

Exam Questions 70-464

Developing Microsoft SQL Server 2012 Databases

<https://www.2passeasy.com/dumps/70-464/>



NEW QUESTION 1

- (Exam Topic 1)

Which data type should you use for CustomerID?

- A. varchar(11)
- B. bigint
- C. nvarchar(11)
- D. char(11)

Answer: D

Explanation:

Invoices.xml

All customer IDs are 11 digits. The first three digits of a customer ID represent the customer's country. The remaining eight digits are the customer's account number.

int: -2^{31} (-2,147,483,648) to $2^{31}-1$ (2,147,483,647) (just 10 digits max)

bigint: -2^{63} (-9,223,372,036,854,775,808) to $2^{63}-1$ (9,223,372,036,854,775,807)

<http://msdn.microsoft.com/en-us/library/ms176089.aspx> <http://msdn.microsoft.com/en-us/library/ms187745.aspx>

NEW QUESTION 2

- (Exam Topic 1)

You need to create the InvoiceStatus table in DB1.

How should you define the InvoiceID column in the CREATE TABLE statement?

- A. InvoiceID bigint
DEFAULT (NEXT VALUE FOR Accounting.InvoiceID_Seq) NOT NULL,
- B. InvoiceID bigint DEFAULT ((NEXT VALUE
FOR Accounting.InvoiceID_Seq OVER
(ORDER BY InvoiceStatusID))) NOT NULL FOREIGN
KEY REFERENCES Accounting.Invoices(InvoiceID),
- C. InvoiceID bigint FOREIGN KEY REFERENCES
Accounting.Invoices(InvoiceID) NOT NULL,
- D. InvoiceID bigint DEFAULT ((NEXT VALUE
FOR Accounting.InvoiceID_Seq
OVER (ORDER BY InvoiceStatusID))) NOT NULL,

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

NEW QUESTION 3

- (Exam Topic 1)

You need to modify InsertInvoice to comply with the application requirements. Which code segment should you execute?

- A. OPEN CERT1;
ALTER PROCEDURE Accounting.usp_InsertInvoice
WITH ENCRYPTION;
CLOSE CERT1;
- B. OPEN CERT2;
ALTER PROCEDURE Accounting.usp_InsertInvoice
WITH ENCRYPTION;
CLOSE CERT2;
- C. ADD SIGNATURE TO Accounting.usp_InsertInvoice
BY CERTIFICATE CERT1;
- D. ADD SIGNATURE TO Accounting.usp_InsertInvoice
BY CERTIFICATE CERT2;

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: D

NEW QUESTION 4

- (Exam Topic 1)

You need to modify the function in CountryFromID.sql to ensure that the country name is returned instead of the country ID. Which line of code should you modify in CountryFromID.sql?

- A. 04
B. 05
C. 06
D. 19

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms186755.aspx> <http://msdn.microsoft.com/en-us/library/ms191320.aspx>

NEW QUESTION 5

- (Exam Topic 1)

You execute IndexManagement.sql and you receive the following error message: "Msg 512, Level 16, State 1, Line 12 Subquery returned more than 1 value. This is not permitted when the subquery follows =, !=, <, <=, >, >= or when the subquery is used as an expression."

You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 18?

- A. WHILE SUM(@RowNumber) < (SELECT @counter FROM @indextable)
B. WHILE @counter < (SELECT COUNT(RowNumber) FROM @indextable)
C. WHILE COUNT(@RowNumber) < (SELECT @counter FROM @indextable)
D. WHILE @counter < (SELECT SUM(RowNumber) FROM @indextable)

Answer: B

NEW QUESTION 6

- (Exam Topic 1)

You are testing disaster recovery procedures.

You attempt to restore DB1 to a different server and you receive the following error message: "Msg 33111. Level 16, State 3, Line 1
Cannot find server certificate with thumbprint
,0xA694FBEA88C9354E5E2567C30A2A69E8FB4C44A9\
Msg 3013, Level 16, State 1, Line 1
RESTORE DATABASE is terminating abnormally."
You need to ensure that you can restore DB1 to a different server. Which code segment should you execute?

- A.

```
RESTORE CERTIFICATE CERT2
FROM FILE='CERT2.CER'
WITH PRIVATE KEY (FILE = 'CERT2.KEY',
DECRYPTION BY PASSWORD='p@ssw0rd1');
```
- B.

```
CREATE CERTIFICATE CERT1
FROM FILE='CERT1.CER'
WITH PRIVATE KEY (FILE = 'CERT1.KEY',
DECRYPTION BY PASSWORD='p@ssw0rd1');
```
- C.

```
CREATE CERTIFICATE CERT2
ENCRYPTION BY PASSWORD='p@ssw0rd1'
WITH SUBJECT = 'EncryptionCertificate';
```
- D.

```
CREATE CERTIFICATE CERT1
ENCRYPTION BY PASSWORD='p@ssw0rd1'
WITH SUBJECT = 'EncryptionCertificate';
```

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

NEW QUESTION 7

- (Exam Topic 2)

Developers report that usp_UpdateSessionRoom periodically returns error 3960.

You need to prevent the error from occurring. The solution must ensure that the stored procedure returns the original values to all of the updated rows.

What should you configure in Procedures.sql?

- A. Replace line 46 with the following code:SET TRANSACTION ISOLATION LEVEL SERIALIZABLE
B. Replace line 46 with the following code:SET TRANSACTION ISOLATION LEVEL REPEATABLE READ
C. Move the SELECT statement at line 49 to line 57.
D. Move the SET statement at line 46 to line 53.

Answer: A

NEW QUESTION 8

- (Exam Topic 2)

You need to recommend a solution to ensure that SQL1 supports the auditing requirements of usp_UpdateSpeakerName.

What should you include in the recommendation?

- A. The Distributed Transaction Coordinator (DTC)
B. Transactional replication
C. Change data capture
D. Change tracking

Answer: A

NEW QUESTION 9

- (Exam Topic 2)

You need to provide referential integrity between the Sessions table and Speakers table. Which code segment should you add at line 47 of Tables.sql?

- A. ALTER TABLE dbo.Sessions ADD CONSTRAINT FK_Sessions_Speakers FOREIGN KEY (SessionID) REFERENCES dbo.Speakers (SpeakerID);
- B. ALTER TABLE dbo.Sessions ADD CONSTRAINT FK_Sessions_Speakers FOREIGN KEY (SpeakerID) REFERENCES dbo.Speakers (SpeakerID);
- C. ALTER TABLE dbo.Speakers ADD CONSTRAINT FK_Speakers_Sessions FOREIGN KEY (SpeakerID) REFERENCES dbo.Sessions (SessionID);
- D. ALTER TABLE dbo.Speakers ADD CONSTRAINT FK_Speakers_Sessions FOREIGN KEY (SessionID) REFERENCES dbo.Sessions (SessionID);

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms189049.aspx> <http://msdn.microsoft.com/en-us/library/ms179610.aspx> <http://msdn.microsoft.com/en-us/library/ff878370.aspx>

NEW QUESTION 10

- (Exam Topic 2)

You need to create the object used by the parameter of usp_InsertSessions. Which statement should you use?

- A. CREATE XML SCHEMA COLLECTION SessionDataTable
B. CREATE TYPE SessionDataTable AS Table
C. CREATE SCHEMA SessionDataTable
D. CREATE TABLE SessionDataTable

Answer: B

NEW QUESTION 10

- (Exam Topic 2)

You need to ensure that if any of the statements in usp_UpdateSpeakerName return an error message, all of the changes executed by usp_UpdateSpeakerName are not committed to the database.

What should you do in Procedures.sql? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Add the following at line 17:
 ROLLBACK TRANSACTION
- B. Add the following at line 05:
 BEGIN TRANSACTION SpeakerUpdate
- C. Add the following at line 05:
 SAVE TRANSACTION SpeakerUpdate
- D. Add the following at line 17:
 ROLLBACK TRANSACTION SpeakerUpdate
- E. Add the following at line 07:
 BEGIN TRANSACTION

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: BD

NEW QUESTION 12

- (Exam Topic 3)

You need to create the object used by the parameter of usp_UpdateEmployeeName. Which code segment should you use?

- A. CREATE XML SCHEMA COLLECTION EmployeesInfo
- B. CREATE TYPE EmployeesInfo AS Table
- C. CREATE SCHEMA EmployeesInfo
- D. CREATE TABLE EmployeesInfo

Answer: B

Explanation:

Example Usage of Table-Valued Parameters (Database Engine)

<http://msdn.microsoft.com/en-us/library/bb510489.aspx>

(Benefits of using Table-Valued Parameters)

/* Create a table type. */

CREATE TYPE LocationTableType AS TABLE (LocationName VARCHAR(50)

, CostRate INT); GO

/* Create a procedure to receive data for the table-valued parameter. */ CREATE PROCEDURE dbo. usp_InsertProductionLocation

@TVP LocationTableType READONLY AS

SET NOCOUNT ON

INSERT INTO AdventureWorks2012.Production.Location (Name

,CostRate

,Availability

,ModifiedDate)

SELECT *, 0, GETDATE() FROM @TVP;

GO

Also:

<http://msdn.microsoft.com/en-us/library/ms175007.aspx>(CREATE TYPE *tabletypename* AS TABLE)

<http://msdn.microsoft.com/en-us/library/ms175010.aspx>(table data types)

Wrong Answers:

<http://msdn.microsoft.com/en-us/library/ms174979.aspx>(CREATE TABLE) <http://msdn.microsoft.com/en-us/library/ms189462.aspx>(CREATE SCHEMA)

<http://msdn.microsoft.com/en-us/library/ms176009.aspx>(CREATE XML SCHEMA COLLECTION)

NEW QUESTION 13

- (Exam Topic 3)

You need to provide referential integrity between the Offices table and Employees table.

Which code segment or segments should you add at line 27 of Tables.sql? (Each correct answer presents part of the solution. Choose all that apply.)

- A. ALTER TABLE dbo.Offices ADD CONSTRAINT
PK_Offices_EmployeeID PRIMARY KEY (EmployeeID);
- B. ALTER TABLE dbo.Employees ADD CONSTRAINT
FK_Employees_Offices FOREIGN KEY (OfficeID)
REFERENCES dbo.Offices (OfficeID);
- C. ALTER TABLE dbo.Employees ADD CONSTRAINT
PK_Employees_EmployeeID PRIMARY KEY (EmployeeID);
- D. ALTER TABLE dbo.Offices ADD CONSTRAINT
FK_Offices_Employees FOREIGN KEY (EmployeeID)
REFERENCES dbo.Employees (EmployeeID);

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: CD

Explanation:

<http://msdn.microsoft.com/en-us/library/ms189049.aspx>

NEW QUESTION 17

- (Exam Topic 3)

You execute usp_SelectEmployeesByName multiple times, passing strings of varying lengths to @LastName. You discover that usp_SelectEmployeesByName uses inefficient execution plans.

You need to update usp_SelectEmployeesByName to ensure that the most efficient execution plan is used. What should you add at line 31 of StoredProcedures.sql?

- A. OPTION (ROBUST PLAN)
B. OPTION (OPTIMIZE FOR UNKNOWN)
C. OPTION (KEEP PLAN)
D. OPTION (KEEPFIXED PLAN)

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms181714.aspx>

NEW QUESTION 19

- (Exam Topic 4)

You need to create a function that will use a SELECT statement in ProductsByProductType.sql. Which code segment should you use to complete the function?

- A. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS @tblInvoices TABLE (ProductID bigint, ProductType varchar(11), CreationDate date)
AS
INSERT INTO @tblInvoices
- B. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS TABLE
AS
RETURN
- C. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS @TblInvoices TABLE (ProductID bigint, ProductType varchar(11), CreationDate date)
AS
- D. CREATE FUNCTION Production.fnProductsByProductType (@ProductType varchar(11))
RETURNS xml
AS
RETURN

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: B

Explanation:

<http://msdn.microsoft.com/en-us/library/ms191320.aspx> <http://msdn.microsoft.com/en-us/library/ms186755.aspx>

NEW QUESTION 22

- (Exam Topic 4)

You are testing disaster recovery procedures.

When you attempt to restore ProductsDB to another server, you receive the following error message: "Msg 33111, Level 16, State 3, Line 5

Cannot find server certificate with thumbprint '0x9D876A3468B911EIBA4CFCBF4724019B\ Msg 3013, Level 16, State 1, Line 5

RESTORE DATABASE is terminating abnormally."

You need to ensure that you can restore ProductsDB to another server. Which code segment should you execute on the other server?

- A. RESTORE CERTIFICATE DBCERT
FROM FILE='DBCERT.CER'
WITH PRIVATE KEY (FILE = 'c:\DBCERT.KEY',
DECRYPTION BY PASSWORD = 'SecretP@ss');
- B. CREATE CERTIFICATE PRODUCTSCERT
ENCRYPTION BY PASSWORD = 'SecretP@ss'
WITH SUBJECT = 'SecurityCertificate';
- C. CREATE CERTIFICATE DBCERT
ENCRYPTION BY PASSWORD = 'SecretP@ss'
WITH SUBJECT = 'SecurityCertificate';
- D. CREATE CERTIFICATE PRODUCTSCERT
FROM FILE='PRODUCTSCERT.CER'
WITH PRIVATE KEY (FILE = 'c:\PRODUCTSCERT.KEY',
DECRYPTION BY PASSWORD = 'SecretP@ss');

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: D

NEW QUESTION 25

- (Exam Topic 4)

You execute IndexManagement.sql and you receive the following error message: "Msg 512, Level 16, State 1, Line 12 Subquery returned more than 1 value. This is not permitted when the subquery follows =, !=, <, <=, >, >= or when the subquery is used as an expression."

You need to ensure that IndexManagement.sql executes properly. Which WHILE statement should you use at line 18?

- A. WHILE SUM(@RowNumber) < (SELECT @counter FROM @indextable)
- B. WHILE @counter < (SELECT SUM(RowNumber) FROM @indextable)
- C. WHILE COUNT(@RowNumber) < (SELECT @counter FROM @indextable)
- D. WHILE @counter < (SELECT COUNT(RowNumber) FROM @indextable)

Answer: D

NEW QUESTION 30

- (Exam Topic 4)

Which code segment should you use to define the ProductDetails column?

- A. ProductDetails xml (DOCUMENT Production.ProductDetailsSchema) NULL
- B. ProductDetails xml NULL
- C. ProductDetails xml (CONTENT Production.ProductDetailsSchema) NULL
- D. ProductDetails varchar(MAX) NULL

Answer: D

NEW QUESTION 32

- (Exam Topic 5)

You need to implement a change to usp_ExportOpenings that meets the integration requirements.

What should you modify in usp_ExportOpenings? (Each correct answer presents part of the solution. Choose all that apply?)

- A. To the end of line 04, add [Opening].
- B. To the end of line 05, add [Opening! title].
- C. To line 10, add FOR XML RAW.
- D. To line 10, add FOR XMLEXPLICIT.
- E. To line 10, add FOR XML AUTO.
- F. To the end of line 04, add [Opening!ELEMENT].
- G. To the end of line 06, add [Opening!salary!ELEMENT].
- H. To the end of line 05, add [Opening!title!ELEMENT].
- I. To the end of line 06, add [Opening! salary].

Answer: ABEI

Explanation:

The AUTO mode generates nesting in the resulting XML by using heuristics based on the way the SELECT statement is specified. You have minimal control over the shape of the XML generated. The nested FOR XML queries can be written to generate XML hierarchy beyond the XML shape that is generated by AUTO mode heuristics.

NEW QUESTION 37

- (Exam Topic 5)

You need to create a script that automates the export of the XML data. The script must meet the integration requirements.

What should you include in the script?

- A. The CREATE SERVER ROLE command and the sp_reassign_proxy, sp_add_job, sp_add_jobstep, and sp_grant_login_to_proxy system stored procedures.
- B. The CREATE CREDENTIAL command and the sp_add_proxy, sp_add_job, sp_add_jobstep, and sp_grant_proxy_to_subsystem system stored procedures.
- C. The CREATE CREDENTIAL command and the sp_reassign_proxy, sp_add_job, sp_add_jobstep, and sp_grant_login_to_proxy system stored procedures.
- D. The CREATE SERVER ROLE command and the sp_add_proxy, sp_add_job, sp_add_jobstep, and sp_grant_proxy_to_subsystem system stored procedures.

Answer: B

NEW QUESTION 39

- (Exam Topic 5)

You need to implement a solution that meets the job application requirements. What should you do?

- A. Create a one-to-one relationship between the Openings table and the Applications table.
- B. Create a one-to-one relationship between the Candidates table and the Applications table.
- C. Add a UNIQUE constraint to the Applications table on the ApplicationID column and CandidateID column.
- D. Add a UNIQUE constraint to the Applications table on the OpeningID column and the CandidateIDcolumn.

Answer: D

NEW QUESTION 42

- (Exam Topic 5)

You need to resolve the performance issues of the usp_ExportOpenings stored procedure. The solution must minimize the amount of hard disk space used.

Which statement should you execute on DB1?

- A. EXEC sp_dboption 'DB1', 'auto create statistics', 'TRUE';

- B. CREATE INDEX IX_Exp_Openings ON Openings(PostDate, FilledDate) INCLUDE (Description, Title, Salary);
C. CREATE INDEX IX_Exp_Openings ON Openings(PostDate) INCLUDE (Description, Title, Salary) WHERE FilledDate IS NULL;
D. EXEC sp_recompile 'usp_ExportOpenings';

Answer: C

NEW QUESTION 46

- (Exam Topic 6)

You need to implement a solution that addresses the performance issues of the usp_GetOrdersByProduct stored procedure. Which statement should you execute?

- A. CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (OrderID, LineItem, UnitPrice, Total, Discount)
- B. CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (LineItem, Quantity, UnitPrice, Total, Discount)
- C. CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
- D. CREATE INDEX IX_OrderDetails_ByProduct
ON OrderDetails (ProductID)
INCLUDE (LineItem, Quantity, UnitPrice, Discount)

- A. Option A
B. Option B
C. Option C
D. Option D

Answer: C

NEW QUESTION 50

- (Exam Topic 6)

You plan to create a stored procedure that inserts data from an XML file to the OrderDetails table. The following is the signature of the stored procedure:

```
CREATE PROCEDURE usp_InsertItems  
@items XML (ValidateOrder)
```

The following is the XSD file used to create the ValidateOrder schema collection:

```
<?xml version="1.0" encoding="UTF-16"?>
<xsd:schema
  xmlns:xsd="http://www.w3.org/2001/XMLSchema" >
<xsd:element name="root">
  <xsd:complexType mixed="true">
    <xsd:sequence>
      <xsd:element name="Product"
        minOccurs="1" maxOccurs="unbounded">
        <xsd:complexType mixed="true">
          <xsd:sequence>
            <xsd:element name="UnitPrice" type="xsd:decimal"
              minOccurs="1" maxOccurs="1" />
            <xsd:element name="Quantity" type="xsd:integer"
              minOccurs="1" maxOccurs="1" />
          </xsd:sequence>
          <xsd:attribute name="lineItem"
            type="xsd:integer" use="required"/>
          <xsd:attribute name="productID"
            type="xsd:integer" use="required"/>
        </xsd:complexType>
      </xsd:element>
    </xsd:sequence>
    <xsd:attribute name="numberItems"
      type="xsd:integer" use="required"/>
  </xsd:complexType>
</xsd:element>
</xsd:schema>
```

You develop a code segment that retrieves the number of items and loops through each item. Each time the loop runs, a variable named @itemNumber is incremented.

You need to develop a code segment that retrieves the product ID of each item number in the loop. Which code segment should you develop?

- A. SET @productID = @items.value'/Root/Product/productID', int)
- B. SET @productID = @items.value'/Root/Product['+ @itemNumber+ ']/@productID', int)
- C. SET @productID = @items.value'/Root/Product['+ @itemNumber+ ']/productID', int)
- D. SET @productID = @items.value'/Root/Product/@productID', int)

Answer: B

NEW QUESTION 51

- (Exam Topic 6)

You need to implement a solution that meets the site requirements. What should you implement?

- A. A non-indexed view on Server1
- B. A non-indexed view on Server2
- C. A distributed view on Server1
- D. A distributed view on Server2

Answer: C

Explanation:

A partitioned view is a view defined by a UNION ALL of member tables structured in the same way, but stored separately as multiple tables in either the same instance of SQL Server or in a group of autonomous instances of SQL Server servers, called federated database servers.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/create-view-transact-sql?view=sql-server-2017>

NEW QUESTION 54

- (Exam Topic 6)

You need to implement a solution that addresses the bulk insert requirements. What should you add to line 08 in usp_ImportOrderDetails?

- A. LASTROW=0.
- B. BATCHSIZE=0.
- C. BATCHSIZE=1000.
- D. LASTROW = 1000.

Answer: C

NEW QUESTION 56

- (Exam Topic 6)

You need to ensure that a new execution plan is used by usp_GetOrdersByProduct each time the stored procedure runs. What should you do?

- A. Execute sp_help 'usp_GetOrdersByProduct'.
- B. Execute sp_recompile 'usp_GetOrdersByProduct'.
- C. Add WITH RECOMPILE to line 03 in usp_GetOrdersByProduct.
- D. Add WITH (FORCESEEK) to line 07 in usp_GetOrdersByProduct.

Answer: C

Explanation:

Ref: [http://msdn.microsoft.com/en-us/library/ms190439\(v=sql.90\).aspx](http://msdn.microsoft.com/en-us/library/ms190439(v=sql.90).aspx)

NEW QUESTION 60

- (Exam Topic 6)

You discover that the usp_GetOrdersAndItems stored procedure takes a long time to complete while usp_AddOrder or usp_AddXMLOrder run. You need to ensure that usp_GetOrdersAndItems completes as quickly as possible.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Set the isolation level of the usp_GetOrdersAndItems stored procedure to SERIALIZABLE.
- B. Execute the ALTER DATABASE Sales SET ALLOW_SNAPSHOT_ISOLATION ON statement.
- C. Set the isolation level of the usp_AddOrder stored procedure to SERIALIZABLE.
- D. Set the isolation level of the usp_GetOrdersAndItems stored procedure to SNAPSHOT.
- E. Set the isolation level of the usp_AddOrder stored procedure to SNAPSHOT.
- F. Execute the ALTER DATABASE Sales SET ALLOW_SNAPSHOT_ISOLATION OFF statement.

Answer: BD

NEW QUESTION 63

- (Exam Topic 6)

You need to modify the Orders table to store the XML data used by the retailers. Which statement should you execute?

- A. ALTER Orders ADD originalOrder XML (ValidateOrder);
- B. ALTER Orders ADD originalOrder XML;
- C. ALTER Orders ADD originalOrder varchar(max);
- D. ALTER Orders ADD originalOrder varbinary(max);

Answer: D

NEW QUESTION 64

- (Exam Topic 7)

You need to modify the stored procedure usp_LookupConcurrentUsers. What should you do?

- A. Use the summary table as an in-memory optimized table with a non-hash clustered index.
- B. Use the summary table as an in-memory optimized table with a non-hash nonclustered index.
- C. Use a type variable instead of the summary table.
- D. Add a clustered index to the summary table.

Answer: A

NEW QUESTION 68

- (Exam Topic 7)

You need to optimize the index and table structures for POSTransaction.

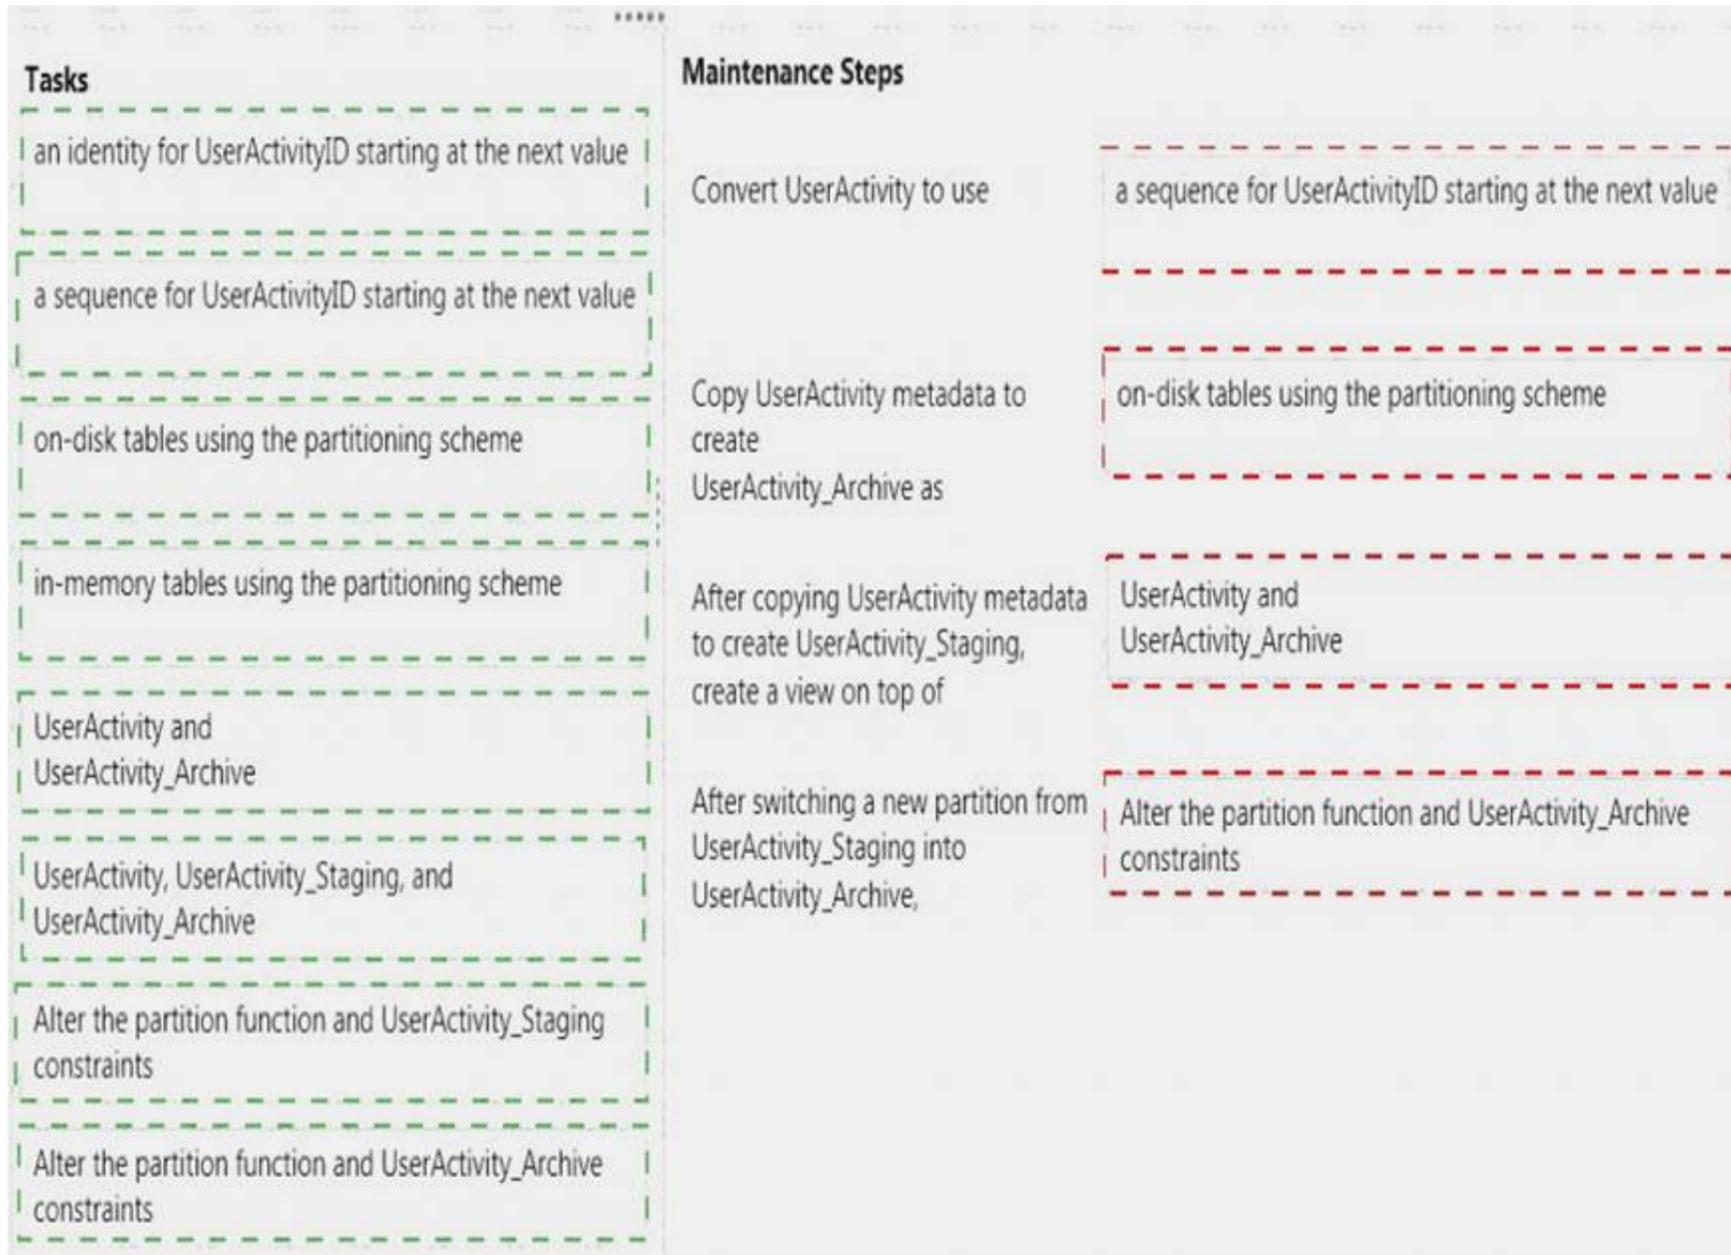
Which task should you use with each maintenance step? To answer, drag the appropriate tasks to the correct maintenance steps. Each task may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Tasks	Maintenance Steps
an identity for UserActivityID starting at the next value	Convert UserActivity to use Task
a sequence for UserActivityID starting at the next value	
on-disk tables using the partitioning scheme	Copy UserActivity metadata to create Task
	UserActivity_Archive as
in-memory tables using the partitioning scheme	After copying UserActivity metadata to create UserActivity_Staging, create a view on top of Task
UserActivity and UserActivity_Archive	
UserActivity, UserActivity_Staging, and UserActivity_Archive	After switching a new partition from UserActivity_Staging into UserActivity_Archive, Task
Alter the partition function and UserActivity_Staging constraints	
Alter the partition function and UserActivity_Archive constraints	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 71

- (Exam Topic 7)

You need to monitor the health of your tables and indexes in order to implement the required index maintenance strategy. What should you do?

- A. Query system DMVs to monitor avg_chain_length and max_chain_lengt
- B. Create alerts to notify you when these values converge.
- C. Create a SQL Agent alert when the File Table: Avg time per file I/O request value is increasing.
- D. Query system DMVs to monitor total_bucket_coun
- E. Create alerts to notify you when this value increases.
- F. Query system DMVs to monitor total_bucket_coun
- G. Create alerts to notify you when this value decreases.

Answer: A

Explanation:

From scenario:

- * You need to anticipate when POSTransaction table will need index maintenance.
- * The index maintenance strategy for the UserActivity table must provide the optimal structure for both maintainability and query performance.

NEW QUESTION 73

- (Exam Topic 7)

You need to modify the stored procedure usp_LookupConcurrentUsers. What should you do?

- A. Add a clustered index to the summary table.
- B. Add a nonclustered index to the summary table.
- C. Add a clustered columnstore index to the summary table.
- D. Use a table variable instead of the summary table.

Answer: A

Explanation:

Scenario: Query the current open micropayments for users who own multiple micropayments by using a stored procedure named usp.LookupConcurrentUsers

NEW QUESTION 78

- (Exam Topic 7)

You need to create the usp.AssignUser stored procedure.

Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks	Answer Area
<pre>IF @StatusID IS NULL RAISERROR (N'The transaction does not exist.',16,1)</pre>	
<pre>WITH NATIVE_COMPILATION, SCHEMABINDING, EXECUTE AS OWNER</pre>	
<pre>CREATE PROCEDURE dbo.usp_AssignUser @UserId int, @POSTransactionId int</pre>	
<pre>WITH (TRANSACTION ISOLATION LEVEL = READ COMMITTED, LANGUAGE = N'us_english')</pre>	
<pre>UPDATE dbo.POSTransaction SET UserId=@UserId WHERE POSTransactionId=@POSTransactio nId END</pre>	
<pre>AS BEGIN</pre>	
<pre>DECLARE @StatusID int SELECT @StatusID=StatusId FROM dbo.POSTransaction WHERE POSTransactionId=@POSTransactionI d</pre>	
<pre>IF @StatusID IS NULL THROW 51000, N'The transaction does not exist.', 1</pre>	
<pre>WITH (TRANSACTION ISOLATION LEVEL = REPEATABLE READ, LANGUAGE = N'us_english')</pre>	
<pre>AS BEGIN ATOMIC</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

```
CREATE PROCEDURE dbo.usp_AssignUser
@UserId int, @POSTransactionId int
```

Box 2:

```
WITH
NATIVE_COMPILATION, SCHEMABINDING,
EXECUTE AS OWNER
```

Box 3:

```
AS
BEGIN ATOMIC
```

Box 4:

```
WITH (TRANSACTION ISOLATION LEVEL =
REPEATABLE READ, LANGUAGE
= N'us_english')
```

Box 5:

```
UPDATE dbo.POSTransaction
SET UserId=@UserId
WHERE POSTransactionId=@POSTransactio
nId
END
```

Box 6:

```
DECLARE @StatusID int
SELECT @StatusID=StatusId
FROM dbo.POSTransaction
WHERE POSTransactionId=@POSTransactionI
d
```

Box 7:

```
IF @StatusID IS NULL
THROW 51000, N'The transaction
does not exist.', 1
```

Note:

* From scenario: The mobile application will need to meet the following requirements:

/Communicate with web services that assign a new user to a micropayment by using a stored procedure named usp_AssignUser.

* Example:

```
create procedure dbo.OrderInsert(@OrdNo integer, @CustCode nvarchar(5))
```

```
with native_compilation, schemabinding, execute as owner as
```

```
begin atomic with
```

```
(transaction isolation level = snapshot, language = N'English')
```

```
declare @OrdDate datetime = getdate();
```

```
insert into dbo.Ord (OrdNo, CustCode, OrdDate) values (@OrdNo, @CustCode, @OrdDate); end
```

```
go
```

* Natively compiled stored procedures are Transact-SQL stored procedures compiled to native code that access memory-optimized tables. Natively compiled stored procedures allow for efficient execution of the queries and business logic in the stored procedure.

* READ COMMITTED versus REPEATABLE READ

Read committed is an isolation level that guarantees that any data read was committed at the moment it is read. It simply restricts the reader from seeing any intermediate, uncommitted, 'dirty' read. It makes no promise whatsoever that if the transaction re-issues the read, will find the same data, data is free to change after it was read.

Repeatable read is a higher isolation level, that in addition to the guarantees of the read committed level, it also guarantees that any data read cannot change, if the transaction reads the same data again, it will find the previously read data in place, unchanged, and available to read.

* Both RAISERROR and THROW statements are used to raise an error in Sql Server.

The journey of RAISERROR started from Sql Server 7.0, whereas the journey of THROW statement has just begun with Sql Server 2012. Obviously, Microsoft is suggesting us to start using THROW statement instead of RAISERROR. THROW statement seems to be simple and easy to use than RAISERROR.

* Explicit transactions. The user starts the transaction through an explicit BEGIN TRAN or BEGIN ATOMIC. The transaction is completed following the corresponding COMMIT and ROLLBACK or END (in the case of an atomic block).

NEW QUESTION 83

- (Exam Topic 7)

You need to implement security for the restore and audit process. What should you do?

- A. Grant the COFFECORP\Auditors group ALTER ANY CONNECTION and SELECT ALL USER SECURABLES permission
- B. Grant the COFFECORP\StoreAgent group ALTER ANY CONNECTION and IMPERSONATE ANY LOGIN permissions.
- C. Grant the COFFECORP\Auditors group CONNECT ANY DATABASE and IMPERSONATE ANY LOGIN permission
- D. Grant the COFFECORP\StoreAgent group CONNECT ANY DATABASE and SELECT ALL USER SECURABLES permissions.
- E. Grant the COFFECORP\Auditors group ALTER ANY CONNECTION and IMPERSONATE ANY LOGIN permission
- F. Grant the COFFECORP\StoreAgent group ALTER ANY CONNECTION and SELECT ALL USER SECURABLES permissions.
- G. Grant the COFFECORP\Auditors group CONNECT ANY DATABASE and SELECT ALL USER SECURABLES permission
- H. Grant the COFFECORP\StoreAgent group CONNECT ANY DATABASE and IMPERSONATE ANY LOGIN permissions.

Answer: A

NEW QUESTION 84

- (Exam Topic 8)

Your company has a main office in London and a branch office in New York.

Your network contains a server named Server5 that has SQL Server 2012 installed. Server5 contains a database named ContentDB and a table named ContentTable.

You add an additional server named Server9 that runs SQL Server 2012.

You need to create a distributed partitioned view. The solution must minimize the amount of network traffic. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Create the view on Server5.
- B. Add Server9 as a linked server.
- C. Create the view on Server9.
- D. Add the Customers table to Server9.
- E. Add Server9 as a Distributor.
- F. Remove the Customers table from Server5.

Answer: ABCD

NEW QUESTION 86

- (Exam Topic 8)

Your network contains a server named Server1 that runs SQL Server 2012. Server1 contains an instance named Instance1. Instance1 contains a database named ContentDatabase.

ContentDatabase uses transaction log backups.

The recovery model of ContentDatabase is set to FULL.

You need to shrink the ContentDatabase_Log log file to 10 MB. The solution must ensure that you can continue to back up the transaction log.

Which three code segments should you execute?

To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

Answer Area

```
DBCC SHRINKFILE (ContentDatabase_Log, 10);  
GO
```

```
ALTER DATABASE ContentDatabase  
SET RECOVERY SIMPLE;  
GO
```

```
ALTER DATABASE ContentDatabase  
SET RECOVERY FULL;  
GO
```

```
ALTER DATABASE ContentDatabase  
SET PAGE_VERIFY CHECKSUM;  
GO
```

```
BACKUP LOG ContentDatabase  
WITH TRUNCATE_ONLY
```

```
DBCC SHRINKFILE (ContentDatabase_Log, 7168);  
GO
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

```
ALTER DATABASE ContentDatabase  
SET RECOVERY SIMPLE;  
GO
```

Box 2:

```
DBCC SHRINKFILE (ContentDatabase_Log, 10);  
GO
```

Box 3:

```
ALTER DATABASE ContentDatabase  
SET RECOVERY FULL;  
GO
```

Note:

* Shrinking a log file to a specified target size

The following example shrinks the log file in the AdventureWorks database to 1 MB. To allow the DBCC SHRINKFILE command to shrink the file, the file is first truncated by setting the database recovery model to SIMPLE.

Transact-SQL

```
USE AdventureWorks2012; GO
```

```
-- Truncate the log by changing the database recovery model to SIMPLE. ALTER DATABASE AdventureWorks2012  
SET RECOVERY SIMPLE; GO
```

```
-- Shrink the truncated log file to 1 MB.
```

```
DBCC SHRINKFILE (AdventureWorks2012_Log, 1); GO
```

```
-- Reset the database recovery model. ALTER DATABASE AdventureWorks2012 SET RECOVERY FULL;  
GO
```

* If the log file does not shrink (after dbcc shrinkfile)

Typically it is the log file that appears not to shrink. This is usually the result of a log file that has not been truncated. You can truncate the log by setting the database recovery model to SIMPLE, or by backing up the log and then running the DBCC SHRINKFILE operation again.

* DBCC SHRINKFILE shrinks the size of the specified data or log file for the current database, or empties a file by moving the data from the specified file to other files in the same filegroup, allowing the file to be removed from the database.

Arguments include: target_size

Is the size for the file in megabytes, expressed as an integer.

NEW QUESTION 90

- (Exam Topic 8)

You are creating a table named Orders.

You need to ensure that every time a new row is added to the Orders table, a table that is used for auditing is updated.

What should you use?

More than one answer choice may achieve the goal. Select the BEST answer.

- A. A Data Definition Language (DDL) trigger
- B. A DEFAULT constraint
- C. A CHECK constraint
- D. A FOREIGN KEY constraint
- E. A data manipulation language (DML) trigger

Answer: E

Explanation:

<http://www.techrepublic.com/blog/programming-and-development/comparing-sql-serverconstraints-and-dmltrig> <http://msdn.microsoft.com/en-us/library/ms178110.aspx>

NEW QUESTION 92

- (Exam Topic 8)

You plan to modify a stored procedure to use temporary data. The stored procedure must meet the following requirements:

▶ Favor physical memory when physical memory is available.

▶ Be able to roll back changes to the temporary data.

You need to recommend which object to add to the stored procedure. Which T-SQL command should you recommend?

- A. CREATE TABLE ##Table...
- B. CREATE TABLE Table...
- C. CREATE VIEW Table...
- D. CREATE PARTITION SCHEME Table...
- E. DECLARE TABLE @ Table...

Answer: A

Explanation:

Temporary Tables

You can create local and global temporary tables. Local temporary tables are visible only in the current session, and global temporary tables are visible to all sessions. Temporary tables cannot be partitioned.

Prefix local temporary table names with single number sign (#table_name), and prefix global temporary table names with a double number sign (##table_name)

NEW QUESTION 94

- (Exam Topic 8)

You administer a SQL Server 2014 instance.

The server is capable of 10000 IO/second (IOPS). During the time period when the second process executes, the disk IO can reach 7000 IOPS, and CPU use can average 30% over the eight processors.

The first process summarizes the day's activity executed by a login of [SummaryReportLogin]. The second process submits transactions executed by a login of [ETLLogin].

A Resource Governor classifier function has been created to return WG_Low for connections from the [ETLLogin] and [SummaryReportLogin].

You need to set up the Resource Group and Workgroup Pools on the instance. You have the following requirements:

▶ Both processes must never use more than 50 percent of the CPU at any one time.

▶ The number of active queries that these processes can execute simultaneously should be limited to a maximum of 10.

▶ The SummaryReportLogin process must always achieve the minimum IOPS required to be minimally affected during executing the ETLLogin processes.

Develop the solution by selecting and arranging the required code blocks in the correct order. You may not need all of the code blocks.

Code Blocks

```
MAX_IOPS_PER_VOLUME=3000
)
```

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
    MAX_DOP = 4
)
USING RP_Low
```

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
    GROUP_MAX_REQUESTS=10
)
USING RP_Low
```

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
    REQUEST_MAX_CPU_TIME_SEC = 100,
    MAX_DOP = 4
)
USING RP_Low
```

```
CREATE RESOURCE POOL RP_Low
WITH
(
    CAP_CPU_PERCENT=50,
    MAX_CPU_PERCENT=30,
```

```
CREATE RESOURCE POOL RP_Low
WITH
(
    AFFINITY_SCHEDULER = (0 to 50),
    MAX_CPU_PERCENT=30,
```

```
CREATE RESOURCE POOL RP_Low
WITH
(
    MAX_CPU_PERCENT=50,

    MAX_IOPS_PER_VOLUME=30
)
```

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

```
CREATE RESOURCE POOL RP_Low
WITH
(
    CAP_CPU_PERCENT=50,
    MAX_CPU_PERCENT=30,
```

Box 2:

```
MAX_IOPS_PER_VOLUME=3000
)
```

Box 3:

```
CREATE WORKLOAD GROUP WG_Low
WITH
(
GROUP_MAX_REQUESTS=10
)
USING RP_Low
```

Note:

CREATE WORKLOAD RESOURCE POOL

* Resource pools. A resource pool, represents the physical resources of the server. You can think of a pool as a virtual SQL Server instance inside of a SQL Server instance.

* Workload groups. A workload group serves as a container for session requests that have similar classification criteria. A workload allows for aggregate monitoring of the sessions, and defines policies for the sessions. Each workload group is in a resource pool.

* CAP_CPU_PERCENT =value

Specifies a hard cap on the CPU bandwidth that all requests in the resource pool will receive. Limits the maximum CPU bandwidth level to be the same as the specified value. value is an integer with a default setting of 100. The allowed range for value is from 1 through 100.

* MIN_IOPS_PER_VOLUME =value

Specifies the minimum I/O operations per second (IOPS) per disk volume to reserve for the resource pool.

* GROUP_MAX_REQUESTS =value

Specifies the maximum number of simultaneous requests that are allowed to execute in the workload group. value must be a 0 or a positive integer.

NEW QUESTION 99

- (Exam Topic 8)

You plan to design an application that temporarily stores data in a SQL Azure database.

You need to identify which types of database objects can be used to store data for the application. The solution must ensure that the application can make changes to the schema of a temporary object during a session.

Which type of objects should you identify?

- A. Common table expressions (CTEs)
- B. Temporary stored procedures
- C. Temporary tables
- D. Table variables

Answer: C

Explanation:

<http://msdn.microsoft.com/en-us/library/ms175972.aspx> <http://msdn.microsoft.com/en-us/library/ms189084.aspx> <http://msdn.microsoft.com/en-us/library/ms175010.aspx> <http://msdn.microsoft.com/en-us/library/bb510489.aspx> <http://msdn.microsoft.com/en-us/library/ms187926.aspx>
<http://zacksfiasco.com/post/2010/01/21/SQL-Server-Temporary-Stored-Procedures.aspx>

NEW QUESTION 100

- (Exam Topic 8)

You have a table named Table1 that contains one million rows. Table1 contains a column named Column1 that stores sensitive information. Column1 uses the nvarchar(16) data type. You have a certificate named Cert1.

You must add a column named Column2 that contains an encrypted version of the data from Column1. You must use two-way encryption. You plan to remove Column1 after you create Column2.

Which five Transact-SQL statements should you run in sequence before you remove Column1? To answer, move the appropriate Transact-SQL statements from the list of Transact-SQL statements to the answer area and arrange them in the correct order.

NOTE: More than one order of answer choices is correct. You will receive credit for any correct orders you select.

Code segments	Answer Area
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = SHA1 ENCRYPTION BY CERTIFICATE Cert1;	
CLOSE SYMMETRIC KEY;	
ALTER TABLE TABLE1 ADD Column2 varbinary(256);	
UPDATE table1 SET Column2 = EncryptByKey(Key_GUID (Key1),Column1);	
ALTER TABLE Table1 ADD Column2 nvarchar(256);	
CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = AES_256 ENCRYPTION BY CERTIFICATE Cert1;	
CREATE CREDENTIAL Cred1 WITH IDENTITY = 'User1', SECRET = 'P@ssword';	
OPEN SYMMETRIC KEY Key1 DECRYPTION BY CERTIFICATE Cert1;	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/security/encryption/encrypt-a-column-of-data?view=sq> <https://docs.microsoft.com/en-us/sql/t-sql/statements/close-symmetric-key-transact-sql?view=sql-server-2017>

NEW QUESTION 101

- (Exam Topic 8)

You discover a sudden increase in processor utilization on a server that has SQL Server installed. You need to correlate server performance and database activity for an extended time period. Which two tools should you use? Each correct answer presents part of the solution.

- A. Activity Monitor
- B. Performance Monitor
- C. SQL Server Profiler
- D. sp_who2
- E. SQL Server Extended Events

Answer: BE

Explanation:

The Performance Monitor side, we have a few SQL Server monitoring tools AKA counters that can be used when troubleshooting CPU performance. The following counters are simple and easy to use:

Processor % Processor Time == < 80% Processor % User Time == < 80% Processor % Privileged time) == < 30% References:

<https://www.sqlshack.com/sql-server-monitoring-tool-for-cpu-performance/>

NEW QUESTION 104

- (Exam Topic 8)

You execute the following code:

```
CREATE TABLE dbo.Customers
(
    id int PRIMARY KEY,
    CustomerName char(10)
)
```

You create a nonclustered index named IX_CustomerName on the CustomerName column. You execute the following query:

```
SELECT * FROM dbo.Customers
WHERE LEFT(CustomerName,1) = 'a'
```

You need to reduce the amount of time it takes to execute the query. What should you do?

- A. Partition the table and use the CustomerName column for the partition scheme.
- B. Replace IX_CustomerName with a clustered index.
- C. Replace LEFT(CustomerName ,1) = 'a' with CustomerName LIKE 'a%'.
- D. Replace LEFT(CustomerName ,1) = 'a' with SUBSTRING(CustomerName ,1,1) - 'a'.

Answer: C

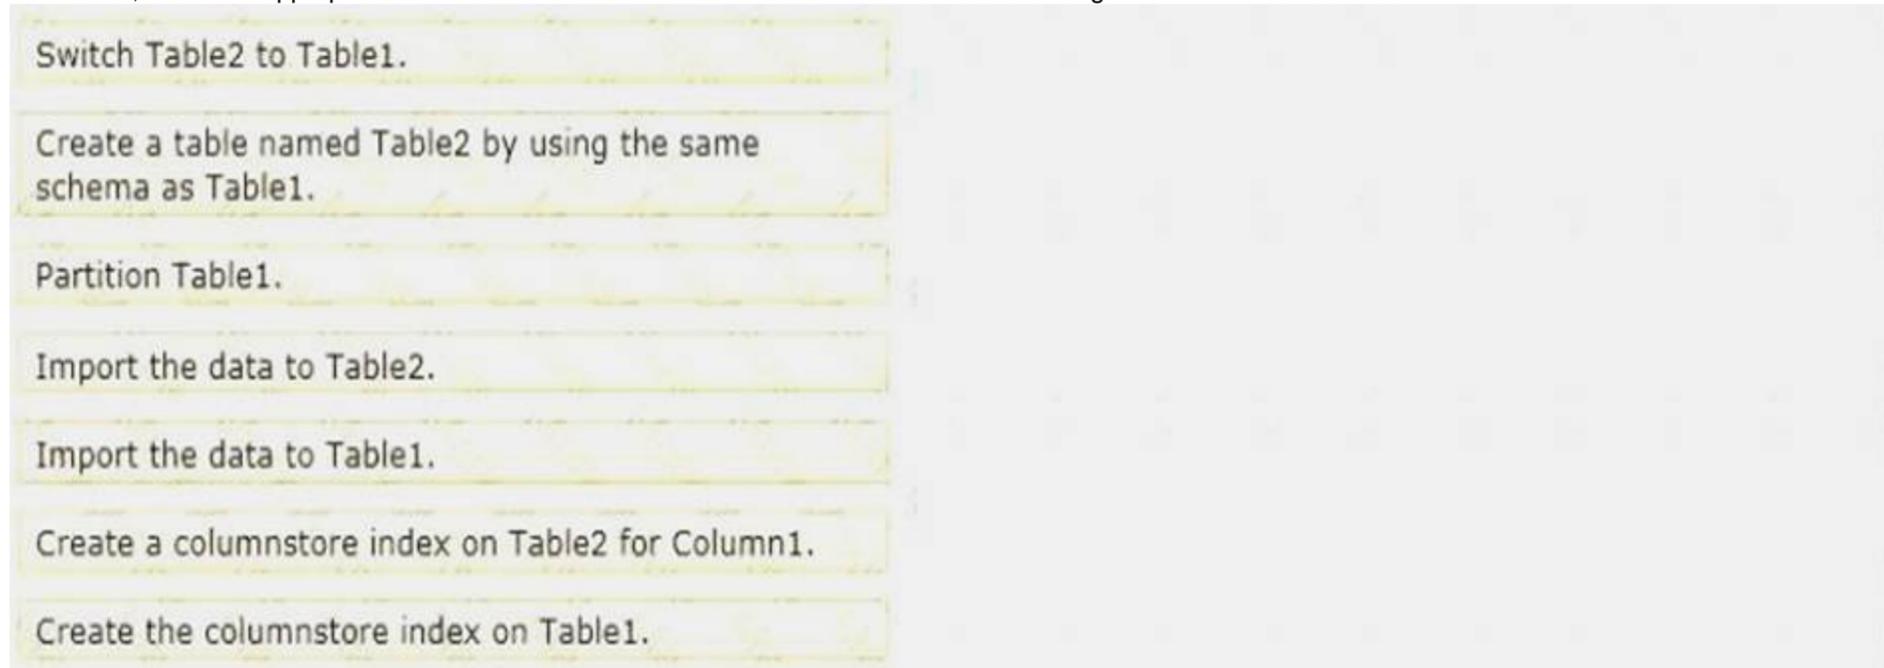
Explanation:

<http://msdn.microsoft.com/en-us/library/ms179859.aspx> <http://msdn.microsoft.com/en-us/library/ms187748.aspx>

NEW QUESTION 108

- (Exam Topic 8)

You have a table named Table1. Table1 has 1 million rows. Table1 has a columnstore index for a column named Column1. You need to import data to Table1. The solution must minimize the amount of time it takes to import the data. What should you do? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



Switch Table2 to Table1.

Create a table named Table2 by using the same schema as Table1.

Partition Table1.

Import the data to Table2.

Import the data to Table1.

Create a columnstore index on Table2 for Column1.

Create the columnstore index on Table1.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Create a table named Table2 by using the same schema as Table1. Note: Table2 is the staging table.

Box 2: Partition Table1

Box 3: Import the data to Table2.

Box 4: Create a columnstore index on Table2 for Column1. Box 5: Switch Table2 to Table1

Note:

* An xVelocity memory optimized columnstore index, groups and stores data for each column and then joins all the columns to complete the whole index. Columnstore indexes can transform the data warehousing experience for users by enabling faster performance for common data warehousing queries such as filtering, aggregating, grouping, and star-join queries.

* Tables that have a columnstore index cannot be updated. There are three ways to work around this problem.

A) To update a table with a columnstore index, drop the columnstore index, perform any required INSERT, DELETE, UPDATE, or MERGE operations, and then rebuild the columnstore index.

B) (applies in this scenario) Partition the table and switch partitions. For a bulk insert, insert data into a staging table, build a columnstore index on the staging table, and then switch the staging table into an empty partition. For other updates, switch a partition out of the main table into a staging table, disable or drop the columnstore index on the staging table, perform the update operations, rebuild or re-create the columnstore index on the staging table, and then switch the staging table back into the main table.

C) Place static data into a main table with a columnstore index, and put new data and recent data likely to change, into a separate table with the same schema that does not have a columnstore index.

Reference: Best Practices: Updating Data in a Columnstore Index

NEW QUESTION 112

- (Exam Topic 8)

You have a SQL Server 2012 database that contains a table named Users. The Users table contains usernames and passwords.

You need to ensure that all new records have a password. Which code segment should you use?

More than one answer choice may achieve the goal. Select the BEST answer.

```
A. ALTER TABLE dbo.Users
    DROP Password;
GO
ALTER TABLE dbo.Users
    ADD Password varchar(30) NOT NULL;
GO

B. ALTER TABLE dbo.Users
    ADD CONSTRAINT CK_Users_Password
    CHECK (Password IS NULL);
GO

C. DROP TABLE dbo.Users;
GO
CREATE TABLE dbo.Users (
    CustID int PRIMARY KEY,
    Name varchar(30),
    Password varchar(30),
    CONSTRAINT CK_Users_Password
    CHECK (Password IS NOT NULL));
GO

D. ALTER TABLE dbo.Users
    ADD CONSTRAINT CK_Users_Password
    CHECK (Password IS NOT NULL);
GO
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

NEW QUESTION 117

- (Exam Topic 8)

You have a database named Database1. You execute the following code:

```
CREATE TABLE dbo.table1
(
    ID int IDENTITY(1,1) NOT NULL PRIMARY KEY,
    FirstName varchar(50) NOT NULL,
    LastName varchar(50) NOT NULL,
    EmailAddress varchar(200) NULL,
    Notes nvarchar(MAX) NULL,
    LastContactDate datetime NULL
)
```

You have the following query. (Line numbers are included for reference only.)

```
01 SELECT FirstName + ' ' + LastName AS Name
02 FROM dbo.table1
03 WHERE Notes LIKE '% call%' AND
04 LastContactDate >= '1/1/2010'
```

Users report that the query takes a long time to complete. You create a full-text index for the Notes column.

You need to recommend changes to the query to reduce the amount of time it takes for the query to complete. Which code segment should you use to replace line 03?

- A. WHERE FREETEXT(notes, '%call%') AND
- B. INNER JOIN FREETEXTTABLE(dbo.table1, notes, 'call') AS t2 ON dbo.table1.ID = t2.key WHERE
- C. WHERE CONTAINS(notes, 'call*') AND
- D. WHERE CONTAINS(notes, '*%call%*>' AND

Answer: A

NEW QUESTION 119

- (Exam Topic 8)

You are a SQL Server 2014 Developer. A database that you work on contains two tables that are defined as follows:

```
CREATE TABLE Product (  
ProductIDint IDENTITY(1,1) PRIMARY KEY,  
ProductNamevarchar(30) NOT NULL,  
LastUpdatedDatesmalldatetime,  
LastUpdatedBynvarchar(128))  
  
CREATE TABLE ProductAudit (  
ProductAuditIDint IDENTITY(1,1) PRIMARY KEY,  
OldProductID int NOT NULL,  
OldProductName varchar(30) NOT NULL,  
UpdatedDatesmalldatetime,  
UpdatedBynvarchar(128))
```

Product is an important table that has sensitive audit requirements. You need to create a trigger that supports the following requirements:

1. Every row that is inserted or updated in Product will reflect its actual LastUpdatedDate and LastUpdatedBy values in the Product table.
 2. Any row that is updated or deleted must write a new record reflecting the OLD values into the ProductAudit table.
 3. Any error that occurs during the course of the trigger's execution must prevent the changes from happening. Develop the solution by selecting and arranging the required code blocks in the correct order.
- You may not need all of the code blocks.

Code Blocks	Answer Area
<pre>DECLARE @OldProductId int, @OldProductName varchar (30) SELECT @OldProductId = ProductId, @OldProductName = ProductName FROM deleted INSERT ProductAudit (OldProductID, OldProductName, UpdatedDate, Update dBy) SELECT @OldProductID, @OldProductName, SUSER_NAME (), GETDATE() UPDATE Product SET LastUpdatedBy = SUSER_NAME(), LastUpdatedDate = GETDATE() FROM Product AS p INNER JOIN inserted AS i ON p.ProductID = i.ProductID UPDATE Product SET LastUpdatedBy = SUSER_NAME(), LastUpdatedDate = GETDATE() FROM Product AS p INNER JOIN inserted AS i ON p.ProductID = i.ProductID INSERT ProductAudit (OldProductID, OldProductName, UpdatedDate, Update dBy) SELECT d.ProductID, d.ProductName, SUSER_NAME (), GETDATE() FROM deleted AS d END COMMIT TRANSACTION IF @@ERROR <> 0 ROLLBACK CREATE TRIGGER ProductAuditTrigger ON Product FOR INSERT, UPDATE, DELETE AS BEGIN</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

* Executing a ROLLBACK TRANSACTION or COMMIT TRANSACTION Transact-SQL statement inside a stored procedure or trigger is possible, but doing so may cause errors.

NEW QUESTION 121

- (Exam Topic 8)

You have a table named Table1 that contains 1 million rows. Table1 contains a column named Column1 that stores sensitive information. Column1 uses the nvarchar(16) data type.

You have a certificate named Cert1.

You need to replace Column1 with a new encrypted column that uses two-way encryption. Which code segment should you execute before you remove Column1? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

Code segments	Answer Area
<pre>ALTER TABLE Table1 ADD Column2 varbinary(256);</pre>	
<pre>CLOSE SYMMETRIC KEY;</pre>	
<pre>CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = SHA1 ENCRYPTION BY CERTIFICATE Cert1;</pre>	
<pre>CREATE CREDENTIAL Cred1 WITH IDENTITY = 'User1', SECRET = 'P@ssw0rd';</pre>	
<pre>ALTER TABLE Table1 ADD Column2 nvarchar(256);</pre>	
<pre>OPEN SYMMETRIC KEY Key1 DECRYPTION BY CERTIFICATE Cert1;</pre>	
<pre>UPDATE table1 SET Column2 = EncryptByKey(Key_GUID ('Key1'),Column1);</pre>	
<pre>CREATE SYMMETRIC KEY Key1 WITH ALGORITHM = AES_256 ENCRYPTION BY CERTIFICATE Cert1;</pre>	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:
 * Use AES_256 for two-way encryption.
 * Use varbinary to store key.
 * CLOSE SYMMETRIC KEY (Transact-SQL)
 Closes a symmetric key, or closes all symmetric keys open in the current session.
 * Example:
 CREATE SYMMETRIC KEY CreditCards_Key11 WITH ALGORITHM = AES_256
 ENCRYPTION BY CERTIFICATE Sales09; GO
 -- Create a column in which to store the encrypted data. ALTER TABLE Sales.CreditCard
 ADD CardNumber_Encrypted varbinary(128); GO
 -- Open the symmetric key with which to encrypt the data. OPEN SYMMETRIC KEY CreditCards_Key11 DECRYPTION BY CERTIFICATE Sales09;
 -- Encrypt the value in column CardNumber using the
 -- symmetric key CreditCards_Key11.
 -- Save the result in column CardNumber_Encrypted. UPDATE Sales.CreditCard
 SET CardNumber_Encrypted = EncryptByKey(Key_GUID('CreditCards_Key11')
 , CardNumber, 1, HashBytes('SHA1', CONVERT(varbinary
 , CreditCardID))); GO

NEW QUESTION 125

- (Exam Topic 8)

You need to identify which nonclustered indexes are unused by queries.

How should you complete the statement? To answer, drag the appropriate values to the correct locations. Each value may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Values

- sys.dm_db_index_operational_stats
- sys.dm_db_index_physical_stats
- sys.dm_db_index_usage_stats
- sys.indexes
- sys.tables



Answer Area

```
SELECT database_id, a.object_id, a.index_id, b.name, a.user_seeks, a.user_scans, a.user_updates
FROM [Value] a
join [Value] b on a.object_id = b.object_id
and a.index_id = b.index_id
where a.index_id >= 2
and a.database_id = db_id()
and (a.user_seeks = 0
or a.user_scans = 0)
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.dm_db_index_usage_stats

sys.dm_db_index_usage_stats shows you how many times the index was used for user queries. It returns counts of different types of index operations and the time each type of operation was last performed in SQL Server.

Box 2: sys.indexes

sys.indexes contains a row per index or heap of a tabular object, such as a table, view, or table-valued function.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-index->

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-indexes-transact-sql>

NEW QUESTION 129

- (Exam Topic 8) You have a SQL Azure database. You execute the following script:

```
CREATE TABLE dbo.Table1
(
    Column1 int PRIMARY KEY,
    Column2 varchar(50) SPARSE NULL
)
```

You add 1 million rows to Table1. Approximately 85 percent of all the rows have a null value for Column2. You plan to deploy an application that will search Column2.

You need to create an index on Table1 to support the planned deployment. The solution must minimize the storage requirements.

Which code segment should you execute?

- A. CREATE INDEX IX_Table1 ON Table1 (Column2) WITH FILLFACTOR=0
- B. CREATE INDEX IX_Table1 ON Table1 (Column1) INCLUDE (Column2)
- C. CREATE INDEX IX_Table1 ON Table1 (Column2) WHERE Column2 IS NULL
- D. CREATE INDEX IX_Table1 ON Table1 (Column2) WHERE Column2 IS NOT NULL

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms188783.aspx> <http://msdn.microsoft.com/en-us/library/cc280372.aspx>

NEW QUESTION 133

- (Exam Topic 8)

You have a SQL Server 2012 database named Database1. Database1 has a data file named Database1_data.mdf and a transaction log named Database1log.ldf. Database1_data.mdf is 1.5 GB. Database1log.ldf is 1.5 terabytes.

A full backup of Database1 is performed every day.

You need to reduce the size of the log file. The solution must ensure that you can perform transaction log backups in the future.

Which code segment should you execute?

To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

The screenshot shows a list of five code segments in a drag-and-drop interface:

- DBCC SHRINKFILE(database1_log,1)
- ALTER DATABASE database1 SET RECOVERY FULL
- ALTER DATABASE database1 SET RECOVERY SIMPLE
- BACKUP LOG database1 WITH TRUNCATE_ONLY
- DBCC SHRINKFILE(database1_data,1)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The screenshot shows the correct sequence of code segments in the answer area, indicated by a red dashed border:

- ALTER DATABASE database1 SET RECOVERY SIMPLE
- DBCC SHRINKFILE(database1_log,1)
- ALTER DATABASE database1 SET RECOVERY FULL

The other code segments from the list are shown in a greyed-out state, indicating they are not part of the correct solution.

NEW QUESTION 136

- (Exam Topic 8)

You are planning two stored procedures named SProc1 and SProc2. You identify the following requirements:

- ▶ SProc1 must return a table.
- ▶ SProc2 must return a status code.

You need to identify which options must be implemented to meet each stored procedure requirement. Which options should you identify?

To answer, drag the appropriate option to the correct requirement in the answer area. (Answer choices may be used once, more than once, or not at all.)

The screenshot shows an interface with two columns: 'Options' and 'Answer Area'.

Options:

- a raise error
- a return value
- a SELECT statement
- a table-valued parameter (TVP)

Answer Area:

- SProc1: Option
- SProc2: Option

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

NEW QUESTION 141

- (Exam Topic 8)

You are creating a table named Orders.

You need to ensure that every time a new row is added to the Orders table, a user-defined function is called to validate the row before the row is added to the table.

What should you use?

More than one answer choice may achieve the goal. Select the BEST answer.

- A. A data manipulation language (DML) trigger
- B. A DEFAULT constraint
- C. A Data Definition Language (DDL) trigger
- D. A CHECK constraint
- E. A FOREIGN KEY constraint

Answer: D

Explanation:

<http://www.techrepublic.com/blog/programming-and-development/comparing-sql-serverconstraints-and-dmltrig> <http://msdn.microsoft.com/en-us/library/ms178110.aspx>

NEW QUESTION 142

- (Exam Topic 8)

You have a database named database1. Each table in database1 has one index per column. Users often report that creating items takes a long time.

You need to perform the following maintenance tasks: What should you use?

To answer, drag the appropriate function to the correct management task in the answer area. (Answer choices may be used once, more than once, or not at all.)

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: sys.dm_db_index_usage_stats

sys.dm_db_index_usage_stats shows you how many times the index was used for user queries. It returns counts of different types of index operations and the time each type of operation was last performed in SQL Server.

Box 2: sys.dm_db_missing_index_details

sys.dm_db_missing_index_details returns detailed information about a missing index; for example, it returns the name and identifier of the table where the index is missing, and the columns and column types that should make up the missing index.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-db-index->

<https://docs.microsoft.com/en-us/sql/relational-databases/system-catalog-views/sys-indexes-transact-sql> [https://technet.microsoft.com/en-us/library/ms345524\(v=sql.105\).aspx](https://technet.microsoft.com/en-us/library/ms345524(v=sql.105).aspx)

NEW QUESTION 146

- (Exam Topic 8)

You use SQL Server 2014 to maintain the data used by applications at your company. You need to run two separate SQL statements.

You must guarantee that the following three things happen:

1. Either BOTH statements succeed or BOTH statements fail as a batch.
2. If an error occurs on the first statement, SQL should not attempt to run the second statement.
3. Error information should be returned to the client. What should you do?

<p><input type="radio"/> A. SET XACT_ABORT ON BEGIN TRY BEGIN TRANSACTION ...Statement 1 ...Statement 2 COMMIT TRANSACTION END TRY BEGIN CATCH ROLLBACK TRANSACTION END CATCH</p>	<p><input type="radio"/> C. SET XACT_ABORT ON BEGIN TRANSACTION ...Statement 1 ...Statement 2 If @@ERROR <> 0 ROLLBACK ELSE COMMIT TRANSACTION</p>
<p><input type="radio"/> B. SET XACT_ABORT OFF BEGIN TRY ...Statement 1 END TRY BEGIN TRY ...Statement 2 END TRY BEGIN CATCH THROW END CATCH</p>	<p><input type="radio"/> D. SET XACT_ABORT ON BEGIN TRY ...Statement 1 if @@ERROR <> 0 GOTO CATCH ...Statement 2 if @@ERROR <> 0 GOTO CATCH END TRY BEGIN CATCH THROW END CATCH</p>

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

Explanation:

* SET XACT_ABORT

When SET XACT_ABORT is ON, if a Transact-SQL statement raises a run-time error, the entire transaction is terminated and rolled back.

When SET XACT_ABORT is OFF, in some cases only the Transact-SQL statement that raised the error is rolled back and the transaction continues processing.

NEW QUESTION 147

- (Exam Topic 8)

You have a SQL Server database.

The recovery model is set to full. The transaction log is backed up every night. You discover that the transaction log never decreases in size.

You execute the DBCC SHRINKFILE statement for the transaction log and you discover that the transaction log is unaffected.

You need to ensure that you can reduce the size of the transaction log. What should you do first?

- A. Truncate the transaction log
- B. Kill long-running transactions
- C. Change the recovery model to bulk-logged
- D. Perform a full backup

Answer: A

Explanation:

The transaction log must be truncated before running the DBCC SHRINKFILE operation.

NEW QUESTION 149

- (Exam Topic 8)

You have a table named Customers that has a clustered index defined on the ID column. You write a script to create a stored procedure.

You need to complete the script for the stored procedure. The solution must minimize the number of locks and deadlocks.

What should you do?

To answer, drag the appropriate option to the correct location in the answer area. (Answer choices may be used once, more than once, or not at all.)

READ COMMITTED	CREATE PROCEDURE Proc1 (@ParamID int)
SERIALIZABLE	AS
WITH(UPDLOCK)	SET TRANSACTION ISOLATION LEVEL
WITH(XLOCK)	BEGIN TRANSACTION
	DECLARE @var as NCHAR(10)
	Select @var = Name
	FROM dbo.Customers
	WHERE ID = @ParamID
	...
	UPDATE dbo.Customers
	SET Name = @var
	WHERE ID = @ParamID
	COMMIT TRANSACTION;
	GO

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Note:

* Optimized bulk load operations on heaps block queries that are running under the following isolation levels: SNAPSHOT

READ UNCOMMITTED

READ COMMITTED using row versioning

* READ COMMITTED

Specifies that statements cannot read data that has been modified but not committed by other transactions. This prevents dirty reads. Data can be changed by other transactions between individual statements within the current transaction, resulting in nonrepeatable reads or phantom data. This option is the SQL Server default.

* SERIALIZABLE (more locks) Specifies the following:

Statements cannot read data that has been modified but not yet committed by other transactions.

No other transactions can modify data that has been read by the current transaction until the current transaction completes.

Other transactions cannot insert new rows with key values that would fall in the range of keys read by any statements in the current transaction until the current transaction completes.

* UPDLOCK

Specifies that update locks are to be taken and held until the transaction completes. UPDLOCK takes update locks for read operations only at the row-level or page-level. If UPDLOCK is combined with TABLOCK, or a table-level lock is taken for some other reason, an exclusive (X) lock will be taken instead.

When UPDLOCK is specified, the READCOMMITTED and READCOMMITTEDLOCK isolation level hints are ignored. For example, if the isolation level of the session is set to SERIALIZABLE and a query specifies (UPDLOCK, READCOMMITTED), the READCOMMITTED hint is ignored and the transaction is run using the SERIALIZABLE isolation level.

* XLOCK

Specifies that exclusive locks are to be taken and held until the transaction completes. If specified with ROWLOCK, PAGLOCK, or TABLOCK, the exclusive locks apply to the appropriate level of granularity.

Reference: Table Hints (Transact-SQL)

NEW QUESTION 151

- (Exam Topic 8)

You plan to create a new table that will contain a column named Salary. Salary will contain highly sensitive data.

Salary must meet the following requirements:

▶ Contain numeric data.

▶ Contain only encrypted data that remains encrypted in memory.

You need to identify which encryption type and data type must be used for Salary. Which encryption type and data type should you identify?

To answer, drag the appropriate encryption type and data type to the correct identifier in the answer area.

Encryption Types

Transparent data encryption (TDE)

Encrypting File System (EFS)

Cell-level encryption

BitLocker Drive Encryption (BitLocker)

Data Types

decimal

varchar

varbinary

money

Answer Area

Encryption Type

Data Type

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Encryption Types

Transparent data encryption (TDE)

Encrypting File System (EFS)

Cell-level encryption

BitLocker Drive Encryption (BitLocker)

Data Types

decimal

varchar

varbinary

money

Answer Area

Encryption Type: Cell-level encryption

Data Type: varbinary

NEW QUESTION 154

- (Exam Topic 8)

You administer a SQL Server 2014 instance.

You have been assigned to determine the cause of frequent long-running transactions that have been tracked to the dbo.Account table, where there are many

cases of blocking and deadlocks. The dbo.Account table contains more than one million rows.

Users and processes frequently search for and update data by using the AccountId column, and less frequently the AccountNumber and GovernmentId columns, all of which contain only unique values. Users frequently get lists of AccountNumber values by searching on Last Name and then First Name.

You need to modify the structure of the dbo.Account table to alleviate the issues.

How should you complete the table definition to reduce contention on the table structure? To answer, drag the appropriate code snippets to the correct locations in the CREATE TABLE statement. Each code snippet may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Note:

Users and processes frequently search for and update data by using the AccountId column (Primary Key Clustered) , and less frequently the AccountNumber (Unique Clustered) and GovernmentId(Unique Clustered) columns, all of which contain only unique values. Users frequently get lists of AccountNumber values by searching on Last Name and then First Name (LastName, Firstname) INCLUDE (AccountNumber).

NEW QUESTION 157

- (Exam Topic 8)

You have a SQL Server 2012 database named database1.

Users report that queries that usually take less than one second to execute, take more than 30 seconds to execute.

You need to view the server resource consumption when the queries are executed. What should you do?

To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Start a SQL Server Profiler trace.

Box 2: Start a data collection by using Performance monitor. Box 3: Save the SQL Profiler trace.

Box 4: Save the Performance Monitor data.

Box 5: Import the performance data into SQL Server Profiler. Note:

* (step1, step 2) Both the Profiler trace and the Performance Monitor logs should be started and stopped at about the same time.

* (step 3, step 4) Once you have completed capturing the data for both tools, you are ready to perform the correlation analysis.

* (step 5) How to Correlate SQL Server Profiler Data with Performance Monitor Data

Correlating Performance Monitor and Profiler data is a straightforward process that simply involves importing both sets of data into Profiler. Start Profiler and load the trace file you want to correlate.

From the main menu of Profiler, select File | Import Performance Data,

* With SQL Server Profiler, we have the tools to identify the causes of such spikes. We can import Performance Monitor log data and compare it directly with Profiler activity. If we see a spike in CPU utilization, we can identify which statement or statements were running at the same time, and diagnose potential problems.

NEW QUESTION 160

- (Exam Topic 8)

You plan to create a database that has multiple tables. The tables will contain product information. Each product has a stock-keeping unit (SKU).

You need to recommend a solution to ensure that each SKU starts with the letters "ADV" and is followed by 10 digits.

The solution must minimize the amount of development effort required. What should you include in the recommendation?

- A. A FOREIGN KEY constraint
- B. A trigger
- C. A user-defined data type
- D. A CHECK constraint

Answer: C

NEW QUESTION 164

- (Exam Topic 8)

You need to implement a solution that meets the data recovery requirements. You update each stored procedure to accept a parameter named @transactionID.

What should you add next to the beginning of each stored procedure?

- A. SAVE TRANSACTION WITH MARK @transactionID
- B. ROLLBACK DISTRIBUTED TRANSACTION @transactionID
- C. BEGIN TRANSACTION WITH MARK @transactionID
- D. COMMIT TRANSACTION @transactionID

Answer: C

NEW QUESTION 166

- (Exam Topic 8)

You have a database that uses the following management views:

- Sys.dm_os_volume_stats
- Sys.dm_db_partition_stats
- Sys.dm_db_file_space_usage
- Sys.fulltext_indexes

You plan to migrate the database to Microsoft SQL Azure. You need to identify which view can be used in SQL Azure. Which view should you identify?

- A. sys.fulltext_indexes
- B. sys.dm_db_file_space_usage
- C. sys.dm_os_volume_stats
- D. sys.dm_db_partition_stats

Answer: D

NEW QUESTION 167

- (Exam Topic 8)

You plan to create a new column in a disk-based table. The column must meet the following requirements: Be able to store images that are larger than 1 MB each.

Be able to access the images from Microsoft .NET Framework applications. You need to recommend which data type must be used in the column. Which data type should you recommend?

More than one answer choice may achieve the goal. Select the BEST answer.

- A. FileStream
- B. nvarchar
- C. image
- D. varbinary

Answer: A

NEW QUESTION 172

- (Exam Topic 8)

You plan to create a custom aggregation function named Median.

You plan to deploy Median to a SQL Server 2014 server named Server1.

You need to ensure that Median can access a web service named WebApp1. The solution must minimize the number of changes made to the database.

You create a Microsoft .NET Framework class that contains the definition of Median. You upload a certificate to Server1.

What three tasks should you perform next?

To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Execute the CREATE AGGREGATE statement.	
Modify the TRUSTWORTHY property of the database.	
Execute the CREATE ASSEMBLY statement.	
Execute the CREATE FUNCTION statement.	
Use the certificate to add a digital signature to the assembly.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Execute the CREATE ASSEMBLY statement.

Box 2: Use the certificate to add a digital signature to the assembly. Box 3: Execute the CREATE AGGREGATE statement.

Note:

* CREATE AGGREGATE

Creates a user-defined aggregate function whose implementation is defined in a class of an assembly in the .NET Framework. For the Database Engine to bind the aggregate function to its implementation, the .NET Framework assembly that contains the implementation must first be uploaded into an instance of SQL Server by using a CREATE ASSEMBLY statement.

NEW QUESTION 175

- (Exam Topic 8)

You have an application that queries a database.

Users report that the application runs more slowly than expected.

You need to identify which queries take the most time to execute. The solution must minimize the use of processor resources on the server.

What should you use?

- A. a SQL Server Profiler trace
- B. a Central Management Server
- C. Query Store
- D. Distributed Replay

Answer: C

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/performance/monitoring-performance-by-using-theque>

NEW QUESTION 179

- (Exam Topic 8)

You have a SQL Azure database.

You need to identify which keyword must be used to create a view that will be indexed. Which keyword should you identify?

- A. SCHEMABINDING
- B. VIEW_METADATA
- C. DISTINCT
- D. DEFAULT

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms187956.aspx> <http://msdn.microsoft.com/en-us/library/ms191432.aspx>

NEW QUESTION 182

- (Exam Topic 8)

Your network contains a SQL Server 2012 instance named SQL1. SQL1 contains a database named DB1. DB1 contains three tables.

The tables are configured as shown in the following table.

Table name	Configuration
Table1	<ul style="list-style-type: none"> Table data will not be updated. The table will contain historical calculations. The table will contain 10 million records.
Table2	<ul style="list-style-type: none"> 20% of the table data will be updated weekly. The table will contain 25 million records.
Table3	<ul style="list-style-type: none"> 40% of the table data will be updated weekly. The table will contain 1 million records.

You plan to create indexes for the tables.

You need to identify which type of index must be created for each table. The solution must minimize the amount of time required to return information from the tables.

Which type of index should you create for each table? To answer, drag the appropriate index type to the correct table in the answer area.

Index Types	Answer Area
Columnstore Index	Table1
Nonclustered Index	Table2
	Table3

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Index Types	Answer Area
Columnstore Index	Table1 Columnstore Index
Nonclustered Index	Table2 Nonclustered Index
	Table3 Nonclustered Index

NEW QUESTION 186

- (Exam Topic 8)

You plan to modify a procedure that contains hundreds of lines of code. The modification must support the following guidelines:

- ▶ Use only tables that are not persistent in the database.
- ▶ Minimize the amount of time required to execute and recompile procedures.

You need to identify which type of table must be used to support the planned modification.

Which type of table should you identify?

- A. A system table
- B. A partitioned table
- C. A table variable
- D. A temporary table

Answer: C

NEW QUESTION 191

- (Exam Topic 8)

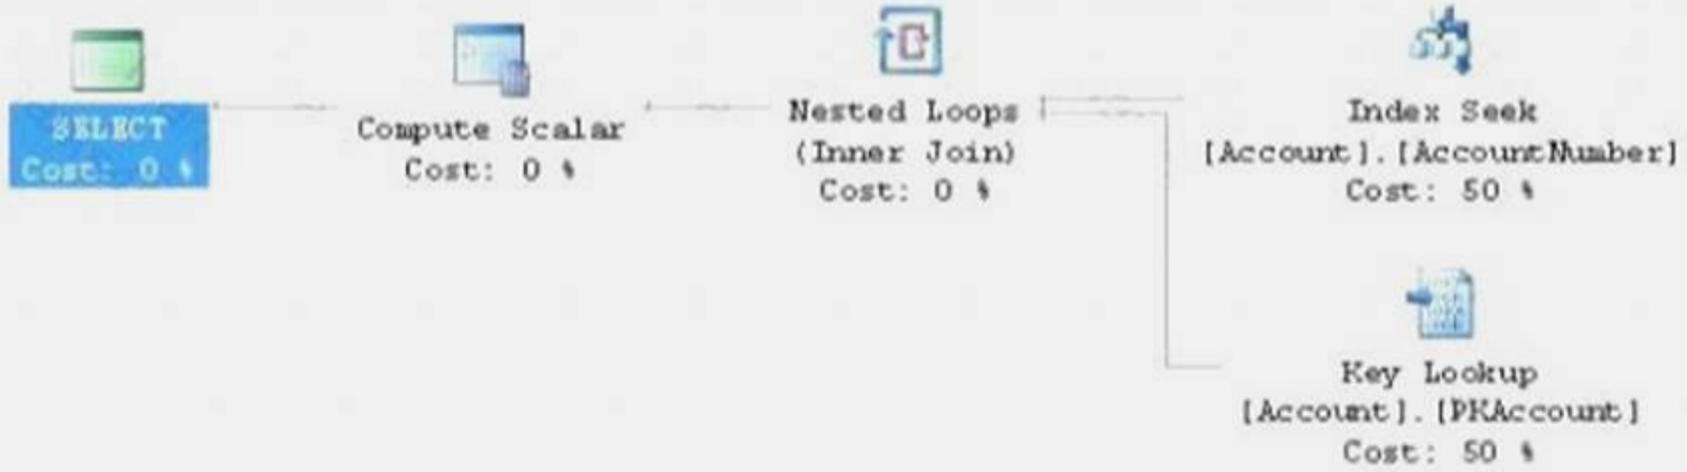
You administer an instance of SQL Server 2014.

You are tasked with tuning a common set of queries. You have the results of several test executions, along with query plans. The schema and the data for all

database object(s) used remain unchanged between executions. The QueryTime column is defined as a computed column that uses the GETDATE() system function. The query plans and results are shown below:

```
SELECT *
FROM dbo.Account
WHERE AccountNumber = 'A10000001'
```

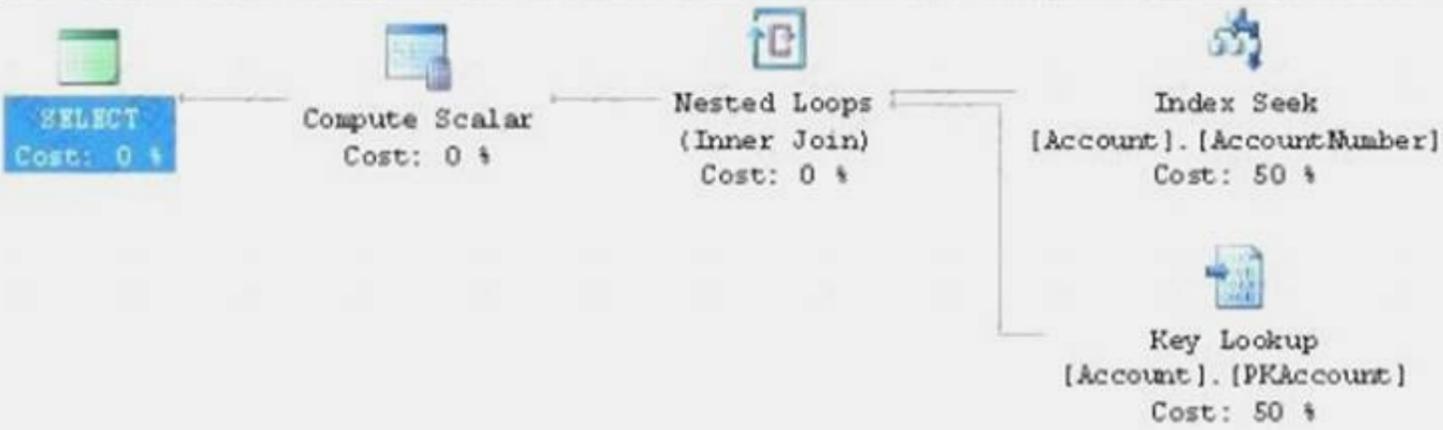
Query 1: Query cost (relative to the batch): 100%
 SELECT * FROM [dbo].[Account] WHERE [AccountNumber]=@1



AccountID	AccountNumber	Name	QueryTime
-----	-----	-----	-----
-----	-	-	----
0F63B176-7257-4480-9D0E-126C45 CEFFF1	A10000001	Don Hall	2014-01-29 18:01:50.923

```
SELECT *
FROM dbo.Account
WHERE AccountNumber IN( 'A10000001', 'Q88700323', 'R00000012' )
GO
```

Query 1: Query cost (relative to the batch): 100%
 SELECT * FROM dbo.Account WHERE AccountNumber IN('A10000001', 'Q88700323', 'R00000012')



AccountID	AccountNumber	Name	QueryTime
----- -----	----- -----	----- -----	----- -----
0F63B176-7257-4480-9D0E-126C45 CEFFF1	A10000001	Don Hall	2014-01-29 20:14:05.660
337227AA-3A4B-4B28-8E02-0ADEAD 06EA10	Q88700323	Darren Parker	2014-01-29 20:14:05.660
C4980E64-874E-4640-8826- BAF35D8FB845	R00000012	Carol Philips	2014-01-29 20:14:05.660

You need to make an initial diagnosis of the situation, based solely on this input
 Which two statements can you make about the performance characteristics of this query? Each correct answer presents a complete solution. Choose two.

- A. The queries would perform better if the index named AccountNumber included the Name and QueryTime column.
- B. The queries would perform worse if the index named AccountNumber included the NameColumn.
- C. The queries would perform better if the index named AccountNumber included the Name column.
- D. The object Account is a table, with an index having a leading column of AccountNumber and a Clustered Index named PKAccount.
- E. The object Account is an indexed view, with an index having a leading column of AccountNumber and a Clustered Index named PKAccount.
- F. The object Account is a view, joining the Account-AccountNumber and Account.PKAccount objects together.

Answer: BD

NEW QUESTION 193

- (Exam Topic 8)

You have a server that has SQL Server 2012 installed.

You need to identify which parallel execution plans are running in serial. Which tool should you use?

- A. Performance Monitor
- B. Database Engine Tuning Advisor
- C. Data Profile Viewer
- D. Extended Events

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/bb677278.aspx> <http://msdn.microsoft.com/en-us/library/bb630282.aspx>
<http://www.sql-server-performance.com/2006/query-execution-plan-analysis/>
<http://www.simple-talk.com/sql/learn-sql-server/understanding-and-using-parallelism-in-sqlserver/>
<http://www.sqlservercentral.com/articles/SQL+Server+2012/At+last%2c+execution+plans+show+true+thread+r>
http://sqlblog.com/blogs/paul_white/archive/2011/12/23/forcing-a-parallel-query-executionplan.aspx http://sqlblog.com/blogs/paul_white/archive/2012/05/02/parallel-row-goals-gone-rogue.aspx <http://msdn.microsoft.com/en-us/library/bb895310.aspx>
<http://msdn.microsoft.com/en-us/library/bb895313.aspx> <http://msdn.microsoft.com/en-us/library/hh231122.aspx>

NEW QUESTION 197

- (Exam Topic 8)

You have a database that is accessed by 300 concurrent users.

You need to log all of the queries that become deadlocked. The solution must meet the following requirements:

- ▶ Provide a representation of the deadlock in XML format.
- ▶ Minimize the impact on the server.

What should you create?

- A. A SQL Server Profiler trace
- B. A script that enables trace flags

- C. A SQL Server Agent job that retrieves information from the sys.dm_tran_active_transactions dynamic management views
- D. A SQL Server Agent job that retrieves information from the sys.dm_tran_session_transactions dynamic management views

Answer: A

Explanation:

Analyze Deadlocks with SQL Server Profiler

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. To trace deadlock events, add the Deadlock graph event class to a trace. This event class populates the TextData data column in the trace with XML data about the process and objects that are involved in the deadlock. SQL Server Profiler can extract the XML document to a deadlock XML (.xdl) file which you can view later in SQL Server Management Studio.

NEW QUESTION 201

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-464 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-464 Product From:

<https://www.2passeasy.com/dumps/70-464/>

Money Back Guarantee

70-464 Practice Exam Features:

- * 70-464 Questions and Answers Updated Frequently
- * 70-464 Practice Questions Verified by Expert Senior Certified Staff
- * 70-464 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 70-464 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year