work as a database administrator at TestKing.com.**

**You need to install a new SQL Server 2005 instance on a newly installed Windows Server 2003 server named TestKing2. This new instance must have the same settings as the existing SQL Server 2005 instance on a server named TestKing1.**

**The only two services that need to be installed are the SQL Server service and the SQL Server agent service.**

**The default installation folder should be used for the new SQL Server instance on TestKing2. You will be able to refer to TestKing1 to view the settings to use.**

**You need to complete the installation wizard on TestKing2.**

**Answer:**

**Explanation: Pending. Send your suggestion to [feedback@testking.com](mailto:feedback@testking.com)**

**QUESTION NO: 3 SIMULATION**

**You work as a database developer at TestKing.com. Users report that queries that are executed against some views are returning incorrect results.**

**You investigate and discover that the problem is being caused by one or more views. You want to ensure that each of these views returns the results described in the following table.**

**Exhibit:**

View name	Required result
<u>HumanResources.vEmployeeTopVactionHours</u>	Returns the employees that have the highest value in the <u>Vacationhours</u> field.
<u>HumanResources.vEmployeeAverageOrMoreVacationHours</u>	Returns the employees that have a value in the <u>VacationHours</u> field that is greater than or equal to the average
<u>Production.vProductMediumPrice</u>	Returns the products that have a value in the <u>ListPrice</u> field that is greater than or equal to 10 and less than 20.
<u>Product.vProductMini</u>	Returns the products that have a value in the <u>Name</u> Field that starts with the text Mini using a case sensitive comparison.

You need to examine the query in each o the four views and then delete the view or views that return incorrect results. You must not delete views that return correct results.

**Answer:**

**Explanation:** Pending. Send your suggestion to [feedback@testking.com](mailto:feedback@testking.com)

#### QUESTION NO: 4 SIMULATION

You work as a database administrator at TestKing.com.

There is an uncommitted transaction in the TestKing database. You must find the process that is associated with this transaction, and you must kill that process. You must not kill any other processes.

You need to kill the process by using the Activity Monitor dialog how within SQL Server Management Studio (SSMS).

**Answer:**

**Explanation:** Pending. Send your suggestion to [feedback@testking.com](mailto:feedback@testking.com)

**QUESTION NO: 5 SIMULATION**

**You work as a database administrator at TestKing.com.**

**A query against the Sales.SalesOrderDetail table is performing poorly. You must create an index that will cause the cost of the query to be as low as possible. You must minimize the space that is used by the index. The index must contain as few columns as possible.**

**There is an existing clustered index on the table. This index must not be dropped.**

**Because the table is on a production server, it is not possible to create different indexes and test their performance.**

**You need to create the required indexes by using the New Index dialog box.**

**Answer:**

**Explanation: Pending. Send your suggestion to [feedback@testking.com](mailto:feedback@testking.com)**

**QUESTION NO: 6 SIMULATION**

**You work as a database administrator at TestKing.com.**

**Users report that insert statements against the view named Production.vExpensiveExpiredProduct do not succeed. You confirm that the insert statements that are being used are valid.**

**The only table that the view uses is named Production.ExpiredProduct. The table has the following definition:**

```
CREATE TABLE Production.ExpiredProduct {  
ExpiredProductID INT IDENTITY CONSTRAINT PK_EXPIREProduct  
PRIMARY KEY  
Name NVARCHAR(50) NOT NULL  
ListPrice MONEY NULL  
}
```

**You need to ensure that insert statements can be complete successfully. You need to apply the fewest possible changes, and you must use SQL Server Management Studio (SSMS).**

**Answer:**

**Explanation:** Pending. Send your suggestion to [feedback@testking.com](mailto:feedback@testking.com)

**QUESTION NO: 7 SIMULATION**

**You work as a database administrator at TestKing.com.**

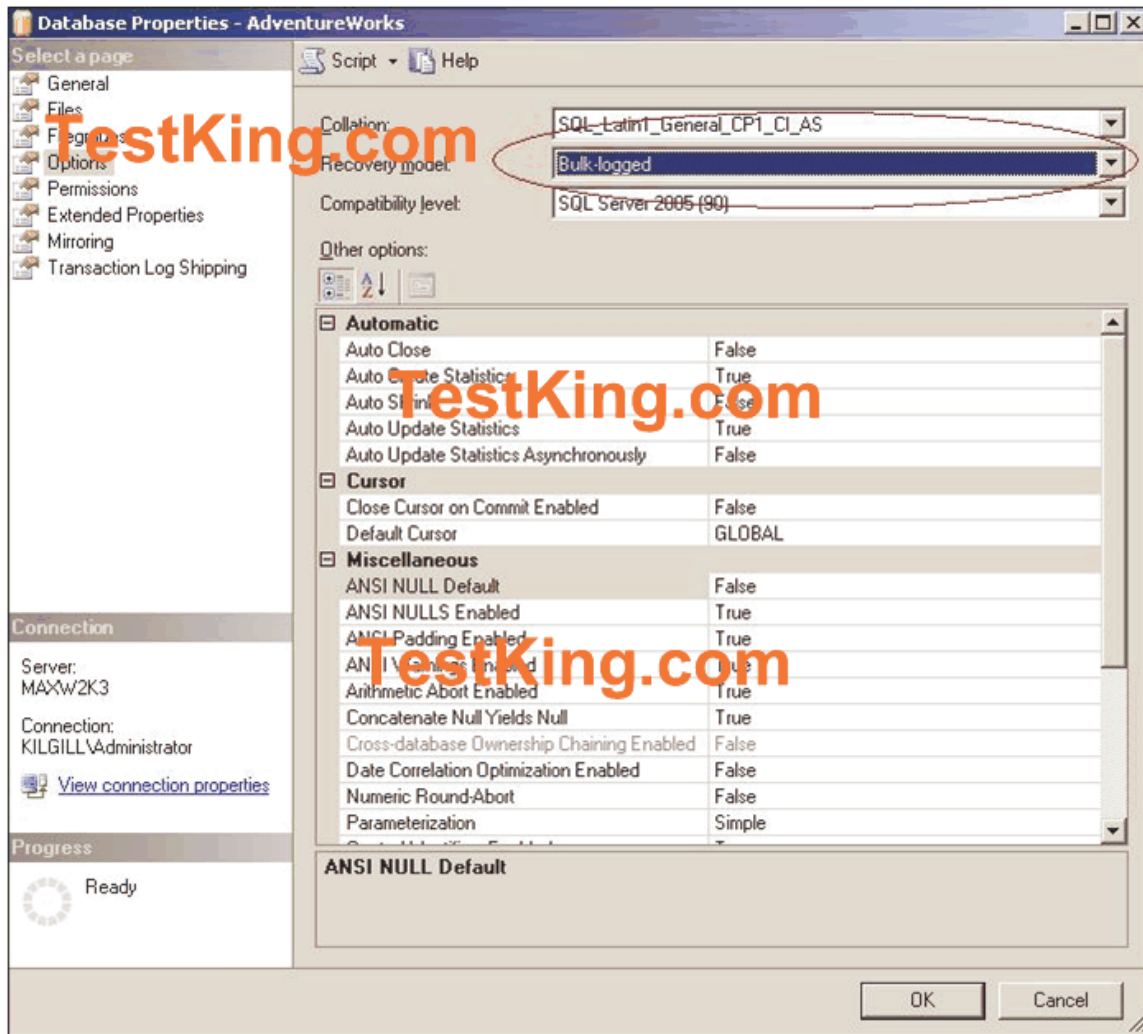
**The TestKing database is the only non-system database on the server. Transaction log backups are part of the backup routine for the TestKing database.**

**Each night, several SQL Server Integration Services (SSIS) packages are run to import 5 million rows of data.**

**You need to minimize the log space that is used by these import operations. You must modify the appropriate server and database settings by using the Database Properties dialog box at the Server Properties dialog box, or both of these dialog boxes, in SQL Server Management Studio (SSMS).**

**Answer:**

**Explanation:**



## QUESTION NO: 8 SIMULATION

You work as a database administrator at TestKing.com.

Users report that they receive an error message when they execute a CLR user-defined function named Production.ufnGetProductionInfo in the TestKing database. The error is occurring because the function is being denied access to the local file system. You must ensure that the least permissions possible to solve the problem are applied.

You need to allow the function to access the local file systems by using SQL Server Management Studio (SSMS).

You must apply the minimum number of possible changes.

**Answer:**

**Explanation: Here is a hint toward a possible solution 8:**

SQL Server Deployment

The resulting assembly must somehow be "installed" in SQL Server before it can be used. This process is known as deployment. As you can see in Figure 6, Visual Studio has options on the Build menu that can be used to deploy the assembly, but there are other steps after this that must be completed using T-SQL. Therefore, let's switch to SQL Server Management Studio (SSMS) and perform the entire deployment process using T-SQL. Refer to Listing 3. As a security measure, the database must first be configured to allow CLR integration. This is accomplished by executing the sp\_configure and reconfigure statements as shown. Next, the assembly must be imported into SQL Server (yes, imported - the assembly is actually stored in the database). The new CREATE ASSEMBLY statement is used for this. Using this statement, we give the assembly a name (MyFuncs), specify the path from which to read the assembly, and indicate the permission level we wish to give to the assembly.

Permission levels are very important. They control the rights that an assembly has with regard to resources outside of SQL Server. The most



restrictive - and recommended - setting is SAFE. It allows no access to external resources. The next level, EXTERNAL\_ACCESS, allows the use of the file system, the network, and the registry. The final level, UNSAFE, should be used with extreme caution. It allows access to any external resource to which the SQL service account has access. This can be very dangerous and should be allowed only for the most trusted assemblies. Whichever setting is selected, it is important to note that not all Framework classes can be used when creating SQL objects. Only those classes marked as safe for SQL Server can be used. Excluded classes include those that could cause SQL Server instability, such as those in the System.Threading namespace. See the Framework documentation for specifics.

When the CREATE ASSEMBLY statement has been executed, the assembly will be imported into the system. The next step is to create a T-SQL wrapper for the user-defined function. Refer again to Listing 3. The wrapper is created by using the new EXTERNAL NAME clause. We start by creating a regular T-SQL function, but instead of writing the function implementation, the EXTERNAL NAME clause is used. This clause specifies the fully qualified name of the .NET function, starting with the assembly name given in the CREATE ASSEMBLY statement.

#### **QUESTION NO: 9 SIMULATION**

**You work as a database administrator at TestKing.com.**

**You are required to create a maintenance plan to rebuild indexes in the TestKing database. The indexes must be rebuilt using a fill factor of 90.**

**The fill factor must be applied to both the leaf level and the balanced tree of the index. While the maintenance plan is running, all tables must be available for queries to be run.**

**You need to accomplish this task by completing the Define Rebuild Index Task page in the Maintenance Plan Wizard.**

**Answer:**

**Explanation:**



#### QUESTION NO: 10 SIMULATION

You work as a database administrator at TestKing.com.

You need to move the TestKing database to a different SQL Server assistance. You must perform a backup that contains all of the changes since the last transaction log backup. You must ensure that no changes can be applied to the database after the backup completes.

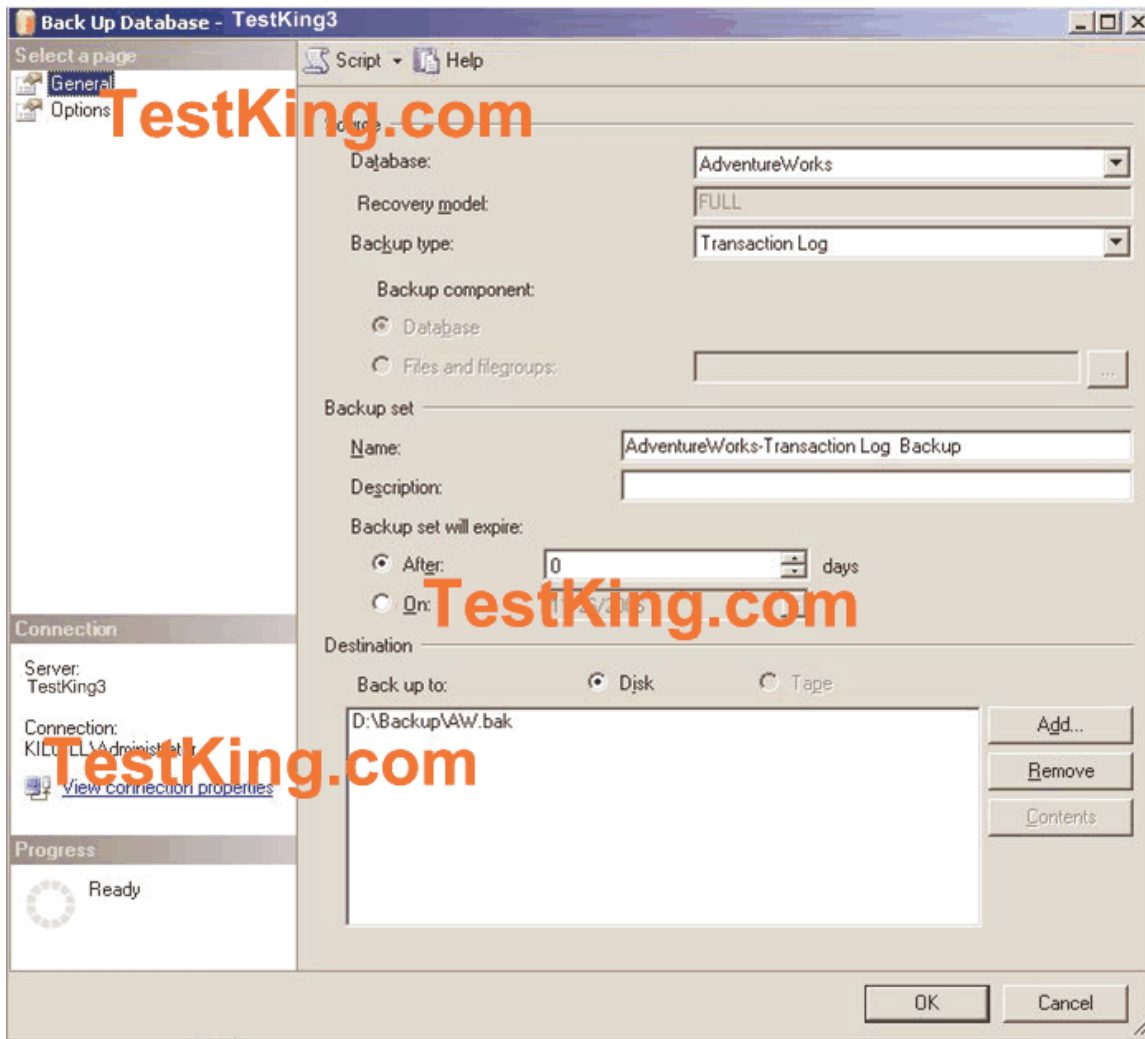
**The backup file must be named AW.bak, and it must be placed in the D:\Backup folder.**

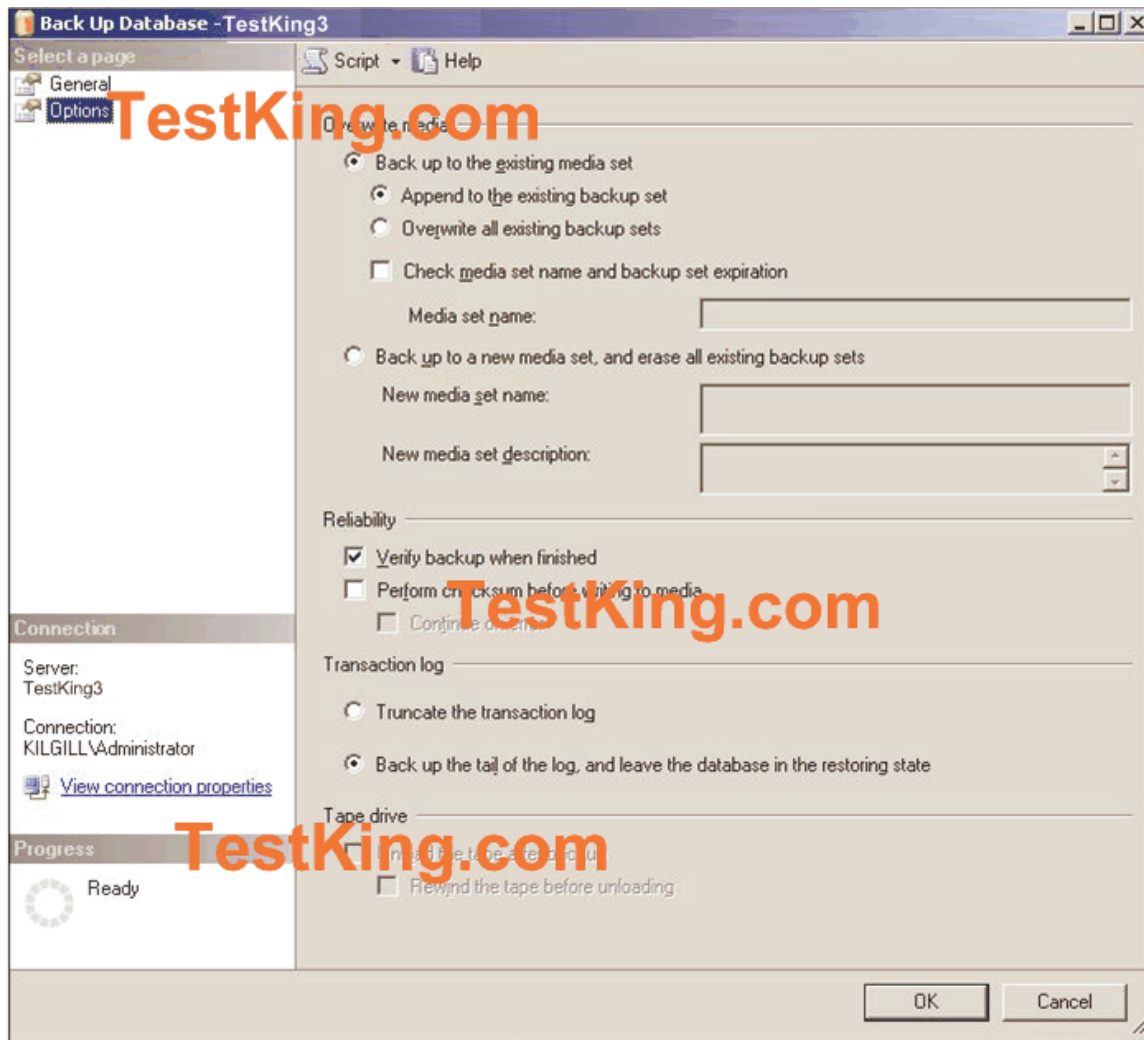
**You need to perform the backup by using the Back Up Database dialog box.**

**Answer:**

**Explanation:**

TestKing.com





**QUESTION NO: 11 SIMULATION**

**Exhibit:**

Volume	RAID level	Free space
C: (Windows volume)	RAID 1	10 GB
D:	RAID 1	20 GB
E:	RAID 0	20 GB

You work as a database administrator at TestKing.com. You need to create a new database named Sales.

The new database must meet the following requirements.

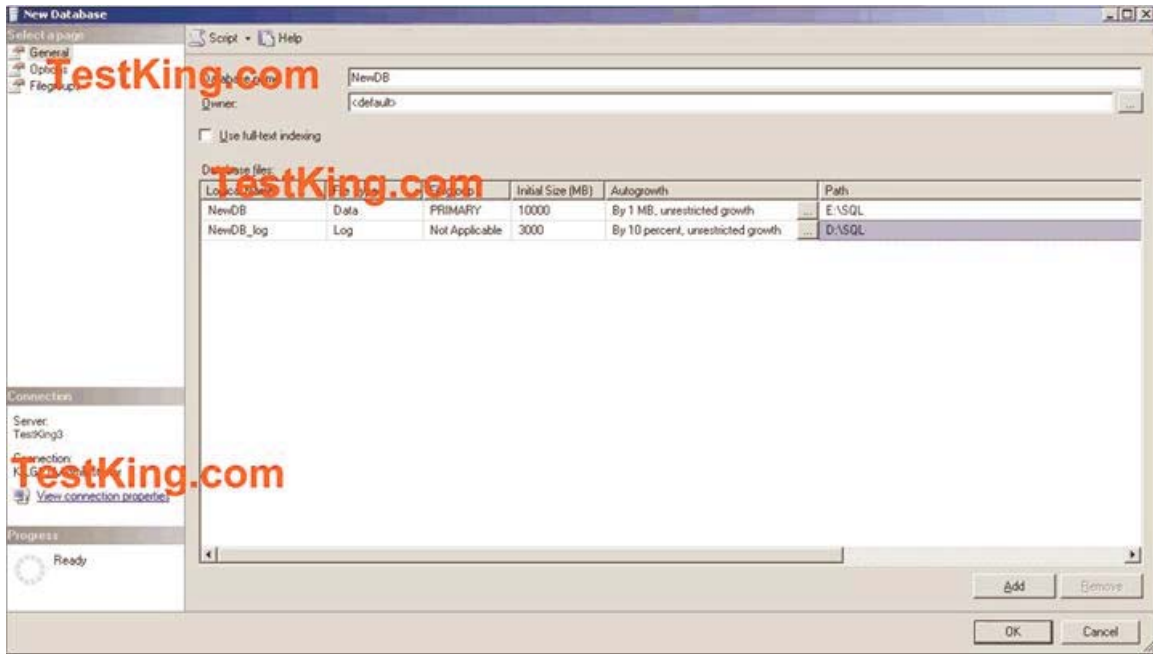
- \* the data file must be 10,000 MB in size
- \* The log file must be 3,000 MB in size.
- \* Both files should have the default settings for automatic file growth and maximum file size.
- \* The transaction log must be stored on a fault-tolerant volume.
- \* The data file, the transaction log file, and the Windows installation must all be located on different volumes.
- \* The database file must be stored in a folder named SQL in the root of each volume.
- \* The transaction log must be set to be automatically truncated

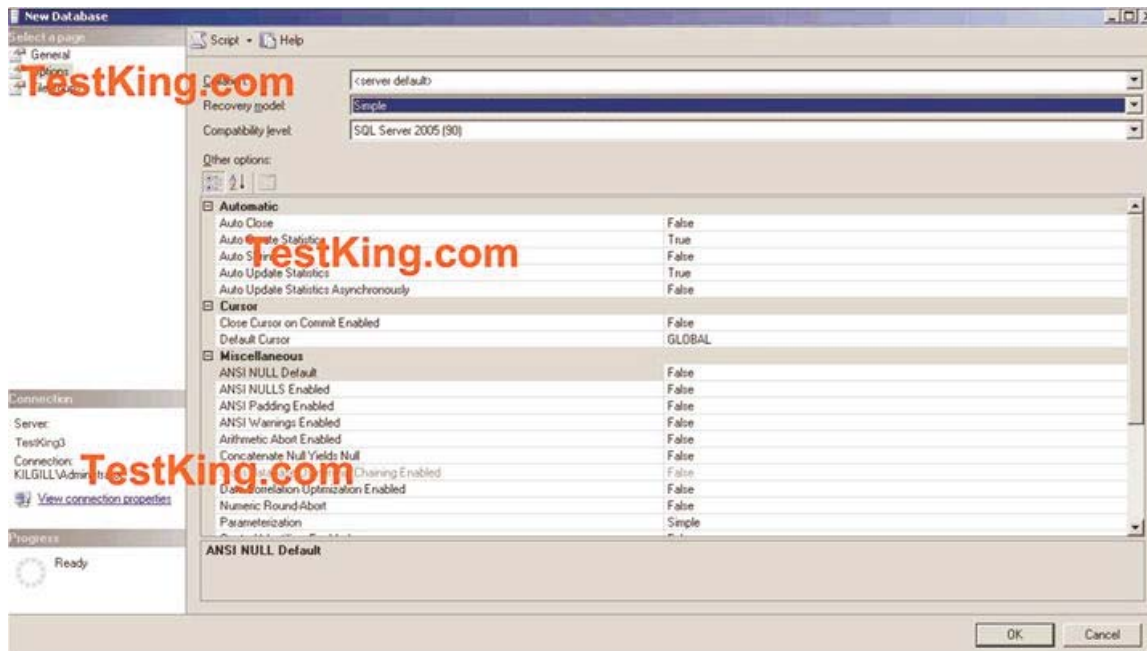
The available volumes are shown in the exhibit.

You need to create the Sales database by using the New Database dialog box.

**Answer:**

**Explanation:**





**QUESTION NO: 12 SIMULATION**

**You work as a database administrator at TestKing.com. TestKing.com has two SQL Server 2005 Servers named TestKing1 and TestKing2.**

**You need to add TestKing2 as a linked server on TestKing1. Only the login named TestKingApp1 should be able to access TestKing2 through TestKing1, A login named TestKingApp2 with the password testking must be used for the linked server connection to TestKing2. You must apply the fewest possible settings.**



**You need to create the linked server by using the New Linked Server dialog box on TestKing1.**

**Answer:**

**Explanation:**

TestKing.com

